Who Knows What They Know?

An Experimental Investigation into the Impact of Explicit Metacognition on Cumulative Cultural Evolution

Juliet A F Robertstone (née Dunstone)

Thesis submitted for the degree of Doctor of Philosophy

Psychology

University of Stirling

23rd October 2020

Abstract

This thesis aims to investigate the cognitive mechanisms underlying human cumulative culture using a novel testing paradigm. Specifically, it experimentally tests the proposal that explicit, metacognitive processes are a requisite capacity for cumulative cultural evolution due to their facilitation of selective copying strategies. This proposal is referred to as the Explicitly Metacognitive Cumulative Culture hypothesis (EMCC, chapter 1). The experimental paradigm used throughout the four studies presented in this thesis aimed to restrict access to explicit, metacognitive processes via the use of dual-tasks. In chapter 2 a series of studies examined the role of task-switching in simple search tasks that required flexible strategies in order to be solved successfully. The final study in chapter 2 also looked at the impact of the same task-switching task on metacognitive monitoring. In chapter 3 the impact of working-memory restriction on selective copying strategies and cumulative score improvement, as well as metacognitive monitoring, was examined.

Overall the findings from this thesis indicate that working memory resources may play a role in facilitating efficient, selective copying strategies. Optimal use of selective strategies resulted in ratcheting over generations, but the addition of concurrent tasks reduced this ratcheting. Metacognitive monitoring did not appear to be affected by a concurrent dual-task. However, there was some indication that metacognitive control strategies could play a role in selective copying, which may not have been tapped by the metacognitive monitoring task.

The novel methods used throughout this thesis have laid important groundwork for future empirical testing of this hypothesis, and the final chapter considers the new research opportunities opened up by these studies.

Acknowledgements

I would like to thank my supervisor Christine Caldwell, for sparking my interest in this research years ago while I was an MSc student, for being so normal when I first met you even though I thought you were basically a celebrity, and for providing 4 years of unwavering support and guidance. I am also extremely grateful to my secondary supervisors: Mark Atkinson for teaching me how to teach myself coding, a gift which has been more valuable than I could have foreseen; and Catherine Grainger for sharing your advice, knowledge and enthusiasm for my research at all stages of my PhD.

To Elizabeth Renner and Rebecca Sharman: thank you for your encouragement, friendship and for showing me that being a woman in academia doesn't necessarily mean you have to grow up.

I am eternally grateful to my fellow RATCHETCOG PhD students, Gemma Mackintosh, Donna Kean, Kirsten Blakey and Charlotte Wilks for providing an endless supply of proofreading, stats-fixing, ranting opportunities, tea breaks, wine sharing and emotional support while we have shared an office and a journey.

To my gal pals, Catherine, Phoebe, Sophie, Kate, Lauren and Vickie: thank you for being a virtual and physical shoulder to cry on, for keeping me spilling over with joy (and tequila) as often as possible, and reminding me that life exists outside of my PhD bubble.

Finally, to my husband Josh. For moving to Scotland so I could embark on this journey, for keeping a roof over (and my sanity inside) my head, for always trying your hardest to understand what I do and for being endlessly enthusiastic and supportive. I could not have done it without you.

A last honourable mention goes to all the pets that have been by my side over the course of my PhD. Thank you Zola, Annamaria, Possum, Langoustine (RIP girls), and Mouse, for making sure my heart was always full.

Preface: Introduction and Thesis Outline

Many species exhibit unequivocal examples of cultural behaviour and can learn this behaviour from conspecifics (Whiten et al., 2016). However, the capacity for cultural traits to accumulate in quality or efficiency over generations – cumulative cultural evolution (CCE), also known as the *ratchet effect* (Tomasello, 1990) is generally regarded to be unique to humans (Dean et al., 2014; Tennie et al., 2009). There is still no consensus on exactly why this is. This thesis aims to investigate one theory: that access to explicit metacognition and system-2 processes is a requisite cognitive capacity for selective social learning strategies in humans, and that these strategies enable cumulative cultural evolution to occur. This thesis introduces this theory as the Explicitly Metacognitive Cumulative Culture Hypothesis (EMCC).

The EMCC is based on theoretical work that posits that explicitly metacognitive social learning strategies allow learners not just to learn from others, but to be explicitly aware of the strategies they are using. This enables them to target their social learning to the most appropriate model, based on their own learning requirements (Heyes, 2016). This is hypothesised to lead to cumulative improvement in traits due to more streamlined reproduction of beneficial traits, with less copying of maladaptive or inefficient practises. A related theory argues that explicit metacognition is beneficial to CCE as it increases capacities to share accurate information about one's own knowledge, in order to facilitate efficient group action (Shea et al., 2014).

These two theories are undoubtedly linked. However, extensive research on the benefits of confidence sharing on group action (Bahrami et al., 2010, 2012; Hertz et al., 2016) and social learning (Olsen, Roepstorff, and Bang, 2019) has already been conducted. In contrast, little to no empirical work has been conducted to examine the benefit of metacognition on selective social learning strategies and CCE. Recent evidence from a long-running programming competition, in which participants were able to use code written by previous contestants shows that selective social learning produces cumulative improvement over time (Miu, Gulley, Laland, & Rendell, 2020). While indicating support for the EMCC, this study did not assess the cognitive mechanisms that enabled the selective copying to occur, or what influenced the individual differences in copying strategies. It is therefore the aim of this thesis to provide the first empirical testing of whether explicit, system-2 processes and metacognition facilitate CCE by providing access to selective social learning strategies.

Metacognitive SLSs are also argued to be themselves a product of cultural evolution (Heyes, 2018a). However, experimental testing of this is unfortunately beyond the scope of a single thesis. The EMCC therefore makes no specific predictions about the origins of explicitly metacognitive SLSs.

In chapter 1 I will outline the EMCC theory in detail and suggest practical ways to test it. In order to test the EMCC, a comparison needs to be made between participants with typical metacognitive capacities and participants without this, in their respective propensity for selective social learning and ratcheting over generations. As metacognition cannot be experimentally removed in order to make this comparison, the methods outlined in this chapter, and used throughout the empirical chapters of this thesis, instead aim to experimentally *restrict access* to explicit metacognition.

Throughout this thesis the definition of 'explicit' is defined using the same criteria as Evans & Stanovich (2013) use in their description of a dual-systems viewpoint: processes are not strictly segregated into one system or another. However, processes are considered to be explicit if they rely on executive functions and are slow or deliberate to execute. Given this assumption, and given the links between executive functions and metacognition (Roebers, 2017; Roebers & Feurer, 2016), the paradigm used throughout this thesis aims to restrict access to explicit metacognition through the use of executive function dual-tasks.

In chapter 2 I will begin testing the EMCC by investigating the impact of executive function dual-tasks on tasks which require the use of selective copying strategies and metacognitive monitoring. As executive functions are not a single unitary skill, the chapter initially assesses multiple different dual-tasks which aim to experimentally restrict different aspects of executive function. It then uses one of these tasks to directly compare the impact of restriction of executive functions on flexible copying strategies, with the impact on fixed strategies requiring the use of a single rule. The same dual-task paradigm is also used to investigate the impact of restricting executive function on metacognitive monitoring accuracy.

In chapter 3 the EMCC hypothesis will be further tested by examining the dual-task impact of cumulative improvement over generations in a more complicated task which still requires selective copying as an optimal strategy. This chapter assesses performance in a task which aims to demonstrate CCE experimentally without running multiple generations of participants. Each participant's *potential* to show an additive ratchet effect, if they were placed in a chain of participants (their potential for ratcheting, or PFR), is examined by exposing each participant to multiple trials containing differing levels of information. This allows for the precise manipulation of the dual-task and control conditions, while still using a task that can demonstrate an accumulation of benefit over generations. This chapter will compare participants' PFR when they have restricted access to their executive resources due to a working memory dual-task, with their performance under control conditions. It will then use the output of that task to simulate multi-generational turnover of the same task, to investigate whether restriction of working memory prevents a ratchet effect from occurring. It will also again examine the impact of a dual-task on metacognitive monitoring, this time using a working memory task.

Finally, in chapter 4 I will summarise the key findings from the experimental chapters, investigate any study limitations and suggest potential avenues for future research into this hypothesis.

Contents

Chapter 1: Outlining the Explicitly Metacognitive Cumulative Culture Hypothesis

The following chapter has been published as part of the Palgrave Communications special issue on Cultural Evolution. It is reproduced here in an identical format to the final submission sent to the journal. The reference for that article is as follows:

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

As chapters 2 and 3 of this thesis are submitted as journal articles, when referring back to this chapter the citation is to ‘Dunstone and Caldwell (2018)’, rather than to ‘chapter 1’.

Introduction

In the current article, we examine the view, recently proposed by Heyes (2016) and Shea et al. (2014), that distinctively human cultural evolution is attributable to capacities for explicit (or Type/System 2) metacognition. Essentially, these accounts argue that an ability to explicitly reflect upon states of knowledge, ignorance, and uncertainty, can fundamentally change the ways we use and share social information, and that these particular processes account for the characteristic forward progress of human culture. In the current review, we aim to evaluate these accounts, considering the evidence for their underlying assumptions as well as the plausibility of mechanistic routes which could potentially link individual-level cognitive processes of explicit metacognition, with population-level outcomes resembling cumulative culture.

Explanations of the Distinctiveness of Human Cumulative Culture

Cumulative cultural evolution is the process by which cultural traits (including behaviours, artefacts and tools) change over multiple episodes of social transmission to become more effective and beneficial to their users (Mesoudi & Thornton, 2018; Caldwell, 2020). In humans, this can lead to cultural traits evolving over many generations which could not have been invented by a single individual. Examples of human cumulative culture span a wide range of domains including abstract conceptual skills such as the cognitive tools provided by mathematical notation and operations (e.g. Bender & Beller, 2014), or survival skills such as lengthy food processing techniques that remove invisible toxins from raw ingredients (e.g. preparation of cycas seeds, Beck, 1992).

Many species have been shown to exhibit cultural traditions in the form of behavioural variation between populations which appears to be maintained by social learning, or evidence of the social diffusion of particular behavioural variants (Whiten et al., 2016). To highlight just a few examples, such evidence has been identified in chimpanzees (Hobaiter et al., 2014), humpback wales (Allen et al., 2013), and great tits (Aplin et al., 2015). Therefore, neither the capacity for culture, nor even the process of cultural evolution across generations (i.e. in the broader sense of any kind of incremental change arising as a consequence of social transmission, in the absence of improving functionality), is restricted to humans. In contrast, the cumulative improvement of traits across generations (sometimes referred to as the cultural “ratchet effect”, e.g. (Tomasello, 1990), is widely regarded to be unique to humans (Dean et al., 2014; Tennie et al., 2009).

Several previous theories have been proposed for why humans have an apparently unique capacity for cultural ratcheting. The most prominent view concerns particular learning mechanisms also proposed to be human unique. Humans are extremely adept at copying others' behaviour (e.g. dubbed "*Homo imitans*" by Meltzoff, 1988) and seem to have a particular talent (or even a compulsion in some situations, Heyes, 2011) to do so. Conversely, non-human animals (henceforth animals) do not seem to exhibit the same proclivities to anything like the same degree. While there is evidence for action imitation in some species (e.g. marmosets: Voelkl & Huber, 2000), this occurs at much lower levels of accuracy, and in far more restricted contexts than in humans who are highly accurate copiers in domain general contexts. This has prompted theorists to propose that capacities for imitation and social learning may have represented critical cognitive developments in human evolution, allowing for cumulative culture (Lewis & Laland, 2012; Tomasello, 1999). Converging lines of evidence from computer models, and tournaments in which alternative strategies compete in simulated mixed populations (Rendell, Fogarty, et al., 2010) have found that the most successful strategies involve high rates of copying, and that model populations predominantly comprising social learners (simulated agents that "acquire a behavior performed by another individual, whether by observation of or interaction with that individual" (Rendell, Boyd, et al., 2010)) are more successful and develop more complex technologies.

There are however a number of reasons to question the notion that particular social learning mechanisms may account for human cumulative culture. Firstly, there is now mounting evidence of imitative abilities in other species (e.g. apes: Whiten et al., 2004, although see Call, Carpenter, & Tomasello, 2005 for evidence counter to these claims) Secondly, in humans, experimental evidence has shown that cumulative culture can arise even when participants are restricted from observing others' actions, such that these learners are forced to rely on non-imitative processes such as emulation of end products (Caldwell & Millen, 2009). In addition, although cumulative culture necessarily involves social transmission as a mechanism for trait heritability, it is important to note that social learning alone cannot account for the increases in trait functionality that exemplify the process. The development of new technologies and behaviours also depends on innovation (Enquist et al., 2008; Lehmann, Feldman, & Kaeuffer, 2010). While copying error or accidental discovery, as well as intentional invention, can be sources of innovations (Caldwell, Cornish, et al., 2016; Henrich et al., 2008), it is clear that increased abilities for high-fidelity copying alone cannot explain cumulative increases in cultural behaviours and artefacts.

These considerations imply that cumulative culture may be explained not only by the mechanisms available to learners, but also the contexts in which these are employed, since it is the selective and strategic use of copying that accounts for its adaptiveness; high fidelity copying may be necessary for cumulative culture to emerge, but it is not sufficient. Indeed, models show that populations of flexible learners that can switch between social and individual learning at critical points outperform populations composed of only social learners, or only individual learners (Ehn & Laland, 2012; Enquist, Eriksson, & Ghirlanda, 2007; Rendell et al., 2010).

Laland (2004) described a number of potential "Social Learning Strategies" (SLSs) which could reflect adaptive rules regarding when to engage in social learning, and who to learn from. These included strategies such as: Copy when Uncertain, Copy the Majority, and Copy Successful Individuals. However, in spite of the fact that such strategies clearly have some potential to explain the selective retention of beneficial traits in learner populations, they nonetheless fail to provide an adequate explanation for the fact that cumulative culture appears to be restricted to humans. This is because a wide range of animals have also been shown to exhibit SLSs. This includes social insects (Smolla et al., 2016), fish (Pike & Laland, 2010) and bats (Jones et al., 2013).

However, Heyes (2018b) has drawn a distinction between social learning strategies that are based on "planetary" decision rules with those based on "cook-like" decision rules. Planetary SLSs, like laws of planetary motion, capture regularities within the observable behaviour of the entities of interest, but the rules are only in the minds of those doing the describing. In contrast, "cook-like" SLSs are more akin to the decision rules used by a cook following a recipe, i.e. they are explicitly represented within the mind of the agent. Heyes (2016; 2018b) has thus argued that it is these *explicitly metacognitive* SLSs that account for the elaborate outcomes of human cumulative culture. Although there is no dissent regarding the fact that animal social learning strategies demonstrate adaptive flexibility, these "cook-like" SLSs are assumed to permit a much higher degree of – insightful – flexibility, potentially optimising the effectiveness of social transmission in a number of different ways (see section *How Might Explicit Metacognition Facilitate Cumulative Culture?*).

The proposal that explicit metacognitive processes may set human social learning apart from that of other animals is compelling, and persuasive theoretical arguments in its favour can be found in Shea et al (2014), Heyes (2016), and Heyes (2018b). In the current review we consider the evidence in support of this explanation (which we refer to here as the Explicitly Metacognitive Cumulative Culture hypothesis, the EMCC), over and above competing alternatives. Recent literature (Heyes, 2018a) elaborates further on this argument to claim that these metacognitive strategies are themselves products of cultural evolution, as well as processes supporting it. Evaluating this extension of the theory is beyond the scope of this current review however, and as such we do not discuss this argument further.

Metacognition and the EMCC as discussed here also encompass capacities for theory of mind, or mentalising. As argued by Carruthers (2009) metacognition and theory of mind are not wholly distinct capabilities, and metacognition about one's own mind may in fact rely upon prerequisite capacities for mind-reading about others' minds. Indeed, some arguments proposing the utility of explicit metacognitive social learning strategies encompass mentalising as much as introspection (e.g. Heyes, 2016, when asking "who knows?"), but see section '*Optimisation of receiver behaviour, due to understanding of others' knowledge states'* for a more detailed discussion of this not uncontroversial distinction.

We begin below by examining a key assumption of the EMCC hypothesis, which is that explicit metacognitive processes are restricted to humans.

Explicit Metacognition as a Uniquely Human Feature

The literature contains numerous claims of metacognitive ability in animals, across a broad range of species (e.g. monkeys: Smith, Redford, Beran, & Washburn, 2009; chimpanzees: Beran et al., 2015; dolphins: Smith et al., 1995; pigeons: Sole, Shettleworth, & Bennett, 2003; rats: Foote & Crystal, 2007; Templer, Lee, & Preston, 2017; and even bees: Perry & Barron, 2013). However, the EMCC rests on the assumption that the experimental paradigms used in these studies are assessing qualitatively different phenomena from the type of metacognition required for cumulative culture, which is assumed to be unique to humans. In this section we examine theories and evidence underlying the assumption that alternative methodologies in metacognition research may be evaluating fundamentally different cognitive processes, and that certain types of metacognition may indeed be manifested only in humans.

Dual Processing Theories of Cognition

As noted above, a critical point in the EMCC is that only humans have conscious access to their social learning decision-making rules, whereas social learning in other animals is driven by automatic processes of which the agents themselves are unaware. It is this difference that is proposed to account for the unusual prevalence of cumulative culture in humans, compared with other species. It is important to note that the EMCC does not imply that *all* human cognition involves conscious access, or indeed that all examples of social learning in humans are based on explicit processes. Rather, the hypothesis draws on theories of human cognition which propose the existence of two systems, or two processing types. We present an overview of this body of literature below.

Theories of dual processes for various aspects of cognition have been relatively widespread since the 1970s (e.g. Wason & Evans, 1974). These theories state that there are two different modes of higher cognitive processing; one which is generally automatic, fast acting, non-conscious and based on associative mechanisms, and one which is conscious, slower to act and rule-based (see Evans & Stanovich (2013), for a summary of attributes typically associated with each of the processing types). These two alternatives are generally referred to as either *Systems* (Systems 1 and 2), or *Process Types* (Type 1 and 2), to capture the automatic (1) and rule-based (2) cases respectively. Although the idea of different *Process Types* offers a less theoretically loaded framework, which is potentially more consistent with a wider range of empirical evidence (e.g. reflecting a continuum, rather than a dichotomy, of alternative cognitive mechanisms), it is the *Systems* label which has been associated with the idea that there may be distinctively human modes of cognition ((Epstein, 1994); (Stanovich, 1999, 2004)). Accordingly, it is the *Systems* label that has been used in the literature relating dual process theories to human cultural evolution (Shea et al., 2014; Heyes, 2016; Heyes, 2018b). In the current review we use both terms.

In relation to the issue of human distinctiveness, some accounts hold that System 1 is phylogenetically ancient, and therefore shared with other animals, whereas System 2 is more recently evolved and likely to be unique to humans. Dual processing theories have been used as a framework for the interpretation of a diverse range of psychological phenomena, from decision making (T. A. Evans, 2007), learning (Dienes & Perner, 1998) and social cognition (Smith & DeCoster, 2000).

Evidence for dual processing comes from dual-task studies, and tasks which apply strict time pressures. This is because System 2 processes are argued to be taxing on executive functions and working memory capacity, as well as generally taking longer. Dual-tasks are designed to put an additional load on finite cognitive capacities. If two tasks require the same cognitive mechanisms this creates a bottleneck in processing, resulting in delayed response or impaired performance in one or both tasks (Pashler, 1994). This means a dual-task can detect what level of processing is being used to complete a task; if task performance is unimpeded by a concurrent working memory or executive function load it is likely to be an automatic (or System 1) process, whereas if working memory load significantly reduces speed or accuracy of responding it is likely to be System 2. Participants under a working memory load have been found to make more incorrect responses based on salient information rather than logical reasoning when completing conjunction fallacy problems, or logic puzzles such as the Wason Selection Task (De Neys, 2006). The application of a strict time pressure may also prevent the use of System 2, as it would not allow the longer processing time needed. This effect has also been found using the Wason Selection Task (Roberts & Newton, 2001).

There is also some neurological evidence for dual processing systems: distinct brain activations for using logic based (System 2) and belief based (System 1) solutions to problems have been found using fMRI (Goel & Dolan, 2003). Additionally, NIRS analysis found that areas implicated in incongruent reasoning trials were not activated when the same tasks were performed under additional cognitive load (Tsujii & Watanabe, 2009). Mcclure, Laibson, Loewenstein, & Cohen (2004) found activation in the prefrontal cortex during reasoning about future monetary rewards but not during immediate decision making. This activation was found in a similar neurological region as areas associated with metacognition and executive functions.

Dual processing theories are by no means universally accepted; see Keren and Schul (2009) or Osman (2004) for some objections. However, critical accounts have generally focused on lack of precise definitions, or evidence of overlap between the proposed dichotomy of characteristics between System 1 and System 2, concluding that the two systems cannot be considered distinct and isolable. However, Evans and Stanovich (2013) have argued that this is merely a poor interpretation of the literature, and that most of the features commonly described as differentiating Type 1 and 2 are just correlates typical of the processing types, and that these should not be expected to operate in a categorical, mutually-exclusive fashion.

Evans and Stanovich (2013) describe their theory of dual processing as a “default-interventionist” view (e.g. p227), meaning that the majority of cognition relies on System 1 processes unless the available response is incorrect or does not meet with the task goals, in which case System 2 will intervene.

Metacognition and Dual Process Distinctions

It is central to the EMCC that it is the agents' conscious access to their decision-making rules that allows human metacognitive processes to generate such highly adaptive outcomes from social learning (e.g. by allowing learners to seek out the most appropriate models dependent on their specific goal, Heyes (2018b), or by allowing those in possession of knowledge to broadcast their degree of confidence or uncertainty as well as their choices, Shea et al. (2014)). It is perhaps unsurprising then that these theories have emphasised the importance of the System 1/System 2 distinction. However, in order to fully understand and evaluate these accounts, it is important to also consider metacognitive processes more generally, and then to turn to the question of how these relate to the dual process framework detailed above.

Metacognition, as originally defined by Flavell (1979), can be thought of as knowledge about one's own cognitive processes, or cognition about cognition. Flavell split this into four separate components: knowledge (of your own cognitive abilities and of learning processes), experiences (current feelings of certainty or doubt), goals (objectives you have in order to achieve your current cognitive task) and actions (behaviours employed to achieve these objectives). Subsequent research has instead divided metacognition into declarative metacognitive knowledge (corresponding to Flavell's knowledge) and procedural metacognition. Procedural metacognition has then been commonly divided into monitoring and control processes (Flavell's experiences and actions respectively) (Roebers, 2017).

The term metacognition therefore encompasses a wide variety of phenomena. Accordingly, it has been used to describe a broad spectrum of findings identified in a wide range of contexts, from detecting tiny changes in perceptual stimuli (Deroy et al., 2016) to deliberately allotting revision time for exams or correcting errors in a piece of written text (Sannomiya & Ohtani, 2015). In addition, some authors consider metacognition to encompass understanding of others' cognition as well as one's own, as part of a wider cognitive capacity for metarepresentation relating to mental states (Kuhn, 2000; Misailidi, 2010). In a large proportion of experimental research, metacognition has typically been operationalised as judgements of confidence in performance of an activity just completed (Judgements of Confidence; JOC), or ratings of prospective performance in an activity about to be completed (Feeling of Knowing; FOK) (Nelson & Narens, 1990)

Given the specific emphasis on explicit/System 2 metacognition within the EMCC, it is worth considering here examples of metacognitive phenomena which could be classified as implicit or System 1 processes, and how these differ from those assumed to implicate explicit or System 2 processes. Any paradigm involving direct report of degree of confidence, doubt or uncertainty, necessarily requires awareness of these states, and therefore would be classified as implicating explicit metacognition. However other, more indirect, methods have also been used as means of evaluating metacognition, particularly within animal studies.

The typical methodological paradigm used to solicit “metacognitive” behavioural responses in the absence of verbal report, involves offering an “opt-out” option within a decision making task, which can be used adaptively by the participant to avoid the risks associated with particularly difficult trials. Adaptive use of the option to opt out is assumed to reflect the subject’s appreciation of their own uncertainty, and therefore metacognition. Such designs have been used to support claims of metacognitive ability in nonhuman primates (e.g. macaques: Smith, Beran, Couchman, & Coutinho, 2008). Alternative methodologies involve ‘information-seeking’ paradigms, where a participant can seek additional information before making a decision, which have also been used to support claims of metacognitive awareness in primates (e.g. Call & Carpenter, 2001).

However, these experiments remain contentious as demonstrations of true metacognitive ability, as adaptive performance could be explained by responses being driven by first-order states of anxiety elicited by the uncertainty of the situation, rather than second-order reflection on the state of uncertainty itself (Carruthers & Ritchie, 2012). This is explained most thoroughly by Carruthers (2008), who has argued that first-order beliefs, along with other basic mechanisms such as signal detection theory, are just as capable of explaining the findings. The account can be summarised thus: the participant is presented with two choices that carry equal valence (the animal is equally motivated to both choose and not choose either option). As soon as a third option (the opt-out or uncertain option) is presented this option automatically becomes the most attractive, especially given its reinforcement history of being associated with a small reward. As this explanation is simpler, in terms of the cognition required by the animal participants, Carruthers argues convincingly that this is a more parsimonious explanation than those ascribing metacognitive capacities to animals. A similar account from Hampton (2009) described how a range of studies claiming to show animal metacognition could be explained by environmental or behavioural cues, or direct competition between a choice to act, or make a “metacognitive” action. These accounts make it clear that although metacognitive introspection could in principle explain the results of the studies in question, the plausibility of such interpretations is seriously challenged by the availability of simpler explanations. Under the dual process framework, the animals’ performance in these studies would therefore likely be classified as System 1, or implicit, metacognition (e.g. Shea et al., 2014).

There are unavoidable challenges involved in establishing whether implicit and explicit metacognitive responses depend on different cognitive processes, especially if our ultimate motivation is to determine whether one is a distinctive feature of human cognition. Adult humans are necessarily capable of both, and we can only use non-verbal measures with animals due to the language requirements of direct assessments of explicit metacognition. However, patterns of emergence during human development potentially provide another source of evidence that could shed light on the relationship between implicit and explicit metacognitive behaviour, and whether implicit adaptive responding can occur in the absence of explicit competence.

Behavioural tests analogous to the “opt-out” paradigms used with animals (described above) have been used to demonstrate implicit metacognitive ability in very young children, from 20 months to around 5 years old. Behavioural tests of metacognitive competence have also included assessment of spontaneous information-seeking prior to committing to a response. These paradigms also potentially provide an insight into implicit reactions to the state of ignorance, without necessarily implicating metacognitive awareness of that state. These studies will not be described in detail here, as reports of implicit metacognitive measures are not directly relevant to the current review. However, please see Bernard, Proust, & Clément (2015), Goupil, Romand-Monnier, & Kouider (2016) or Call & Carpenter (2001) amongst others for examples of the paradigms in question.

In spite of the early development of such behavioural responses to uncertainty (analogous to the evidence from animals), evidence of explicit metacognitive understanding only appears to emerge later. Although verbal reports can be readily obtained from preschool aged children, studies requiring them to verbalise their own state of knowledge nonetheless indicate that they have difficulty doing this and that when they do they show a pervasive bias towards overestimation of their own knowledge and performance (e.g. Rohwer, Kloo, & Perner, 2012).

The earliest examples of accurate performance based on explicit measures of metacognition come from children of around four to five years old. For example, Rohwer, Kloo and Perner (2012) found that only children older than five could provide reports about what they did not know in a partial exposure task. Cultice, Somerville and Wellman (1983) also found accurate explicit metacognitive responding in children aged four and five years old, asked to name familiar individuals from their photograph. When children were unable to spontaneously recall the name themselves, they could respond with reasonable predictive accuracy to the question: "If I told you a lot of names, do you think you would know or remember which one was her name?".

The adaptive responses to uncertainty identified in children younger than four years old (including those aged one and two, e.g. Goupil et al., 2016, and Call & Carpenter, 2001) appears strikingly similar to the behaviour of animals in opt-out and information seeking studies. However, this kind of competence appears to precede the ability to provide explicit, accurate evaluations of states of knowledge, which apparently only develops some years later. This would therefore seem to corroborate accounts which propose that successful performance on the alternative task types is underpinned by different processes, and that the animal studies therefore do not provide evidence of explicit metacognition. This is consistent with the EMCC's assumption that System 2 metacognitive capacities are specific to humans.

How Might Explicit Metacognition Facilitate Cumulative Culture?

As an explanation for distinctively human cumulative culture, the EMCC rests on two fundamental assumptions. The first of these is the corresponding distinctiveness of explicit metacognition (as examined in the preceding section). The second of these is that the resulting reflective awareness of states of knowledge, ignorance and uncertainty (identified as the defining feature of explicit metacognition) offers significant benefits with regard to the optimisation of social information use, in ways that could explain the ratchet-like advances which distinguish human culture from the traditions of other species. Having considered the first of these premises in the preceding sections, we now turn to the second. What basis is there, either evidential or logical, for believing that explicit metacognition might enable cumulative culture? What are the potential routes by which this might occur? We hope that by clarifying the potential links between explicit metacognition and cumulative culture we can identify areas where evidence is lacking, with a view to informing future research efforts investigating the EMCC.

Explicit metacognition could potentially enable cumulative culture in a number of different ways. Below, we categorise the potential benefits as arising from receiver behaviour, or sender behaviour. Within both of these categories benefits could arise as a consequence of more effective representation of one's own knowledge state, or that of others. It should be noted at this point that the existing accounts of the EMCC focus respectively on optimisation of sender behaviour due to understanding of own knowledge state (Shea et al., 2014), and optimisation of receiver behaviour due to understanding of others' knowledge states (Heyes, 2016), both detailed below.

Optimisation of receiver behaviour, due to understanding of own knowledge state

In much the same way that metacognitive awareness is assumed to facilitate academic performance (Dunlosky & Metcalfe, 2009) it is possible that it could similarly enable cumulative culture for reasons that are not inherently linked to how an agent understands or interacts with others. Awareness of one's own knowledge state would allow learners to seek out new information when necessary, and recognise when updating their knowledge might be beneficial. This awareness might also be crucial when acquiring a new skill or knowledge is likely to require a protracted period of effortful practice before mastery is achieved. Essentially, we would predict that such awareness would result in social information being used in a much more optimal fashion than would otherwise be possible, encouraging highly strategic social information seeking, as well as direction of effort towards innovation when social information sources are judged to be inadequate. To our knowledge, the role of this kind of reflective awareness in directing one's own learning has not been investigated within the social learning and cultural evolution literature. However, some authors have alluded more tangentially to the importance of self-focussed strategic effort in social learning. For example, Galef (2013) has stated: "in the case of skiing, there is no learning to do an act from seeing it done. Rather, there is learning by observation that an act is possible.... [A] novice can … select from within her available repertoire of movements.... Then, over time, she can bring that first approximation into greater accord with the demonstrated act." (p125). Galef (2013) also suggests that such learning may be particularly important for cumulative culture.

As noted above, this would be a route by which explicit metacognition might be critical to generating cumulative culture without the effects being restricted to *social* learning specifically. In the accounts of both Shea et al. (2014) and Heyes (2016), explicit metacognition is assumed to facilitate cumulative culture because it helps agents make inferences about others' knowledge (Heyes, 2016) or provide information to others (Shea et al., 2014). It is perhaps not surprising that, in attempting to explain a phenomenon which itself certainly does depend on social learning, authors have focused on explanations which would specifically facilitate that type of learning over and above others. However, we would suggest that explicit metacognition might potentially facilitate efficient use of any kind of vicariously acquired information, as well as helping in any situation where habitual or automatic responses may need to be overridden due to the availability of up-to-date, or situation-specific, information which indicates that these are not appropriate. This account would be consistent with the assumptions of the "default interventionist" views of dual-process cognition described earlier, which posit that System 2 intervenes only when automatic System 1 processes are inadequate for the task in question. Thus, explicit decision rules, reasoning processes and learning strategies are likely to be intrinsically associated with situations where default responses (based on personal reinforcement history and/or genetically inherited behavioural biases) will be ineffective.

Although such situations will be by no means restricted to contexts involving social information use, the need to override default automatic and habitual responses may be a prevalent feature of these contexts. Consider a situation in which a new possibility becomes apparent to an agent through vicarious exposure to another's behaviour; for example, the agent might observe that plentiful food resources, such as tubers, could be found underground. Taking full advantage of this new information might necessitate an immediate switch in foraging strategy, overriding habitual responses which have been directly reinforced on multiple occasions. Although similar exposure to new information might occur outside of social contexts (e.g. if tubers were to be revealed as an incidental outcome of disturbance of the ground surface), the behaviour of others is perhaps particularly likely to provide information of immediate utility. Furthermore, once transgenerational accumulation of knowledge was in evidence, social sources would then effectively become repositories of particularly valuable information that might be otherwise hard to acquire. Therefore the benefits of this type of learning might be most apparent in social contexts, even though the learning mechanisms themselves would be general-purpose ones, not specifically adapted for use in social contexts.

Essentially, the suggestion here is that System 2 metacognition may be critical due to the high "executive function" demands of the type of social learning likely to be involved in cumulative culture. Overlap between the concepts of executive function and metacognition have been acknowledged in the existing literature (e.g. Roebers, 2017). Indeed, some research effort has already been targeted at the question of whether executive function limitations (specifically difficulties with inhibition) might explain the absence of cumulative culture in chimpanzees (e.g. Davis et al., 2016). We would see such an explanation as falling under the umbrella of the broader EMCC, within this particular category of optimisation of receiver behaviour due to understanding of own knowledge state.

Optimisation of sender behaviour, due to understanding of own knowledge state.

Understanding of one's own state of knowledge can also potentially facilitate cumulative culture by influencing sender behaviour, increasing the likely benefit to others of oneself as a source of social information. Access to one's own level of confidence or uncertainty means that this information can be conveyed to others, alongside actual behavioural decisions. This would then allow others to make more strategic use of that social information, weighting information more heavily when a source reports confidence, or disregarding conflicting information when a source reports high levels of uncertainty. It is this aspect of the EMCC that forms the focus of Shea et al.'s (2014) argument. There is some experimental evidence suggesting that this kind of metacognitive communication does indeed improve the efficacy of social information use. For example Bahrami et al. (2010) studied pairs of participants completing a low-level perceptual decision-making task. When members of a pair had similar visual acuity, they performed better as a pair than they did individually, as long as they were given the opportunity to communicate freely. The authors concluded that this benefit was attributable to the participants providing accurate estimates of their own confidence level within their communication.

Optimisation of receiver behaviour, due to understanding of others' knowledge states

Although there is still considerable debate over whether metacognition relating to one's own mind involves the same processes as metacognition regarding the mind of others (e.g. see Carruthers, 2009), when it comes to *explicit* metacognition, it certainly seems likely that understanding one's own mind, and understanding those of others, are likely to be linked, given the degree of reflective awareness involved (even if the specifics of which understanding comes first may be unclear; see the various models outlined in Carruthers, 2009). Nonetheless, it is worth noting that arguments in support of the EMCC that place the emphasis on explicit understanding of *other's* minds (e.g. Heyes, 2016), are using the term metacognition in a context that, in other areas of the literature, would be regarded as non-standard, and possibly even controversial (e.g. Nichols & Stich, 2003). Furthermore, the literature previously discussed in this review, which relates to the question of whether particular types of metacognition may be unique to humans, may not be strictly relevant in addressing this particular interpretation of the EMCC, since an ability to evaluate one's own confidence or uncertainty may or may not predict one's understanding of others as mental agents. However, if anything, it is probably much easier to make an argument that an explicit understanding of others' minds is restricted to humans. "Theory of mind" (e.g. Premack & Woodruff, 1978), has been a focus of much empirical enquiry and many theoretical analyses in both comparative and developmental psychology, and therefore we do not intend to reiterate findings or conclusions in depth here. But a number of accounts have proposed separate systems for mindreading as a means to reconcile behavioural findings suggesting some tracking of other's mental states in toddlers and animals (e.g. Krupenye et al., 2016; Southgate, Senju, & Csibra, 2007), with consistent evidence that explicit understanding of others' beliefs does not develop until around the age of four in children (e.g. Wellman, Cross, & Watson, 2001), as well as the failure of nonhuman apes in an equivalent nonverbal analogue task (Call & Tomasello, 1999). Apperly and Butterfill's (e.g. Apperly & Butterfill, 2009) two-systems account is perhaps the most high profile of the theories that have been proposed to reconcile these findings (although others exist, e.g. Perner & Roessler, 2012). However, here it suffices to note that it is a relatively widespread view that implicit and explicit tests of understanding of others' mental states may be measuring different processes.

It follows fairly logically, then, to conclude that a System 2, or explicit, understanding of others' mental states might give an agent a significant advantage in their use of social information, allowing them to use this more flexibly and in accordance with the most up-to-date information about who is likely to be an effective model (in line with Heyes's distinction between cook-like and planetary-like decision rules). However, in spite of the convincing rationale for this potential advantage, to our knowledge no empirical studies to date have tested whether an explicit understanding of social sources as mental agents confers benefits over and above implicitly represented strategies. For example, it might be expected that with advancing age, children become capable of using social information in increasingly sophisticated ways, perhaps overriding general purpose biases and heuristics when new information about others' actual knowledge comes to light. For example, recognising that the actions of a single knowledgeable individual are likely to be more valuable than the same number of actions from multiple uniformed individuals. In the absence of such evidence, this particular assumption of how the EMCC might operate may well be plausible, but it nonetheless remains highly speculative.

Optimisation of sender behaviour, due to understanding of others' knowledge states.

Explicit understanding of others' mental states might also bring about changes in sender behaviour, as well as that of the receiver. Even in animals, social learning is not necessarily restricted to the use of inadvertent cues acquired from others as a consequence of incidental observation of behaviours performed only in the interests of the actor themselves. Therefore senders can play an active role in social transmission, and the finer details of how they do so may be significant. In animals, behaviour that functions to teach others has been documented in a number of different species (e.g. meerkats: Thornton & McAuliffe, 2006; ants: Franks & Richardson, 2006). However, as with animals' social learning "strategies", this is a further example where the adaptive function of the behaviour is assumed not to be driven by the agent's *understanding* of that function. This is therefore very different from teaching as it would normally be interpreted in humans, which would generally be expected to implicate some degree of recognition of the part of the teacher, regarding the potential effect of the behaviour on another's knowledge or skill level. The question pertinent to the EMCC then, is whether sender behaviour can facilitate the learning of others much more effectively when senders have an understanding of others' states of knowledge or ignorance. As with the mechanism described in the previous section, a logical argument for this can be constructed with very little difficulty. At the very least, such an understanding would open out the potential contexts within which teaching could occur, whereas (to extend the analogy) "planetary" teaching behaviour would be expected to be restricted to contexts involving an extended selection history (including species-typical behaviours, such as particular predation skills as in meerkats, or well-defined categories of episodic knowledge, such as routes to food sources as in ants). Caldwell, Renner and Atkinson (2017) have previously argued that intentional teaching may be particularly valuable for supporting cumulative culture, since almost by definition cumulative culture is likely to involve novel behavioural variants that are not part of the species-typical repertoire.

In addition to broadening the contexts across which teaching can occur, an understanding of others' minds may also render teaching behaviour far more effective, due to the ability to gauge one's own behaviour in response to the apparent needs of the learner. Teachers can selectively show or perform particular features of what is to be transmitted, with a view to making this maximally informative, based on their own understanding of what might benefit a learner. Furthermore, an understanding of the mind of the learner also allows for adjustments to be made online during teaching, in direct response to the learner's level of success. Mistakes can be corrected, or misunderstandings clarified, and redundancy can be avoided by skipping elements already mastered. A similarly high level of responsiveness might be unlikely in the absence of sensitivity to the meaning of potential cues to knowledge and competence.

There is some literature documenting developmental changes in teaching behaviour in young children which appears consistent with this. Ronfard & Corriveau (2016) studied how children aged between three and five years old taught a game to puppet characters that had demonstrated differing levels of competence. They found that children's ability to monitor the relative accuracy of the puppets improved with age, and that older children tailored their instruction more precisely to the apparent needs of the learner, more often directly addressing the specific errors of individual puppets. This finding is consistent with the assumption that increasing awareness of others' mental states can facilitate transmission by altering sender behaviour.

Optimisation of the sender-receiver interaction due to understanding of minds of self and other

It should be noted that the above categorisations are not intended to be regarded as mutually exclusive; indeed it would be surprising in some cases if they operated in complete isolation from one another. In addition, for each of these two categorisations, it might be expected that benefits arising from the interaction between the two alternative mechanisms could be more than the sum of their individual parts (i.e. the combination of smart behaviour on the part of both sender and receiver, or the combination of understanding of one's own knowledge in relation to others', might be particularly effective in generating cumulative culture). For example, an interaction between an experienced individual (who is motivated to impart their knowledge) and a naïve partner (who is motivated to learn), will likely be most effective when each recognises the other's motivation. In the categorisations detailed above we have only discussed communication in the context of sender behaviour, but communication on the part of the receiver may also have a powerful role to play once there is a mutual appreciation of a shared motivation. This allows the receiver to effectively communicate what the *sender* may need to know, in order to provide the most effective guidance. Clearly, such bidirectional cooperative interactions involve high levels of flexibility in the behaviour of both the sender and receiver, informed by their understanding of both their own, and their partner's, state of knowledge. For the receiver to effectively communicate their needs, this is likely to include not just a representation of the sender's mental state, but a representation of the sender's representation of the receiver's mental state (second order theory of mind, e.g. Perner & Wimmer, 1985), which they are in a position to correct, update, or augment.

In such contexts the breakdown of roles into “sender” and “receiver” becomes significantly less clear-cut. Consistent with this, it has been shown that *both* members of a pair can improve their performance on certain tasks through two-way information sharing (Bahrami et al., 2010). Bahrami et al. found that such benefits only occurred when participants were able to communicate freely, and thus share their confidence levels in addition to their own initial best guess, consistent with the idea that these benefits arise due to metacognitive competence relating to both communicating one’s own level of knowledge, and the interpretation of others’. There may be particular value in being able to interpret another’s knowledge state relative to one’s own, in ways that make each interacting agent simultaneously both a provider of social information (through their influence on another’s success level), and also a beneficiary (through their own improved performance).

It needs to be acknowledged that metacognition is not infallible; people are often under- or over-confident when rating their performance (for example see Metcalfe & Dunlosky, 2008; Miller & Geraci, 2011). However, this shortcoming in self-regulation may be overcome by the shared nature of *explicit* metacognition: Bang et al., (2017) found a collective benefit in making collective perceptual judgements when ‘poorly calibrated’ groups of participants (groups where the more confident members were not the more accurate or skilled members) matched their confidence levels. This may suggest that explicitly sharing metacognitive information about confidence, such as in the scenarios described by Shea et al. (2014), would help to counteract negative effects of poor metacognitive accuracy on personal decision making.

How can the EMCC be tested?

Currently, evidence for the EMCC remains very limited. The accounts proposed by Shea et al. (2014) and Heyes (2016) are built on indirect inference, drawing links between apparent differences in metacognitive awareness in humans versus animals, and the plausibility of metacognition facilitating cumulative culture (sometimes supported by evidence suggesting that outcomes of social learning may be influenced by the availability of metacognitive information). However, in order to effectively evaluate these proposals, more direct evidence is now required. Firstly, there is a need for further empirical evidence that experimentally manipulates the availability of metacognitive resources and/or information, in order to look for direct impact on outcomes of social learning. Secondly, there is a need for studies that fully operationalise cumulative culture, as opposed to studying single transmission events, or looking at interactions only at the level of the dyad. We would expect the combination of these methods to elicit results that showed a reduction in ratchet-like behaviour over generations. That is, methods that are shown to produce accumulation of improvement over generations under normal conditions would no longer show this accumulation when those tasks are carried out in conditions that prevent access to or use of system-2.

Identifying empirical evidence of a causal link between explicit metacognition and cumulative culture is critical in order to establish that the EMCC has some explanatory power over and above other speculative explanations of cumulative culture in terms of other apparently uniquely human features. There are a multitude of features that differentiate humans from other animals, and it is often not difficult to make an argument for the involvement of a particular cognitive or behavioural trait in cumulative culture. What will distinguish such proposals is the availability of empirical evidence that convincingly demonstrates a causal link between the feature or trait in question and outcomes of social learning.

Brain imaging techniques may be used to identify if there are correlations between brain regions activated when using adaptive social learning strategies and capacities for cumulative culture, and those activated when making explicit metacognitive judgements; the EMCC would predict strong correlations in these areas. However, this would not provide direct evidence of a causal link between explicit metacognition and cumulative culture.

Studies which experimentally manipulate the availability of metacognitive resources or information are therefore required to test the EMCC. This is likely to be considerably easier to do for studies investigating the effects of sender behaviour, compared with those focusing on the abilities of the receiver. Accordingly, studies already exist (e.g. Bahrami et al., 2010, described previously) which have experimentally manipulated opportunities for communication, and therefore the potential for sharing metacognitive information, which demonstrate positive impacts on the effectiveness of social information. However, it is much harder to manipulate the extent to which a receiver can employ explicit metacognition in their interpretation of others' behaviour, since it is not possible to simply remove human capacities for metacognition. Nonetheless, we can envisage at least two potentially fruitful avenues of investigation which would allow some insight into the effects of the availability of metacognitive resources on the part of the receiver. The first of these would involve the use of dual-task methods, described previously. The EMCC specifically implicates System 2 involvement, and it should be possible to block or impede the involvement of System 2 using dual-tasks that also place demands on executive function. This could therefore act as a proxy for restricting explicit metacognition directly, with the expectation that reduced access to explicit processing would restrict participants’ ability to interpret (and also share) social information in ways that could be critical for generating cumulative culture.

The expected outcomes of such tasks would be a reduced capacity to make the requisite social learning decisions required for cumulative culture to emerge, and therefore a reduction or absence of ratcheting in tasks that would ordinarily have been shown to produce a ratchet effect in the laboratory.

A further promising approach would be to investigate the effects of developmental changes in metacognitive competence on performance in social learning paradigms, by studying both in young children of a range of ages. The EMCC predicts strong correlations between the emergence of explicit metacognitive competence, executive function capacities and proficiency in strategic social learning tasks.

It should be noted however, that none of the approaches discussed would allow direct manipulation of the involvement of explicit metacognition. Whilst dual-task methods offer potential for experimental manipulation, this would be premised on an assumption that these functioned to block explicit metacognition. Interpretation of such results would therefore be strengthened considerably by the existence of additional evidence validating this assumption, which to our knowledge has yet to be tested. We know of no studies to date which have investigated the involvement of System 2 processing (i.e. as assessed through evidence of interference under dual-task conditions) in explicit reports of metacognition such as judgements of confidence and feelings of knowing (JOC and FOK). Such tests would provide key evidence in evaluating the EMCC, which would inform both theory and method.

Neuroimaging approaches are also somewhat limited in their scope as although they may demonstrate the involvement of brain areas associated with explicit metacognition, they are not necessarily able to show whether participants are unable to produce ratcheting effects in cultural evolution tasks without the involvement of these areas.

Evidence from developmental approaches, whilst offering insights into the potential for cumulative culture both before and after the development of explicit metacognition, would be necessarily pseudo-experimental, involving no attempt to experimentally manipulate the variable of interest, making it much more difficult to identify a causal association. Nonetheless, if relationships were to be found between individual-level measures of metacognitive ability, and individual-level measures of social learning proficiency, especially if these persisted when controlling for age, this would provide fairly convincing support for the EMCC. Certainly, in spite of their respective limitations, both dual-task and developmental approaches, and to some extent neuroimaging approaches, have the potential to provide much stronger support than circumstantial evidence of common exclusivity to humans.

In addition, we would also suggest that a truly robust test of the EMCC would involve laboratory simulation of cumulative culture (e.g. Caldwell & Millen, 2008; 2009), rather than the study of single transmission events, or dyadic interactions. If ultimately the EMCC aims to explain the (group-level) phenomenon of human cumulative culture, it is critical to show that the involvement or otherwise of explicit metacognition does actually impact on the degree to which learning benefits accrue over multiple generations (e.g. Caldwell 2020). Thus, experimental designs using transmission chain or microsociety paradigms (e.g. Mesoudi & Whiten, 2008) would provide a key source of evidence in evaluating the EMCC.

Finally, we also propose that research should be targeted at identifying which of the routes described in Section *How Might Explicit Metacognition Facilitate Cumulative Culture?* account for any link found between explicit metacognition and cumulative culture. We have argued that in principle all are plausible. However, a full account would specify which of these (whether in isolation or combination) appeared to be critical to supporting ratchet effects in cultural evolution.

Conclusion

To date, there is as yet no generally accepted theory explaining the apparent uniqueness of human cumulative culture. The theories recently proposed by Heyes (2016, 2018b) and Shea et al. (2014), which implicate the use of explicit metacognition and System 2 cognition (or Type 2 processes) have the potential to provide a convincing account of distinctively human culture. Here we have used the term the Explicitly Metacognitive Cumulative Culture hypothesis (EMCC), to refer to any view proposing that System 2 processes allow human learners to use metacognition in ways that facilitate social learning. We have also proposed a number of different routes by which System 2 metacognition might have potential to enable cumulative culture, through optimising the behaviour of either the sender or receiver behaviour, based on an explicit understanding of the mental states of either oneself or others.

We have established that, to date, there has been little or no empirical work directly testing these proposals. Indirect evidence is available which provides some support for the view that the implicit metacognitive competence identified in animals depends on processes distinct from explicit metacognition. There is also some support for the view that information transmission may become more effective with increasingly metacognitive competence (at least on the part of the sender), and that having the opportunity to communicate metacognitive confidence levels, in addition to task responses themselves, can also increase benefits of social learning. However, there are significant gaps in the literature, particularly from the point of view of establishing the mechanistic links (see section 2.3) between apparently distinctively human explicit metacognition, and the evolutionary anomaly of cumulative culture. In particular, we see a need for studies involving laboratory tasks which operationalise the group-level phenomenon of cumulative culture, rather than focussing on single transmission events. We have also highlighted that dual-task methods, understood to restrict the use of System 2, have not as yet been exploited within the literature on social learning and cultural evolution, and that these offer a potentially powerful tool for experimentally manipulating the availability of cognitive resources needed for explicit metacognitive processing. We have further suggested that developmental research in human children could shed valuable light on this topic, as children's advancing metacognitive competence offers a natural experiment permitting investigation of the resulting effects on the efficacy of social learning, through increasingly flexible and sophisticated behaviour (whether in the role of sender or receiver).

In conclusion therefore, we consider that the EMCC has considerable promise as a potential explanation for the elaborateness of human culture in relation to the behavioural traditions of other animals. Further research is now warranted in order to test key assumptions and flesh out the details of the links.

Chapter 2: Testing Capacities for Flexible Social Learning Using a Win-Stay, Lose-Shift Paradigm

The Explicitly Metacognitive Cumulative Culture hypothesis (EMCC; Chapter 1) theorises that cumulative cultural evolution may rely on human-unique cognitive capacities for system-2 processes or explicit metacognition. These cognitive capacities are hypothesised to be a requirement for cumulative culture as they allow for efficient and adaptive use of selective copying strategies (Heyes, 2016). The previous chapter suggested that, to empirically test the EMCC, one should aim to examine performance in cultural evolution tasks in populations whose metacognitive skills are expected to be lower than those of typically developed modern human adults. While metacognition cannot be removed from human participants, dual-task paradigms can be used to temporarily reduce access to explicit processes, and by extension metacognition. The studies presented throughout this thesis therefore make use of dual-task paradigms to investigate capacities for cumulative cultural evolution when access to explicit, metacognitive processes are restricted.

The first set of studies presented in this thesis investigate the individual necessary components that are required for cumulative cultural evolution to occur. Participants' abilities to make efficient, flexible learning decisions while under additional cognitive load was examined.

The following chapter is in three parts. Firstly, in chapter 2A the theoretical reasoning behind the methodological approach is explained. I also describe an initial pilot study (P) that I ran in order to establish which dual-task paradigm was the most effective for investigating this research question. In chapter 2B (Experiment 1) and chapter 2C (Experiment 2) I go on to describe further studies run using the paradigm outlined in 2A. To avoid repetition of methods, a detailed methods and procedure is given in chapter 2A. The procedure sections in 2B and 2C are minimal and will largely refer to the procedure given in 2A, with detailed focus on changes from this original methodology only.

At the time of submission, the contents of the following chapter were under revision at *PLOS One* and published as a pre-print on *PsyArXiv*. As submitted papers some of the methods and analyses are separated into supplementary information. This information has been placed back into the text in chronological order to allow for natural flow of the text. Although at the time of thesis submission the final revised article had not been submitted, where possible the suggested reviewer revisions have been incorporated into this chapter. The following chapter is therefore as close as possible to its submitted counterpart.

References for the submitted articles are given below:

Dunstone, J., Atkinson, M., Grainger, C., Renner, E., & Caldwell, C. A. (Under Revision). Flexible social learning strategies are harder than the sum of their parts. *PLOS One.*

Dunstone, J., Atkinson, M., Grainger, C., Renner, E., & Caldwell, C. A. (2020, August 25). Flexible social learning strategies are harder than the sum of their parts. https://doi.org/10.31234/osf.io/xcyu9

As chapter 3 is also a submitted article, references to this chapter given in chapter 3 are cited as 'Dunstone et al., (2020)' rather than 'Chapter 2'.

Chapter 2A: Pilot Testing to Establish a Dual-Task Paradigm

Introduction

The Explicitly Metacognitive Cumulative Culture hypothesis (EMCC; Dunstone & Caldwell 2018) theorises that cumulative cultural evolution may rely on human-unique cognitive capacities for system-2 processes or explicit metacognition. These cognitive capacities are hypothesised to be a requirement for cumulative culture as they allow for efficient and adaptive use of selective copying strategies (Heyes, 2016). To investigate this theory, participants' abilities to make efficient, flexible learning decisions while under additional cognitive load were examined.

Cumulative cultural evolution (CCE) is the process by which cultural traits (including behaviours, artefacts and tools) change over generations to become more effective and beneficial to their users (Caldwell, Atkinson, et al., 2016). This is generally thought to be unique to humans (Dean et al., 2014), and can lead to cultural traits evolving over many generations which could not have been invented by a single individual without access to guidance from, or previous attempts made by, earlier generations.

Although it is not possible to capture all of cumulative cultural evolution in a single cognitive capacity, and few (if any) authors would claim to have found a single trait which was responsible for uniquely human cumulative culture (Mesoudi & Thornton, 2018), many theories suggest potential abilities that could help facilitate this capacity in humans. One current theory for why humans possess this unique capacity asserts that we have the ability to strategically apply social learning in adaptive contexts to ensure optimal performance, retention of beneficial information and innovation away from unhelpful or maladaptive elements of traits. The theory also proposes that this is possible because humans are explicitly aware of their learning strategies, and can strategically apply them in the most appropriate contexts. These learning strategies have been dubbed *explicitly metacognitive* social learning strategies (SLSs) by Heyes (2016, 2018) as they require thinking about mental states (knowledge of one's own knowledge, and sometimes also that of others). Explicitly metacognitive SLSs have been hypothesised to facilitate cumulative culture by allowing learners to direct social learning towards the most appropriate models (Heyes, 2018b), communicating metacognitive perspectives to in-group members (Shea et al., 2014) and making more effective use of information available in the environment (Dunstone & Caldwell, 2018).

Many social learning strategies, also referred to as learning biases or learning heuristics, have been identified (e.g. Kendal et al., 2018; Laland, 2004). However, these tend to refer to processes that are automatic, and often shared by non-human animals (henceforth animals). These SLSs are referred to as ‘planetary’ by Heyes (2016) as they can be described by observers, but generally do not exist in the mind of the animal or human carrying them out (a bit like laws of planetary motion). What distinguishes explicitly metacognitive SLSs is that they are ‘cook-like’ (Heyes, 2016) in that the rules to be followed are in the mind of the actor rather than (or as well as) the mind of the observer, similar to a cook following a recipe. Cook-like SLSs can be explicitly or verbally described by the agent using them.

It is the goal of this research to investigate these hypotheses by experimentally testing whether the capacity for explicitly metacognitive social learning strategies could indeed play an instrumental role in human cumulative culture. This will be assessed by attempting to restrict participants’ access to system-2 cognitive resources (see section *Dual-systems and Dual-tasks*) while they are required to make learning decisions that we assume to be necessary to generate cumulative culture. These learning decisions are argued to make CCE possible by improving upon actions already performed or observed.

Testing Capacities for Cumulative Culture

Multiple definitions of CCE have been given in the literature. Mesoudi and Thornton (2018) have evaluated these multiple definitions and identify four core elements that must be present for a particular case to be classified as CCE:

1) innovation or change to an established or observed behaviour
2) social learning of that behaviour by conspecifics
3) objective improvement of fitness (whether this is true biological fitness or the notion of ‘cultural fitness’) due to the socially learned behaviour change or innovation
4) repeated social learning of the new behaviour so that it outlives the original generation in which it appeared.

These criteria indicate that one necessary (but not sufficient) element of CCE is the capacity to retain the beneficial elements of a behaviour, but ignore or change the non-beneficial elements. At each stage of learning this would require a basic decision to either repeat an action, or to perform a different action which may result in a more beneficial outcome. While in naturalistic settings this may provide a learner with near infinite possibilities of learning choices, this learning decision is essentially a win-stay, lose-shift (WSLS) paradigm taking place: observe successful behaviour (win) and do the same (stay), or observe unsuccessful behaviour (lose) and do something different (shift). This paradigm can therefore be experimentally modelled with a WSLS task.

Capturing individual learner behaviour in this way contrasts with experimental tasks which aim to capture CCE as a population-level process, over multiple learner generations. Methods such as the one used here can however be extended in ways that permit assessment of the potential for CCE in contexts where experimentally examining generational turnover may be problematic (Caldwell et al., 2020).

The validity of the assumption that CCE relies upon the ability to adaptively switch between copying (staying) and exploring (shifting) can be tested by comparing participant performance in mixed blocks of trials where participants must flexibly choose between copying and exploring, with fixed blocks of trials in which participants only ever explore or only ever copy. The expected outcome of such comparisons would be that single trial-type conditions (always explore or always copy) would be rapidly automatized due to the same action being repeated every time, whereas the flexible trial conditions would not to the same extent. This would be taken to reflect reliance on automatic compared to explicitly metacognitive SLSs, respectively (due to highly practised, automatised actions being processed in a system-1 manner (Evans & Stanovich, 2013)).

Dual-Systems and Dual-Tasks

A key aspect of the hypothesis under investigation is that the metacognitive social learning strategies are explicitly applied, and used reflectively, with conscious access by the learner (Heyes 2016, 2018). It is this explicit use that likely differentiates explicitly metacognitive social learning strategies from the social learning strategies, learning biases or heuristics employed by animals such as those outlined in Kendal et al. (2018).

The concepts of implicit and explicit cognitive processes are often discussed without much consideration for specific theoretical definition (Gómez et al., 2017). To avoid this terminological black-boxing, this discussion will use dual processing theories (as summarised in Evans and Stanovich [2013]) to draw the distinction between implicit and explicit processes, with implicit processes being generally confined to system-1, or type-1, processing and explicit processes belonging to system-2 or type-2. The explicit nature of the processes means they rely on the use of executive functions – core cognitive processes that are theorised to be key to system-2 processing. Metacognitive monitoring and control processes have also been shown to have strong correlates with executive function capacities in adults and children (see an extensive review by Roebers (2017) for details).

This explicitly metacognitive hypothesis for the evolution of cumulative culture (EMCC) can therefore be tested with the use of dual-tasks; secondary tasks completed concurrently with a main experimental task that put additional load onto certain executive functions. Due to the expectation of a processing bottleneck when multiple tasks are competing for the same executive function resources (Pashler, 1994), if a concurrent executive function task impacts the ability to effectively apply efficient WSLS strategies, then executive function resources are implicated in the use of those strategies. Executive function dual-tasks have been used, for example, to assess metacognitive involvement in an opt-out task (Coutinho et al., 2015), executive function requirements of Theory of Mind reasoning (Bull et al., 2008) and the implicit or explicit nature of certain processes such as perspective taking (Qureshi et al., 2010).

The following studies will therefore use executive function dual-tasks to attempt to restrict access to metacognitive reasoning while participants complete a WSLS decision making task. The pilot investigates multiple different executive function dual-tasks, performed concurrently with a basic search task. Experiment 1 explores the impact that a dual-task has on participants' ability to apply a flexible WSLS strategy, using executive functions as a proxy for metacognition. Experiment 2 explores the impact of dual-tasks further, and also directly tests for the effect of dual-task interference on explicit metacognitive judgements.

The results of these studies are not intended to demonstrate the presence, or lack of, direct cumulative improvement over many generations. However, they are intended to model whether specific cognitive capacities allow or prevent learning decisions to be made that objectively improve upon observed behaviour.

Executive functions (EF) can be split into 3 main areas: storing, retrieving and updating information from working memory; inhibiting automatic responses to stimuli; and task switching (Miyake et al., 2000). Although correlated, these are distinct cognitive functions that affect different aspects of behaviour, as shown by confirmatory factor analyses (Miyake et al. 2000; Spiess, Meier, & Roebers, 2015). When designing dual-task methods to impede executive functions it would therefore be unwise to simply pick a dual-task without evaluating how the different EFs might interact with the main task.

Before a more substantial sample was recruited, multiple pilot tasks were tested on a smaller number of participants. The purpose of doing this was to compare the effect of multiple different executive function dual-tasks on a binary search task (see below), which required the use of a WSLS strategy in order to perform correctly. This adopts a methodology similar to that used by Bull, Phillips and Conway (2008), who tested multiple different executive functions (inhibition, switching and updating) with two ToM tasks, as well as a range of controls, to systematically assess the impact of executive function on ToM processing.

The aim of this pilot was to establish which dual-task paradigm (if any) caused the greatest interference with reaction times (RTs) in the binary search task, with the intention of using this task combination with a larger sample size. Comparisons are drawn between change in RTs relative to no dual-task, and to a control dual-task that is equally physically demanding but with no EF requirement. This is to ensure that dual-task interference is only attributed to competing EF demands where appropriate. The task(s) that elicited a significantly slower reaction time when under EF load than when in control conditions would then be considered for further testing.

Methods

Participants:

Participants were recruited at the University of Stirling and took part in partial fulfilment of a course requirement to participate in research studies. Two participants chose to receive cash remuneration instead and were paid at a rate of £5/hour. Sixty three participants took part (eight male; 55 female, mean age: 20, range: 17-45). Of these, four were excluded due to computer errors meaning they could not finish the task. Seventeen of the total participants completed the training phase of the experiment but did not score above the inclusion threshold for the full testing phase and so left the experiment after they had completed training (see *procedure* for details). In total 42 participants (seven male; 35 female, mean age:20.1, range: 17-40) completed the full procedure. Only data from this group of 42 has been included in the below analysis. All participants had normal or corrected to normal vision and hearing. Participants all gave written consent to take part and were aware that they were free to withdraw from the study at any time. Ethical approval for the study was given by the University of Stirling General University Ethics Panel (reference GUEP 111).

Apparatus:

Participants were tested using a desktop computer running Windows 8 with a standard mouse, a Black Box Toolkit 4 button response pad button box and Sony MDR-Pro over-ear headphones.

Task Design:

All tasks were written and run in Psychopy version 1.84.2. Code to run all of the tasks described below can be found here: https://osf.io/bfg3x/.

All participants completed the same main task, the binary search task detailed below. The binary search task was completed either on its own (baseline condition), or with one of five different dual-task methods (four audio, one visual), designed with the intention to place demands on executive function. Individual dual-tasks were between-subjects variables, with each participant only completing one dual-task. Each task designed to impair executive function (EF-task) was paired with a matched task that required the same motor response to be performed but with a substantially reduced executive function load (control task). This was a within-subjects variable, with each participant completing an EF-task and a control task, counterbalanced for order. Participants in the baseline condition completed only the search task. The baseline condition was not included as an additional within-participants variable in order to avoid fatigue and/or practise effects from participants having to complete a large number of trials of the same main task.

Brief descriptions of each dual-task are given in table 1, with detailed descriptions given below.

Binary Search Task

This was a simple two alternative forced-choice (2AFC) binary search task, intended to assess how quickly participants could use vicariously presented information to make decisions. Participants played as a penguin avatar, and the task required participants to find desired hidden objects (fish) behind coloured shapes presented on screen, and to avoid undesired objects (sharks). Participants played against the computer (represented as a robot avatar) to find as many fish as possible.

Participants were shown an information trial in which, after a brief pause, the computer selected a subset of the shapes presented on screen and displayed the hidden object(s). After a second pause the participant was then presented with the same shapes again and was required to make their own selections with the goal of finding the desired object(s). The information trial was either successful (fish were found) or unsuccessful (sharks were found). Participants were always instructed to find the fish and avoid the shark, and were awarded one point on trials in which they correctly selected the stimuli that displayed the fish. Reward location was fixed, so to perform optimally participants needed to employ a win-stay, lose-shift (WSLS) strategy. The task therefore models the most basic requirement for cumulative cultural evolution discussed above (in *Testing Capacities for Cumulative Culture*).

Detailed task description of Search Task

For pilot (P) & E1 trials the task was a binary choice, with the computer and participant making one selection from two shapes. In E2 two selections were required from four shapes. On each trial, after a 2 second, ready screen participants were shown an information trial in which the computer selected one (P & E1) or two (E2) of the shape stimuli on the screen to reveal the object(s) underneath. The colour and shape of stimuli were pseudo-randomly determined. The number of sides was randomly selected from a choice of 3-6 (P & E1; for E2 the choice also included a circle), with the caveat that on each trial no two shapes could have the same number of sides. The colour of the shapes was randomly selected from a pre-specified list of 16 colours. The colours were chosen from a colour chart (such as the example in figure 1), with the aim of choosing colours that appeared far enough apart on the chart that the colours were immediately distinguishable. On each trial no two shapes could be the same colour.

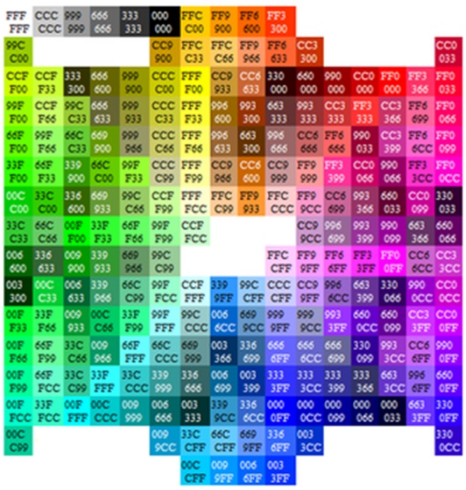

Figure 1: Example colour chart used to select colours that were as visually different to each other as possible

Table 1: Brief descriptions of each EF dual-task, and the matched control dual-task. Task 1 is based on the dual-task used by Qureshi, Apperly, & Samson (2010); tasks 2-4 are based on the dual-tasks used by Bull, Phillips and Conway (2008), and task 5 is based on the task used by Coutinho et al. (2015) which was argued to inhibit metacognitive responding in humans and monkeys.

Dual-task	EF Task	Control Task
1: *Incongruence*	Participants listened to a series of auditory tones (either 1 or 2) through headphones and responded with the opposite number (i.e. either 2 or 1) of mouse clicks.	Identical to the EF task, but participants responded with the same number of clicks as tones.
2: *Inhibition* ***(withholding)***	Participants listened to a series of auditory tones through headphones and responded to each tone with a mouse click, except under certain specified conditions when they were required to withhold their response.	Identical to the EF task, but participants responded to every tone with a mouse click with no conditions requiring withholding a response.
3: *Switching*	Participants listened to a series of auditory tones through headphones and responded to each tone with a pre-specified mouse response. After an auditory cue they were required to switch to a different pre-specified mouse response.	Identical to the EF task, but participants would continue with the same mouse response for the duration of the task.
4: *Updating*	Participants listened to a series of auditory tones, in sets of either 1 or 2 tones, through headphones. They responded by clicking the mouse once more than the number of tones heard in the previous set.	Identical to the EF task, but participants responded to the number of tones heard in the current block.
5: Working Memory	Participants were presented with two numbers on screen, one of which was clearly much larger in font size than the other, and one of which was larger in numerical value than the other. Participants were shown the two numbers briefly, and asked to remember them until cued to respond with either the number that is large in font size, small in font size, large in value, or small in value.	The same format as the EF task, but participants were just shown fixation crosses rather than asked to remember numbers. Instead of being asked to recall values participants were asked a very simple arithmetic question.

Shape colours were selected from a limited list rather than randomly generated to prevent colours that were technically different but functionally indistinguishable from being used on the same trial. Colours were also intentionally chosen to be salient against the task background (see figure 2). The information trial was either fully successful (a fish was found; two fish in E2) or fully unsuccessful (a shark was found; 2 sharks in E2). In E2 there was also the option for a trial to be partially successful (one fish and one shark were found). For the selective copying conditions, in P & E1 there was a 50% chance of showing a successful or unsuccessful information trial, and whether the computer selected the shape on the left or the right of the screen with an equal balance of all possible success and side options. In E2 there was an equal balance between fully successful, fully unsuccessful and partially successful trials with an equal balance of all possible stimulus pairings. There was therefore an equal balance of successful and unsuccessful selections on each side of the screen. These were presented in a fully random order. In the WS and LS conditions all trials were fully successful or fully unsuccessful.

The information trial was displayed by either a fish or a shark appearing in the centre of the shape stimulus after a random delay of between 800-2400ms (in line with previous reaction time literature, e.g Cook & Bird [2011]). After each information trial the visual counter of the computer's score next to its avatar would increase by the number of fish revealed in that trial (0-2). When the task was not intended to impose a memory load (Pilot, E1, and some conditions of E2) all shapes, the object(s) found and their location would remain visible to the participant for the duration of the following test trial. When the task was intended to impose a memory load (some conditions of E2) the objects found during the information trial would disappear as the test stimuli appeared.

Participants then completed a test trial by making their own selection from the same shape stimuli which were presented again directly underneath the information stimuli. In P & E1 the test stimuli appeared as soon as the information trial selection had been displayed. In E2 the test stimuli appeared after a random delay between 800-2400ms, to allow an equal delay when the memory load was present or absent. The participants' aim was to find all the fish as quickly as possible. Responses were collected using a Black Box toolkit 4 button response pad. In P and E1 participants were instructed to only use the two outermost buttons, with the left button corresponding to the left stimulus and the right button corresponding to the right stimulus. In E2 the buttons left to right corresponded to the stimuli left to right. After selecting a shape, either a fish or a shark was displayed in the centre of that shape stimulus. If successful, the visual counter of the participant's score next to the penguin avatar would increase by one. In E2 a second selection would then need to be made. On partially successful trials the reward structure was fixed so that if employing a correct WSLS strategy there was a 50% chance of finding a fish or a shark in either of two available locations. This screen was displayed for a further 1s, and then participants were given immediate feedback about their performance on a screen that told them if they had been successful, and if so told them their reaction time (in seconds) and a running total of their score so far. If unsuccessful they were again given immediate feedback and a reminder that they were looking for fish and aiming to avoid sharks, but were not given their reaction time or a running total of their score. This screen was displayed for 2s and then either the working memory test trial (in P, working memory condition) or a new trial of the choice task would commence. In E2 a 3s time limit was imposed on the first response. If a selection was not made within 3s the trial would end and a message would display on screen reading 'trial missed, response must be given within 3 seconds' and one point would be removed from the total.

Detailed description of each distractor task

Dual-task 1: Incongruence

EF task: Participants listened to a series of auditory tones, in sets of either 1 or 2 tones, through headphones and were instructed to respond to each set of tones with the opposite number (i.e. either 2 or 1) of mouse clicks. Tone sets were played every 3 seconds. There was a 50% chance of a single tone and a 50% chance of a double tone, presented in a fully random order.

Control task: Identical to the EF task, but participants were instructed to respond with the same number of clicks as tones.

This is similar to the executive function distractor task used by Qureshi, Apperly and Samson (2010).

Dual-task 2: Inhibition

EF task: Participants listened to a series of auditory tones, in sets of either 1, 2 or 3 tones, through headphones and were instructed to respond to each tone with the same number of mouse clicks as tones, unless the number of tones matched a pre-specified number for which they were instructed to withhold their response. Whether this pre-specified number was 1, 2 or 3 was randomly determined at the start of the testing session. Tone sets were played every 3 seconds. The number of tones played was split at 40%, 40%, 20% with the number that should have responses withheld presented less often.

Control task: Identical to the EF task, but participants were instructed to respond to every tone with the same number of mouse clicks as tones, with no numbers requiring withholding a response.

This is similar to the 'inhibition' executive function distractor task used by Bull, Phillips and Conway (2008).

Dual-task 3: Switching

EF task: Participants listened to a series of auditory tones, in sets of either 1 or 2 tones, through headphones and were instructed to respond to each tone with the same number of mouse clicks as tones. After an auditory cue they were required to switch to responding by clicking once more than the number of tones heard. Participants were instructed to switch between clicking the same number of times as the tones heard and clicking once more than the tones heard each time they heard the auditory switch cue. Tone sets were played every 3 seconds. There was a 50% chance of a single tone and a 50% chance of a double tone, presented in a fully random order. Switch cues were played at an average rate of 1 cue for every 5 tone sets played, presented at pseudo-random intervals: the series of tones could not begin with a switch cue, and a switch cue could not be played immediately after another switch cue.

Control task: Identical to the EF task, but participants were instructed to continue to respond to each tone with the same number of mouse clicks as tones heard for the duration of the task and no switch cues were played.

This is similar to the ‘switching’ executive function distractor task used by Bull, Phillips and Conway (2008).

Task 4: Updating

EF task: Participants listened to a series of auditory tones, in sets of either 1 or 2 tones, through headphones. They were instructed to respond by clicking the mouse once more than the number of tones heard in the *previous* set i.e. if the tone sequence 2, 1, 2... was heard the participant should respond with the click sequence 0,3,2… (0: as the first set has no predecessor, 3: in response to the first set which contained 2 tones, 2: in response to the second set which contained 1 tone). Participants were instructed not to click in response to the first tone. Tone sets were played every 3 seconds. There was a 50% chance of a single tone and a 50% chance of a double tone, presented in a fully random order.

Control task: Identical to the EF task, but participants were instructed to respond by clicking once more than the number of tones heard in the current block.

This is similar to the ‘updating’ executive function distractor task used by Bull, Phillips and Conway (2008).

Task 5: Working Memory

EF Task: This dual-task was presented between trials of the main task. At the start of each trial participants were given a visual instruction of "remember the numbers". They were then presented with two numbers on screen, one of which was clearly much larger in font size than the other, and one of which was larger in numerical value than the other. On 50% of trials the large value number was in larger font, and on the remaining 50% the large value number was smaller in font size. These trials were presented in a fully random order. The numbers were pseudo-randomly selected from the digits 2-9, with the constraint that the numbers could not be the same. The numbers were displayed for 3s and then masked with a white square for 0.5s. After each trial of the main task participants were presented with a screen with 2 unfilled white squares and a question in the middle. This question would ask which number was larger/smaller in size or which number was larger/smaller in value and the participant was instructed to select either the left or the right square to relate to the numbers presented at the beginning of the trial. The choice of question was randomly picked from a possible 4. There was no time limit for responses to this question. After making a response participants were given immediate feedback about whether they had been correct.

Control task: The format of the task was the same. However, instead of being shown numbers to remember at the beginning of the task participants were instructed to look at 2 fixation crosses, placed in the same screen position as the numbers in the EF Task. After each trial of the main task participants were then given a very simple numeracy task and had to say which of 2 numbers presented on the screen was larger or smaller in value. These numbers were randomly picked from a list of large and small numbers with no overlap between the lists. All numbers were single digits between 1-9. The question was a random choice of the 2 questions (larger or smaller).

This is modelled on the executive function distractor task used by Coutinho et al., (2015).

Procedure:

Participants were tested individually, although the majority of participants took part in the same room and at the same time as another participant. When two participants took part at once both participants were verbally instructed by the experimenter that they were acting entirely independently and were not competing with one another. Participants were facing away from each other so could not see each other's screens. They were either both taking part in silent tasks or were both wearing headphones so audio distraction from the other participant was minimal. Both participants began the task at the same time and if one participant finished in a quicker time than the other they were escorted quietly out of the testing room to ensure the remaining participant was not disturbed.

Before beginning trials participants received detailed on-screen task instructions which included examples of the audio sounds they would be exposed to and the visual stimuli they were looking for. Participants then completed 4 practise trials of their audio dual-task on its own, followed by 2 practise trials of the binary search task. Participants in the *baseline* condition completed 2 trials of practise of the binary search task only, and participants in the *working memory* dual-task condition completed 2 trials of the binary search task and dual-task together. After this brief practise, participants were then required to complete 16 trials of pre-test training in each condition (EF load and control), with the order of conditions counterbalanced across participants (*baseline* participants did 2 blocks of 16 trials of the same task) in full dual-task conditions. For the audio dual-tasks the number of trials was determined by the time taken to complete the binary search task. For the working memory dual-task participants completed 16 trials per block.

To ensure full focus was given to both tasks an inclusion criterion of 75% accuracy was set for both tasks, with a larger participation reward available to participants that completed the full study. Participants that scored below 75% accuracy in the training round (averaged over both blocks) in one or both tasks left the study at this point and received a smaller reward (a reduced participation fee, or fewer research participation tokens, relative to participants completing the full study). Participants were made explicitly aware of this inclusion criterion when signing up to take part in the study. They were also reminded of this again when giving consent to take part, and then once more as part of the written instructions for the study.

Participants who passed the training round then completed a further 48 trials of testing in each condition in full dual-task conditions (*baseline* participants did 2 blocks of 48 trials of the same task). For the audio dual-tasks the number of trials was determined by the time taken to complete the binary search task. For the *working memory* dual-task participants completed 48 trials per block.

Results

The aim of this study was to select a task or tasks that proved effective in restricting participants' ability to effectively use social information to make decisions. The results summarised here are therefore fairly brief, and focus only on whether significant differences in the response times in the binary search task were found between the baseline condition and each dual-task condition, and between the EF and Control blocks of each dual-task condition.

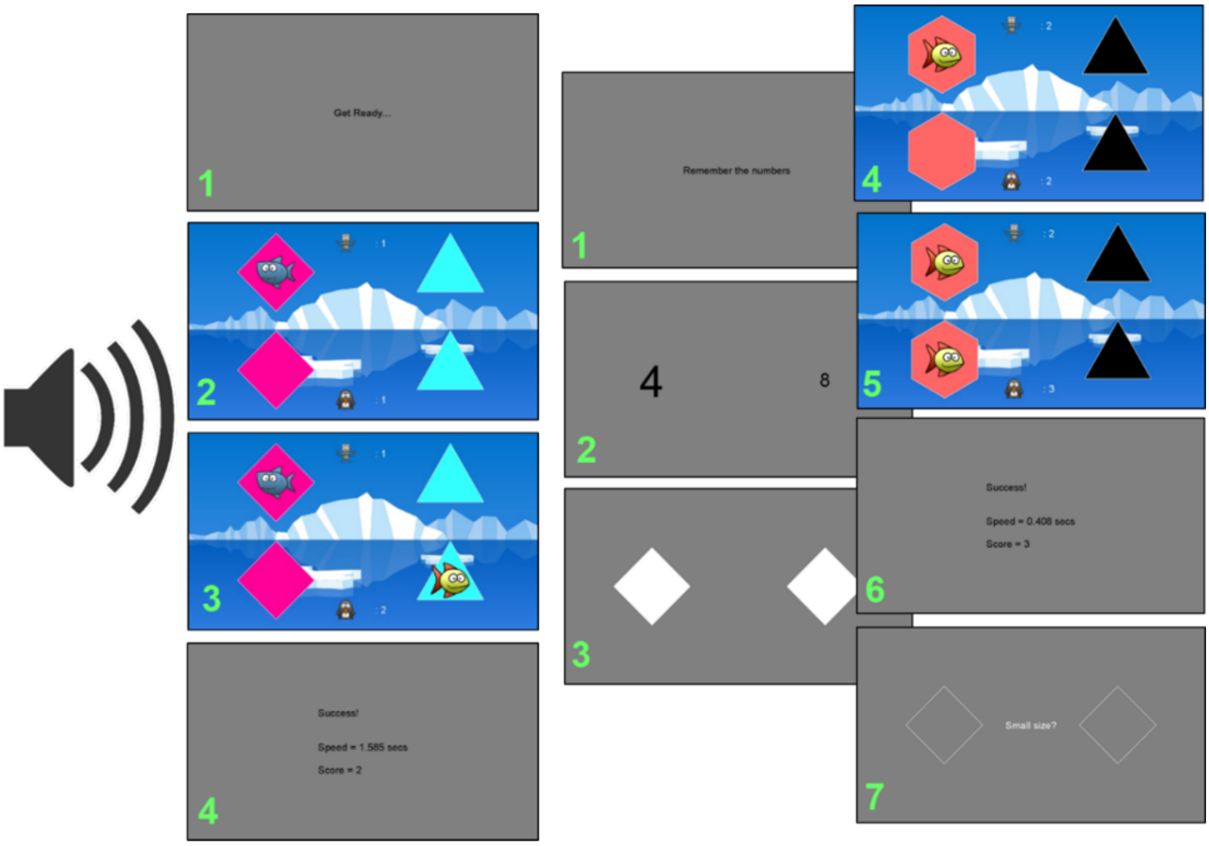

Figure 2:

LEFT: *Example trial of the binary search task being completed with an audio dual-task. This is an example of an unsuccessful information trial (shark is displayed).* ***RIGHT:*** *Example trial of the binary search task being completed with the working memory dual-task. This is an example of a successful information trial (fish is displayed). In both cases the test trial is successful (stimulus showing a fish is selected).*

The study aim was to recruit 8 participants into each dual-task condition, and 8 participants into the baseline. This was achieved for all tasks apart from *updating*. This was due to large numbers of participants in the *updating* condition scoring lower than the inclusion criterion during training (12 out of 14 tested without software issues), meaning data collection was halted for this condition. The results summarised below have therefore been conducted with the *updating* condition omitted.

The variable of interest was the reaction time of all responses given in the main task. Response time was used rather than accuracy as accuracy was predicted to be at ceiling levels in the task. Figure 3 shows the mean RT in the binary search task for each dual-task type. Accuracy in this task was approximately at ceiling across all task conditions (accuracy range: 96.4%-98.9%), so is not analysed here.

Outliers

Overall 186 outliers were removed from the data for very long or very quick reaction times. Outlier removal is important for these data in order to screen out responses that were made so quickly they may reflect an accidental button push or repetitive pressing without paying attention to the task, or so slowly they may be due to external distraction unrelated to the task. Each outlier represents a single trial. A relatively broad inclusion criterion (3 x Median Absolute Deviation (MAD) from the mean [Leys et al., 2013]) was used to ensure genuinely long reaction times caused by dual-task interference, which were the expected outputs of the study, were not artificially trimmed. Outliers were removed per participant, per task type in line with recommendations from Ratcliff (1993) regarding outlier removal for reaction time datasets with high inter-participant variation. 182 outliers were removed from the upper end of the response distribution and 4 were removed from the bottom end. This represents 4.7% of the total data.

Difference from baseline

Each task's difference from the baseline task was calculated using a linear mixed effects model for each task with one fixed effect of block condition (baseline, control or executive function block) and one random effect of participant ID. *P*-values were estimated from the resultant t-statistics with degrees of freedom being the number of observations minus the number of fixed parameters in the model (Baayen et al., 2008). For the *incongruence* and *withholding* tasks the model was not significantly better than a null model (incongruence: χ^2=0.288, p=.866; withholding: χ^2=5.18, p=.075) and the tasks showed no significant differences from the baseline condition in either block. Both the *switching* and *working memory* task models were significantly better than the null model (switching: χ^2=43.6, p<.001; working memory: χ^2=21.2, p<.001) and both tasks showed significantly longer reaction times in each block compared to the baseline task. The results of each model are presented in table 2.

Table 2: Difference between average reaction time in milliseconds of each task from the baseline condition (details of models given in brackets)

TASK	CONTROL DIFFERENCE FROM BASELINE (MS)	EF DIFFERENCE FROM BASELINE (MS)
INCONGRUENCE	9.05 (SE=27.0, t(1467)=0.335, p=.738)	5.772 (SE=27.0, t(1467)=0.214, p=.831)
WITHHOLDING	8.88 (SE=25.2, t(1473)=0.352, p=.25)	24.8 (SE=25.3, t(1473)=0.980, p=.327)
SWITCHING	22.28 (SE=24.7, t(1453)=0.901, p<.001)	80.8 (SE=24.7, t(1453)=3.27, p<.001)
WORKING MEMORY	45.7 (SE=31.5, t(1490)=1.45, p<.001)	74.9 (SE=31.5, t(1490)=2.38, p<.001)

Difference between control task and EF task

The difference between the control and executive function blocks of each task were calculated using a linear mixed effects model for each task with one fixed effect of block condition (control or executive function block) and one random effect of participant ID. For the *incongruence* task the model was not significantly better than a null model (χ^2=0.179, p=.672) and the task showed no significant differences between the executive function and control blocks. The *withholding, switching* and *working memory* task models were all significantly better than the null model, although *withholding* was only marginally better (withholding: χ^2=4.10, p=.043; switching: χ^2=26.6, p<.001; working memory: χ^2=16.3, p<.001) and all three tasks showed significantly longer reaction times in the executive function block compared to the control block (this is again marginal for the *withholding* task). The results of each model are presented in table 3.

Table 3: Mean RT in ms to 3.s.f (SD in brackets) of each task block. Difference column shows results of linear mixed effects models testing if blocks were significantly different to each other.

TASK	CONTROL	EF	DIFFERENCE
INCONGRUENCE	506 (113)	502 (114)	b=-3.29, SE=7.80, t(721)=-0.422, p=.673
WITHHOLDING	505 (112)	521 (111)	b=15.9, SE=7.83, t(727)=2.03, p=.043
SWITCHING	518 (125)	577 (177)	b=58.6, SE=11.3, t(708)=5.20, p<.001
WORKING MEMORY	542 (107)	572 (124)	b=29.2, SE=7.18, t(744)=4.06, p<.001

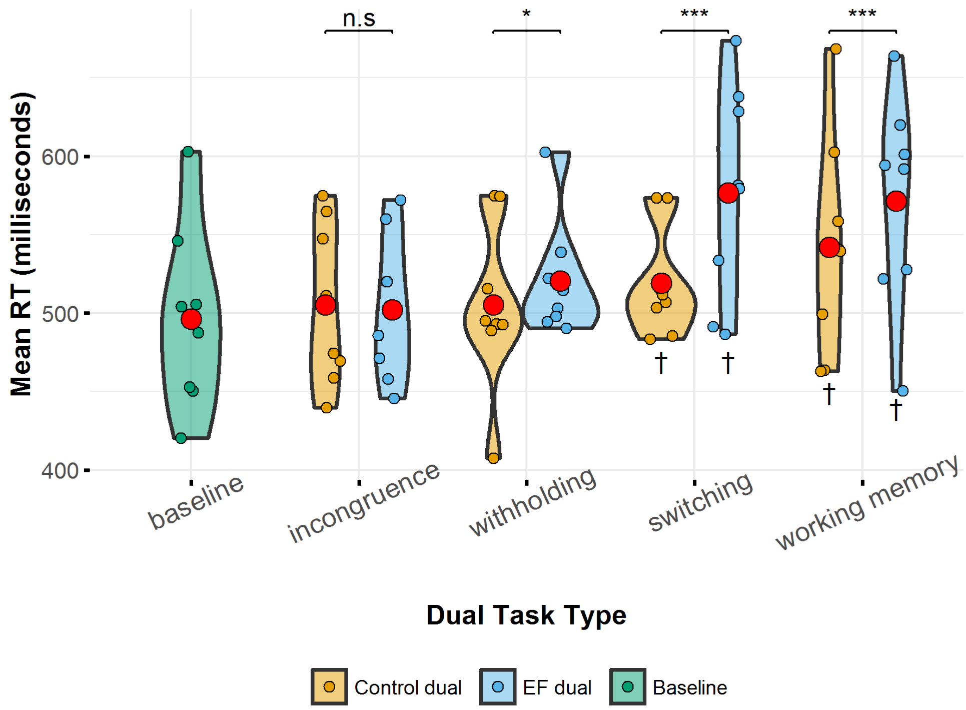

Figure 3: RT for each task type split by dual-task block. Red dots indicate mean for each block. † indicates block is significantly different to the baseline condition. Brackets indicate significance of difference between control and EF block within a task.

Discussion

The results presented above suggest the *switching* or *working memory* tasks could both be used as viable dual-tasks to experimentally restrict participant ability to make rapid decisions using social information, as they both produced a significant difference between the control and executive function blocks, and the executive function blocks were significantly different from the baseline condition. This slowing down of responses in these conditions indicates that these executive functions are involved in making inferences with social information, although it should be noted that this task did not involve genuine social information. Both tasks also showed significant differences between the control block and the baseline task. This is to be expected, as doing anything at the same time as the main task is likely to increase reaction times, and highlights the importance of including a control dual-task as well as a baseline and a test condition. Although there was also a marginal significant difference between control and EF blocks in the *withholding* condition, as neither block of this task presented a difference from the baseline condition this would not be an appropriate dual-task to continue to use within this testing paradigm. Previous literature implicates *updating* and *inhibition* most strongly in metacognition tasks (Roebers 2017). The different result found here may highlight the importance of empirical testing to establish sound methods for data collection. However, this could also indicate that although executive functions are playing a formative role in the learning decisions required for CCE to occur, metacognition may not be implicated.

The *incongruence* task showed no significant RT difference, when completed under control and executive-function tasks, suggesting the executive function load these tasks placed on participants was not large enough to impair performance on the binary search task, or that the type of EF required to complete the *incongruence* task is not required to make WSLS decisions.

As the data from the *updating* task are so limited, a definitive conclusion about this task type cannot be drawn. It is clear from this specific set of results that the task in the way it was designed for this study was too challenging for the majority of participants. However, that does not mean tasks utilising the updating component of executive function would never be useful as dual-tasks.

While *switching* and *working memory* both showed significant results, only 1 dual-task method at a time was sought to take forward for further testing. This was to ensure all participants will have their executive functions impaired in the same way, to ensure comparable data across the entire task. For the subsequent studies presented here just the *switching* task is used. This task was chosen as it was faster to complete than the *working memory* task, allowing for more trials to be completed within the same testing period. The *working memory* task may be used in future testing.

Chapter 2B: Investigating Dual-Task Interference on a Simple Selective Search-Task - Experiment 1

Dual-task interference may be observable due to a reduced capacity to apply sustained attention to a rule that is currently required and change the response type accordingly, i.e., the reduced ability to re-evaluate on each trial whether the presented information should be copied or not. This would predict significant interference from a dual-task when a selective strategy was required, but not when the same response type was required for every trial. System-2 processing is therefore implicated in the use of a selective response strategy, analogous to a basic selective SLS.

Experiment 1 therefore compared mixed blocks of testing, where participants had to apply a selective response strategy (as used in the pilot), with blocks requiring single response-types (always copy or never copy). The always-copy and never copy groups showed participants information trials which were always successful or always unsuccessful, respectively, and therefore required responses of repetitive, reinforceable behaviours that do not require an SLS to be employed. Both always-copy and never-copy controls were included to differentiate between the difficulty increases of switching compared to copying, and flexibly copying compared to either copying or switching.

As only the selective blocks required participants to continuously update and switch between response strategies, it was hypothesised that in this condition participants' performance would be negatively impacted by the dual-task, whereas interference would be minimal when participants were just using a repeated application of a single rule. It should be noted that this task, particularly in the single response type blocks, was designed to be simple for participants to do. This was partly to ensure that any effects found were not caused by confounding effects of main task difficulty or other cognitive factors that have not been strictly controlled for. The two control conditions were intended to capture a level of decision making where executive functions are *not* expected to be implicated. Pilot testing (see previous section, 2A) demonstrated that dual-task interference *did* have a significant negative effect on responses in selective blocks.

Methods

Participants:

Participants were recruited at the University of Stirling and took part in exchange for research participation tokens which were required for course completion. Twenty-four participants received cash remuneration instead and were paid at a rate of £5/hour rounded up to the nearest pound. 166 participants took part (57 male; 109 female, mean age: 21.2, age range: 16-51). Of these, 45 participants completed the training phase of the experiment but did not score above the inclusion threshold of 75% accuracy (see *procedure)* for the full experimental phase and so left the experiment after they had completed training. A further one participant was excluded because during debrief they informed the experimenter that they had not understood the task and had not been completing it correctly. In total 120 participants (47 male, mean age: 21.1, range: 16-49) completed the full study and were included in the dataset reported in the analyses below. All participants had normal or corrected to normal vision and hearing. Participants all gave written consent to take part and were aware that they were free to withdraw from the study at any time. Ethical approval for the study was given by the University of Stirling General University Ethics Panel (reference GUEP 111A).

Apparatus:

Participants were tested using a desktop computer running Windows 8 with a standard mouse, a Black Box Toolkit 4 button response pad button box and Sony MDR-Pro over-ear headphones.

Procedure:

Participants completed the same binary search task as used in the pilot, in conjunction with the *switching* dual-task only, as well as the matched control dual-task. The procedure remained the same, although more test trials were completed and participants were now split into three groups, each of which required the use of a different copying strategy in order to be successful. Participants again completed 16 trials of training prior to the full study.

Participants who passed the training round then completed a further 72 trials of testing in each condition in full dual-task conditions. For the audio task the number of trials was determined by the time taken to complete the binary search task.

Participants were randomly split into 3 equal groups which determined which information trial types they received (the training trials for all participants were set as though they were in the selective copying group):

WSLS (Selective Copying group): In this group there was a 50% chance of observing a successful or unsuccessful information trial, and whether the computer selected the shape on the left or the right of the screen. There was therefore an equal balance of successful and unsuccessful selections on each side of the screen. These were presented in a fully random order.

WS (Copy-All group): In this group 100% of information trials were successful. There was still an equal balance of whether the shape on the left or the right was selected, presented in a fully random order.

LS (Avoid-All group): In this group 100% of the information trials were unsuccessful. There was still an equal balance of whether the shape on the left or the right was selected, presented in a fully random order.

Overall, 45 participants were excluded based on their training round score (WSLS group: 18, WS group: 16, LS group: 11).

We predicted that reaction times would be slower in the block where participants completed an EF dual-task, compared to the control task block. We also predicted that the impact of the dual-task would be greater in the WSLS, (selective copying) group than in either of the WS or LS groups. No specific predictions were made about the difference in RTs between the WS and LS groups.

Results

Outliers

Overall 1279 outliers were removed from the data for extremely long or extremely quick reaction times (taken as 3 x MAD from the mean, see *Outliers* in 2A). 1211 outliers were removed from the upper end of the response distribution and 68 were removed from the bottom end. This represents 7.4% of the total data.

Analysis

Binary Search Task

Figure 4 shows the overall reaction time for each strategy group by block condition. Accuracy in the search task was at ceiling (accuracy range: 98.1%-98.6%). Mean reaction times for each strategy group and block condition are given in table 4.

Table 4: Mean RTs (milliseconds, SD in brackets) for each strategy group and block condition in E1

GROUP MEMBERSHIP	BLOCK CONDITION	MEAN (SD) RT (ms)
WSLS	EF Dual	564 (206)
	Control Dual	518 (154)
LS	EF Dual	512 (200)
	Control Dual	462 (144)
WS	EF Dual	467 (172)
	Control Dual	429 (139)

Reaction times were analysed using a linear mixed effects model with fixed effects of: group (WSLS, LS or WS), block condition (EF dual-task or control) and item number (within block), and their interactions. Participant ID was included as a random effect. The WSLS group in the control condition was taken as the baseline and *p*-values were estimated from the resultant t-statistics with degrees of freedom being the number of observations minus the number of fixed parameters in the model (Baayen et al., 2008).

The model was a significantly better fit than the null equivalent ($\chi^2(11)=531$, $p<.001$). There were significant effects of group membership (LS faster than WSLS: b=-53.7, SE=22.3, t(15994)=-2.41, *p*=.032; WS faster than WSLS: b= -77.6, SE=22.3, t(15994)=-3.48, *p*=.001; both corrected for multiple comparisons), executive function block condition, with the EF block slower than the control: (b=52.6, SE=7.80, t(15994)=6.75, *p*<.001) and item number, with reaction time decreasing as trials increase: (b=-0.438, SE=0.132, t(15994)=-3.31, *p*<.001). A post-hoc Bonferroni-Holm correction of all pairwise comparisons of group membership indicated that the overall difference between the single response-type groups, LS and WS, was not significant (b=-23.9, SE=22.3, z=-1.07, p=.284). The interaction between group membership and block condition was not significant for either single response-type group compared to the selective copying group (LS: p=.483, WS: p=.425). Taken together these results indicate that the selective strategy condition of the binary choice task was more challenging than either of the single-response conditions and, although the *switching* task had a negative impact on response times in each group, it did not have a *greater effect* on the selective strategy group, relative to the other groups.

Response time after successful or unsuccessful information trials was analysed using a linear mixed effects model with one fixed effect of success in the information trial and one random effect of participant ID. This model was a significantly better fit than the null equivalent ($\chi^2(1)=79.4$, $p<.001$). Taking all groups together, responses were significantly faster after successful information trials (wins) than after unsuccessful information trials (losses) ($b=-34.8$, $SE=3.9$, $t(15999)=-8.92$, $p<.001$).

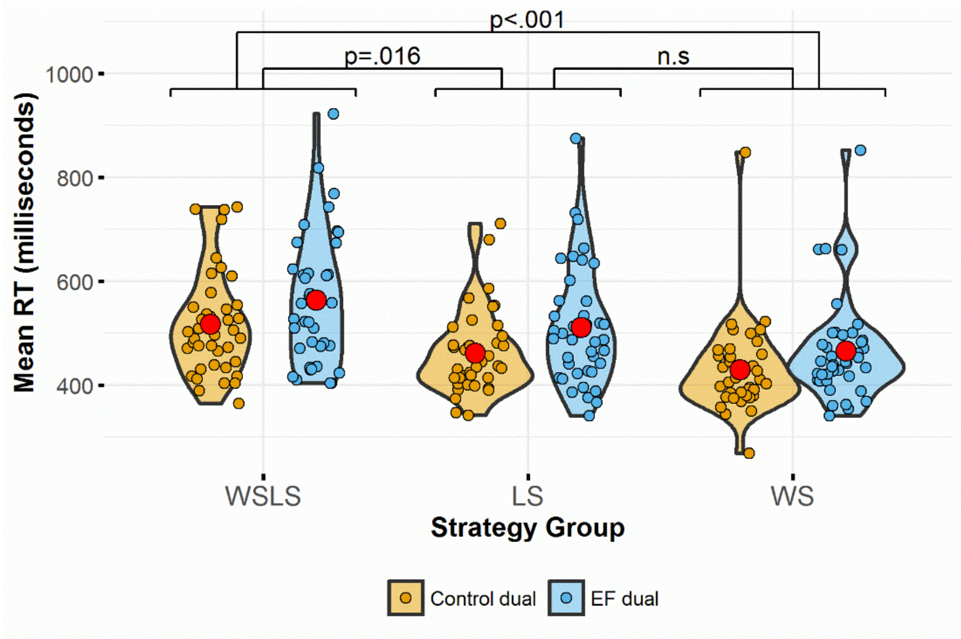

Figure 4: Reaction times in the binary search task for each group and block condition. Red dots indicate the mean in each block. Labelled brackets indicate significance level of comparison between strategy groups overall. The overall RT in the selective copy group (WSLS) was significantly slower than both the always copy (WS) and never copy (LS) groups. There was no difference overall between the LS and WS groups. The difference between the control and EF block condition was the same for each strategy group.

Looking at the selective strategy group only, a linear mixed effects model was constructed with fixed effects of block condition, success in the information trial and their interactions. The control block was taken as the baseline. Participant ID was included as a random effect. This model was significantly better than the null equivalent ($\chi^2(3)=204$, $p<.001$). Reaction times were significantly faster after successful than after unsuccessful information trials ($b=-35.2$, $SE=5.6$, $t(5384)=-6.29$, $p<.001$) and significantly slower in the executive function compared to the control block ($b=45.8$, $SE=5.70$, $t(5384)=8.03$, $p<.001$). There was no significant effect of the interaction between block condition and information trial success ($p=.975$). Figure 5 shows mean reaction times by information trial type.

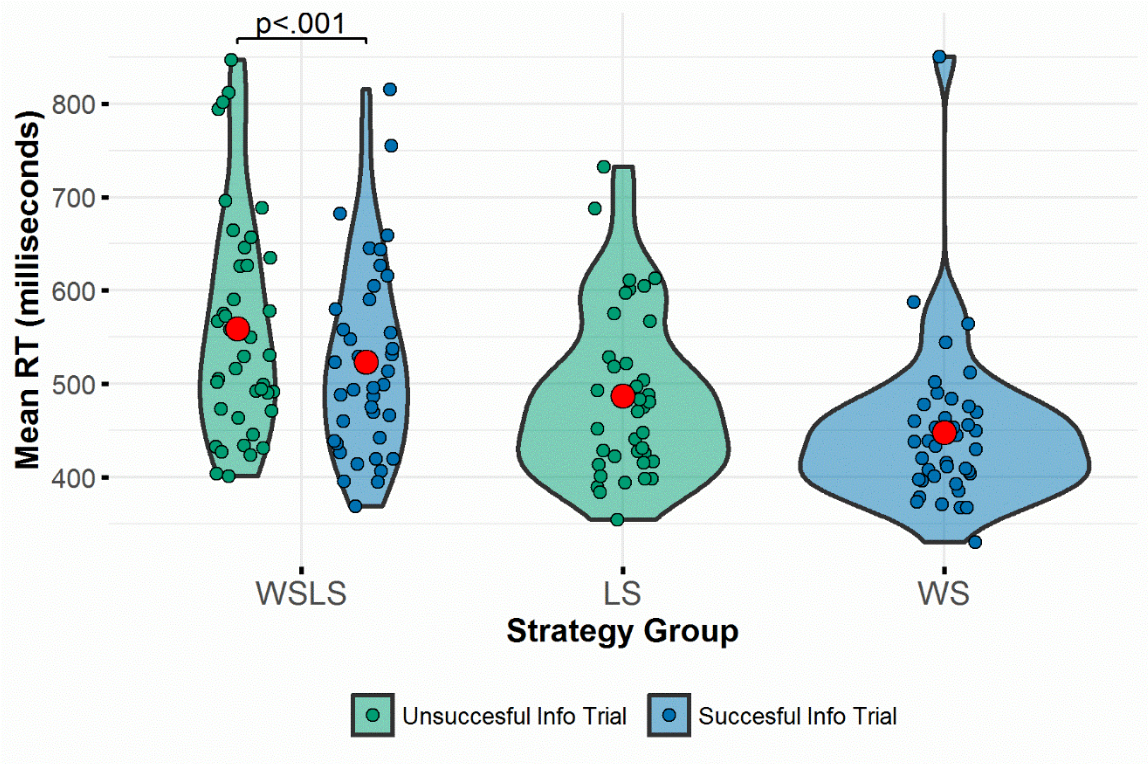

Figure 5: Reaction times based on observing successful and unsuccessful information trials. Red dots show mean in each block or condition. Brackets indicate significance level of difference between blocks.

Executive Function Dual-task

Participants all completed a different number of trials of the dual-task, as their total trial number depended on the speed at which they completed the main task. Analysis of the dual-task performance was therefore capped at 124 trials, as this was the minimum number of trials completed by all participants in either block (range EF block: 124-147; range control block: 153-169).

Participant accuracy was at ceiling in the control condition, and significantly above chance in the executive function condition: performance in each block of each condition was significantly above a chance level of 50% as shown by binomial testing (p<.001 for all conditions).

Success on each trial of the dual-task (shown in table 5) was analysed using a binomial linear mixed effects model with fixed effects of group, block condition, scaled item number and their interactions, and participant ID as a random effect. This model was significantly better than the null model ($\chi^2(11)=4881$, p<.001). Accuracy was significantly lower in the EF block (b=-2.38, SE=0.095, z=-25.1, p<.001) and accuracy got lower as trial number increased, in the EF block only (b=-2.53, SE=0.228, z=-11.1, p<.001; see figure 7). There was a significant interaction between group, block condition and item number (b=1.01, SE=0.325, z=3.10, p=.002). Post-hoc analysis using the emtrends function in R showed that accuracy declined *over item number* significantly less in every group in the control block compared to every group in the EF block ($p<.001$ for all comparisons). In the EF block the WSLS group showed significantly more decline over item number than the WS and LS groups ($p=.001$ and $p=.009$ respectively) but there was no difference between the LS and WS groups ($p=.999$). However, a post-hoc Bonferroni-Holm correction of all pairwise comparisons of group membership indicated there was no significant difference in overall accuracy between groups ($p>.99$ for all comparisons), indicating that the EF impact on the different groups was indeed similar. There was either no offloading of the dual-task impact onto the audio task rather than the search task, or the amount of offloading was the same across strategy groups.

Table 5: Mean accuracy in the audio switching task for each group and block condition

GROUP MEMBERSHIP	BLOCK CONDITION	MEAN ACCURACY
WSLS	EF Dual	78.0%
	Control Dual	97.1%
LS	EF Dual	82%
	Control Dual	97.2%
WS	EF Dual	73.6%
	Control Dual	96.7%

Additional Statistical Analyses

Participant accuracy was at ceiling levels in the control block and significantly above chance for the EF block in all conditions ($p<.001$ for all conditions) – see figure 6.

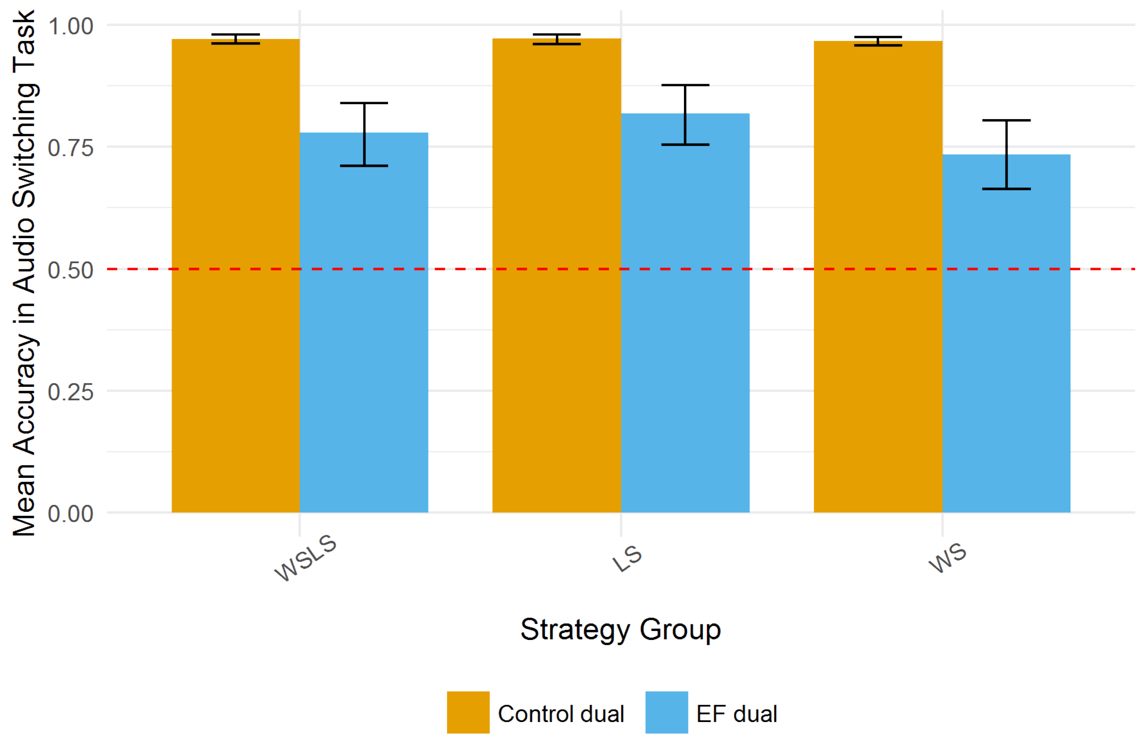

Figure 6: Mean accuracy in switching task per condition. Error bars represent bootstrapped 95% confidence intervals. Red line indicates chance performance.

The interaction between block condition and item number is shown in figure 7. The decline in accuracy in the EF block is expected to be due to the fact that accidentally missing or not hearing a single switching cue could mean a participant believed they were responding accurately, but was actually wrong on all trials after that point, unless another cue was missed.

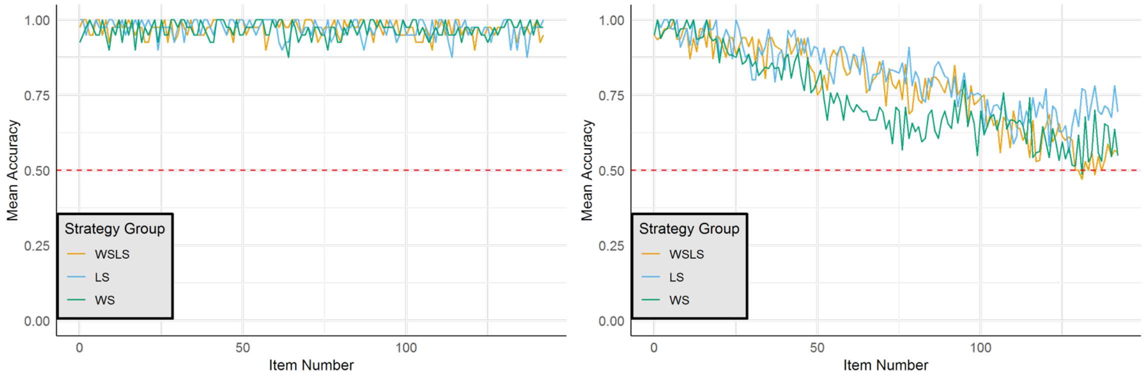

Figure 7: Decrease in mean accuracy in dual-task in executive function block (right panel) compared to consistent high accuracy in the control block (left panel). Dashed line indicates chance performance.

Discussion

Overall there was a significant effect of the dual-task on the binary search task, with reaction times in executive function dual-task blocks significantly slower than in the control blocks. However, the impact of the dual-task was the same across all group conditions, with no significant differences in how much the concurrent task slowed down performance.

This may indicate that there is an EF requirement even in such a simple task, causing performance disruption even in conditions where it was not predicted. Even though the tasks were intended to be very simple for participants, the impact of the executive function task may show that even basic decision-making draws on system-2 resources.

However, given the same level of interference was found across conditions it may also indicate that, although the WSLS condition required less predictable responses than the other conditions, all conditions were extremely easy for participants. Ceiling effects may therefore have obscured any potential differences between conditions in the impact of the EF task relative to the control task. The significant effect of item number on reaction times suggests a training effect in the main task and thus supports this conclusion, as it could indicate that participants reached ceiling levels of performance in the main task before the end of each block. The decrease in RTs also suggests that responses to the main task were becoming increasingly well-practiced, which is likely to decrease dependence on system-2 resources with greater task experience. As the RTs in the search task decreased, accuracy in the concurrent audio task also decreased. This may suggest some offloading of task demands onto the concurrent task. However, the amount of offloading appears to remain the same across strategy groups.

Although there was no difference in how much the dual-task affected each strategy group, there was a significant difference between overall task performance in the three different conditions. Reaction times in the WSLS group, the only condition requiring behaviour analogous to a social learning strategy, were significantly slower than those in either of the non-selective conditions. Responses were also significantly faster after successful information trials (wins) than unsuccessful information trials (losses), overall and in the WSLS condition. This indicates that participants found responses after 'wins' easier than responses after 'losses', which would suggest that the lose-shift condition should be the most challenging, as it is comprised only of unsuccessful information trials. The finding that the flexible strategy was slowest therefore indicates that flexible strategy use is significantly more challenging than applying a simple additive combination of each component action within the strategy; the flexible nature of the strategy itself is what makes it challenging.

Chapter 2C: Investigating Dual-Task Interference on a Simple Selective Search-Task - Experiment 2

In order to establish whether true differences in dual-task interference between using a selective copying strategy and repeatedly applying a consistent rule were masked by ceiling effects in the task, further data were collected for experiment 2 using a task that was more challenging for participants. The more challenging task followed the same basic structure but now two selections, out of four possible options, were made on the information and test trials, rather than one out of two. The aim of this was to increase the requirement to make selective choices, in the selective strategy condition, as on some trials two different strategies (both WS and LS) would be required at once in order to perform optimally at the task. The strategy requirements in the repeated copying (WS group from E1) condition, however, would remain the same.

The simultaneous application of two information use strategies (retain useful elements and deviate from detrimental elements of a single demonstration), increases the ecological validity of the task, as it models more realistically the cognitive requirements of learning situations that may lead to cumulative culture. For example, if observing a conspecific foraging successfully in some locations and unsuccessfully in others, selective copying of only the successful locations combined with deviating from the unsuccessful locations would be required in order to outperform the observed model. In copying situations such as this it is unlikely that the observer would be able to act simultaneously with the model (due to e.g. social status, availability of foraging locations or tools). Consequently, there will also be some memory requirement involved to retain information about what should be copied and what should be deviated from. An additional memory load condition was therefore also introduced, to increase the overall difficulty and ecological validity of the task. This was added for both the selective strategy condition and the repeated copying condition.

A direct measure of metacognition was also tested in E2. This enabled a direct test of whether the aim of restricting access to participants' metacognitive reasoning, as outlined in section 1.2, was indeed fulfilled by the EF dual-task. It also provided a measure of metacognitive efficiency that could be correlated with proficiency in using a selective copying strategy; we predicted this relationship to be positive.

Methods

Participants

Participants were recruited at the University of Stirling and took part either in exchange for course credit or cash remuneration at a rate of £5/hour rounded up to the nearest pound. One-hundred and fifty-one participants took part (39 male; 112 female, mean age: 20.4, age range: 17-35). 57 participants completed the training phase of the experiment but did not complete the full testing phase (see below for a break-down of exclusions and accuracy thresholds required to pass all tasks). In total 94 participants (23 male; 71 female, mean age: 20.2, range: 17-35) completed the full study and were included in the dataset reported in the analyses below. All participants had normal or corrected to normal vision and hearing. Participants all gave written consent to take part and were aware that they were free to withdraw from the study at any time. All participants were task naïve, as participants that took part in the pilot or E1 were not able to sign up for E2. Ethical approval for the study was given by the University of Stirling General University Ethics Panel (reference GUEP 468).

57 participants completed the training phase of the experiment but did not complete the full testing phase:

- 6 were excluded due to experimenter or technical error during testing
- 1 was excluded as a fire alarm sounded during the second block of testing and the building needed to be evacuated
- 1 was taken unwell during the second testing break and left voluntarily.
- The remaining 49 participants were excluded as they did not pass their initial training round.
- 3 participants did not fully understand the instructions, due to language barriers which were only made known to the experimenter after the task ended. This meant they were often responding unintentionally incorrectly to trials, despite giving full attention and effort to the task.
- 6 participants failed on the basis of their accuracy scores on the metacognition task.
- 2 participants failed as they did not respond to any of the audio tones and did not respond to prompts from the experimenter to remind them of the task instructions.
- 1 participant failed based on their search task score.
- The remaining 37 participants failed based on their audio-task accuracy.

Audio task accuracy from participants who failed the training round was generally well below the required level to pass with a mean only slightly above chance performance (see figure 8).

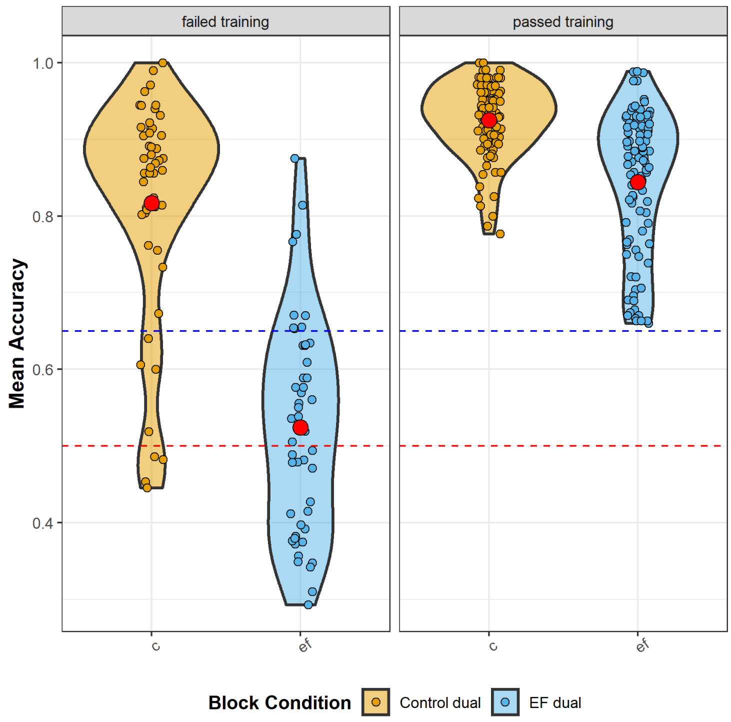

Figure 8: Accuracy in the training round of the audio dual-task for participants who passed and failed training. Red dots indicate mean of each block. Red dashed line indicates chance performance, blue dashed line indicates passing threshold. Some participants appear to be above the passing threshold but still appear in the 'failed training' section as they may have passed the accuracy threshold in one block, but a pass in all blocks & conditions was required to continue to full testing.

Apparatus

Participants were tested using a desktop computer running Windows 8 with a standard mouse, a Black Box Toolkit 4 button response pad button box and Goji over-ear active noise cancelling headphones. All tasks were written and run in Psychopy version 1.84.2 (Peirce et al., 2019).

Procedure:

The procedure remained very similar to that used in Experiment 1 with changes described below. All participants completed both the search task and the metacognition task described below in conjunction with the switching dual-task used in Experiment 1. Whether participants completed the search task or the metacognition task first was counterbalanced across participants. Both tasks were completed alongside the concurrent *switching* dual-task and control dual-task used in E1. The order of the switching and control blocks was again counterbalanced across participants. Due to the low pass rate in Experiment 1, a minor change was made to the switching task to lower the memory requirements. The requirement of the participant to keep track of previous switching behaviour was removed. The switching task now had two different switch cues, one to indicate that participants should start adding one to the tones they heard and one to indicate that they should click the same number of times as tones they heard. These different cues were noticeably distinct, and logically set to ensure participants should remember which cue was which: the cue to start adding one to tones consisted of 2 short buzzes and the cue to stop adding one was a single buzz. Participants were informed about the meaning of these cues during the instructions of the task and played an example of both cues as well as given a small number of practise trials using them before the training round.

The passing thresholds for the training rounds were also updated for E2. For both the metacognition and WSLS search tasks a performance threshold of 75% was set as a requirement for passing the initial training round. For the WSLS search task this was intended to be quite generous as previous testing had shown ceiling levels of accuracy in this task. The accuracy threshold was only intended to act as a motivator for participants to give the task sufficient attention to ensure they could continue on to the full testing rounds of the experiment.

A score threshold of 75% in the metacognition task allowed for an accuracy of around 65% in the perceptual task, slightly lower than that found in initial pilot testing of the task, and then to perform at around chance at the confidence judging stage. This was to ensure participants with naturally lower levels of metacognition were not screened out, while ensuring participants who were performing at chance in the visual perception task, due to lack of concentration or poor visual acuity (and therefore whose metacognitive judgments were either not informative or always at the lowest level of confidence) were.

Difficulty level in the visual perception task had been set through piloting at a level that meant participants were expected to score on average around 75% accuracy (see below).

For the audio dual-task the performance threshold was set at 65% accuracy. This needed to be reached in both the executive function and control rounds. This was to ensure participants were performing significantly above chance in both tasks (based on a predicted 80 – 85 trials completed in each condition of the training rounds). Chance level was presumed to be 50% accuracy for both EF and Control rounds as participants should always have been choosing from one of two response options: click once or twice on non-switching trials, and click twice or three times on switching trials. This accuracy threshold assumed participants had a clear understanding of the task instructions and were attempting to respond appropriately, and was essential to include to ensure data collected were accurately reflecting behaviour under dual-task conditions. If not attempting, or not understanding how, to respond in an accurate manner participants would not be considered to be taking part under genuinely dual-task conditions and data from those participants would therefore be misleading and inaccurate.

Participants were required to pass all three tasks in order to progress to the full testing rounds.

As Experiment 2 was longer than the previous two experiments, five-minute breaks were added between the training round and first test round, and between the first and second blocks of testing to ensure participants did not experience task fatigue. During the breaks participants were given an unrelated colouring activity, were offered water to drink and were permitted to use the toilet if required. They were asked not to communicate with other participants or the experimenter and were not permitted to use their phones. During the breaks a visual counter informed the participant of the remaining time and a short alarm sounded at the end of the break and prompted the participant to restart the experiment by pressing a button. All participants completed both breaks for the full five minutes each, and no participants overran their break by more than a few seconds.

Search Task

The procedure used was similar to that used for the binary search task in experiment 1. However, the number of conditions was reduced so a comparison was only made between a selective copy and an always copy condition. The difficulty of the task was increased so that participants were now required to make two selections out of four rather than one out of two. This meant that in the selective copy condition there was also more variability of information trial types to include an equal balance of trials which revealed two successful stimuli, two unsuccessful stimuli, and one successful and one unsuccessful stimulus. In the always copy condition the information trial would always show two successful stimuli. Participants completed 18 trials in the training phase followed by 54 trials in each block of the experimental phase.

A memory load was also included for 50% of participants in each strategy group. With the memory load present the information trial would not remain visible while the test trial selections were being made. A time limit was also added to all trials, so if the first response was not made within three seconds the trial would time out and one point would be removed from the participant's total score. See figure 9 for an example of the procedure.

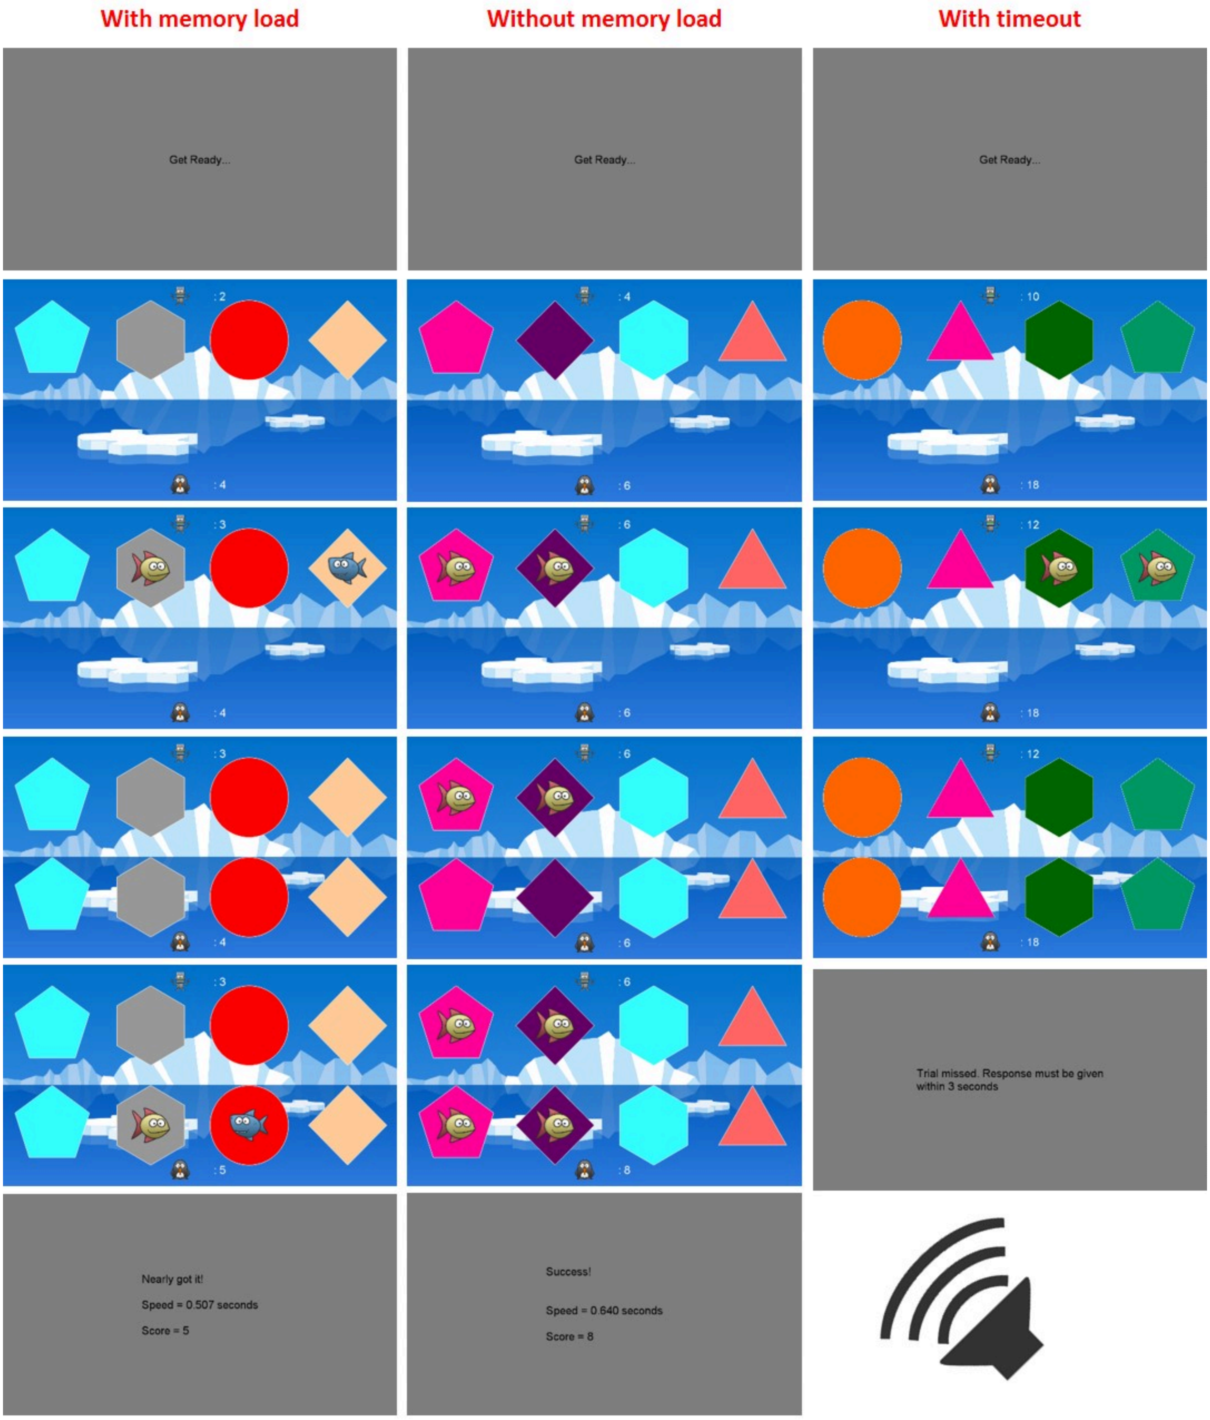

Figure 9: Example trials of the search task. In all conditions the task was completed alongside the audio dual-task.
Left: *an information trial with one succesful (fish) and one unsuccesful (shark) stimuli, with an additional memory load. The test trial shows the participant has used a correct strategy and repeated the successful information, and shifted away from the unsuccesful information, although due to chance they have selected one unrewarded stimulus.* ***Middle:*** *an information trial with two successful stimuli, with no memory load. The test trial shows two successful stimuli are selected.* ***Right:*** *a trial with a timeout for no response.*

Metacognition Task

The metacognition task assessed metacognitive monitoring ability, and consisted of a two alternative forced choice (2AFC) task in which participants were required to decide which of two patches of white dots was sparser.

Participants were shown two circular patches of white dots on a grey background, with the question 'Which circle has fewer dots?' written beneath it. To generate the dot circles one circle was randomly generated with a density of between 500 and 800 dots. A second circle was generated by then adding or removing a fixed percentage of the dot density, based on the difficulty level. There were 7 difficulty levels, each twice as difficult as the level before. Level 1 had a difference of 64% between the patch densities, so if one patch contained 500 dots the other would contain either 820 or 180 dots. The difference between the patches reduced by half at each difficulty level until level 7 which had a difference of 1% between the patches, so if one patch contained 500 dots the other would contain either 505 or 495 dots. The top end of the difficulty scale was designed to be functionally impossible to ensure participants would be guessing on a certain number of trials, to encourage them to use the full scale of confidence. Overall, the difficulty levels were set through pilot testing, to try to achieve accuracy levels of around 75% in the task. There was an even balance of all 7 difficulty levels in each round, presented in a fully random order. The same difficulty levels were used across participants.

Reward structure in the task was fixed so that there was a 50% chance of the target patch (the patch with the lower density) being on the left of the screen, so the target should have appeared on the left or right side of the screen an equal number of times.

Participants were required to select one of two buttons on the button box to choose the patch on the left or the right of the screen. They needed to make their selection within 5 seconds, otherwise the trial would end and a message would display on screen reading 'Trial missed, response must be given within 5 seconds' and one point would be removed from the total. If the trial was missed an opportunity would not be given to report a confidence rating.

Once a selection had been made participants were asked to rate their confidence from 1-4 using one of the four buttons on the button box, with the buttons from left to right representing the numbers from 1-4. No time limit was imposed on the confidence rating. After a response had been given a red 'x' would appear on the scale at the confidence level selected for 1s, and then a score screen would appear showing how many points had been gained or lost. Points were gained for correct patch discriminations and lost for incorrect patch discriminations, and the number of points gained or lost corresponded to the confidence level given. For example, a correct selection rated with confidence 3 would gain 3 points, and an incorrect selection rated with confidence 2 would lose 2 points. This symmetrical points structure was used to try to motivate participants to avoid rating responses they were unsure of with high confidence, due to the risk of losing a high number of points (although see *Discussion* below for the potential negative impact of this scoring system). For negative scores the score screen would show in red text, and for positive scores it would show in white. During trials a running total score would be shown in the top right corner of the screen. This score could be negative, and negative running scores were displayed in red.

The metacognition task was completed concurrently with the switching dual-task and control dual-task. Participants completed 28 training trials and 56 test trials in each block, with an equal number of trials at each difficulty level, presented in a random order. See figure 10 for an exmaple of one trial of the metacognition task.

Experiment Two Predictions

We predicted that, as in experiment 1, the EF-dual-task would make RTs in the search task slower. We again predicted more dual-task interference in the WSLS group than the WS group, due to the requirement to use a selective strategy in the WSLS group. This is in line with EMCC prediction that the use of selective strategies relies on explicit processes. We did not make specific predictions regarding the addition of a memory load to the search task. The EMCC would predict that a task that reduced capacities for selective copying would also have a negative impact on metacognitive efficiency. We therefore predicted the *switching* dual-task would have a negative impact on metacognitive efficiency in the metacognition task, in line with the predictions for the search task

Results

Search Task

Outliers

Overall 531 outliers (taken as 3 x MAD from the mean, see *Outliers* in Pilot) were removed from the data for very long or very short reaction times. 473 outliers were removed from the upper end of the response distribution and 58 were removed from the bottom end. This represents 5.2% of the total data. Outliers have only been removed from participants' first responses. This is because some second responses were extremely quick as participants were using a response strategy of deciding on both of their selections before making either and then pressing both buttons almost simultaneously, so removing these trials would create an incorrect representation of the data. If R1 was removed on any particular trial, the entire trial was removed from the dataset so R2 was also removed.

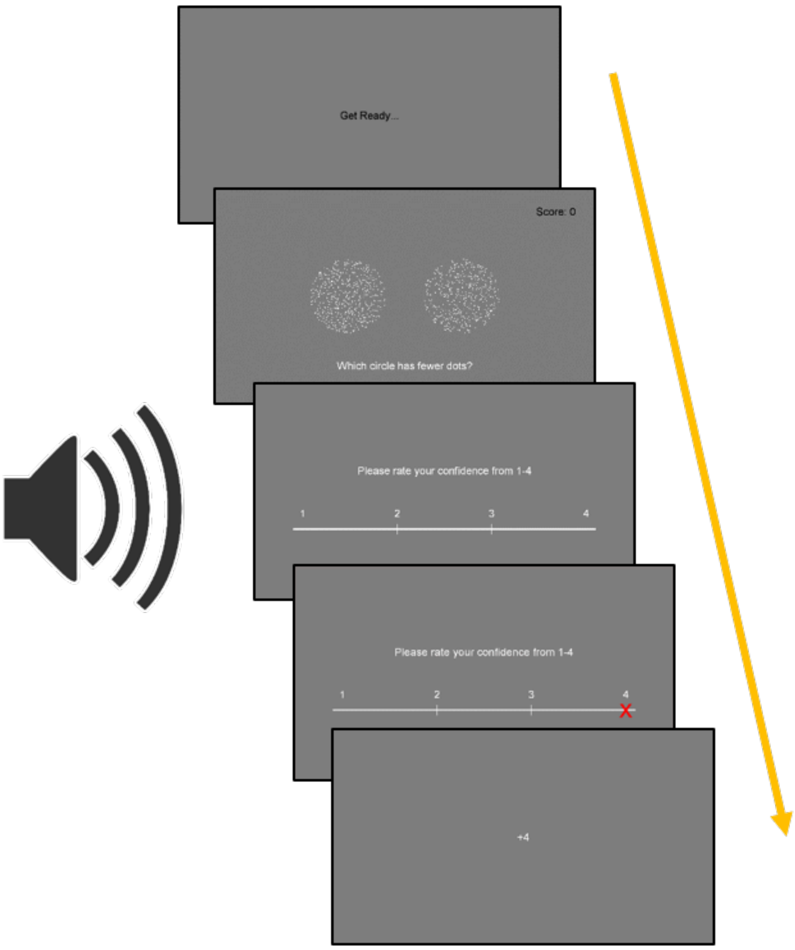

Figure 10: An example of a trial of the metacognition task

.

Analysis

Correct WSLS strategy adoption was at ceiling with correct strategy use nearly 100% of the time across conditions (range: 99.5%-100%).

Response times (shown in table 6) were analysed using a linear mixed effects model with fixed effects of group (WSLS or WS), block condition (EF dual-task or control dual-task), memory load, response number (first (R1) or second (R2) response), item number, level of success of the information trial (whether it scored 0, 1 or 2 points) and the interactions between group, block condition, memory load and response number. Participant ID was included as a random effect. The WSLS group, in the control condition with no memory load was taken as the baseline and *p*-values were estimated from the resultant t-statistics with degrees of freedom being the number of observations minus the number of fixed parameters in the model (Baayen et al., 2008). R2 is measured from R1 rather than stimulus onset.

The model was significantly better that its null equivalent ($\chi^2(17)=1123$, *p*<.001).

Table 6: Mean RTs (milliseconds, SD in brackets) for each strategy group, block condition and memory load condition in E2

STRATEGY GROUP MEMBERSHIP	BLOCK CONDITION	MEMORY LOAD CONDITION	MEAN (SD) RT1 (ms)	MEAN (SD) RT2 (ms)
WSLS	EF Dual	Memory Load	341(228)	225(272)
	Control Dual	Memory Load	307(186)	221(237)
WS	EF Dual	Memory Load	335(185)	251(189)
	Control Dual	Memory Load	326(176)	230(189)
WSLS	EF Dual	No Memory Load	307(220)	309(238)
	Control Dual	No Memory Load	303(205)	305(275)
WS	EF Dual	No Memory Load	331(216)	191(246)
	Control Dual	No Memory Load	314(186)	185(228)

The EMCC would predict that the dual-task would have more of an impact in the WSLS group, so would therefore predict longer RTs in the EF block of the selective group, when compared to the control and to the WS group. Longer RTs would also be predicted in the EF block compared to the control block, and when under an additional working memory load. These predictions are supported by the significant interaction between group, block condition and memory load (b=-35.8, SE=16.5, t(19247)=-2.17, p=.030) (see figure 15). There was an additional significant main effect of item number, indicating faster responses as trial number increased: (b=-0.787, SE=0.094, t(19247)=-8.38, p<.001).

There were also multiple significant interactions. For clarity, each interaction is described in detail below. Post-hoc comparisons of the interactions were carried out using estimated marginal means (using the emmeans package in R).

Group, block condition and memory load (see above):
In the WSLS group, both block conditions showed an RT decrease when a memory load was present, compared to there being no memory load. This was more pronounced in the control block.

Conversely, in the WS group both blocks showed an RT *increase* with a memory load present (see figure 15).

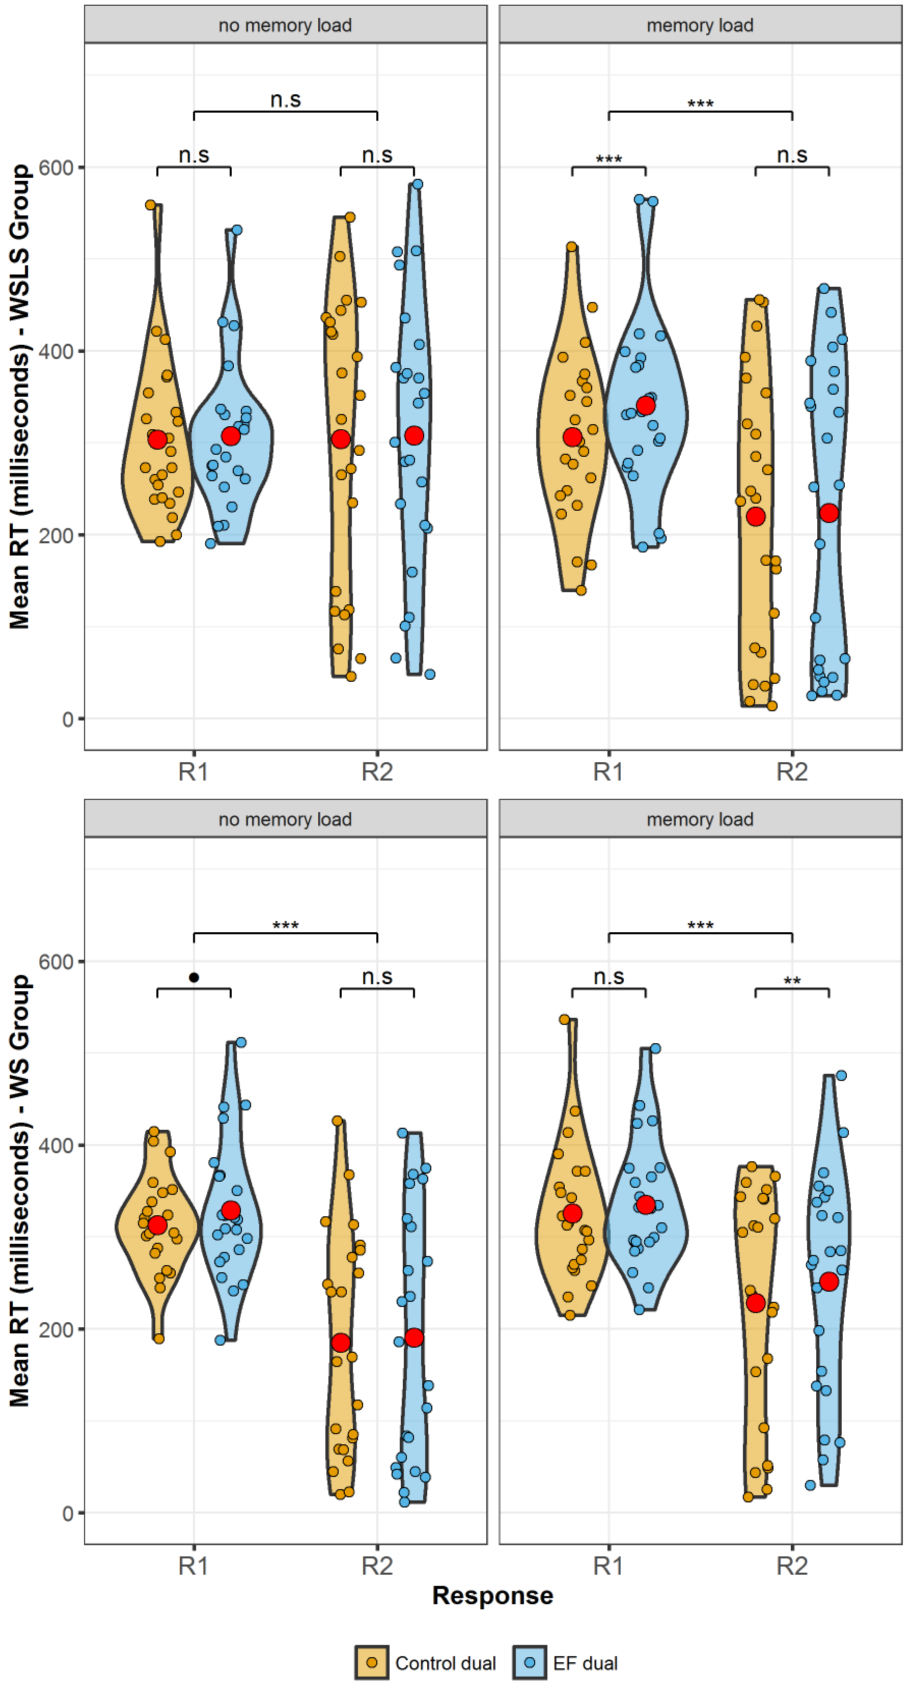

Figure 11: Response times at R1 and R2 for the WSLS group (top) and WS group (bottom), split by block condition and memory load

Block condition and memory load (b=29.5, SE=11.5, t(19247)=2.56, p=.011):
When there was no memory load the difference between the executive function and control blocks approached significance (b=-7.33, SE=4.14, t(19157)=-1.77, p=.077), but when a memory load was present the EF block was significantly slower than the control block (b=-17.5, SE=4.13, t(19156)=-4.25, p<0.001). This indicates that the dual-task had a significant impact on RTs only when combined with the additional memory load.

Group and response number (b=-130, SE=11.7, t(19247)=-11.2, p<.001):
At R1 both groups had similar RTs (b=-15.6, SE=18.2, z=-0.856, p=.394) but at R2 the WS group had significantly faster RT than the WSLS group (b=46.7, SE=18.2, z=2.56, p=.012), indicating a greater speed increase between responses in the WS group compared to the WSLS group (see figure 16).

Memory load and response number (b=-88.0, SE=11.5, t(19247)=-7.64, p<.001):
There was a significant decrease in RT between R1 and R2 overall both with and without a memory load (p<.001 for both memory conditions), but the decrease was larger when there was a memory load present (see figure 17).

Group, memory load and response number (b=120, SE=16.5, t(19247)=7.28, p<.001), and *group, block condition, memory load and response number* (b=52.1, SE=23.4, t(19247)=2.23, p=.026):
In the WS group there was a significant decrease in RT between R1 and R2, both with and without a memory load and in both blocks (p<.001 for both). This decrease was also significant when split between the EF and control blocks (again p<.001 for both blocks). However, in the WSLS group this RT decrease only applied when a memory load was present; when there was no memory load there was no decrease in RT (R1-R2, overall: b=-1.42, SE=5.76, t(19156)=-0.246, p=.805; control block: b=-1.74, SE=8.13, t(19156)=-0.214, p=0.803; R1-R2, EF block: b=-1.10, SE=8.16, t(19156)=-0.314, p=.893) (see figure 18).

Overall, the interactions indicate that although the dual-task had a similar effect on the different strategy groups, the addition of a memory load significantly changed the way participants behaved in the different strategy groups. Across both strategy groups the detrimental impact of the dual-task on RTs was also contingent on the memory load being present.

Figure 11 shows the reaction times at R1 and R2, split by block condition, memory load and group.

Metacognition Task

Mean accuracy in the visual discrimination of the stimuli in the metacognition task was approximately 76% (control block: 75.9%; EF block: 76.3%), close to the desired accuracy based on the difficulty levels set during piloting. However, this mean accuracy level reflects a wide range of accuracy, with some participants performing much higher than expected and some participants performing much lower (control range: 62.3-87.5%; EF range: 64.3-91.1%). A repeated measures ANOVA showed no significant difference in discrimination accuracy between control and EF blocks (F(93)=0.291, p=.591).

Participants in both blocks tended to rate their confidence highly, with an aversion to using the lowest confidence level even at the hardest difficulty level (see figure 12).

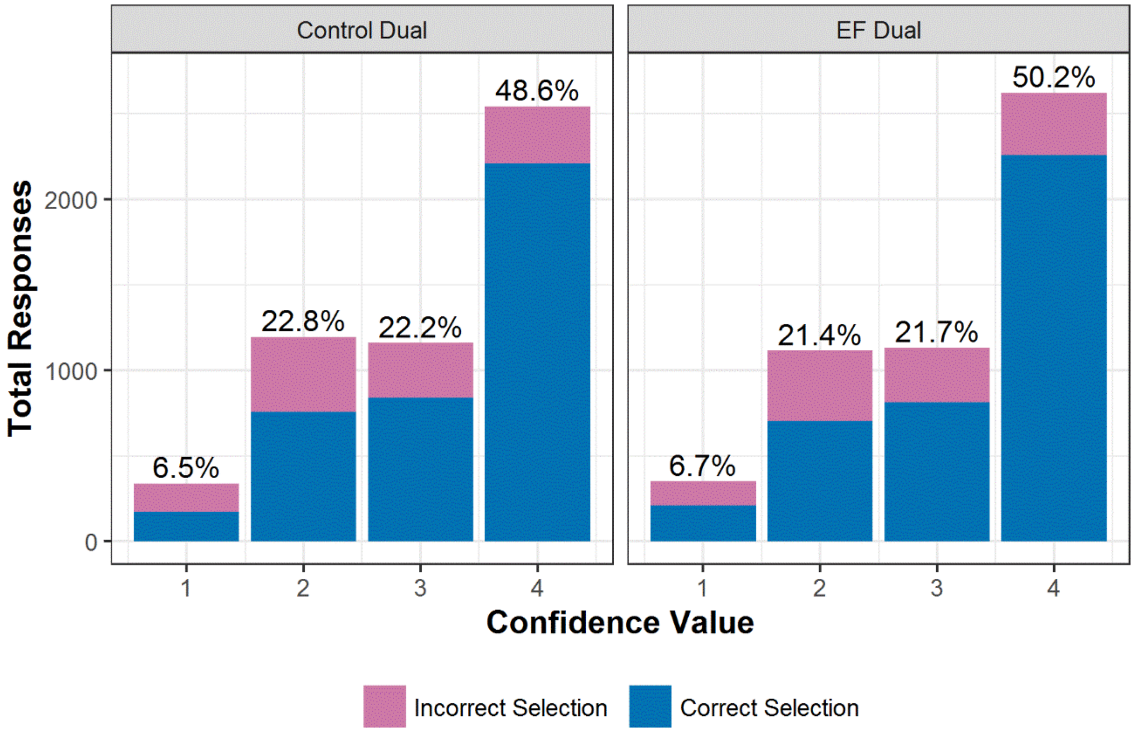

Figure 12: Total responses given at each level of the confidence scale, split by the accuracy of the visual discrimination. Percentages indicate the percentage of all responses in that block at that confidence level.

Success on each trial was analysed using a binomial general linear mixed effects model with fixed effects of block condition, scaled confidence and scaled difficulty level and their interactions. Participant ID was included as a random variable and the control block was taken as the baseline. The model was significantly better than the null equivalent (x2(7)=2329, p<.001). There were significant effects of confidence level (b=4.45, SE=0.395, z=11.3, p<.001) and the interaction between confidence and difficulty level (b=-2.50, SE=0.293, z=-8.56, p<.001). This indicates increased success when participants were more confident, although the increase does not apply equally to all difficulty levels (see figure 13). This increase is particularly pronounced from confidence level 1 to level 2, although this may be largely due to the low number of responses given with confidence level 1.

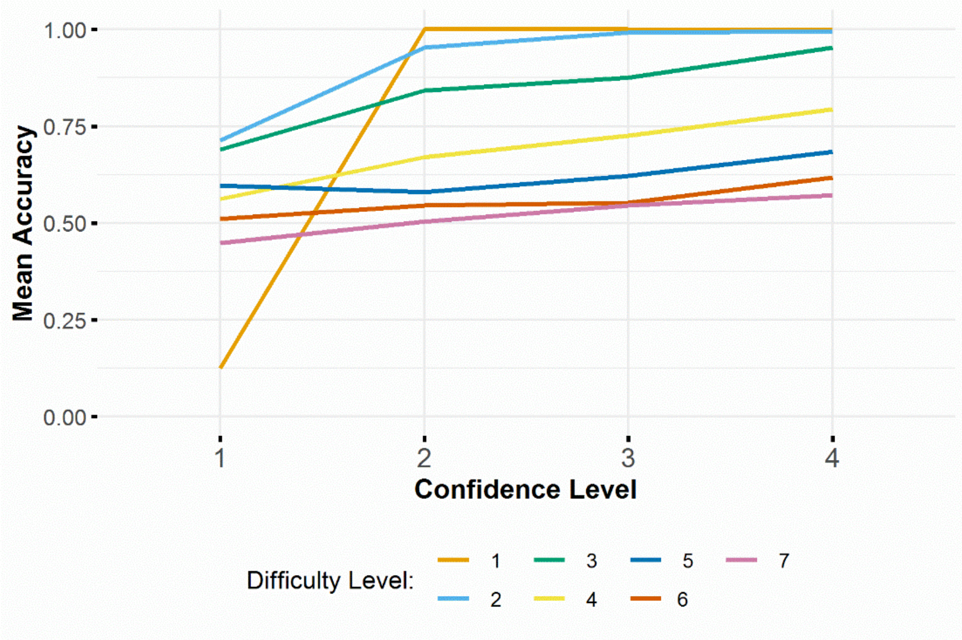

Figure 13: Accuracy in metacognition task at each confidence level, split by trial difficulty. Difficulty Level 1, least difficult; Difficulty Level 7, most difficult.

The impact of the dual-task on participants' metacognition was analysed using the metaSDT package in R (Barrett et al., 2013). A single score of metacognitive efficiency (m-ratio) for each participant in each block was calculated as metacognitive sensitivity (meta-d'$_b$: balanced model fit) divided by overall sensitivity in the visual task (d'). Meta-d' is defined as the type-I sensitivity (accuracy in the visual discrimination task) that would be found if all of a participant's type-II ratings (subsequent confidence ratings) were considered to be optimal (Fleming & Lau, 2014). This measure aims to give a bias free measure of metacognition which is not affected by performance in the visual task or a bias towards over- or under-confidence (Maniscalco & Lau, 2012). A repeated measures ANOVA showed no significant difference in m-ratio between control and EF blocks (mean control: 1.13, mean EF: 1.03, $F(93)=0.663$, $p=.418$) (see figure 14).

Correlation between search task and metacognition task

The mean difference in RT between the EF and control blocks of the search task was calculated for each participant. This RT difference reflected the impact of the dual-task. Overall, there was no correlation between dual-task impact and baseline metacognitive efficiency (taken as m-ratio for the control block only) for either R1 or R2 (R1: $r=.011$, $t(92)=0.102$, $p=.919$; R2: $r=.070$, $t(92)=0.675$, $p=.501$), indicating no relationship between greater metacognitive efficiency and impact of the dual-task on speed of applying a flexible copying strategy. This may indicate that, in contrast with the predictions of the EMCC, increased metacognitive monitoring accuracy does not necessarily lead to an increased proficiency in applying a flexible strategy.

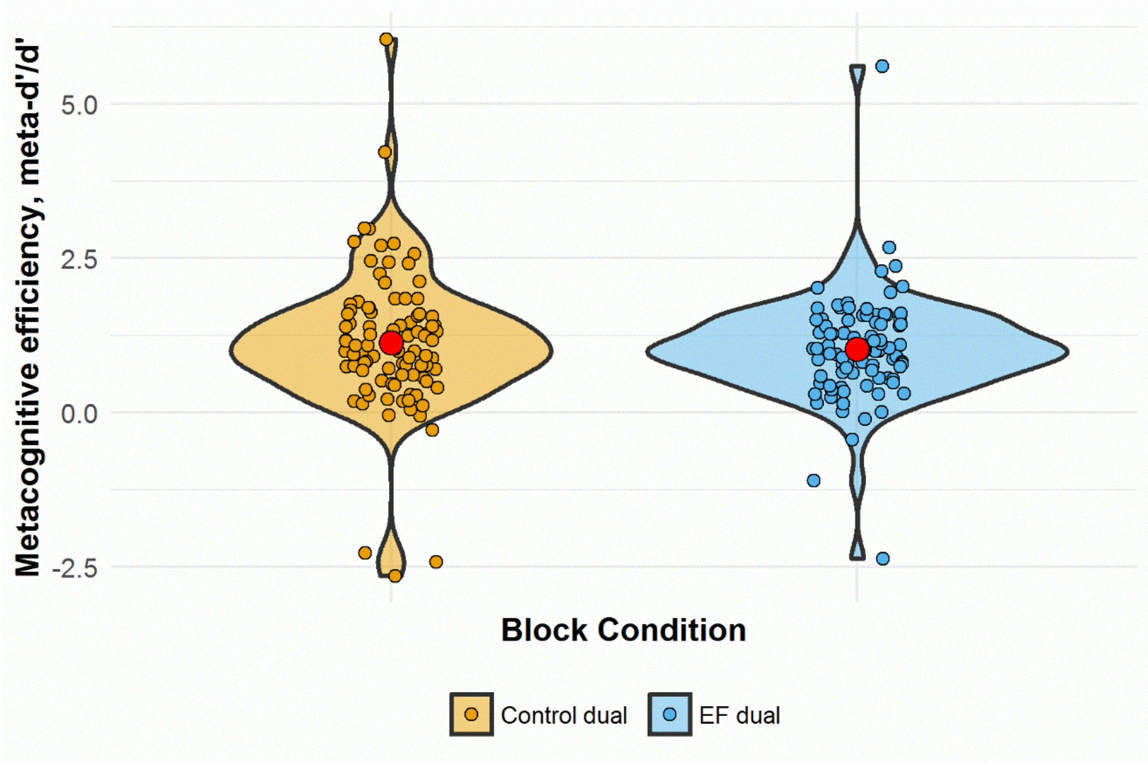

Figure 14: M-ratio (metacognitive efficiency) for the control and EF blocks. Red dots indicate the mean for each block.

Audio Dual-task

Participants all completed a different number of trials of the dual-task, as their total trial number depended on the speed at which they completed the search and metacognition tasks, and which task they were doing. Analysis of the dual-task performance was therefore capped at 84 trials, as this was the minimum number of trials completed by all participants in either block in both tasks (range metacognition EF block: 84-126; metacognition control block: 107-129; search EF block: 97-132; search control block: 123-139. The numbers of trials in the EF blocks are lower due to approximately one in six trials being a switch cue and not being counted in the analysis).

Participant accuracy was at or close to ceiling in all conditions; accuracy range 86.5%-95.7%.

Success on each trial of the audio task was analysed using a binomial linear mixed effects model with fixed effects of block condition, main task and their interactions, and scaled item number. Participant ID was included as a random effect. This model was significantly better than the null model ($\chi^2(4)=850$, p<.001). Accuracy was significantly lower in the EF block (b=-1.21, SE=0.057, z=-21.3, p<.001) but significantly higher during the search task compared to the metacognition task (b=0.219, SE=0.069, z=3.16, p=.002). Accuracy got lower as trial number increased, (b=-0.421, SE=0.071, z=-5.91, p<.001). There was also a significant interaction between main task and block condition (b=0.249, SE=0.085, z=2.91, p=.004), with a greater accuracy difference between control and EF blocks for trials completed alongside the metacognition task than the search task.

Additional Analyses

Figures 15-17 give a visual representation of some of the interactions from E2.

Group, block condition and memory load:

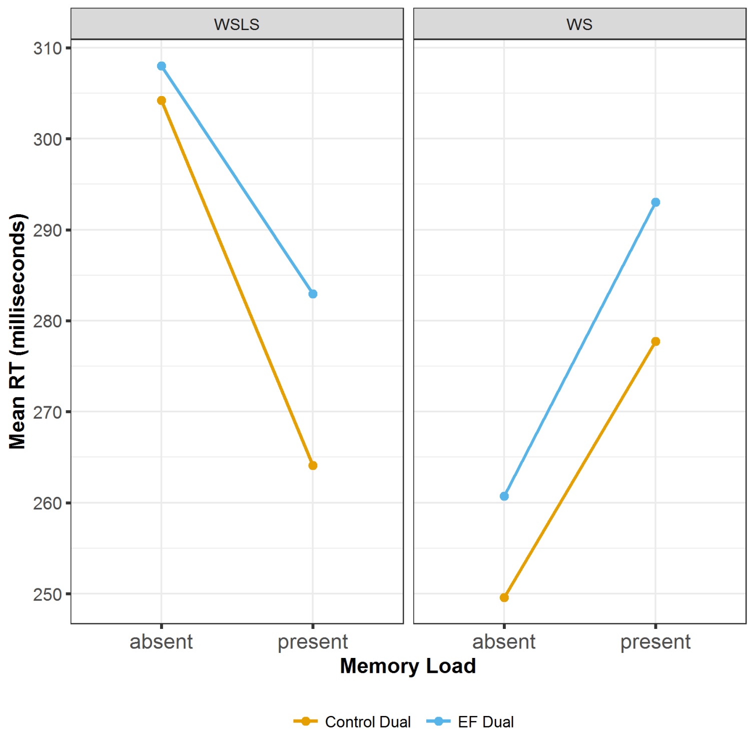

Figure 15: Change in RT with memory load, split by block condition and group

Group and Response Number:

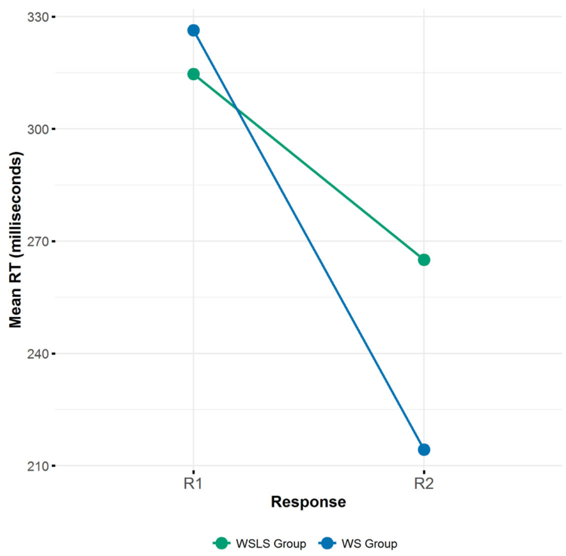

Figure 16 (left): Speed increase between R1 and R2 for each strategy group

Memory load and response number:

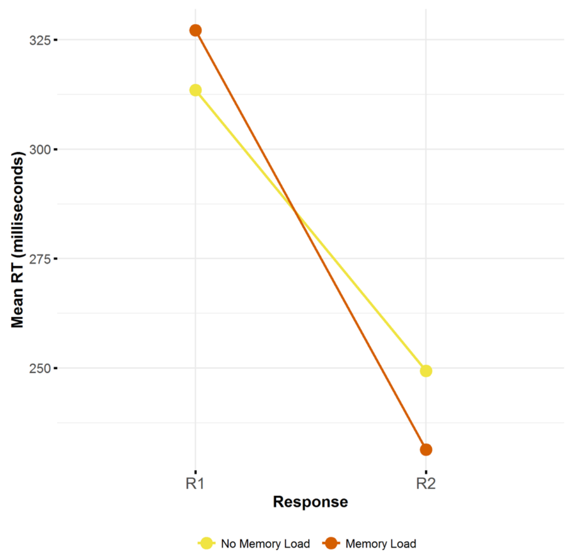

Figure 17 (right): Reaction time decrease comparison between memory load and no memory load conditions.

Group, block condition, memory load and response number:

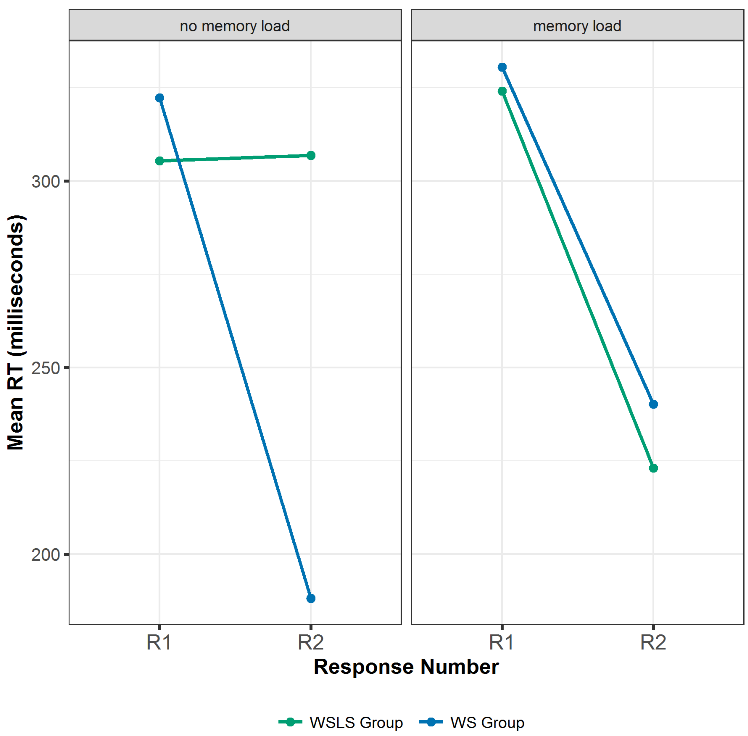

Figure 18: Speed increase from R1 to R2, split by memory load and strategy group.

Discussion

The data show no overall effect of the dual-task on response times in the search task, despite the search task being more challenging than the binary task used in E1. There was also no significant difference found between groups overall. However, there was a significant difference between the groups in terms of how much the memory load affected RTs. The WS, always-copy group appeared to be relatively unaffected by the additional memory load: although there was an overall increase in RTs with an additional memory load this was not significant and the difference between R1 and R2 remained the same both with and without a memory load. In contrast, the WSLS, selective-copy group showed a significant increase in RT in the EF block at R1, followed by a significant overall decrease in RT from R1 to R2, when the memory load was applied.

This may indicate that participants were aware that the information trial would not remain visible when making their selections, and so R1 and R2 are planned before making R1. This then leads to a very quick R2, as participants are able to select both stimuli almost simultaneously. The increased RT in the EF-block only, indicates an EF-cost to this 'planning to remember'. When no memory load is present this planning is not required, so there is no RT increase at R1 but also no decrease at R2. This response strategy is in itself metacognitive, as it involves monitoring and control of the response behaviour. This may indicate some involvement of metacognition in efficiently applying flexible copying strategies, and in overcoming the challenges posed by working memory loads.

The finding that the always-copy condition shows a consistent RT decrease between responses, whereas the selective-copy condition RTs are somewhat dependent on the availability of working-memory, mirrors the findings from Experiment 1 that indicate that selective copying is more challenging than always or never copying. This increased difficulty of applying a strategy may rely more on working-memory than task-switching executive resources.

The data from the metacognition task appear to show no difference in metacognitive accuracy when under additional executive function load. There was also no correlation between control metacognitive efficiency and the impact of the dual-task on the search task. These findings are in contrast to the predictions of the EMCC, and suggest limited involvement of metacognitive control in successfully using the type of flexible copying strategy hypothesised to be requisite for CCE. However, the outcome of the metacognition task may be impacted by the tendency of participants to make much more use of the highest confidence level than might be expected if participants were responding with honest confidence ratings. Although the meta-d' method of assessing metacognitive efficiency is designed to be robust against type-II response bias, the measure can be affected by very high or low values of type-II hit and false alarm rates (Barrett et al., 2013). The tendency to only respond with the higher end of the confidence scale may also have been exacerbated by the payoff structure in the metacognitive rating task: with a symmetrical points payoff structure, overall payoff in the task will tend to be higher (assuming an average success rate higher than chance levels of 50%) if a high rating is always given for confidence. This may have allowed participants with low metacognitive sensitivity to still score highly by applying a blanket rule of 'select high confidence'.

The m-ratio scores may also have been artificially inflated due to the task design that had multiple varying difficulty levels. This design had the desired effect of ensuring some trials were much more difficult than others and therefore all participants had a combination of correct and incorrect answers. However, Rahnev & Fleming (2019) suggest that staircase procedures or tasks with multiple different difficulty levels can exaggerate estimates of metacognitive efficiency, as there can be direct comparison between trials which are definitely very hard with trials that are definitely very easy. This could be alleviated in future testing by using a single, fixed difficulty level and awarding points for metacognitive accuracy in a different way, for example by using a strictly proper scoring rule. A strictly proper scoring rule assigns points not simply for accuracy, but for the combination of accuracy and confidence. A high score can therefore be obtained by correct responses rated with high confidence, but also by incorrect responses rated with low confidence. Similarly, low scores would be assigned on correct trials rated with low confidence or incorrect trials rated with high confidence.

There was a significant decrease in concurrent audio task accuracy for both main task conditions, especially the metacognition task. This may indicate participants were offloading some of the cognitive demands of the EF block to the concurrent audio distractor rather than the main task, with more offloading during the metacognition task than the search task. This suggests that there is EF involvement in both main tasks, but the offloading has masked some of the differences between the EF and control blocks.

The above factors mean that definitive conclusions about the relationship between executive processes and metacognitive control efficiency cannot be drawn at this stage. However, this experiment did find some evidence for system-2 involvement in applying flexible copying strategies rapidly, especially when under additional working memory load, and the use of a metacognitive *control* strategy was implicated. Future testing could investigate the involvement of explicit metacognition in using flexible copying strategies by amending the metacognition task with the changes noted above, employing a task to measure impact on metacognitive control strategies directly, and ensuring equal or higher motivation was given to the ongoing task to avoid offloading of cognitive costs.

Conclusion

The hypothesis presented at the outset of the study, that system-2 executive resources and explicit metacognition are required to make the learning decisions needed for cumulative cultural evolution to occur (the EMCC), was tested by applying a concurrent executive load to participants completing a visual search task (E1 & E2). A visual discrimination task followed by a metacognitive judgement was also tested in conjunction with an executive dual-task (E2).

The results indicated that dual-task methods that put additional load on the *switching* executive function effectively restricted access to the cognitive resources required to make rapid, accurate decisions in a simplified social-learning context. The switching task had a significant impact on very simple decision making that did not require flexible strategy use, as well as on the application of a flexible strategy. These flexible strategies were shown to be more challenging for learners, due to longer RTs in E1, and fewer RT decreases in E2, which may or may not be due to higher executive function demands. The *switching* dual-task used in E2 removed the working memory demands that were present in E1 by removing the requirement for participants to keep track of previous switch cues. Given that the dual-task in E2 did not generate an overall RT increase in the search task, this suggests that the *switching* interference found in E1 may have been caused by working memory demands, rather than switching demands. This is supported by the significant memory load interference in E2. The different impact of the working memory load on the different strategy groups in E2 also indicates system-2 involvement in making selective copying decisions, as a reliance on working memory is a hallmark of system-2 in most dual-systems theories (Evans and Stanovich, 2013).

The results from the E2 metacognition task indicated that dual-task interference did not impair explicit metacognitive judgements. Due to confounding factors of task design, these results are not conclusive. However, some evidence was found for the use of metacognitive monitoring and control strategies when copying flexibly and overcoming additional working memory loads in the E2 search task.

The results summarised above offer partial support to the EMCC. While executive function interference with explicit metacognitive efficiency was limited, this may reflect the uneven use of confidence levels and inflated metacognition scores within the metacognition task, rather than a genuine absence of dual-task interference. Competition for system-2 executive function and working memory resources did have a detrimental impact on the ability to complete the search task. However, in some contexts this dual-task interference was not limited to the flexible strategy condition. Evidence was also found for the involvement of metacognitive control strategies in overcoming additional memory loads to use flexible strategies efficiently. This implicates the role of system-2 in making ecologically valid copying decisions, which may be beneficial for cumulative cultural evolution. There was evidence of offloading of cognitive demands for both main tasks, also indicating system-2 involvement which may have been masked by unequal focus being given to the main and concurrent tasks. Future testing should be able to ascertain whether this system-2 involvement is in fact a consequence of explicit metacognition.

Chapter 3: Testing the Impact of Restricted Working Memory Access on *Potential For Ratcheting*

Chapters 2A-2C began to investigate the Explicitly Metacognitive Cumulative Culture hypothesis (EMCC; Chapter 1). This hypothesis states that it is access to explicit, system-2 processes and metacognition that allow for uniquely human cumulative culture. The series of studies presented throughout Chapter 2 began to investigate this by using dual-task methods to restrict access to explicit processes while concurrent cultural evolution and metacognition tasks were completed.

The findings from chapters 2A-2C suggested that dual-tasks that placed an additional load on executive function resources did not have a significant impact on simple selective copying strategies, over and above the impact on non-selective strategies that made use of blanket copying. However, there was some indication that increased demands on memory prompted participants to complete the task in different ways, including the use of metacognitive control strategies. There was also evidence throughout each section in chapter 2 that the main task used was very easy for participants. This may have influenced the results, if the task was simple enough that it could still be completed efficiently even with restricted access to executive functions. The search task used throughout Chapter 2 also did not allow for ratcheting when under cognitive load, a key component of testing the EMCC, to be examined.

In the following chapter I therefore investigate the impact of a working-memory dual-task, on a task that investigated selective copying over multiple different levels of information. This makes the task more challenging than the one used in Chapter 2. It also makes it possible to examine participants' *potential for ratcheting,* and it can be used to simulate generational turnover to examine the impact of the dual-task over multiple generations. As in chapter 2C, this chapter also examines the impact of a concurrent working memory task on a metacognitive monitoring task. However, the format and scoring system of the metacognition task have been updated to try to mitigate some of the inconclusive results found in 2C.

In addition to the acknowledgements made at the start of this thesis, I would like to additionally thank the Newcastle Life Science Centre for giving me the opportunity to recruit and test participants in the study presented in the following chapter. I would also like to thank Dr Nicolas Cladière for assistance with calculating the measure of selection distance, and my research assistant Phoebe Abruzzese for her valuable help collecting data at the Life Centre.

The contents of the following chapter are currently under review at *Cognitive Science*. As a submitted paper a small section of the analysis is separated into supplementary information. This information has been placed back into the text in chronological order, but otherwise the following chapter is presented as it is in its submitted form. The reference for the submitted article is:

Dunstone, J., Atkinson, M., Renner, E., & Caldwell, C.A. (under review). Restricted Access to Working Memory Does Not Prevent Cumulative Score Improvement in a Cultural Evolution Task

Introduction:

Cumulative Cultural Evolution (CCE) is the process by which cultural traits are transmitted through generations of cultural agents and, crucially, amended by successive generations to become more effective, efficient or beneficial for their users (Caldwell, Atkinson, et al., 2016; Mesoudi & Thornton, 2018). This process has been dubbed 'the ratchet effect' due to the unidirectional progress of improvement (Tomasello, 1990) and is generally considered to be a process that is unique to, or at least qualitatively distinctive in, humans (Dean et al., 2014; Tennie et al., 2009; Mesoudi & Thornton, 2018). Much research within cultural evolution therefore aims to identify the human-unique cognitive capacities that enable this gradual accumulation over generations to occur.

Human propensities for social learning and high-fidelity imitation have been proposed as candidate requirements for CCE (Lewis & Laland, 2012; Tomasello, 1999). However, increasing evidence of social learning in non-human animals (henceforth animals) as diverse as bees (Alem et al., 2016), chimpanzees (Hobaiter et al., 2014) and humpback wales (Allen et al., 2013) as well as high-fidelity copying in chimpanzees and orangutans (Horner & Whiten, 2005; Whiten et al., 2004) brings the notion of these capacities as sufficient requirements for CCE into question.

A current theory argues that explicit, system-2 processes (discussed in detail below) which enable *strategic* learning from others, rather than simply the capacity to imitate, may explain the cumulative improvement of traits over generations (Dunstone & Caldwell, 2018; Heyes, 2016). This theory has been dubbed the Explicitly Metacognitive Cumulative Culture Hypothesis (EMCC; Dunstone & Caldwell 2018). It posits that, by the time they reach adulthood, typically developing modern humans have an explicit awareness of which information is beneficial, and should therefore be copied and retained, and which information can be discarded. This allows for not only the retention and transmission of cultural traditions, but their cumulative evolution over time. Such an awareness is believed to require a metacognitive understanding of one's own knowledge and skill level, and an explicit awareness of the social learning strategy (Laland, 2004) to be applied in a particular context.

Measuring Cumulative Cultural Evolution

Demonstrating cumulative cultural evolution in the lab has traditionally been done using transmission chain methods (see Kirby, Cornish, & Smith, 2008; Zwirner & Thornton, 2015). However, this method is very labour intensive due to the large number of participants required, especially if comparing across multiple experimental conditions (for example, Caldwell & Millen (2009) tested 700 participants which may be unfeasible for some studies). If wanting to make developmental or inter-species comparisons the large sample size required may present methodological barriers that are hard, if not impossible, to overcome due to a lack of available participants (although see Horner, Whiten, Flynn, & de Waal, 2006; Reindl & Tennie, 2018; Sasaki & Biro, 2017).

Caldwell et al. (2020) proposed an alternative method of testing for the capacities for cumulative cultural evolution by testing the propensity of individuals to improve upon information they are given at any one time. Caldwell et al. called this the individual's *potential for ratcheting* (PFR). If presented with various pieces of information reflecting varying levels of success, a participant that was able to perform better (e.g. score more points in a task) after observing a better demonstration (such as one that was higher scoring) would be displaying some PFR. If they were consistently able to outperform the demonstration, at multiple levels of success, this may reflect a level of PFR that was more analogous with cumulative cultural evolution. Being exposed to higher scoring demonstrations, that score higher than would be expected by chance or random exploration, is the experimental equivalent of individuals being exposed to a trait or tool that has been modified by a conspecific to be more useful or productive than in its natural state in the environment. For example, exposure to knapped stone tools as opposed to unmodified rocks.

This paradigm has been used to test the potential for ratcheting in tufted capuchins (*Sapajus apella*) (Kean et al., in prep) and young children (Wilks et al., in press).

Cumulative cultural evolution can be conceptualised as a process of searching an adaptive landscape for an improved solution to a problem. Indeed, to a certain extent most tasks can be framed as search tasks (Hills et al., 2015; Wu et al., 2018). To test for PFR in a particular population, a task in which participants search a defined space to exploit known rewards and explore for unknown rewards can therefore function as a generalisable abstraction of a natural foraging task. There are many examples in the literature of search tasks being used to demonstrate cumulative cultural evolution. Mackintosh et al., (in prep) showed evidence of CCE when intentional information transfer was used in a grid search task with a structured reward space. Derex & Boyd (2016) found cumulative score improvement in a computer-based task requiring participants to choose between 'ingredients' to create a 'remedy' after which they were scored on the effectiveness of their remedy. There was a defined scoring system unknown to participants, meaning each trial was a virtual search of all possible ingredient combinations. Similarly, in a computer-based task that required participants to create tools using combinations of raw materials that could only be combined in pre-defined (but unknown to participants) ways, the number of tools created increased cumulatively over generations (Edmiston et al., 2018).

Grid tasks that don't involve searching for rewards in fixed locations are also used in the cultural evolution literature. Kempe, Gauvrit, & Forsyth (2015) observed score improvement in both children and adults in a task which required participants to recreate a pattern on a 10x10 grid. However, score improvement in children was the result of the children significantly reducing the complexity of the pattern they were meant to reproduce, potentially due to their limited working memory capacity. Similarly, the score increase observed in baboons and children in an anti-copying task (Saldana et al., 2019) resulted from a simplification of the input.

Measuring EMCC

To assess the validity of the EMCC hypothesis, we need to be able to test whether an observed behaviour relies on explicit processes. Here, when referring to explicit processes, we are specifically referring to the system-2 processes as defined by Evans and Stanovich (2013) in their summary of dual-processing theories. The defining element of Evans and Stanovich's conceptualisation of system-2 is that it relies on working memory. Whether PFR relies upon system-2 process can therefore be investigated by comparing performance in a PFR task with and without access to working memory resources. Wilks et al. (in press) found that PFR in young children was only evident when there was no memory requirement to the task. Under greater memory demands PFR was not present, although there was some indication that children were able to make greater use of higher (compared to lower) success demonstrations from age 6, when working memory skills are argued to be more developed (Roman et al., 2014). Task interference from memory demands was also found in a search task that required using a simple selective social learning strategy (Dunstone et al. 2020).

Restricting access to working memory resources via a dual-task paradigm has been argued to impede metacognitive responding in an uncertainty monitoring task (Coutinho et al., 2015) and a confidence rating task (Maniscalco & Lau, 2015). Additionally, a working memory load applied concurrently to a theory of mind (ToM) task reduced responses requiring belief-processing in adult participants (Schneider et al., 2012). Under some accounts, ToM is believed to rely on the same cognitive processes as metacognition, with some arguing that metacognition is actually just turning the use of one's ToM capacity on oneself (Carruthers, 2009).

Following this line of reasoning, restricting access to working memory via a dual-task paradigm should therefore be a viable method by which we can restrict usage of metacognitive processing. If trying to assess the involvement of explicit metacognitive processes on PFR, a dual-task paradigm which taxes working memory can therefore be used as a proxy to restricting metacognition directly.

The following study aimed to experimentally test the EMCC hypothesis by restricting access to participants' working memory resources while they completed a task that evaluated their PFR. Participant scores were then used to model the expected outcome if the same task was iterated over many generations. We predicted that, under a working memory load, participants would use social information less efficiently, resulting in a reduced PFR in individuals and less ratcheting over simulated generations.

Participants completed a grid-search task under dual-task conditions that applied an additional working memory load. Participants were further split by cue type: half of participants completed the task with visible cues in which squares revealed during the information trial were shown during each trial, and half with transient cues in which the reward values of squares were shown briefly and then disappeared (i.e. introducing a memory requirement to the grid search task itself). This was included to add an additional, ecologically valid, memory load to the copying strategy requirement. In many 'real world' copying scenarios the behaviour to be copied will not remain visible while it is being copied. It may not be possible to act simultaneously with the model to be copied due to limited foraging locations or tool availability. Additionally, some behaviours may not leave an obvious physical trace of having been completed, such as if a resource being foraged does not show signs of being depleted after a foraging event (e.g. water sources).

The study also aimed to investigate the impact of the same concurrent working memory load on metacognitive efficiency (defined as metacognitive sensitivity divided by overall task sensitivity – see results section). If metacognitive efficiency is significantly negatively impacted by the restriction of working memory access, this would lend support to the EMCC, and potentially shed light on the mechanisms involved in any link found between PFR and working memory load. However, if PFR is reduced under working memory load but metacognition is not this may implicate the role of system-2, but not necessarily metacognitive, processes in cumulative cultural evolution.

Methods

Design

Participants were randomly allocated to one of two main task conditions – Grid Search and Metacognition. Participants assigned to the grid search task were then randomly assigned to one of two conditions (see *Grid Search Task* below). A between-subjects design was used in order to keep the participation period short enough for testing in public. We aimed for a total participation time of around 15 minutes per participant to avoid task fatigue, low uptake or high drop-out rates (the latter two due to potential participants not wanting to take a long time out of a paid-entry science centre visit to participate).

Equipment

Participants completed all the tasks on touchscreen Lenovo Yoga laptops. All interaction between the participant and the task was via the touchscreen. No keyboard or mouse were provided. Participants wore Goji over-ear active noise-cancelling headphones. The task was written in Psychopy 3.0.3 and run in Psychopy 1.84.2 (Peirce et al., 2019). Code to run the tasks is available via the OSF: https://osf.io/dhvmy/.

Procedure

For both tasks participants were asked to touch the computer screen to begin, and were then taken through a series of on-screen instructions. They were then asked 4 understanding check questions to ensure the instructions had been taken in as intended. If any of these questions were answered incorrectly the correct answer was displayed on screen. All participants scored 50% or higher in the understanding check. Participants were then given an opportunity to ask any questions before beginning the task.

At one testing location (Centre for Life only, see below), after completing both blocks of the task participants were invited to write their scores up to be placed on a daily scoreboard that was visible to all participants and visitors to the centre. This was voluntary and participants were not required to provide their names.

Grid search Task

Each trial of the Grid Search Task consisted of an information trial in which the participant was shown 5 selections on a 5x5 grid of white squares. These appeared automatically and were displayed for 1 second before the participant was required to make their own selections, with the aim being to find rewarded squares in the grid. Participants were randomly assigned to participate with either transient cues or visible cues (between subjects):

- Transient cues (GST): vicarious information from the information trial disappeared before the participant was required to make their own selection from the grid

- Visible cues (GSV): vicarious information from the information trial remained visible (but with muted colours) while the participant was required to make their own selection from the grid.

The selections in the information trial showed between 0-5 rewards (six different levels of success) presented as green (rewarded) and red (unrewarded) squares. There were an equal number of trials with each number of rewards presented across the task, presented in a fully random order. Half a point was available for each correct selection made and there was a total of 5 green squares located randomly in each grid, so each grid was worth up to 2.5 points. Immediate feedback was given to participants about their correctness, with the selected square changing colour and an audio cue playing to indicate either a rewarded or unrewarded selection: rewarded selections made by the participant turned the grid square dark green and a 'ping' sound was played, unrewarded selections turned the square red and produced a 'pop' sound.

See figure 19 for an example of a trial.

Metacognition Task

On each trial two patches of static dots were presented on screen and participants were asked to rate which patch had a higher density by selecting a point on a scale under their selected patch. The patches were randomly generated on each trial. The target patch was always 8% more dense than the foil patch, which would be generated with a dot-density between 600 and 800 dots per patch. Whether the target appeared on the left or right of the screen was randomised for each trial, so there was an approximately equal split of targets on the left and right. This difficulty level aimed to achieve a discrimination accuracy of around 72% (based on a previous task which used 7 different difficulty levels and found 72% accuracy at the 8% level: Dunstone et al.,2020). A single difficulty level rather than a range of levels or titrated difficulty based on participant accuracy was used based on findings from Rahnev & Fleming (2019) that suggest staircase procedures can inflate estimates of metacognitive efficiency.

The scale under each patch ran from 0-100 (0 in the centre, 100 at the far left/right), with participants instructed to rate their confidence along this scale. Verbal markers were given at point 100, 50 and zero on each scale saying "definitely left/right", "50% left/right" and "Guessing left/right", respectively. The further along the scale the selection was made, the more sure the participant was that the selected patch was correct. Participants could tap anywhere on the scale and this would automatically be rounded to the nearest 10. The confidence rating would then be converted to a decimal value (for example 40% confidence = 0.4).

After participants had made their rating, the scale with the selected confidence rating would remain visible and the participant was required to tap a statement at the top of the screen that asked them to confirm they were selecting a specific side with a specific level of confidence.

Points on each trial were awarded between 0-1 for the confidence ratings, scored using a strictly proper scoring rule: points=1-(accuracy-confidence)^2. This means points were not awarded solely for correctness, but for the combination of accuracy and confidence. High scores were awarded for high confidence ratings given to correct responses and to low confidence ratings given to incorrect responses. Low scores were awarded to correct answers rated with low confidence and incorrect answers rated with high confidence. Participants were informed in the instructions that responses would be scored based on a combination of accuracy and confidence and told they could achieve the highest scores by rating their confidence accurately. Immediate feedback was given to participants about their score, and an audio cue indicated whether the selection was correct (‘ping’) or incorrect (‘pop’).

See figure 19 for an example of a trial.

Distractor Tasks

Each trial of both the grid search task and the metacognition task (herein referred to as ‘main tasks’) was sandwiched with a trial of a concurrent working memory distractor task, or a control distractor task.

Working Memory Distractor Task

Before the onset of the main task trial, two single-digit numbers were presented on screen, one in a much larger font size than the other. The participant was instructed to remember both the size of the text and the value of the numbers. The numbers were then masked, and the trial of the main task commenced. After the trial of the main task was completed a memory probe question was asked, with participants required to recall either the large or small font size, or the large or small number value. Participants were asked to tap the side of the screen that contained the correct value from the start of the trial.

Control Distractor Task

The control task followed the same task structure, but before a main task trial, instead of being shown two numbers to remember participants were shown two fixation crosses. After the main task trial, instead of a memory probe question two numbers were presented and a simple arithmetic question asking which of the two numbers presented was either larger or smaller was asked. Participants were asked to tap the correct answer.

Participants with both cue types completed a block of trials with both the working memory and control distractor tasks. Blocks were counterbalanced, so half of participants completed the control block first and half completed the working memory block first.

For both the working memory and control block, 2 points were awarded for a correct response to the distractor task, and 2 points were lost for an incorrect response. This feedback was immediate via a visual score update and an audio cue ('ping' for correct, and a gameshow style 'buzz' for incorrect). A running total of the participant's score was visible throughout the game.

Each trial of each main task had a 5-second time limit, after which the trial would timeout and move onto the next trial, skipping any remaining part of that trial. For example, if a participant had only selected two of their five selections in the grid search task within the time limit the trial would end and the memory probe question would not be asked. For trials that timed-out, only the points accrued in the first 5 seconds would be awarded. The exception to this was if during the metacognition task a participant had made a confidence rating but did not tap to confirm their selection within the 5 seconds. In these instances the trial would continue as normal after a visual prompt that the confirmation had been missed and the working memory trial would be presented. Time-limits were included firstly to ensure that participants were using their immediate working memory to recall the numbers at the end of each trial, and secondly to ensure the overall running time of the study remained short (see *design* above).

Participants

Data were collected both in a public science centre and a university lab. All Participants gave written consent to take part. Ethical approval for the study was granted by the University of Stirling General University Ethics Panel (GUEP 600-602).

Public sample

All participants in the public data collection period were tested at the International Centre for Life in Newcastle-upon-Tyne, UK. Participants were not financially compensated for their time and participation was voluntary. All participants were informed they could stop participating at any time. One hundred and seventy four participants (79:95 male:female, mean age=39.9, sd=16.17) took part, split across the two main tasks (116 in the grid-search task and 58 in the metacognition task). One participant chose to take part in both tasks. A sample size of 40 participants per condition (for a total of 120 participants) was pre-registered (shorturl.at/lszA8). This was exceeded as the sample was reached earlier than expected, but the full data collection period was still completed out of courtesy to the science centre and on the basis that a larger sample could only improve the reliability of the results. The analysis presented in the results section uses the full dataset, as preliminary analysis showed both the full and reduced data sets produced models with the same significant effects.

Eleven participants were excluded in line with the exclusion criteria stated in the pre-registration document, for either choosing to leave before the task was completed or due to technical problems that occurred while they were taking part. A further four participants were excluded based on BPS guidelines for giving informed consent: one did not check the boxes to confirm consent was given, two did not put names on their consent forms and one participant was later discovered to have learning difficulties which may have meant that they were not capable of giving informed consent.

Lab sample

A large number (nine of fifty-five) of the public sample participants in one of the grid-search task conditions appeared to have not clearly understood the task instructions, as they always copied the hint exactly and never explored any of the other locations in the grid (see *Grid Search Task* above). The task instructions were therefore updated slightly before running further participants using the Grid Search Task again in the lab (an updated pre-registration document was made: shorturl.at/lnxGL). The metacognition task was not repeated in the lab.

Participants were recruited at the University of Stirling and took part in exchange for research participation tokens which were required for course completion. All participation was voluntary and participants were informed that they could stop testing at any time. Ethical approval for the study was given by the University of Stirling General University Ethics Panel (reference GUEP600). Ninety participants (13:73:1 male: female: non-binary, mean age=21.2, sd=5.06) took part. These were split evenly across two testing conditions (45 in each). One participant was removed from each condition: one as they chose to withdraw their data after participating and one as they skipped almost 40% of trials in one block of testing. All exclusions were in line with the pre-registration document, although the registered sample size was slightly exceeded due to the nature of the study sign-up system allowing for some overbooking.

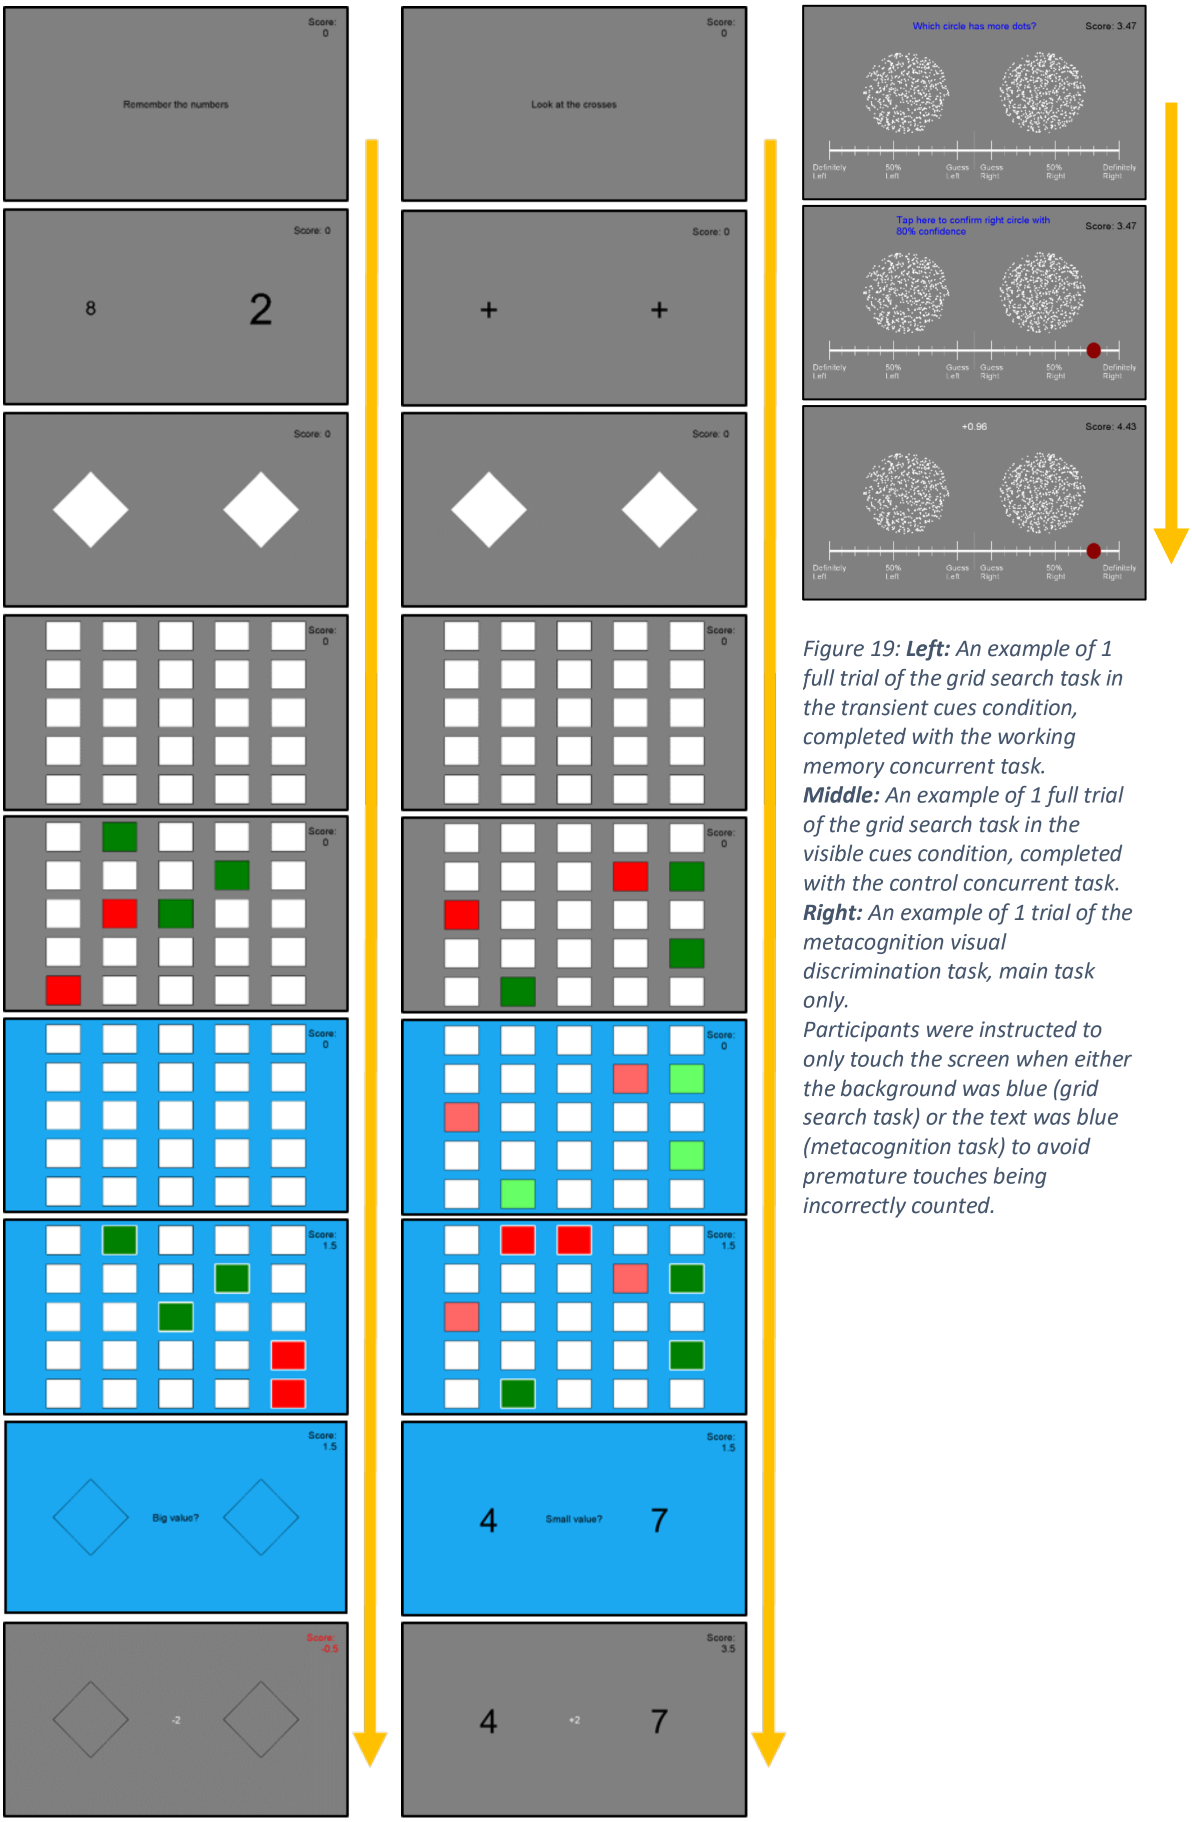

Figure 19: ***Left:*** *An example of 1 full trial of the grid search task in the transient cues condition, completed with the working memory concurrent task.*
Middle: *An example of 1 full trial of the grid search task in the visible cues condition, completed with the control concurrent task.*
Right: *An example of 1 trial of the metacognition visual discrimination task, main task only.*
Participants were instructed to only touch the screen when either the background was blue (grid search task) or the text was blue (metacognition task) to avoid premature touches being incorrectly counted.

Overall the number of participants per condition is given in table 7.

Table 7: Distribution of participants in different tasks, conditions and testing locations

Task	Centre for Life Sample	Lab Sample	Total Sample
Grid Search - Transient Cues	55	44	99
Grid Search – Visible Cues	53	44	97
Metacognition	51	0	51
Total	159	88	**247**

Results

PFR Task

Analysis was carried out on the combined sample from the Centre for Life and the lab, as the testing location was shown to have no significant impact on scores, outperformance or strategy use in the grid-search task (p=.151, p=.116 and p=.324 respectively). All analysis below is for all test trials that were completed without a timeout, with the first practise trial removed for every participant. For all models *p*-values were estimated from the resultant t-statistics with degrees of freedom being the number of observations minus the number of fixed parameters in the model (Plonsky, 2015). All models were significantly better than their null equivalent. The baseline for all models presented below is the control block with visible cues, testing at the Centre for Life.

Overall Score

Overall score on each trial (see figure 20) was analysed using a linear mixed effects model with fixed effects of block condition (control or working memory, herein WM), cue type (transient or visible), number of rewards in the information trial (rewards or 'hits') and testing location (Centre for Life or lab), and the interactions between block condition, cue type and rewards. Rewards was included as a random slope and participant ID as a random effect. Accuracy in the concurrent distractor task trial was not included as a random effect as this produced a singular fit. Significantly fewer points were scored with the transient cue type (b=-0.146, SE=0.056, t(8995)=-2.61, p=.009) and significantly more points were scored on trials where the information trial showed more rewards (b=.728, SE=0.014, t(8995)=50.6, p<.001). There were no significant effects of block condition (although this approached significance: b=-0.058, SE=0.034, t(8995)=-1.69, p=.092) or testing location (b=-0.055, SE=0.039, t(8995)=1.44, p=.151) and no significant interactions (b≤|0.075|, SE≤0.048, t(8995)≤|1.57|, p≥.117).

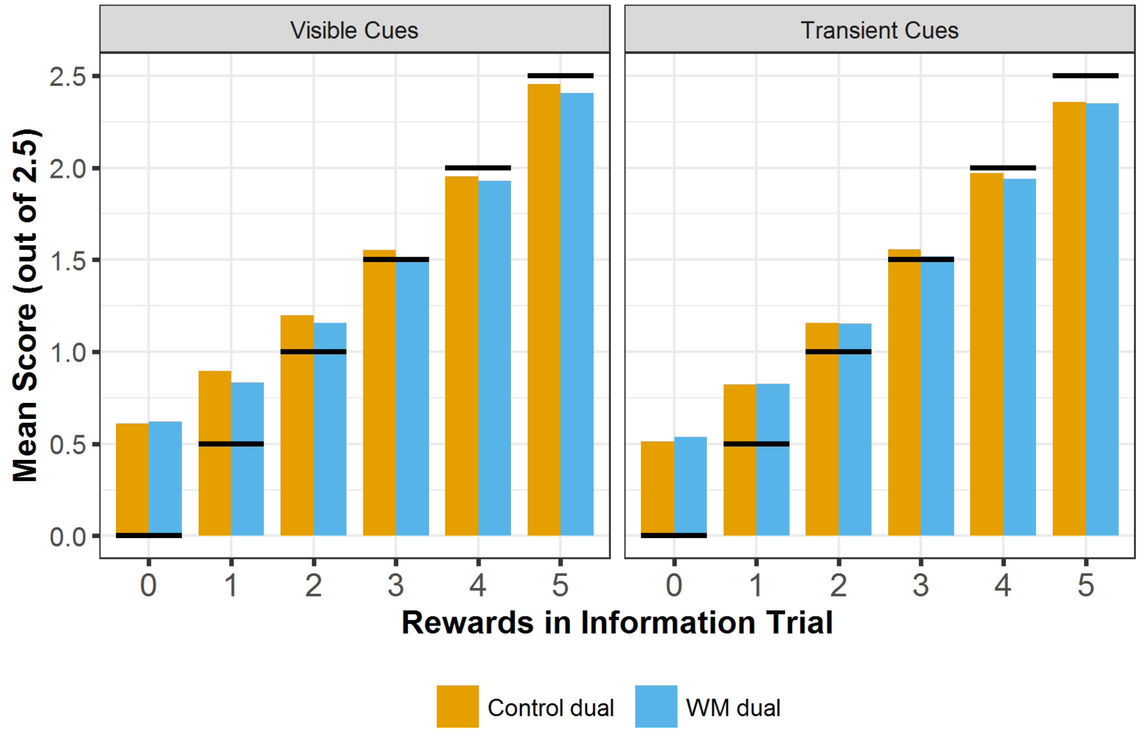

Figure 20: Mean Score gained at each number of rewards shown in the information trial hint, split by block condition and cue type. Black horizontal lines indicate the information trial score at each reward level. This is the score that would be required to pass in order to outperform the information trial. Across conditions participants easily outperform the information trial when the information trial shows few rewards, but struggle to even match performance when the information trial is high scoring.

Outperformance of the Demonstration

A measure of outperformance was calculated for each trial, which quantified each participant's potential to outperform the information trial on any given trial. This measure was included to assess not just whether participants scored more highly after observing better information, but whether they could display *potential for ratcheting* analogous to cumulative culture, by improving upon the information they have seen. This is calculated by subtracting the information trial score from the expected score on each trial. The expected score was calculated by assigning points based on strategy use: repeating a 'hit' gains half a point in the same manner as during game play, repeating the selection of a square shown to be unrewarded (a 'miss') gains 0 points and exploration of the grid scores points proportionally to the likelihood of finding a hit in the remaining 20 unexplored grid squares. This score takes into account that a participant may be using a correct strategy and still not find rewarded squares in the grid, due to the element of chance required to make successful grid explorations.

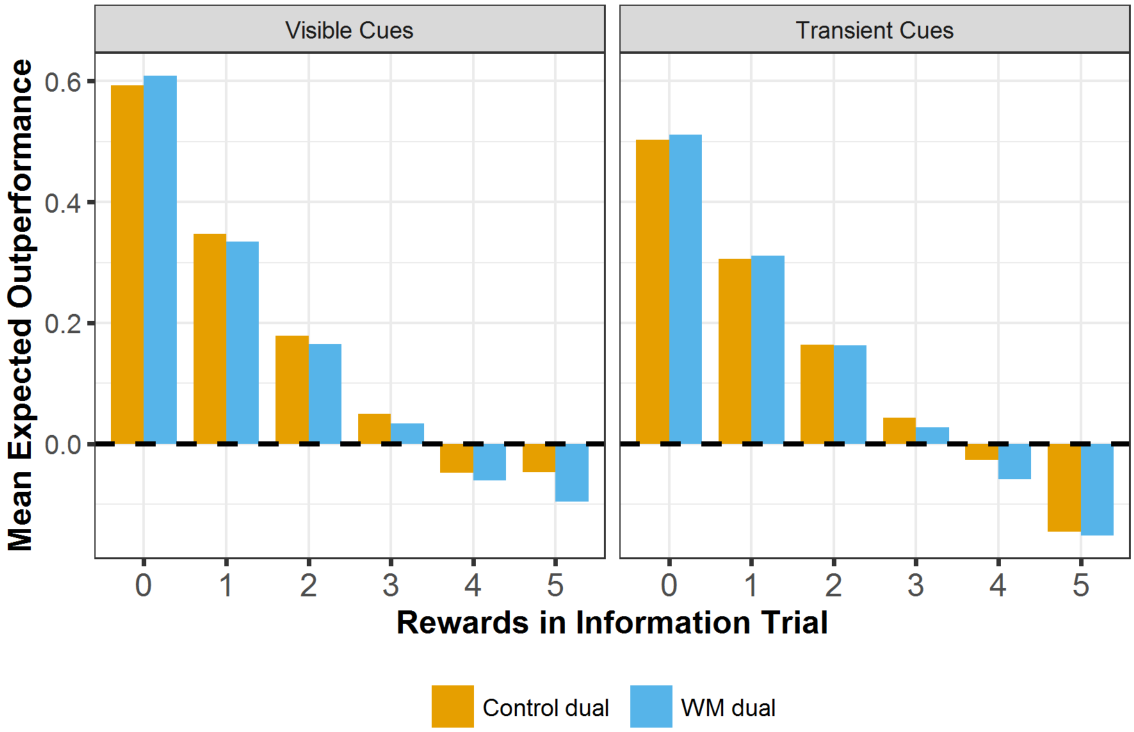

Figure 21: Mean outperformance of the information trial at each hint level. Dashed line indicates performance required to match the score of the hint. Bars below this line indicate mean performance was lower than the information trial. Based on the strategies used, participants in both blocks can outperform the information trial when it shows few rewards, but struggle to even match performance when the information trial is high scoring.

Take, for example, a trial in which the information trial showed 3 rewards. There would be 2 rewards, or 'hits' left to find in the grid out of the 20 remaining unexplored squares. If the participant repeated 3 'hits', repeated 1 'miss' and made 1 grid exploration, their expected score for that trial would be (3*0.5)+(1*0)+(1*(0.5*(2/20))) = 1.5+0+0.05 = 1.55. The information trial score for that trial would be 1.5, so the outperformance score would be 0.05. Mean outperformance in each condition is shown in figure 21.

Outperformance on each trial was analysed using a linear mixed effects model with fixed effects of block condition (control or WM), cue type (transient or visible), number of rewards in the information trial (rewards) and testing location (Centre for Life or lab), and the interactions between block condition, cue type and rewards. Rewards was included as a random slope and participant ID as a random effect. Participants were significantly less able to outperform the model when taking part with transient cues (b=-0.050, SE=0.022, t(8995)=-2.29, p=.022), significantly less able to outperform the model on trials where the information trial showed more rewards (b=-0.129, SE=0.006, t(8995)=-20.4, p<.001), and there was a significant interaction between block condition and number of rewards (b=-0.009, SE=0.003, t(8995)=-2.62, p=.009). Post-hoc analysis using the emtrends function in R indicates that this interaction is due to significantly less decline in outperformance in the control block as rewards increase, as compared to the working memory block (b=0.007, SE=0.002, z=3.14, p=.002). There were no significant effects of block condition (b=0.008, SE=0.010, t(8995)=0.848, p=.396), testing location (b=0.027, SE=0.017, t(8995)=1.57, p=.116), or the interactions between block and cue type, cue type and rewards or the three way interaction between block, cue type and rewards (b≤|0.005|, SE≤0.014, t(8995)≤|0.621|, p≥.535).

Strategy Use

Correct strategy use (repeat rewarded selections and avoid unrewarded selections, ranging from 0-5 on each trial) and was analysed using a linear mixed effects model with fixed effects of block condition (control or WM), cue type (transient or visible), number of rewards in the information trial (rewards) and testing location (Centre for Life or lab), and the interactions between block condition, cue type and rewards. Rewards was included as a random slope and participant ID as a random effect. The WM block had significantly higher correct strategy use than the control block (b=0.080, SE=0.027, t(8995)=2.93, p=.003), and the transient cue type showed significantly lower correct strategy use than the visible cue type (b=-0.708, SE=0.189, t(8995)=-3.75, p<.001). There was a significant interaction between block condition and rewards (b=-0.030, SE=0.009, t(8995)=-3.35, p<.001) and between cue condition and the number of rewards (b=0.129, SE=0.041, t(8995)=3.17, p=.002).

Post-hoc testing using the emtrends function in R shows that, although correct strategy use in the WM block is higher than in the control block when rewards are low, it increases at a significantly slower rate compared with the control block (b=0.024, SE=0.006, z=3.80, p<.001). Strategy score increased significantly more with increased rewards for transient cues only (b=-0.014, SE=0.040, z=-3.36, p<.001). When cues were transient, correct strategy use increased significantly with number of rewards, whereas when cues were visible it was consistently close to ceiling (see figure 22).

The difference in strategy use between cue types seems to be driven by a substantial number of participants in the transient-cues condition that always copied the information trial exactly, even if it contained no rewarded squares. Fifteen percent of participants taking part with transient cues copied the information trial exactly (blanket copying), in at least 50% of trials that showed fewer than 5 rewards (blanket copying when shown 5 rewards is the correct strategy). In the visible cues condition this figure is only 4.1% of participants. This may indicate that some participants were treating the grid-search element of the task, rather than the solely the working memory task, as a memory test (although see the discussion for more information).

The reverse strategy, avoiding the squares revealed in the information trial entirely, even when they showed multiple (or even all) of the rewards for that trial, was also used but by a much smaller subset of participants; 1% of participants in the transient cues condition and 5% of participants in the visible cues condition avoided the information cues entirely on at least 50% of trials that showed 1 or more rewards (avoiding the cues entirely when shown 0 rewards is the correct strategy).

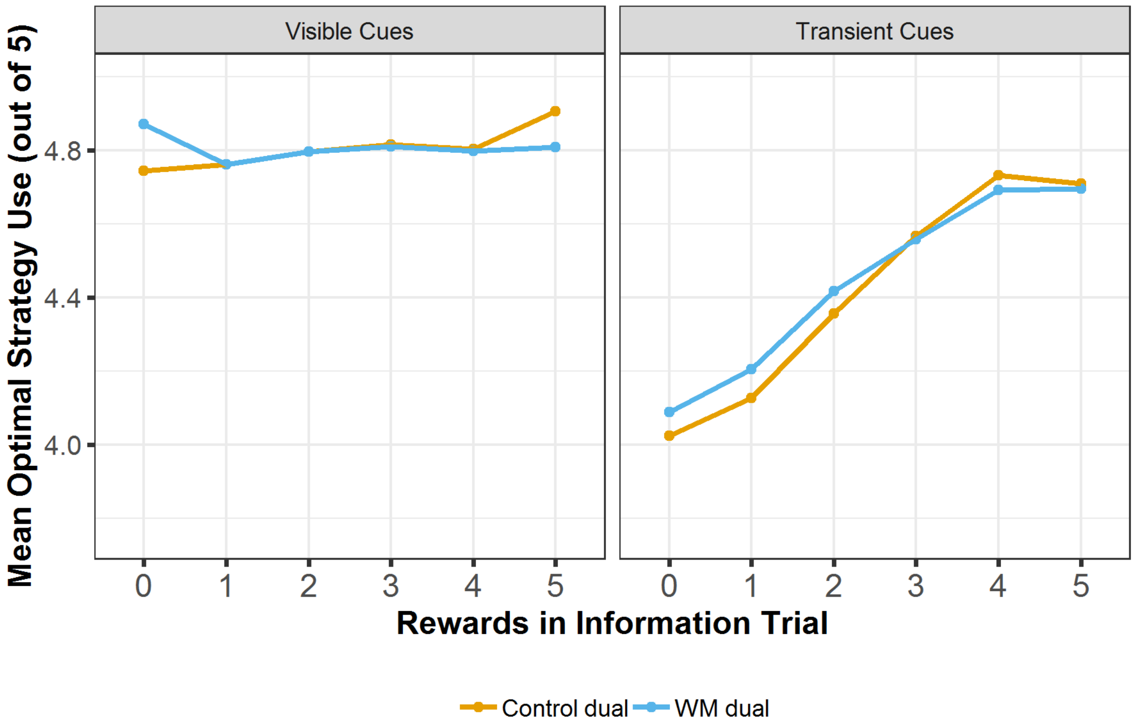

Figure 22: Optimal strategy use at each level of rewards shown in the information trial hint, split by block condition and cue condition

Analysis of Errors Made

The type of errors made in different conditions was analysed used a poisson linear mixed effects model with fixed effects of error type (omission errors, when not repeating rewarded selections, or commission errors of repeating non-scoring squares), block condition and cue type and their interactions. Participant ID was included as a random effect. There were significantly more errors made in the WM condition compared to the control condition (b=0.313, SE=0.095, z=3.31, p<.001). There was also a significant interaction between error type and block condition (b=-0.783, SE=0.145, z=-5.40, p<.001). This was due to more omission errors in the WM blocks compared to the control blocks, and more commission errors in the control blocks compared to the WM blocks (see table 8). There were also significantly more commission errors when cues were transient compared to when cues were visible (b=1.65, SE=0.127, z=13.1, p<.001). Finally, there was a significant 3-way interaction between error type, block condition and cue type (b=0.438, SE=0.180, z=2.43, p=.015). Post-hoc analysis of the interaction using the emtrends function indicates that when the cues remained visible there were differences in error types between blocks: participants made significantly more omission than commission errors in the WM block (b=0.671, SE=0.106, z=6.34, p<.001) but there was no difference in error types in the control block (b=-0.113, SE=0.099, z=-1.14, p=.256). When cues were transient both blocks had significantly higher rates of commission compared to omission errors (b≤-1.42, SE≤0.079, z≤-19.7, p<.001). Mean total errors by participant are shown in figure 23.

Table 8: Mean total errors by participant, split by cue condition, block condition and error type

Cue Condition	Block Condition	Mean Omission Errors (total per participant)	Mean Commission Errors (total per participant)
Visible Cues	*Control*	2.08	2.32
Visible Cues	*WM*	2.84	1.45
Transient Cues	*Control*	1.84	10.8
Transient Cues	*WM*	2.40	9.74

Clustering of Selections

A measure of selection distance was calculated for each set of grid selections (the five grid squares displayed in the information trial, or selected by the participant, on each trial). This took pairwise distances between each square selected in the grid on any particular trial and summed each pair to give a total value. This total value is the selection distance for each set of selections. Larger total values indicate the selections are more spaced out, and smaller values indicate the selections are more clustered. Figure 24 shows the mean selection distance for each condition at each information trial reward level.

Selections made by participants were significantly more clustered than the grid squares displayed in the information trials, as shown by a repeated measures ANOVA (information trial: 53.1, participant selections: 49.1; (F(389)=247, p<.001)). The grid squares displayed in each information trial were always selected at random, so more clustering in the participant selections indicates participants were selecting grid squares that were closer together than would be expected by chance.

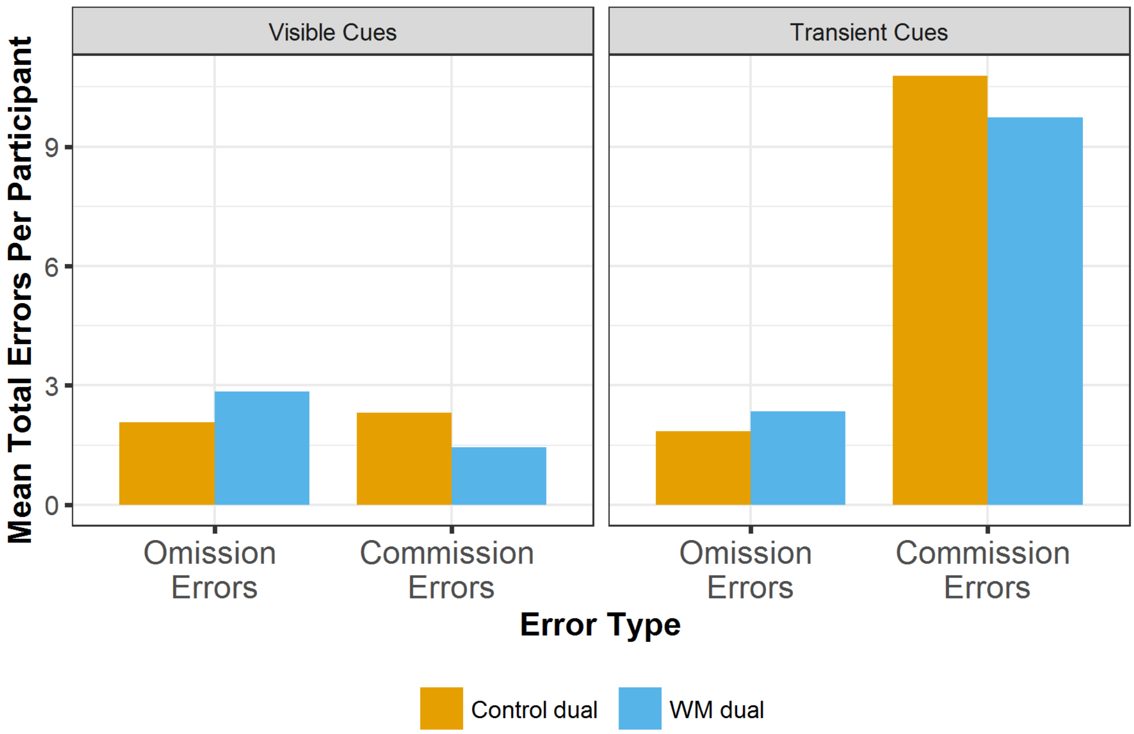

Figure 23: Mean total errors per participant, per block. Significantly more errors were made in the WM block. When cues are visible participants make more omission errors compared to commission errors in the WM block, but there is no difference between error types in the control block. When cues are transient participants make more commission errors compared to omission errors in both blocks.

The selection distance for each set of participant selections was analysed using a linear mixed effects model with fixed effects of block condition, cue type, rewards in the information trial, the grid distance of the information trial squares and the interactions between block condition, cue type and rewards. Participant ID was included as a random effect. Selections were significantly more clustered in the WM block (b=-1.53, SE=0.481, t(8995)=-3.18, p=.001). Participant selections were significantly less clustered as number of rewards in the information trial increased (b=2.11, SE=0.111, t(8995)=18.9, p<.001), and if the information trial was less clustered (b=0.377, SE=0.011, t(8995)=34.3, p<.001). The difference between WM and control selections was greater when cues were visible (b=1.41, SE=0.670, t(8995)=2.10, p=.035), although when the number of rewards in the information trial was high, the difference was greater when cues were transient (b=-0.445, SE=0.220, t(8995)=-2.03, p=0.043) (see figure 24 for both interactions). Cue type did not have a significant effect on selection distance (b=1.30, SE=0.773, t(8995)=1.69, p=.092).

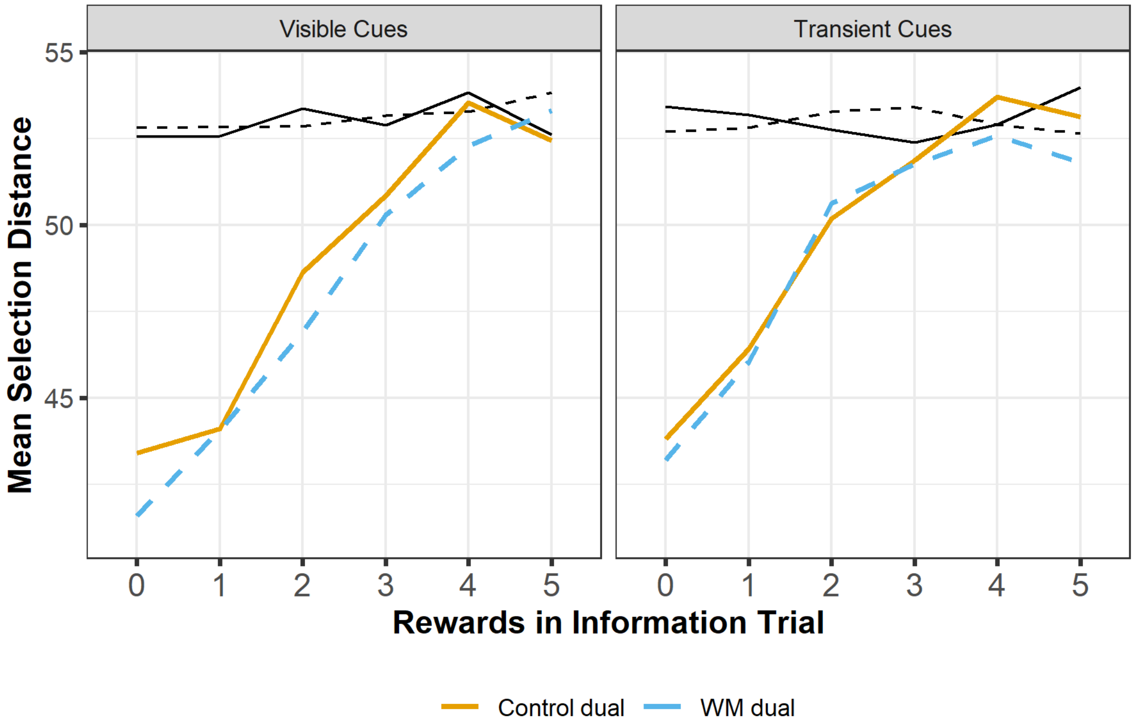

Figure 24: Mean selection distance at each reward level, split by block condition and cue type. Black lines indicate the selections displayed in the information trial. Solid lines are the control block, dashed lines are the WM block. Participants made selections that were more clustered (lower selection distance) than the information trial and participant selections were more clustered in the WM block than the control block. The more rewards shown in the information trial, the more the participant selections resembled the clustering of the information trial.

Simulated Transmission Chain

We simulated whether agents that behaved in the same way as real participants would show a ratchet effect in a transmission chain. This would indicate whether the experimental manipulations used in the task prevented cumulative improvement of scores over multiple generations. The simulation was run in PsychoPy 3.0.3 (Peirce et al., 2019). Code to run the simulation is provided on the OSF: https://osf.io/dhvmy/.

At generation 1 of the simulation an agent would receive a score based on the score of a randomly selected real participant from a trial that was shown 0 rewards. That simulated score would then form the input for an agent in generation 2. For example, if the score at generation 1 was 0.5 points due to 1 reward being found in the grid, the score of the agent at generation 2 would be drawn from a real participant on a trial where the information trial showed 1 rewarded square. The simulation ran for 40 generations and was repeated 5,000 times. Figure 25 shows the average score at each generation across all runs of the simulation. Optimal behaviour (always repeat rewarded squares and never repeat unrewarded squares) and random behaviour (select 5 random grid squares on every trial) are also shown on the graph.

The simulation shows that even though the task was simple, it was not trivial. Optimal behaviour still required around 40 generations of participants to reach the maximum possible score and random behaviour never found on average more than 1 rewarded square per trial. Agents did not perform optimally in any condition, and differences between the conditions are minimal. However, agents behaving as if they were in the control blocks did reach a higher maximum score, and reached this marginally sooner, than agents with reduced working memory access.

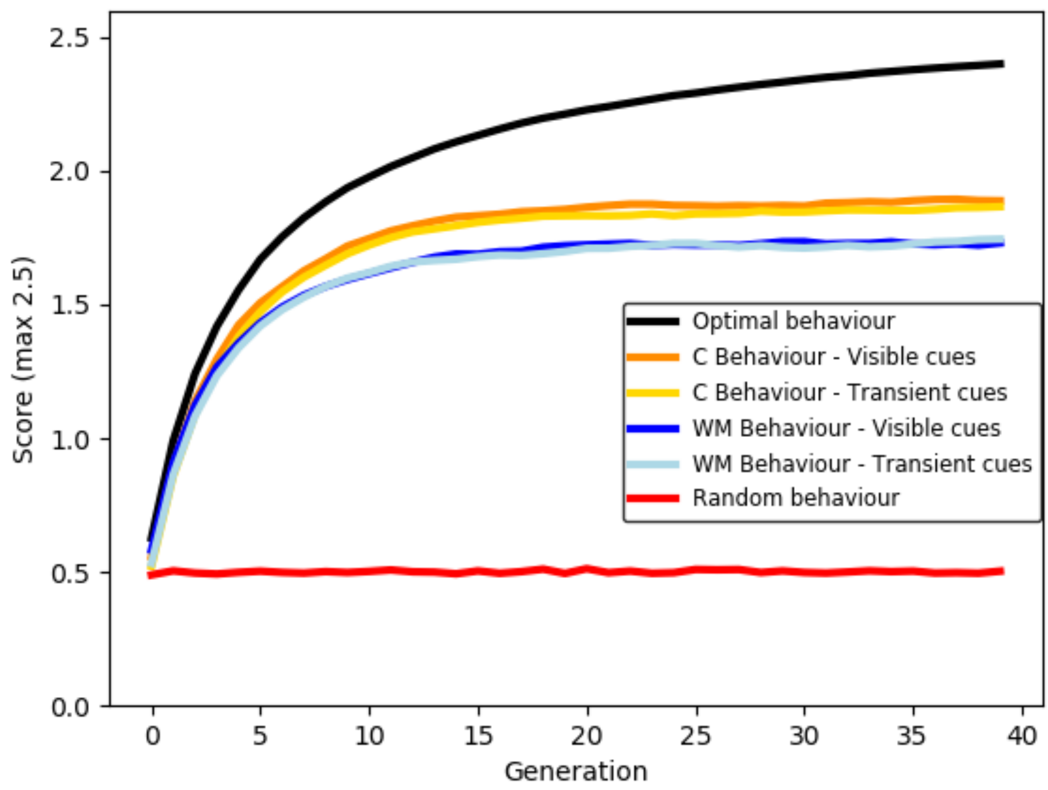

Figure 25: Simulated scores at each generation of a transmission chain created using real participant scores for each condition, compared with optimal and random behaviour.

Metacognition Task

Metacognitive accuracy was around 75% (Control block: 75.5%, WM block: 73.1%), although there was substantial individual participant variation in discrimination sensitivity (control block range: 62.5%-91.7%, WM block range: 54.2%-91.7%).

One participant was removed from the analysis due to floor effects, as their performance in both blocks was below 60% accuracy, which was not significantly above chance in either block (as shown by a binomial test: $p \geq .541$). Eight participants scored below 60% accuracy in the WM block. Two participants scored above 90% in the WM block and two scored above 90% in the control block. However, none of these participants were removed from analysis as the floor and ceiling effects were not consistent across blocks. Their unusually high or low performance in certain blocks may therefore have reflected real effects of the experimental manipulation.

A repeated measures ANOVA showed no significant difference in discrimination accuracy between control and WM blocks ($F(49)=2.0$, $p=.164$).

Participants in both blocks tended to rate their confidence highly, with much higher rates of responding at the higher end of the confidence scale (see figure 26).

Success on each trial was analysed using a binomial general linear mixed effects model with fixed effects of block condition, confidence, target side (whether the correct stimulus to pick was on the left or right) and the interaction between block condition and confidence. Participant ID was included as a random variable and the control block was taken as the baseline. The model was significantly better than the null equivalent (χ^2 (4)=43.9, $p<.001$). Participants were significantly more successful when they were more confident ($b=0.834$, $SE=0.251$, $z=3.312$, $p<.001$), although this result may be affected by the very low number of responses given at the lower end of the confidence scale (see figure 26). There was also a significant effect of target side ($b=0.409$, $SE=0.096$, $z=4.24$, $p<.001$), with a higher chance of success on each trial if the target stimulus was on the right. This shows a strong right-side bias in responses, despite an approximately equal balance of the left and right stimuli being correct.

The impact of the dual-task on participants' metacognition was analysed using the metaSDT package in R (Barrett et al., 2013). A score of metacognitive efficiency (m-ratio) for each participant in each block was calculated as metacognitive sensitivity (meta-d'_b) divided by overall sensitivity in the visual task (d'). Meta-d' is defined as the type-I sensitivity (accuracy in the visual discrimination task) that would be found if all of a participant's type-II ratings (confidence ratings) were considered to be optimal (Fleming & Lau, 2014). This measure aims to give a bias free measure of metacognition which is not affected by performance in the visual task (Maniscalco & Lau, 2012). Metacognitive efficiency, as measured by m-ratio, was fairly low, around 50% for both blocks, and a repeated measures ANOVA showed no significant difference between control and WM blocks (mean control: 0.544, mean WM: 0.406, $F(49)=0.174$, $p=.679$).

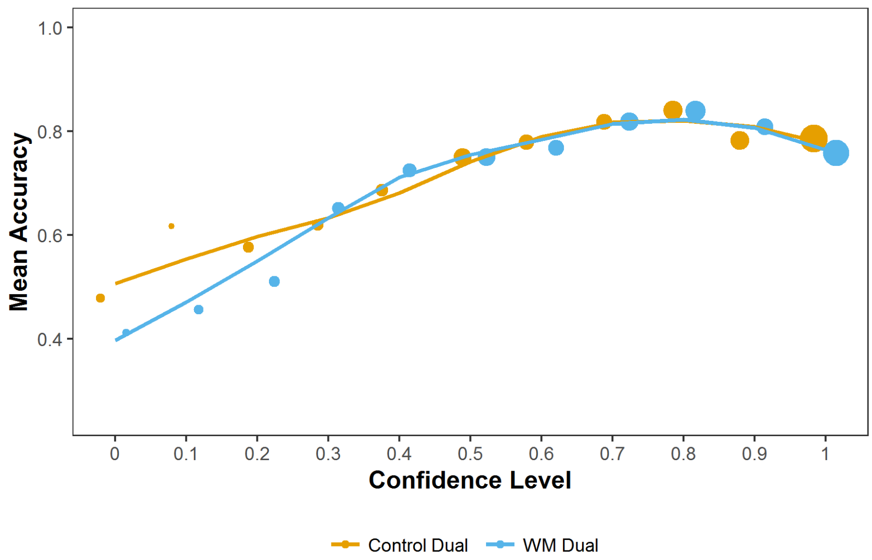

Figure 26: Mean accuracy of patch discrimination at each confidence level. Size of points indicates the number of responses given at that confidence level, with larger dots indicating more overall responses.

Concurrent Distractor Tasks

For the grid-search main task, accuracy in the distractor task was almost at ceiling in the control block for both cue conditions, and high in both cue conditions of the WM block (see table 9).

Table 9: Mean accuracy in the concurrent working memory task when completed alongside the grid search task

Block	Cue Type	Mean Distractor Task Accuracy
Control	*Visible Cues*	96.6
Control	*Transient Cues*	95.8
WM	*Visible Cues*	80.3
WM	*Transient Cues*	76.0

A binomial linear mixed effects model with fixed effects of block condition, cue type and their interaction, and participant ID included as a random effect, found accuracy was significantly lower in the WM block compared to the control block ($b=-2.01$, $SE=0.130$, $z=-15.4$, $p<.001$). There were no significant effects of cue type or the interaction between block condition and cue type ($b\leq|0.163|$, $SE\leq0.188$, $z\leq|0.869|$, $p\geq.385$). This model was significantly better than the null equivalent ($\chi^2(3)=769$, $p<.001$).

In the metacognition task, accuracy in the distractor task was also at ceiling in the control block and high in the WM block (see table 10).

Table 10: Mean accuracy in the concurrent working memory task when completed alongside the metacognition task

Block	Mean Distractor Task Accuracy
Control	97.9
WM	80.1

A binomial linear mixed effects model with a fixed effect of block condition and participant ID included as a random effect, found accuracy was significantly lower in the WM block compared to the control block (b=-2.55, SE=0.223, z=-11.4, p<.001). This model was significantly better than the null equivalent ($\chi^2(3)=220$, p<.001).

As the concurrent task accuracy was significantly lower in the WM block, additional analyses were run on a subset of the data only including trials in which the distractor task was answered correctly. This was in order to analyse data from trials in which participants were more likely to be actually using their working memory to answer the question, rather than focusing solely on the main task.

If the results from the reduced data set were different from using the full dataset, it may indicate that significant offloading of the main task demands onto the concurrent memory task was occurring. The analysis was run using models with the same fixed and random effects as in the original analysis. Unless otherwise stated, all model results stated below describe the same significant effects and interactions as the full dataset.

The model analysing points scored showed significantly fewer points were scored in the transient cues condition (b=-0.150, SE=0.058, t(7839)=-2.61, p=.009) and significantly more points were scored on trials where the hint showed more rewards (b=0.727, SE=0.014, t(7839)=52.7, p<.001).

The model analysing outperformance of the information indicated higher levels of outperformance in the visible cues condition (b=-0.047, SE=0.022, t(7839)=-2.12, p=.034) and significantly less outperformance with higher rewards in the hint (b=-0.128, SE=0.006, t(7839)=-20.8, p<.001). In contrast to the model with the full dataset, the interaction between block condition and number of rewards in the hint was not significant, although it did approach significance (b=-0.006, SE=0.003, t(7839)=-1.84, p=.066). Post-hoc analysis of the interaction, however, reveals that there is a still a significant difference between the rate of decline in outperformance between the control and WM blocks. The WM block again shows a faster decline as rewards increase, with increased *under*performance when the information trial shows more rewarded squares (b=0.006, SE=0.002, z=2.35, p=.019).

Re-analysis of optimal strategy use showed the EF block had significantly higher correct strategy use than the control block (b=0.091, SE=0.029, t(7839)=3.17, p=.002), and the transient cue condition showed significantly lower correct strategy use than the visible cues condition (b=-0.692, SE=0.191, t(7839)=-3.62, p=<.001). There was a significant interaction between block condition and rewards (b=-0.030, SE=0.009, t(7839)=-3.21, p=.001) and between cue condition and the number of rewards (b=0.127, SE=0.041, t(7839)=3.11, p=.002).

Analysis of the errors made in the grid search task again showed significantly more errors were made in the WM block (b=0.313, SE=0.094, z=3.31, p<.001). As in the full data set, there were significant interactions between error type and block (b=-0.783, SE=0.145, z=-5.40, p<.001), between error type and cue type (b=1.65, SE=0.127, z=13.1, p<.001) and the three way interaction between error type, block condition and cue type (b=0.438, SE=0.180, z=2.43, p=.015).

Analysis of clustering of selections again showed significantly more clustering in the WM block compared to the control block (b=-1.52, SE=0.519, t(7848)=-2.94, p=.003). Selections were significantly less clustered when the information trial showed more rewards and when the information trial was less clustered (p<.001 for both). All interactions remained the same, although some of the effects were more marginal, likely due to the reduced size of the data set. The full dataset showed that cue condition did not significantly impact the clustering of selections. This effect remained in the reduced data set (p=.082).

The metacognition task re-analysis showed significantly higher accuracy on trials with high confidence (b=0.935, SE=0.253, z=3.69, p<.001) and a strong right-side bias (b=0.484, SE=0.102, z=4.76, p<.001).

Given the similarity of these results as compared to the full datasets, any impact from offloading of main task demands to the concurrent working memory task was considered to be minimal.

Discussion

Grid Search Task and Simulation

Results showed that participants were able to score significantly higher on the grid-search task when the information trial hint contained more rewarded squares, and scored lower overall in the transient cues condition. The increase in scores after observing higher-scoring information was not affected by the memory load imposed by the dual-task. This showed that adults were able to demonstrate a 'potential for ratcheting' (PFR) (Caldwell et al., 2020), even in dual-task conditions.

The decrease in outperformance with an increase of information trial rewards is to a certain extent expected, as the higher the number of rewards observed, the fewer rewards remain in the grid, so the less likely grid exploration is to be successful. However, the rate of decline was significantly steeper when under additional working memory load, due to more *under*performance at the higher hint levels. This is reflected in the interaction analysis of correct strategy use, which shows better strategy use in the WM block when the number of rewards in the information trial is low but worse strategy use when rewards are high (this is discussed further below). Additionally, significantly more errors were made in the WM blocks and when cues were transient, both conditions that imposed additional demands on participants' memory. This may be due to higher memory requirements when the information trial had a higher number of rewards. Taken together, these findings suggest some working memory involvement in individuals displaying high levels of PFR. This suggests some support for the EMCC, as it implicates memory requirements (and therefore system-2 involvement) in the capacity to consistently outperform demonstrated information.

The level of underperformance at high information trial reward levels revealed substantial errors across all conditions. However, the analysis of error types suggests that additional working memory loads affected the type of errors participants made. Significantly more omission errors compared to commission errors were made in the WM blocks, and significantly more commission errors compared to omission errors occurred in the control blocks. This may indicate that the benefit of working memory access for PFR is in preventing errors, specifically omission errors, rather than guiding efficient exploration of unknown space. The errors analysis, therefore, is not consistent with the idea that system-2 processes are beneficial for CCE due to facilitation of selective copying. Rather, the analysis suggests there may be some benefit of working memory access for high-fidelity copying.

There is therefore mixed support for the EMCC. Although the specific predictions made in the introduction are not met, system-2 involvement, that relies on high memory capacity, may still be beneficial to CCE even if the benefits are not due to an increased propensity for selective copying.

The significant effect indicating a higher number of errors in the WM block appears to be at odds with the higher levels of optimal strategy use seen in the WM block. This apparent paradox is caused by the significant interaction between block condition and error type. This interaction is driven largely by a small number of participants who made substantially more of one type of error depending on the block condition. However, these participants were not treated as outliers (and thus removed from the data) for two reasons. Firstly, the outlying behaviour seems to be caused by the experimental manipulation between blocks. Secondly, many participants made a similar number of errors under different cue conditions, so the large number of errors made by some participants only present as outliers when the data is split by both error type and cue type.

There was a significant difference in correct strategy use between cue conditions. When cues were visible (GSV) there was consistently high use of the correct strategy in both control and WM blocks. However, when cues were transient (GST) correct strategy use in both blocks increased significantly with the number of hints in the information trial. This difference in strategy use was driven largely by a sub-group of participants in GST that frequently fully copied the information trial hint. These participants copied all the locations in the information trial, rather than exploring other grid locations, even when the information trial showed few or no rewarded squares. This strategy use is unlikely to be explained by a poor understanding of how to complete the task, as explicit instructions were given informing participants that they should look for 5 green squares and that they could search the entire grid for them. All participants scored 50% or higher in the understanding check. A practise trial was also included in both blocks. If participants re-selected unrewarded squares during this practice they were reminded before beginning their test trials that they should search the whole grid. In addition to this, the task instructions were updated between public and lab testing to make it more unambiguous that the entire grid should be searched for the rewards, but the blanket copying strategy persisted in the lab.

This strategy may have been employed as the task was framed as a game to test working memory, to ensure participants put as much effort into the working memory distractor task as they did to grid search or metacognition tasks. Participants may therefore have been deliberately trying to replicate the information trial rather than score as many points as possible if they thought that was the goal of the task. However, this misconception should have been removed by the task instructions and practise trials.

Another explanation for this strategy use may reflect the more challenging nature of the task in the transient cues condition. The additional cognitive load of remembering the grid locations of the information trial squares after they were hidden, as well as remembering which squares to repeat and which to avoid, may have been more challenging for some participants. They may have therefore employed a copying strategy of just repeating all observed squares to ensure that any rewards were always repeated, at the expense of missing out on further exploration. This is somewhat similar to the search strategy employed by the adult participants tested by Schulz, Wu, Ruggeri, & Meder (2019), who had a tendency to exploit known rewards rather than explore unknown space.

Support for this explanation may be found in the over-imitation literature. Schleihauf & Hoehl (2020) propose that the use of system-1 processing may lead to blanket copying, via a learned heuristic to save cognitive resources. Although the blanket copying strategy was used when the executive function load was both present and absent, it may be the case that the additional memory load introduced by the transient cues was sufficient to trigger system-1 processing. This would indicate strong involvement of memory resources in efficient selective copying. However, as shown by the simulation, the difference between cue types was almost indistinguishable after the task was iterated over multiple generations. This result therefore offers mixed support for the EMCC as, despite an indication of system-2 processes being required for selective copying to occur, similar levels of ratcheting are possible even when blanket copying is prevalent within the population. This finding contrasts with that suggested by the errors analysis, which indicates that WM resources are beneficial for high-fidelity copying but not for efficient exploration. This may reflect the difference in the type of memory load that applies to each condition. The WM block restricts access to working memory via the application of a concurrent memory load. In contrast, the transient cues add an element of memory recall to the main task itself. The difference in the impact of these two types of memory load can be observed in the simulation. The suboptimal strategy used in response to transient cues is only adopted by a subset of the overall population, so although it causes large changes to their behaviour these individual anomalies are smoothed out when part of a large multi-generational population. In contrast, the impact of the WM block is minimal for each participant, but this minimal change in behaviour is amplified over multiple generations of transmission, resulting in a marginal difference between the control and WM blocks. The implication that follows for the EMCC could be that, while explicit working memory resources are beneficial for both aspects of selective copying strategies (high-fidelity copying and efficient exploration), within a variable population of adults changes in high fidelity copying may be more likely to result in changes to the accumulation of beneficial modifications. However, this outcome may be different in a population that had generally lower working memory capacities, such as young children, and therefore might be more prone to use a sub-optimal strategy when exposed to transient cues.

The 5-second time limit for each trial, although being sufficient to make all 5 selections with some deliberation, may also have caused participants to use their system-1 for all trials, even in the control condition. This could explain the lack of differences between blocks, and could also be a cause of the sub-optimal ratcheting observed across conditions in the simulation. If the trial time-limit explains some of the gap between optimal behaviour and the control blocks in the simulation, it could indicate a substantial role for system-2 processes, if not specifically working memory, in cumulative ratcheting over generations. If so, this might offer more robust support for the explicit, although not necessarily the metacognitive, element of the EMCC. Indeed, simulated multi-generational transmission of a similar search task with no time limit restriction showed that simulated populations of participants rapidly reached the maximum possible score (Mackintosh, Atkinson, & Caldwell, in prep; transparent condition). Further study using this paradigm with the time-limits removed could potentially establish if this interpretation is substantiated (being mindful of the reasons time-limits were initially included – see *procedure* above).

Overall, the simulation showed that the behaviour of participants that show a potential for ratcheting, in a task that is not transmitted over generations, does lead to an increase in scores when that behavioural data is used to simulate multi-generational transmission. This confirms the validity of the PFR task in testing for the capacity for cumulative culture, especially as the simulation of optimal responding produces sufficient score accumulation that the maximum available score is reached. However, the similarity between all four experimental conditions despite the difference in strategies used by groups of individuals within those conditions indicates that, contrary to our predictions, restricted access to working memory does not necessarily prevent ratcheting across generations.

It also demonstrates that populations with very different behavioural strategies may produce observable behaviour that is very similar. The different pattern of error types found between WM blocks (more omission errors) and control blocks (more commission errors) can also explain some of the similarity between blocks when iterated over multiple generations. If working memory involvement does indeed play a significant role in efficient copying, as suggested earlier, the ratcheting displayed here indicates alternative copying strategies may be involved in order to compensate when access to working memory resources is unavailable.

The simulation also illustrates the likely population-level impact of participants' failure to act optimally, across all conditions. This was partly due to the sub-optimal copying strategy used by a small number of people in the visible cues condition (GSV) and a large number of people in the transient cues condition (GST). Given how challenging it is to find all the rewards in the space even with errorless behaviour, the simulation demonstrates how having even a small percentage of agents within a population following a sup-optimal strategy can have large impacts on the overall population behaviour. There were also a small number of participants who shifted away from rewarded squares shown to them in the hint, predominantly in GSV. The difference between the type of sub-optimal strategy used in the different cue conditions may explain the similarity between the outcome of the simulation for both conditions: score improvement may have been slowed by redundant repeating of non-scoring squares (commission errors) in GST, whereas it may have increased faster but showed more score decreases due to omission errors in GSV. The presence of even a very few shifts away from rewarded squares, coupled with the difficulty in finding all 5 rewards within the grid, prevents the overall population from reaching the maximum score.

It is worth drawing attention to the contrast between our findings and those from other studies of cultural evolution that have utilised task paradigms which examine copying specifically (i.e. tasks in which the optimal response is always to repeat, in contrast with search tasks such as the one used in the current study, in which rewarded selections should be repeated and unrewarded ones avoided). Studies using copying tasks have often shown steep increases in scores over generations due to participants' tendencies to form clusters of selections. This is noticeably prevalent in participants that have limited (due to age [Kempe et al., 2015] or species [Claidière et al., 2014]) or restricted (due to processing bottlenecks [Atkinson, Dunstone and Caldwell, under review]) working memory capacity. Clustering allows for more faithful replication by subsequent generations due to the increased regularity of the search space, so the task becomes easier over generations. The tendency to shape the input to be more easily learnable is a feature of many studies of cultural evolution, particularly in tasks which have conventional (i.e. copy this image) goals, and likely reflects weak learning biases of the participants (Tamariz & Kirby, 2015). The present study, in contrast, has an instrumental (i.e. score as many points as possible) goal and, as the nature of the grid search task required searching for rewards in fixed locations, the form of the information trial was never shaped by a previous participant to become more easily copied. In fact, as the information trials become more high scoring, finding the remaining rewarded squares becomes more challenging as the size of the unknown search space remains the same but the number of rewards to be found decreases. This is therefore more analogous to many examples of human cumulative culture in the 'real world'; the more accumulated knowledge required to form a trait, the more difficult it is to reproduce faithfully and the harder it is to improve upon (Mesoudi, 2011). This pattern is observed in miniature in the results presented above. In the grid search task participants struggle to match the information trial when it scores more highly. Similarly, even optimal performance in the simulation shows rapid score accumulation while there are still multiple rewards to be found in the grid but much slower accumulation when there are very few rewards remaining in the same size search space. This again highlights the value of using this method to assess capacities for CCE under different experimental manipulations.

This simulation did not allow for clusters to persist over generations, as subsequent generational 'input' was always formed from the randomly generated information trials of the original task and reward locations were fixed. This is one limitation of using a PFR method rather than true multi-generational transmission. However, some clustering was observed in the grid-search task. Participants made selections that were more clustered than the random selections displayed in the information trials and made more clustered selections in the WM blocks than in the control blocks. This indicates that, as in a task with a copy goal, participants have a weak bias to cluster their responses. The increased clustering in the WM blocks therefore suggests that, in line with the predictions of the EMCC, restricted memory access causes deficits in accumulation of benefits over generations. This deficit may be due to the restriction of explicit processes leading to greater reliance on system-1 biases, in turn leading to reduced ratcheting over generations.

Metacognition Task

There was an increase in accuracy with higher confidence, showing participants did have some metacognitive awareness. However, average metacognitive efficiency across conditions was fairly low. There was also no significant difference between experimental conditions, indicating that increased demands on working memory resources did not affect efficient metacognitive responding. Konishi et al., (2020) similarly found that metacognitive efficiency was either not affected or actually enhanced when under dual-task settings. However, the study by Konishi et al. made use of a staircase procedure in some conditions which has been found to enhance metacognitive responding. It also did not directly test working memory impacts on metacognitive ability, but instead asked participants to make metacognitive judgements on two tasks simultaneously.

To establish whether the outcome of the metacognition task supports the EMCC it should be considered in conjunction with the results of the PFR task and simulation. Contrasting results across the two main tasks provide little in the way of support for the EMCC, as our findings suggest that metacognition was not affected by the same processes as participants' potential for ratcheting.

One conclusion drawn from the PFR task may be that working memory is beneficial for high-fidelity copying, which in turn is beneficial for cumulative cultural evolution. The finding of no deficit in metacognitive monitoring with reduced memory capacity, however, indicates little support for the metacognitive aspect of the EMCC as metacognitive efficiency was unaffected by the dual-task manipulation.

As the simulation showed that some ratcheting can occur even with restricted working memory access, the finding that metacognitive monitoring accuracy was not inhibited by the memory load does not rule out the role of metacognition in facilitating that ratcheting. However, it still offers limited support for the EMCC as it shows no evidence that metacognition and ratcheting are disrupted by the same task manipulations.

A final interpretation is that the PFR results were influenced by the time constraints on each trial of the main task. Time limits may have prompted participants to use their system-1 processing throughout, which may explain the sub-optimal responding from participants even in the control block. As the same time limit was also used during the metacognition main task, this interpretation may explain why metacognitive efficiency was similarly low in both testing blocks. A replication of the metacognition task with no time limits, or extended time limits, could provide more evidence to support or refute the EMCC. If restriction of explicit processes throughout the experiment was reducing differences between the control and WM blocks, this may indicate a substantial role for explicit processes in ratcheting over generations and metacognitive monitoring, even if working memory specifically is not implicated.

Two additional considerations should also be explored when interpreting the metacognition task data.

Firstly, the outcome of the present study might have been affected by the strong biases towards responding at higher confidence levels and responding to the right-side stimulus. The side bias was so strong in some participants it appears that they may have been responding on the right not because they genuinely thought the right-hand stimulus was the target but because simply because they used their right hand to respond and this biased them towards touching the right hand side of the screen.

This may also have affected the bias towards high confidence levels: if some participants were always selecting with their right hand and so always selecting the right-side edge of the screen they would also have been selecting from the higher end of the confidence scale for the right-hand stimulus. A small number of participants had a bias towards selecting the left side, which could also result in higher confidence values given for the left-stim simply as a result of selecting more often on the left edge of the screen. This combination of right- and left-side biases may reflect the handedness of participants in the sample, but this cannot be confirmed as handedness was not collected in the participant information.

Although the meta-d' method of assessing metacognitive efficiency is designed to be robust against type-I and type-II response bias, the measure can be affected by very high or low values for both type-I and type-II hit and false alarm rates (Barrett et al., 2013). With such a strong bias towards one response side and the high end of the confidence scale, the type-II false-alarm rate for some participants would be very high. This may then indicate that the lack of a difference between conditions is exacerbated by meta-d' measures that have been influenced by extreme inequality in the responses. This potential limitation could be combatted in future testing by changing the orientation of the stimuli to make it less likely that participants would always respond with the same part of the scale.

Secondly, the metacognition task investigated only metacognitive monitoring accuracy. It may be the case that the experimental manipulations had an impact on metacognitive *control* processes that were not detected, leading to an underestimation of dual-task impact on metacognition. It is easily conceivable that while metacognitive monitoring is beneficial in assessing the level of one's own knowledge, it is metacognitive control that enables the use of the selective copying strategies that are argued to be a requisite for CCE. This interpretation would therefore suggest that future testing paradigms should examine the role of both metacognitive monitoring and control on potential for ratcheting.

Conclusion

This study aimed to establish whether having reduced working memory resources negatively impacted capacities for cumulative cultural evolution in humans. We experimentally tested the EMCC (the hypothesis that system-2 executive resources and explicit metacognition are required for cumulative cultural evolution to occur). This was examined by manipulating access to working memory while participants' *Potential for Ratcheting* was tested, via use of a dual-task paradigm. The results indicate partial support for the EMCC: there is some evidence of explicit processes facilitating efficient copying, and the restriction of explicit resources leading to reduced capacity to perform optimally in a grid-search PFR task. No clear evidence of metacognition being impacted by the same task manipulations was found.

The impact of reduced working memory resources on metacognitive monitoring appeared to be minimal. This is in contrast to previous findings (e.g. Maniscalco & Lau, 2015), but may have been distorted by short time constraints within the task or strong response-side biases in responding. Differences in metacognitive control may also have occurred but not been detected by the task. This finding can therefore neither conclusively support nor refute the metacognitive element of the EMCC.

In the grid search task there was some evidence for decreased working memory resources leading to a reduced ability to copy selectively and efficiently. When more copying and less exploration was required, the use of optimal search strategies was generally lower in the WM block compared to the control block. Additionally, participants made more errors in the WM block, particularly errors in accurate copying of known rewards. Participants in the transient cues condition may also have been using a learned system-1 heuristic to search the grid, due to the additional working memory load imposed by the disappearing cues. Finally, participants were more likely to cluster their responses in the WM blocks compared to the control blocks, which may indicate increased reliance on an implicit bias. These findings suggest that selective social learning strategies may be more efficient if the copier has full access to their working memory resources.

When iterated over multiple generations, differences between the blocks and conditions were no longer apparent. Similar levels of ratcheting were seen with both full working memory access and substantial working memory loads, although ratcheting was also sub-optimal in all experimental conditions. The results presented above therefore indicate that, although beneficial for efficient copying, abundant working memory resources are not a necessary requirement for cumulative score improvement over generations. However, short time-constraints on each trial of the grid-search task may have impacted participants' access to explicit processes, even in the control block, contributing to an artificially large gap between the control block and optimal responding, and possibly supporting the idea of a far more substantial role for explicit processes in ratcheting potential.

Worthwhile directions for future research on this topic might include direct investigation of the effect of varying time constraints on each trial, and comparisons between metacognitive control and monitoring processes.

Chapter 4: General Discussion

Summary of Aims and Hypotheses

This thesis aimed to add to the body of literature investigating the origins of uniquely human cognition. It aimed to provide the first empirical evidence to evaluate theories that propose system-2 processes, that underpin selective social learning strategies, are a requisite cognitive capacity for cumulative culture.

Despite a large body of literature on cumulative cultural evolution (CCE), to date there has not been a satisfactory explanation for the uniqueness of this capacity in humans. However, in recent years a group of theories have been written that suggest explicit, system-2 processes may be responsible for human capacities for cumulative culture. These theories suggest that explicitly metacognitive processes provide humans with two main benefits:

1. To be able to verbally state not only their current mental state or the state of their environment, but to be able to articulate their level of confidence in that state. This is argued to provide benefits to coordinating group action, facilitating CCE by means of improving the overall abilities of groups (Shea et al., 2014).
2. To be able to evaluate potential social learning situations and apply deliberate, selective strategies to only obtain beneficial information from your surroundings or conspecifics, and to innovate if beneficial information is not available (Heyes, 2016; this thesis, chapter 1). This is argued to facilitate CCE by ensuring that traits that are copied are beneficial, thus maladaptive or unhelpful traits are less likely to spread and the overall functionality of traits can improve.

The first benefit has been investigated extensively by Bahrami and colleagues (Bahrami et al., 2010, 2012; Hertz et al., 2016) who have shown that explicit sharing of confidence can improve coordination and overall performance in group tasks, and mediate the effects of 'equality bias' - assigning too much credit to an unreliable group member (Mahmoodi et al., 2015). Olsen, Roepstorff, and Bang (2019) have also shown that participants who displayed higher levels of metacognitive accuracy benefitted more from social learning and social interaction in an experimental task.

The second benefit, dubbed the Explicitly Metacognitive Cumulative Culture Hypothesis (EMCC) (chapter 1), had, prior to the work in this thesis, not been empirically tested. A recent study by Miu et al, (2020) found evidence to suggest that flexible social learning has tangible, cumulative benefits in a real-world setting. However, this study was not able to investigate *how* this flexible learning occurs or the reasons for individual variation in copying styles. This thesis therefore aimed to begin specifically testing the theory that it is access to explicit, metacognitive processes that facilitates the flexible learning that results in accumulation of benefits.

Overall, this thesis aimed to answer the question: does explicit metacognition facilitate cumulative cultural evolution in adult humans? However, this question cannot be directly answered as a direct comparison cannot be made between participants that have access to a modern adult human level of metacognition, and participants that don't, as metacognitive capacities cannot be totally removed experimentally. The research question we aimed to answer in this thesis was therefore modified to ask: does *restricted access* to explicit metacognition *reduce capacities for* cumulative cultural evolution in adult humans?

Assumptions

The methods I used throughout this thesis relied on the following assumptions or theoretical approaches:

Firstly, when talking about 'explicit processes' I adopted the Evans & Stanovich (2013) stance on dual-systems theories. This means I did not assume strict rules or segregation of what is and is not an explicit process, but assumed explicit processes rely on working memory, and are slow and deliberate to execute. A *default interventionist* (J. S. B. T. Evans, 2007) viewpoint is assumed, which presumes that system-1, heuristic processes will be used as a default unless system-2, analytic process intervene.

Secondly, I have been, at least initially, relying on the links between executive function and metacognition to use executive functions as a proxy for metacognition. This is not intended to imply they are one and the same. But, as metacognition relies on executive functions (Roebers, 2017; Roebers & Feurer, 2016), if executive functioning is not available this would arguably impair metacognitive processes. This assumption has been challenged by some of our empirical findings.

Thirdly, I have measured CCE not by increased task performance over generations, but by individual participant performance in tasks that are designed to test their potential to improve upon individual pieces of information they have received. This has been dubbed an individual's *potential for ratcheting* (PFR; Caldwell et al., 2020). This is explained in detail in chapter 3. This means I was primarily interested in whether participants were able to use a flexible social learning strategy to copy beneficial information and avoid non-beneficial information by exploring other available locations. This strategy would, in an open-ended, iterated task, result in cumulative improvement over generations as beneficial solutions are retained and actively sought out over non-beneficial ones. The efficient use of a flexible social learning strategy is therefore considered a necessary (although not sufficient) requirement for CCE.

Chapter 2 Summary

In chapter 2 I investigated whether restricted access to executive functions (EF) could negatively affect participants' abilities to selectively copy information provided by a computer opponent. This would give an indication of whether participants were able to use flexible strategies efficiently even when their executive function resources were unavailable. To investigate this I used a series of tasks with a dual-task paradigm. Participants completed a simple visual search task in which they needed to flexibly switch between copying and not copying their computer opponent. This task was completed alongside a concurrent distractor task that aimed to put an additional load on their executive function resources.

Initially (chapter 2, section A) I aimed to establish which, if any, dual-task would prevent efficient selective strategy use. Of the tasks trialled, tasks which restricted access to *task-switching* and *working memory* executive functions both negatively impacted the speed of responding in the search task.

I then compared the impact of a *task-switching* dual-task on a flexible strategy with the impact on two single-response strategies: always copy and never copy (chapter 2, section B). The *switching* task was chosen over the *working memory* task as it was faster to complete so it allowed more trials to be completed within the same time duration of testing. The results showed significant dual-task interference in all conditions, but no difference between conditions in how much impact the dual-task had.

In order to rule out ceiling effects influencing the results found in 2B, in chapter 2 section C I made the search task more challenging by increasing the number of stimuli to search and including an additional memory load for half of the participants. This change also meant that on some trials participants would be required to use a flexible strategy within a trial, rather than only between trials. Participants also completed a metacognitive monitoring task in conjunction with the same *switching* dual-task.

The findings from chapter 2 suggested that, while reduced access to *switching* executive functions can have a significant negative impact on flexible copying, this can also have a negative impact on blanket copying that requires no strategy. As such, one explanation of the dual-task effects found in 2A and 2B is that they may just reflect the increasing difficulty of doing *any* task with reduced EF resources. However, flexible strategies did prove to be more cognitively demanding than a simple combination of the two component strategies the flexible strategy was built on.

Between 2B and 2C the *switching* distractor task was changed slightly. I removed an additional memory component that required participants to keep constant track of the previous switch cue, as results showed that accuracy in the distractor task decreased rapidly as trials progressed. I hypothesised that this may have been due to accidentally missing a single cue to switch response strategies, which would lead to incorrect responding even with full effort applied to the task, and that this decline in accuracy may have been contributing to the high failure rate in the task. When removing this working memory component from the distractor, the impact of the dual-task was also removed. This indicates that competition for *switching* executive function resources did not have a significant impact on participants' ability to use a flexible strategy. A separate working memory load did produce significant differences in strategy use, prompting the apparent use of the metacognitive control strategy of 'planning to remember': when responding on trials that had an EF-load, a flexible strategy requirement, and additional memory requirements participants would plan both responses before making either, and give them together. This initial delay in responding resulted in overall faster responses. It therefore appears that *working memory* rather than *task-switching* may be an important factor in the capacity to flexibly switch between different copying strategies. The use of a metacognitive strategy in the search task was in direct contrast to the findings from the metacognition task, which showed no difference between control and EF-load conditions. Taken together, these results may indicate that the EF interference found in the search task resulted from memory requirements, rather than metacognitive requirements, of selective social learning.

Overall, chapter 2 offered limited support for the EMCC. While a multigenerational effect was not looked for in this chapter (see 'assumptions' above) support for the hypothesis would be suggested if restriction of executive functions reduced participants' capacities to use flexible strategies over and above the use of non-flexible strategies. This was not found in 2B. It was also not found in the way I predicted in 2C. The dual-task paradigm did not produce a difference between strategy groups, or in metacognitive responding. However, the addition of a memory component did produce differences in how participants approached the task and revealed a metacognitive control strategy. This may indicate that the lack of dual-task interference on processes assumed to be necessary requirements for CCE was due to the dual-task itself, rather than because system-2 processes are not a prerequisite for CCE in humans. My subsequent experiment therefore aimed to approach the EMCC from a different angle. A working memory dual-task was used, in order to assess the impact of a different aspect of executive function on selective SLSs. A different main task, that allowed for investigation of selective copying at multiple different difficulty levels and allowed for examination of multi-generational effects, was used. These differences in experimental design also allowed for investigation of the EMCC without the confounds of main-task ceiling effects.

Chapter 3 Summary

In chapter 3 I used a working memory dual-task in conjunction with a grid-search task that assessed participants' *potential for ratcheting* (PFR). This meant assessing individuals' capacities to perform better in a task after seeing better quality information, and to outperform information they have observed. Compared to the task used in chapter 2, this method had the benefit of examining whether participants could successfully apply an appropriate copying strategy when the task search space was much larger. It also gave a wider range of information quality to be copied, which meant I could:

- Directly assess whether scores were higher after observing higher quality information
- Assess whether participants were able to improve upon the information they had observed
- Simulate how scores would increase if the task was iterated over multiple generations

Being able to assess multiple different levels of information, and simulate the outcome of an iterated task, allows for greater insights into the possible effects of system-2 cognition in 'real world' cases of CCE.

Each of the above factors was compared with and without direct access to working memory resources. This aimed to establish whether access to working memory, and therefore system-2 processes, are a requisite for a ratchet effect to occur. The same working memory dual-task was also used in conjunction with a metacognitive monitoring task.

The *working memory* (WM) distractor task, as used in chapter 2A, was used here rather than the *switching* task used in 2B and 2C. This was partly due to the findings from chapter 2 that suggested working memory may play a greater role in selective social learning than other executive functions such as switching. This task was also chosen due to the practical reasons of finding a dual-task that fit well with the PFR grid-search task. Any of the audio distractor tasks used in 2A that required constant mouse responses would have been very difficult for participants to complete alongside the grid-search task, simply due to coordinating both hands while making the number of selections that were required to complete both tasks simultaneously. Audio tasks would also have been very difficult to complete in a public setting, due to the unavoidable external noise in environments such as science centres.

The findings from chapter 3 showed that pressures from the working memory dual-task had limited impact on participants' capacity to perform better after observing better information or to improve upon their information trial: PFR did not seem to be selectively affected by the WM dual-task in comparison to the control task. However, the memory constraint introduced by the cue condition did have a significant impact on the strategies participants used to search the grid. The working memory distractor also affected the number and type of errors participants made. However, when iterated over 40 simulated generations, differences between conditions in the final scores reached were no longer apparent. Cumulative score improvement was observed in all conditions, with minimal differences between control and test conditions, although in all cases this accumulation was still far below optimal.

Accuracy scores and metacognitive efficiency in the metacognition task were also not negatively affected by the dual-task paradigm. However, participant engagement in the task and other confounding factors, like a very strong bias for responding on one side of the screen, may have impacted this result. Future testing using this paradigm could alleviate this side bias by placing the stimuli in a different orientation, for example vertically rather than horizontally stacked.

These findings indicate that restricted access to working memory resources does not restrict potential for ratcheting. Reduced working memory access may increase the number of errors made in efficient copying but, as shown by the simulation, this does not necessarily prevent ratcheting over generations. It also did not restrict metacognitive monitoring ability (see below for more detail). Chapter 3 therefore also offered little support for the EMCC.

Methodological Limitations

Lack of Impact on Metacognitive Efficiency

Overall, despite working memory load having an impact on response strategy in both the search task (chapter 2) and the grid-search task (chapter 3), and the combination of the dual-task and a working memory load prompting a metacognitive control strategy in 2C, the dual-task appeared to have little to no impact on metacognitive efficiency, as measured by the metacognition task.

This may suggest that the dual-tasks used were not sufficiently inhibiting of executive resources to produce a significant difference between blocks of the metacognition task. This theory is feasible in 2C, as dual-task interference in the search task was only found in conjunction with an additional working memory load. However, it is unlikely to be the case in chapter 3, as the task used is so similar to that used by Coutinho et al. (2015), that found concurrent completion of a dual-task produced significant deficits in accuracy in an opt-out metacognition task. Alternatively, it may be that even the control rounds of the dual-tasks were too hard, blocking access to system-2, and meaning system-1 was therefore in use all the time. This again could be plausible in chapter 2, as while the control block did not require dedicated access to a specific executive function, it did require constant attention. Again, however, in chapter 3 this theory seems unlikely as the control block of the working memory dual-task required no concurrent cognitive effort as responses were only required between trials.

It could be the case that while access to explicit metacognition is beneficial for ratcheting, cognitively modern humans can use other aspects of their cognition to achieve the same ends if required. This would mean that experimentally restricting access to metacognition had limited effects on PFR. Alternatively, it might be the case that explicit metacognition is uniquely beneficial to CCE in some real social learning contexts, where the benefit of confidence sharing is tangible due to improvements in social coordination (see above). However, this too might not necessarily require metacognition, just a shared method for conferring information and certainty.

A final reason that may have influenced the dual-tasks' effect on the metacognition task could be differences between metacognitive monitoring and control. The metacognition tasks used in chapters 2 and 3 measured participants' metacognitive monitoring accuracy – how sure they were of the selection they were making. However, it offered no opportunity to mitigate any uncertainty with a behavioural strategy - metacognitive *control*. It is conceivable that involvement of explicit metacognition in CCE relies at least as much on the capacity to modify one's own behavioural learning strategies, as it does on an awareness of one's own knowledge state. Alternatively, it may be that the dual-tasks used throughout this thesis inhibit metacognitive *control* processes, and as such are unsuitable for inhibiting the metacognitive *monitoring* processes required for the confidence tasks used in 2C and chapter 3. Indeed, the 'planning to remember' strategy observed in 2C indicates a metacognitive control strategy employed to mitigate the impact of increased working memory and executive function task loads.

Impact of Time Restrictions

A confounding factor that may have impacted all of the tasks is time pressures applied to each trial. In every task either a specific time-limit was applied, or participants were told to answer quickly. This pressure to respond rapidly may have been sufficient to trigger system-1 processing, as system-2 processes are slower and more effortful (Evans & Stanovich, 2013). The allotted time on each trial was always sufficient to allow some deliberation, but being told there was a time limit may have meant some participants had slightly heightened anxiety levels which made them answer more quickly than they needed to or made them rush. If participants were using system-1 processing regardless of the presence of the dual-task, it could explain several important findings from the studies presented in chapters 2 and 3.

It may explain the lack of significant differences between control and EF blocks in some tasks. Metacognitive efficiency was generally fairly poor, which may reflect that participants' metacognition was restricted in both task blocks. Manipulations of the main-task, rather than the dual-task (such as the additional memory loads in 2C and the transient cues condition in chapter 3) may also have caused differences in strategy use if participants' use of one system over the other was affected by the main task rather than the dual-task.

Consistent use of system-1 could also potentially explain why performance across conditions was sub-optimal in chapter 3, and why, although there was some ratcheting over simulated generations it was far from perfect. The sub-optimal performance in the transient-cues condition was largely caused by the large number of participants that copied redundant information from the hint, which could be linked to the use of system-1 processing. Schleihauf & Hoehl (2020) propose a dual-systems theory of over-imitation, in which blanket-copying of unnecessary or redundant information may be due to a learned but automatic copying heuristic. Participants may have been using this heuristic in the transient cues condition due to the short time limit and additional requirement to remember which grid squares were revealed in the hint. Although the time-limit was the same in the visible cues condition, this condition required less information to be processed in the given time. This therefore gives participants more opportunity to realise there is potential to score higher if avoiding squares known to be unrewarded, and to employ this strategy over an automatic blanket-copying heuristic.

However, time limits on each trial were important elements of the task design. For the tasks in chapter 3 a short time limit was required. With unlimited time participants would potentially no longer be using their immediate working memory to recall the numbers at the end of each trial. Additionally, pilot testing for the tasks in chapter 2 showed some participants would take such a long time to respond to some trials, when no time limit was in place, that they far exceeded the expected duration of the study. Speaking to participants during debriefing revealed this was often due to over-thinking what was intended to be a simple task. Running the study with the potential for over-running the predicted study length was not a viable option for testing in a science centre, as participants that felt the study was too long would often leave without finishing. Pilot testing for chapter 3 in Glasgow Science Centre and a science event at the University of Stirling showed that many participants would ask to leave before the end if the task took longer than around 10-15 minutes. Even with the imposed time limits on trials, around 5% of participants that took part at the Centre for Life left before they had completed the task in full and had to be removed from the dataset in line with the exclusion criteria.

A large number of participants in the grid search task (chapter 3) missed a large number of trials because they would *only* copy the rewarded squares shown in the hint and wouldn't explore the rest of the grid at all or reselect the unrewarded squares from the hint, despite the fact this made the task time out on those trials. This was despite full on-screen instructions being given, an understanding check being passed and a practise trial at the start of each block. On tasks with no rewarded hint-squares some participants were reluctant to touch the screen at all. Removing a time-limit may have eradicated this problem. However, I think this would have come at the cost of extremely long delays on early trials, especially on trials with no or few rewarded squares in the hint, leading to an increased rate of participants choosing to end participation early.

An alternative way of testing the EMCC, that would eradicate some of the issues described above while still manipulating access to metacognition, could be to compare the PFR of participants that had received metacognitive training with those that had not. Metacognitive training has been shown to be effective at improving metacognitive monitoring skills, across task domains (Carpenter et al., 2019). Participants that were given feedback on their metacognitive ratings in a visual discrimination task showed significant improvements in their metacognitive monitoring accuracy when subsequently tested on both the same task paradigm and a novel, untrained task. Improvements were not observed in a control group that received feedback on their discrimination accuracy only. A testing paradigm can be easily imagined in which participants complete similar main tasks to the ones in chapters 2 and 3, that aim to measure participant's PFR, in which a test group receives prior metacognitive training and a control group does not. Group scores could then be compared to see if enhancements in metacognitive monitoring provided enhancements in ratcheting task performance, either through an increase in scores obtained or a steeper increase in cumulative score improvement. This paradigm could remove potential confounds of dual-tasks being too challenging or not challenging enough, as well as the potential of participants using alternative strategies to complete the task even with their metacognition restricted. It would also remove time limit confounds by removing the requirement for time-limits on each trial altogether for a number of reasons. Firstly, if no working memory manipulation was occurring between trials, there would be no pressure to keep trials short. Secondly, as there would be no dual-task manipulation between trials, the length of participation in each testing session could be reduced, so the risk of the testing session running to an unacceptable length would be lower. Finally, participants that had willingly undergone the metacognitive training prior to participation might be more likely to participate in a long task.

It should be noted that this alternative method is not a perfect solution. It is not necessarily the case that improved metacognitive skills would improve ratcheting potential. Modern humans without additional metacognitive training have the capacity for CCE, so it may not be possible to enhance metacognitive monitoring ability sufficiently that it improved PFR in an experimental setting. Trying to restrict access to metacognition or reduce metacognitive accuracy via the use of dual-tasks therefore seemed a logical first step in testing the EMCC. Training similar to that used by Carpenter et al. (2019) would also not remove any of the limitations discussed above in distinguishing between metacognitive monitoring and control.

Novel Contributions in a Relatively New Field

The field of study linking explicit metacognition (Heyes, 2016; Shea et al., 2014), or explicit processes more generally (Heyes et al., 2020) to cumulative cultural evolution is still in its relative infancy. There is, to my knowledge, no existing empirical work directly investigating the impact of explicit processes or metacognition on CCE in the lab. The precedent upon which to base the methodological decisions used throughout this thesis therefore had to be taken from studies with different goals and research questions. This meant the paradigms used may have proven less effective than they have in other fields, due to the different constraints that apply in different testing contexts. For example, Coutinho et al. (2015) found significant interference from a working memory dual-task, that inspired the distractor task used in chapter 3, on an opt-out uncertainty monitoring task. However, Nicholson et al., (2019) have recently shown that opt-out paradigms such as the one used by Coutinho et al. measure first-order decision making and do not require the type of explicitly metacognitive representations that explicit judgements of confidence rely on.

Additionally, best practise for how to test certain aspects of the hypothesis has not remained constant over the course of the data collection period. For example, the finding that metacognition tasks using multiple different difficulty levels may be detrimental to accurate recording (Rahnev & Fleming, 2019) was published after the studies in chapter 2 had been completed. The evidence to suggest metacognitive training can produce benefits in metacognitive efficiency across domains (Carpenter et al., 2019) was also not published until after the testing paradigm used throughout this thesis had already been designed and begun to be executed.

This indicates that, although the results presented in this thesis do not offer support for the hypothesis outlined in chapter 1, they should not necessarily be interpreted as evidence against it.

Overall the results of the studies presented in this thesis do suggest some conclusive findings.

The use of a flexible copying strategy, that requires switching between copying and not copying, is more cognitively demanding than either non-flexible strategy (always copy or never copy). This implies that it is the flexible nature of the strategy that makes it challenging. The finding that flexible copying is more demanding for participants than applying blanket copying strategies is not evidence on its own that flexible social learning strategies (SLSs) are required for cumulative cultural evolution. However, iterating the grid-search task over many simulated generations showed that optimal flexible copying behaviour produced a ratchet effect in task scores.

The results from chapters 2 and 3 also indicated that working memory demands have an impact on the use of flexible SLSs by influencing the copying strategies used by participants. This may result in the use of metacognitive control strategies and system-1 learning heuristics for blanket copying, employed to mitigate the impact of the working memory load. The results indicated that some ratcheting could occur even with restricted access to working memory.

Metacognitive monitoring skills were not implicated in the use of selective SLSs required for CCE. However, there may have been substantial interference from the dual-task manipulation on metacognitive *control* that was not detected by the tasks used in this thesis.

Directions for Future Research

Having laid the groundwork for study in this field, this thesis suggests fruitful avenues for further research into this question. These build on the methods used throughout this thesis, but aim to amend them in order to mitigate the limitations summarised above.

The results summarised above suggested that a short time limit on trials may have prompted the use of system-1 processing even under control conditions. Future testing could make use of the grid-search paradigm used in chapter 3, but modify the task so there was either no time limit imposed on each trial, or a sufficiently long time limit that participants never felt rushed. This would ensure no confounding interference from time constraints on the main-task manipulations.

As noted above, the metacognition task used throughout this thesis only tested participants' metacognitive *monitoring* accuracy. Future research in this area should therefore at a minimum also investigate the impact of task manipulations on metacognitive *control.* Further study that makes use of a dual-task paradigm should aim to design dual-tasks that reliably restrict access to metacognitive monitoring and control processes. In this way, it may be distinguished which, if either, aspect of metacognition is beneficial for participants' potential for ratcheting.

An alternative, or additional, testing paradigm would be to manipulate metacognitive skill through training, rather than restricting access via a dual-task. The expectation would not be that no ratchet effect was observed in participants without additional metacognitive training, as those participants would still have the typical cognitive skills of modern humans. Rather, participants with their metacognitive skills enhanced through training may be able to achieve incremental score increases more quickly than those who had not received training. If this finding was translated to an open-ended 'real world' scenario that was not bounded by a maximum score, it would imply that populations with greater metacognitive skills were able to reach successive stages of trait evolution earlier, and therefore ratchet further within the same time frame, than populations with lower metacognitive efficiency.

Concluding Remarks

In this thesis I aimed to provide empirical support for the hypothesis that explicit metacognition and system-2 processing were necessary factors for cumulative cultural evolution in humans.

Overall, the results presented in this thesis indicate that flexible social learning strategies are more challenging than fixed strategies. Additionally, if these strategies are used optimally ratcheting over generations can occur, but task constraints that prevent use of system-2 processes may prevent optimal behavioural strategies. Although metacognitive monitoring was not significantly affected by the task manipulations, there is some indication that metacognitive control strategies play a role in selective copying. The novel methods used throughout this thesis have also laid the groundwork for future empirical testing of this hypothesis. I have therefore suggested additional testing methods which could provide more evidence for or against the EMCC.

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*, 1–28. https://doi.org/10.1371/journal.pbio.1002564

Allen, J., Weinrich, M., Hoppitt, W., & Rendell, L. E. (2013). Network-Based Diffusion Analysis Reveals Cultural Transmission of Lobtail Feeding in Humpback Whales. *Science*, *340*(6131), 485–488. https://doi.org/10.1126/science.1231976

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/10.1038/nature13998

Apperly, I. A., & Butterfill, S. A. (2009). Do humans have two systems to track beliefs and belief-like states? *Psychological Review*, *116*(4), 953–970. https://doi.org/10.1037/a0016923

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, *59*(4), 390–412. https://doi.org/10.1016/j.jml.2007.12.005

Bahrami, B., Olsen, K., Bang, D., Roepstorff, A., Rees, G., & Frith, C. (2012). Together, slowly but surely: the role of social interaction and feedback on the build-up of benefit in collective decision-making. *Journal of Experimental Psychology. Human Perception and Performance*, *38*(1), 3–8. https://doi.org/10.1037/a0025708

Bahrami, B., Olsen, K., Latham, P. E., Roepstorff, A., Rees, G., & Frith, C. (2010). Optimally Interacting Minds. *Science*, *329*(February), 1081–1086.

Bang, D., Aitchison, L., Moran, R., Herce Castanon, S., Rafiee, B., Mahmoodi, A., Lau, J. Y. F., Latham, P. E., Bahrami, B., & Summerfield, C. (2017). Confidence matching in group decision-making. *Nature Human Behaviour*, *1*(6), 0117. https://doi.org/10.1038/s41562-017-0117

Barrett, A. B., Dienes, Z., & Seth, A. K. (2013). Measures of metacognition on signal-detection theoretic models. *Psychological Methods*, *18*(4), 535–552. https://doi.org/10.1037/a0033268

Beck, W. (1992). Aboriginal preparation of Cycas seeds in Australia. *Economic Botany*, *46*(2), 133–147. https://doi.org/10.1007/BF02930628

Bender, A., & Beller, S. (2014). Mangarevan invention of binary steps for easier calculation.

Proceedings of the National Academy of Sciences, *111*(4), 1322–1327. https://doi.org/10.1073/pnas.1309160110

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/10.1016/j.cognition.2015.05.023

Bernard, S., Proust, J., & Clément, F. (2015). Procedural metacognition and false belief understanding in 3- to 5-year-old children. *PLoS ONE*, *10*(10). https://doi.org/10.1371/journal.pone.0141321

Bull, R., Phillips, L. H., & Conway, C. A. (2008). The role of control functions in mentalizing: Dual-task studies of Theory of Mind and executive function. *Cognition*, *107*(2), 663–672. https://doi.org/10.1016/j.cognition.2007.07.015

Caldwell, C. A. (2020). Using Experimental Research Designs to Explore the Scope of Cumulative Culture in Humans and Other Animals. *Topics in Cognitive Science*, *12*(2), 673–689. https://doi.org/10.1111/tops.12391

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/10.1002/wcs.1516

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/10.1177/0963721416641049

Caldwell, C. A., Cornish, H., & Kandler, A. (2016). Identifying innovation in laboratory studies of cultural evolution: rates of retention and measures of adaptation. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *371*(1690), 20150193-. https://doi.org/10.1098/rstb.2015.0193

Caldwell, C. A., & Millen, A. E. (2008). Experimental models for testing hypotheses about cumulative cultural evolution. *Evolution and Human Behavior*, *29*(3), 165–171. https://doi.org/10.1016/j.evolhumbehav.2007.12.001

Caldwell, C. A., & Millen, A. E. (2009). Social Learning Mechanisms and Cumulative Cultural Evolution. *Psychological Science*, *20*(12), 1478–1483. https://doi.org/10.1111/j.1467-9280.2009.02469.x

Caldwell, C. A., Renner, E., & Atkinson, M. (2017). Human Teaching and Cumulative Cultural Evolution. *Review of Philosophy and Psychology*, *9*(4), 751–770. https://doi.org/10.1007/s13164-017-0346-

3

Call, J., & Carpenter, M. (2001). Do apes and children know what they have seen? *Animal Cognition*, *3*(4), 207–220. https://doi.org/10.1007/s100710100078

Call, J., Carpenter, M., & Tomasello, M. (2005). Copying results and copying actions in the process of social learning: Chimpanzees (Pan troglodytes) and human children (Homo sapiens). *Animal Cognition*, *8*(3), 151–163. https://doi.org/10.1007/s10071-004-0237-8

Call, J., & Tomasello, M. (1999). A Nonverbal False Belief Task : The Performance of Children and Great Apes. *Child Development*, *70*(2), 381–395. https://doi.org/10.1111/1467-8624.00028

Carpenter, J., Sherman, M. T., Kievit, R. A., Seth, A. K., Lau, H. C., & Fleming, S. M. (2019). Domain-general enhancements of metacognitive ability through adaptive training. *Journal of Experimental Psychology: General*, *148*(1), 51–64. https://doi.org/10.1037/xge0000505

Carruthers, P. (2008). Meta-cognition in animals: A skeptical look. *Mind and Language*, *23*(1), 58–89. https://doi.org/10.1111/j.1468-0017.2007.00329.x

Carruthers, P. (2009). How we know our own minds: the relationship between mindreading and metacognition. *The Behavioral and Brain Sciences*, *32*(2), 121–138; discussion 138-182. https://doi.org/10.1017/S0140525X09000545

Carruthers, P., & Ritchie, J. B. (2012). The emergence of metacognition: Affect and uncertainty in animals. In M. J. Beran, J. Brandl, J. Perner, & J. Proust (Eds.), *Foundations of Metacognition*. Oxford University Press.

Claidière, N., Smith, K., Kirby, S., & Fagot, J. (2014). Cultural evolution of systematically structured behaviour in a non-human primate. *Proceedings of the Royal Society B: Biological Sciences*, *281*(1797), 1–9. https://doi.org/10.1098/rspb.2014.1541

Cook, J., & Bird, G. (2011). Social attitudes differentially modulate imitation in adolescents and adults. *Experimental Brain Research*, *211*(3–4), 601–612. https://doi.org/10.1007/s00221-011-2584-4

Coutinho, M. V. C., Redford, J. S., Church, B. A., Zakrzewski, A. C., Couchman, J. J., & Smith, J. D. (2015). The interplay between uncertainty monitoring and working memory: Can metacognition become automatic? *Memory & Cognition*, *43*(7), 990–1006. https://doi.org/10.3758/s13421-015-0527-1

Cultice, J. C., Somerville, S. C., & Wellman, H. M. (1983). Preschoolers' Memory Monitoring: Feeling-of-Knowing Judgments. *Child Development*, *54*(6), 1480–1486.

Davis, S. J., Vale, G. L., Schapiro, S. J., Lambeth, S. P., & Whiten, A. (2016). Foundations of cumulative culture in apes: Improved foraging efficiency through relinquishing and combining witnessed behaviours in chimpanzees (Pan troglodytes). *Scientific Reports*, *6*(April), 1–12. https://doi.org/10.1038/srep35953

De Neys, W. (2006). Automatic-heuristic and executive-analytic processing during reasoning: Chronometric and dual-task considerations. *Quarterly Journal of Experimental Psychology (2006)*, *59*(6), 1070–1100. https://doi.org/10.1080/02724980543000123

Dean, L. G., Vale, G. L., Laland, K. N., Flynn, E., & Kendal, R. L. (2014). Human cumulative culture: A comparative perspective. *Biological Reviews*, *89*(2), 284–301. https://doi.org/10.1111/brv.12053

Derex, M., & Boyd, R. (2016). Partial connectivity increases cultural accumulation within groups. *Proceedings of the National Academy of Sciences of the United States of America*, *113*(11), 2982–2987. https://doi.org/10.1073/pnas.1518798113

Deroy, O., Spence, C., & Noppeney, U. (2016). Metacognition in multisensory perception. *Trends in Cognitive Sciences*, *20*(10), 736–747. https://doi.org/10.1016/j.tics.2016.08.006

Dienes, Z., & Perner, J. (1998). *A theory of implicit and explicit knowledge*. *1999*, 735–808.

Dunlosky, J., & Metcalfe, J. (2009). *Metacognition*. SAGE Publications.

Dunstone, J., Atkinson, M., Grainger, C., Renner, E., & Caldwell, C. A. (2020, August 25). Flexible social learning strategies are harder than the sum of their parts. https://doi.org/10.31234/osf.io/xcyu9

Dunstone, J., & Caldwell, C. A. (2018). Cumulative culture and explicit metacognition: a review of theories, evidence and key predictions. *Palgrave Communications*, *4*(1), 1–11. https://doi.org/10.1057/s41599-018-0200-y

Edmiston, P., Derex, M., & Lupyan, G. (2018). Cumulative improvements in iterated problem solving. *CogSci 2018*, 318–323.

Ehn, M., & Laland, K. N. (2012). Adaptive strategies for cumulative cultural learning. *Journal of Theoretical Biology*, *301*, 103–111. https://doi.org/10.1016/j.jtbi.2012.02.004

Enquist, M., Ghirlanda, S., Jarrick, A., & Wachtmeister, C.-A. (2008). Why does human culture increase exponentially? *Theoretical Population Biology*, *74*(1), 46–55. https://doi.org/10.1016/j.tpb.2008.04.007

Enquist, Magnus, 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/10.1525/AA.2007.109.4.727.728

Epstein, S. (1994). Integration of the cognitive and the psychdynamic unconscious. *American Psychologist*, *49*(8), 709–724. https://doi.org/10.1037//0003-066X.49.8.709

Evans, J. S. B. T. (2007). On the resolution of conflict in dual process theories of reasoning. *Thinking and Reasoning*, *13*(4), 321–339. https://doi.org/10.1080/13546780601008825

Evans, T. A. (2007). *Hypothetical thinking: Dual processes in reasoning and judgement*. Psychology Press.

Evans, T. A., & Stanovich, K. E. (2013). Dual-process theories of higher cognition: Advancing the debate. *Perspectives on Psychological Science*, *8*(3), 223–241. https://doi.org/10.1177/1745691612460685

Flavell, J. H. (1979). Metacognition and cognition monitoring: a new area of cognitive developmental inquiry. *American Psychologist*, *34*(10), 34: 906-911.

Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, *8*, 443. https://doi.org/10.3389/fnhum.2014.00443

Foote, A. L., & Crystal, J. D. (2007). Metacognition in the Rat. *Current Biology*, *17*(6), 551–555. https://doi.org/10.1016/j.cub.2007.01.061

Franks, N. R., & Richardson, T. (2006). Teaching in tandem-running ants. *Nature*, *439*(7073), 153–153. https://doi.org/10.1038/439153a

Galef, B. G. (2013). Imitation and local enhancement: Detrimental effects of consensus definitions on analyses of social learning in animals. *Behavioural Processes*, *100*, 123–130. https://doi.org/10.1016/j.beproc.2013.07.026

Goel, V., & Dolan, R. J. (2003). Explaining modulation of reasoning by belief. *Cognition*, *87*, B11–B22. https://doi.org/10.1016/S0010-0277(02)00185-3

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/10.1080/02109395.2016.1268389

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/10.1073/pnas.1515129113

Hampton, R. R. (2009). Multiple demonstrations of metacognition in nonhumans: Converging evidence or multiple mechanisms? *Comparative Cognition & Behavior Reviews*, *4*, 17–28. https://doi.org/10.3819/ccbr.2009.40002

Henrich, J., Boyd, R., & Richerson, P. J. (2008). Five misunderstandings about cultural evolution. *Human Nature*, *19*(2), 119–137. https://doi.org/10.1007/s12110-008-9037-1

Hertz, U., Romand-Monnier, M., Kyriakopoulou, K., & Bahrami, B. (2016). Social influence protects collective decision making from equality bias. *Journal of Experimental Psychology: Human Perception and Performance*, *42*(2), 164–172. https://doi.org/10.1037/xhp0000145

Heyes, C. (2011). Automatic imitation. *Psychological Bulletin*, *137*(3), 463–483. https://doi.org/10.1037/a0022288

Heyes, C. (2016). Who Knows? Metacognitive Social Learning Strategies. *Trends in Cognitive Sciences*, *20*(3), 204–213. https://doi.org/10.1016/j.tics.2015.12.007

Heyes, C. (2018a). *Cognitive Gadgets*. Harvard University Press.

Heyes, C. (2018b). Enquire within: cultural evolution and cognitive science. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *373*(1743). https://doi.org/10.1098/rstb.2017.0051

Heyes, C., Bang, D., Shea, N., Frith, C. D., & Fleming, S. M. (2020). Knowing Ourselves Together: The Cultural Origins of Metacognition. *Trends in Cognitive Sciences*, *24*(5), 349–362. https://doi.org/10.1016/j.tics.2020.02.007

Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., Couzin, I. D., Bateson, M., Cools, R., Dukas, R., Giraldeau, L. A., Macy, M. W., Page, S. E., Shiffrin, R. M., Stephens, D. W., & Wolfe, J. W. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, *19*(1), 46–54. https://doi.org/10.1016/j.tics.2014.10.004

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), e1001960. https://doi.org/10.1371/journal.pbio.1001960

Horner, V., & Whiten, A. (2005). Causal knowledge and imitation/emulation switching in chimpanzees (Pan troglodytes) and children (Homo sapiens). *Animal Cognition*, *8*(3), 164–181. https://doi.org/10.1007/s10071-004-0239-6

Horner, V., Whiten, A., Flynn, E., & de Waal, F. B. M. (2006). Faithful replication of foraging techniques along cultural transmission chains by chimpanzees and children. *Proceedings of the National Academy of Sciences*, *103*(37), 13878–13883. https://doi.org/10.1073/pnas.0606015103

Jones, P. L., Ryan, M. J., Flores, V., & Page, R. A. (2013). When to approach novel prey cues? Social learning strategies in frog-eating bats. *Proceedings of the Royal Society B: Biological Sciences*, *280*(1772). https://doi.org/10.1098/rspb.2013.2330

Kean, D., Renner, E., Atkinson, M., Kuroshima, H., Wright, C. I., Derrick, A., & Caldwell, C. A. (Manuscript in Preparation). *Potential for Cumulative Culture in Capuchin Monkeys (Sapajus apella) in a Simulated Transmission Chain Study*.

Kempe, V., Gauvrit, N., & Forsyth, D. (2015). Structure emerges faster during cultural transmission in children than in adults. *Cognition*, *136*, 247–254. https://doi.org/10.1016/j.cognition.2014.11.038

Kendal, R. L., Boogert, N. J., Rendell, L. E., Laland, K. N., Webster, M., & Jones, P. L. (2018). Social Learning Strategies: Bridge-Building between Fields. *Trends in Cognitive Sciences*, *0*(0), 1–15. https://doi.org/10.1016/j.tics.2018.04.003

Keren, G., & Schul, Y. (2009). Two is not always better than one. *Perspectives on Psychological Science*, *4*(6), 533–550. https://doi.org/10.1111/j.1745-6924.2009.01164.x

Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: an experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences of the United States of America*, *105*(31), 10681–10686. https://doi.org/10.1073/pnas.0707835105

Konishi, M., Compain, C., Berberian, B., Sackur, J., & de Gardelle, V. (2020). Resilience of perceptual metacognition in a dual-task paradigm. *Psychonomic Bulletin and Review*. https://doi.org/10.3758/s13423-020-01779-8

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/10.1126/science.aaf8110

Kuhn, D. (2000). Metacognitive Development. *Current Directions in Psychological Science*, *9*(5), 178–181. https://doi.org/10.1111/1467-8721.00088

Laland, K. N. (2004). Social learning strategies. *Animal Learning & Behavior*, *32*(1), 4–14. https://doi.org/10.3758/BF03196002

Lehmann, L., Feldman, M. W., & Kaeuffer, R. (2010). Cumulative cultural dynamics and the coevolution of cultural innovation and transmission: An ESS model for panmictic and structured populations. *Journal of Evolutionary Biology*, *23*(11), 2356–2369. https://doi.org/10.1111/j.1420-9101.2010.02096.x

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/10.1098/rstb.2012.0119

Leys, C., Ley, C., Klein, O., Bernard, P., & Licata, L. (2013). Detecting outliers: Do not use standard deviation around the mean, use absolute deviation around the median. *Journal of Experimental Social Psychology*, *49*(4), 764–766. https://doi.org/10.1016/j.jesp.2013.03.013

Mackintosh, G., Atkinson, M., & Caldwell, C. A. (Manuscript in Preparation). *A demonstration of methodology for examining cumulative culture evolution without social transmission*.

Mackintosh, G., Atkinson, M., Renner, E., & Caldwell, C. A. (Manuscript in Preparation). *Intentional information sharing promotes cumulative culture relative to inadvertent behavioural cues: an experimental demonstration*.

Mahmoodi, A., Bang, D., Olsen, K., Zhao, Y. A., Shi, Z., Broberg, K., Safavi, S., Han, S., Nili Ahmadabadi, M., Frith, C. D., Roepstorff, A., Rees, G., & Bahrami, B. (2015). Equality bias impairs collective decision-making across cultures. *Proceedings of the National Academy of Sciences of the United States of America*, *112*(12), 3835–3840. https://doi.org/10.1073/pnas.1421692112

Maniscalco, B., & Lau, H. C. (2012). A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness and Cognition*, *21*(1), 422–430. https://doi.org/10.1016/j.concog.2011.09.021

Maniscalco, B., & Lau, H. C. (2015). Manipulation of working memory contents selectively impairs metacognitive sensitivity in a concurrent visual discrimination task. *Neuroscience of Consciousness*, *2015*(1), 1–13. https://doi.org/10.1093/nc/niv002

Mcclure, S. M., Laibson, D. I., Loewenstein, G., & Cohen, J. D. (2004). Separate Neural Systems Value Immediate and Delayed Monetary Rewards. *Science*, *306*(August 2016), 503–507. https://doi.org/10.1126/science.1100907

Meltzoff, A. N. (1988). The human infant as Homo imitans. In T. Zenall & G. Galef (Eds.), *Social learning: psychological and biological perspectives* (pp. 318–340). Psychology Press.

Mesoudi, A. (2011). Variable Cultural Acquisition Costs Constrain Cumulative Cultural Evolution. *PLoS*

ONE, *6*(3). https://doi.org/10.1371/journal.-pone.0018239

Mesoudi, A., & Thornton, A. (2018). What is cumulative cultural evolution? *Proceedings of the Royal Society B*. https://doi.org/10.1098/rspb.2018.0712

Mesoudi, A., & Whiten, A. (2008). The multiple roles of cultural transmission experiments in understanding human cultural evolution. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *363*(1509), 3489–3501. https://doi.org/10.1098/rstb.2008.0129

Metcalfe, J., & Dunlosky, J. (2008). Metamemory. In J. Byrne (Ed.), *Learning and Memory: A Comprehensive Reference Volume 2* (First Edit, pp. 349–362). Elsevier.

Miller, Tyler M., & Geraci, L. (2011). Unskilled but aware: Reinterpreting overconfidence in low-performing students. *Journal of Experimental Psychology: Learning Memory and Cognition*, *37*(2), 502–506. https://doi.org/10.1037/a0021802

Misailidi, P. (2010). Children’s Metacognition and Theory of Mind: Bridging the Gap. In A. Efklides & P. Misailidi (Eds.), *Trends and Prospects in Metacognition Research* (pp. 279–292). Springer US. https://doi.org/10.1007/978-1-4419-6546-2

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*(23).

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/10.1006/cogp.1999.0734

Nelson, T. O., & Narens, L. (1990). Metamemory: A Theoretical Framework and New Findings. *The Psychology of Learning and Motivation*, *26*, 125–173. https://doi.org/10.1016/S0079-7421(08)60053-5

Nichols, S., & Stich, S. P. (2003). *Mindreading: An integrated account of pretence, self-awareness, and understanding other minds.* Oxford University Press.

Nicholson, T., Williams, D. M., Grainger, C., Lind, S. E., & Carruthers, P. (2019). Relationships between implicit and explicit uncertainty monitoring and mindreading: Evidence from autism spectrum disorder. *Consciousness and Cognition*, *70*(January), 11–24. https://doi.org/10.1016/j.concog.2019.01.013

Olsen, K., Roepstorff, A., & Bang, D. (2019). Knowing whom to learn from : individual differences in metacognition and weighting of social information. *PsyArXiv*, *March 22*. https://doi.org/10.31234/osf.io/jqheu

Osman, M. (2004). An evaluation of dual-process theories of reasoning. *Psychonomic Bulletin & Review*, *11*(6), 988–1010. https://doi.org/10.3758/BF03196730

Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin*, *116*(2), 220–244. https://doi.org/10.1037/0033-2909.116.2.220

Peirce, J., Gray, J. R., Simpson, S., Macaskill, M., Höchenberger, R., Sogo, H., Kastman, E., Lindeløv, J. K., & Peirce, J. (2019). PsychoPy2 : Experiments in behavior made easy. *Behavior Research Methods*, *3*, 195–203.

Perner, J., & Roessler, J. (2012). From infants' to children's appreciation of belief. *Trends in Cognitive Sciences*, *16*(10), 519–525. https://doi.org/10.1016/j.tics.2012.08.004

Perner, J., & Wimmer, H. (1985). John thinks that Mary thinks that: attribution of second order beliefs by 5-year-old to 10-year-old children. *Journal of Experimental Child Psychology*, *39*(3), 437–471. https://doi.org/10.1016/0022-0965(85)90051-7

Perry, C. J., & Barron, A. B. (2013). Honey bees selectively avoid difficult choices. *Proceedings of the National Academy of Sciences of the United States of America*, *110*(47), 19155–19159. https://doi.org/10.1073/pnas.1314571110

Pike, T. W., & Laland, K. N. (2010). Conformist learning in nine-spined sticklebacks' foraging decisions. *Biology Letters*, *6*(4), 466–468. https://doi.org/10.1098/rsbl.2009.1014

Plonsky, L. (2015). *Advancing Quantitative Methods in Second Language Research*. Routledge.

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

Qureshi, A. W., Apperly, I. A., & Samson, D. (2010). Executive function is necessary for perspective selection, not Level-1 visual perspective calculation: Evidence from a dual-task study of adults. *Cognition*, *117*(2), 230–236. https://doi.org/10.1016/j.cognition.2010.08.003

Rahnev, D., & Fleming, S. M. (2019). How experimental procedures influence estimates of metacognitive ability. *Neuroscience of Consciousness*, *2019*(1), 1–42. https://doi.org/10.1093/nc/niz009

Ratcliff, R. (1993). Methods for Dealing with Reaction-Time Outliers. *Psychological Bulletin*, *114*(3), 510–532. https://doi.org/10.1037/0033-2909.114.3.510

Reindl, E., & Tennie, C. (2018). *Young children fail to generate an additive ratchet effect in an open-ended construction task*. 1–22. https://doi.org/10.1371/journal.pone.0197828

Rendell, L. E., Boyd, R., Cownden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., & Laland, K. N. (2010). Why copy others ? Insight from the social learning strategies tournament. *Science*, *328*(5975), 208–213. https://doi.org/10.1126/science.1184719.Why

Rendell, L. E., Fogarty, L., & Laland, K. N. (2010). Rogers' paradox recast and resolved: Population structure and the evolution of social learning strategies. *Evolution*, *64*(2), 534–548. https://doi.org/10.1111/j.1558-5646.2009.00817.x

Roberts, M. J., & Newton, E. J. (2001). Inspection times, the change task, and the rapid-response selection task. *The Quarterly Journal of Experimental Psychology. A, Human Experimental Psychology*, *54*(4), 1031–1048. https://doi.org/10.1080/713756016

Roebers, C. M. (2017). Executive function and metacognition: Towards a unifying framework of cognitive self-regulation. *Developmental Review*, *45*, 31–51. https://doi.org/10.1016/j.dr.2017.04.001

Roebers, C. M., & Feurer, E. (2016). Linking Executive Functions and Procedural Metacognition. *Child Development Perspectives*, *10*(1), 39–44. https://doi.org/10.1111/cdep.12159

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/10.1111/j.1467-8624.2012.01830.x

Roman, A. S., Pisoni, D. B., & Kronenberger, W. G. (2014). Assessment of Working Memory Capacity in Preschool Children Using the Missing Scan Task. *Infant and Child Development*, *23*(6), 575–587. https://doi.org/10.1038/jid.2014.371

Ronfard, S., & Corriveau, K. H. (2016). Teaching and preschoolers' ability to infer knowledge from mistakes. *Journal of Experimental Child Psychology*, *150*, 87–98. https://doi.org/10.1016/j.jecp.2016.05.006

Saldana, C., Fagot, J., Kirby, S., Smith, K., & Claidière, N. (2019). High-fidelity copying is not necessarily the key to cumulative cultural evolution: A study in monkeys and children. *Proceedings of the Royal Society B: Biological Sciences*, *286*(1904). https://doi.org/10.1098/rspb.2019.0729

Sannomiya, M., & Ohtani, K. (2015). Does a dual-task selectively inhibit the metacognitive activities in text revision? *Thinking Skills and Creativity*, *17*, 25–32. https://doi.org/10.1016/j.tsc.2015.04.002

Sasaki, T., & Biro, D. (2017). Cumulative culture can emerge from collective intelligence in animal groups. *Nature Communications*, *8*, 15049. https://doi.org/10.1038/ncomms15049

Schleihauf, H., & Hoehl, S. (2020). A dual-process perspective on over-imitation. *Developmental Review*, *55*(January), 100896. https://doi.org/10.1016/j.dr.2020.100896

Schneider, D., Lam, R., Bayliss, A. P., & Dux, P. E. (2012). Cognitive Load Disrupts Implicit Theory-of-Mind Processing. *Psychological Science*, *23*(8), 842–847. https://doi.org/10.1177/0956797612439070

Schulz, E., Wu, C. M., Ruggeri, A., & Meder, B. (2019). Searching for Rewards Like a Child Means Less Generalization and More Directed Exploration. *Psychological Science*, *30*(11), 1561–1572. https://doi.org/10.1177/0956797619863663

Shea, N., Boldt, A., Bang, D., Yeung, N., Heyes, C., & Frith, C. (2014). Supra-personal cognitive control and metacognition. In *Trends in Cognitive Sciences* (Vol. 18, Issue 4, pp. 186–193). https://doi.org/10.1016/j.tics.2014.01.006

Smith, E. R., & DeCoster, J. (2000). Dual-Process Models in Social and Cognitive Psychology: Conceptual Integration and Links to Underlying Memory Systems. *Personality and Social Psychology Review*, *4*(2), 108–131. https://doi.org/10.1207/S15327957PSPR0402_01

Smith, J D, Redford, J. S., Beran, M. J., & Washburn, D. A. (2009). Rhesus monkeys (Macaca mulatta) adaptively monitor uncertainty while multi-tasking. *Animal Cognition*, *13*(1), 93–101. https://doi.org/10.1007/s10071-009-0249-5

Smith, J D, Schull, J., Strote, J., McGee, K., Egnor, R., & Erb, L. (1995). The uncertain response in the bottlenosed dolphin (Tursiops truncatus). *Journal of Experimental Psychology. General*, *124*(4), 391–408. https://doi.org/10.1037/0096-3445.124.4.391

Smith, J David, Beran, M. J., Couchman, J. J., & Coutinho, M. V. C. (2008). The Comparative Study of Metacognition: Sharper Paradigms, Safer Inferences. *Psychonomic Bulletin & Review*, *15*(4), 679–691. https://doi.org/10.3758/PBR.15.4.679

Smolla, M., Alem, S., Chittka, L., & Shultz, S. (2016). Copy-when-uncertain: bumblebees rely on social information when rewards are highly variable. *Biology Letters*, *12*(6), 208–213. https://doi.org/10.1098/rsbl.2016.0188

Sole, L. M., Shettleworth, S. J., & Bennett, P. J. (2003). Uncertainty in pigeons. *Psychonomic Bulletin & Review*, *10*(3), 738–745. https://doi.org/10.3758/BF03196540

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/10.1111/j.1467-9280.2007.01944.x

Spiess, M. A., Meier, B., & Roebers, C. M. (2015). ProspectiveMemory, executive functions, and metacognition are already differentiated in young elementary school children: Evidence from latent factor modeling. *Swiss Journal of Psychology*, *74*(4), 229–241. https://doi.org/10.1024/1421-0185/a000165

Stanovich, K. E. (1999). *Who is rational? Studies of individual differences in reasoning*. Psychology Press.

Stanovich, K. E. (2004). *The robot's rebellion: Finding meaning the age of Darwin*. University of Chicago Press.

Tamariz, M., & Kirby, S. (2015). Culture: Copying, compression, and conventionality. *Cognitive Science*, *39*(1), 171–183. https://doi.org/10.1111/cogs.12144

Templer, V. L., Lee, K. A., & Preston, A. J. (2017). Rats know when they remember: transfer of metacognitive responding across odor-based delayed match-to-sample tests. *Animal Cognition*, 1–16. https://doi.org/10.1007/s10071-017-1109-3

Tennie, C., Call, J., & Tomasello, M. (2009). Ratcheting up the ratchet: on the evolution of cumulative culture. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *364*(1528), 2405–2415. https://doi.org/10.1098/rstb.2009.0052

Thornton, A., & McAuliffe, K. (2006). Teaching in wild meerkats. *Science*, *313*(5784), 227–229. https://doi.org/10.1126/science.1128727

Tomasello, M. (1990). Cultural transmission in tool use and communicatory signaling of chimpanzees? In S. Parker & K. Gibson (Eds.), *"Language" and Intelligence in Monkeys and Apes: Comparative Developmental Perspectives* (pp. 274–311). Cambridge University Press.

Tomasello, M. (1999). The Human Adaptation for Culture. *Annual Review of Anthropology*, *28*, 509–529. https://doi.org/10.1017/S0140525X0003123X

Tsujii, T., & Watanabe, S. (2009). Neural correlates of dual-task effect on belief-bias syllogistic reasoning: A near-infrared spectroscopy study. *Brain Research*, *1287*, 118–125. https://doi.org/10.1016/j.brainres.2009.06.080

Voelkl, B., & Huber, L. (2000). True imitation in marmosets. *Animal Behaviour*, *60*(2), 195–202.

https://doi.org/10.1006/anbe.2000.1457

Wason, P. C., & Evans, T. A. (1974). Dual processes in reasoning? *Cognition*, *3*(2), 141–154. https://doi.org/10.1016/0010-0277(74)90017-1

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

Whiten, A., Caldwell, C. A., & Mesoudi, A. (2016). Cultural diffusion in humans and other animals. *Current Opinion in Psychology*, *8*, 15–21. https://doi.org/10.1016/j.copsyc.2015.09.002

Whiten, A., Horner, V., Litchfield, C. A., & Marshall-Pescini, S. (2004). How do apes ape? *Learning & Behavior*, *32*(1), 36–52. https://doi.org/10.3758/BF03196005

Wilks, C. E. H., Rafetseder, E., Renner, E., Atkinson, M. A., & Caldwell, C. A. (In Press). Cognitive Prerequisites for Cumulative Culture are Context-Dependent: Children's Potential for Ratcheting Depends on Cue Longevity. *Journal of Experimental Child Psychology*.

Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, *2*(12), 915–924. https://doi.org/10.1038/s41562-018-0467-4

Zwirner, E., & Thornton, A. (2015). Cognitive requirements of cumulative culture: teaching is useful but not essential. *Scientific Reports*, *5*(November 2015), 16781. https://doi.org/10.1038/srep16781