Flexible social learning strategies are harder than the sum of their parts

Abstract

The use of 'explicitly metacognitive' learning strategies has been proposed as an explanation for uniquely human capacities for cumulative culture. Such strategies are proposed to rely on the use of explicit, system-2 cognitive processes, and to enable advantageous selective copying. To investigate the plausibility of this theory, we investigated participants' ability to make flexible learning decisions under executive function (EF) resource load.

Over two studies adult participants completed a simple win-stay lose-shift (WSLS) paradigm task, intended to model a social learning situation where vicarious information can be used to inform response choice, by copying rewarded responses and avoiding those that are unrewarded. This was completed alongside a concurrent EF distractor task. Participants were split into three copying conditions: those that needed to use a flexible copying, WSLS strategy, those that should always copy (information trial always revealed the target), and those that should always do the opposite (information trial always revealed the non-target).

The first study found significant effects of EF load and copying condition; the condition demanding a flexible strategy was more challenging than those requiring the use of one rule consistently. In addition, consistently copying was less challenging than consistently avoiding stimuli selected in the information trial. However, each information condition was equally affected by competing executive function demands. The second study found some differences between selectively copying and always copying, as well as effects of memory load, suggesting the impact of EF load on selective copying may stem from working memory requirements. No impact of EF load was found on participants' metacognitive ability.

These results suggest that learning decisions are underpinned by the use of executive functions even at a very basic level, and that selective copying strategies are more challenging than a combination of their component parts. We found minimal evidence that selective copying strategies relied on executive functions any more than consistent copying or deviation. However, task experience effects suggested that ceiling effects could be masking differences between conditions which might be apparent in other contexts. The results also suggest that such contexts might include situations in which information from a social source must be retained in memory.

Keywords: *Cumulative Culture, cultural evolution, metacognition, executive function, dual-processing, dual-tasks*

1. 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' ability to make efficient, flexible learning decisions while under additional cognitive load was 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 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 within their lifetime.

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 possible 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, 2018), 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 SLS 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 SLS 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 SLS 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 1.2) 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.

1.1 Testing Capacities for Cumulative Culture

Although multiple definitions of CCE have been given in the literature, many focus on different aspects of cultural evolution that could be assessed differently. 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 copy or 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 cumulative cultural evolution 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)).

1.2 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 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 Roebbers (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).

This study will therefore use executive function dual tasks to attempt to restrict access to metacognitive reasoning while participants complete a WSLs decision making 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.

Pilot Testing to establish a dual task paradigm

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, & Roebbers, 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 method (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, mean age: 20, range: 17-45). Of these, four were excluded due to computer errors meaning they could not finish the task. A further 17 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, 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. 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 in the supplementary information.

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: <i>Incongruence</i>	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: <i>Inhibition</i>	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: <i>Switching</i>	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: <i>Updating</i>	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: <i>Working Memory</i>	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, and the testing trial was a very simple arithmetic question.

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 were presented with an information trial which showed two shape stimuli on screen. After a brief pause the task revealed the value (successful – a fish was shown -or unsuccessful – a shark was shown) of one of the two stimuli. Participants were then required to make one selection from the two stimuli presented again (see figure 1, and the supplementary information). 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 a WSLs strategy was always successful. The task therefore models the most basic requirement for cumulative cultural evolution discussed above (section 1.1).

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.

Figure 1:
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).

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

278

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

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

289 The variable of interest was the reaction time of all responses given in the main task. Response time
290 was used rather than accuracy as accuracy was predicted to be at ceiling levels in the task. Figure 2
291 shows the mean RT in the binary search task for each dual-task type. Accuracy in this task was
292 approximately at ceiling across all task conditions (accuracy range: 96.4%-98.9%), so is not analysed
293 here.

294

295 *Outliers*

296 Overall 186 outliers were removed from the data for very long or very quick reaction times. Each
297 outlier represents a single trial. A relatively broad inclusion criterion (3 x Median Absolute Deviation
298 (MAD) from the mean; Leys, Ley, Klein, Bernard, & Licata, 2013) was used to ensure genuinely long
299 reaction times caused by dual task interference, which were the expected outputs of the study, were
300 not artificially trimmed. Outliers were removed per participant, per task type in line with
301 recommendations from Ratcliff (1993) regarding outlier removal for reaction time datasets with high
302 inter-participant variation. 182 outliers were removed from the upper end of the response
303 distribution and 4 were removed from the bottom end. This represents 4.7% of the total data.

304

305 *Difference from baseline*

306 Each task's difference from the baseline task was calculated using a linear mixed effects model for
307 each task with one fixed effect of block condition (baseline, control or executive function block) and
308 one random effect of participant ID. *P*-values were estimated from the resultant *t*-statistics with
309 degrees of freedom being the number of observations minus the number of fixed parameters in the
310 model (Baayen, Davidson, & Bates, 2008). For the *incongruence* and *withholding* tasks the model was
311 not significantly better than a null model (incongruence: $\chi^2=0.288$, $p=.866$; withholding: $\chi^2=5.18$,
312 $p=.075$) and the tasks showed no significant differences from the baseline condition in either block.
313 Both the *switching* and *working memory* task models were significantly better than the null model
314 (switching: $\chi^2=43.6$, $p<0.001$; working memory: $\chi^2=21.2$, $p<0.001$) and both tasks showed significantly
315 longer reaction times in each block compared to the baseline task. The results of each model are
316 presented in table 2.

317

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 ($\chi^2=0.179$, $p=.672$) 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 (withholding: $\chi^2=4.10$, $p=.043$; switching: $\chi^2=26.6$, $p<.001$; working memory: $\chi^2=16.3$, $p<.001$) and all three tasks showed significantly longer reaction times in the executive function block compared to the control block. 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 2: 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. 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 at 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 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.

2. 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 also 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, with blocks requiring single response-types (always copy or never copy) which are 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 mixed 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 blocks, was designed to be simple for participants to do. This was partly to ensure that any effects found are 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 Supplementary Information II) demonstrated that dual-task interference did have a significant negative effect on responses in mixed 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, 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 a 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 strategic copying group):

WSLS (Selective Copying group): In this group there was a 50% chance of showing a rewarded or unrewarded 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 rewarded and unrewarded 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 rewarded. 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 unrewarded. 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 also completed an EF dual task. 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 Pilot). 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 3 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

STRATEGY 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 ($\chi^2(11)=531, p<0.001$). There were significant effects of strategy 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 conditions, 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 strategic 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.

Figure 3: 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.

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) ($b=-34.8$, $SE=3.9$, $t(15999)=-8.92$, $p<.001$).

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 4 shows mean reaction times by information trial type.

Figure 4: 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).

Table 5: Mean accuracy in each strategy group and block condition

STRATEGY 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%

Success on each trial of the dual task 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 SI figure

S3). 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 strategy group in the control block compared to every strategy 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 between overall accuracy between strategy 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.

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.

3. 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 would be required at once in order to perform optimally at the task. The strategy requirements in the repeated copying 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, 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 supplementary information for a break-down of exclusions and accuracy thresholds required to pass all tasks). In total 94 participants (23 male, 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. Ethical approval for the study was given by the University of Stirling General University Ethics Panel (reference GUEP 468).

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. The search task and metacognition task were each completed in blocks, the order of which was counterbalanced across participants. Due to the low pass rate in Experiment 2, 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 task now had two different switch cues: one to indicate that participants should use each of the two pre-specified mouse responses (see supplementary information for more details). 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 5 for an example of the procedure.

Figure 5: 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. They were then asked to rate their confidence in that decision on a scale from 1-4

(1 being least confident, 4 being most confident), using the 4 buttons of the button box. Points were gained for correct responses and lost for incorrect responses, with the number of points gained or lost equal to the confidence rating given. Immediate feedback about the number of points gained or lost, and therefore the correctness of the response, was given after each trial and a running total of points was given in the top right corner of the screen throughout the experiment. This running score could be negative. A selection needed to be made within five seconds or the trial would time out and one point would be taken from the total. No time limit was imposed on the confidence rating. The metacognition task was completed in conjunction with the switching dual task.

There were 7 difficulty levels, with each level twice as hard as the level before it (set by percentage difference in the density of the patch, see SI for full details). This was to ensure some trials could definitely be answered correctly and some would have to be answered based on a random guess between the two stimuli, to ensure a mixture of correct and incorrect answers and to encourage use of the full scale of confidence ratings. The difficulty levels were set during piloting to aim for an estimated average accuracy of around 75% in the task. The same difficulty levels were used across participants. Participants completed 28 training trials and 56 test trials, with an equal number of trials at each difficulty level, presented in a random order.

Figure 6: An example of a 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

673 requirement to use a selective strategy in the WSLs group. This is in line with EMCC prediction that
674 the use of selective strategies relies on explicit processes. We did not make specific predictions
675 regarding the addition of a memory load to the search task. The EMCC would predict that a task that
676 reduced capacities for selective copying would also have a negative impact on metacognitive
677 efficiency. We therefore predicted the *switching* dual-task would have a negative impact on
678 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.

Analysis

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)

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

Response times 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 than its null equivalent ($\chi^2(17)=1123, p<.001$).

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 S5).

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 S5).

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 S6).

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 S7).

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 S8).

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

Response

Control dual EF dual

Figure 7: Response times at R1 and R2 for the WSLs group (top) and WS group (bottom), split by block condition and memory load.

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 8).

Figure 8: 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 ($\chi^2(7)=2329$, $p<.001$). 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 9). 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 9: Accuracy in metacognition task at each confidence level, split by trial difficulty. Difficulty Level 1, least difficult; Difficulty Level 7, most difficult.

770

771 The impact of the dual task on participants' metacognition was analysed using the metaSDT package
 772 in R (Barrett et al., 2013). A single score of metacognitive efficiency (m-ratio) for each participant in
 773 each block was calculated as metacognitive sensitivity (meta- d'_b : balanced model fit) divided by
 774 overall sensitivity in the visual task (d'). Meta- d' is defined as the type-I sensitivity (accuracy in the
 775 visual discrimination task) that would be found if all of a participant's type-II ratings (subsequent
 776 confidence ratings) were considered to be optimal (Fleming & Lau, 2014). This measure aims to give a
 777 bias free measure of metacognition which is not affected by performance in the visual task or a bias
 778 towards over- or under-confidence (Maniscalco & Lau, 2012). A repeated measures ANOVA showed
 779 no significant difference in m-ratio between control and EF blocks (mean control: 1.13, mean EF: 1.03,
 780 $F(93)=0.663$, $p=.418$) (see figure 10).

Figure 10: M-ratio (metacognitive efficiency) for the control and EF blocks. Red dots indicate the mean for each block.

781

782 Correlation between search task and metacognition task

783 The mean difference in RT between the EF and control blocks of the search task was calculated for
 784 each participant. This RT difference reflected the impact of the dual-task. Overall, there was no
 785 correlation between dual-task impact and baseline metacognitive efficiency (taken as m-ratio for the
 786 control block only) for either R1 or R2 (R1: $r=.011$, $t(92)=0.102$, $p=.919$; R2: $r=.070$, $t(92)=0.675$,
 787 $p=.501$), indicating no relationship between greater metacognitive efficiency and impact of the dual
 788 task on speed of applying a flexible copying strategy. This may indicate that, in contrast with the
 789 predictions of the EMCC, increased metacognitive monitoring accuracy does not necessarily lead to
 790 an increased proficiency in applying a flexible strategy.

791 Audio Dual Task

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

799 Participant accuracy was at or close to ceiling in all conditions (see figure S4); accuracy range 86.5%-
 800 95.7%.

801 Success on each trial of the audio task was analysed using a binomial linear mixed effects model with
 802 fixed effects of block condition, main task and their interactions, and scaled item number. Participant
 803 ID was included as a random effect. This model was significantly better than the null model
 804 ($\chi^2(4)=850$, $p<.001$). Accuracy was significantly lower in the EF block ($b=-1.21$, $SE=0.057$, $z=-21.3$,
 805 $p<.001$) but significantly higher during the search task compared to the metacognition task ($b=0.219$,
 806 $SE=0.069$, $z=3.16$, $p=.002$). Accuracy got lower as trial number increased, ($b=-0.421$, $SE=0.071$, $z=-$
 807 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.

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.

5. 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.

Acknowledgements

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme under grant agreement No. 648841 RATCHETCOG ERC-2014-CoG.

Declarations of interest: none

References

- 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>
- 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>
- 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., 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>
- 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>
- 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>
- 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>
- 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>
- 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>
- 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 Psicología*, 38(1), 37–62. <https://doi.org/10.1080/02109395.2016.1268389>
- 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. (2018). 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>
- 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>
- Laland, K. N. (2004). Social learning strategies. *Animal Learning & Behavior*, 32(1), 4–14.

970 <https://doi.org/10.3758/BF03196002>

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

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

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

979 Pashler, H. (1994). Dual-task interference in simple tasks: Data and theory. *Psychological Bulletin*,
 980 116(2), 220–244. <https://doi.org/10.1037/0033-2909.116.2.220>

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

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

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

990 Ratcliff, R. (1993). Methods for Dealing with Reaction-Time Outliers. *Psychological Bulletin*, 114(3),
 991 510–532. <https://doi.org/10.1037/0033-2909.114.3.510>

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

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

998