CellPress

Opinion

Who Knows? Metacognitive Social Learning Strategies

Cecilia Heyes[1,2,*]

To make good use of learning from others (social learning), we need to learn from the right others; from agents who know better than we do. Research on social learning strategies (SLSs) has identified rules that focus social learning on the right agents, and has shown that the behaviour of many animals conforms to these rules. However, it has not asked what the rules are made of, that is, about the cognitive processes implementing SLSs. Here, I suggest that most SLSs depend on domain-general, sensorimotor processes. However, some SLSs have the characteristics tacitly ascribed to all of them. These metacognitive SLSs represent 'who knows' in a conscious, reportable way, and have the power to promote cultural evolution.

It Pays to Know Who Knows

Learning from others (social learning) can be a wonderful thing, but only if you learn from the right others. Whether you are trying to blend in with a new social group or tackling a new piece of technology, there is no point in copying another ingénue. Indeed, if the learning itself is laborious, or mistakes are expensive (one false move produces ostracism or a factory reset), taking a lead from the wrong people can be even more costly than going it alone. Sticking with old habits, or trying new things until you hit on something that works (asocial learning) can be better than copying the clueless (Box 1) [1,2]. To make social learning work for us, we need to keep track of who knows.

Over the past 10 years, researchers from a range of disciplines have been looking closely at transmission biases [3,4] or SLSs [5–7]. These are rules, such as 'copy when uncertain', 'copy the majority', and 'copy the most successful', that rational agents should use to maximise the efficiency of their social learning. Using mathematical modelling, behavioural ecologists and economists have devised a large number of SLSs, and found that these rules often predict when and who we copy in laboratory experiments [5,6,8]. And the rules do not only apply to people. The behaviour of an impressive array of animals, including rats [9], sticklebacks [10], fruit flies [11], and frog-eating bats [12], also conforms to SLSs.

Fortified by these achievements, ecologists and economists have proposed that SLSs provide a key to understanding two fundamental features of life on Earth: the evolution of social learning and human culture [5,13,14]. The first of these proposals is straightforward: social learning is ubiquitous in the animal kingdom [15]; even ants [16] and caddisfly larvae [17] can learn from others. The ubiquity of social learning would be a puzzle if creatures opted for social rather than asocial learning at random, and were just as likely to copy ignorant as knowledgeable conspecifics. In that case, social learning would amount to nothing more than indiscriminate recycling of behaviour. On average, it would not improve a creature's efficiency or increase its reproductive fitness [18]. Therefore, research on SLSs helps to explain the evolution of social learning by showing that it could be, and often is, launched and targeted in ways that make it adaptive. Research on SLSs shows that it is not the tool itself (social learning) that is valuable, but the way it is used.

Trends

When agents learn selectively from others, they are using SLSs.

Mathematical modelling indicates that these strategies enhance the efficiency of social learning in humans and other animals.

SLSs are described as if they were rules, such as 'copy when uncertain' and 'copy the most successful', that are voluntarily deployed.

SLSs are thought to have contributed to the emergence of human culture.

Here I propose a dual-system psychological framework for research on the mechanisms and functions of SLSs.

[1]All Souls College, University of Oxford, Oxford, OX1 4AL, UK
[2]Department of Experimental Psychology, University of Oxford, Oxford, OX1 4AL, UK

*Correspondence:
cecilia.heyes@all-souls.ox.ac.uk
(C. Heyes).

CrossMark

Box 1. Social and Asocial Learning Depend on the Same Psychological Processes

This article concerns the psychological processes controlling the deployment of social and asocial learning. Previous research, examining the intrinsic properties of these two types of learning, has shown that they exploit different sources of information: cues provided by other agents (social learning), and cues from the inanimate environment (asocial learning). However, four lines of evidence converge on the conclusion that information from these sources is processed by the brain using the same computations [15]: (i) Social and asocial learning ability co-vary. Among birds [68] and among primates [69], species that perform well in tests of social learning also tend to perform well in tests of asocial learning, and this positive correlation is present even when statistical analyses control for body mass, brain volume, and phylogeny. (ii) Social learning in solitary animals. In laboratory tests, animals such as octopuses [70] and red-footed tortoises [71] (Figure I), which lead solitary lives in the wild, prove themselves to be adept at learning from social cues. (iii) Common varieties. Social and asocial learning each come in the same three varieties (learning about single stimuli, about relations among stimuli, and about relations between stimuli and responses, or actions and outcomes) [15], and each type of social and asocial learning has been found in a wide range of species [16,72]. For example, social learning about actions and outcomes, which was once thought to be uniquely human, has now been found in birds [73,74]. (iv) Associative learning of social value. Studies of human decision-making, combining mathematical modelling with functional brain imaging, have found that the same computations, based on the calculation of prediction error, are involved in processing information from social partners (social learning) and personal experiences of reward (asocial learning) [26,75]. Processing of social and asocial cues is sometimes carried out in different brain areas [76]. However, evidence is emerging that, rather than being rigidly specialised for social and asocial learning, each brain area may be capable of processing social and asocial cues, and switches back and forth according to which type of cue is currently more relevant for action [77,78]. Together, these four lines of evidence suggest that, intrinsically, both social and asocial learning depend on domain-general, or 'generic' [75], processes of learning.

Trends in Cognitive Sciences

Figure I. Social Learning in Solitary Animals. The red-footed tortoise *Chelonoidis carbonaria* engages in very little social interaction under free-living conditions, but is capable of social learning in laboratory tests [71]. Reproduced with permission from Peter Baumber and Anna Wilkinson.

Glossary

Associative learning: mechanisms that encode information about relations between events in excitatory and inhibitory links between sensory and motor representations. Traditionally, these mechanisms have been studied using conditioning procedures. The computations that control link formation depend on prediction error [64,65].

Domain-general: in this article, a process is domain-general if it applies the same computations to information from social and asocial sources. Therefore, a process mediating a SLS is domain-general if it selects between social and asocial inputs, among social inputs, and among asocial inputs, in the same way.

Implicit (System 1) versus explicit (System 2) cognition: implicit cognition, or System 1, is minimally dependent on working memory, and typically processes information rapidly and in parallel. We are seldom conscious of its operations, and many of them involve sensorimotor processes. By contrast, explicit cognition, or System 2, is highly dependent on working memory, and typically slow, serial, rule based, and conscious [66,67].

Implicit (System 1) versus explicit (System 2) metacognition: metacognition is a special kind of cognition that represents properties of cognitive processes (e.g., their speed or reliability), rather than properties of the world (e.g., the location or value of an object). Implicit metacognition represents cognitive processes in System 1 (e.g., rapidly, in parallel, with minimal dependence on working memory), whereas explicit metacognition represents cognitive processes in System 2 (e.g., slowly, serially, and in a conscious, reportable form).

Metacognition: the use of metacognitive representations. These represent properties of cognitive processes; for example, the reliability of a perceptual representation [32].

'When' and 'who' SLSs: 'when' SLSs are rules specifying the conditions in which an agent should engage in social learning rather than persist in using an established behaviour or engage in asocial learning. For example, copy when

The second proposal, that SLSs help to explain human culture, is harder to understand. Indeed, it seems to conflict with the first proposal. If a range of animals use SLSs, and SLSs promote culture, why is human culture so much richer than that of other animals? There is some evidence that nonhuman primates and cetaceans show geographical variation in behaviour that is not due to genetic or ecological differences [19,20], but why do rats, sticklebacks, fruit flies, and frog-eating bats not show this kind of cultural diversity? Why do these species and many others, equipped with SLSs, not show cumulative cultural change, that is, the accretion of wisdom over generations, through social learning, to produce sophisticated technology, elaborate social practices, and vast libraries of knowledge about the world? There can be no doubt that many

human qualities, especially language, contribute to making us 'odd' animals [3]. The question here is whether the use of SLSs is one of those qualities and, if so, how it contributes to making humans such peculiar, cultural animals.

Here, I suggest that some SLSs have the power that has been ascribed to all SLSs: the power to promote culture. The culture-promoting SLSs are based on a special kind of **metacognition** (see Glossary); psychological processes that produce explicit, reportable representations of 'who knows'. These explicitly metacognitive SLSs are used only by humans. All other SLSs, in human and nonhuman animals, are supported by simpler, general-purpose psychological mechanisms. I argue that this deep divide, between metacognitive and other SLSs, has not been recognised before because research on SLSs has been dominated by ecology and economics, disciplines concerned with what agents do, not what they think [21]. When recent developments in cognitive science are brought to bear on SLSs, it becomes apparent that humans are the only animals that target their social learning by asking 'who knows?'.

Animals Are Like Planets

Planetary motion conforms to certain rules, for example, 'the orbit of a planet is an ellipse with the Sun at one of two foci'. However, planets do not understand these rules or implement them deliberately. Planetary motion is law-like because it is driven by consistent internal and external forces, but the rules of planetary motion are in the minds of scientists, not in the minds of planets. Similarly, I suggest that the behaviour of nonhuman animals can be described and predicted by SLSs, but the strategies are in the minds of scientific observers, not of the animals themselves.

The language used in research on SLSs implies that they are rules in the minds of agents rather than scientific observers. To say that an agent 'uses' a strategy implies a voluntary act of deployment. To characterise the strategy itself in words, such as 'copy when uncertain', implies that what is deployed is a rule encoded in a reportable form, and 'copy' implies that the rule relates specifically to decisions about social learning. Even if few researchers would endorse these claims explicitly, the language of SLSs is potentially misleading. A recent analysis of the evidence indicates that the psychological forces driving animals to engage in social rather than asocial learning, and to copy one model rather than another, are not voluntarily deployed, domain-specific rules (Box 2) [22]. Instead, these forces comprise processes of **associative learning** that operate automatically, rather than in a controlled way, and are **domain-general**; that is, they select not only between social and asocial sources of information, and among social cues, but also among asocial or inanimate sources of information. They are products of biological evolution (natural selection operating on genetic variants) but they did not, as the language used in economic and ecological analyses implies, evolve specifically for the guidance of social learning. Their evolutionary purpose is to direct all learning towards objects and events that are most likely to carry useful information.

People Are (Sometimes) Like Cooks

Associative processes are powerful determinants of human as well as animal behaviour [23–26]. Therefore, it is likely that these domain-general processes underpin much of the human behaviour that conforms to SLSs [6,27–31]. However, I suggest that humans sometimes use explicitly metacognitive SLSs (Box 3). These strategies, and only these, have the characteristics that previous research has tacitly ascribed to all SLSs, found in animals and humans. They are consciously represented, reportable, domain-specific rules. These rules are metacognitive in that they represent properties of the cognitive processes of the rule-user and of other agents. They are distinct from other explicit metacognitive rules, such as those used to allocate study time during examination preparation, or to decide how much to bet on a perceptual judgement, because they specify conditions in which it is advisable to engage in social rather than asocial learning, and, in the former case, from whom one should learn.

uncertain. 'Who' SLSs are rules specifying the type of agent to be copied. For example, copy older individuals [5].

Box 2. Frog-Eating Bats Copy When Dissatisfied

Heyes and Pearce [22] reviewed the most striking examples of SLSs in animals and found that they could be explained by domain-general processes of attention and associative learning. For example, in one of the experiments they examined, frog-eating bats initially learned to approach one of two auditory stimuli, A or B, for food [12] (Table I). For A-trained bats (the design was counterbalanced), food was always available from the loudspeaker playing A. In the second phase, these 'observer' bats were again able to retrieve food from the loudspeaker playing A, but now food could also be retrieved from the other loudspeaker, playing B. During this phase, food was available on the loudspeaker playing A on 100% of the trials for group 100-social, and on 50% of the trials for group 50-social and for group 50-solitary. The trials for group 100-social and group 50-social involved a second bat, a 'model', which was released at the same time as the observer, and which had been trained to find food on the loudspeaker playing B, but not A. Subsequent testing revealed a strong preference for A over B in groups 50-solitary and 100-social, but a preference for B over A in group 50-social.

The behaviour of the bats in this study conformed to the rule 'copy when dissatisfied'. Observation of a model that went to B encouraged the observers to reverse their original preference for A over B (social learning), when the rate of reward at A had recently declined ('dissatisfied' group 50-social), but not when it had been sustained ('satisfied' group 100-social). However, associative learning is sufficient to explain the behaviour of the bats. For example, on the associative account, when the bats had a choice between A and B on test, group 50-social approached A less than group 100-social and, therefore, by default, approached B more, because, during the second phase of training, approaches to A had been intermittently rewarded for group 50-social and consistently rewarded for group 100-social. Compared with consistent reward, intermittent reward yields a weaker association between the cue and the positive event that follows and, therefore, intermittent reward does not make a cue as attractive as consistent reward [79]. In support of this domain-general explanation, Jones and colleagues [12] found that, in a final test where B was presented without A, groups 50-social and 100-social were equally likely to approach B.

Table I. Key Elements of the Design Used by Jones *et al.* [12], Illustrated by the Groups of Bats Initially Trained to Approach Cue A

Phase of Training	Group	% Reward at A	% Reward at B
1	All	100	0
2	100-social	100	100 + model feeding at B
	50-social	50	100 + model feeding at B
	50-solitary	50	100 no model

Some of the labels currently in use make it obvious that SLSs can have metacognitive content, that is, that they can relate to 'who knows'. For example, 'copy when uncertain' clearly implies that agents should be biased in favour of social learning to the extent that they lack confidence in their own knowledge; to the extent that their own knowledge, or the process yielding that knowledge, is metacognitively represented as lacking precision or reliability. Other labels need unpacking to reveal their potential metacognitive content. For example, 'copy the majority' becomes 'copy the majority when the majority is likely to know best'.

When humans use explicitly metacognitive SLSs, they are like cooks rather than like planets. Cooks know the rules to which their behaviour conforms, and the conformity of their behaviour is due, in part, to their knowledge of the rules. They can tell you that 'a good cook never fries until they see blue smoke arise', and this explicit knowledge is part of what makes a cook wait until the oil is hot before putting food in the pan.

No research to date has been dedicated to finding out what kinds of psychological process mediate SLSs. However, the results of several studies provide evidence that, in adult humans, SLSs are sometimes implemented by explicit metacognitive rules (Box 4).

Recipes Are Socially Learned

Research on metacognition in general indicates that explicitly metacognitive rules are learned [32,33]; this learning typically depends on social interaction [34–36] and, consequently, there is marked cultural variation in explicit metacognition [37–40]. For example, children learn by instruction to use 'semantic clustering' to retrieve the names of animals from memory [34]

Box 3. Explicit Metacognition

Metacognition is broadly defined as cognition-about-cognition or thinking-about-thinking. **System 2** or **explicit metacognition** is a special kind of thinking-about-thinking [32] (Figure I). In common with **System 1** or **implicit metacognition**, it represents properties of cognitive processes (e.g., their speed or reliability) rather than properties of the world (e.g., the location or value of an object). However, unlike implicit metacognition, explicit metacognition represents properties of cognitive processes in a conscious, reportable form. With explicit metacognition, an agent can think and say 'I'm sure' or 'I'm doubtful'.

More broadly, within dual systems models of the mind [66,67], and related theories [80,81], explicitly metacognitive representations are part of a cognitive system that tends to handle problems slowly and serially. Its functioning depends on working memory, and correlates with differences between people in general intelligence. By contrast, implicit metacognition, and many of the cognitive processes represented by explicit metacognition, including the processes of associative learning, are part of a cognitive system that tends to handle problems rapidly and in parallel, and that is minimally dependent on working memory. Explicit metacognition can be regarded as a component of System 2 without embracing controversial aspects of dual systems models [82,83]; for example, without assuming that the features of System 2 cognition are perfectly correlated, or that System 1 operation cannot be characterised using (planetary) rules [32,66].

Explicit metacognition is typically studied by cognitive scientists using judgements of learning and confidence judgements. When working towards an examination, students use explicitly metacognitive judgements of their own prior learning to exclude from future study materials they have already assimilated, and to prioritise material they have nearly, but not quite, mastered [84,85]. When making perceptual decisions, for example, about the presence or orientation of an object in a stimulus array, people use explicitly metacognitive confidence judgements to decide how much they should bet on the accuracy of their decisions, and to communicate the reliability of their decisions to cooperation partners [33,35,86]. The power of explicit metacognition to influence the operation of implicit processes is illustrated by studies showing that explicit metacognitive beliefs about self-control modulate eating behaviour, procrastination during an examination period, and performance on simple laboratory tasks requiring the inhibition of habitual responses [87–89].
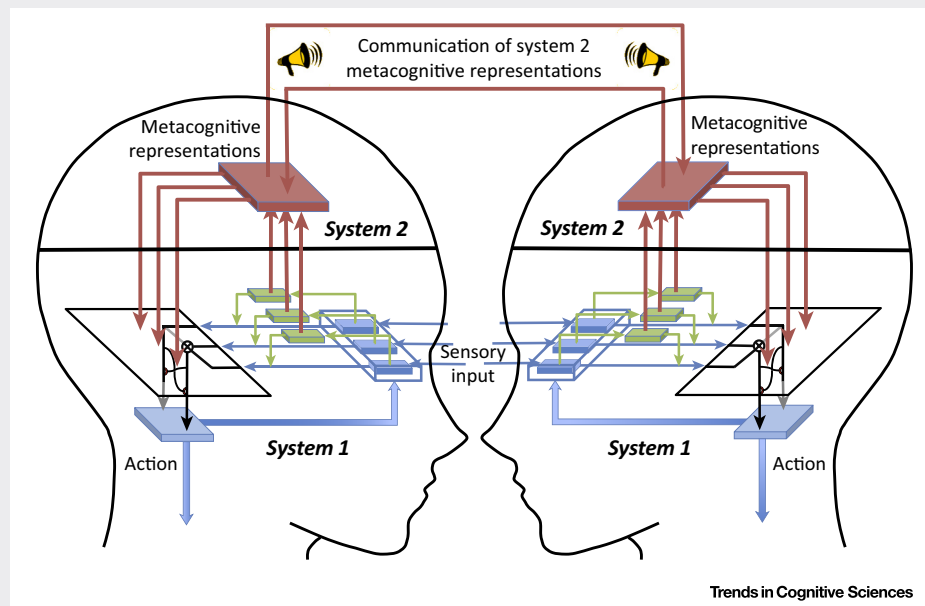

Figure I. System 2 Metacognition. Shea *et al.* [32] argued that explicit or 'System 2' metacognition is for cognitive control across two agents. System 2 metacognitive representations are derived from information in System 1, but they are in a form available for report. For example, the reliability of a sensory signal can be reported in terms of confidence. When agents are cooperating, these reports can be used to optimise control by, for example, giving more weight to the more confident observer [90]. Reproduced, with permission, from [32].

(e.g., think of birds first, then mammals. . .), and adults learn through social interaction explicitly to metarepresent their confidence in ways that make two heads better than one [35,36].

Given these findings from research on metacognition in general, one would expect explicitly metacognitive SLSs also to be products of learning through social interaction, and to vary across

Box 4. Explicitly Metacognitive SLSs

There is evidence that both **'when' and 'who' SLSs** can be explicitly metacognitive. In a study relating to 'when' strategies, people were asked to make a preliminary decision in foraging and perceptual tasks, and an explicit judgement of their confidence in that decision, before being given the opportunity to use social information to make a final decision [91]. The participants' confidence judgements were accurate (they had lower confidence in wrong rather than right preliminary decisions) and, crucially, they were increasingly likely to use social information as their confidence declined, suggesting that they deliberately applied the rule 'copy when uncertain'. In a study relating to 'who' strategies, people given the opportunity to copy members of their own or an opposing team were able to report the strategies they had adopted in a post-test questionnaire [92]. Similarly, in another study, participants preferentially copied popular choices when other people's payoffs were visible, but preferentially copied unpopular choices when their payoffs were invisible [6]. A potential explanation for this unpopularity bias is that people were deliberately avoiding an 'information herd' [21,93]; their social learning was guided by explicit knowledge, acquired from others, that majorities know best only when they have access to information about payoffs.

Further evidence that SLSs can be explicitly metacognitive comes from an exciting new field of cognitive neuroscience with great potential to illuminate SLSs. This field uses computational modelling of behavioural and neuroimaging data to identify the psychological and neural processes mediating social and asocial learning [76,94–96]. As yet, these techniques have not been used deliberately to investigate SLSs, but they have already provided evidence that social and asocial learning depend on the same psychological processes (Box 1, main text), and that, in adult humans, 'who' SLSs can be explicitly metacognitive. For example, in a foraging task, people made use of social information (advice about which of two options to choose) to the extent that they believed the advisor to be motivated to help rather than to mislead them [95]. These beliefs were explicitly stated, and the basic effect (co-variation between the advisors' incentives and the participants' use of social information) disappeared when participants were told that the advisors did not know which option they were recommending. Therefore, these results suggest that the participants used an explicitly metacognitive strategy, such as 'copy when the model intends to help'.

cultures. Evidence consistent with the first of these predictions indicates that preschool children are more like planets than like cooks. When children under 4- or 5-years of age observe the nonverbal actions of others, their behaviour conforms to rules such as 'copy when uncertain' [41], 'copy the in-group' [42], and 'copy prestigious individuals' [43,44], but I have argued in two reviews of the empirical evidence that their choices can be explained by domain-general psychological processes [45,46]. This suggests that domain-specific SLSs do not come online until relatively late in development, when there has been plenty of opportunity for them to be learned through social interaction. Evidence consistent with the second prediction shows that there is marked cross-cultural variation among the SLSs used by adults [47–51]. For example, in contrast with Westerners, Fijians are less likely to seek advice from people with more formal education [50] and, in contrast with Britons, people from mainland China engage in more social learning, and their social learning is less dependent on uncertainty [47].

Concluding Remarks

This article asks, for the first time, what SLSs are made of, about the cognitive mechanisms that direct social learning to the right 'others'. It proposes that most SLSs are implemented by domain-general sensorimotor processes (planetary SLSs), while a small proportion, found only in humans, are explicitly metacognitive (cook-like SLSs). This dual systems account provides a framework for future research on SLSs (see Outstanding Questions), and casts further light on the roles of SLSs in the evolution of social learning and of human culture.

Economists and ecologists write about SLSs as if they were all domain-specific rules produced by genetic evolution; rules about the use of social information (not information in general), produced by the slow process of natural selection operating on genetic variants. If this were true, it is likely that SLSs would be highly inflexible over time. By contrast, I have suggested that all animal behaviour, and much human behaviour, conforming to SLSs, is based on domain-general processes of associative learning. These learning processes are products of genetic evolution, but, because they are minimally canalised or genetically constrained, they allow SLSs to change rapidly over time. For example, if the younger members of a population discovered some clever foraging techniques [52], associative learning could quickly convert 'copy older

individuals' to 'copy younger individuals' in that population [22]. Therefore, the idea that animals are like planets (their SLSs are based on associative learning) casts further light on the ubiquity of social learning in the animal kingdom. Social learning is a valuable tool because it can be used, not only selectively, but also with a selectivity that can be adjusted rapidly over time to track changes in the social and asocial environment.

My proposal that most SLSs are based on associative learning implies that they are smarter (more supple and adaptive) than was previously thought. However, if planetary SLSs are so smart, what are the advantages of cook-like, explicitly metacognitive SLSs; of rules that are domain specific and socially learned? I suggest that cook-like SLSs have the edge because they can isolate 'who knows' with greater accuracy and precision than planetary SLSs. They can be more veridical and specific in identifying sources of superior knowledge because metacognitive SLS can be based on collective wisdom accumulated over generations [53]. Planetary SLSs change as a function of the user's own, recent experience; for example, if a monkey finds that information from females has yielded higher payoffs recently, it will turn its attention from males to females [22,54]. By contrast, because they are socially learned, metacognitive SLSs can distil the experience of many agents over an extended period of time. For example, 'copy the majority only when payoffs are visible' (Box 4) [6] is based on the experience of many people in a range of environments, and anyone who has learned from others about the dangers of herding is likely to stick with this rule even when they encounter exceptions. Furthermore, metacognitive SLSs can comprise hierarchies of rules [55], in which some of the conditions are socially defined activities and roles. For example, when uncertain about a health problem, copy the medicinal plant expert with highest prestige; when uncertain about a boat-building problem, copy the boat builder with the largest fleet. Or, in contemporary Western society, when you do not know much about maths or science, place more trust in science accompanied by mathematical equations [49].

Evidence that nonhuman animals use explicit metacognition in solitary tasks [56,57], although controversial [32,58], raises the possibility that some can also represent 'who knows' in social tasks. However, the advantages of cook-like SLSs depend on them being constructed from the experience of many agents. Consequently, for nonhuman animals to reap the benefits of metacognitive SLSs, they would need not only to be capable of representing 'who knows', but of communicating these beliefs to others. It is not inconceivable that this could be done without language, but it is unlikely that it could be done on any significant scale, and in a way that allows wisdom about 'who knows' to accumulate over time.

SLSs that identify 'who knows' with high accuracy and precision could not only benefit the individual, but also promote cultural evolution by making it safer for learners to: (i) copy one or a small number of models; (ii) copy them blindly, without introducing innovations based on asocial learning; and (iii) invest in the development and use of high-fidelity copying mechanisms, such as teaching and imitation [46,59]. It is now widely acknowledged that these high-fidelity mechanisms are not sufficient for cultural evolution, but it remains likely that they contribute [60,61]. All three of these effects (i–iii) would enhance 'parent–offspring relations' [62] (the heritability of differentially fit cultural variants) which is one of the three, major conditions for Darwinian evolution [63].

Previous discussions of the relation between SLSs and human culture have also emphasised the importance of high-fidelity copying mechanisms, such as language, teaching, imitation, and mental time travel [13,14]. These discussions imply that there is some connection between SLSs and high-fidelity copying, but they do not specify the nature of this connection, or explain why SLSs are special; how their role in the emergence and operation of high-fidelity copying differs from the roles of other psychological capacities, such as perceiving, remembering, and learning. By contrast, the dual systems account explains the relation between SLSs and human culture,

including the emergence of high-fidelity copying. It suggests that SLSs of one, comparatively rare type (the cook-like SLSs, based on explicit metacognition) improve the accuracy and precision with which agents can track 'who knows', and thereby create conditions in which high-fidelity mechanisms can evolve, develop, and operate.

I began with a question: If a wide range of animals use SLSs, and SLSs promote culture, why is human culture so much richer than that of other animals? I have argued that the answer is that the SLSs found in other animals, including rats, sticklebacks, fruit flies, and frog-eating bats, do not promote cumulative cultural change because they are based on general-purpose cognitive processes. These processes make planetary SLSs impressively supple, but do not enable them to achieve the accuracy and precision of metacognitive SLSs. Metacognitive SLSs, found only in humans, focus social learning on knowledgeable agents so precisely that they encourage blind, high-fidelity copying, a key requirement for cumulative cultural change. Metacognitive SLSs, the ones that really concern 'who knows', are an integral part of what makes us such peculiar, cultural animals.

References

1. Hills, T.T. *et al.* (2015) Exploration versus exploitation in space, mind, and society. *Trends Cogn. Sci.* 19, 46–54
2. Cohen, J.D. *et al.* (2007) Should I stay or should I go? How the human brain manages the trade-off between exploitation and exploration. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 362, 933–942
3. Richerson, P.J. and Boyd, R. (2005) *Not by Genes Alone: How Culture Transformed Human Evolution,* University of Chicago Press
4. Muthukrishna, M. *et al.* (2015) The when and who of social learning and conformist transmission. *Evol. Hum. Behav.* 37, 10–20
5. Rendell, L. *et al.* (2011) Cognitive culture: theoretical and empirical insights into social learning strategies. *Trends Cogn. Sci.* 15, 68–76
6. Wisdom, T.N. *et al.* (2013) Social learning strategies in networked groups. *Cogn. Sci.* 37, 1383–1425
7. Legare, C.H. and Nielsen, M. (2015) Imitation and innovation: The dual engines of cultural learning. *Trends Cogn. Sci.* 19, 688–699
8. Rendell, L. *et al.* (2010) Why copy others? Insights from the social learning strategies tournament. *Science* 328, 208–213
9. Galef, B.G. (2009) Strategies for social learning: testing predictions from formal theory. *Adv. Study Behav.* 39, 117–151
10. Webster, M.M. and Laland, K.N. (2015) Public information use by foraging ninespine sticklebacks: social learning or an unlearned social influence on travel direction? *Behaviour* 152, 1569–1584
11. Battesti, M. *et al.* (2015) Biased social transmission in Drosophila oviposition choice. *Behav. Ecol. Sociobiol.* 69, 83–87
12. Jones, P.L. *et al.* (2013) When to approach novel prey cues? Social learning strategies in frog-eating bats. *Proc. Biol. Sci.* 280, 20132330
13. Laland, K.N. and Rendell, L. (2013) Cultural memory. *Curr. Biol.* 23, R736–R740
14. Fogarty, L. *et al.* (2012) Mental time travel, memory and the social learning strategies tournament. *Learn. Motiv.* 43, 241–246
15. Heyes, C. (2012) What's social about social learning? *J. Comp. Psychol.* 126, 193
16. Dawson, E.H. *et al.* (2013) Learning by observation emerges from simple associations in an insect model. *Curr. Biol.* 23, 727–730
17. Wisenden, B.D. *et al.* (1997) Learned recognition of predation risk by Enallagma damselfly larvae (Odonata, Zygoptera) on the basis of chemical cues. *J. Chem. Ecol.* 23, 137–151
18. Rogers, A.R. (1988) Does biology constrain culture. *Am. Anthropol.* 90, 819–831
19. Whiten, A. *et al.* (1999) Cultures in chimpanzees. *Nature* 399, 682–685
20. Whiten, A. (2011) The scope of culture in chimpanzees, humans and ancestral apes. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 366, 997–1007
21. Raafat, R.M. *et al.* (2009) Herding in humans. *Trends Cogn. Sci.* 13, 420–428
22. Heyes, C. and Pearce, J.M. (2015) Not-so–social learning strategies. *Proc. Biol. Sci.* 282, 20141709
23. Cook, R. *et al.* (2014) Mirror neurons: from origin to function. *Behav. Brain Sci.* 37, 177–192
24. Heyes, C. (2012) Simple minds: a qualified defence of associative learning. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 367, 2695–2703
25. Hogarth, L. *et al.* (2013) Associative learning mechanisms underpinning the transition from recreational drug use to addiction. *Ann. N. Y. Acad. Sci.* 1282, 12–24
26. Behrens, T.E.J. *et al.* (2008) Associative learning of social value. *Nature* 456, 245–249
27. Nakahashi, W. *et al.* (2012) Adaptive social learning strategies in temporally and spatially varying environments. *Hum. Nat.* 23, 386–418
28. Hoppitt, W. and Laland, K.N. (2013) *Social Learning: An Introduction to Mechanisms, Methods, and Models,* Princeton University Press
29. Molleman, L. *et al.* (2014) Consistent individual differences in human social learning strategies. *Nat. Commun.* 5, 3570
30. Wood, L.A. *et al.* (2013) Whom do children copy? Model-based biases in social learning. *Dev. Rev.* 33, 341–356
31. Mesoudi, A. (2009) How cultural evolutionary theory can inform social psychology and vice versa. *Psychol. Rev.* 116, 929
32. Shea, N. *et al.* (2014) Supra-personal cognitive control and metacognition. *Trends Cogn. Sci.* 18, 186–193
33. Timmermans, B. *et al.* (2012) Higher order thoughts in action: consciousness as an unconscious re-description process. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 367, 1412–1423
34. Hurks, P.P.M. (2012) Does instruction in semantic clustering and switching enhance verbal fluency in children? *Clin. Neuropsychol.* 26, 1019–1037

35. Mahmoodi, A. *et al.* (2013) Learning to make collective decisions: the impact of confidence escalation. *PLoS ONE* 8, e81195

36. Bahrami, B. *et al.* (2012) Together, slowly but surely: the role of social interaction and feedback on the build-up of benefit in collective decision-making. *J. Exp. Psychol. Hum. Percept. Perform.* 38, 3

37. Heine, S.J. *et al.* (2001) Divergent consequences of success and failure in japan and north america: an investigation of self–improving motivations and malleable selves. *J. Pers. Soc. Psychol.* 81, 599

38. Mayer, A. and Träuble, B.E. (2013) Synchrony in the onset of mental state understanding across cultures? A study among children in Samoa. *Int. J. Behav. Dev.* 37, 21–28

39. Li, J. (2003) US and Chinese cultural beliefs about learning. *J. Educ. Psychol.* 95, 258

40. Güss, C.D. and Wiley, B. (2007) Metacognition of problem-solving strategies in Brazil, India, and the United States. *J. Cogn. Cult.* 7, 1–25

41. Williamson, R.A. *et al.* (2008) Prior experiences and perceived efficacy influence 3-year–olds' imitation. *Dev. Psychol.* 44, 275

42. Buttelmann, D. *et al.* (2013) Selective imitation of in-group over out-group members in 14-month-old infants. *Child Dev.* 84, 422–428

43. McGuigan, N. (2013) The influence of model status on the tendency of young children to over-imitate. *J. Exp. Child Psychol.* 116, 962–969

44. Chudek, M. *et al.* (2012) Prestige-biased cultural learning: bystander's differential attention to potential models influences children's learning. *Evol. Hum. Behav.* 33, 46–56

45. Heyes, C. (2015) When does social learning become cultural learning? *Dev. Sci.* Published online November 6, 2015. http://dx.doi.org/10.1111/desc.12350

46. Heyes, C. Born pupils? Natural pedagogy and cultural pedagogy. Perspect. Psychol. Sci. (in press)

47. Mesoudi, A. *et al.* (2015) Higher frequency of social learning in China than in the West shows cultural variation in the dynamics of cultural evolution. *Proc. Biol. Sci.* 282, 20142209

48. Toelch, U. *et al.* (2014) Individual consistency and flexibility in human social information use. *Proc. Biol. Sci.* 281, 20132864

49. Eriksson, K. (2012) The nonsense math effect. *Judgment Decis. Making* 7, 746–749

50. Henrich, J. and Broesch, J. (2011) On the nature of cultural transmission networks: evidence from Fijian villages for adaptive learning biases. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 366, 1139–1148

51. Efferson, C. *et al.* (2007) Learning, productivity, and noise: an experimental study of cultural transmission on the Bolivian Altiplano. *Evol. Hum. Behav.* 28, 11–17

52. Kawamura, S. (1959) The process of sub-culture propagation among Japanese macaques. *Primates* 2, 43–60

53. Heyes, C. (2012) Grist and mills: on the cultural origins of cultural learning. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 367, 2181–2191

54. van de Waal, E. *et al.* (2010) Selective attention to philopatric models causes directed social learning in wild vervet monkeys. *Proc. Biol. Sci.* Published online June 10, 2010. http://dx.doi.org/10.1098/rspb.2009.2260

55. McElreath, R. *et al.* (2008) Beyond existence and aiming outside the laboratory: estimating frequency-dependent and pay-off-biased social learning strategies. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 363, 3515–3528

56. Beran, M.J. *et al.* (2013) Language-trained chimpanzees (*Pan troglodytes*) name what they have seen but look first at what they have not seen. *Psychol. Sci.* 24, 660–666

57. Beran, M.J. *et al.* (2015) Go when you know: Chimpanzees' confidence movements reflect their responses in a computerized memory task. *Cognition* 142, 236–246

58. Kepecs, A. and Mainen, Z.F. (2012) A computational framework for the study of confidence in humans and animals. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 367, 1322–1337

59. Catmur, C. *et al.* (2009) Associative sequence learning: the role of experience in the development of imitation and the mirror system. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 364, 2369

60. Caldwell, C.A. and Millen, A.E. (2009) Social learning mechanisms and cumulative cultural evolution is imitation necessary? *Psychol. Sci.* 20, 1478–1483

61. Dean, L.G. *et al.* (2014) Human cumulative culture: a comparative perspective. *Biol. Rev.* 89, 284–301

62. Godfrey-Smith, P. (2012) Darwinism and cultural change. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 367, 2160–2170

63. Lewontin, R.C. (1970) The units of selection. *Annu. Rev. Ecol. Syst.* 1, 1–18

64. Pearce, J.M. (2008) *Animal Learning and Cognition: An Introduction,* Psychology Press

65. Schultz, W. and Dickinson, A. (2000) Neuronal coding of prediction errors. *Annu. Rev. Neurosci.* 23, 473–500

66. Evans, J.S.B.T. and Stanovich, K.E. (2013) Dual-process theories of higher cognition advancing the debate. *Perspect. Psychol. Sci.* 8, 223–241

67. Kahneman, D. (2011) *Thinking, Fast and Slow,* Macmillan

68. Reader, S.M. and Laland, K.N. (2002) Social intelligence, innovation, and enhanced brain size in primates. *Proc. Natl. Acad. Sci. U.S.A.* 99, 4436–4441

69. Reader, S.M. *et al.* (2011) The evolution of primate general and cultural intelligence. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 366, 1017–1027

70. Fiorito, G. and Scotto, P. (1992) Observational learning in *Octopus vulgaris*. *Science* 256, 545–547

71. Wilkinson, A. *et al.* (2010) Social learning in a non-social reptile (*Geochelone carbonaria*). *Biol. Lett.* 6, 614–616

72. Leadbeater, E. (2015) What evolves in the evolution of social learning? *J. Zool.* 295, 4–11

73. Dorrance, B.R. and Zentall, T.R. (2002) Imitation of conditional discriminations in pigeons. *J. Comp. Psychol.* 116, 277

74. Saggerson, A.L. *et al.* (2005) Imitative learning of stimulus-response and response-outcome associations in pigeons. *J. Exp. Psychol. Anim. Behav. Processes* 31, 289

75. Garvert, M.M. *et al.* (2015) Learning-induced plasticity in medial prefrontal cortex predicts preference malleability. *Neuron* 85, 418–428

76. Behrens, T.E.J. *et al.* (2007) Learning the value of information in an uncertain world. *Nat. Neurosci.* 10, 1214–1221

77. Nicolle, A. *et al.* (2012) An agent independent axis for executed and modeled choice in medial prefrontal cortex. *Neuron* 75, 1114–1121

78. Rushworth, M.F.S. *et al.* (2013) Are there specialized circuits for social cognition and are they unique to humans? *Curr. Opin. Neurobiol.* 23, 436–442

79. Pearce, J.M. and Collins, L. (1987) An evaluation of the associative strength of a partially reinforced serial CS. *Q. J. Exp. Psychol.* 39, 273–293

80. Norman, D.A. and Shallice, T. (1986) Attention to action: willed and automatic control of behaviour. In *Consciousness and Self-Regulation* (Davidson, R.J., Schwartz, G.E. and Shapiro, D., eds), pp. 1–18, New York, Plenum Press

81. Dehaene, S. and Naccache, L. (2001) Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition* 79, 1–37

82. Keren, G. and Schul, Y. (2009) Two is not always better than one a critical evaluation of two-system theories. *Perspect. Psychol. Sci.* 4, 533–550

83. Kruglanski, A.W. and Gigerenzer, G. (2011) Intuitive and deliberate judgments are based on common principles. *Psychol. Rev.* 118, 97

84. Metcalfe, J. (2009) Metacognitive judgments and control of study. *Curr. Dir. Psychol. Sci.* 18, 159–163

85. Yue, C.L. *et al.* (2013) When disfluency is—and is not—a desirable difficulty: the influence of typeface clarity on metacognitive judgments and memory. *Mem. Cogn.* 41, 229–241

86. Fleming, S.M. *et al.* (2012) Metacognition: computation, biology and function. *Philos. Trans. R. Soc. Lond. B: Biol. Sci.* 367, 1280–1286

87. Rigoni, D. *et al.* (2013) When errors do not matter: weakening belief in intentional control impairs cognitive reaction to errors. *Cognition* 127, 264–269

88. Brass, M. *et al.* (2013) Imaging volition: what the brain can tell us about the will. *Exp. Brain Res.* 229, 301–312

89. Job, V. *et al.* (2010) Ego depletion: is it all in your head? Implicit theories about willpower affect self-regulation. *Psychol. Sci.* 21, 1686–1693

90. Bahrami, B. *et al.* (2010) Optimally interacting minds. *Science* 329, 1081–1085

91. Morgan, T.J.H. *et al.* (2011) The evolutionary basis of human social learning. *Proc. Biol. Sci.* 279, 653–662

92. Apesteguia, J. *et al.* (2007) Imitation: theory and experimental evidence. *J. Econ. Theory* 136, 217–235

93. Banerjee, A.V. (1992) A simple model of herd behavior. *Q. J. Econ.* 197, 797–817

94. Cook, J.L. *et al.* (2014) The social dominance paradox. *Curr. Biol.* 24, 2812–2816

95. Diaconescu, A.O. *et al.* (2014) Inferring on the intentions of others by hierarchical Bayesian learning. *PLoS Comput. Biol.* 10, e1003810

96. Toelch, U. *et al.* (2013) The neural underpinnings of an optimal exploitation of social information under uncertainty. *Soc. Cogn. Affect. Neurosci.* 9, 1746–1753