



## Review

**Cite this article:** Heyes C, Pearce JM. 2015 Not-so-social learning strategies. *Proc. R. Soc. B* **282**: 20141709.  
<http://dx.doi.org/10.1098/rspb.2014.1709>

Received: 9 July 2014

Accepted: 11 December 2014

**Subject Areas:**

behaviour, cognition, evolution

**Keywords:**

associative learning, cultural evolution, evolution of cognition, selective learning, social learning strategies, transmission bias

**Author for correspondence:**

Cecilia Heyes

e-mail: [cecilia.heyes@all-souls.ox.ac.uk](mailto:cecilia.heyes@all-souls.ox.ac.uk)

## Not-so-social learning strategies

Cecilia Heyes<sup>1</sup> and John M. Pearce<sup>2</sup>

<sup>1</sup>All Souls College and Department of Experimental Psychology, University of Oxford, High Street, Oxford OX1 4AL, UK

<sup>2</sup>School of Psychology, Cardiff University, Tower Building, 70 Park Place, Cardiff CF10 3AT, UK

Social learning strategies (SLSs) are rules specifying the conditions in which it would be adaptive for animals to copy the behaviour of others rather than to persist with a previously established behaviour or to acquire a new behaviour through asocial learning. In behavioural ecology, cultural evolutionary theory and economics, SLSs are studied using a ‘phenotypic gambit’—from a purely functional perspective, without reference to their underlying psychological mechanisms. However, SLSs are described in these fields as if they were implemented by complex, domain-specific, genetically inherited mechanisms of decision-making. In this article, we suggest that it is time to begin investigating the psychology of SLSs, and we initiate this process by examining recent experimental work relating to three groups of strategies: *copy when alternative unsuccessful*, *copy when model successful* and *copy the majority*. In each case, we argue that the reported behaviour could have been mediated by domain-general and taxonomically general psychological mechanisms; specifically, by mechanisms, identified through conditioning experiments, that make associative learning selective. We also suggest experimental manipulations that could be used in future research to resolve more fully the question whether, in non-human animals, SLSs are mediated by domain-general or domain-specific psychological mechanisms.

## 1. Introduction

When is it a good idea to copy others? The term ‘social learning strategies’ (SLSs) originated 10 years ago to refer to rules specifying the conditions in which it might be adaptive for non-human animals to copy the behaviour of others rather than to persist with a previously established behaviour or to learn a new one through ‘asocial learning’—direct interaction with the inanimate environment [1]. Examples of the many SLSs that have been proposed include *copy when asocial learning is costly*, *copy successful individuals* and *copy the majority* [2]. The term ‘SLS’ and the lists of potential rules have been valuable both in stimulating and in summarizing intriguing empirical work on the conditions in which non-human animals (henceforth ‘animals’) engage in social learning [3]. Research on SLSs has also drawn attention to the crucial fact that all learning is selective. In the social domain, some potential informants (here called ‘models’) will provide more valuable information about the environment than others. Adaptation would then be best served by observers learning selectively from the individuals whose behaviour is most informative [4]. However, we believe that, after a decade in which SLSs have been investigated using ‘the phenotypic gambit’ [3]—from a purely functional perspective—it is time to look more closely at what they are made of, and to ask about the psychological mechanisms as well as the adaptive functions of SLSs (i.e. about the causes as well as the consequences of selectivity in social learning).

For several reasons, the psychological mechanisms we propose are based on the principles of associative learning. These principles have been used with great success to account for the phenomena of asocial learning [4], and can account for many aspects of social learning [5,6]. Moreover, the formal manner in which the principles have been presented means that it is possible to derive clear predictions from them concerning the various SLSs considered below. Finally, the ease with which the principles of associative can be embodied within connectionist networks provides a ready route for investigating the neural basis of both asocial and social learning.

In emphasizing the importance of psychological mechanisms we are not calling for a return to the days when research on social learning was dominated by taxonomy; by attempts to delineate types of social learning with labels such as ‘stimulus enhancement’, ‘local enhancement’, ‘emulation’ and ‘imitation’. Taxonomies of this kind typically delineated social learning *effects*, not social learning *mechanisms*; they focused on the behavioural products of social learning, but were silent about the psychological (and neurological) processes yielding these outcomes [5].

To make the case for mechanism-focused research, we examine focal experiments relating to three categories of SLS: *copy when alternative unsuccessful*, *copy when model successful* and *copy the majority*. These categories subsume all the various proposals that have been made concerning *when* social learning should take place, and *who* should be copied during such learning. The focal studies are recent, carefully conducted, representative examples of empirical work on SLSs, which relate to a range of species and tasks. In each of the three sections that follow, we ask whether the experimental findings could be due to the operation of domain-general psychological mechanisms—mechanisms that select among asocial cues, as well as among social cues, and between social and asocial cues—or whether the findings suggest that the selectivity of social learning is due to domain-specific psychological mechanisms—mechanisms that are dedicated to selecting among social cues and between social and asocial cues. Because research on SLSs has not been focused on mechanisms, the available data often do not tell us whether domain-general or domain-specific mechanisms are more likely to have mediated an experimental effect. In these cases, we do not recommend the use of a meta-empirical principle, such as Morgan’s Canon, to decide which hypothesis is correct [7]. Rather, we suggest specific experimental strategies that could more fully resolve the question. In the final section, we reflect on what our analysis implies about the sustainability of the phenotypic gambit as a research strategy, and, more broadly, about the evolution of social learning.

## 2. Copy when alternative unsuccessful

A number of SLSs suggest that animals are more likely to copy when the alternative to copying—persistence with a previously established behaviour or asocial learning of a new behaviour—is proving, or is likely to prove, unsuccessful. These strategies include *copy when dissatisfied* (which may subsume *copy when established behaviour is unproductive*), *copy when uncertain* (which may subsume *copy when prior information is outdated*) and *copy when asocial learning is costly* [3]. It is not always easy to understand why researchers have related a particular empirical study to a particular SLS. For example, many experimental manipulations, like many naturally occurring conditions, are likely to make animals dissatisfied, uncertain and subject to high costs. In these cases, SLS assignment is fairly arbitrary. However, following the conventions of the field, in this section, we consider experiments that have been said to relate to dissatisfaction, uncertainty and costly asocial learning, respectively.

A recent study of frog-eating bats provides an example of *copy when dissatisfied* [8]. In the first phase of this experiment (table 1), individual ‘observer’ bats learned to approach one of two auditory cues, A or B, for food. For

**Table 1.** Key elements of the design used by Jones *et al.* [8], 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

example, for A-trained bats, food was always available from the loudspeaker playing A. In the second phase, the observers were again able to retrieve food from the loudspeaker that played A, but now food could also be retrieved from the other loudspeaker, while it played B. The two cues were presented antiphonally. During this stage, food was made 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’, who was released at the same time as the observer, and who had been trained to find food on the loudspeaker playing B, but not A. Subsequent testing, in which each observer bat was presented with A played through one loudspeaker and B through the other, revealed a strong preference for A over B in group 50-solitary, and a preference for B over A in group 50-social. Thus, being trained in the presence of a bat that reliably went to B encouraged the observers in group 50-social to reverse their original preference for A over B. According to Jones *et al.* [8], this reversal was a consequence of group 50-social adopting the strategy of *copy when dissatisfied*, where the dissatisfaction arose from the introduction of the 50% schedule associated with A. Support for this interpretation is provided by the additional finding that during the test, group 100-social, which presumably did not experience dissatisfaction, did not relinquish its original preference for A.

When considering the results for group 50-social and group 100-social, it must be borne in mind that cue A had been intermittently paired with food for the former and consistently paired with food for the latter. Conditioning experiments, guided by associative learning theory, have shown that intermittently rewarded cues are typically less attractive than consistently rewarded cues [9]. Therefore, in the presence of both A and B together, it is not surprising that group 50-social was more willing than group 100-social to approach cue B. That is, the results from these groups may have occurred not because social learning had resulted in B being more attractive for group 50-social than group 100-social, but because A had been more reliably paired with food for group 100-social than group 50-social. In support of this proposal, it is noteworthy that during a final test with B by itself, there was no difference between these two groups.

Turning now to the results from group 50-solitary and group 50-social, this contrast suggests that observing the model approach B made B more attractive to the observer bats in group 50-social. However, there is no good reason

for believing this social learning was augmented in group 50-social by the dissatisfaction engendered by the intermittent food associated with A. For this conclusion to be justified, it would have to be shown that a similar social learning effect was *not* evident in groups trained in the same way as 50-social and 50-solitary, but with food available at A on every trial (100-social and 100-solitary). Thus, it appears that the dissatisfaction effect (group 50-social versus 100-social) was not necessarily social, and that the social learning effect (group 50-social and group 50-solitary) was not necessarily modulated by dissatisfaction.

A prominent example of *copy when uncertain* can also be explained by domain-general psychological processes modulating asocial learning [10]. In this experiment, rats consumed a mixture of two distinctively flavoured novel diets, cinnamon and cocoa (A plus B), or just one of these diets (A or B), before they were injected with an emetic. They were later allowed to interact with a model rat that had consumed one of two further diets, anise or marjoram (X or Y). This interaction was intended to result in a socially acquired preference for the food consumed by the model, through the experience of the food being paired with the smell of the model's breath [11]. In a subsequent choice test between X and Y, observers that had consumed the compound (A plus B) showed a stronger preference for the model's diet than rats that had consumed a single diet (A or B) at the beginning of the experiment. This was described as an example of *copy when uncertain*. It was assumed that consuming two flavours before toxicosis led to uncertainty about which of them caused the nausea. This uncertainty was then assumed to encourage rats to acquire future food preferences socially, or to place more reliance on socially acquired preferences. However, at the level of psychological mechanisms, it is possible that the effect was generated by the observers' attitude towards the alternative diet (Y), rather than towards the diet consumed by the model (X). In terms of associative learning theory, the consumption of A plus B before toxicosis will permit the generalization of the acquired aversion to a greater range of flavours, including X and Y, than when A or B alone is paired with toxicosis. At the same time, in both groups, any aversion that generalizes to X will be countered by the subsequent socially acquired attraction to this food [5,11]. Because the aversion to Y can be expected to be stronger in the A plus B than the A or B rats, the former will then be expected to express a stronger preference for X over Y than the latter. Thus, stimulus generalization, rather than social learning enhanced by uncertainty, could have been responsible for the results from the foregoing experiment.

Other studies using the same food preference paradigm, introduced by Galef, have shown that protein-deprived rats show a stronger preference for a diet encountered on a conspecific model's breath than protein-replete rats, and similarly enhanced 'Galef effects' for rats maintained on diluted or unpalatable diets, or under uncomfortable housing conditions [11]. All of these results have been classified as examples of *copy when dissatisfied* or *copy when uncertain* because, intuitively, it seems reasonable to describe rats maintained under these impoverished conditions as dissatisfied or uncertain. However, if one considers the psychological processes mediating the effects, it is possible that the 'dissatisfied' and 'uncertain' labels will turn out to be misleading. For example, each of these impoverished conditions could enhance socially acquired food preferences by increasing

the frequency and/or duration of the observer rats' interaction with their models. As noted above, the Galef effect depends upon the observer experiencing some property of the food in conjunction with a component of the model's breath [11]. As a consequence, any factor that promotes interaction with a demonstrator who has recently eaten the food in question will enhance the socially acquired attraction for that food. However, dissatisfaction and uncertainty do not always promote more intensive interaction with conspecifics, and more intensive interaction is not always caused by dissatisfaction or uncertainty. Therefore, across the full range of conditions that might be thought to induce dissatisfaction or uncertainty, these intuitive, folk psychological labels may provide poor predictors of the magnitude of social learning effects. Illustrating this point, Lindeyer *et al.* [12] recently found that, in adulthood, rats that had received relatively little licking and grooming from their mothers (low LG) showed weaker socially enhanced food preferences than rats that had received more licking and grooming (high LG). As the authors pointed out, previous research has indicated that low-LG rats are more risk-sensitive and anxious—at an intuitive level, they are more 'uncertain'—than high-LG rats. Therefore, this result conflicts with a prediction based on the intuitive formula *copy when uncertain*.

Now we turn to our final focal example of *copy when alternative unsuccessful*, relating to the subcategory *copy when asocial learning is costly*. According to this hypothesis, animals will rely on the knowledge acquired through asocial learning when there is a low risk of danger, but rely on social learning when this risk is high. Webster & Laland [13] tested this hypothesis with minnows who first learned asocially that food was available from feeder A and feeder B (experiment 1). Each observer was then confined to an observation box and allowed to watch conspecifics shoaling in the vicinity of feeder B rather than feeder A, before being given a choice between the two feeders under low, medium or high predation risk. The higher the level of risk, the stronger was the preference for B over A. This enhanced preference for the location where shoaling had taken place, as risk increased, was also seen in observers that had found food at A, but not at B, prior to observation (experiment 2).

Given that the experimental manipulation was introduced after conspecific observation, it is clear that this effect of predation risk was due to processes that mediate performance rather than learning; the use, rather than the acquisition, of public information. In principle, these performance processes could be domain-specific. For example, test behaviour might be regulated by a psychological process that is dedicated to 'gating' socially acquired information; that opens the gate, allowing such information to control behaviour, when predation risk or other indicators of costly asocial learning are high, and closes the gate when they are low. However, an alternative, domain-general explanation is also available. According to this account, initial training had two effects. (i) Through appetitive conditioning, a form of asocial learning, the minnows developed a tendency to search for food in location A, and location B in the case of experiment 1. (ii) Through social learning, the minnows acquired a tendency to approach location B, where conspecifics had previously been seen to shoal. There is abundant evidence that the rate at which an animal responds for food can be weakened considerably if an aversive stimulus, such as a cue for shock, is presented. This conditioned suppression effect has been

observed in a variety of species, including goldfish [14]. As the level of threat increases, therefore, it is likely that it will exert a similar, suppressive effect and reduce any tendency to search for food in A, or anywhere else, and thus allow the tendency to approach B, solely because of its association with a shoal of conspecifics, to manifest itself fully.

This domain-general explanation is consistent with the results of two studies that have failed to find support for *copy when asocial learning is costly*. In these studies, rats that encountered a distinctive diet on the breath of a conspecific model did not show an enhanced preference for the model's diet when subsequently tested under high rather than low predation risk [15,16]. They did, however, show a reduced consumption of food during the test trials, as would be expected if high predation risk serves to suppress appetitive behaviour.

### 3. Copy when model successful

In §2, we argued that effects categorized as *copy when alternative unsuccessful* could be due to domain-general psychological processes; mechanisms that mediate learning in general, not just social learning. Perhaps this is not surprising. After all, *copy when alternative unsuccessful* implies that it is the efficiency of the 'alternative' that regulates the acquisition and use of social information, and the alternative is often asocial learning—the domain in which theories of associative learning have been developed using conditioning procedures. In this section, we consider examples of *copy when model successful*. This label implies that it is the pay-offs associated with the social option, rather than the primary domain of associative learning, that drive the acquisition and use of social information. Therefore, one might expect to find stronger evidence here that SLSs are mediated by domain-specific psychological mechanisms.

Some SLSs within the category *copy when model successful* suggest that animals estimate model success directly, by registering the pay-offs the model is receiving for their current behaviour (the behaviour that the observer is deciding whether or not to copy). Others suggest that animals estimate model success indirectly, using age, kin status, social status or gender as proxies for success. Strategies of the latter kind would be adaptive in ecologies where, for example, older or higher-status individuals are more likely to 'know the ropes'.

Horner *et al.* [17] reported an effect of a model's social status or 'prestige', an indirect indicator of success, in chimpanzees. After observing a high-status model, A, depositing tokens in a striped container for food reward, and simultaneously a low-status model, B, depositing tokens in a spotted container for food reward (counterbalanced), the observer chimpanzees were more likely to deposit their own tokens in the container used by A. This study has novelty value because SLSs have rarely been studied in chimpanzees, but the results are tantalizing for several reasons: the effect is small; models A and B differed, not only in social status, but also in age, and possibly also in size; and, crucially, the study gave us very little information about the behaviour of the models while they were being observed. This omission is important because there is evidence that animals find it is easier to attend to, and hence learn about, a stimulus when it is presented alone, rather than with another stimulus [18]. Given that model B was subordinate to model A, and the two were performing in a competitive

environment (only one model at a time could be supplied with a token or a food reward), it is possible that B deferred to A in a way that would make B's behaviour less likely to command the attention of the observers. For example, B may have waited until A had deposited many of her tokens, and the novelty of the spectacle had worn off for the observers, or tended to make her deposits just after A, when A's attention was focused on her container, and the observers' attention was focused on A's food reward. If B behaved in a way that was less likely to command attention, the 'status' effect observed in this experiment could be due to domain-general processing; all learning, not just social learning, demands attention to the to-be-learned stimuli [4]. Moreover, if the observers were less likely to copy B than A because they paid less attention to B than to A, the effect would be absent in conditions where status is not confounded with attention. For example, if A and B were observed successively rather than observed simultaneously, it is possible the 'status' effect would disappear, because observers would be able to pay full attention to each model. Of course, if the 'status' effect remained, then this might mean nothing more than that, even in isolation, high-status chimpanzees are given more attention than those of low status.

Another study of primates shows that, even when a social learning effect can be linked more firmly with an indirect indicator of model success, it can be explained with reference to domain-general psychological processes. In this study, wild vervet monkeys repeatedly observed a dominant female or a dominant male retrieving food from a box via one of two doors [19]. When subsequently allowed access to the box themselves, observers of female models showed a stronger preference for the door used by the model than did observers of male models. (In vervets, females are the philopatric sex; they remain in their natal group all their lives, whereas males migrate to another group when they are sexually mature. This is therefore a putative example of *copy when model successful* because it is possible that, on average, females are more knowledgeable with respect to the local environment than males.) Furthermore, using careful measures of attention, this study found that observers were more likely to look at female than male models at the moment when they were opening the box. This result suggests that, at the level of psychological mechanisms, the female models were more effective than the males because they commanded more attention. In principle, this could be because vervets are born with a learning mechanism that includes a learning rate parameter, the value of which varies with the sex of the model and/or another indicator of average success. However, it is at least equally plausible that vervets learn to attend more to females than males as a result of experience in which female behaviour provides a better predictor of reward, and that this 'learning to attend' occurs via the same mechanisms that modulate attention to inanimate stimuli. The properties of these domain-general mechanisms have been revealed by experiments in which animals have received discrimination training with some stimuli that were relevant to the delivery of reward, and some that were irrelevant. When these stimuli are used in a new task, learning about the previously relevant stimuli progresses more readily than about the previously irrelevant stimuli [20]. Such an outcome shows that through experience animals come to pay more attention, and hence learn more readily about relevant rather than irrelevant stimuli [21]. To find out whether vervets and other animals learn to attend to particular categories of

model (e.g. females versus males) via these domain-general mechanisms, one could conduct similar transfer experiments. For example, if animals have the opportunity to learn in one task that the behaviour of type A models is a better predictor of reward or 'success' than the behaviour of type B models, do the observers attend more to type A models in a subsequent task of a different kind?

In our final example of *copy when model successful*, Coolen *et al.* [22] asked whether nine-spined sticklebacks can use a direct, rather than an indirect, indicator of success in deciding whether to copy. While confined in a transparent chamber, individual nine-spined sticklebacks were allowed simultaneously to observe two groups of conspecific models, one feeding from a 'rich' feeder and another from a 'poor' feeder (experiments 2 and 4). The rich feeder delivered several blood worms six times in 10 min, whereas the poor feeder delivered several blood worms on only two or three occasions in 10 min. After this observation phase, the models were removed from the tank, and the observers were released from the transparent chamber. Upon release, the observers were more likely to make their first approach to the rich feeder, and they spent longer in the vicinity of the rich than the poor feeder. This preference for the rich over the poor feeder was interpreted by Coolen *et al.* as evidence that the nine-spined sticklebacks had made 'a judgement as to the relative profitability of resources on the basis of the success of others' (p. 2417), but there are two obstacles to this interpretation. First, the sticklebacks may have preferred the rich feeder because more blood worms, rather than more feeding success, had been observed—visually or via olfaction—in the vicinity of the rich than the poor feeder. To eliminate this possibility, it would be necessary to examine the choice behaviour of control groups that had observed the rich and poor feeders releasing blood worms while conspecifics shoaled nearby but did not feed on the worms. Second, even if the observers' preference for the rich feeder depended on their observation of the models' feeding behaviour, it can be explained readily by associative learning theory. If they ever fed in a group prior to the experiment, the fish will have experienced a correlation between feeding themselves (a primary reinforcer), and observing others feeding. This experience will have established the sight of other fish feeding as a secondary reinforcer; a stimulus which, when paired with other cues—such as those identifying a feeder—would make the other cues attractive. Thus, when the observer fish were released into the tank, they would be more likely to approach the rich than the poor feeder because the rich feeder had been paired more than twice as often with secondary reinforcement (i.e. the sight of other fish feeding). In that case, although the observer fish could be said to have made 'a judgement as to the relative profitability of resources on the basis of the success of others', they could equally well be said to have shown observational conditioning, and the two descriptions would lead to different predictions about their behaviour in other conditions. For example, on the associative account outlined above, one would expect the effect to depend on fish having the opportunity to feed in groups before the experiment.

In a later study [23], Webster & Laland attempted to test the observational conditioning hypothesis by increasing illumination of the tank after the observers had watched the models feeding but before they were tested. It was assumed that this change would eliminate the effects of observational conditioning through generalization decrement. This manipulation,


however, had no effect on the magnitude of the observer sticklebacks' preference for the rich over the poor feeder, which makes it difficult to draw clear theoretical conclusions from the experiment. For example, the extent of the stimulus change may have been insufficient to engender detectable generalization decrement. Furthermore, it is not clear why Webster & Laland assumed that the effects of observational conditioning alone should be susceptible to generalization decrement. Therefore, the null results of this experiment indicate that the change in illumination was not great enough to have an impact on choice performance, but do not tell us whether or not the learning process was associative.

#### 4. Copy the majority

The final category of SLS to be considered here, *copy the majority*, is especially interesting because population-level modelling suggests that one of the SLSs in this group, known as *disproportionately copy the majority* or *hyperconformity*, plays a key role in making behaviour homogeneous within social groups, and therefore in promoting cultural evolution [24]. It hardly needs to be said that the effectiveness of social learning will be influenced by the proportion of the population that exhibits the behaviour to be copied. The more animals within a group that demonstrate a particular activity, the more likely it will be that an observer will experience an opportunity to learn about it. However, an exaggerated or disproportionate tendency to copy the majority would go beyond this influence and rapidly promote uniformity of behaviour within groups. The emergence of such excessive conformity would present a puzzle at the psychological level, because it would show that something more than the frequency with which a particular activity is observed determines the effectiveness of learning about that activity.

As far as we are aware, only one study, by Pike & Laland [25], purports to show that animals have a disproportionate tendency to copy the majority. This study used procedures similar to those of Coolen *et al.* [22] (described above). Pike & Laland first allowed nine-spined sticklebacks to learn through their own efforts that feeder A delivered more worms than feeder B, and then divided the fish into six groups. Three groups observed, respectively, three, four or five members of a group of six conspecific models shoaling (moving in close spatial proximity to one another) and feeding at B, and the remainder of the models shoaling and feeding at A, when A delivered fewer worms than B (groups SF 3–3, SF 2–4, SF 1–5). The other three groups were the same except that no food was delivered during observation, and therefore the models were observed shoaling but not feeding at A and B (groups S 3–3, S 2–4, S 1–5).

Figure 1*a* shows the results from the groups that had originally observed models shoaling and feeding at A and B. The most striking feature of these results is that the magnitude of the difference in mean preference for B over A in SF 1–5 compared with SF 2–4 is nearly six times greater than that between SF 2–4 and SF 3–3. It was this feature that led Pike & Laland to present these data as evidence of disproportionate copying of the majority. However, if one looks at figure 1*b*, which shows the equivalent results for the groups that observed models shoaling without feeding, a similar pattern can be seen. Group S 1–5 showed a substantially stronger preference for B over A, and yet, on average, neither S 2–4 nor S



**Figure 1.** Mean preference for option B (see text for details). Redrawn from reference [25].

3–3 showed any preference between B and A. Thus, the results in figure 1b demonstrate that the tendency to approach location B in preference to A increased disproportionately as the ratio of fish originally observed shoaling at B relative to A increased. Presumably, this influence was also effective in the groups that observed shoaling and feeding, and, by itself, can account for the trend observed in figure 1a. Consequently, the stronger preference for B over A shown by each of the three SF groups relative to their S counterparts confirms that social learning about feeding was effective, but it remains to be demonstrated that this kind of social learning increases disproportionately as the proportion of the members of a group displaying feeding increases. Furthermore, to the extent that the especially strong preference for B in SF 1–5 was due to the shoaling component of the models' behaviour, it may not have been disproportionate with respect to the 'active ingredient' of the model stimulus. For example, it is possible that attention to a shoal, or the reinforcing power of a shoal stimulus, depends on the amount or type of activity within the shoal, and that activity patterns change disproportionately when the shoal size increases from four to five. In this case, the effect would not be replicated in tasks or species where the size and the activity pattern of the majority were unconfounded. Thus, to provide a more compelling demonstration of disproportionate copying of the majority, it would be necessary to employ tasks where the size and the activity pattern of the majority were dissociated. This might be achieved by giving larger majorities more space in which to move around than smaller majorities, or by presenting models successively rather than in groups.

## 5. Implications

In this article, we have examined some of the most interesting and persuasive studies of SLSs in animals. In the growing field from which these studies have emerged, researchers write about SLSs as if the following propositions were true. (1) Social and asocial learning depend on different psychological mechanisms. (2) The mechanisms that mediate social learning evolved genetically (via selection operating on genetic variance) in response to pressures from the social environment. (3) Animals can decide to use social learning, asocial learning or both to guide their behaviour in any given situation. (4) Animals use domain-specific mechanisms (SLSs) to decide whether to use social learning, asocial learning or both. (5) These domain-specific mechanisms involve reasoning (conscious or unconscious) and concepts

such as 'copying', 'certainty', 'success' and 'majority'. (6) The domain-specific decision mechanisms (SLSs) evolved genetically.

Building on previous discussions of (1) and (2) [5,6], our survey of research on SLSs casts doubt on all of these propositions, and raises the following possibilities for investigation in future research. (1') At least in animals, social and asocial learning are mediated by the same associative processes. Social learning differs from asocial learning at the level of cues, not of processes; learning is called 'social learning' when one or more of the to-be-learned cues is carried by, or instantiated in, another agent. (2') The domain-general associative mechanisms that mediate learning about both social and asocial cues evolved genetically, but as a means of tracking predictive relationships between events, not specifically in response to selection pressure from the social environment. (3') Social and asocial learning are not distinct processes that can be switched on and off independently. (4') The mechanisms that make learning selective—for example, modulate the extent to which animals learn about social and asocial cues—are described by associative learning theory [4]. Thus, learning is modulated by the salience of the events involved, and the amount of attention that is paid to them; it is also modulated by the temporal relationship between one event and another, and by the contingency between the two events. (5') Because associative mechanisms make learning selective, they can be described in this minimal sense as 'decision processes' or as learning 'strategies', but they do not involve conscious or unconscious reasoning. (6') The associative mechanisms that make all learning selective are genetically inherited. However, at least in animals, there are no mechanisms dedicated to determining whether the learner should depend on social or asocial learning; no specifically *social* learning strategies. Therefore, it is unnecessary to ask to what extent such strategies are products of genetic evolution, cultural evolution and/or learning.

Leading investigators have made clear that, as part of the phenotypic gambit, they discuss SLSs 'as if' propositions 1–5 were true, but this language does not represent genuine theoretical commitments. For example, Hoppitt & Laland [3] suggest that 'as a reasonable first approximation, research into learning strategies can proceed through functional considerations without any commitment to mechanism, and researchers may assume that it does not matter whether animals adopt such strategies as a consequence of evolved psychological mechanisms, learning, culture, or some combination of processes' (p. 200). This may well be correct as a 'first approximation'. It has certainly led many researchers to take an interest in SLSs, and to the generation of new

data on social learning in animals. However, we suspect that the current practice—using language that implies strong commitment to certain mechanisms, while suggesting that functions can be studied without concern for mechanisms—is now in danger of leading the field astray.

At the most fundamental level, the problem relates to generalization and prediction. When SLSs are described as if they were domain-specific, genetically evolved mechanisms (as if 1–6 were true) reporting a particular experimental effect as showing, for example, that ‘species X uses the SLS *copy the majority*’ promotes certain expectations. It leads one to expect species X—and perhaps other closely related species, but not more distantly related animals—to exhibit behaviour that fits the description *copy the majority* whenever they are exposed to more models exhibiting one behavioural variant rather than another. But if the psychological mechanisms that make social learning selective are *really* domain-general and taxonomically general associative mechanisms, these expectations will not be fulfilled. Under the same conditions, a very broad range of species—not just those that are closely related to X—will show similar behaviour, and under different conditions—for example, when models are observed successively rather than simultaneously, or when the minority behaviour is especially salient—even species X will not *copy the majority*. Thus, if research on SLSs continues to ignore mechanisms, there is a risk that it will produce systematically

misleading generalizations about the conditions in which animals do and do not engage in social learning.

More generally, until research on SLSs begins to enquire about mechanisms, it will have very limited potential to tell us about the evolution of behaviour, cognition and culture. Obviously, a research strategy that simply assumes, at least in its use of language, that behavioural plasticity evolves genetically through the accretion of domain-specific cognitive mechanisms is ill-suited to casting light on whether this assumption is valid. This general problem is especially pressing in relation to the evolution of culture. Most research on SLSs is explicitly aimed at uncovering the evolutionary roots of the human capacity for cumulative cultural evolution. However, it is possible that the SLSs—or ‘transmission biases’ [24]—supporting human cultural evolution are rooted in entirely different mechanisms from those that yield SLS effects in animals. For example, the latter could be domain-general products of genetic evolution, whereas the former are domain-specific products of cultural evolution. Whatever the merits of this particular suggestion, which is discussed elsewhere in relation to research on SLSs in human adults and children [26,27], we urge those who study SLSs in animals to build on their important work to date by conducting experiments that investigate, not just when and who animals copy, but how and why they make those decisions.

## References

- Laland KN. 2004 Social learning strategies. *Anim. Learn. Behav.* **32**, 4–14. (doi:10.3758/BF03196002)
- Rendell L *et al.* 2011 Cognitive culture: theoretical and empirical insights into social learning strategies. *Trends Cogn. Sci.* **15**, 68–76. (doi:10.1016/j.tics.2010.12.002)
- Hoppitt W, Laland KN. 2013 *Social learning: an introduction to mechanisms, methods, and models*. Princeton, NJ: Princeton University Press.
- Pearce JM. 2008 *Animal learning and cognition: an introduction*. Hove, UK: Psychology Press.
- Heyes C. 2012 What’s social about social learning? *J. Comp. Psychol.* **126**, 193–202. (doi:10.1037/a0025180)
- Heyes C. 1994 Social learning in animals: categories and mechanisms. *Biol. Rev.* **69**, 207–231. (doi:10.1111/j.1469-185X.1994.tb01506.x)
- Heyes C. 2012 Simple minds: a qualified defence of associative learning. *Phil. Trans. R. Soc. B* **367**, 2695–2703. (doi:10.1098/rstb.2012.0217)
- Jones PL, Ryan MJ, Flores V, Page RA. 2013 When to approach novel prey cues? Social learning strategies in frog-eating bats. *Proc. R. Soc. B* **280**, 20132330. (doi:10.1098/rspb.2013.2330)
- Pearce JM, Collins L. 1987 An evaluation of the associative strength of a partially reinforced serial CS. *Q. J. Exp. Psychol.* **39**, 273–293.
- Galef Jr BG, Dudley KE, Whiskin EE. 2008 Social learning of food preferences in ‘dissatisfied’ and ‘uncertain’ Norway rats. *Anim. Behav.* **75**, 631–637. (doi:10.1016/j.anbehav.2007.06.024)
- Galef BG. 2009 Strategies for social learning: testing predictions from formal theory. *Adv. Stud. Behav.* **39**, 117–151. (doi:10.1016/S0065-3454(09)39004-X)
- Lindeyer CM, Meaney MJ, Reader SM. 2013 Early maternal care predicts reliance on social learning about food in adult rats. *Dev. Psychobiol.* **55**, 168–175. (doi:10.1002/dev.21009)
- Webster MM, Laland KN. 2008 Social learning strategies and predation risk: minnows copy only when using private information would be costly. *Proc. R. Soc. B* **275**, 2869–2876. (doi:10.1098/rspb.2008.0817)
- Geller I. 1964 Conditioned suppression in goldfish as a function of shock reinforcement schedule. *J. Exp. Anal. Behav.* **7**, 345–349. (doi:10.1901/jeab.1964.7-345)
- Galef Jr BG, Whiskin EE. 2006 Increased reliance on socially acquired information while foraging in risky situations? *Anim. Behav.* **72**, 1169–1176. (doi:10.1016/j.anbehav.2006.05.003)
- Galef BG, Yarkovsky N. 2009 Further studies of reliance on socially acquired information when foraging in potentially risky situations. *Anim. Behav.* **77**, 1329–1335. (doi:10.1016/j.anbehav.2009.01.038)
- Horner V *et al.* 2010 Prestige affects cultural learning in chimpanzees. *PLoS ONE* **5**, e10625. (doi:10.1371/journal.pone.0010625)
- Mackintosh NJ. 1971 An analysis of overshadowing and blocking. *Q. J. Exp. Psychol.* **23**, 118–125. (doi:10.1080/0033557143000121)
- van de Waal E, Renevey N, Farve CM, Bahary R. 2010 Selective attention to philopatric models causes directed social learning in wild vervet monkeys. *Proc. R. Soc. B* **282**, 20092260. (doi:10.1098/rspb.2009.2260)
- Dopson JC, Esber GR, Pearce JM. 2010 Differences in the associability of relevant and irrelevant stimuli. *J. Exp. Psychol. Anim. Behav. Process.* **36**, 258. (doi:10.1037/a0016588)
- George DN, Pearce JM. 1999 Acquired distinctiveness is controlled by stimulus relevance not correlation with reward. *J. Exp. Psychol. Anim. Behav. Process.* **25**, 363. (doi:10.1037/0097-7403.25.3.363)
- Coolen I, Bergen YV, Day RL, Labnd KN. 2003 Species difference in adaptive use of public information in sticklebacks. *Proc. R. Soc. Lond. B* **270**, 2413–2419. (doi:10.1098/rspb.2003.2525)
- Webster MM, Laland KN. 2013 The learning mechanism underlying public information use in ninespine sticklebacks (*Pungitius pungitius*). *J. Comp. Psychol.* **127**, 154. (doi:10.1037/a0029602)
- Richerson PJ, Boyd R. 2005 *Not by genes alone: how culture transformed human evolution*. Chicago, IL: University of Chicago Press.
- Pike TW, Laland KN. 2010 Conformist learning in nine-spined sticklebacks’ foraging decisions. *Biol. Lett.* 20091014.
- Heyes C. In preparation. The cultural evolution of social learning strategies.
- Heyes C. Submitted. Selective trust: when does social learning become cultural learning?