

Research

How copying affects the amount, evenness and persistence of cultural knowledge: insights from the social learning strategies tournament

L. Rendell^{1,*}, R. Boyd², M. Enquist^{3,4}, M. W. Feldman⁵,
L. Fogarty¹ and K. N. Laland^{1,*}

¹*Centre for Social Learning and Cognitive Evolution, School of Biology, University of St Andrews, Westburn Lane, St Andrews, Fife KY16 9TS, UK*

²*Department of Anthropology, University of California, Los Angeles, CA 90095, USA*

³*Department of Zoology, Stockholm University, 11691 Stockholm, Sweden*

⁴*Centre for the Study of Cultural Evolution, Stockholm University, 11691 Stockholm, Sweden*

⁵*Department of Biological Sciences, Stanford University, Stanford, CA 94305, USA*

Darwinian processes should favour those individuals that deploy the most effective strategies for acquiring information about their environment. We organized a computer-based tournament to investigate which learning strategies would perform well in a changing environment. The most successful strategies relied almost exclusively on social learning (here, learning a behaviour performed by another individual) rather than asocial learning, even when environments were changing rapidly; moreover, successful strategies focused learning effort on periods of environmental change. Here, we use data from tournament simulations to examine how these strategies might affect cultural evolution, as reflected in the amount of culture (i.e. number of cultural traits) in the population, the distribution of cultural traits across individuals, and their persistence through time. We found that high levels of social learning are associated with a larger amount of more persistent knowledge, but a smaller amount of less persistent expressed behaviour, as well as more uneven distributions of behaviour, as individuals concentrated on exploiting a smaller subset of behaviour patterns. Increased rates of environmental change generated increases in the amount and evenness of behaviour. These observations suggest that copying confers on cultural populations an adaptive plasticity, allowing them to respond to changing environments rapidly by drawing on a wider knowledge base.

Keywords: social learning; cultural evolution; simulation modelling

1. INTRODUCTION

Understanding the evolution of human culture is one of the greatest challenges facing science. The gulf between the complexity of human culture and cognition and that observed in other animals is so vast that to many it has seemed unbridgeable. Nonetheless, evolutionary links are there to be found. Most behavioural biologists now acknowledge, for example, the existence of diverse behavioural traditions observed in other apes and monkeys [1–3], the surprisingly complex cognition of corvids [1], and the impressive collective decision-making and rich social behaviour of insect societies [2,3]. Yet, the fact remains that however much we talk up these phenomena, there is a chasm

between our achievements and theirs. If one accepts, as we do, the argument that chimpanzees, and for that matter monkeys, whales, birds and fishes, have some semblance of culture, then one must acknowledge that the ‘culture’ of nonhuman animals is very different from our own.

Such reasoning leads to two kinds of question. First, in what ways do the processes that underlie human culture differ from those observed in other animals, such that they can create such distinct patterns of behavioural, social and technological complexity. Second, how did those processes that underlie human culture evolve out of the kind of rudimentary capabilities observed in other animals? In simple terms, we can ask ‘what is the gap?’ and ‘how can we bridge it?’ The traditional routes to addressing these questions are well represented in the other contributions to this issue [4–7]. Here we take a different approach, drawing on and extending the insights into the evolution of culture that emerged from the social learning strategies tournament [8]. In this paper, we will first review the setup and results of this tournament, and then use the same model

* Authors for correspondence (ler4@st-andrews.ac.uk; knl1@st-andrews.ac.uk).

Electronic supplementary material is available at <http://dx.doi.org/10.1098/rstb.2010.0376> or via <http://rstb.royalsocietypublishing.org>.

One contribution of 26 to a Discussion Meeting Issue ‘Culture evolves’.

framework to investigate how the strategies that were successful in the tournament affect cultural evolution.

Our tournament was a competition designed to understand the most effective means to learn in a complex, changeable world. Similar tournaments have proved an effective means for investigating the evolution of cooperation [9]. The background to the tournament was the observation that while social learning is central to the capacity for human culture, and while human culture is widely thought to be responsible for our success as a species [10], it remains something of a mystery as to why individuals profit by copying others and how best to do this (note that throughout this paper, we use the term ‘copying’ as synonymous with social learning in the broad sense—any form of socially contingent learning by which individuals come to do what others already do [11]). At first sight, social learning appears advantageous because it allows individuals to avoid the costs of trial-and-error learning. However, theoretical work shows that this advantage can be offset if social learning leads individuals to acquire inappropriate or outdated information in non-uniform and changing environments [10,12–15]. Current theory suggests that to avoid these errors individuals should be selective in when and how they use social learning [10,16], and that natural selection should favour the best ‘social learning strategies’—psychological mechanisms that specify when individuals copy and from whom they learn [17]. Formal theoretical analyses [10,15,18–21] and experimental studies [22–24] have explored a small number of strategies. However, for a more authoritative understanding the relative merits of a large number of alternative social learning strategies must be assessed. To address this, we organized a computer tournament, in which strategies competed in a complex and changing simulation environment, with a €10 000 prize awarded to the winner [8]. In this article, we extend our earlier analyses of the tournament results to consider how the strategies that did well affect the *amount*, *evenness* and *persistence* of cultural traits.

Among the striking differences between human and animal culture is the sheer amount of culture that humans possess. Here, the *amount* of culture refers to the number of cultural traits that a population knows about, or actually expresses in behaviour, which we call the amount of knowledge and behaviour, respectively. We explore the relationship between these quantities and the learning strategies deployed in a population, expressed as the amount and type (social or asocial) of learning used.

We also consider how evolutionarily successful strategies affect the *evenness* of culture, which we define as the flatness or uniformity of the distributions of knowledge and expressed behaviour across a population. Logic suggests that copying should increase the behavioural evenness exhibited by a population, since copying generates homogeneity in exhibited behaviour, but it is less clear whether and how copying will affect the evenness of acquired knowledge. Once again, we use the output from the tournament to shed light on this issue.

Finally, we examine how these strategies affect the *persistence* of cultural traits, which we define here as

the average length of time individual traits persist, given that they became known or expressed by at least one individual, either in the knowledge or the expressed behaviour of at least one individual in a population. Human culture is uniquely cumulative, with each generation building upon the cultural knowledge of the previous generations [25]. Cumulative culture requires cultural traits to persist for long enough to allow refinements or elaborations of acquired knowledge [26]. The cultures of other animals are frequently characterized by ‘lightning traditions’, which rapidly sweep through a group of animals, and then are replaced as quickly, with little sign of any accumulation of knowledge [27]. Conversely, humans today possess knowledge first acquired many thousands of years ago. It is plausible that the preservation of acquired knowledge over long periods of time creates the opportunity for refinement, elaboration and diversification [26], and that this again explains some of the uniqueness of human culture. We use the data generated by the tournament to explore how the longevity of cultural knowledge is affected by the proportion of learning that is copying. We consider how the level of copying affects the average persistence of both exploited behaviour and behavioural knowledge in the population’s repertoire.

One of the challenges facing a developing theory of cultural evolution is to link the small-scale, social-learning decisions of individuals to the creation and subsequent evolution of the collections of knowledge, tradition, language and behaviour that characterize populations at the level analysed in other contributions to this issue [28,29]. As we describe below, the tournament is effective both because it proposes specific means by which copying may be implemented strategically to enhance copying efficiency, and because it illustrates the population-level consequences of such strategies. In the sections below we first summarize the methods and findings of the social-learning strategies tournament [8]. Second, we present analyses of tournament simulations designed to shed light on how copying affects the longevity, evenness and amount of cultural knowledge in the virtual ‘cultures’ constructed by the strategies that did well in our tournament. Finally, we will attempt to collate these insights to shed light on both the field of cultural evolution and, more generally, understanding of the evolution of culture.

2. THE TOURNAMENT

(a) *Methods*

To investigate why copying is so widespread in the animal kingdom, and in humans in particular, we organized a computer simulation tournament [8]. Entrants were challenged to specify a strategy to enable agents to survive and prosper in a simulated environment.

In this simulated environment 100 agents could learn about 100 different behavioural acts, each with a distinct payoff, drawn from an exponential distribution. This learning problem belongs to a class termed ‘restless multi-armed bandits’ commonly used in a variety of

fields to study learning [30,31]. The environment varied, as payoffs changed with a fixed probability, p_c , per simulation iteration.

The goal for entrants to the tournament was to design a simple piece of software (submitted either in code or pseudo code) that specified the circumstances under which individuals with their strategy should learn asocially (play INNOVATE), learn socially (play OBSERVE) or perform an act (play EXPLOIT). INNOVATE returned accurate information about the payoff of a randomly selected behaviour previously unknown to the agent. OBSERVE returned noisy information about the behaviour and payoff currently being demonstrated in the population by n_{observe} other agents, selected at random from those playing EXPLOIT (i.e. at least one individual in the population needed to be performing a behaviour pattern for it to be observable by others). Playing OBSERVE could return no behaviour if none was demonstrated or if a behaviour that was already in the agent's repertoire was observed. Finally, playing EXPLOIT performed an act from the individual's repertoire, chosen by the agent (more specifically, the strategy controlling that agent), and the agent received the associated payoff—this was the only way in which an agent could actually acquire payoffs.

OBSERVE moves were error prone in two ways. There was a small probability of acquiring a behaviour different from that being observed, $p_{\text{copyActWrong}}$, in which case an individual added a random act from the 99 unperformed acts to its repertoire. This act was still associated with the payoff received by the observed individual. Independently, the payoff associated with each act observed was subject to a normally distributed random error with mean 0 and standard deviation $\sigma_{\text{payoffError}}$.

Entrants were given some information about the simulation environment, and strategies had access to agents' own personal histories. They were free to estimate parameters like n_{observe} , p_c , and the errors associated with each move using this information. Each agent also possessed a behavioural repertoire, empty at birth. The agent could only acquire knowledge, with which it could then acquire payoffs through playing EXPLOIT, by adding to its repertoire using either of the two learning moves OBSERVE or INNOVATE.

The evolutionary dynamics of the tournament simulation were modelled as a death–birth process with each individual having a per-iteration probability of dying, fixed at 1/50, giving each agent an expected lifespan of 50 rounds. After each death, individuals were selected to reproduce in proportion to their mean lifetime payoff (p , the sum of the payoffs gained by playing EXPLOIT divided by the number of iterations the individual had been alive) and their offspring replaced dying individuals. The probability of individual z reproducing was $p_z/\Sigma p$, where Σp was the summed mean lifetime payoff of the population in that iteration. Offspring generally inherited their parent's strategy, but could mutate and so carry a different strategy. This mutation rate (set at a probability of 1/50 per birth) allowed new strategies to invade the population.

The tournament was run in two stages, a pairwise round-robin stage and a melee that included the top

10 strategies from the first stage. Each pairwise contest consisted of 10 simulations in which agents performing strategy A were introduced (using the mutation process described above) into a population containing only strategy B and 10 simulations in which strategy B was introduced into a population containing only strategy A. In each simulation, the dominant strategy was introduced and run for 100 rounds without mutation so that agents could establish their behavioural repertoires. Mutation was then introduced, providing the second strategy with the opportunity to invade, and simulations were run for a further 10 000 rounds. The mean frequency of a strategy over the last 2500 simulation rounds was its score for that simulation. These scores were then averaged over the 20 simulations, and this average recorded as the contest score for that strategy in that contest. These simulations were run with the parameter set ($p_c = 0.01$, $p_{\text{copyActWrong}} = 0.05$, $\sigma_{\text{payoffError}} = 1$, $n_{\text{observe}} = 1$). A further set of simulations was run across a range of conditions using the 24 highest ranked strategies and the top 10 strategies were picked from among these. This top 10 set proceeded to the melee round.

In the melee simulation there were two sets of conditions, one systematic and one random. For the systematic condition set, we selected a number of values for each of the four parameters, p_c , n_{observe} , $p_{\text{copyActWrong}}$ and $\sigma_{\text{copyPayoffError}}$. Fifty simulations were run with each of the 280 possible combinations of these parameter values giving 14 000 simulations. To check that the results of this process were not unduly affected by the specific parameter values we chose, we also ran random conditions, where parameter values were chosen at random from statistical distributions. For biological plausibility these distributions were weighted towards lower values of p_c , n_{observe} and $p_{\text{copyActWrong}}$. We selected 1000 unique sets of parameter values and ran a single simulation with each set of values. Systematic and random analyses gave identical returns on the ranked performance of the 10 strategies, computed across all simulations and based on their frequency in the last quarter of each simulation.

(b) Results

The most striking outcome of the tournament was the success of strategies that relied heavily on copying when learning, in spite of the absence of a fixed additional cost to asocial learning. Copying paid under a surprisingly broad range of conditions, even when it was highly error prone, even when only a single individual was copied, and even when copying revealed no information about the payoff to the copied behaviour. The effectiveness of social learning was observed in spite of the absence of an inherent cost to asocial learning, and in the presence of costs to social learning. We also found that the presence of the copy error, $p_{\text{copyActWrong}}$, was an important source of new information for strategies relying almost exclusively on social learning. The results showed that when the value of $p_{\text{copyActWrong}}$ was set to 0, the positive effect of reliance on social

learning on mean individual fitness reversed, becoming strongly negative ($r = -0.30$, $p < 0.001$). Thus, when there is no copy error, high levels of social learning are associated with reduced average individual fitness in the population.

Our analyses suggested that copying pays because other individuals filter behaviour, making adaptive information available for others by only performing the highest-payoff behaviour in their repertoires. This means that even random copying is typically far more efficient than trial-and-error learning, because copiers select from a subset of the most effective actions. This helps us to understand why copying is widespread in animals [32].

However, to be successful in the tournament, strategies had to do far better than copy at random. Successful strategies restricted learning and maximized exploiting, timing bouts of copying for when payoffs drop, such that they acquired new knowledge that enhanced their performance after a change in the environment. The winning strategy, *discountmachine* (entered by Dan Cownden and Tim Lillicrap), evaluated current information based on its age, and judged how valuable it would be in the future, a form of mental time travel that greatly increased its learning efficiency. Copying only increased the mean individual fitness of individuals in the population when conducted in the efficient manner exhibited by the best performing strategies. This may help explain why human culture, but not that of animals, has led to demographic success, an observation that previous analytical theory has found difficult to explain [13]. Conceivably, only humans have the psychological attributes to be able to copy with this kind of efficiency.

3. THE DYNAMICS OF CULTURAL KNOWLEDGE AND BEHAVIOUR: FACTORS AFFECTING AMOUNT, EVENNESS AND PERSISTENCE

(a) *Methods*

The initial tournament analysis [8] focused primarily at the individual level, asking how and why individuals using certain strategies performed well. However, the tournament framework also allowed us to explore factors such as how copying affects the amount of cultural knowledge at the population level. Each simulation contains 100 individuals, and each of those individuals has, at any one time, a set of behaviour patterns in its repertoire (almost always a subset of the 100 possible behaviour patterns defined in the multi-armed bandit). These combined repertoires thus constitute the combined knowledge of that population. The population can also be characterized by the set of behaviour patterns it is performing, provided at least one individual is playing EXPLOIT at a given time. This distinction between things individuals know about and things individuals actually do, between knowledge and behaviour, is not often captured in theoretical studies of cultural evolution, and cultural evolution theory has been criticized precisely for this reason [33]. It is, however, a core feature of the tournament model that we can distinguish between behaviour and knowledge. Here, we investigate these two aspects of culture independently, running

simulations that kept track of three simple measures that together characterized the knowledge and behaviour present in these virtual cultures.

The first quantity we tracked was simply the number of different behaviour patterns present, which we expressed as a proportion of the 100 possible behaviour patterns defined by the multi-armed bandit. We measured both the proportion (out of 100) of possible behaviour patterns that were known—i.e. that were present in the repertoire of at least one individual—and the proportion that were actually performed in an EXPLOIT move by at least one individual. We call these proportions the *amount* of knowledge and behaviour, respectively; they measure the number of behaviour patterns known about, or performed, by a population at a given time.

Second, to describe what we term *evenness* in the context of this paper, we measured the flatness of the frequency distribution of behaviour patterns across the population using Pielou's evenness index [34]. This is a measure used in quantifying species evenness in ecological communities based on the Shannon–Wiener diversity index, and is given by

$$\mathcal{J} = \frac{-\sum_{i=1}^S p_i \ln p_i}{\ln(S)}, \quad (3.1)$$

where S is the number of species present in a sample, and p_i is the relative frequency of species i in the sample. In our case, we are using it as a measure of the distribution of behaviour patterns, so S is the number of possible behaviour patterns (100) and p_i is the number of individuals in the population at that iteration with that behaviour. The value of \mathcal{J} can range from 0 to 1, the latter representing maximum evenness. For example, consider a sample of 100 individuals, in an environment where five behaviour patterns were possible. Maximum evenness would be if 20 individuals each performed one behaviour ($p = [0.2 \ 0.2 \ 0.2 \ 0.2 \ 0.2]$, $\mathcal{J} = 1$), while minimum evenness would be if 100 individuals performed one behaviour and none performed the others ($p = [1 \ 0 \ 0 \ 0 \ 0]$, $\mathcal{J} = 0$). An uneven distribution of behaviour with most individuals choosing the same behaviour could come about through conformism, the strategy of preferentially copying the most commonly seen behaviour, the importance of which in human culture is an area of active debate [35,36]. However, we were interested to use our tournament model to explore the extent to which apparently conformist outcomes, such as an uneven distribution of behaviour, could arise in the absence of explicitly conformist strategies being deployed by individual agents [37].

Finally, we measured the rate of cultural turnover by calculating what we term the *persistence* of knowledge and expressed behaviour. Persistence was the average number of continuous iterations that behaviour patterns were known (as before, present in the repertoire of at least one individual) or for which they were performed (as before, being chosen in an EXPLOIT move by at least one individual), given that they had become known or expressed by at least one individual (such that persistences of zero did not occur).




Figure 1. Performance of entered strategies. Plot shows ranked overall strategy scores in the final stage of the tournament (*cWYTLWPD* = *copyWhenYoungThenLearnWhenPayoffsDrop* and *wTGGTGS* = *whenTheGoingGetsToughGetScrounging*). Scores are averaged over all final stage simulations. Inset plot shows mean scores for all 104 entered strategies in the pairwise round; error bars show standard deviation. The dotted black line on the inset indicates the 10 highest scoring strategies. Figure 1a adapted from Rendell et al. [1].

We concentrated our analysis on the top ten strategies (figure 1) as, being highly effective, we reasoned, these are the strategies most likely to occur at high frequency in nature and so it is their behaviour that is of greatest relevance here. We gathered data on amount, evenness and persistence by running two batches of tournament simulations. The first recreated the *random melee* section of the tournament, where the top ten strategies competed simultaneously across varied simulation parameters, to investigate how variation in learning strategies affects cultural dynamics. Varying simulation parameters reflect different assumptions about the environmental conditions. The cultural measures we analysed could change in response to this variation through both the way strategies themselves changed their behaviour in response to varying parameter values, and also the way in which the parameter values themselves altered the simulation dynamics by changing the error rates and relative cost associated with social learning. We were interested in the consequences of these strategic shifts on cultural dynamics at the population level. The second set of simulations took each strategy in turn and ran simulations with only that strategy present, using the same parameter values as the *pairwise* tournament phase. We then repeated this exercise twice, in each case making a major alteration to the model assumptions. First, we set the probability of copying the wrong behaviour when playing OBSERVE, $p_{\text{CopyActWrong}} = 0$, removing the assumption that copying errors could introduce new behaviour into a population. Secondly, we set the standard deviation of a normally distributed error applied to payoffs returned by OBSERVE, $\sigma_{\text{CopyPayoffError}} = 100$, making information about the payoffs of behaviour acquired by social learning too inaccurate to be useful. We then compared the characteristics of the virtual cultures constructed by the single strategies under these varied assumptions.

(b) Results

This was not a standard simulation study in which one changes a parameter to analyse how this parameter causally affects some measure in the results, rather it was a correlation study, where both dependent and independent variables were outcomes of variation in other variables. Several noteworthy results emerged. We first characterized the learning approach of the mixed strategy populations simply as the average proportion of learning in each round that was social (i.e. the number of individuals playing OBSERVE divided by the number of individuals playing either learning move—OBSERVE plus INNOVATE), and then examined how this single feature affected the population level measures described above (figure 2). In general, we found that there was a step change in the relationships we examined, that, for the parameter sets explored, occurred after the proportion of OBSERVE exceeded around 0.7.

High levels of social learning were associated with increased amounts of knowledge until the proportion of OBSERVE exceeded approximately 0.7, after which we observed a ceiling effect with populations generally knowing all possible behaviour patterns (figure 2a). This is a counterintuitive result because, compared to reliance on asocial learning, a reliance on social learning must reduce the amount of new information entering a population [38]. We interpret this apparent paradox as being due to two factors. It results in part from the assumption that copy error could introduce new behaviour into a population, hence more social learning leading to more copying errors leading to more knowledge, an interpretation which is supported by the results of our switching off copying errors, presented below. Secondly, more copying also means that behaviour patterns are more likely to be retained within the population, since multiple copies of any given variant are more likely to be generated. This can increase the amount of cultural knowledge in the population, because any knowledge that is built up over time is much less likely to be lost (figure 2e) thereby compensating for the reduced introduction of variants that would have occurred through innovation. However, for increases in copying to be associated with increases in the amount of culture there has to be a source of new variants present, either stemming from innovation (even at low levels) or copy-error.

Conversely, high levels of social learning were associated with reduced numbers of expressed behaviour patterns, especially when the proportion of OBSERVE exceeded approximately 0.7 (figure 2b). Accompanied by the concurrent reduction in the evenness of behaviour apparent in figure 2d, this observation suggests that in populations containing effective strategies and with high levels of social learning, individuals concentrate on performing relatively few behaviour patterns even though they have knowledge of many. Those few behaviour patterns performed are drawn from those with the best payoffs available at that time as many individuals choose to adopt them.

As these patterns in the amount of culture were counterintuitive, we were curious as to whether they could be explained either by simply the amount of learning that strategies did (quantified as the




Figure 2. How effective learning strategies affect culture at the population level. Plots show quantitative measures of the (a,b) amount, (c,d) evenness, and (e,f) longevity of knowledge (a,c,e), and behaviour (b,d,f), plotted against the proportion of learning that was social, in mixed-strategy populations running under the *random melee* tournament conditions. Each point represents the average value across all iterations for a single 10 000 iteration simulation run. Note the order of magnitude difference in y-axis scales between (e) and (f).

proportion of moves that were either OBSERVE or INNOVATE), or directly from the variation in simulation parameter values. We first fitted simple GLMs with identity link and normal error [39] to the data on the amount of knowledge, using as predictors the average proportion of learning in each simulation, which we term $p(\text{learn})$, and the values of the four simulation parameters (p_c , n_{observe} , $\sigma_{\text{copyPayoffError}}$ and $p_{\text{copyActWrong}}$). We used these predictors, including all first-order interactions, in a model selection analysis, testing all possible combinations to see which gave the best fit, as determined by the lowest AIC value. Once the best model had been identified, we then compared its AIC with that of an identical model but with the average proportion of learning that was social (which we term $p(\text{OBSERVE})$) as an additional predictor. The best model without $p(\text{OBSERVE})$ contained all predictors and interactions save $\sigma_{\text{copyPayoffError}} * p_{\text{copyActWrong}}$ and $\sigma_{\text{copyPayoffError}} * p(\text{learn})$, with an AIC of -7013.5 . The same model with $p(\text{OBSERVE})$ added as a predictor had an AIC of -7093.4 , some 79.9 units smaller, comfortably exceeding 3, the rule of thumb generally used to indicate a better fitting model [40]. This analysis shows that while other factors certainly had an effect, the proportion of learning that is social still

independently explains a significant amount of variation in the amount of knowledge.

High levels of social learning did not affect the evenness of knowledge (figure 2c), but drastically reduced the evenness of behaviour (figure 2d). Once the proportion of OBSERVE exceeded approximately 0.7 we saw populations where most individuals performed just one or two behaviour patterns.

High levels of social learning (especially greater than approx. 0.7) were sometimes associated with very large increases in the persistence, or longevity, of knowledge in these populations (figure 2e). Given that the average lifespan of individuals in these populations is 50 iterations, it is clear that increased levels of social learning can, in this model, lead to knowledge that far outlives its original innovators, lasting for tens and sometimes hundreds of generations. Conversely, in populations with high levels of social learning, persistence of behaviour was reduced, because effective strategies were quick to stop exploiting behaviour that did not return payoffs as high as expected, leading to a low average persistence (as the most effective strategies were the ones that did most social learning, populations with high levels of social learning are probably dominated by those strategies).




Figure 3. How the rates of environmental change and copying error affect the amount and evenness of expressed behaviour when effective strategies are present. Plots (a,b) show the proportion of the 100 possible behaviour patterns being used (i.e. by playing EXPLOIT) in mixed-strategy populations running under the *random melee* tournament conditions, as a function of (a) the rate of environmental change, and (b) the standard deviation of error in the estimated payoff of a behaviour learned socially (i.e. by OBSERVE). Plots (c,d) show the evenness of the distribution of behaviour patterns being used, in the same populations, as a function of (c) the rate of environmental change, and (d) the standard deviation of error in the estimated payoff of a behaviour learned socially. Each point represents the average value across all iterations for a single 10 000 iteration simulation run.

We found that only the evenness and amount of behaviour (not knowledge) were affected by varying the simulation parameters (figure 3), while persistence was not affected at all. We found that increased rates of environmental variation were associated with both greater evenness and more expressed behaviour patterns (figure 3a,c), even though the evenness and amount of knowledge were not affected. It is not inevitable that an increased amount of expressed behaviour is linked to increased evenness—it would be perfectly possible, for example, for a population to contain one very popular behaviour and a lot of relatively unpopular ones—but in these simulations they do appear to be linked, a result of individuals diverging in their estimates of what is currently the best behaviour to be deploying. Surprisingly, the extent to which social learning was error-prone appeared to have no effect on knowledge and behaviour at the population level (figure 3b,d), even when it reached extreme levels such as 0.5.

We also examined how two important assumptions in our tournament structure may have affected our results. These assumptions were first, that social learning carries a probability of learning the wrong behaviour (copy error) and thereby introduces new knowledge into a population, and second, that social learners are capable of learning the payoff associated with a given behaviour (payoff information). It is easy to envisage situations where these might not apply—for example, in the first case where strict error-correcting mechanisms are deployed to maintain

fidelity (reproduction of classical music, or ballet steps), and in the second case, when the behaviour is one with delayed or opaque payoffs (e.g. planting seeds). We ran simulations containing single strategies explicitly, considering the effect of no copy error and no payoff information (figure 4).

We found that these factors had strong effects on the amount of behaviour and knowledge. Both were reduced in the absence of copying errors but increased in the absence of useful payoff information (figure 4). They had either no effect (in the absence of copy error) or relatively little effect (in the absence of payoff information) on the evenness of either behaviour or knowledge, except that a lack of payoff information slightly increased the evenness of both knowledge and behaviour. Switching off copy error profoundly affected the persistence of behaviour, increasing it up to seven-fold for some strategies, but did not affect the persistence of knowledge. A lack of payoff information appeared to have little general effect on persistence, although in some strategies it appeared to increase the persistence of knowledge. These effects are generally intuitive. The switching off of copying errors reduces the supply of new information into a population when that population learns mostly by social learning, which is thought to be an important weakness of social learning in general [10,38], resulting in reduced amounts of knowledge and behaviour. Individuals in these populations were forced to rely heavily on that reduced set of behaviour patterns they did know about, resulting in an increased




Figure 4. The effect of information transmission parameters on (a) amount, evenness and (b) persistence of knowledge and behaviour (note differing x -axis scales between panels). Boxplots show the change from the standard conditions used in the tournament when (white bars) the probability of copying the wrong behaviour when playing OBSERVE, $p_{\text{copyActWrong}} = 0$ or when (grey bars) the social learning of payoffs is so error-prone as to be uninformative (the standard deviation of a normally distributed error applied to payoffs returned by OBSERVE, $\sigma_{\text{copyPayoffError}} = 100$). Values are taken from averages over ten 10 000 iteration simulations in each of which only one of the 10 best tournament entries was present.

persistence of behaviour in a population. Finally, increased evenness of behaviour in the absence of payoff information probably reflects a reduced ability by the strategies in question to identify the currently optimal behaviour.

4. DISCUSSION: WHAT DOES THE TOURNAMENT IMPLY FOR CULTURAL EVOLUTION?

Analysis of the social learning strategies tournament [8] revealed that social learning is a more effective means of knowledge gain than asocial learning under most biologically plausible conditions, even when it is highly error prone. Copying pays because the copied individual is usually rational in performing his or her most effective (i.e. highest-payoff) behaviour, such that only the subset of high-performance behaviour patterns are exploited, thereby inadvertently filtering behaviour for the copier. This is probably why copying is widespread in nature, since even random copying is more effective than trial-and-error learning, and effective copying can be performed in the absence of cognitively challenging learning rules. Social learning thus can have benefits that may have been previously underappreciated [41]. However, this does not detract from the utility of strategic copying. The most successful strategies were highly selective about when they copied, which they achieved by engaging in bouts of copying only when payoffs dropped significantly, and evaluating information based on its

age. It was only when copying was conducted in this efficient manner that it increased mean fitness within the population. This may explain why a reliance on social learning has led to demographic success in humans but not in other animals; conceivably, only we copy with sufficient efficiency and fidelity to give the demographic success observed in human history [10].

The analyses of population-level cultural characteristic produced a number of interesting insights. Noting that humans possess a great deal more cultural knowledge than other animals, we suggested that this might go hand-in-hand with our greater reliance on social learning compared to other animals. Our results were consistent with this hypothesis in establishing a positive relationship between the proportion of learning that was OBSERVE and the amount of cultural knowledge in the population (figure 2a). While, in principle, this relationship need not be a causal one, the aforementioned analyses, plus further simulations presented in the electronic supplementary material, lead us to the view that it probably is. Under some circumstances increased reliance on copying can increase the amount of culture present in the population because the transmission of knowledge between individuals allows it to outlive any single individual, thereby reducing the rate at which a population loses knowledge. However, it is also the case that the relationship between reliance on social learning and amount of knowledge is partly, but not solely, a product of selection between alternative strategies,

perhaps favouring strategies associated with large repertoire sizes. Learning strategies can also be characterized by the amount of learning they do, and our statistical analyses showed that this also affected the amount of culture present in a positive way, which is an intuitive result. The amount of learning was, however, a less powerful predictor of our data than the amount of social learning, which is perhaps less intuitive. High levels of copying were associated with complete knowledge saturation, such that all theoretically possible behaviour was represented in the repertoire of at least one individual in the population. This knowledge was retained in spite of the fact that enhanced copying led to a much smaller average proportion of behaviour actually being expressed at any point in time (figure 2*b*). In the electronic supplementary material we present the results of simulations that raise the number of possible behaviour patterns to 10 000, conducted to gain a deeper understanding of the analyses presented here. The objective of allowing such a large number of behaviour patterns was to prevent any ceiling effects associated with knowledge saturation. The simulations also utilized a single strategy to remove the possibility of selection among strategies. Even under such conditions, we find that an increased reliance on social learning can be associated with increases in the amount of cultural knowledge, by virtue of it allowing behaviour to persist for longer than it otherwise would.

This relationship is not monotonic, however, as we find a peak in the amount of culture at some probability of social learning less than unity. In the extreme, when there is no copy error and only social learning, the amount of knowledge collapses to virtually zero.

We think these results are best understood in conjunction with the striking effect of copying on the retention of knowledge within a population (figure 2*e*). A heavy reliance on social learning led to an average duration of knowledge within populations that was several orders of magnitude longer than for populations reliant on asocial learning. We observed a threshold level of copying, above which cultural knowledge could be retained for many hundreds of generations (given an average lifespan of 50 iterations). We suggest that it is this retention of knowledge that allows knowledge repertoires to approach saturation when populations rely heavily on social learning, because even with very low levels of innovation, if knowledge persists then it will, over time, accumulate. The caveat here is that there must be a source of new variation present, either through copying errors or occasional innovation. Once again, there is a marked contrast with the strong negative effect of copying on the persistence of behaviour (figure 2*f*). High levels of copying lead to rapid turnover in the behaviour patterns exploited, without a concurrent loss of knowledge from the population's repertoire. These observations have parallels in human populations, where we simultaneously witness fads and fashions that change rapidly, representing high turnover in behaviour, and the retention of cultural knowledge over millennia. By illustrating how copying enhances the persistence of knowledge, our analyses explain why these observations should not be viewed as conflicting.

These observations help to resolve the conundrum that, while logic dictates that copying promotes behavioural homogeneity, humans have accumulated large amounts of cultural knowledge. As expected, individuals converge, through copying, on high-payoff behaviour, such that increasing levels of copying reduce the proportion of behaviour exploited to the high-performance end of the spectrum. However, in a changing environment, with individuals born naive, inevitably there is some low-level exploitation of poor-performing behaviour too. It would seem that copying may sometimes promote the maintenance of cultural knowledge in a population, by preventing such knowledge from being lost when individuals die. This illustrates, once again, the striking adaptive advantages of social learning: individuals rapidly converge on, and reap the benefits of, exploiting high-payoff behaviour, yet at the same time high-copying populations retain high levels of cultural knowledge, conferring the plasticity to switch behaviour when environments change.

As expected, we also found that copying typically reduced the evenness of exploited behaviour, because it leads to a small number of high-performance behaviour patterns being disproportionately performed (figure 2*d*). Indeed, even in the absence of an explicit 'conformist learning' rule being deployed by individual agents, we witnessed some semblance of conformity emerging at the population level, reflected in lower evenness of the distribution of behaviour (figure 2*d*) a finding consistent with other analyses [42]. This did not, however, greatly affect the size of the cultural knowledge base (figure 2*a*), since, as we have described, it allowed more knowledge to be retained over a greater period. Our tournament simulations focused on a single focal population, and we envisage that, were multiple populations to be involved, the large cultural knowledge base promoted by copying would lead to extensive cultural diversity between populations.

We also considered how some other parameters in our analyses affected the amount, persistence and evenness of culture. One of the cleanest and most intuitive relationships was between the rate of environmental change and amount of culture, where greater rates of turnover in the environment led to a greater range of behaviour patterns being performed (figure 3*a*). High rates of change also resulted in more even distributions of behaviour, as no single high-performance behaviour was persistently optimal in the changing world (figure 3*c*). Unlike much previous theory, which has suggested that a reliance on social learning can sometimes hinder the adaptive tracking of temporally changing environments [13,14,43], a heavy reliance on social learning did not compromise the ability of agents in our tournament to adjust to changing environments. This probably reflects the greater biological realism of the tournament over analytical models, since only in the former do individuals possess a repertoire of behaviour. Possessing knowledge of multiple behaviour patterns allows individuals to switch rapidly to an alternative high-performing action when changes in the environment reduce the payoff to the current behaviour. This flexibility dramatically reduces the costs of copying.

Perhaps surprisingly, copy error had little effect on the amount of expressed behaviour (figure 3*b,d*), except at the extreme where there was no copy error at all (figure 4). Error-free copying reduced the amount of both cultural knowledge and behaviour, since new behaviour could no longer be introduced by this route. While it is no surprise that without copy error behaviour patterns persisted for longer (figure 4), it is less intuitive that this would lead to individual actions being exploited for longer (on average, twice as long). This reflects the fact that a great deal of copying (53% of all OBSERVE moves in the first tournament phase) failed to introduce new behaviour into individuals' repertoire, as individuals observed behaviour patterns that they already knew about. In the tournament, copy error increases the chances that individuals will acquire new behaviour when they play OBSERVE. These findings reinforce our view that copy error may be an important source of adaptive variation within natural populations [44].

One of the more surprising findings from the tournament was that copying paid even when copiers had no information about the payoff associated with the copied behaviour [8]. Like many others [30], we had assumed that one advantage to copying would be that it allowed individuals to home in on high-performance behaviour; and indeed, this assumption may yet be correct. However, copying offers advantages over trial-and-error learning even if observers receive no payoff information, because even blind copiers benefit from the aforementioned filtering of behaviour by the copied agents [10]. This insight could help explain the extreme reliance of children on imitation, leading them faithfully to copy even superfluous actions in a demonstrated task [45]. When children copy adults, they are typically taking advantage of decades of information filtering by the adult, making it on average simply more efficient to take their word for it. We also find that, in the absence of payoff information, greater amounts of cultural knowledge are retained and exploited (figure 4), since the potential to be selective in the acquisition and performance of behaviour is reduced by a lack of payoff information. This also leads to greater evenness in behaviour across the population.

The tournament has proved an effective means of exploring a number of questions and paradoxes concerning cultural evolution. By illustrating the striking utility of copying across such a broad range of conditions, and drawing attention to the adaptive filtering performed by individual agents, it helps to explain why social learning is widespread in nature. By isolating the factors that lead strategies to be successful, it allows us to make a series of predictions as to the patterns of strategic copying likely to be observed in nature (e.g. copying should increase when payoffs drop, but rapidly drop off once effective behaviour is found). A focus on the winning strategy leads us to the insight that mental time travel, combined with the ability to estimate rates of environmental change, may be a vital feature of human copying, since it allows individuals to assess the probable utility of current information in the future. Our analyses also help explain how a highly culturally

dependent species like humans might accumulate large amounts of cultural knowledge, when copying leads to behavioural homogeneity. Provided copying errors or innovation introduce new behavioural variants, copying can simultaneously increase the knowledge base of a population, and reduce the range of exploited behaviour to a core of high-performance variants. Similar reasoning accounts for the observation that copying can lead to knowledge being retained over long periods of time, yet trigger rapid turnover in behaviour. Low-level performance of sub-optimal behaviour is sufficient to retain large amounts of cultural knowledge in copying populations, over long periods. Indeed, a high level of copying is associated with the retention of cultural knowledge being increased by several orders of magnitude.

These observations suggest that copying confers an adaptive plasticity on cultural populations, allowing them to respond to changing environments rapidly by drawing on a deep knowledge base. In biological evolution the rate of change is positively related to genetic diversity [46], and formal analyses suggest a similar relationship between the rate of cultural evolution and the amount of cultural variation [10,18]. Accordingly, we might envisage that populations heavily reliant on culture would rapidly diverge behaviourally, exploiting the rich levels of variation retained in their knowledge base. Our tournament suggests that the ecological and demographic success of our species, our capacity for rapid change in behaviour, our cultural diversity, our expansive knowledge base, and the sheer volume of cultural knowledge we exhibit, may all be direct products of the heavy, but smart, reliance of our species on social learning.

We are very grateful to Kimmo Eriksson whose discussion and comments greatly improved our manuscript. The authors would like to acknowledge the use of the UK National Grid Service (www.grid-support.ac.uk) in carrying out this work. We thank all those who entered the tournament for contributing to its success. We are also very grateful to Robert Axelrod for providing advice and support with regard to the tournament design. This research was supported by the CULTAPTATION project (European Commission contract FP6-2004-NESTPATH-043 434).

REFERENCES

- 1 Emery, N. J. & Clayton, N. S. 2004 The mentality of crows: convergent evolution of intelligence in corvids and apes. *Science* **306**, 1903–1907. (doi:10.1126/science.1098410)
- 2 Sumpter, D. J. T. 2006 The principles of collective animal behaviour. *Phil. Trans. R. Soc. B* **361**, 5–22. (doi:10.1098/rstb.2005.1733)
- 3 Conradt, L. & List, C. 2009 Group decisions in humans and animals: a survey. *Phil. Trans. R. Soc. B* **364**, 719–742. (doi:10.1098/rstb.2008.0276)
- 4 Whiten, A. 2011 The scope of culture in chimpanzees, humans and ancestral apes. *Phil. Trans. R. Soc. B* **366**, 997–1007. (doi:10.1098/rstb.2010.0334)
- 5 Perry, S. 2011 Social traditions and social learning in capuchin monkeys (*Cebus*). *Phil. Trans. R. Soc. B* **366**, 988–996. (doi:10.1098/rstb.2010.0317)
- 6 Henrich, J. & Broesch, J. 2011 On the nature of cultural transmission networks: evidence from Fijian villages for adaptive learning biases. *Phil. Trans. R. Soc. B* **366**, 1139–1148. (doi:10.1098/rstb.2010.0323)

- 7 van Schaik, C. P. & Burkart, J. M. 2011 Social learning and evolution: the cultural intelligence hypothesis. *Phil. Trans. R. Soc. B* **366**, 1008–1016. (doi:10.1098/rstb.2010.0304)
- 8 Rendell, L. et al. 2010 Why copy others? Insights from the social learning strategies tournament. *Science* **328**, 208–213. (doi:10.1126/science.1184719)
- 9 Axelrod, R. 1980 Effective choice in the prisoner's dilemma. *J. Conflict Resolution* **24**, 3–25. (doi:10.1177/002200278002400301)
- 10 Boyd, R. & Richerson, P. J. 1985 *Culture and the evolutionary process*. Chicago, IL: Chicago University Press.
- 11 Whiten, A., Horner, V., Litchfield, C. A. & Marshall-Pescini, S. 2004 How do apes ape? *Learn. Behav.* **32**, 36–52.
- 12 Enquist, M. & Ghirlanda, S. 2007 Evolution of social learning does not explain the origin of human cumulative culture. *J. Theor. Biol.* **246**, 129–135. (doi:10.1016/j.jtbi.2006.12.022)
- 13 Rogers, A. 1988 Does biology constrain culture? *Am. Anthropol.* **90**, 813–819. (doi:10.1525/aa.1988.90.4.02a00030)
- 14 Feldman, M. W., Aoki, K. & Kumm, J. 1996 Individual versus social learning: evolutionary analysis in a fluctuating environment. *Anthropol. Sci.* **104**, 209–231.
- 15 Enquist, M., Eriksson, K. & Ghirlanda, S. 2007 Critical social learning: a solution to Rogers' paradox of non-adaptive culture. *Am. Anthropol.* **109**, 727–734. (doi:10.1525/aa.2007.109.4.727)
- 16 Giraldeau, L.-A., Valone, T. J. & Templeton, J. J. 2002 Potential disadvantages of using socially acquired information. *Phil. Trans. R. Soc. Lond. B* **357**, 1559–1566. (doi:10.1098/rstb.2002.1065)
- 17 Laland, K. N. 2004 Social learning strategies. *Learn. Behav.* **32**, 4–14.
- 18 Cavalli-Sforza, L. L. & Feldman, M. W. 1981 *Cultural transmission and evolution: a quantitative approach*. Princeton, NJ: Princeton University Press.
- 19 Henrich, J. & McElreath, R. 2003 The evolution of cultural evolution. *Evol. Anthropol.* **12**, 123–135. (doi:10.1002/evan.10110)
- 20 Borenstein, E., Feldman, M. W. & Aoki, K. 2008 Evolution of learning in fluctuating environments: when selection favors both social and exploratory individual learning. *Evolution* **62**, 586–602. (doi:10.1111/j.1558-5646.2007.00313.x)
- 21 Wakano, J. Y., Aoki, K. & Feldman, M. W. 2004 Evolution of social learning: a mathematical analysis. *Theor. Popul. Biol.* **66**, 249–258. (doi:10.1016/j.tpb.2004.06.005)
- 22 Galef Jr, B. G. 2009 Strategies for social learning: testing predictions from formal theory. *Adv. Study Behav.* **39**, 117–151. (doi:10.1016/S0065-3454(09)39004-X)
- 23 Pike, T. W. & Laland, K. N. 2010 Conformist learning in nine-spined sticklebacks' foraging decisions. *Biol. Lett.* **6**, 466–468. (doi:10.1098/rsbl.2009.1014)
- 24 Pike, T. W., Kendal, J. R., Rendell, L. E. & Laland, K. N. 2010 Learning by proportional observation in a species of fish. *Behav. Ecol.* **21**, 570–575. (doi:10.1093/beheco/arq025)
- 25 Tomasello, M. 1994 The question of chimpanzee culture. In *Chimpanzee cultures* (eds R. W. Wrangham, W. C. McGrew, F. B. M. de Waal & P. G. Heltne), pp. 301–317. Cambridge, MA: Harvard University Press.
- 26 Enquist, M., Strimling, P., Eriksson, K., Laland, K. & Sjostrand, J. 2010 One cultural parent makes no culture. *Anim. Behav.* **79**, 1353–1362. (doi:10.1016/j.anbehav.2010.03.009)
- 27 Laland, K. N., Richerson, P. J. & Boyd, R. 1993 Animal social learning: towards a new theoretical approach. In *Perspectives in ethology*, vol. 10, *Behavior and evolution* (eds P. P. G. Bateson, P. H. Klopfer & N. S. Thompson), pp. 249–277. New York, NY: Plenum Press.
- 28 Currie, T. E. & Mace, R. 2011 Mode and tempo in the evolution of socio-political organization: reconciling 'Darwinian' and 'Spencerian' evolutionary approaches in anthropology. *Phil. Trans. R. Soc. B* **366**, 1108–1117. (doi:10.1098/rstb.2010.0318).
- 29 Gray, R. D., Atkinson, Q. D. & Greenhill, S. J. 2011 Language evolution and human history: what a difference a date makes. *Phil. Trans. R. Soc. B* **366**, 1090–1100. (doi:10.1098/rstb.2010.0378).
- 30 Schlag, K. H. 1998 Why imitate, and if so, how? *J. Econ. Theory* **78**, 130–156. (doi:10.1006/jeth.1997.2347)
- 31 Gross, R., Houston, A. I., Collins, E. J., McNamara, J. M., Dechaume-Moncharmont, F. X. & Franks, N. R. 2008 Simple learning rules to cope with changing environments. *J. R. Soc. Interface* **5**, 1193–1202. (doi:10.1098/rsif.2007.1348)
- 32 Leadbeater, E. & Chittka, L. 2007 Social learning in insects—from miniature brains to consensus building. *Curr. Biol.* **17**, R703–R713. (doi:10.1016/j.cub.2007.06.012)
- 33 Cronk, L. 1995 Commentary on Laland KN, Kumm J & Feldman MW Gene-culture coevolutionary theory: a test case. *Curr. Anthropol.* **36**, 147–148.
- 34 Smith, B. & Wilson, J. B. 1996 A consumer's guide to evenness indices. *Oikos* **76**, 70–82. (doi:10.2307/3545749)
- 35 Eriksson, K., Enquist, M. & Ghirlanda, S. 2007 Critical points in current theory of conformist social learning. *J. Evol. Psychol.* **5**, 67–87. (doi:10.1556/JEP.2007.1009)
- 36 Henrich, J. & Boyd, R. 1998 The evolution of conformist transmission and the emergence of between group differences. *Evol. Hum. Behav.* **19**, 215–241. (doi:10.1016/S1090-5138(98)00018-X)
- 37 Franz, M. & Matthews, L. J. 2010 Social enhancement can create adaptive, arbitrary and maladaptive cultural traditions. *Proc. R. Soc. B* **277**. (doi:10.1098/rspb.2010.0705)
- 38 Rendell, L., Fogarty, L. & Laland, K. N. 2010 Roger's paradox recast and resolved: population structure and the evolution of social learning strategies. *Evolution* **64**, 534–548. (doi:10.1111/j.1558-5646.2009.00817.x)
- 39 Faraway, J. H. 2006 *Extending the linear model with r: generalized linear, mixed effects and nonparametric regression models*. Boca Raton, FL: Chapman & Hall/CRC.
- 40 Burnham, K. P. & Anderson, D. R. 2002 *Model selection and multimodel inference: a practical information-theoretic approach*, 2nd edn. New York, NY: Springer.
- 41 Rieucou, G. & Giraldeau, L.-A. 2011 Exploring the costs and benefits of social information use: an appraisal of current experimental evidence. *Phil. Trans. R. Soc. B* **366**, 949–957. (doi:10.1098/rstb.2010.0325)
- 42 Matthews, L. J., Paukner, A. & Suomi, S. J. 2010 Can traditions emerge from the interaction of stimulus enhancement and reinforcement learning? An experimental model. *Am. Anthropol.* **112**, 257–269. (doi:10.1111/j.1548-1433.2010.01224.x)
- 43 Boyd, R. & Richerson, P. J. 1988 An evolutionary model of social learning: the effects of spatial and temporal variation. In *Social learning: psychological and biological perspectives* (eds T. Zentall Jr & B. G. Galef Jr), pp. 29–48. Hillsdale, NJ: Lawrence Erlbaum.
- 44 Henrich, J. & Boyd, R. 2002 On modeling cognition and culture: why cultural evolution does not require replication of representations. *J. Cogn. Cult.* **2**, 87–112. (doi:10.1163/156853702320281836)
- 45 Lyons, D. E., Damrosch, D. H., Lin, J. K., Macris, D. M. & Keil, F. C. 2011 The scope and limits of overimitation in the transmission of artefact culture. *Proc. R. Soc. B* **366**, 1158–1167. (doi:10.1098/rspb.2010.0335)
- 46 Fisher, R. A. 1930 *The genetical theory of natural selection*. Oxford, UK: Clarendon Press.