Experimental studies of modern human social and individual learning in an archaeological context: People behave adaptively, but within limits

Alex Mesoudi

Department of Anthropology *and*
Centre for the Coevolution of Biology and Culture
Durham University
Dawson Building, South Road
Durham, DH1 3LE
United Kingdom

Current affiliation:
Human Behaviour and Cultural Evolution Group
Biosciences, College of Life and Environmental Sciences
University of Exeter Cornwall Campus
Penryn, Cornwall TR10 9FE
United Kingdom

Email: a.mesoudi@exeter.ac.uk

Word count: 8233

Abstract

It has been proposed that one reason for the success of *Homo sapiens* is our advanced learning abilities. Theoretical models suggest that complex cultural adaptations can arise from an optimal mix of (i) individual learning that is of sufficient accuracy plus (ii) social learning that is of sufficiently high fidelity and is payoff-biased. Here I review the findings of a series of experimental studies of human learning, designed to simulate the kind of technology-based tasks that our ancestors would have faced. Results of these studies support the predictions of the models, and show that contemporary humans' learning strategies are broadly adaptive. Performance typically improved through effective individual learning and payoff-biased social learning. The latter crucially allowed participants to escape low-fitness locally optimal artifact designs and jump to higher-fitness designs, assuming a realistic multimodal adaptive landscape underlying artifact fitness. On the other hand, people also exhibited predictable flaws in their learning, such as the copying of neutral traits exhibited by successful models along with their functional traits (i.e. cultural hitchhiking), and an unwillingness to share information with others under certain circumstances.

Keywords: cultural evolution; cumulative culture; individual learning; social learning

1. Introduction

In a relatively brief period of evolutionary time, our species has successfully colonised and inhabited virtually every terrestrial environment on the planet, from the driest deserts to frozen tundra, from high-altitude mountain ranges to remote island chains, such that we now account for about eight times the biomass of all other wild terrestrial vertebrates combined (Hill et al. 2009). Other hominin species such as the Neanderthals have gone extinct, possibly due in part to the success of *Homo sapiens*, while our closest living relative species, chimpanzees, are limited to a few small, scattered populations across Africa. What accounts for the extraordinary evolutionary success of our species?

One possibility, proposed by the Replacement of Neanderthals by Modern Humans (RNMH) Project (Akazawa 2012), is that anatomically modern *Homo sapiens* possessed superior learning abilities compared to their fellow hominins and other primates. This hypothesis has its roots in theoretical modelling work in the field of cultural evolution going back several decades, which has linked evolutionary rates of change and phenotypic adaptation to learning strategies (Aoki et al. 2011; Aoki et al. 2005; Boyd and Richerson 1985; Cavalli-Sforza and Feldman 1981; Rogers 1988). A primary focus of these models has been the interplay between individual (or asocial) learning, in which novel solutions to problems are invented by a single individual, and social learning (or cultural transmission), in which solutions are copied from one or more other individuals in the population. The latter can take on different modes, such as vertical transmission from one's biological parents (Cavalli-Sforza and Feldman 1981), conformist transmission of the most popular solution in one's group (Henrich and Boyd 1998), or payoff/prestige biased transmission in which the most successful/prestigious individual in one's group is preferentially copied (Boyd and Richerson 1985; Henrich and Gil White 2001).

Although the results of these models are varied, a general finding seems to be that some mix of individual and social learning is adaptive in fluctuating environments that change too rapidly for innate, genetic responses to evolve, yet not so rapid that previous generations' solutions to problems are out-of-date (Aoki et al. 2005; Boyd and Richerson 1988). Moreover, if individual learning is sufficiently accurate, and social learning is of sufficiently high fidelity and is payoff-biased such that adaptive solutions are preferentially copied, then this mix of social and individual learning can result in cumulative cultural evolution (Aoki et al. 2012; Ehn and Laland 2012; Enquist et al. 2008; Mesoudi 2011b; Powell et al. 2009). Just as cumulative genetic evolution can result in complex genetic adaptations such as eyes or wings, cumulative cultural evolution can similarly generate complex cultural adaptations that most likely underlie our species' success, from bow-and-arrows, kayaks and celestial navigation to agriculture, airplanes and quantum physics (Richerson and Boyd 2005).

Did anatomically modern humans uniquely possess an optimal mix of sufficiently accurate individual learning plus sufficiently high fidelity, payoff-biased social learning? Was one of these ingredients missing in other hominin, or other primate, species? It is, of course, extremely difficult to infer the learning abilities of extinct hominin species from the incomplete and often ambiguous artifactual record. We can, however, test these predictions in contemporary humans. If groups of people solve problems in the way predicted by the aforementioned theoretical models, then we can be more confident in the validity of those models, and more confident in asserting that our species' learning capacities are evolutionarily adaptive. Just as importantly, if people do *not* behave as predicted (e.g. if they eschew payoff-biased social learning in favour of, say, conformist or random copying), then this requires modification of the assumptions of the models and/or modification of the original hypothesis that modern humans possess adaptive learning capacities.

With this aim in mind, in this paper I will review the results of a series of experimental studies conducted by myself and collaborators that have probed the learning abilities of contemporary humans when faced with a novel and complex task – what we have dubbed the *Virtual Arrowhead Task* – that is designed to resemble technology found in the material record. Hopefully, the findings of these experiments can inform both theoretical models of the evolution of human learning capacities, and interpretation of the often ambiguous archaeological record. This is not to say that experimental simulations are a perfect tool: far from it. While they offer many advantages, such as the ability to control extraneous conditions, manipulate

variables, replicate findings and generate complete behavioural datasets, they are limited by their lack of external validity, such as their short time spans, lower incentives, restricted social interaction and the assumption that the behaviour of contemporary humans can be extrapolated to that of past people. I therefore conclude with an extended discussion of the limitations and real-life applications of experimental methods in this context.

2. The Virtual Arrowhead Task

The Virtual Arrowhead Task was originally designed by myself and archaeologist Michael O'Brien to capture the key aspects of North American projectile points (Mesoudi and O'Brien 2008a, b), although we have since used it to explore the learning of complex technology in general (Atkisson et al. 2012; Mesoudi 2008, 2011a). One limitation of many of the theoretical models of cultural evolution discussed above, as well as some experimental tests of such models (e.g. McElreath et al. 2005), is that the 'task' or 'problem' that must be solved is unrealistically simple: often it is assumed that individuals can exhibit just one of two possible discrete traits, with one of those traits giving a higher payoff than the other trait as specified by the state of the environment. Even the simplest of human technology, however, comprises multiple component traits, some of which might be continuous (e.g. the length or width of a handaxe: Lycett and von Cramon-Taubadel 2008), others discrete but with more than two states (e.g. arc-shaped vs. curved vs. triangular base shapes of projectile points: O'Brien et al. 2001); some might be functional (e.g. the thickness or length of arrowheads: Cheshier and Kelly 2006) and some might be functionless (e.g. decorative patterns on canoes: Rogers and Ehrlich 2008). The overall 'cultural fitness' of an artifact will be a combination of these component trait values, each of which interacts with one another, as well as with the skill of the manufacturer/user, and stochastic factors such as weather conditions.

We therefore sought to design a task that was simultaneously complex enough to give us insights about how people solve real-life technology-based problems, and simple enough to be able to inform the theoretical models described above and yield tractable findings. In our task (see Mesoudi and O'Brien 2008a for a full description), participants in small groups of 5-6 each design an arrowhead via a computer program (**Figure 1**). This virtual arrowhead is composed of three continuous traits (Height, Width and Thickness), which can each take any value from 1-100 arbitrary units, and two discrete traits (Shape and Colour), which can each take one of four categorical values. Over a series of trials (or 'hunts'), participants can improve their arrowhead by either individual trial-and-error learning, by directly altering the values of one or more of the traits, or social learning, by copying the design of another group member. The form of this social learning (e.g. payoff bias, conformity) can be manipulated.

On each hunt the participant tests their arrowhead in a virtual hunting environment, receiving a score in calories out of 1000. The closer their design is to one or more hidden optimal designs pre-specified by us using fitness functions, the higher the score ('fitness' is used here to refer to cultural fitness of an artifact, which may, or may not, correspond to the biological fitness of the individual using that artifact). The overall fitness of the arrowhead is given by the sum of the separate fitness functions for the constituent traits (**Figure 2**). The discrete trait Shape has a step fitness function, with the four shapes randomly assigned either 100%, 90%, 66% or 33% of the maximum possible fitness from that trait. Colour is neutral and does not contribute to fitness in any way. The continuous traits (Height, Width and Thickness) each have bimodal fitness functions. For each, one randomly chosen value gives 100% of the fitness contribution (the global optimum), and another random value gives 66% of that maximum (the local optimum).

When added together, these bimodal functions generate a multimodal adaptive landscape (Wright 1932), where each coordinate represents a different arrowhead design and the height of the landscape represents the fitness of that design. With three bimodal traits there are $2^3=8$ peaks in our adaptive landscape, with each peak varying in its maximum payoff. For example, an arrowhead with Height, Width and Thickness all at their globally optimal values gives the full 1000 calories; an arrowhead with Height and Width at their global optima and Thickness at its local optimum gives a slightly lower maximum payoff; an arrowhead with Height, Width and Thickness all at their local optima gives the lowest maximum payoff. Given that most

real-life problems can typically be solved in multiple ways, with some solutions better than others, this is likely to be representative of real-life technological fitness (Boyd and Richerson 1992). Note, however, that participants were told nothing about these fitness functions (just as, presumably, real-life hunter-gathers have no a priori knowledge of the effectiveness of most of the technology they use). Finally, there is always a small random error in the score, simulating stochastic conditions such as weather or prey availability.

After each hunt, participants are informed of their score out of 1000 calories. Participants go through three seasons of hunting, with each season comprising 30 hunts. Optimal values change between seasons, but not during seasons, and participants are informed about both of these facts. During each season the participant can see their cumulative score (the sum of the scores on every hunt up to that point), and in group conditions their relative rank compared to other group members' cumulative scores. Motivational reward has varied across the studies described below: in some studies participants were rewarded monetarily based on their absolute score, in others based on their relative rank, and in others no monetary reward is given at all (interestingly, no obvious differences have been observed across these different motivational regimes).

This task is intended to capture the key aspects of most complex technology, including that used by both modern humans and Neanderthals around the time of their coexistence: a technology composed of multiple constituent traits (some continuous and some discrete, some functional and some neutral), that is cognitively opaque (there is no obvious, intuitive relation between an artifact and its effectiveness: Gergely and Csibra 2006) and which has multiple locally optimal alternative designs (i.e. a multimodal adaptive landscape). In a series of studies we have explored how contemporary humans engage with this task, with the following key findings.

3. Key findings

3.1 People are effective individual learners, but can get stuck on local optima

While much theoretical modelling work has looked at a diverse range of social learning strategies (Laland 2004), individual learning is often under-theorised in models, where it is often assumed that individuals come up with the correct solution to a problem with some fixed probability. We were interested in opening this 'black box' and exploring the strategies that people use when engaging in individual learning.

When playing alone, participants on average show effective individual learning. **Figure 3** shows that mean score increases over successive hunts, plateauing to a level significantly higher than that of the starting (random) design. Analyses of these data revealed that participants appear to engage in a simple but effective reinforcement learning, or 'win-stay-lose-shift', strategy (Mesoudi and O'Brien 2008a, b): pick a trait at random (e.g. Width), modify the trait (e.g. increase Width), if the payoff increases then keep modifying the trait in that way (e.g. increase Width further); if the payoff decreases then do the opposite (e.g. decrease Width). This is repeated until the payoff no longer changes, at which point the whole process is repeated for the next trait. In terms of the multimodal adaptive landscape, this simple hill-climbing algorithm results in the participant converging on the nearest peak in the landscape.

Formally, we can define two parameters in this strategy: d , which we defined as the number of traits that a participant changed on a single hunt ($0 \leq d \leq 5$), and c , the amount by which a continuous trait is modified during one hunt ($0 \leq c \leq 99$). If more than one continuous trait was changed in a hunt then c represents the mean of these traits, and we focus on the continuous traits because these are responsible for most of the improvement and variation in payoffs. Empirically, our participants typically had a d of 1 (mean=1.43, median=1, mode=1) and a c of 5 (mean=9.50, median=5, mode=5), meaning that on each hunt they changed one trait by 5 units. To test our hypothesised individual learning strategy, an agent-based model was constructed that followed the rules specified above with $d=1$ and $c=5$ (Mesoudi and O'Brien 2008b). As shown in **Figure 3**, the simulated values match well with the actual data from participants, reaching virtually identical end points and showing a similar gradual increase then plateau.

Interestingly, the participants do best relative to the simulation during early hunts (see hunts 4, 5 and 6 in

Figure 3). Further analyses showed that this is because c was not, in fact, constant across all hunts, as suggested by the slightly higher mean of 9.50. During earlier hunts, when participants generally had low scores, they responded by increasing c , i.e. making larger modifications to their arrowheads. Consequently, score was negatively and significantly correlated with c ($r_s = -0.368$, $p < 0.01$). In a multimodal adaptive landscape this is an adaptive individual learning strategy: if your score is low, you are most likely to be in a low-fitness valley, and large modifications may well transport you to a higher-fitness part of the landscape. If your score is high, then modifications should be small, otherwise you may move off your peak and into a valley.

Note also from **Figure 3** that the maximum mean score at hunt 30 of around 750-800 calories, which appears to have levelled off at a kind of equilibrium, falls quite short of the maximum possible 1000 calories. This, again, is because of the multimodal adaptive landscape. The individual learning strategy followed by our participants, and simulated in the model, leads participants uphill from a random starting point to the top of the nearest peak. This might be the globally optimal peak, but equally might be one of the seven other locally optimal but globally suboptimal peaks. So even though participants saw that their score was less than the maximum of 1000, the majority chose to stick with their pretty-good-but-not-perfect design, what Simon (1956) called 'satisficing'. This represents a disadvantage of pure individual learning in a multimodal adaptive landscape: independent individual learners can get stuck on locally optimal, but globally suboptimal peaks.

Note also that this individual learning strategy was employed for a range of randomly generated optimal arrowhead designs, with these random optima changing between seasons and across studies. Participants did not exhibit any intuitive notion of what an effective arrowhead design looked like, or if they did, it was (i) different for each participant given that they each started at different points in the landscape (see Mesoudi & O'Brien 2008b), and (ii) quickly overridden when the a priori intuitively good arrowhead design was found to perform poorly in the experiment. In this case, then, general-purpose learning rules override any pre-existing intuitive content biases or cultural attractors (Sperber 1996) regarding projectile point characteristics (at least in our non-expert participants; we explicitly excluded archaeology students and amateur replica-arrowhead-makers from the studies to avoid too specialised knowledge).

3.2 People use payoff-biased social learning to jump to higher-fitness designs

We can now ask how social learning, and in particular payoff-biased social learning, changes participants' performance on the task. Payoff-biased social learning was implemented by allowing participants to view the arrowhead design of another member of their group, given information about those group members' cumulative scores up to that point. When this is allowed, either after a long period of individual learning (Mesoudi and O'Brien 2008a) or concurrently with individual learning (Mesoudi 2008), participants readily engage in payoff-biased social learning, copying the design of the most successful person in their group rather than copying a random group member or continuing with individual learning. The result of payoff-biased social learning is a significant jump in the mean score relative to individual learners (**Figure 4**).

Payoff-biased social learning is adaptive here because it allows participants to abandon their locally optimal designs and jump, almost instantaneously, to the globally optimal peak, or at least the highest peak found by anyone in the group. Payoff-biased social learning has this effect almost by definition, because participants who have found higher peaks will have higher scores, and they are preferentially copied. To confirm that the multimodal shape of the adaptive landscape was responsible for the advantage of social learning, it was shown that (i) there were significantly more participants with designs at or near a locally optimal peak immediately before social learning is allowed than after, and conversely, significantly more participants at globally optimal peaks after social learning than before (Mesoudi and O'Brien 2008a), and (ii) when the adaptive landscape was made unimodal (by removing the local optima from the fitness functions for Height, Width and Thickness shown in **Figure 2**, to create a single globally optimal design/peak), the advantage of social learning disappeared, and individual learners achieved mean scores identical to multimodal social learners (Mesoudi 2008).

Moreover, just as a participant's individual learning strategy changed in response to the participant's score, so too did their social learning. The lower a participants' score, the more use they made of social information (Mesoudi 2008). This was indicated by a significant and negative correlation ($r=-0.29$, $p<0.001$) between participants' scores and a measure of social influence, defined as the amount by which a participant changed their existing arrowhead to make it more similar to the arrowhead of the participant who they had chosen to view.

This performance-dependent payoff-biased social learning, or "copy-successful-individuals-when-behaviour-is-unproductive" (Laland 2004), is again adaptive. Boyd and Richerson (1995) showed that this flexible and selective learning strategy of engaging in social learning only when individual learning is particularly costly or difficult is one way of solving 'Rogers' paradox' (Rogers 1988). Rogers suggested that social learners can be seen as 'information scroungers' free-riding on the costly efforts of individual learners (or 'information producers'), with a net result that a mixed population of social and individual learners will never have a higher mean fitness than a population solely comprised of individual learners. Boyd and Richerson (1995) showed that making learners selective, engaging in social learning only when individual learning is costly or difficult, removes this problem, allowing social learning to evolve and mean fitness to increase. That our participants behave in this way, only engaging in social learning when their scores are low (which we can infer is because they are finding individual learning difficult), is encouraging. It is also encouraging that other studies using different tasks have found similar effects, such as Morgan et al.'s (2012) finding that the lower a participant's confidence in their performance, the more they rely on social learning.

3.3 Payoff-biased social learning is preferred to other forms of social learning

The Virtual Arrowhead studies discussed so far compared individual learning with payoff-biased social learning, with alternative social learning strategies difficult or impractical for participants to use. In a recent study (Mesoudi 2011a), participants were given the option to engage in three additional strategies: random copying (copying the arrowhead of a randomly-chosen fellow group member), conformity (in which continuous traits were divided into 10-unit intervals, i.e. 1-10, 11-20, 21-30..., and the conforming participant is assigned the mid-value of the most popular interval in their group) and averaging (in which participants were assigned the arithmetic mean of everyone in the groups' values for each trait, similar to Boyd & Richerson's [1985] blending inheritance), along with payoff bias (copying the arrowhead of the highest-scoring group member) and individual learning (directly changing the traits with no social influence) as before.

Payoff-biased social learning was the clear favourite compared to the other social learning strategies. Across all hunts played by all participants, 78% involved individual learning, 19% payoff-biased social learning, and only around 1% each of conformity, random copying and averaging (Mesoudi 2011a). Again, this choice of social learning strategy is adaptive in the multimodal adaptive landscape implemented here. As shown using agent-based models simulating each of these strategies (Mesoudi and O'Brien 2008b), only payoff-biased social learning outperforms individual learning, due to the aforementioned reason that individual learners stuck on locally optimal peaks can jump to higher-fitness peaks found by more successful group members. Random copying also allows participants to jump peaks, but to a random, not necessarily high, peak. Conformity allows participants to jump to the most popular peak, but again there is no reason that this most-popular peak is the highest (unless payoff-bias has already acted). Averaging is particularly bad, as the mean trait value of several peaks is likely to be mid-way between all of them, i.e. in a valley.

3.4 Payoff biased social learning leads to "cultural hitchhiking"

In sections 3.2 and 3.3 we saw how payoff-biased social learning is adaptive, allowing participants to jump to high-fitness peaks in the multimodal adaptive landscape. Yet this social learning strategy also comes with disadvantages. As noted above, the Colour trait was neutral and had no effect on the score shown to participants. Despite this, Colour was copied by our participants just as faithfully as the other functional traits during payoff-biased social learning (Mesoudi and O'Brien 2008a).

This was measured by calculating pair-wise inter-trait correlations across all participants in a group, i.e. the correlation between all participants' Height and Width, the correlation between all participants' Height and Thickness, and so on. Following periods of individual learning, these inter-trait correlations were found to be quite low, around $r=0.1-0.3$. This is to be expected, as different individual learners would diverge to different peaks in the adaptive landscape, thus reducing between-participant similarity. Once payoff-biased social learning was permitted, however, the inter-trait correlations increased significantly to around $r=0.3-0.9$. This is because all participants in the group copied the single most-successful group member (apart from that most-successful participant him or herself, of course, who could not copy themselves), thus all participants ended up with extremely similar arrowheads. Colour showed the same pattern of inter-trait correlations as the other traits, indicating that Colour was copied along with the other functional traits in a complete package.

This is an example of a neutral trait hitchhiking on functional traits, and represents the downside of payoff-biased social learning: while copying a successful individual will on average lead to the acquisition of adaptive behaviour, occasional neutral or even maladaptive traits might also be copied. This hitchhiking was explored formally by Boyd and Richerson (1985) as 'indirect' bias, which encompasses payoff-biased social learning, in which successful people are preferentially copied, and prestige bias, in which people with high social status are preferentially copied (which may, or may not, correspond to objective measures of success in tasks such as hunting). If this 'cultural hitchhiking' is an intrinsic side-effect of payoff- or prestige-biased social learning then we might expect neutral traits to be common in the archaeological record, and not necessarily seek functional explanations for all observed traits (see also Bentley et al. 2004; Dunnell 1978; Neiman 1995).

A more recent study illustrates further the power of prestige bias. Henrich and Gil-White (2001) suggested that people often identify from whom to copy based on quite minimal and subtle cues of prestige, such as looking times. Highly prestigious individuals should be looked at more by others than less prestigious individuals because they are good sources of information, and so looking times might constitute a cheap and quick cue regarding who to copy. Atkisson et al. (2012) tested this prediction in the Virtual Arrowhead Task, presenting participants with objective success information – the scores of other group members – as before, but also fictional looking time information concerning how long each group member had chosen to view every other group members' arrowheads. Even though the looking times were fictional and therefore useless, this marker of prestige was used at least as much as the objective success information when participants were choosing from whom to copy. So again, while this prestige-biased social learning may be broadly adaptive, it can easily misfire.

Further studies might look more systematically at the conditions under which we would expect neutral or maladaptive to hitchhike via generally prestigious models. We might predict hitchhiking to be particularly prevalent when it is difficult to directly assess the efficacy of different traits. In the arrowhead task, the constant random error in feedback likely obscured the fact that Colour had no systematic effect on payoffs; future studies might vary the size of this feedback error to determine whether hitchhiking disappears below some error threshold. Whether maladaptive traits hitchhike, meanwhile, is likely dependent on their cost relative to the fitness benefits of the adaptive traits exhibited by prestigious demonstrators. One might predict that highly maladaptive traits would not spread beyond an initial accidental copying event after which their negative effects are detected (although cases such as kuru [Durham 1991] or celebrity-driven copycat suicides [Mesoudi 2009] might suggest otherwise).

3.5 Informational access costs block social learning

In the Virtual Arrowhead experiments discussed so far, participants could freely view other participants' arrowhead designs. This is unlikely to hold true for all real-life situations, however. Henrich and Gil-White (2001) suggested that even though many hunter-gatherer societies are relatively egalitarian, highly skilled individuals will often receive material benefits (e.g. food) or non-material benefits (e.g. status) from letting others watch them engage in their skilled activity. Stout (2002) found that knowledge of stone tool production in adze makers of Indonesian Irian Jaya was carefully protected through the use of highly selective apprenticeships. Similarly, in industrialised societies, it is commonplace for highly skilled or

knowledgeable people, from car mechanics to lawyers, to set prices for access to their skills or knowledge. Moreover, the level of skill and knowledge often covaries with their price: more knowledgeable lawyers set higher fees than less knowledgeable lawyers, for example. These prices can be seen as 'informational access costs', which potential social learners must pay in order to access social information.

In one study, I therefore added informational access costs to the Virtual Arrowhead Task (Mesoudi 2008). Each participant could set their own access cost, in terms of calories, that other group members had to pay in order to view their arrowhead design. These costs were added and subtracted to the participants' actual cumulative scores. For example, if Participant 1 set an access cost of 450 calories and Participant 2 chose to copy Participant 1, then 450 calories would be deducted from the cumulative score of Participant 2 and 450 calories would be added to the cumulative score of Participant 1.

As expected, participants with higher scores set higher informational access costs than participants with lower scores. Participants were clearly aware, then, that their fellow participants will engage in payoff-biased social learning and preferentially copy the highest-scoring participant, such that their information would be in the highest demand and therefore be most valuable. But unexpectedly, rather than seeking to profit from the access costs of potential copiers, the highest-scoring participants in the group typically set excessively high access costs (mean=2500 calories, although ranging up to 23,000 calories) which no other group member was willing to pay. Consequently, the frequency of social learning dropped, and the frequency of payoff-biased social learning dropped to almost zero. At the group level, the overall increase in mean score illustrated in **Figure 4** disappeared, and groups of social learners with informational access costs performed no better than groups of individual learners.

In a sense, this use of informational access costs to block social learning is a product of the competitive nature of the task as it was set up in that study. Participants were informed not only of their absolute score but also their relative rank in their group (although participants in this particular study were unpaid, they seemed to be motivated primarily by rank rather than absolute performance). It was therefore in the interest of high-scoring participants to maintain their advantage by protecting their high quality information. If the incentives were to be changed such that participants are only shown or rewarded for their absolute performance and not provided with information about relative performance, then access costs might be lower and copying more frequent (although there would still be no positive incentive to sharing one's information, just no negative consequence). Alternatively, if groups rather than individuals are rewarded for their overall relative group score, then we might expect more information sharing to occur between group members (but not with members of other groups, if permitted). Adding environmental change might also encourage information sharing even in the most individually competitive situation, as participants might seek to profit from their high-quality information before it becomes out-of-date.

Nevertheless, this study is valuable in demonstrating that people (at least Western people) are not indiscriminately egalitarian with their information. Indeed, the apprenticeships observed by Stout (2002), as well as other institutions such as guilds, might be seen as following the same principles, with high-quality skills and knowledge protected from outsiders.

4. Limitations and applications

There are, of course, many limitations of laboratory experiments. Generally, experiments lack 'external validity', the degree to which the experimental situation resembles the real-life situation of interest. This is true of all experiments, but particularly so when seeking to simulate past technological change in traditional societies, as we are here. The computer-based task described above is obviously a highly abstracted and simplified version of real-life artifact design practiced by past hunter-gatherers. The task lacks any kind of motor activity and physical object affordances. The incentive (a few pounds or dollars) is very different to the incentive to feed oneself and one's family. The participants – typically Western college students – are different in many ways to the long-dead hunter-gatherers responsible for manufacturing artifacts found in the archaeological record. The time-frame is very different: an hour or so in the experiment versus years or decades acquiring the skills needed to manufacture complex artifacts such as

arrowheads or handaxes. So too is the social structure: a closed and small group of unrelated strangers in the experiments versus a much larger kin-based society with overlapping generations, migration from other groups, and so on.

All of these limitations should be recognised. Yet experiments make up for their obvious lack of external validity by having high 'internal validity', the degree to which they afford experimental control (Mesoudi 2007). In experiments we can isolate and manipulate specific variables in order to test their causal effect; we can randomly assign participants into different conditions in order to test hypotheses; we can re-run situations in multiple groups to determine whether observed effects are robust or historically contingent; and we can obtain complete and unbiased data regarding our participants' behaviour. None of these are possible with historical or ethnographic methods for both practical and ethical reasons. Archaeologists cannot 're-run' history or manipulate key variables to see how history would have changed in response to that variable, and seldom have uninterrupted or unbiased historical data sets. Ethnographers cannot randomly assign contemporary hunter-gatherers into different control and experimental societies to see how a key variable affects behaviour. Essentially, historical and observational methods are limited in being correlational, whereas experiments can test causal hypotheses.

Experiments can therefore be seen as a useful bridge between theoretical models and historical/ethnographic methods. The key point is that these methods should be used in combination. Theoretical models and experiments that are not informed by real-life historical and observational data will simply reflect the uninformed and probably incorrect intuitions of the modeller/experimenter. Conversely, historical and observational data alone cannot be used to test causal hypotheses due to their non-interventionist and correlational nature.

This interplay is hopefully illustrated in our previous application of the Virtual Arrowhead Task to a specific archaeological case study. Bettinger and Eerkens (1999) documented how projectile points from the Great Basin region of the south-western United States from around 300-600 AD exhibited systematic differences between two sites. In one site, in central Nevada, the inter-trait correlations were very high, indicative of a small number of uniform types. In eastern California, in contrast, inter-trait correlations were significantly lower, such that there were no systematic links between the dimensions of different arrowheads. Having ruled out any differences in prey or material type between the two sites, Bettinger and Eerkens (1999) suggested that the difference lay in learning strategies: prehistoric Nevada featured strong payoff- or prestige-biased social learning, such that hunters copied a small number of designs exhibited by a few high-status individuals, whereas prehistoric California featured much more individual learning, which increased variation as different hunters experimented in different ways.

As noted above, our experimental simulation supported this hypothesised scenario (Mesoudi and O'Brien 2008a): when our participants were allowed to engage in payoff-biased social learning then inter-trait correlations increased (like in Nevada), and when our participants had to rely on individual learning then inter-trait correlations were low (like in California). This supports Bettinger and Eerkens' (1999) hypothesis, and shows that it is consistent with actual human behaviour.

Yet we also showed that this hypothesis only works under certain assumptions that were not specified by Bettinger and Eerkens (1999). For example, the hypothesis only works under the assumption of a multimodal adaptive landscape. If there is a single optimal point design, then individual learners will converge on this design, and inter-trait correlations will remain high. Indeed, independent work testing the functional characteristics of projectile points suggests that multiple locally optimal designs are a reasonable assumption. Cheshier and Kelly (2006) found that long, thin points were easier to aim and hit prey with but less likely to result in a kill due to the small wounds they create, whereas thick, wide points were harder to fire but more likely to result in a kill because they created a larger wound. Here we have at least two optima: one maximising firing power, the other maximising the likelihood of a kill.

Moreover, our experimental programme suggests possible reasons why prehistoric Nevada might have featured more social learning than prehistoric California. Perhaps individual learning was more costly in

Nevada due to its harsher environment making social learning more adaptive, or perhaps informational access costs were higher in California therefore blocking social learning. These hypotheses, suggested by our experiments, can hopefully guide further archaeological research. In sum, the interplay of theoretical models, archaeological data and lab experiments provides a richer understanding of the past than any one of these methods alone.

5. Conclusions

The aim of this series of studies (Atkisson et al. 2012; Mesoudi 2008, 2011a; Mesoudi and O'Brien 2008a, b) was to test the predictions of theoretical models concerning the adaptiveness of contemporary humans' learning strategies, using a complex task designed to be representative of real-life human technology. Participants in small groups designed virtual arrowheads via individual and social learning, while we manipulated key variables such as the form of the underlying fitness functions, the possible social learning strategies permitted, the cost of individual learning, and whether social information was free or costly to access.

Our findings demonstrated that people approached this task in a broadly adaptive manner. They used a simple but effective reinforcement-based individual learning strategy that improved their payoff by leading them to a locally-optimal arrowhead design. They engaged in payoff-biased social learning in preference to alternative and less effective social learning strategies such as conformity, random copying and averaging, with this payoff-biased social learning uniquely allowing participants to jump from low-fitness locally optimal designs to high-fitness globally optimal designs that had been found by more successful group members. At a larger scale, payoff-biased social learning is especially likely to lead to cumulative cultural evolution (Aoki et al. 2012; Enquist et al. 2008; Mesoudi 2011b; Powell et al. 2009) by selectively preserving and building on effective cultural traits. It is therefore encouraging that our participants readily and preferentially engaged in this particular social learning strategy.

Moreover, both individual and social learning flexibly responded to the participants' performance in real-time. When participants were performing poorly, they made larger changes to their arrowhead when learning individually, and they were more likely to engage in payoff-biased social learning. This latter 'selective learning' – copying others only when individual learning is costly or difficult – has been shown to be adaptive relative to a mix of pure individual learners and pure social learners, allowing our participants to avoid the detrimental effect of information scrounging (Boyd and Richerson 1995).

Yet there were also flaws in our participants' learning strategies. Payoff-biased social learning was indiscriminate such that participants readily copied functionless traits from successful individuals alongside their functional traits. Indirect cues to prestige, such as looking times, were used as guides to who to copy as much as objective measures of success, even when it was inappropriate to do so, which may exacerbate the spread of neutral or even maladaptive traits. Finally, when participants were allowed to set access costs that others had to pay in order to see their arrowhead, they used these to block all social learning. At a population level, this may be detrimental to the overall preservation and accumulation of knowledge, and highlights how the cooperative motivation to share information on the part of the demonstrator is just as important as the social learners' choice of who to copy.

A comparison of contemporary humans' learning abilities with those of prehistoric hominins (either anatomically modern humans or Neanderthals) is beyond the scope of this paper, and will be left to those expert in interpreting the archaeological record. It is instructive, however, to compare the results of these studies with similar learning studies of chimpanzees. Some studies suggest that, in contrast to our human participants, chimpanzees are less likely to switch to superior solutions to tasks. Marshall-Pescini and Whiten (2008), for example, found that chimpanzees will readily copy and use a quite-good method for extracting honey from a puzzle box (sticking a wand into the box and licking honey off the end) but, when shown an even better method (using the wand to open the top of the box to expose all of the honey), fail to switch to this superior solution (see also Hrubesch et al. 2009). This stands in contrast to our participants, who readily abandoned their own arrowheads and switched to superior designs. This lack of payoff-biased

social learning in chimps might explain why their cultural traditions remain non-cumulative (Tennie et al. 2009), if they fail to selectively copy and switch to superior traits.

On the other hand, more recent studies suggest that chimpanzees *will* switch to superior methods if they are dissatisfied with their current payoff (Dean et al. 2012; Yamamoto et al. 2013), suggesting that they do exhibit some form of payoff-biased social learning. Dean et al. (2012) attributed a lack of cumulative culture in chimpanzees instead to a lack of teaching, imitation and/or prosociality. The latter finding in particular might be of particular importance. Chimpanzees have been shown to be inordinately self-interested, failing to share food with others even when there is no cost to sharing (Jensen et al. 2007; Silk et al. 2005). As we showed in our studies using informational access costs, a lack of cooperation can severely block social learning. Human cumulative culture may therefore be intimately tied to our cooperative motivations (Dean et al. 2012; Hill et al. 2009; Mesoudi and Jensen 2012).

Assuming that chimpanzees are closer behaviourally to the common ancestor of chimpanzees and humans that lived around 6 million years ago (which is, admittedly, a contestable assumption), we can speculate that somewhere in the hominin lineage the capacities for high fidelity and flexible payoff-biased social learning, tied to cooperative motivations to allow individuals to copy one other, evolved and facilitated the emergence of cumulative cultural adaptations. As illustrated in Section 4, it is possible to detect the signatures of different learning strategies in the archaeological record, as we did in the Great Basin by inferring payoff-biased social learning from high inter-trait correlations and individual learning from low inter-trait correlations. Perhaps the same might be possible with earlier material culture to determine, say, whether Neanderthals exhibited payoff-biased social learning. The appearance of culturally hitchhiking neutral or maladaptive traits might also serve as an indication of payoff-biased social learning. In sum, hopefully the further interplay of lab experiments and theoretical models, along with comparative studies of non-human primates and the archaeological study of prehistoric hominin material culture, will lead us to a better understanding of our species' success story.

References

- Akazawa T (ed) (2012) Replacement of Neanderthals by modern humans (RNMH) project series, Vol 1. RNMH Project Group, Kochi University of Technology, Tokyo
- Aoki K, Lehmann L, Feldman MW (2011) Rates of cultural change and patterns of cultural accumulation in stochastic models of social transmission. *Theor Popul Biol* 79:192-202
- Aoki K, Wakano JY, Feldman MW (2005) The emergence of social learning in a temporally changing environment: A theoretical model. *Curr Anthropol* 46:334-340
- Aoki K, Wakano JY, Lehmann L (2012) Evolutionarily stable learning schedules and cumulative culture in discrete generation models. *Theor Popul Biol* 81:300-309
- Atkisson C, Mesoudi A, O'Brien MJ (2012) Adult learners in a novel environment use prestige-biased social learning. *Evol Psychol* 10:519-537
- Bentley RA, Hahn MW, Shennan SJ (2004) Random drift and culture change. *Proc R Soc B* 271:1443-1450
- Bettinger RL, Eerkens J (1999) Point typologies, cultural transmission, and the spread of bow-and-arrow technology in the prehistoric Great Basin. *Am Antiq* 64:231-242
- Boyd R, Richerson PJ (1985) Culture and the evolutionary process. University of Chicago Press, Chicago, IL
- Boyd R, Richerson PJ (1988) An evolutionary model of social learning: The effects of spatial and temporal variation. In: Zentall T, Galef BG (eds) *Social learning*. Erlbaum, Hillsdale, NJ, pp 29-48
- Boyd R, Richerson PJ (1992) How microevolutionary processes give rise to history. In: Nitecki M, Nitecki DV (eds) *History and evolution*. State University of New York Press, Albany, pp 178-209
- Boyd R, Richerson PJ (1995) Why does culture increase human adaptability? *Ethology and Sociobiology* 16:125-143
- Cavalli-Sforza LL, Feldman MW (1981) Cultural transmission and evolution. Princeton Univ. Press, Princeton
- Cheshier J, Kelly RL (2006) Projectile point shape and durability: The effect of thickness:length. *Am Antiq* 71:353-363
- Dean LG, Kendal RL, Schapiro SJ, Thierry B, Laland KN (2012) Identification of the social and cognitive processes underlying human cumulative culture. *Science* 335:1114-1118
- Dunnell RC (1978) Style and function: A fundamental dichotomy. *Am Antiq* 35:305-319
- Durham WH (1991) Coevolution: Genes, culture, and human diversity. Stanford University Press, Stanford.
- Ehn M, Laland KN (2012) Adaptive strategies for cumulative cultural learning. *J Theor Biol* 301:103-111.
- Enquist M, Ghirlanda S, Jarrick A, Wachtmeister CA (2008) Why does human culture increase exponentially? *Theor Popul Biol* 74:46-55
- Gergely G, Csibra G (2006) Sylvia's recipe: The role of imitation and pedagogy in the transmission of cultural knowledge. In: Enfield NJ, Levenson SC (eds) *Roots of human sociality: Culture, cognition, and human interaction*. Berg, Oxford,
- Henrich J, Boyd R (1998) The evolution of conformist transmission and the emergence of between-group differences. *Evol Hum Behav* 19:215-241
- Henrich J, Gil White FJ (2001) The evolution of prestige. *Evol Hum Behav* 22:165-196
- Hill K, Barton M, Hurtado AM (2009) The emergence of human uniqueness. *Evol Anthropol* 18:187-200
- Hrubesch C, Preuschoft S, van Schaik C (2009) Skill mastery inhibits adoption of observed alternative solutions among chimpanzees (*Pan troglodytes*). *Animal Cognition* 12:209-216
- Jensen K, Call J, Tomasello M (2007) Chimpanzees are rational maximizers in an ultimatum game. *Science* 318:107-109
- Laland KN (2004) Social learning strategies. *Learn Behav* 32:4-14
- Lycett SJ, von Cramon-Taubadel N (2008) Acheulean variability and hominin dispersals: A model-bound approach. *J Archaeol Sci* 35:553-562
- Marshall-Pescini S, Whiten A (2008) Chimpanzees (*Pan troglodytes*) and the question of cumulative culture: An experimental approach. *Animal Cognition* 11:449-456
- McElreath R, Lubell M, Richerson PJ, Waring TM, Baum W, Edsten E, Efferson C, Paciotti B (2005) Applying evolutionary models to the laboratory study of social learning. *Evol Hum Behav* 26:483-508
- Mesoudi A (2007) Using the methods of social psychology to study cultural evolution. *J Soc Evol Cult Psychol* 1:35-58

- Mesoudi A (2008) An experimental simulation of the 'copy-successful-individuals' cultural learning strategy: Adaptive landscapes, producer-scrounger dynamics and informational access costs. *Evol Hum Behav* 29:350-363
- Mesoudi A (2009) The cultural dynamics of copycat suicide. *PLOS ONE* 4:e7252.
- Mesoudi A (2011a) An experimental comparison of human social learning strategies: Payoff-biased social learning is adaptive but underused. *Evol Hum Behav* 32:334-342
- Mesoudi A (2011b) Variable cultural acquisition costs constrain cumulative cultural evolution. *PLOS One* 6:e18239
- Mesoudi A, Jensen K (2012) Culture and the evolution of human sociality. In: Vonk J, Shackelford T (eds) *The oxford handbook of comparative evolutionary psychology*. Oxford University Press, Oxford, UK, pp 419-433
- Mesoudi A, O'Brien MJ (2008a) The cultural transmission of Great Basin projectile point technology I: An experimental simulation. *Am Antiq* 73:3-28
- Mesoudi A, O'Brien MJ (2008b) The cultural transmission of Great Basin projectile point technology II: An agent-based computer simulation. *Am Antiq* 73:627-644
- Morgan TJH, Rendell LE, Ehn M, Hoppitt W, Laland KN (2012) The evolutionary basis of human social learning. *Proc Roy Soc B* 279:653-662.
- Neiman FD (1995) Stylistic variation in evolutionary perspective - inferences from decorative diversity and interassemblage distance in Illinois woodland ceramic assemblages. *Am Antiq* 60:7-36
- O'Brien MJ, Darwent J, Lyman RL (2001) Cladistics is useful for reconstructing archaeological phylogenies: Palaeoindian points from the southeastern United States. *J Archaeol Sci* 28:1115-1136
- Powell A, Shennan S, Thomas MG (2009) Late Pleistocene demography and the appearance of modern human behavior. *Science* 324:1298-1301
- Richerson PJ, Boyd R (2005) *Not by genes alone*. University of Chicago Press, Chicago
- Rogers AR (1988) Does biology constrain culture? *Am Anthropol* 90:819-831
- Rogers DS, Ehrlich PR (2008) Natural selection and cultural rates of change. *PNAS* 105:3416
- Silk JB, Brosnan SF, Vonk J, Henrich J, Povinelli DJ, Richardson AS, Lambeth SP, Mascaró J, Schapiro SJ (2005) Chimpanzees are indifferent to the welfare of unrelated group members. *Nature* 437:1357-1359
- Simon HA (1956) Rational choice and the structure of the environment. *Psychological Review; Psychological Review* 63:129-138
- Sperber D (1996) *Explaining culture: A naturalistic approach*. Oxford University Press, Oxford.
- Stout D (2002) Skill and cognition in stone tool production. *Curr Anthropol* 43:693-722.
- Tennie C, Call J, Tomasello M (2009) Ratcheting up the ratchet: On the evolution of cumulative culture. *Phil Trans R Soc B* 364:2405-2415
- Wright S (1932) The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proceedings of the Sixth International Congress of Genetics* 1:356-366
- Yamamoto S, Humle T, Tanaka M (2013) Basis for cumulative cultural evolution in chimpanzees: Social learning of a more efficient tool-use technique. *PLOS ONE* 8:e55768.

Figures

Figure 1 - A screenshot of the Virtual Arrowhead Task. Participants can choose to directly change the traits in the box at the top (individual learning) or copy the design of another participant in the box on the left (social learning). Feedback is given in calories depending on how close the design is to one or more hidden optimal designs.

Height (1-100):
Shape:
Thickness (1-100):

Width (1-100):
Color:

PREVIOUS GROUP

PLAYER#	SCORE
Player 1	20,875
Player 2	17,537
Player 3	17,150
Player 4	16,426
Player 5	13,637
Player 6	16,472

The number of calories you received is shown below. Also shown below is your Season Score, the running total of all of your hunts, the number of Hunts Left and the current Season.

Now press NEXT HUNT to begin designing your arrowhead

01000

Season: 1

Calories (/1000): 561

Hunts Left: 30
Season Score: 561

Figure 2 - Fitness functions for the constituent traits. The overall fitness of an arrowhead was given by the sum of these fitness functions. The continuous traits (Height, Width and Thickness) had bimodal functions, generating a multimodal adaptive landscape. A fifth trait, Colour, was neutral and did not affect arrowhead fitness. From Mesoudi and O'Brien (2008a)

Figure 3 - The mean score of a sample of individual learners ($N=27$) showing a gradual increase and plateauing over successive hunts (black line with squares), along with the simulated performance of the reinforcement learning strategy with $d=1$ and $c=5$ (red line with triangles).

Figure 4 - The mean score of individual learners (black line with squares) compared to payoff-biased social learners (blue line with circles). The latter could copy one another only during the last five hunts, during which their score significantly increased relative to the individual learners, who could not copy on any hunt.