ELSEVIER

Evolution and Human Behavior 29 (2008) 350–363

Evolution and Human Behavior

An experimental simulation of the "copy-successful-individuals" cultural learning strategy: adaptive landscapes, producer–scrounger dynamics, and informational access costs

Alex Mesoudi*

Department of Social and Developmental Psychology, University of Cambridge, Free School Lane, CB2 3RQ Cambridge, United Kingdom

Initial receipt 29 January 2008; final revision received 20 April 2008

Abstract

An experimental simulation of cultural evolution was conducted using the virtual arrowhead task. Participants designed "virtual arrowheads" and tested them in "virtual hunting environments", improving their designs through either individual trial-and-error learning or by copying the design of another participant. A previous study using this task [Mesoudi, A., & O'Brien, M. J. (2008). The cultural transmission of Great Basin projectile point technology I: An experimental simulation. *American Antiquity*, 73, 3–28.] found that a cultural learning strategy of "copy-successful-individuals" was significantly more adaptive than individual learning. The present study explored the robustness of this finding using the same task but under different conditions. It was found that (a) individual learning was significantly more adaptive in a unimodal adaptive landscape than in a multimodal adaptive landscape, suggesting that the adaptive advantage of cultural learning would disappear in unimodal environments; (b) the adaptive advantage of copy-successful-individuals was maintained when cultural learning was permitted at regular intervals, despite the increased opportunity for information scroungers to inhibit exploration of the environment, because participants flexibly switched between individual and cultural learning depending on circumstances; (c) allowing participants to set access costs that other participants must pay in order to view their designs severely curtailed the use of cultural learning and especially the copy-successful-individuals strategy.

Keywords: Adaptive landscapes; Copy-successful-individuals; Cultural evolution; Cultural learning; Cultural transmission; Gene-culture coevolution; Indirect bias; Producer–scrounger dynamics; Informational access costs; Prestige bias; Social learning; Virtual arrowhead task

1. Introduction

Over the past 30 years, the field of cultural evolution (or gene culture coevolution/dual inheritance theory) has sought to incorporate culture into evolutionary analyses of human behaviour by treating culture as an evolutionary process that operates in parallel to genetic evolution (Boyd & Richerson, 1985, 2005; Cavalli-Sforza & Feldman, 1981; Henrich & McElreath, 2003; Laland & Brown, 2002; Mesoudi, Whiten, & Laland, 2004, 2006; Richerson & Boyd, 2005). The majority of this work has involved the construction and analysis of formal mathematical models of cultural processes (e.g., Cavalli-Sforza & Feldman, 1981; Boyd & Richerson, 1985, 2005) drawing on the methods of theoretical population genetics. Such models have, for example, explored the conditions under which culture itself would be favoured by genetic evolution (Boyd & Richerson, 1985, 2005), the properties of different modes of cultural transmission, such as vertical, oblique, and horizontal transmission (Cavalli-Sforza & Feldman, 1981), and the origin and consequences of various cultural forces, such as conformity (Henrich & Boyd, 1998) and prestige bias (Henrich & Gil White, 2001).

Mathematical models, however, are only as good as their assumptions, and these assumptions, as well as the models' predictions, need to be empirically tested. One way of doing this is by using laboratory experiments (Mesoudi, 2007), and

---

* Tel.: +44 01223 334577.
E-mail address: am786@cam.ac.uk.
URL: http://amesoudi.googlepages.com.

1090-5138/$ – see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.evolhumbehav.2008.04.005

several studies have recently sought to simulate cultural evolution and cultural transmission in the lab using experimental methods (Baum, Richerson, Efferson, & Paciotti, 2004; Caldwell & Millen, 2008; Efferson, Lalive, Richerson, McElreath, & Lubell, 2008; Kameda & Nakanishi, 2002, 2003; McElreath et al., 2005; Mesoudi & O'Brien, 2008; Mesoudi, Whiten, & Dunbar, 2006). In one of these studies, Mesoudi and O'Brien (2008) experimentally simulated a cultural learning strategy of "copy successful individuals" [see Laland (2004); also resembling Boyd & Richerson's (1985) "indirect bias" or Henrich & Gil-White's (2001) "prestige bias"] using the "virtual arrowhead" experimental task. In this task, participants played a computer game in which they designed a technological artifact (an arrowhead) either by individual trial-and-error learning or by copying successful fellow participants. Mesoudi and O'Brien found that allowing participants to preferentially copy the designs of successful models resulted in significantly improved performance relative to individual learning controls, suggesting that this copy-successful-individuals cultural learning strategy is significantly more adaptive than individual learning.

Here I present a further experiment using the same virtual arrowhead task that is designed to explore the robustness of this earlier finding. Specifically, the present study asks whether the copy-successful-individuals strategy remains adaptive when (a) the shape of the adaptive landscape that determines a design's fitness is changed from multimodal to unimodal, thus making individual learning easier; (b) when cultural learning is permitted at regular intervals throughout the experiment rather than only during the last few trials, thus allowing greater opportunity for the emergence of free-riding information scroungers; and (c) when potential models can set informational access costs that potential learners must pay in order to access the models' information. Before discussing the theoretical basis of these issues, the following section provides more details of the virtual arrowhead experimental task.

2. The virtual arrowhead experimental task

Mesoudi and O'Brien (2008) employed a computer-based experimental task in which participants designed their own virtual arrowheads. [The aim of that study was to simulate a specific archaeological scenario of past cultural evolution proposed by evolutionary archaeologists Bettinger and Eerkens (1999); however, these archaeological details are not directly relevant to the present study and are not discussed further.] Participants are told to imagine themselves as hunters in the distant past and told that they must design an arrowhead that they will then use to go hunting for food. They are told that their hunting success depends on the design of their arrowhead: Some arrowhead designs give higher returns than other arrowhead designs, and it is their task to find the best/optimal arrowhead design for their particular hunting environment.

The participants enter values for five attributes of their arrowhead: length, width, thickness, shape, and color. The first three attributes are continuous, ranging from 1 to 100 arbitrary units, and the last two are discrete, each taking one of four different values. Once the participants have entered values for all five attributes, they can test their design by going on a hunt (trial), during which they are given feedback on the success of their arrowhead design. This feedback is given in calories, ranging from 1 to 1000, which is partly determined by the participants' chosen attribute values via experimenter-set fitness functions. These fitness functions may vary in shape (e.g., unimodal or bimodal: see Section 3.1) or, alternatively, attributes may be neutral (e.g., color) and not affect fitness. This mix of continuous and discrete, functional and neutral attributes is intended to provide a more realistic simulation of cultural (especially technological) evolution compared with previous cultural evolution experiments (e.g., Efferson et al., 2008; Kameda & Nakanishi 2002, 2003; McElreath et al., 2005), which employ a relatively simple task featuring a single dichotomous attribute (e.g., which one of two crops to plant). There is also random error in the feedback, where the feedback displayed to participants was randomly drawn from a normal distribution with a mean of the actual fitness of the participant's arrowhead design and a standard deviation of 5 calories. This small random error makes individual learning somewhat unreliable, which is probably more representative of most real-life learning tasks than an assumption of perfect environmental feedback. During a typical hour-long experiment, the participant plays three seasons of hunting, with each season comprising 30 hunts. There are therefore 30 opportunities to improve one's arrowhead design during each season.

In order to explore the effect of cultural/social learning, participants play the virtual arrowhead game on networked PCs. The experimenter can allow participants to view the arrowhead design of one or more other participants in the same group via the on-screen interface. In Mesoudi and O'Brien (2008), participants in groups of six had to learn individually for the first 25 hunts of each season before being allowed to copy each other during the last five hunts of each season. Specifically, we simulated a copy-successful-individuals cultural learning strategy, in which participants could see the cumulative score of every other participant in their group and choose to copy one of those fellow group members. The strategy followed by the majority of participants was to copy the arrowhead of the participant with the highest score. As shown in Fig. 1A, participants who could employ the copy-successful-individuals cultural learning strategy significantly outperformed individual controls (who did not engage in cultural learning at any time) during the last five hunts when the former were allowed to copy one another.

3. Aims of the current study

The experiment reported here was designed to explore the robustness of the aforementioned finding that a copy-successful-individuals cultural learning strategy is more adaptive than individual learning. The virtual arrowhead task described above was modified in several ways, designed to test the adaptiveness of the copy-successful-individuals strategy under a wider range of conditions. Aspects of the task that did not change were the season/hunt structure (three seasons of hunting, 30 hunts per season), the number of attributes (three continuous, two discrete), the method of input and feedback, the group sizes (groups of five or six plus individual controls), and the form of cultural learning (copy-successful-individuals, where participants have access to other group members' cumulative scores and can copy on the basis of this information).

3.1. Unimodal vs. multimodal adaptive landscapes

Boyd and Richerson (1992) have argued that much cultural evolution, just like much biological/genetic evolution, operates within adaptive landscapes of varying shapes. An "adaptive landscape" is a concept commonly used in evolutionary biology (Arnold, Pfrender, & Jones, 2001; Wright, 1932) to represent the design space of all possible combinations of multiple phenotypic characters, where the height of the landscape represents fitness. "Peaks" in the landscape represent phenotypes of high fitness, while "valleys" represent phenotypes of low fitness. The same principles are likely to apply to technological artifacts, given that many artifacts are similarly determined by multiple characters each of which contribute to the artifact's success or cultural fitness. This leads to the prediction:

> "If the adaptive [landscape] has a unique maximum, then every population will evolve to the same equilibrium mean phenotype, independent of its starting position. On the other hand, if there is more than one local maximum, different equilibrium outcomes are possible depending on initial conditions." (Boyd & Richerson, 1992, p. 191)

Mesoudi and O'Brien (2008) implemented the latter, a multimodal adaptive landscape with multiple locally optimal peaks. Our adaptive landscape represented all possible arrowhead designs, or all possible combinations of all attribute values, where the height of the landscape represents the payoff to the participant of that arrowhead design. We made the landscape multimodal by giving each of the three continuous attributes (length, width, and thickness) bimodal fitness functions, as shown in Fig. 1B. (See Appendix A for fitness equations. The two discrete attributes did not affect the shape of the adaptive landscape.) Each attribute therefore had two locally optimal values, one local and one global; deviation from either of these optima reduced the feedback score, but the global optimum gave a higher score than the local optimum. The score seen by participants was proportional to the sum of all attributes' fitness contributions, therefore giving a multimodal adaptive landscape with eight peaks of varying fitness. For example, the highest peak is found where length, width, and thickness are all at their global optima, a slightly lower peak where length and width are at global optima and thickness is at its local optimum, and so on. There are $2^3=8$ peaks in total. Mesoudi and O'Brien predicted that individual learning would result in different participants getting stuck on locally optimal but globally suboptimal fitness peaks, whereas copy-successful-individuals cultural learning allows participants to "jump" from their low-fitness peak to a higher-fitness peak found by a more successful member of their group. As shown in Fig. 1A, we did indeed find that cultural learners outperformed individual learners, consistent with this prediction. Hence it appeared that the multimodal adaptive landscape was instrumental in making the copy-successful-individuals strategy more adaptive than individual learning.

The first aim of the present study was to further test this explanation for the adaptive advantage of copy-successful-individuals. If individual learning is indeed hampered by multimodal fitness environments because individuals get stuck on locally optimal peaks, then removing those peaks should significantly improve individual learners' perfor-

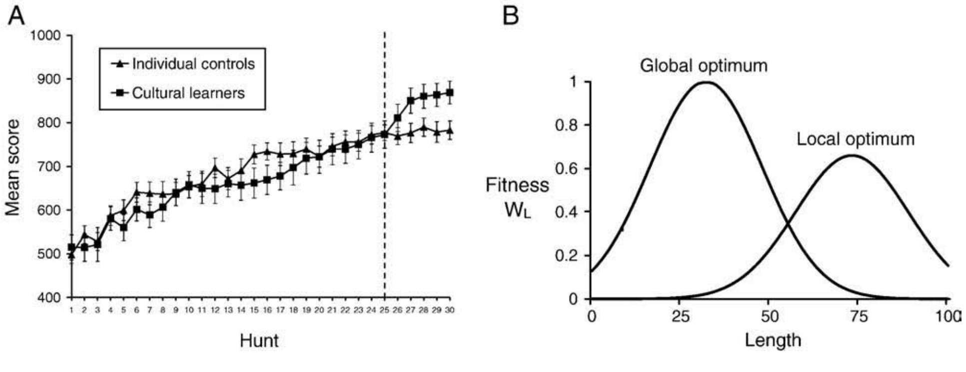

Fig. 1. (A) Mean score per participant in Mesoudi and O'Brien (2008), comparing individual learners with cultural learners. The latter group could engage in copy-successful-individuals cultural learning only during the final five hunts, to the right of the vertical dotted line. Error bars show standard error. (B) An example bimodal fitness function for one of the continuous arrowhead attributes, length. Here, a length of 30 gave a maximum fitness contribution of $W_L=1$, while a length of 75 gave a locally optimal but globally suboptimal fitness of $W_L=0.66$. Length, width, and thickness each had bimodal fitness functions like this, with peaks at different *x*-axis values.

mance. Thus, in the first season of hunting, individual learners experienced a unimodal fitness environment with a single optimum/peak. It is predicted that these unimodal individual learners will easily find this single optimum, thus perform just as well as the cultural learners of Mesoudi and O'Brien (2008), eliminating the adaptive advantage of the copy-successful-individuals strategy.

3.2. Terminal vs. periodic cultural learning

The opportunity for cultural learning in Mesoudi and O'Brien (2008) was limited to the final 5 hunts of a 30-hunt sequence and followed a much longer period of individual learning (henceforth referred to as "terminal cultural learning"). In reality, however, people are not forced to undergo a period of obligatory individual learning before being allowed to engage in cultural learning. Rather, they can switch between individual and cultural learning as they see fit. In the present experiment, participants therefore engaged instead in "periodic cultural learning", where they could copy other participants every fifth hunt after the first five hunts (Hunts 6, 11, 16, 21, and 26). It is predicted that making cultural learning periodic throughout the season will favour the emergence of "information scroungers"—participants who forego lengthy and costly individual learning and instead consistently free-ride on the individual learning efforts of other participants ("information producers") in the group. Whereas terminal cultural learning has an enforced period of individual learning during which the adaptive landscape can be more extensively explored by all group members, periodic cultural learning allows information scroungers to consistently free-ride throughout the season, thus inhibiting exploration of the adaptive landscape and reducing the chances that the globally optimal design is found. Hence, it is predicted that periodic cultural learners will have lower mean scores than terminal cultural learners.

In support of this prediction, Kameda and Nakanishi (2002) found experimentally that participants divided themselves into information producers and information scroungers in the manner suggested above, and that these two groups coexisted at equilibrium. Moreover, a model analysed by Rogers (1988) suggests that a mix of individual and cultural learners can never have a mean fitness greater than a population of all individual learners because of the presence of information scroungers in the former. This might be used to predict that periodic cultural learners will in the present study perform no better than individual controls. However, the reason why information scroungers are predicted to reduce group fitness is different in the present study compared to Rogers' model and Kameda and Nakanishi's experiment. In the latter, the reason that information scroungers reduce group fitness is environmental change. Because information scroungers cannot track novel environmental change, group fitness is reduced when scroungers copy outdated information from each other. In a constant environment, on the other hand, once information producers have acquired the optimal behaviour, then information scroungers can copy this optimal behaviour with no reduction in group fitness (in fact, group fitness should increase given the assumed lower learning costs of scrounging). In the present experiment, the environment does not change, but it is predicted that information scroungers will emerge and reduce group fitness for a different reason: the multimodal adaptive landscape. Information scroungers will reduce group fitness by reducing the number of participants learning individually, thus reducing the chances that one of the higher fitness peaks will be found. Note also that the detrimental effect of information scroungers would persist even after equilibrium is reached, because this equilibrium is less likely to be the globally optimal, high-fitness design.

3.3. Informational access costs

An unrealistic aspect of all of the cultural evolution experiments mentioned so far is that potential models have no choice whether to allow potential cultural learners to copy them. It is much more likely that, at least amongst non-kin, successful or attractive models might set an "access cost" that others must pay in order to gain access to their knowledge. Henrich and Gil-White (2001) have argued that information is commonly treated as a good that is traded between models and learners, with low-status individuals giving gifts or deference to high-status (prestigious) individuals in exchange for their expert knowledge, as is commonly observed in the ethnographic record.

In the third season of the present experiment, participants were allowed to set access costs that other participants had to pay them in order to view their arrowhead design. These costs were expressed in the same units as the hunting feedback (calories) and were added and deducted from each player's total score in real time. It is predicted that access costs will reduce the frequency of cultural learning and especially the copy-successful-individuals strategy, because successful players will set high access costs in an attempt to protect their high-quality information from potential competitors. Hence, groups in which access costs can be set are predicted to have lower mean scores than groups in which access costs cannot be set. On the other hand, some successful players might set prices that are not excessively high, thus receiving modest fees from several other players and increasing their score further. If the buyers receive information that increases their score by more than the price they paid, then their costs will be offset, and the mean score of the entire group would increase. Thus, access costs might generate an "information goods" market in which the information-rich get even richer, while a trickle-down effect makes the information-poor slightly less poor. Given the lack of formal theoretical work in the cultural evolution literature on access costs, however, this aspect of the study remains somewhat exploratory.

4. Methods

Details of participants and procedure for Mesoudi and O'Brien (2008) can be found in that publication. The following describes the methods and participants for the new experiment only.

4.1. Participants

Seventy-five participants of mean age 16.51 years (SD=0.58) participated in the experiment. Participants were unpaid and completed the study as part of an A-Level Psychology class at a sixth form (11th grade) college in Cambridge, UK. The participants (64% female) had little experience of participating in psychology experiments and no knowledge of the theory discussed above. Seventeen participants were individual controls. The other 58 formed eight groups of 6 (n=48) and two groups of 5 (n=10). Due to computer failure, for participants in one session (three groups of 6, n=18) the final season of hunting was not completed, so Season 3 data are available for 40 participants only. However, there is no reason why these participants' data would differ from the data of participants who completed the entire study.

4.2. Procedure

The experiment was run in a large computer classroom. Each participant sat at a physically separate, networked computer and completed the experiment entirely via the on-screen interface. Participants were randomly allocated to a group or to be an individual control. Participants were instructed not to talk or communicate with each other except via the computer program, and not to write anything down. The entire study lasted 45–60 min.

4.3. Task

The aim of the virtual arrowhead experiment is to design a "virtual arrowhead" and test this arrowhead design in a "virtual hunting environment", all via a computer program (programmed in Borland C++ Builder 6.0; code available upon request). Participants entered values for five arrowhead attributes: length, width, and thickness (continuous attributes each ranging from 1 to 100 arbitrary units), shape and color (each taking one of four discrete values). Once valid values of each attribute had been entered, participants could press the HUNT button to receive feedback in calories, ranging from 1 to 1000. The experiment was divided into three seasons each comprising 30 hunts (trials). The stated objective was to find the optimal values of each attribute that would give the highest possible score (1000 calories). This score was calculated using fitness functions (see Appendix A), although participants were told nothing about these functions. There was also random error in the feedback, with the score seen by participants randomly taken from a normal distribution with a mean given by the fitness of their arrowhead design and a standard deviation of 5 units. Participants were informed about the presence of random error in feedback, but not its magnitude. During each season, the running total of all calories received during every hunt up to that point was displayed on screen, giving a cumulative score, and participants in groups were informed how they ranked within their group with respect to this cumulative score. At the start of each season the cumulative score was reset to zero and the optimal values were changed, although optimal values did not change within each season. Participants were informed that optimal values changed between seasons and not during seasons.

Individual controls played all three seasons with no opportunity to learn from other participants; hence these participants relied entirely on individual trial-and-error learning. Participants in groups ("cultural learners") could, every five hunts (Hunts 6, 11, 16, 21, and 26), see on screen a list of other participants in their five- or six-person group, anonymously labelled Player 1, Player 2, etc. Displayed next to each of these labels was each player's cumulative score. Each participant could click on a button next to one of the other group members' labels and that person's attribute values were displayed on their screen next to their own values. The participant could then modify their own attribute values to match the model's attribute values if they wished. There was no obligation at any point to either view another group member's design or view the most successful group member's design, or even if viewed to change one's own arrowhead design to match the model's. Video clips of participants playing an individual learning hunt and a cultural learning hunt of the virtual arrowhead task can be found in the Electronic Supplementary Material.

4.4. Design

Each season differed in order to test the predictions made above. During Season 1, the individual controls experienced unimodal fitness functions (see Section 3.1 and Appendix A). Each attribute had a single optimal value (what was formerly the global optimum), giving a single optimal point design. This condition was then compared with equivalent data from Mesoudi and O'Brien (2008) in which individual learners faced a multimodal environment. Cultural learners in all three seasons, and individual controls in Seasons 2 and 3, experienced bimodal fitness functions, as in Mesoudi and O'Brien. Unlike in Mesoudi and O'Brien, cultural learning in all three seasons was periodic (see Section 3.2), i.e., permitted every fifth hunt (Hunts 6, 11, 16, 21, and 26). Hence, a comparison could be made between periodic and terminal cultural learning, both occurring within a multimodal environment. During Season 3, cultural learners were asked to set a price in calories that other players would have to pay in order to view their arrowhead design (see Section 3.3). Participants set this price immediately before each cultural learning hunt. During these hunts each player's price was displayed next to their label and cumulative score. If a participant chose to copy another player in Season 3, then the model's price was automatically and immediately deducted from the learner's cumulative score. Following each cultural learning hunt, participants were informed of how many other players had copied them and how many calories they will consequently receive (number of players copying them multiplied by their access cost) and they saw their score increase by this amount.

4.5. Statistical analyses

In order to compare different experimental conditions, mixed analyses of variance (ANOVAs) were performed typically on the last five hunts of each season. Hunt was a within-groups factor with five levels, and experimental group was a between-groups factor with two levels (cultural learners vs. individual controls). The last five hunts were singled out to enable a comparison with the terminal cultural learning of Mesoudi and O'Brien (2008), which only occurred during these final five hunts, and because equilibrium had typically been reached by these last five hunts (see Electronic Supplementary Material).

5. Results

5.1. How does the shape of the adaptive landscape affect the adaptiveness of individual learning?

Fig. 2 shows the mean score of the unimodal individual learners from Season 1, with the multimodal individual learners and multimodal cultural learners from Mesoudi and O'Brien (2008) for comparison. Unimodal individual learners from the present study outperformed multimodal individual learners from Mesoudi and O'Brien, a difference that was significant over the final five hunts [$F(1,45)=6.28$, $p<.016$]. This indicates that individual learning was significantly more adaptive in unimodal environments than in multimodal environments. Against this background, the advantage observed by Mesoudi and O'Brien for copy-successful-individuals disappears, given that this strategy does not differ significantly from unimodal individual learning over the last five hunts [$F(1,33)=0.13$, nonsignificant (ns)]. In fact, the mean score at the final hunt for unimodal individual learners (866.28) was very similar to the mean score at the final hunt for the terminal/multimodal cultural learners (869.09), and both were much larger than the equivalent value for multimodal individual learners (785.72). Hence, the hypothesis that individual learning is easier/more adaptive in unimodal environments than multimodal environments, such that the observed advantage of cultural learning disappears, is supported.

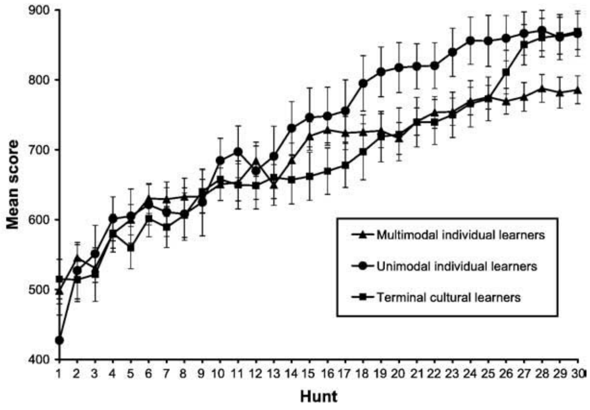

Fig. 2. Mean scores of unimodal individual learners from the present study, multimodal individual learners from Mesoudi and O'Brien (2008), and multimodal/terminal cultural learners from Mesoudi and O'Brien. The latter could engage in cultural learning only during the final five hunts. Error bars show standard error.

It should be noted that the two fitness environments (unimodal and multimodal) differed in their average payoff irrespective of the shape of the adaptive landscape. The average payoff from the multimodal environment (i.e., the average score from every value of every attribute) was 565 calories, while the equivalent average payoff from the unimodal environment was 430 calories (imagine removing the smaller normal distribution from Fig. 1B and the resulting drop in average fitness across all values of length). This means that unimodal individual learners would have received lower scores than would multimodal individual learners if they chose attribute values purely at random. Indeed, this can be observed in the very first hunt in Fig. 2. This a priori disadvantage for unimodal individual learners was rapidly overcome, however, given that unimodal individual learners eventually outperformed multimodal individual learners, and makes the fitness advantage seen for unimodal individual learners more impressive.

5.2. How do periodic cultural learners compare with terminal cultural learners and individual learners?

Fig. 3 shows, separately for Seasons 1 and 2, the mean score for periodic cultural learners, terminal cultural learners, and individual learners (all of whom learned in a multimodal fitness environment, such that none of the differences discussed below can be attributed to differences in fitness functions). For Season 1 (Fig. 3A), the mean cumulative score (the sum of the scores from all 30 hunts) was significantly higher for periodic cultural learners than individual learners [independent samples *t* test: $t(86)=2.38$, $p<.02$]. However, there was no significant difference in mean score between periodic cultural learners and individual learners during the final five hunts [$F(1,86)=2.52$, ns]. So groups of periodic cultural learners during Season 1 showed superior performance over the entire 30 hunts, but did not end up at equilibrium with higher scores than individual controls. This can be contrasted with terminal cultural learners, who did not show superior performance over individual learners until the final five hunts. Periodic and terminal cultural learners can also be directly compared after equivalent numbers of cultural learning opportunities (e.g., the first cultural learning opportunity for terminal cultural learners was Hunt 26, the equivalent for periodic cultural learners was Hunt 6). For Season 1, the terminal cultural learners

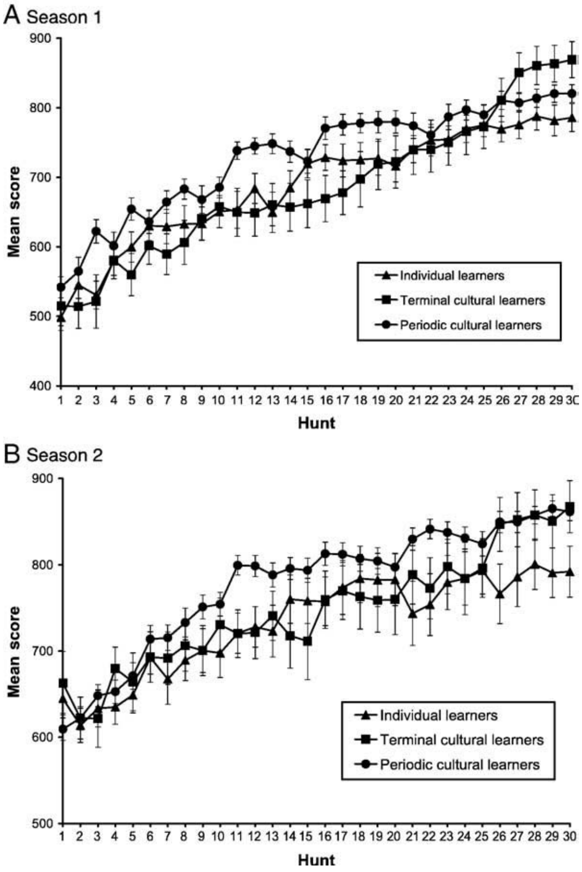

Fig. 3. (A) Mean scores for periodic cultural learners from Season 1 of the present study, individual learners from Mesoudi and O’Brien (2008), and terminal cultural learners from Mesoudi and O’Brien. All three groups experienced the same multimodal environment. Error bars show standard error. (B) Equivalent data for Season 2 of both studies.

significantly outperformed the periodic cultural learners at every equivalent cultural learning hunt [t’s(74)>2.13, p’s<.015]. Hence, the prediction that periodic cultural learning results in lower fitness than terminal cultural learning and performs no better than individual learning is upheld for Season 1. For Season 2 (Fig. 3B), the mean cumulative score was significantly higher for periodic cultural learners than individual controls [$t(86)=2.73$, $p<.01$], and periodic cultural learners significantly outperformed individual controls over the last five hunts [$F(1,73)=8.49$, $p<.005$]. Comparing equivalent cultural learning opportunities for periodic and terminal cultural learners during Season 2, only the first opportunity significantly differed in score [$t(74)=3.94$, $p<.001$], with the others showing no significant difference [t’s(74)<1.98]. So the prediction that periodic cultural learning results in lower fitness than terminal cultural learning and does not outperform individual learning was not supported for Season 2.

It was argued that periodic cultural learners would perform less well than terminal cultural learners because information scroungers in the former would inhibit exploration of the adaptive landscape, making it less likely that groups of periodic cultural learners will find the global optimum than terminal cultural learners, who have an enforced period of individual learning exploration before cultural learning is allowed. This can be tested by comparing the number of attributes that are at their global optima for terminal vs. periodic learners at equilibrium, predicting that the former should have more. The number of continuous attributes that were ±5 units of the global optimum at Hunt 30 was tallied, giving a number between 0 and 3 for each participant. For Season 1, there was a non-significant trend for terminal cultural learners to have found more global optima than the periodic cultural learners (Mann–Whitney test: $z=1.74$, $p=.082$). For Season 2, there was no difference between terminal and periodic cultural learners ($z=0.46$, $p=.643$). Hence, there is suggestive but not conclusive evidence that there was less exploration of the adaptive landscape by the periodic cultural learners in Season 1, resulting in lower fitness than terminal cultural learners, as predicted. For Season 2 there was no difference in exploration between the two groups, and periodic cultural learners performed equally as well as the terminal cultural learners, contrary to predictions.

To explore this unexpected finding further, more detailed analyses were carried out on the participants’ use of social information. Recall that cultural learners were not obligated to view another participant’s design, nor view the most successful group member, nor change their own design to match the model’s when they had viewed. Fig. 4A shows the frequency of cultural learners who chose to view another player’s design, further divided into those who viewed the most successful member of their group vs. those who viewed a different group member. Across all viewing opportunities in Seasons 1 and 2, a majority of participants chose to view during every viewing opportunity (irrespective of who they viewed). Except for the first cultural learning hunt of Season 1, the majority of participants during Seasons 1 and 2 who viewed another player’s design chose to view the most successful group member. Across all hunts of Season 1, 76.8% of participants who viewed chose to view the most successful group member (85.5% excluding Hunt 6). This choice significantly differed from that expected if participants were selecting models at random [$\chi^2(5)=469.00$, $p<.001$]. Season 2 showed the same pattern, with a majority of views targeting the most successful group member (81.4%), which again significantly differed from random choice [$\chi^2(5)=498.44$, $p<.001$]. Aside from the first opportunity of Season 1, which might be discounted as a practice trial during which the participants learned how to use the viewing feature of the program, the majority of cultural learners who viewed during Seasons 1 and 2 employed the copy-successful-individuals strategy.

However, we still cannot be sure that participants who viewed another participant’s design actually changed their

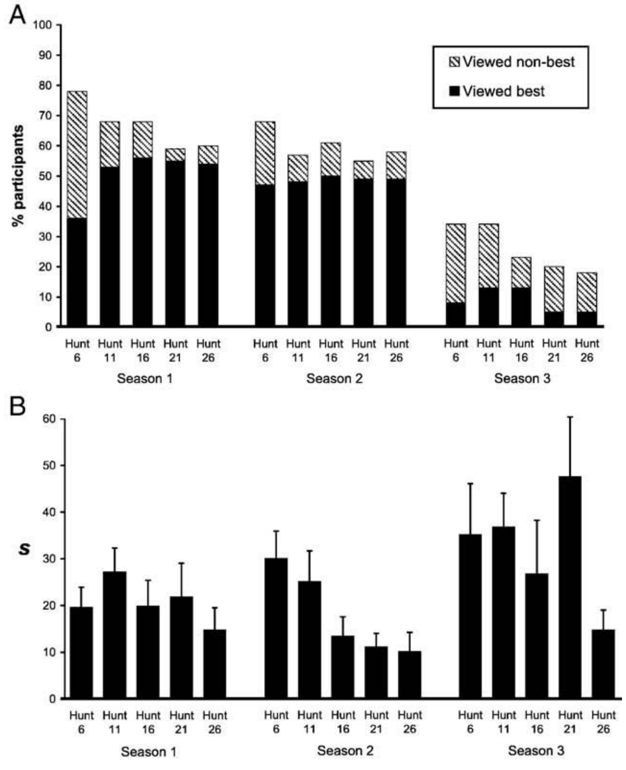

Fig. 4. (A) Percentages of periodic cultural learners who viewed another participant's arrowhead design during each cultural learning opportunity. Bars are divided into instances in which participants copied the most successful member of their group vs. instances in which another member of the group was copied. (B) Mean social influence, s, the degree to which a cultural learner modified their attribute values to more closely match the attribute values of a viewed model. Error bars show standard error.

own design to match the models'. A measure of "social influence", s, was therefore calculated, defined as the degree to which a cultural learning participant who viewed another participant's design changed their own attribute values to be more similar to the attribute values of the model that they viewed (see Appendix B for details of how s was calculated). Positive values of s indicate that a participant made their design more similar to the model's, while negative values of s indicate movement away from the model's design. The size of s indicates the magnitude of this change. When $s=0$, there is no change in the participant's design relative to the model's. Fig. 4B shows that s was positive during every cultural learning hunt of Season 1 and consequently significantly differed from zero [one-sample t test: $t(199)=8.66$, $p<.001$], and was also positive during Season 2, again significantly different from zero [$t(181)=7.96$, $p<.001$]. Most of this Season 2 change took place during the first two cultural learning opportunities.

The within-participant consistency of s can be used to determine whether the periodic cultural learners were polymorphic or monomorphic (Kameda & Nakanishi, 2002). Polymorphism is where some participants consistently engage in individual learning (with $s=0$) and others consistently engage in cultural learning (with $s>0$), such that each participant's values of s would be relatively consistent over trials. Monomorphism is where every participant engages in either individual learning or cultural learning with a fixed probability, such that their values of s are inconsistent over trials. An index of the within-participant consistency of s over trials was therefore calculated (see Appendix B). This index ranged from 0 (*entirely inconsistent*) to 4 (*entirely consistent*). For Season 1, the mean consistency was 2.34 (SD=1.28), which was significantly different from a perfect consistency of 4 [$t(57)=9.86$, $p<.001$]; for Season 2, the mean consistency was 2.45 (SD=1.23), again significantly different from 4 [$t(57)=9.60$, $p<.001$]. This indicates a deviation from pure polymorphism (index=4) and suggests that participants were switching between individual learning and cultural learning over the course of the seasons.

The inconsistency in social information use suggests that participants may have been flexibly switching between individual and cultural learning according to circumstances. Boyd and Richerson (1995) found analytically that populations of flexible learners, who engage in individual learning when individual learning is easy/cheap and cultural learning when individual learning is difficult/expensive, do better than both populations of all individual learners and polymorphic populations of individual and cultural learners who cannot switch strategies. Perhaps the participants here were similarly switching from individual to cultural learning when they found that individual learning was not effective. This would predict negative correlations between the cues of performance that the participants received (cumulative score and rank) and their use of social information (s): Participants who receive feedback that they are performing badly within the group (low score/rank) should be more likely to use social information (high s). For Season 1 (discounting Hunt 6, see above), there was a significant and negative correlation between score and s ($r=-.13$, $p<.05$) and between rank and s ($r_s=-.19$, $p<.005$). Season 2 also showed significant (and stronger) negative correlations between score and s ($r=-.29$, $p<.001$) and between rank and s ($r_s=-.26$, $p<.001$). The fact that Season 2 correlations were higher than Season 1 correlations might explain why periodic cultural learning more effectively outperformed individual controls during Season 2 than Season 1, and matched terminal cultural learning in the former but not the latter. In summary, the advantage of periodic cultural learners over individual controls shown in Fig. 3B might be attributed to a combination of two cultural learning strategies described by Laland (2004): "copy successful individuals" and "copy when behaviour is unproductive", yielding the composite strategy, "copy successful individuals when behaviour is unproductive".

5.3. Do informational access costs reduce the frequency of cultural learning?

Fig. 5A shows the effect of access costs on mean score during Season 3, with individual learners from Season 3 and cultural learners from Season 2 (who could not set access costs) for comparison. There was no significant difference between the access-cost-cultural-learners and individual learners during the final five hunts of Season 3 [$F(1,55)=0.24$, ns]. There was also no difference between the mean cumulative scores of the individual controls and access-cost-cultural-learners [$t(55)=0.73$, ns]. The Season 2 no-access-cost-cultural-learners, by contrast, significantly outperformed the Season 3 access-cost-cultural-learners during the final five hunts [$F(1,39)=15.59$, $p<.001$].

Season 3 cultural learners set an average access cost of 664.84 calories for others to access their designs, although this varied widely across participants (SD=2410.18; range, 0–23,000). These access costs had the effect of reducing the frequency of cultural learning (Fig. 4A). On average, participants viewed another players' design 1.35 times (out of a maximum of 5) during Season 3, compared to 3.75 for Season 1 [significantly higher, paired samples *t* test: $t(39)=8.39$, $p<.001$] and 3.14 during Season 2 [significantly higher: $t(39)=8.23$, $p<.001$]. However, when other players *were* viewed, they were more likely to be copied, as shown by higher values of s during Season 3 in Fig. 4B. Within-groups ANOVAs showed that s was significantly higher during Season 3 than during Season 1 [$F(1,19)=8.80$, $p<.01$] and Season 2 [$F(1,19)=9.21$, $p<.01$]. Hence, participants were more likely to use information that they paid for compared to information they received for free.

Fig. 4A also shows that participants in groups were much less likely to view the most successful group member's design during Season 3 than during Seasons 1 and 2. Only 36% of participants who viewed chose to view the most successful group member during Season 3 (compared to 76.8% for Season 1 and 81.4% for Season 2). Instead, most participants (44%) chose to copy the *second* best group member. These choices were significantly different from chance [$\chi^2(5)=51.28$, $p<.001$]. The reason that participants switched from copying the most successful group member to copying the second best is because the most successful group member set much higher access costs than other group members (Fig. 5B). The top-ranked player in each group set access costs that were on average 10 times higher than the access costs of lower-ranked group members, and high-ranked players generally set higher access costs (group rank significantly and positively correlated with access cost: $r_s=.39$, $p<.001$). High-ranked participants also received higher payments as a result of being viewed, such that payment received positively correlated with group rank ($r_s=.39$, $p<.001$), and high-ranked players tended to pay less to view other players' designs, such that group rank negatively correlated with price paid ($r_s=-.19$, $p<.01$).

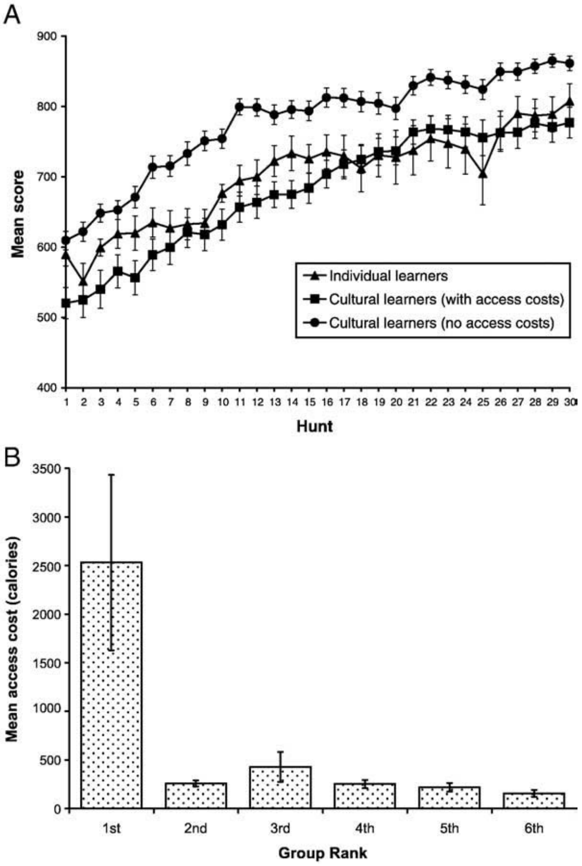

Fig. 5. (A) Mean scores for periodic cultural learners during Season 3, who could set access costs, periodic cultural learners during Season 2, who could not set access costs, and individual learners during Season 3. (B) Mean access costs set by participants during all cultural learning opportunities of Season 3. Error bars show standard error.

5.4. Additional analyses

Additional analyses can be found in the Electronic Supplementary Material, which show that (a) participants modified continuous attributes more than discrete attributes but were not sensitive to different fitness weights of different continuous attributes; (b) there was some evidence that participants kept their values from the previous season when starting a new season despite being told that optimal values change between seasons, but this carry-over quickly disappeared as participants learned in the new environment; (c) most participants reached equilibrium within the allotted 30 hunts; and (d) there were no batch effects between the present study and that of Mesoudi and O'Brien (2008). None of these additional findings were found to affect the conclusions presented above.

6. Discussion

To briefly summarise the results, it was found that (a) individual learning was significantly more adaptive in unimodal adaptive landscapes than in multimodal adaptive landscapes, such that the advantage of the copy-successful-individuals cultural learning strategy found by Mesoudi and O'Brien (2008) disappears; (b) periodic cultural learning, featuring regular cultural learning opportunities throughout the experiment, was (in Season 2) significantly more adaptive than individual learning, and performed equally as well as the terminal cultural learning in which cultural learning was permitted only on the final few trials; (c) allowing participants to set informational access costs that others must pay to view their designs significantly reduced the frequency of cultural learning and especially the use of the copy-successful-individuals strategy, turning cultural learners effectively into individual learners. The following sections discuss each of these findings in turn.

6.1. The shape of the adaptive landscape can affect cultural evolution

Individual learning was found to be significantly more adaptive in unimodal adaptive landscapes, in which there is a single optimal arrowhead design, than in multimodal adaptive landscapes, in which there are multiple locally optimal arrowhead designs of different fitness. In unimodal environments, simple reinforcement learning will always lead to the best possible arrowhead design from any starting point. [A reinforcement learning algorithm might be "increase/decrease an attribute by a certain amount; if your score increases then increase/decrease that attribute further; if your score decreases then decrease/increase the attribute", which was shown by Mesoudi and O'Brien (in press) to generate comparable results to the individual learners in Mesoudi and O'Brien (2008).] In multimodal environments, by contrast, learners can get stuck on locally optimal but globally suboptimal peaks, reducing the mean fitness of the population. Copy-successful-individuals cultural learning allows individuals to jump from locally optimal peaks found via individual learning to the globally optimal peak found by a more successful member of the population. However, if the environment is unimodal, then this advantage of cultural learning is eliminated.

Consistent with this, the unimodal individual learners of the present study achieved mean fitness equivalent to the multimodal cultural learners of Mesoudi and O'Brien (2008). Hence, the adaptive advantage of copy-successful-individuals cultural learning disappeared when individual learning was unimodal. It should be noted that these conditions are not strictly comparable because there was no unimodal cultural learning group. While this can be tested in future studies, it is highly unlikely that unimodal cultural learners would outperform unimodal individual learners. Models (Boyd & Richerson, 1995) and experiments (McElreath, et al., 2005; Mesoudi & O'Brien, 2008) both suggest that making individual learning easier, as it is in unimodal environments, makes cultural learning relatively less adaptive.

This finding is important because actual cultural evolution is likely to occur in multimodal adaptive landscapes (Boyd & Richerson, 1992; Mesoudi & O'Brien, 2008) in which there are several stable, locally optimal artifact designs or mental representations of varying fitness. Previous cultural evolution experiments (e.g., Kameda & Nakanishi, 2002, 2003; McElreath et al., 2005; Efferson et al., 2008) have used relatively simple dichotomous cultural traits (e.g., which of two crops to plant), one of which has higher fitness than the other. While simplicity is often desirable in order to draw tractable conclusions, a full understanding of cultural evolution will require experimental simulations using cultural traits with multiple continuous and discrete attributes and multimodal adaptive landscapes, as is possible with the virtual arrowhead task. Future studies might systematically vary aspects of the multimodal fitness environment, such as the number of peaks, the difference in height of the different peaks, or non-linear fitness relationships between attributes.

6.2. Periodic cultural learning performs just as well as terminal cultural learning

It was predicted that periodic cultural learning, where cultural learning is possible at regular intervals throughout the experiment, would be less adaptive relative to individual learning than terminal cultural learning, where cultural learning is allowed only in the last few hunts. Terminal cultural learning has an enforced period of individual learning during which the fitness environment can be effectively explored by all participants in the group, with a greater probability that a high-fitness peak is found. Periodic cultural learning, on the other hand, does not have this extended period of individual learning. Information scroungers who always engage in cultural learning can potentially free-ride on the individual learning of others throughout the entire season, resulting in fewer individuals exploring the fitness environment and lower group fitness.

While results from Season 1 supported this hypothesis to some degree, results from Season 2 did not. The latter might be considered more representative of our participants' actual decision-making given that during Season 1 participants may have been learning how to play the game rather than how to use the available information effectively. In Season 2, periodic cultural learners significantly outperformed individual controls and achieved comparable performance as terminal cultural learners. It was argued that this was because participants did not separate into pure individual learners (information producers) and pure cultural learners (information scroungers) as expected. Instead, they flexibly switched between individual learning (when they were performing well) and cultural learning (when they were performing poorly), therefore following a "copy-successful-individuals-when-behaviour-is-unproductive" cultural learning strategy. This replicates an analytical finding of Boyd and Richerson (1995) and an experimental finding of Kameda and

Nakanishi (2003) that populations of flexible learners outperform both populations of pure individual learners and mixed populations of pure individual learners and pure cultural learners. The fact that this was observed in both Kameda and Nakanishi's study and the present study despite differences in task, participants, and cultural learning strategies suggests that flexible switching between cultural and individual learning may be a general and adaptive feature of human cognition, as suggested by Boyd and Richerson's model.

6.3. Informational access costs significantly reduced the use of cultural learning

During Season 3 cultural learners could set access costs that other participants had to pay to view their designs. Successful participants tended to set very high access costs and prevented other group members from copying their designs, significantly reducing the frequency of cultural learning and especially the copy-successful-individuals strategy. Consequently, Season 3 access-cost-cultural-learners effectively became individual learners and performed no better than actual Season 3 individual controls, and significantly worse than Season 2 cultural learners who could not set access costs.

There was no evidence for the alternative prediction that successful players might set access costs that are high enough to tempt several other players to buy their information, thus increasing both the successful player's score (via the multiple fees received) and the other players' scores (via the better quality information). This secondary prediction was somewhat informal, however, and formal game theoretic analyses are needed to identify equilibrium access costs and access frequencies. Perhaps if the experiment were extended over further seasons of hunting, then market dynamics would converge on equilibria featuring more modest access costs. On the other hand, perhaps the successful participants here were behaving optimally in blocking all cultural learning because the lack of environmental change meant that successful participants could protect their advantage knowing that their valuable information will remain valuable in the future. In contrast, fluctuating environments would mean that the valuable information possessed by successful players might not remain valuable in the future, providing an incentive to sell it before it becomes outdated. Extending the number of trials and introducing environmental fluctuation might therefore generate fluctuating and unpredictable information markets. We might also, under these alternative conditions, observe a form of indirect bias (Boyd & Richerson, 1985) in which a participant mistakenly buys information from a less successful model. This less successful model's score would increase due to the fee received, which would erroneously signal to other participants that this model's information is worth buying. The other participants would buy the model's information, further increasing the model's score, thus generating a maladaptive information cascade. Finally, altering the task motivation from individuals competing against one another to a more cooperative situation might also affect access costs. For example, having separate groups of participants compete against one another might result in low access costs for in-group participants and high access costs for out-group participants. Hence, cultural learning would occur predominantly within groups, generating low within-group variation and high between-group variation and therefore providing the preconditions for cultural group selection (Richerson & Boyd, 2005).

The ability to set access costs is probably a more realistic assumption and better reflects actual cultural transmission than allowing free access to information. Ethnographic evidence suggests that successful individuals in traditional societies are paid in deference or material goods for access to their expert knowledge (Henrich & Gil White, 2001). Similar phenomena occur in industrial societies; for example, successful lawyers set higher prices for access to their legal knowledge than do less successful lawyers. These cases suggest that real-life access costs are not so high as to block all cultural learning, as was found here, although as noted above there are likely to be many differences between real situations and the experimental task (e.g., environmental fluctuation). Future field studies might attempt to quantify these naturally occurring access costs and the conditions under which they occur, providing data that could inform future experiments.

6.4. Cumulative cultural evolution

As well as allowing tests of the adaptiveness of specific cultural learning strategies, the virtual arrowhead task also constitutes a simulation of cumulative cultural evolution (Boyd & Richerson, 1996; Tomasello, Kruger, & Ratner, 1993), in which innovations or modifications from different individuals are accumulated over successive generations. Although many species have been shown to exhibit geographically variable cultural traditions (e.g., Whiten et al., 1999), it is thought that only humans have the capacity for cumulative culture and that it is this property that underlies the unique complexity and diversity of human culture (Tomasello, 1999). In the present experiment, cultural learners pooled their individually acquired knowledge to produce artifacts that were, under certain conditions, functionally better than artifacts produced by individual controls, indicative of cumulative cultural evolution. However, it is possible that single individuals could, if they were lucky and skilful, find the globally optimal arrowhead design that gives the maximum 1000 calories. Future studies using the virtual arrowhead task might simulate the more stringent definition of cumulative cultural evolution in which the accumulated product of several individuals could not have been invented by a single individual alone, which is typical of most modern technology in industrial societies, such as cars or computers, and probably much of the technology of pre-industrial societies also (Henrich, 2008). For example, different participants or groups might be asked to combine separate artifacts (e.g., bows, arrows) to create a compound artifact (e.g., the bow and arrow) similar to the experimental simulations of Insko et al. (1983) of intergroup production and trade. Finally, adapting these cultural evolution experiments for use with non-human species, along the lines of recent cultural transmission experiments with chimpanzees (e.g., Whiten, Horner, & de Waal, 2005), and comparing the results with those seen for human participants might reveal the cognitive capacities that underlie cumulative culture.

6.5. Wider implications

In the real world, many individuals and organisations face problems similar to those faced by the participants in this experiment, i.e., whether to engage in costly trial-and-error learning or whether instead to adopt the practices of another individual or organisation. If the latter, then the question becomes who to copy (the most successful, the majority, etc.) and when to copy (when performing poorly, when uncertain, etc.). For example, businesses can either innovate by engaging in costly research and development or they can copy the practices of other more successful businesses. Governments can experiment with their policies or they can copy the successful policies of other governments. Consumers can personally test every available brand or they can buy the most popular brand. Each of these examples describe a different cultural learning strategy (copy successful individuals, copy successful behaviours, and copy the majority, respectively), each of which is likely to have different long- and short-term consequences for profits, welfare, happiness, etc. Cultural evolution experiments are useful in telling us not only what individuals should do in such situations, but also what people actually do. The present study suggests that a cultural learning strategy that intuitively seems effective—preferentially copying successful individuals—is only effective when environments have multiple optima of varying fitness and is only employed when access to others' information is free. And while some theoretical models suggest that cultural learning would be hampered by the emergence of free-riding information scroungers, the present study suggests that people avoid this by flexibly switching between individual and cultural learning, only copying others when they are doing poorly. Future mathematical models, experiments, and field studies will no doubt offer many more insights into processes of cultural evolution, both scientific insights into how people behave under different circumstances and practical insights into how individuals, businesses, and governments can improve their decision-making.

Acknowledgments

I am grateful to Mark Funnell and Julian Oldmeadow for organising access to participants, the students of Hills Road Sixth Form College, Cambridge, for participating in the study, and Tatsuya Kameda and an anonymous reviewer for valuable comments on the manuscript. This research was supported by a Mellon Foundation Fellowship held by the author.

Appendix A. Fitness equations

The fitness of an arrowhead, W ($0 \leq W \leq 1000$), was determined by the participant's entered attribute values of length, width, thickness, and shape (a fifth attribute, color, had no effect on fitness), according to Eq. (1):

$$W = 1000(0.275\ W_{\mathrm{L}} + 0.25\ W_{\mathrm{W}} + 0.35\ W_{\mathrm{T}} + 0.125\ W_{\mathrm{S}}) \tag{1}$$

where W_{L}, W_{W}, W_{T}, and W_{S} are the individual fitness contributions of length, width, thickness, and shape, respectively, each ranging from 0 to 1. Hence, these attributes are differentially weighted in descending order of importance of thickness, length, width, and shape. Participants received feedback based only on the single number, W, and were not informed of separate fitness contributions of the attributes. The three continuous attributes, length, width, and thickness, each had bimodal fitness functions (Fig. 1B) comprising two overlapping normal distributions with means O_1 and O_2, representing the global optimum and local optimum, respectively. The global optimum gave the maximum fitness contribution of the attribute (e.g., W_{L}=1), while the local optimum gave two-thirds that value (e.g., W_{L}=0.66). Each attribute had different optima, e.g., for length, O_{L1} and O_{L2}, for width O_{W1} and O_{W2} and so on. The fitness contribution from each of these optima (using length as an example), W_{L1} and W_{L2}, were given by Eqs. (2) and (3):

$$W_{\mathrm{L1}}(X_{\mathrm{L}}, O_{\mathrm{L1}}) = P_1 \exp\left[-(X_{\mathrm{L}}/100 - O_{\mathrm{L1}}/100)^2/(2\sigma)\right] \tag{2}$$

$$W_{\mathrm{L2}}(X_{\mathrm{L}}, O_{\mathrm{L2}}) = P_2 \exp\left[-(X_{\mathrm{L}}/100 - O_{\mathrm{L2}}/100)^2/(2\sigma)\right] \tag{3}$$

where X_{L} is the participant's entered value of length, P_1 and P_2 are the maximum fitness contributions of the two optima (P_1=1, P_2=0.66), and σ is the standard deviation of the normal distribution (σ=0.025). The overall fitness contribution of length, W_{L}, was the greater of the two values W_{L1} and W_{L2} according to Eq. (4):

$$W_{\mathrm{L}} = \begin{cases} W_{\mathrm{L1}} \text{ if } W_{\mathrm{L1}} > W_{\mathrm{L2}} \\ W_{\mathrm{L2}} \text{ if } W_{\mathrm{L2}} > W_{\mathrm{L1}} \end{cases} \tag{4}$$

Equivalent expressions for the fitness contributions of width and thickness are obtained by substituting different subscripts. The fitness contribution of shape was determined by ranking each of the four discrete shapes (Shape 1, Shape 2, Shape 3, and Shape 4) into a predetermined random order, with the highest ranked value giving a fitness contribution of

W_S=1, the second W_S=0.9, the third W_S=0.66 and the last W_S=0.33. The actual feedback shown to participants, W', was calculated by adding random error to W. Hence, W' was drawn randomly from a normal distribution with mean W and standard deviation ε, where ε=5. Individual controls in Season 1 experienced unimodal fitness functions, where fitness was given by a single normal distribution [i.e., Eqs. (1) and (2) only, ignoring Eqs. (3) and (4)].

Appendix B. Calculation of social influence, *s*

To obtain the measure of social influence, *s*, two quantities were calculated for each continuous attribute. First, the absolute difference between the viewer's attribute value and the model's attribute value immediately before cultural learning, representing deviation from the model before cultural learning. Second, the absolute difference between the viewer's value after cultural learning and the model's value before cultural learning (which was the value that was seen by the viewer), representing deviation from the model after cultural learning. The second difference was subtracted from the first difference, giving the degree to which the deviation between viewer and model decreased as a result of cultural learning. This measure was calculated separately for each of the three continuous attributes, and these three values were summed to give the overall social influence, *s*. Because they were not represented on continuous scales and were not modified as frequently as the continuous attributes, shape and color were not used to calculate *s*. The within-participant consistency of *s* was calculated as follows (based on Kameda & Nakanishi, 2002, pp. 384–385). First, *s* was categorised according to its sign: positive, zero, or negative. Then, for each participant (and separately for each season), the number of times that the participant kept the same sign of *s* in consecutive cultural learning hunts was counted. For example, if a participant had values of *s* of 22, 0, −12, 11, and 18 for Hunts 6, 11, 16, 21, and 26, respectively, then they would get a consistency score of 1 (for the final two hunts of consecutive positive *s* values). The consistency index therefore ranged from 0 (entirely inconsistent, e.g., *s*={−6, 0, 23, 0, 7}) to 4 (entirely consistent, e.g., *s*={0, 0, 0, 0, 0} or *s*={21, 8, 16, 33, 5}).

Appendix C. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.evolhumbehav.2008.04.005.

References

Arnold, S. J., Pfrender, M. E., & Jones, A. G. (2001). The adaptive landscape as a conceptual bridge between micro- and macroevolution. *Genetica*, 112–113, 9–32.

Baum, W. M., Richerson, P. J., Efferson, C. M., & Paciotti, B. M. (2004). Cultural evolution in laboratory microsocieties including traditions of rule giving and rule following. *Evolution and Human Behavior*, *25*, 305−326.

Bettinger, R. L., & Eerkens, J. (1999). Point typologies, cultural transmission, and the spread of bow-and-arrow technology in the prehistoric Great Basin. *American Antiquity*, *64*, 231−242.

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. Chicago: University of Chicago Press.

Boyd, R., & Richerson, P. J. (1992). How microevolutionary processes give rise to history. In M. Nitecki, & D. V. Nitecki (Eds.), *History and Evolution* (pp. 178−209). Albany: State University of New York Press.

Boyd, R., & Richerson, P. J. (1995). Why does culture increase human adaptability? *Ethology and Sociobiology*, *16*, 125−143.

Boyd, R., & Richerson, P. J. (1996). Why culture is common, but cultural evolution is rare. *Proceedings of the British Academy*, *88*, 73−93.

Boyd, R., & Richerson, P. J. (2005). *The origin and evolution of cultures*. Oxford: Oxford University Press.

Caldwell, C. A., & Millen, A. E. (2008). Experimental models for testing hypotheses about cumulative cultural evolution. *Evolution and Human Behavior*, *29*, 165−171.

Cavalli-Sforza, L. L., & Feldman, M. W. (1981). *Cultural transmission and evolution*. Princeton: Princeton University Press.

Efferson, C., Lalive, R., Richerson, P. J., McElreath, R., & Lubell, M. (2008). Conformists and mavericks: The empirics of frequency-dependent cultural transmission. *Evolution and Human Behavior*, *29*, 56−64.

Henrich, J. (2008). A cultural species. In M. J. Brown (Ed.), *Explaining culture scientifically*. Seattle: University of Washington Press.

Henrich, J., & Boyd, R. (1998). The evolution of conformist transmission and the emergence of between-group differences. *Evolution and Human Behavior*, *19*, 215−241.

Henrich, J., & Gil White, F. J. (2001). The evolution of prestige. *Evolution and Human Behavior*, *22*, 165−196.

Henrich, J., & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology*, *12*, 123−135.

Insko, C. A., Gilmore, R., Drenan, S., Lipsitz, A., Moehle, D., & Thibaut, J. W. (1983). Trade versus expropriation in open groups. *Journal of Personality and Social Psychology*, *44*, 977−999.

Kameda, T., & Nakanishi, D. (2002). Cost–benefit analysis of social/cultural learning in a nonstationary uncertain environment. *Evolution and Human Behavior*, *23*, 373−393.

Kameda, T., & Nakanishi, D. (2003). Does social/cultural learning increase human adaptability? *Evolution and Human Behavior*, *24*, 242−260.

Laland, K. N. (2004). Social learning strategies. *Learning and Behavior*, *32*, 4−14.

Laland, K. N., & Brown, G. R. (2002). *Sense and nonsense*. Oxford: Oxford University Press.

McElreath, R., Lubell, M., Richerson, P. J., Waring, T. M., Baum, W., Edsten, E., Efferson, C., & Paciotti, B. (2005). Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior*, *26*, 483−508.

Mesoudi, A. (2007). Using the methods of social psychology to study cultural evolution. *Journal of Social, Evolutionary and Cultural Psychology*, *1*, 35−58.

Mesoudi, A., & O'Brien, M. J. (2008). The cultural transmission of Great Basin projectile point technology I: An experimental simulation. *American Antiquity*, *73*, 3−28.

Mesoudi, A., & O'Brien, M. J. (in press). The cultural transmission of Great Basin projectile point technology II: An agent-based computer simulation. *American Antiquity*.

Mesoudi, A., Whiten, A., & Dunbar, R. I. M. (2006). A bias for social information in human cultural transmission. *British Journal of Psychology*, *97*, 405−423.

Mesoudi, A., Whiten, A., & Laland, K. N. (2004). Is human cultural evolution Darwinian? Evidence reviewed from the perspective of The Origin of Species. *Evolution*, *58*, 1−11.

Mesoudi, A., Whiten, A., & Laland, K. N. (2006). Towards a unified science of cultural evolution. *Behavioral and Brain Sciences*, *29*, 329−383.

Richerson, P. J., & Boyd, R. (2005). *Not by genes alone*. Chicago: University of Chicago Press.

Rogers, A. (1988). Does biology constrain culture? *American Anthropologist*, *90*, 819−831.

Tomasello, M. (1999). *The cultural origins of human cognition*. Cambridge, MA: Harvard University Press.

Tomasello, M., Kruger, A. C., & Ratner, H. H. (1993). Cultural learning. *Behavior and Brain Sciences*, *16*, 495−552.

Whiten, A., Goodall, J., McGrew, W. C., Nishida, T., Reynolds, V., Sugiyama, Y., Tutin, C., Wrangham, R. W., & Boesch, C. (1999). Cultures in chimpanzees. *Nature*, *399*, 682−685.

Whiten, A., Horner, V., & de Waal, F. B. M. (2005). Conformity to cultural norms of tool use in chimpanzees. *Nature*, *437*, 737−740.

Wright, S. (1932). The roles of mutation, inbreeding, crossbreeding and selection in evolution. *Proceedings of the Sixth International Congress of Genetics*, *1*, 356−366.