Original Article

Humans copy rapidly increasing choices in a multiarmed bandit problem

Ulf Toelch[a,b,*], Matthew J. Bruce[b], Marius T.H. Meeus[c], Simon M. Reader[b,*]

[a]Innovation and Environmental Sciences, Utrecht University, Heidelberglaan 2, P.O. Box 80115, 3508 TC Utrecht, The Netherlands
[b]Behavioural Biology, Department of Biology and Helmholtz Institute, Utrecht University, The Netherlands
[c]Department of Organization Studies, Center for Innovation Research, Tilburg University, The Netherlands

Abstract

Conformist social learning, the tendency to acquire the most common trait in a group, allows individuals to rapidly acquire established beneficial traits from a multitude of options. However, conformist strategies hinder acquisition of novel advantageous behavior patterns, because such innovations are by definition uncommon. This raises the possibility that proxy cues of the success of novel traits may be utilized to identify and acquire advantageous innovations and disregard failing options. We show that humans use changes in trait frequency over time as such a cue in an economic game. Participants played a three-alternative forced choice game (i.e., a multi-armed bandit), using social information to attempt to locate a high reward that could change location. Participants viewed temporal changes in how many players chose each option in two successive rounds. Participants supplemented conformist strategies with a "copy-increasing-traits" strategy. That is, regardless of the traits absolute population frequencies, participants' choices were guided by changes in trait frequencies. Thus, humans can detect advantageous innovations by monitoring how many individuals adopt these over time, adopting traits increasing in frequency, and abandoning traits decreasing in frequency. Copying rapidly increasing traits allows identification and acquisition of advantageous innovations, and is thus potentially key in facilitating their early diffusion and cultural evolution.
© 2010 Elsevier Inc. All rights reserved.

Keywords: Innovation; Social learning; Conformist transmission; Cultural evolution

1. Introduction

Cumulative cultural evolution describes a process, apparently unique to humans, in which knowledge is accumulated and socially transmitted with the results exceeding individual lifetime achievements. Human culture rests largely upon the capabilities to produce novel behavior patterns, products, and ideas (henceforth "traits") and to learn such traits from others (Boyd & Richerson, 1985; Cavalli-Sforza & Feldman, 1981). Balancing such innovation and social learning is a crucial factor for the emergence of cumulative cultural evolution. However, the specific characteristics that enable humans to accumulate and build on existing knowledge are unknown. One problem is the detection of novel advantageous behavior in a sea of traits. In order to identify beneficial traits, individuals could employ a variety of transmission biases including direct, prestige, and conformist biases (Henrich, 2001; Laland, 2004). Alternatively, individuals could copy models at random or invent traits de novo (Bentley, Hahn, & Shennan, 2004; Reader & Laland, 2003).

Direct transmission biases are based on individual evaluation of observed behavioral variants, such as evaluation of the palatability of food items or efficiency of a food processing technique, but cannot be utilized when trait utility is difficult to judge (Henrich, 2001). Prestige biases result in copying of traits possessed by successful or prestigious individuals, biases that will be advantageous when the successful individuals' traits also lead to success for the copier, but that may result in the copying of inappropriate traits (Henrich & Gil-White, 2001; Offerman & Schotter, 2009). Moreover, identification of successful individuals may be difficult or time-consuming. Conformist biases involve the disproportionate copying of traits present

* Corresponding authors. Cognitive Neurobiology, Humboldt University, Dorotheenstrasse 94, 10117 Berlin, Germany.
  E-mail addresses: ulf.toelch@charite.de (U. Toelch), s.m.reader@gmail.com (S.M. Reader).

in a majority of the population. Conformist biases have the advantage that they can be employed even when information is not available or difficult to gather on individual success or trait characteristics. Since majority traits are likely to be advantageous, particularly in stable environments, conformist social learning will provide copiers with a suite of adaptive traits (Henrich & Boyd, 1998). Moreover, conformist transmission will increase similarities within groups and differences between groups (Boyd & Richerson, 1985; Efferson, Lalive, Richerson, McElreath, & Lubell, 2008; Henrich & Boyd, 1998). However, conformist transmission cannot result in the spread of novel beneficial innovations that are at a low frequency, and will even lead to the loss of such traits from groups (Eriksson, Enquist, & Ghirlanda, 2007). Without novel or modified traits being introduced into populations, cumulative cultural evolution cannot occur and will stagnate. Moreover, empirical studies fail or only partly show conformist transmission (Coultas, 2004), even under conditions where conformist transmission would be advantageous (Efferson et al., 2008). This has lead to considerable interest in social learning biases or strategies that will facilitate the spread of advantageous innovations while protecting against the spread of hazardous innovations (e.g., Laland, 2004; Schlag, 1999).

Conformist transmission is thus likely to be complemented by alternative learning strategies that allow for the acquisition of novel, low frequency traits. Analysis of child name frequency changes over time suggests adoption may be driven by both absolute name frequencies and temporal changes in name frequencies (Berger & Le Mens, 2009; Gureckis & Goldstone, 2009). Here we test experimentally whether changes in trait frequencies over time are perceived as salient cues for acquisition. Traits that rapidly increase in frequency indicate both acquisition feasibility and high value, and hence potentially indicate that acquisition would be advantageous. We used a three alternative forced choice paradigm (or multi-armed bandit) to examine whether humans employed such cues and how different social cues were integrated. Participants viewed the frequency of choices of a group of previous players at two points in time: an initial round and a follow-up round. To focus on the use of social cues, we provided participants no direct feedback on the payoffs that choices received. Thus, only social information was available, in the form of others' choices. We predicted that participants base their choices on both (1) absolute observed frequencies, i.e. employ conformist strategies, and (2) the change in frequencies over time, i.e., copy rapidly increasing traits. Here, we focus on the identification of individual strategies that impact on the diffusion of innovations. Previous studies document considerable differences between individuals in their social information use (Efferson et al., 2008; Toelch et al., 2009), raising the possibility that individuals will flexibly employ multiple learning strategies and that a mix of strategies will be observed within a population.

2. Methods

2.1. Participants

Participants were 23 adults [15 female, eight male, mean age (±S.D.): 24.4±8.1 years], undergraduate students recruited by posters. Participants were paid on average €5 (ca. $7 US) and were informed that their precise remuneration depended upon personal performance (with payment ranging from €3 to €7). Before beginning, participants gave informed written consent and received written instructions (see electronic supplementary material) with three multiple-choice questions to check understanding. If players answered any questions incorrectly the experimenter explained the unclear point to the participant and asked whether any questions remained. After the experiment, participants completed a brief questionnaire, were paid in private, and debriefed. Procedures were approved by the Universitair Medisch Centrum Medical Ethics Review Committee (protocol 06-672) and comply with American Psychological Association ethical guidelines and the principles expressed in the Declaration of Helsinki.

2.2. Procedure

Participants played a computer game for 25–30 minutes, choosing one of three available options (card decks) per round. The decks were labeled RED, BLUE, GREEN but were not colored. One deck yielded a higher reward than the other two (5–9 versus 2–6 points). The reward of the deck chosen by the player was revealed in Phase 1 of the game only (see below). Points were integer values with equal probability for each value to be drawn. The high value deck changed position once every five rounds, on average. For example, the deck yielding the highest rewards might change from the RED to GREEN deck. This change was not announced, but players had been instructed about the rate of change and that one deck was better than the other two. Thus players had to discover a change in the highest-yielding deck from the change in received rewards. If they discovered they were no longer choosing the highest-yielding deck they knew that the highest-yielding deck must now be one of the other two decks. Since rewards were variable, it was not always immediately obvious which deck was yielding the highest rewards. The exact variance of the decks was not revealed to players.

The game consisted of two phases (Fig. 1). In Phase 1 (60 rounds) individual information only was available: a message told players their points received on the previous round. In Phase 2, players received no information on points but, instead, saw the choices of 20 other players (henceforth "demonstrators") who had played Phase 1 of the game before. Demonstrator data were collated from a preliminary study. Demonstrators were undergraduate students who participated as part of a course, and behaved similarly to participants (see below). Participants thus made choices based on the real decisions of previous players from the same

Phase 1

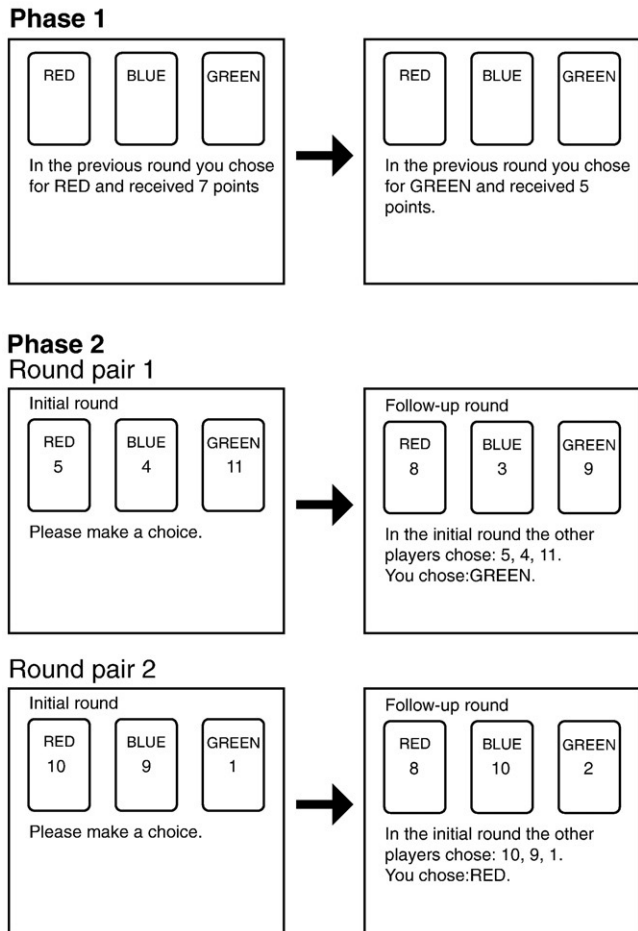## Phase 2
Round pair 1

Round pair 2

Fig. 1. Outline of experimental procedure and example of information presented to players. Arrows indicate sequence of screens. (A) Two rounds of Phase 1, where individual information only was presented in the form of points scored. Played for 60 rounds. (B) Two round-pairs of Phase 2, where social information only was presented in the form of the choices of other players that had previously played the same game. Played for 75 round pairs. In this example, on the first follow-up round (B: upper right panel) the green deck is chosen by most players (termed the *majority deck*) but the red deck increased in frequency fastest. See main text for a detailed description.

study population. We used this demonstrator data rather than an artificial dataset to avoid unwarranted deception of participants and to ensure that the demonstrator data represented a sample of human choice patterns. Participants were further informed that demonstrators played exactly the same game as they did in Phase 1, that they would see player decisions over two consecutive rounds, and the high-reward deck did not change between these two rounds. Participants viewed 75 such pairs of rounds, and made a choice on both rounds of each pair. For example, players might view 5, 4 and 11 players choosing Decks 1, 2 and 3, respectively, and then make a choice ("initial round"). In the "follow-up round," they saw the choices demonstrators had made after receiving feedback on their choices, e.g., 8, 3 and 9 players choosing Decks 1–3. In this example, most demonstrators are choosing Deck 3 but Deck 1 is increasing in frequency

fastest. On the follow-up round, participants also viewed information from the initial round, to eliminate any need to memorize the frequencies (Fig. 1). The deck with the highest number of demonstrator choices ("majority deck") changed between initial and follow-up round in one third of the rounds on Phase 2. This one third probability was based on the collated data of the demonstrators. Final payment depended on the final 30 decisions from Phase 1 and all decisions in Phase 2, and players knew this. We excluded the first 30 decisions to allow participants unconstrained exploration of the game.

2.3. Analysis

Phase 1 gave players experience with the game with individual information only. In Phase 2, participants had to decide, based on the frequency distributions of previous players, which of three options they would choose. We first tested whether players' choices deviated from random choice via 95% confidence intervals. We then examined switches between decks from the initial to follow-up round, since these allow us to distinguish between different frequency dependent strategies such as conformist transmission and copy increasing traits. Switches refer to whether players changed decks between rounds: note that the identity of the majority deck could also change between rounds, and thus players might switch decks but choose the majority deck on both rounds. A choice for the majority deck in both the initial and follow-up round would be termed a *conformist strategy*. A choice for the increasing and not the majority deck on the follow-up round would be termed a *copy-increasing-traits strategy*.

We examined switching behavior using a Generalized Linear Mixed Model (GLMM) with switching as a binary dependent variable (0=no switch, 1=switch) and player identity as a random effect affecting the intercept (Bolker et al., 2009). We assumed a beta binomial distribution (to correct for over-dispersion) with logit link function (Gelman & Hill, 2007). Independent variables were the frequency change in demonstrator choices for the deck players chose in the initial round, the demonstrator frequency change for the deck players chose in the follow-up round, and whether players chose a majority deck in the initial round. We incorporated player gender in exploratory analyses, but this variable was non-significant and was thus removed from subsequent analyses. All analyses were conducted in R 2.9.2 (R Core Development Team, 2009). The GLMM was calculated with the lmer function of the package lme4. We test effects for significance by examination of confidence intervals, since calculation of degrees of freedom and thus p values is problematic when accounting for the hierarchical structure of our data with GLMM (for discussion of this topic, see Bolker et al., 2009; Gelman & Hill, 2007). Since the perceived value of individual information could influence the use of social information, we tested for a correlation between the performance in Phase 1 and the

proportion of majority choices in the initial round during Phase 2.

3. Results

In the initial round, players chose the majority deck significantly more than chance (Table 1), although individuals differed considerably, with the range of choices for the majority deck 31−92%. There was no significant correlation between the number of correct choices players made in Phase 1 and the choices for the majority deck in the initial round (Kendall rank correlation, τ=0.12, n=23, p=.44). We divided the follow-up round data into four categories, split according to whether players had (1) chosen a majority deck on the initial round and (2) switched decks between initial and follow-up rounds. In three of the four categories, players significantly preferred the majority deck (Table 1): only players that chose a majority deck in the initial round and then switched decks did not show a significant preference for the majority deck. Thus, individuals had a strong preference to copy-the-majority but also made choices for non-majority decks.

We investigated why players made non-majority choices by examining influences on the decision to switch decks. Players were more likely to switch decks when decreasing numbers of demonstrators chose the players' initial choice, and increasing demonstrator numbers chose their follow-up choice (Table 2). Picking a majority deck in the initial round reduced the probability of switching decks. For example, when demonstrator frequencies did not change between the initial and follow-up round, a player that chose a non-majority versus majority deck on the initial round had a 45% versus 12% probability of switching decks [inverse-logit transformation of −0.2 (row 1 of Table 2)=0.45; inverse-logit of (−0.2)+(−1.83) (rows 1, 4 of Table 2)=0.12]. However, demonstrator frequency changes had stronger effects on switching probability when players picked a majority deck in

Table 1
Participants generally prefer the card deck chosen by the majority of demonstrators ("majority deck") over the other two decks

Behavior from initial to follow-up round	Choice on initial round	
	Majority deck	Non-majority deck
No switch	**88.9[a]±12.2***	**47.0[b]±20.4***
Switched decks	51.0[c]±20.4	**71.6[d]±18.4***
Percent switchers	39.2±19.9	45.2±20.3

Upper two rows: percent choices (means±95% CI) for the majority deck on follow-up round, divided according to choices on the initial round and whether participants switched decks between the initial and follow-up round. Participants chose the majority deck on 65.9±19.4% of occasions on the initial round, significantly more than expected by chance (33%). Significant preferences for the majority deck are asterisked and in bold. Superscripts indicate chance expectations based on random choice: a: 66%; b: 16.5%; c: 33%; d: 41.25% (66% chance that one remaining deck is the majority deck, 16.5% chance that the other is the majority deck: (0.5×0.66)+(0.5×0.165)=0.4125. Lower row (italics): percentage of occasions (means±95% CI) that participants switched decks.

Table 2
The tendency to switch deck choice between initial and follow-up rounds depended on the absolute number of demonstrator choices and the change in demonstrator choices that participants observed

Explanatory variable	Estimated coefficient±95% CI
Intercept	**−0.20±0.008***
Change in frequency initial choice (FI)	**−0.19±0.001***
Change in frequency follow-up choice (FF)	**0.32±0.002***
Majority deck initial round (M)	**−1.83±0.006***
Interaction FI:M	**−0.99±0.003***
Interaction FF:M	**0.67±0.003***

Positive coefficients indicate an increased likelihood that a player will change decks. Participants were thus less likely to switch when their initial choice was for the majority deck (M, row 4). Participants were more likely to switch with larger decreases in demonstrator choices for the participant's initial choice (FI, row 2). Participants were also more likely to switch with larger increases in demonstrator choices for the participant's follow-up choice (FF, row 3). Interaction effects (rows 5 and 6) are explained in the main text. Figures are coefficient estimates (logit transformed) and 95% CI from a GLMM. Significant factors (i.e., the CIs do not encompass zero) are asterisked and in bold.

the initial round than when they picked a non-majority deck. That is, there were interaction effects (Table 2), with a steeper relation between switching probability and demonstrator frequency change in the majority versus non-majority case, such that more switching was observed in the majority versus non-majority case with very large shifts in demonstrator frequency. For example, if two fewer demonstrators chose the player's initial choice and four more demonstrators chose the player's follow-up choice, a player that chose a non-majority versus majority deck on the initial round had a 81% versus 98% probability of switching decks [inverse-logit of (−0.2)+(−0.19*−2)+(0.32*4) (rows 1–3 of Table 2)=0.81]. In summary, player switching behavior was influenced by changes in the frequency of demonstrator choices.

There were considerable differences between players in their propensity to switch decks, estimated by the random effect of player identity on the intercept in our statistical model. The mean back-transformed (inverse logit) model intercept across players was 0.45 (±S.D.=0.2), ranging from .09, indicating a player that almost never switched decks, to .8, a player changing decks frequently. This raises the question of whether particular individuals pursued avoid-decreasing/copy-increasing-traits strategies or whether the observed effects were the result of occasional use of these behaviors across the population. To examine whether specific individuals utilized a copy-increasing-traits strategy, we analyzed the 16 rounds where ambiguity between options was highest, i.e., two decks had the same number of demonstrators (seven or eight) in the follow-up round. Participants consistently employing a copy-increasing-traits strategy were predicted to choose a deck with increasing numbers of demonstrators and not a deck that was a majority deck on the initial or follow-up round. We scored a player as employing a copy-increasing-traits strategy when at least 12 of 16 choices were for the increasing deck (the point at which the 95% CI of draws from a binomial distribution exceeds

random choice). By this conservative method, we identified five players using a copy-increasing-traits strategy. Thirteen players noted in a questionnaire after the experiment that monitoring frequency changes formed part of their strategy (Appendix 1). Player 19, for example, stated: "…*pick the deck the most people changed to in the follow-up trial*." These 12 players included those five players (Players 7, 10, 12, 14, 19) identified statistically as employing a copy-increasing-traits strategy.

Demonstrators were as successful as participants in Phase 1, indicating that both groups were similarly motivated and likely employed similar strategies (proportion correct choices in the last 30 rounds of Phase 1: demonstrators ($n=20$): mean±S.D.=0.38±0.19; participants ($n=23$): 0.44±0.13, Wilcoxon test; $W=164$, $p>.1$).

4. Discussion

Players made choices according to both the overall frequency of observed decisions and the change over time in these frequencies. Players exhibited a conformist bias on the initial round, their choices biased towards the option chosen by the majority of demonstrators, and they maintained this conformist bias when demonstrator choices were approximately constant. However, players were also nonconformists on occasions, choosing non-majority options, driven to change by decreases in demonstrators choosing their current deck and increasing demonstrator numbers choosing another deck. The strength of this effect was driven by the relative change in demonstrator frequency between rounds, indicating that participants deviated from copying the majority if the change between rounds was large enough. To our knowledge, this is the first experimental description of a "copy-increasing-traits" social learning strategy.

A copy-increasing-traits strategy would speed the diffusion of innovations, potentially increasing the rate of cultural evolution. Decreases in trait frequencies indicate a potentially outdated trait to avoid, while increases in trait frequencies indicate a potentially useful trait to acquire. Moreover, observed abandonment of a trait by others could sensitize individuals to novel possibilities, lowering the threshold to change behavior and increasing the rate of innovation, even when direct copying does not occur. The copy-increasing-traits strategy complements the existing palette of social learning strategies (Boyd & Richerson, 1985; Laland, 2004) and, like prestige and conformist biases (Boyd & Richerson, 1985; Henrich & Boyd, 1998), would lead to directed, nonrandom spread of innovations.

The rate of increase may allow individuals to estimate the potential advantages of a novel trait. Steep increases in frequency can indicate that a trait is highly beneficial, or rapid adoption is necessary, although such patterns may also indicate a short-lived fad that is less likely to persist (Berger & Le Mens, 2009). Copy-increasing strategies will have the strongest utility when traits are easily perceivable, i.e., traits exhibited frequently in public or that are unusually visible. The strategy requires that changes in trait frequencies be monitored over time, raising the possibility that such changes are salient to humans and easily remembered. However, this memory requirement could be lifted in certain circumstances, for instance when information is available on past behavior or when trait characteristics like appearance or proficiency indicate how long a trait has been possessed. For example, food debris could indicate past dietary choices, or newly made tools could be visually distinctive from old ones. Fast-rising traits are explicitly marked in some contexts, such as music sales charts (e.g., the UK singles chart), stock markets (e.g., top risers and fallers are listed in the FTSE 100), and internet radio and video sharing sites (e.g., www.youtube.com, www.last.fm). Stock investors preferentially buy so-called "attention-grabbing" stocks (Barber, Odean, & Zhu, 2009). A copy-increasing-traits strategy may be particularly valuable in contexts with many options but few beneficial traits, limiting the possibility for individual discovery of beneficial traits because of, for example, limited time for individual exploration (Salganik, Dodds, & Watts, 2006). The strategy can be viewed as an early-detection mechanism for beneficial traits, thus complementing conformist biases: rapidly increasing traits that continue their rise will become a majority trait. Moreover, a copy-increasing strategy could ease decisions between two traits at a similarly high population frequency. In such ambiguous cases, at least five players in our experiment used the increasing frequency cue to decide between two equivalent majority options. The strategy also complements prestige biases since a trait's speed of increase can indicate how easy it is to adopt. Demonstrator success alone might be a misleading cue for copying since demonstrator traits may be difficult to adopt.

While here we demonstrate a copy-increasing-traits strategy over a short timescale, and the examples cited above of marked increasing traits also influence relatively short-term decisions, it is an open question to what extent the copy-increasing-traits strategy is employed in longer-term decisions. In realms of human social learning such as the adoption of farming technologies and crops, hunting equipment, or symbolic traits, changes in trait frequency over time will be visible and thus a possible cue for adoption (Dow, Reed, & Olewiler, 2009; Mesoudi & O'Brien, 2008; Rogers & Ehrlich, 2008; Rogers, 1995). Like conformist transmission, a copy-increasing strategy will result in accelerating rate of adoption over time. However, we note that transmission dynamics alone may be insufficient to distinguish between conformist and copy increasing strategies, necessitating a focus on the cues utilized in social learning (for discussion see Franz & Nunn, 2009; Henrich, 2001; Laland, Richerson, & Boyd, 1996; Lefebvre, 1995; Reader, 2004; Rogers, 1995). A copy-increasing-traits strategy may be particularly advantageous under circumstances favoring the rapid adoption of innovations, such as when groups are exposed to novel environments through migration, environmental change, or niche construction

(Boyd & Richerson, 1985; Lefebvre, Reader & Sol, 2004; Sol, Bacher, Reader, & Lefebvre, 2008).

Copying increasing traits risks acquisition of a trait only tested by a minority of the population, and thus may be a flexibly employed strategy or a strategy associated with particular classes of individuals. Clearly, if all individuals consistently and exclusively employed copy-increasing strategies, no information on trait utility would be available. The strategy could thus accelerate both adaptive and maladaptive informational cascades (Bikhchandani, Hirshleifer & Welch, 1992). Further, traits acquired via this strategy could be subject to a higher probability that a trait is lost ("low stickiness"), and be replaced more quickly than traits acquired by other methods such as direct teaching. That is, increasing trait frequency may simply be used as a tip-off that a trait should be investigated as a possibility for long-term adoption. We observed considerable differences between players, suggesting that players employed different strategies, a finding backed up by the players' self-reports. Moreover, the performance of players on Phase 1 did not predict how often players will copy the majority. This underlines the finding from previous studies (McElreath et al., 2005; Toelch et al., 2009) that humans flexibly employ a range of strategies, with individual differences alongside context and performance cues important determinants of the strategy employed. It is unlikely that a single strategy can explain the diffusion of innovations and cumulative cultural evolution. Instead, a mix of strategies that can change over time, including individual assessment of trait outcomes, will determine the outcome of such processes. Identifying different learning and assessment strategies and the circumstances under which they are employed, and integrating these findings with other factors such as motivation, ecology, and demography, will be crucial for our understanding of cumulative cultural evolution (Dow et al., 2009; Kirby, Cornish & Smith, 2008; Powell, Shennan & Thomas, 2009).

Acknowledgments

Appendix 1: Players' verbatim descriptions of their strategies in phase 2 of the game. Italics (our emphasis) mark passages that describe copy-increasing or avoid-decreasing traits strategies

Player Description

1. Stay risk neutral, just choose the middle number of lastround.

2. I assumed that the decision of other players where based on the same strategy as mine for the first part *because there were rounds when you could see the number of selections increasing for a certain deck. So I guided myself after this selections but not totally.* I often changed selection of the decks even if another one would have a high number of selections.

3. I didn't use any strategy, played how i wanted it.

4. Here I stuck at the last chosen deck, the one with a high percentage of the players. Once again *I allowed a big drop only once,* sticking with the deck. On a second time I switched to the deck that seemed to attract the most players two rounds in a row. When in doubt, I remained at the former deck.

5. I kept exactly the same as the first phase. I chose decks different from the "suggested" ones in some occasions because they would just click in a row in a single deck. I chose to hop around hitting at most twice in each deck.

6. I tried to picked up the decks with high score.

7. Initial round: just random chosen colors *Follow up round: looking how many people changed to another one, most times I choose the color where most people changed to.*

8. Watch what the other players did, then with those numbers try to decide which dek was the higher one.

9. No strategy...

10. In the beginning i just chose what most people chose, but then i decided that was stupid so I thought of a good strategy. In the initial round i chose what most people had chosen, but then *in the follow up round i chose the deck that gained people.* So if in the initial round the deck was chosen 10 times and in the follow up round it was chosen 14 times i think it was a good deck because they decided to chose that deck again, so i chose it too.

11. I started clicking on the deck where the most people had clicked on. I kept clicking on that deck for a couple of times. Then I moved to another deck, the one where the most people clicked on and staid with that deck for a couple of times.

12. In the initial round I chose mostly the deck that most people chose. *In the follow up round I chose the deck in which the number was higher than it was in the initial round.*

13. In de second part of the game after a while I thought that if in the *Last round the number of choices by an other player became higher than the number of choices by players in the Initial round.* That is the best choice because a player don't choose a deck when the point of the cards became lower.

14. At the initial part I picked the deck with a medium amount of people that choose the deck, *in the follow-up round I picked the deck which had the highest difference in peple choosing it.* So if 11

people choose something in the first round and 14 in de second I choose that deck. If the difference of two decks was the same I went for the deck with the highest number.

15. More with knowledge, but also a little bit just "gambling".

16. Following the pattern of people. If high numbers stayed high I assumed these decks were scoring high and people were staying. *When people left their decks I assumed they were scoring low.*

17. I was choosing the deck wich have been chosen by the most players. The high number it influenced me.

18. Most of the time, I choose the deck with the highest decision rate of other players, especially if it was high in the follow up round. I tried to figure out wich deck was the best deck, *and kept watching changes in decision rates of the other decks to find out when the best deck changes position.*

19. Pick just any deck at the first trial (in my case the green one usually, because green is appealing for some obscure reason and I didn´t keep track of how many times the best deck would switch location). *Then pick the deck the most people changed to in the follow-up trial*, or just the deck with the majority of people (which weren´t always the same, in which case I randomly picked one of the two).

20. Follow the largest group in the initial rounds, but try not to switch decks to much and *try to find a group that does not shrink to much on every follow-up round.*

21. Did not think, just did it.

22. I always choose the color with the highest rate. *After that, I watched how many people stayed or switched.* When people stayed at a color that has a high rate, it means that a lot of people thought that they choose the pair with the highest rate. *If people switched, I watched how many people swithed.* If 5 out of 7 people (on one color) switch, that indicated a lower score then when 5 of of 10 people switch. This strategy makes that I usually choose the color with the highest rate.

23. This one felt a bit more tricky because I didn´t know how many points I was recieving with each round. In the initial round I would choose randomly. *For the second round, I based my choice off of how many people appeared to change.* Ex. If the initial round had 6,5,9 and the second round had 5,9,6 then I would choose the deck in round 2 containing 9 choices because I assumed no one had changed their choice from that deck, while one person from the left deck and 3 people from the right deck did change their choice. So basically I based my decision off of which deck appeared to have the fewest people changing their minds because if all the participants stayed on this deck it obviously had a good value.

Appendix A. Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.evolhumbehav. 2010.03.002.

References

Barber, B. M., Odean, T., & Zhu, N. (2009). Systematic noise. *Journal of Financial Markets, 12*, 547–569, doi:10.1016/j.finmar.2009.03.003.

Bentley, R. A., Hahn, M. W., & Shennan, S. J. (2004). Random drift and culture change. *Proceedings of the Royal Society B: Biological Sciences, 271*, 1443–1450, doi:10.1098/rspb.2004.2746.

Berger, J., & Le Mens, G. (2009). How adoption speed affects the abandonment of cultural tastes. *Proceedings of the National Academy of Sciences of the United States, 106*, 8146–8150, doi:10.1073/pnas.0812647106.

Bikhchandani, S., Hirshleifer, D., & Welch, I. (1992). A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy, 100*, 992–1026.

Bolker, B. M., Brooks, M. E., Clark, C. J., Geange, S. W., Poulsen, J. R., & Stevens, M. H. H., et al. (2009). Generalized linear mixed models: a practical guide for ecology and evolution. *Trends in Ecology & Evolution, 24*, 127–135, doi:10.1016/j.tree.2008.10.008.

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

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

Coultas, J. (2004). When in Rome... An evolutionary perspective on conformity. *Group Processes & Intergroup Relations, 7*, 317–331, doi:10.1177/1368430204046141.

Dow, G. K., Reed, C. G., & Olewiler, N. (2009). Climate reversals and the transition to agriculture. *Journal of Economic Growth, 14*, 27–53.

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.

Eriksson, K., Enquist, M., & Ghirlanda, S. (2007). Critical points in current theory of conformist social learning. *Journal of Evolutionary Psychology, 5*, 67–87, doi:10.1556/JEP.2007.1009.

Franz, M., & Nunn, C. L. (2009). Network-based diffusion analysis: A new method for detecting social learning. *Proceedings of the Royal Society B: Biological Sciences, 276*, 1829–1836.

Gelman, A., & Hill, J. (2007). *Data analysis using regression and multilevel/hierarchical models.* Cambridge University Press.

Gureckis, T. M., & Goldstone, R. L. (2009). How you named your child: Understanding the relationship between individual decision making and collective outcomes. *Topics in Cognitive Science, 1*, 651–674, doi:10.1111/j.1756-8765.2009.01046.x.

Henrich, J., & Gil-White, F. (2001). The evolution of prestige—Freely conferred deference as a mechanism for enhancing the benefits of cultural transmission. *Evolution and Human Behavior, 22*, 165–196.

Henrich, J. (2001). Cultural transmission and the diffusion of innovations: Adoption dynamics indicate that biased cultural transmission is the predominate force in behavioral change. *American Anthropologist, 103*, 992–1013.

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

Kirby, S., Cornish, H., & Smith, K. (2008). Cumulative cultural evolution in the laboratory: An experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences, 105*, 10681–10686, doi:10.1073/pnas.0707835105.

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

Laland, K. N., Richerson, P. J., & Boyd, R. (1996). Developing a theory of animal social learning. In: C. M. Heyes, & B. G. Jr Galef (Eds.), *Social Learning in Animals: The Roots of Culture* (pp. 129–154). London: Academic Press.

Lefebvre, L. (1995). The opening of milk-bottles by birds: Evidence for accelerating learning rates, but against the wave-of-advance model of cultural transmission. *Behavioural Processes, 34*, 43–53.

Lefebvre, L., Reader, S. M., & Sol, D. (2004). Brains, innovations and evolution in birds and primates. *Brain, Behavior and Evolution, 63*, 233–246.

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

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.

Offerman, T., & Schotter, A. (2009). Imitation and luck: An experimental study on social sampling. *Games and Economic Behavior, 65*, 461–502.

Powell, A., Shennan, S., & Thomas, M. G. (2009). Late pleistocene demography and the appearance of modern human behavior. *Science, 324*, 1298–1301, doi:10.1126/science.1170165.

R Development Core Team. (2009). R: A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing.

Reader, S. M., & Laland, K. N. (2003). *Animal innovation*. Oxford: Oxford University Press.

Reader, S. M. (2004). Distinguishing social and asocial learning using diffusion dynamics. *Learning and Behavior, 32*, 90–104.

Rogers, D. S., & Ehrlich, P. R. (2008). Natural selection and cultural rates of change. *Proceedings of the National Academy of Sciences USA, 105*, 3416–3420.

Rogers, E. M. (1995). *Diffusion of innovations*. 4th edn. New York: Free Press.

Salganik, M. J., Dodds, P. S., & Watts, D. J. (2006). Experimental study of inequality and unpredictability in an artificial cultural Market. *Science, 311*, 854–856, doi:10.1126/science.1121066.

Schlag, K. H. (1999). Which one should I imitate? *Journal of Mathematical Economics, 31*, 493–522.

Sol, D., Bacher, S., Reader, S. M., & Lefebvre, L. (2008). Brain size predicts the success of mammal species introduced into novel environments. *The American Naturalist, 172*, S63–S71.

Toelch, U., van Delft, M. J., Bruce, M. J., Donders, R., Meeus, M. T. H., & Reader, S. M. (2009). Decreased environmental variability induces a bias for social information use in humans. *Evolution and Human Behavior, 30*, 32–40, doi:10.1016/j.evolhumbehav.2008.07.003.