



Cognitive Science 44 (2020) e12903

© 2020 The Authors. *Cognitive Science* published by Wiley Periodicals LLC on behalf of Cognitive Science Society (CSS). All rights reserved.

ISSN: 1551-6709 online

DOI: 10.1111/cogs.12903

Inferring Behavior From Partial Social Information Plays Little or No Role in the Cultural Transmission of Adaptive Traits

Mark Atkinson, Kirsten H. Blakey, Christine A. Caldwell

Psychology, University of Stirling

Received 30 December 2019; received in revised form 8 June 2020; accepted 25 August 2020

Abstract

Many human cultural traits become increasingly beneficial as they are repeatedly transmitted, thanks to an accumulation of modifications made by successive generations. But how do later generations typically avoid modifications which revert traits to less beneficial forms already sampled and rejected by earlier generations? And how can later generations do so without direct exposure to their predecessors' behavior? One possibility is that learners are sensitive to cues of non-random production in others' behavior, and that particular variants (e.g., those containing structural regularities unlikely to occur spontaneously) have been produced deliberately and with some effort. If this non-random behavior is attributed to an informed strategy, then the learner may infer that apparent avoidance of certain possibilities indicates that these have already been sampled and rejected. This could potentially prevent performance plateaus resulting from learners modifying inherited behaviors randomly. We test this hypothesis in four experiments in which participants, either individually or in interacting dyads, attempt to locate rewards in a search grid, guided by partial information about another individual's experience of the task. We find that in some contexts, valid inferences about another's behavior can be made from partial information, and these inferences can be used in a way which facilitates trait adaptation. However, the benefit of these inferences appears to be limited, and in many contexts—including some which have the potential to make inferring the experience of another individual easier—there appears to be no benefit at all. We suggest that inferring previous behavior from partial social information plays a minimal role in the adaptation of cultural traits.

Keywords: Cultural evolution; Cumulative culture; Social inference; Social learning

Correspondence should be sent to Mark Atkinson, School of Management, University of St Andrews, The Gateway, North Haugh, St Andrews, KY16 9RJ, UK. E-mail: mda20@st-andrews.ac.uk

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

1. Introduction

Many cultural traits, such as those involved in preparing food, locating resources, determining causal relationships, or even designing scientific experiments, are fundamentally “search tasks,” in that they involve searching for and selecting behaviors from a vast number of possible options, in search of behaviors which result in desirable outcomes. For these tasks, users have to balance the benefits of exploiting familiar actions for known rewards with exploring novel actions with unknown payoffs (Hills, Todd, Lazer, Redish, & Couzin, 2015; Wu, Schulz, Speekenbrink, Nelson, & Meder, 2018). Human culture is unique in that many such traits, with repeated transmission from individual to individual and generation to generation, change in ways that make them increasingly beneficial to their users. As a result of this, later generations are able to exploit some inventions and behaviors that might also have been preferred by their predecessors, but which were not used simply because they had not yet been discovered. Thus, benefits of experience appear to accrue over successive generations within human populations such that it is possible for individuals to exploit traits discovered as a consequence of others’ cumulative exploration efforts (e.g., Henrich & McElreath, 2003; Tennie, Call, & Tomasello, 2009; Tomasello, 1999).

By contrast, there is little evidence of this occurring in cases of social transmission in other animals, including species recognized to be capable of imitative learning, and those exhibiting population-specific behavioral traditions in natural environments (e.g., Dean, Vale, Laland, Flynn, & Kendal, 2014; Tennie et al., 2009; though see, e.g., Caldwell et al., 2020; Laland & Hoppitt, 2003; Mesoudi & Thornton, 2018; Schofield, McGrew, Takahashi, & Hirata, 2018, for further discussion regarding the lack of evidence for cumulative culture in non-humans). However, in many respects, it is the human anomaly of accumulating benefits that requires explanation, more so than its apparent absence in other species. To appreciate this point, consider that each generation can be regarded as being exposed to a broadly equivalent “snapshot” of information about the behavior of others. That snapshot is limited by individual lifespans, the availability of potential demonstrators, and that the amount of effort devoted to the development of one’s own skill is necessarily limited, even if only by the fact that it can only realistically increase at the expense of other potentially valuable opportunities for exploration or exploitation. However, in spite of the window of opportunity being bounded and relatively constant, benefits from that learning apparently continue to accrue. This suggests that the social information to which later generations are exposed is in itself more valuable to potential learners, relative to equivalent social information which would have been available from observation of members of earlier generations.

Although we are inclined to take it for granted on account of its very ubiquity, this phenomenon of increasingly beneficial traits arising through social transmission—usually referred to as cumulative cultural evolution (e.g., Boyd & Richerson, 1996; Mesoudi & Thornton, 2018; Tomasello, 1999)—is therefore highly remarkable. Human cumulative culture is sometimes explained as a consequence of high-fidelity transmission of traits (e.g., Dean et al., 2014; Lewis & Laland, 2012), and certainly this must be part of the

story. However, even if traits are transmitted with relatively high fidelity, how is it that, in refining and modifying observed behaviors, newcomers typically avoid producing behaviors that have been abandoned in earlier generations upon discovery of more profitable alternatives? Avoiding adopting such behaviors is necessary for traits to typically become increasingly beneficial to the user, and such avoidance appears to occur in spite of the fact that learners have had no direct or indirect exposure to these earlier behaviors (and therefore often have no knowledge of the relative payoffs involved). This is particularly noteworthy given that cumulative cultural evolution is expected to generate traits which would be increasingly improbable as individual discoveries (i.e., that an individual would be increasingly unlikely to discover without social information). Therefore, errors of transmission would typically be expected to operate in the opposite direction, creating pressure for simplification and reproducibility. What then can account for the preservation, and even continued improvement, of behaviors which would be unlikely to be produced spontaneously, without much apparent “backwards slippage” (e.g., Tennie et al., 2009) reducing benefits to the user? How, for example, are technological tools altered in ways which increase their efficiency when a learner does not have access to a historical record of earlier devices and their payoffs? Or how does a researcher discover a more optimal solution to a problem than one they observe in light of the “file-drawer problem,” resulting in their not having access to alternative solutions already found by other researchers to be in some way inferior?

One possibility, which we explore here, is that learners can detect systematic regularities in others’ behavior, recognize that they have been produced deliberately and with some effort, and use them to generalize beyond the behavior they observe. Humans are sensitive to structural regularities in received information in a variety of modalities (as evidenced by the statistical learning literature; see, e.g., Fiser & Aslin, 2001; Frank, Tenenbaum, & Gibson, 2013; Frost, Armstrong, & Christiansen, 2019; Saffran & Kirkham, 2018), and so if non-random behavior is attributed to an informed strategy on the part of a demonstrator, then the learner might be able to infer that apparent avoidance of certain behavioral variants may mean that these have already been sampled and rejected, either by that demonstrator themselves or by one of their (cultural) predecessors. It is easy to see that such an ability might potentially facilitate the process whereby newcomers can apparently pick up a search where their predecessors left off, with some protection against backwards slippage. This could potentially prevent, or at least delay, performance plateaus that might otherwise be expected if learners simply modified socially inherited behaviors through random search strategies which treated all possible unknowns as equally worthy of personal exploration.

Such a mechanism could be useful for many types of cultural trait, but it could be particularly important for traits with certain characteristics. This would include those which leave no lasting physical record from which a learner could determine the behaviors of many or all previous generations. It will also apply to traits which are opaque, in that the function or effects of particular aspects of a behavior are difficult or impossible to determine. This may be the case for certain manufacturing processes, complex technologies, meal preparations, or the tying of complex knots, to give a few examples, where a

behavior may involve the learner completing a number of non-intuitive sub-goals. Alternatively, or additionally, this mechanism may be particularly useful for traits for which it is difficult or impossible to predict the payoffs which may be received from adopting a particular behavior, which may be the case when adopting a foraging strategy in an environment which offers little explicit indications of where desirable resources may be located. However, even if this mechanism could be especially useful in these cases, it could still benefit other types of cumulative cultural trait as well. For example, a learner could save having to examine a (possibly prohibitively vast) physical record, or avoid making (possibly very numerous) calculations or predictions about the payoffs a particular behavior could receive, even where it would be theoretically possible to do so.

If such inferences do turn out to allow learners to make significantly more effective use of social information in ways that could allow them to overcome performance plateaus, then this could also explain why this phenomenon does not tend to be observed in other species. As detailed above, learners would need to make and use inferences about others' experience or knowledge, to avoid wasting potentially valuable search time exploring options that have already been investigated and rejected in prior generations. And in spite of recent evidence of great apes predicting others' behavior on the basis of those agents' beliefs (Kano, Krupenye, Hirata, Tomonaga, & Call, 2019; Krupenye, Kano, Hirata, Call, & Tomasello, 2016), such understanding appears to be "implicit," that is, it is not believed to be accessible for use in other contexts. Such a capacity therefore would function only to generate expectations about others' likely behavior, and it could not be used to inform one's own behavioral decisions outside of the context of immediate interaction with that particular individual.

Here, we investigated the role of inferring previous behaviors in the cultural transmission of search tasks, in line with established methodologies used to study the evolution of cultural traits in controlled conditions (see Caldwell et al., 2020; Caldwell, Atkinson, & Renner, 2016, for a review). We used a variant of a search task paradigm we have employed previously in the study of the mechanisms of cumulative cultural evolution in adults (Atkinson et al., submitted; Mackintosh et al., in prep.), children (Wilks et al., in prep.; Wilks et al., submitted), and non-human primates. Modeling the naturalistic process of cultural transmission, our experimental design focused on the actions, and those actions' effects on a transmitted behavior, of an individual learner (an "Observer"), who is exposed to some information about the behavior of another (a "Demonstrator"), and then acquires and modifies that behavior. We used human participants as Demonstrators so as to obtain social information which exhibited naturalistic variation, as opposed to constructing the social information ourselves. In a task similar to a game of "battleships," the Demonstrators searched a grid of tiles with the aim of maximizing the payoffs generated by their selections. They did so by selecting tiles that increased task score ("hits"), and avoiding those that used up search opportunities without increasing the score ("misses"). They were given multiple attempts at the task to allow them to develop increasingly beneficial behaviors as a result of their own trial-and-error learning. We anticipated (correctly) that these ordered variants of the behaviors would therefore possess one of the key features of cumulative cultural traits: They would improve, in the sense

that the payoffs to the user (here, their task score) increased, with earlier behavioral variants being abandoned in favor of those leading to higher payoffs. The Observers, who were given the same task as the Demonstrators of achieving as high a score as possible by selecting hits and avoiding misses, were only provided with the final variant(s) of the behavior. They therefore had social information which they could benefit from when adopting and modifying the observed behavior, but they were not explicitly provided with the full history of the trait (i.e., information of all the variants sampled and rejected by a Demonstrator). To allow our investigation to focus on whether the Observers made and used valid information about the Demonstrator's behavior that they were not explicitly provided with, we made using the explicitly provided social information as trivial as possible. The social information remained visible for inspection while the Observer was making their selections, and the hits and misses of the social information were clearly marked. The Observer therefore had no memory constraints to contend with when using the social information, and the relationship between elements of the behavior and their contribution to the overall payoff was completely transparent. Over four experiments, each considering a different context of task completion and transmission of social information, we investigated whether Observers would (a) infer behavior from partial social information of a Demonstrator's experience of a task (i.e., determine behavior over and above the partial social information they were directly exposed to) and (b) use those inferences to facilitate their own performance on the task. If they did, then this would suggest that such a mechanism could, at least in principle, play a role in cumulative cultural evolution. Further work would then of course still be necessary to confirm that it played a role beyond a single step of cultural transmission, as well as establishing the extent to which this mechanism plays a role in more naturalistic contexts.

In Experiment 1, we gave participants (the Observers) only limited partial information about another's experience of a problem and found no evidence that they made and used valid inferences about the other person's behavior from that information. In Experiment 2, we aimed to assess whether participants would carry out the inferential step in a context in which we believed it would be easier for them to do so. We adapted Experiment 1 so that the information received by the participants was more likely to contain transparent structural regularities from which further information about another's behavior could be inferred while also increasing the overall amount of information transmitted. In this second context, we found that participants did make and use valid inferences about another's behavior. The benefit received from those inferences, however, was marginal, with many participants showing no evidence of having made any inferences at all. In Experiments 3 and 4, we explored two further contexts which had the potential for reducing the challenge of making and using valid inferences of another's behavior. Experiment 3 involved a differently structured reward space in which previous participants were likely to have behaved in more predictable ways, and Experiment 4 involved dyadic interaction and so offered the potential for participants to coordinate their search strategies. Despite the greater potential for individuals to behave in more predictable and structurally transparent ways in each of these experiments, however, we found no evidence that the participants made and used valid inferences of another's behavior. We end with a

discussion of the implications of the results of all four experiments for the cultural transmission of adaptive traits.

2. Experiment 1: Exposure to minimal partial information

We assessed the ability for participants (Observers) to make and use valid inferences about others' past experience of a problem from exposure to partial information about their behavior. In this case, this partial information was a single attempt at a problem made by another participant (a Demonstrator).

For consistency and ease of comparison with Experiments 2–4 which follow, below we present the methodology and results for this experiment referring only to a single Observer condition. However, we also collected data for a second Observer condition which can be compared with the one we report here, and it was the approach with both Observer conditions that we preregistered with the Open Science Framework prior to data collection: <https://osf.io/9heqv>. The second Observer condition was essentially a control condition, in which no structural information about the Demonstrator's search strategy was available, and so it was not possible for the Observer to make valid inferences about previous Demonstrator behavior. For completeness, we describe the preregistered version of this experiment in Supporting Information Section 1, although the additional data and analyses do not alter the conclusions we reach below.

2.1. Methods

The task closely resembled the popular childhood game of “battleships,” but with a 20×20 board and 20 ships of size 1×1 . More specifically, all participants were given the goal of maximizing the payoffs generated through their search of a grid, by selecting tiles that increased task score (“hits”), and avoiding those that used up search opportunities without increasing the score (“misses”). Participants were grouped in pairs: The first participant of each pair was assigned to the Demonstrator role, and completed the task in the absence of any social information; the second was assigned to the Observer role, and completed the task while exposed to part of the behavior of the Demonstrator.

Each participant completed three different search problems, with each problem consisting of searching a grid of 20×20 tiles. For each problem, a randomly allocated 20 of the 400 tiles were designated scoring tiles (i.e., hits), and selecting those added 1 point to the participant's score without decreasing the number of search attempts remaining. Note that the random allocation of the rewards meant that the location of one reward could not be predicted from any other reward(s). This allowed us to specifically investigate the use of inferences relating to Demonstrator search strategy, without any confounding factor of there also being some (potentially inferable) structure in the reward space.

The other 380 tiles were designated as non-scoring (misses), and selecting those did not increase the task score, but used up a search attempt. Participants (whether Demonstrators or Observers) had 10 opportunities to search the grid in each problem, and to

earn as high a score as they could within these 10 “expeditions.” In each of the 10 expeditions, participants had up to 20 “search attempts”: They kept making selections until they reached their limit of 20 misses (or found all 20 hits). It was therefore possible, if extremely unlikely, for a participant to make 39 selections (19 misses and 20 hits or 19 hits and 20 misses) in a single expedition. The number of search attempts remaining within the current expedition was shown on screen, along with the cumulative score for the entire experiment. When selected, scoring tiles were marked with a green hexagon, while selected non-scoring tiles were marked with a red circle. Once selected, a tile could not be selected again within the same expedition. See Fig. 1 panel (i) for screen as viewed by the participant at the start of an expedition, and panel (ii) for how the experiment could appear part way through an expedition. Note that this example is primarily intended to illustrate the methodology: Such a transparently systematic search of the grid was not typical for the majority of the participants.

For each new expedition of a single problem, the scoring tiles remained in the same positions. However, all tiles returned to their unselected state, so the location of the hits and misses found by the participant were once again concealed (i.e., at the start of each expedition, the screen would revert to, e.g., panel (i) in Fig. 1, but with the score at the top right of the screen recording the participant’s cumulative score at that point). Selecting scoring tiles, whether selected in a previous expedition or not, always added 1 point to the participant’s score without decreasing the number of search attempts remaining. As the participant did not have direct access to the selections of their previous expeditions, they were (implicitly) incentivized to search the grid in a systematic, that is, non-random, way, so as to keep track of which tiles they had selected in previous expeditions and not re-select tiles previously found to be unrewarded.

After participants completed their 10 expeditions for one problem, they moved on to the next problem. Participants were given three problems to allow them to potentially benefit from task experience, in that they had the opportunity to develop their search strategy, or use entirely different ones, as problem number increased. To make the transition between different problems more salient, and so stress that the location of the rewarded tiles in one problem was independent of the location of the rewarded tiles in another, the (unselected) tiles were colored differently for each problem (white, gray, or black). Participants were also explicitly informed that one problem had ended and the next was to begin.

As noted above, participants assigned to the Demonstrator role attempted each task with no prior information about whether any of the tiles were scoring or non-scoring. Each Demonstrator’s performance was then used to generate the information provided to a second participant assigned to the Observer role.

Observer participants were told that another participant had taken part in the task, and that their task was the same, that is, that they had 10 expeditions for each of the three problems, and were aiming to get as high a score as possible. Observer participants were then given full information about the selections that were made by the Demonstrator on the Demonstrator’s final expedition for each problem. The Demonstrator selections from that final expedition were marked on the grid. These selections were given one after another in the order in which they were selected. Scoring tiles were marked with the

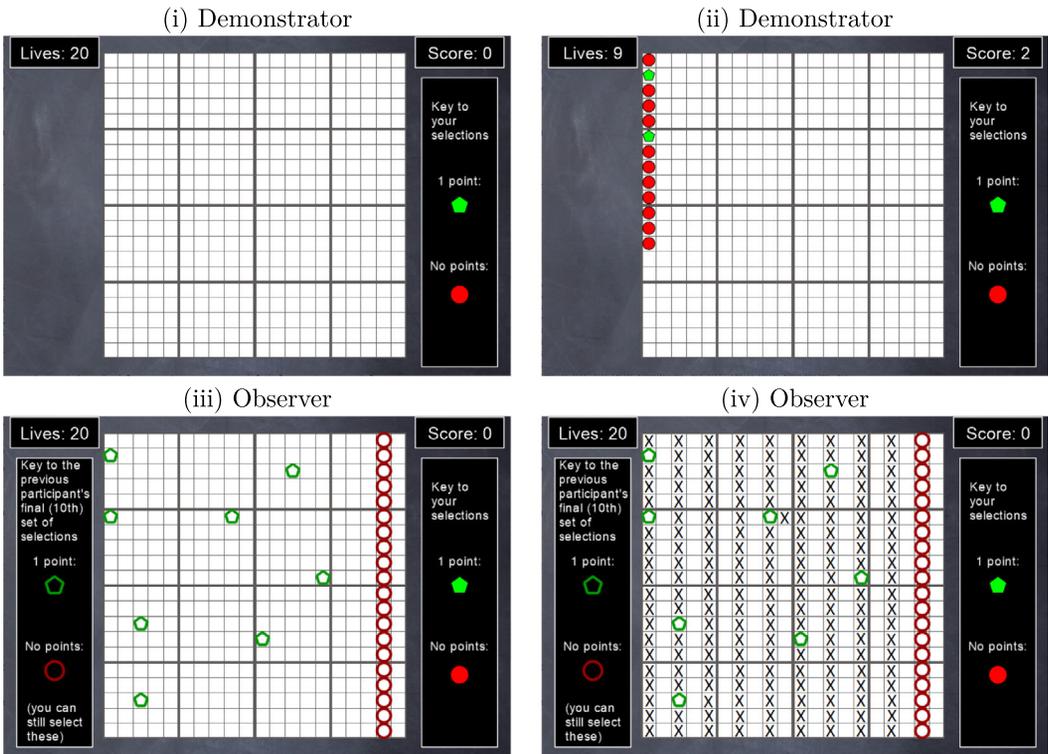


Fig. 1. Example screenshots from hypothetical participants in Experiment 1. Panel (i): Demonstrator at the start of an expedition. Panel (ii): Demonstrator partway through an expedition. Panel (iii): Observer at the start of an expedition. Panel (iv): Observer at the start of an expedition (as in panel (iii)), with additional crosses illustrating tiles which the Demonstrator had already found to be unrewarded, but which the Observer received no explicit information about. Above-chance avoidance of these crosses (which were not made visible to the participant at any point) would suggest that the Observer was able to infer the search behavior of the Demonstrator (at least to some extent) and avoid “redundant selections”: misses that the Demonstrator had already selected, even though the Observer did not have direct access to this information. In the example shown here, the Demonstrator had broadly used a search strategy of searching every other column for undiscovered hits, in addition to selecting all hits discovered in previous expeditions. Note that such transparently systematic Demonstrator behavior was not typical for the majority of our participants.

outline of a green hexagon, and non-scoring tiles marked with the outline of a red circle. This same information about the Demonstrator’s selections (excluding the information about the order in which they were selected within the expedition) was displayed on the grid for all 10 of the Observer’s expeditions for that particular problem. See Fig. 1 panel (iii) for an example of what the task could potentially look like to an Observer at the start of an expedition (following a particularly transparently systematic Demonstrator search). Panel (iv) illustrates the Demonstrator expedition which the Observer would be explicitly informed of (the final Demonstrator expedition for that problem) alongside crosses which indicate an example of the additional information from the first nine Demonstrator

expeditions which the Observer would not be explicitly informed of (the Observer would never see any crosses).

The Observer could make use of information they received however they wished. They were free to select tiles which were indicated as hits, which would always earn a score of 1 point. They were also free to select the tiles indicated as misses, which would always score 0 points, as well as reducing the number of search attempts remaining within the current expedition. Note that as the locations of all the hits from the Demonstrator's last expeditions were marked on screen, it should have been trivial for the Observer to get a higher cumulative score than the Demonstrator. As noted above, this allowed us to investigate whether Observers could make and use valid inferences from partial social information without adding a confound of, for example, a memory bottleneck which could affect the use of the partial social information received (we do investigate the role of exposure time and memory bottlenecks on cultural evolution outcomes using a similar task in Atkinson et al., submitted, however).

See Fig. 2 for an illustration of the relationship between the Demonstrator and Observer for a given pair for participants. We assumed that the Demonstrators (themselves asocial learners) would have behaved in particular ways in their own completion of the task. Specifically, we expected that when given a series of opportunities to explore the same grid, these participants would have developed non-random search strategies which allowed them to avoid redundancy across their own attempts (i.e., minimizing re-selection of grid points already identified as a “miss” on a previous attempt). Hence, for participants placed in the role of Observer, who were exposed to the final search attempt of one of the Demonstrators, it may have been possible for them to have avoided redundant selections—selections already made and found to be unrewarded by the Demonstrator—above the level that would be expected by chance, if they could infer the systematic approach of their predecessor with some degree of accuracy.

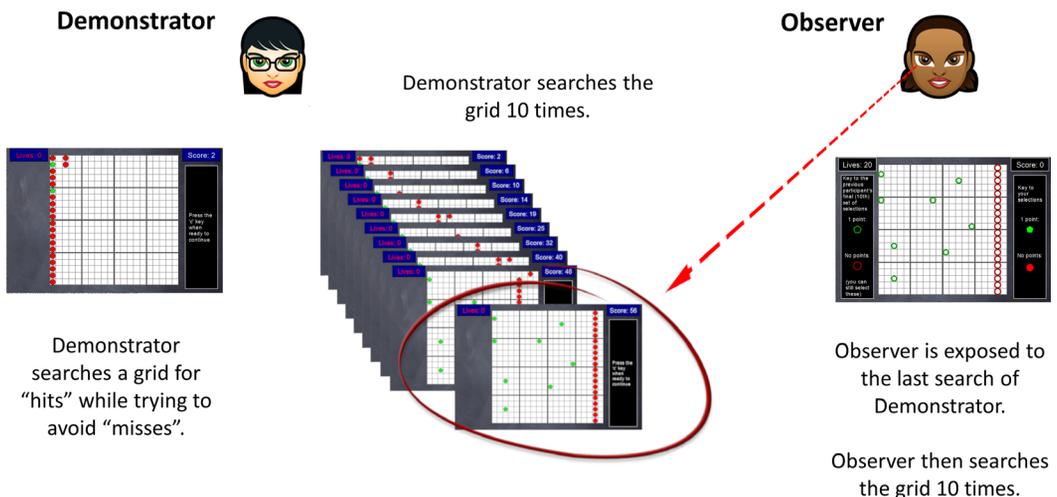


Fig. 2. Illustration of the relationship between Demonstrator and Observer in Experiment 1.

For participants assigned to the Observer role, we defined a category of selections—“redundant” selections—to capture those selections which were redundant from the perspective of the social inheritance history, that is, non-scoring tiles that were also selected by the Demonstrator predecessor at some point during the Demonstrator’s expeditions exploring the same problem, but which the receiver themselves were not directly exposed to in the social information they received (i.e., they were not selected in the final Demonstrator expedition). If the Observer managed to perfectly infer the search behavior of the Demonstrator from the partial social information transmitted to them, then they could avoid making any redundant selections.

In our analyses, we considered two key variables: For all participants (both Demonstrators and Observers), we assessed the selection of scoring tiles (note that this is equivalent to assessing participant score for Experiment 1, though cf. Experiments 2–4); for the Observers, we assessed the selection of redundant tiles.

The experiment was written and run using PsychoPy 1.84.2 (Peirce et al., 2019).

2.1.1. Participants

Forty adult participants (mean age 20.7, range 18–31; 23 females and 17 males) were recruited at the University of Stirling and took part in exchange for either research participation tokens required for course completion (35 participants) or £3 in cash (five participants).

2.2. Results

Statistical analyses (for all the experiments we report here and in the Supporting Information) involved generalized linear mixed effects modeling and were carried out using R (R Core Team, 2013) and *lme4* (Bates et al., 2013). Models with binomial dependent variables used logit link. Models with “maximal” random effects structures (Barr, Levy, Scheepers, & Tily, 2013) were considered in the first instance, with random slopes, followed by random intercept terms removed as necessary to address singular fit or non-convergence issues. $p < .05$ was taken as statistically significant, and for non-logit-linked models p values were estimated from the resultant t-statistics, taking an upper bound for the degrees of freedom as the number of observations minus the number of fixed parameters in the model (Baayen, Davidson, & Bates, 2008).

2.2.1. Selection of scoring tiles

Across all three problems, the mean number of scoring tiles selected by the Demonstrators was 133.6 ($SD = 39.33$), compared to 278.2 ($SD = 48.80$) for the Observers. We constructed a logit-linked mixed model with whether a selected tile was a scoring tile or not as (binary) dependent variable, and participant role (Demonstrator vs. Observer; treatment coded), expedition number (centered), and their interaction as fixed effects. Participant identity within pair membership, and problem number were included as random intercepts.

Observers selected more scoring tiles than Demonstrators ($b = 0.826$, $SE = 0.045$, $z = 18.56$, $p < .001$), and the selection of scoring tiles increased with expedition

number ($b = 0.160$, $SE = 0.008$, $z = 19.86$, $p < .001$). The effect of expedition number was more pronounced for the Demonstrators ($b = -0.111$, $SE = 0.010$, $z = -11.30$, $p < .001$).

2.2.2. Redundant selections

We assessed Observer selection of redundant selections by comparing the number of redundant selections actually made to the number we would have expected if the selections were made at random. If the Observers selected fewer redundant selections than this chance level, then this would indicate that they were inferring the location of these redundant selections, and searching the grids in such a way so as to avoid them. The mean (of individual Demonstrator means) of *potential* redundant selections was 63 ($SD = 9.9$; range 48–84). Note that if the Observer was perfectly able to make and use valid inferences about the Demonstrator's behavior, they would be able to avoid *all* potential redundant selections.

To account for Observers selecting the same unrewarded tiles on multiple expeditions (e.g., due to their not being able to accurately remember their selections from previous expeditions), we compared the size of the set of redundant selections made at least once with the number of redundant selections we would expect if the number of unique unrewarded Observer selections (not including any unrewarded selections of tiles included in the social information received from the Demonstrator) were made at random. Selecting fewer than this chance level would indicate that they were making and using valid inferences about the Demonstrator's behavior to avoid unrewarded selections already sampled by the Demonstrator.

For the Observers, the mean (mean of the means of each participant's performance over the three problems they encountered) number of unique redundant selections was 78.7 ($SD = 8.72$), on average 20.0 *more* than the expected chance level: 58.7 ($SD = 10.07$). The difference between the actual unique redundant selections and the expected chance level is illustrated, alongside the equivalent results for Experiments 2–4, in Fig. 3.

We constructed a linear mixed model to investigate the number of redundant selections made by the Observers, comparing them to the number of redundant selections which would be expected if the unique number of Observer selections were made randomly. The fixed effect compared the actual number of redundant selections with the expected values, treatment coded with the expected values as the baseline. Participant identity was included as random intercept. The actual number of redundant selections made at least once by the Observers was more than would have been expected if the number of unique Observer selections were made randomly ($b = 20.001$, $SE = 2.980$, $t_{119} = 6.712$, $p < .001$).

To investigate the role of expedition number, we constructed an additional model with selection of a redundant selection as (binary) dependent variable, expedition number (centered) as fixed effect, and participant identity as random intercept. Selections of redundant selections reduced with expedition number ($b = -0.025$, $SE = 0.006$, $z = -4.154$, $p < .001$).

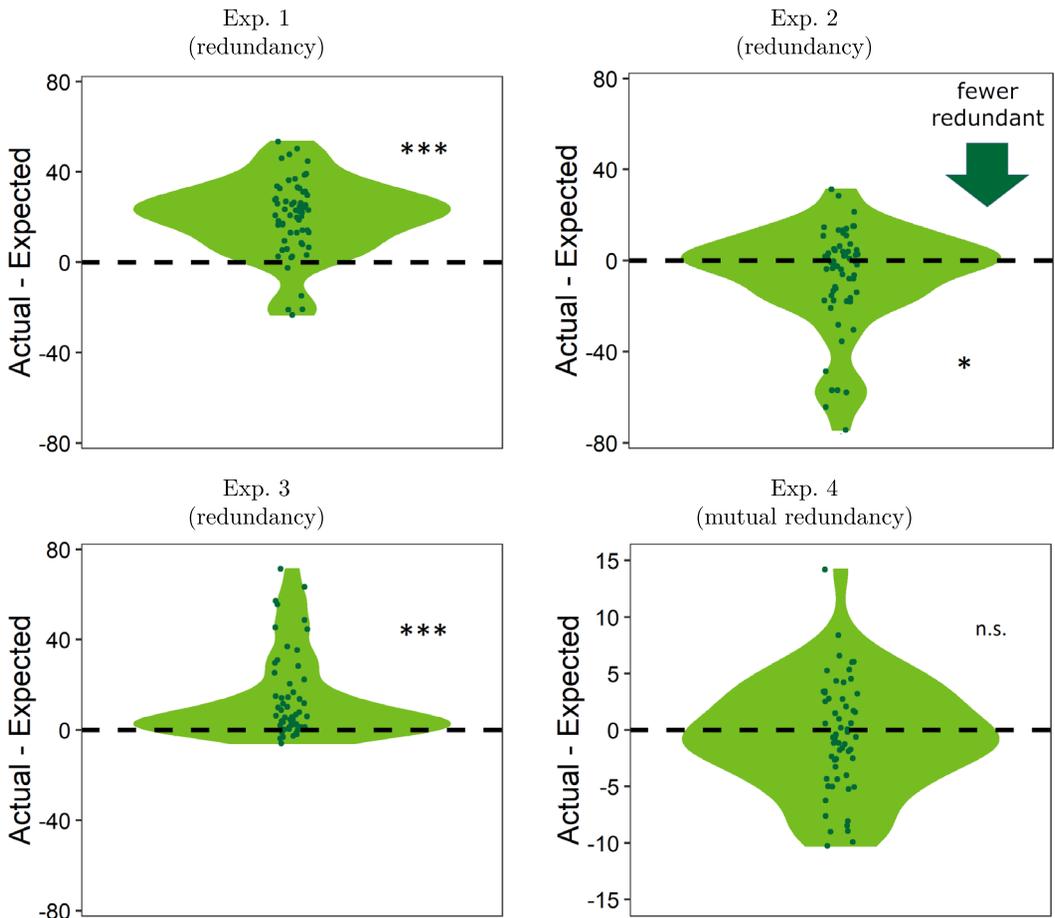


Fig. 3. Difference between actual and chance unique redundant selections for Experiments 1–3, and between actual and chance mutually redundant selections for Experiment 4. Each point represents Actual–Expected number of selections for each problem, each Observer participant providing the data for three points. Lower Actual–Expected indicates that redundant (Experiments 1–3) or mutually redundant (Experiment 4) selections were avoided to a greater extent than would have been expected if the participants made the same number of unique selections randomly. In Experiments 1 and 3, there were a greater number of redundant selections than the chance level. In Experiment 2, there were fewer redundant selections than the chance level. In Experiment 4, there was no difference.

2.3. Discussion

As expected, both Demonstrators and Observers selected more scoring tiles in later expeditions for each problem, indicating that they were using their own past experience of the problem when making their selections (i.e., by remembering where they had searched before and remembering which of their previous selections were hits and/or which of their previous selections were misses). Observers selected more scoring tiles than Demonstrators, indicating that Observers were making use of the social information

available to them (i.e., by selecting the tiles they were informed were hits and/or avoiding the tiles they were informed were misses).

Contrary to what might be expected, however, Observers made *more* redundant selections than would have been expected if they made the same number of unique unrewarded selections at random. There is therefore no evidence here that the Observers were able to successfully make and use valid inferences about the Demonstrators' past experiences. Instead, there was a significant overlap in the unrewarded tiles selected by the Demonstrators and Observers, with a wide range in the number of times individual grid locations were searched in the experiment overall. The most commonly searched locations were searched over three times as often as the least commonly searched locations. See Supporting Information Section 3 for illustrative histogram and heat maps. Regardless of role, the participants would appear to have had a bias toward selecting some tiles more than others. We return to the effect of participants having such shared "egocentric" biases, and how they may hinder making and using inferences about others' behavior, in Section 6.

In repeating rewarded behavior directly observed, the Observer behavior observed here is consistent with that likely to preserve cumulative benefits of behaviors arising from repeated cultural transmission. And as the Observers did not select *only* redundant selections among the tiles they did not explicitly receive information about, this behavior could still result in individual exploration leading to novel, beneficial, modifications to observed behavior. However, in failing to infer Demonstrator behavior and by selecting a greater number of redundant selections than we would expect by chance, this behavior would be more likely to result in a performance plateau than if they managed to infer previous behaviors from partial information.

When speculating why we did not see any evidence of the Observer inference here, it is important to consider how there would have been two (broadly) separate steps involved if our Observers had made and used valid inferences about another's behavior. One possibility therefore is that the Observers may not have been able to make any valid inferences about Demonstrator behavior, whether they attempted to or not. A second possibility is that Observers could have made valid inferences about Demonstrator behavior, but did not, or that they did make such inferences, but then did not use them. Explanations for a lack of evidence of inference here may involve either aspects of the Demonstrator behavior, or aspects of the Observer behavior, or both.

The Observers may not have made and used valid inferences about the Demonstrator behavior here due to the social information they received not allowing easy enough and/or accurate enough inference of another's behavior. This may have been due to there being too little information transmitted, and that an increased amount of social information received by the Observers would have led to greater avoidance of redundant selections. Alternatively, or additionally, the behavior of the Demonstrators may have been such that there was minimal transparent structural information available for the Observers to detect. One possibility is that the Demonstrators searched the grids in structured yet highly idiosyncratic ways. This may have led to structural regularities existing, but their only being sufficiently transparent to the Demonstrators themselves,

with the Observers being unable to detect them. The overlap between Demonstrator and Observer selections, however, would suggest this is unlikely. Instead, though the Demonstrators appeared to have had a tendency toward selecting some tiles more than others in general, they were otherwise making their selections in a fairly unstructured, that is, pseudo-random, way. The mean proportion of Demonstrator unrewarded selections after the first expedition were tiles that they had already selected and found to be non-scoring on at least one previous expedition was 23% ($SD = 5.9\%$). This suggests that though the Demonstrators were reselecting previously discovered rewarded tiles (as indicated by selection of scoring tiles increasing with expedition number), they were not optimally tracking the tiles they had previously found to be unrewarded. If this is the case, the inferential step would have been particularly challenging for the Observers. Exploratory analysis of the proportion of these Demonstrator reselections of unrewarded tiles also found no evidence of an effect of problem number (see Supporting Information Section 2.1). There was therefore no suggestion that Demonstrator task experience led to more structured search strategies, and so no evidence that the later problems could have been less challenging for the Observers to infer Demonstrator behavior.

Other explanations relate more specifically to the Observers, and their egocentric biases and motivations involved in the task. Given that the primary goal for the participants was maximizing their score, the Observers may not have devoted much effort to inferring previous behavior from partial information, even though being successful at this could have led to their increasing their scores further. This could have been due to the perceived benefits of carrying out the inference being outweighed by the perceived costs of doing so, some assessment of the difficulty of inferring the Demonstrator behavior from the available information, or the participants simply not considering it at all. Alternatively, or additionally, the Observers may not have made use of any valid inferences due to this entailing Observers deviating from behavior they may otherwise have preferred. As seen by the overlap between Demonstrator and Observer selections, and the tendency for tiles in some positions to be selected more often than others in general, if Demonstrator behavior was correctly inferred, making the most use of those inferences would likely have involved the Observer deviating from their own preferred selections. They may have been unable, or unwilling, to do this. And as indicated by redundant selections decreasing with expedition number, some tiles explored by the Demonstrators were also particularly likely to be explored first by the Observers. This prioritization of particular tiles suggests making and using inferences of Demonstrator behavior may have been especially unlikely at the start of an Observer's search. Additional exploratory analysis into redundant tile selection found no effect of problem number (see Supporting Information Section 2.2), and so there is no suggestion that Observers (not) making and using inferences about Demonstrator behavior was affected by task experience.

In Experiment 2, we investigated whether participants would make and use valid inferences about another's behavior in a context in which we believed it would be easier for them to do so.

3. Experiment 2: Exposure to more extensive partial information

As in Experiment 1, we assessed the ability for Observers to make and use inferences about others' past experience of a problem from exposure to partial information about their behavior. To investigate a context in which inferring another's behavior would likely be easier and of greater benefit, this follow-up study departed from Experiment 1 in two main ways. First, we increased the amount of Demonstrator information available to the Observer, providing the information from the last three expeditions as opposed to just the last one. This gave the Observer both information about a greater number of Demonstrator selections overall from which to infer search behavior, and also between-expedition dynamic information which we expected would make it easier for the Observer to infer the trajectory of the Demonstrator's search strategy. Second, we introduced a penalty (for both Demonstrators and Observers) for when a participant selected an unrewarded tile which they had already selected (and found to be unrewarded) on a previous expedition. As discussed above, the Demonstrators of Experiment 1 did not appear to have been keeping track of previously discovered unrewarded selections optimally. With little penalty for reselecting non-scoring tiles, they may have been exploring the grid in a fairly unstructured, that is, pseudo-random, way. By adding a reselection penalty for unrewarded tiles here, we expected the Demonstrators to search the grid in a more transparently structured way, due to a motivation to better track their unrewarded selections from previous expeditions. Including such a reselection penalty may better reflect many more naturalistic contexts, where there would likely be substantial costs (in time and energy, for example) to an individual in repeating behaviors they had already produced and found not to be beneficial.

We made two key predictions. First, we predicted that Observers would outperform Demonstrators, that is, that the presence of some social information from the Demonstrator would enable the Observer to select a greater number of hits. Second, we predicted that Observers, when exposed to only partial information about Demonstrator behavior, would demonstrate that they had been able to make and use valid inferences about the Demonstrator behavior which they were not directly exposed to. We predicted that they would make fewer selections already made and found to be unrewarded (i.e., "misses") by the Demonstrator than would be expected if the Observer were making their selections randomly.

This experiment was registered with the Open Science Framework prior to data collection: <https://osf.io/jquz8>.

3.1. Methods

The methodology for Experiment 2 followed that of Experiment 1, but with the following changes.

First, we increased the amount of Demonstrator information observed by the Observer, giving them full information about the selections that were made by the Demonstrator on the Demonstrator's final three expeditions for each problem (as opposed to

the final one expedition in Experiment 1). This not only increased the amount of information the Observers received overall but also provided between-expedition procedural information about Demonstrator searches. Before the Observer's first expedition for each problem, they were shown each of the Demonstrator's 8th, 9th, and 10th expedition's selections. To highlight first that there were seven expeditions for which the Observer received no information, they were shown seven screens with the message "Search number 1 is hidden from you," and so on. For each of the Demonstrator's 8th, 9th, and 10th expeditions, the Observer was shown the Demonstrator selections one after another in the order in which they were selected. Each of the Demonstrator's selections was marked on the grid, with scoring tiles marked with the outline of a green hexagon, and non-scoring tiles marked with the outline of a red circle. This same information about the Demonstrator's selections (including which of the last three of the Demonstrator's expeditions they were from, as indicated by the opacity of the shape outlines, but excluding the information about the order in which they were selected within each expedition) was displayed on the grid for all 10 of the Observer's expeditions for that particular problem.

Second, to increase the likelihood that the Demonstrator would search the grid in a transparently structured way and so increase the likelihood that the Observer would have been able to make inferences about the Demonstrator's search history over and above the partial information they received, we introduced a penalty for when one of the 380 non-scoring tiles was selected which had already been selected by the participant in a previous expedition. Rather than 1 point for hit and 0 points for a miss, as in Experiment 1, there were 10 points for a hit, 0 points for the first time a non-scoring selection is made, and -1 point for subsequent selections of a non-scoring tile.

Finally, to highlight that the allocation of the scoring tiles was made randomly, the participant was shown the scoring tiles—those indicated by a green hexagon—being "shuffled": They were shown 50 random allocations of 20 scoring tiles in succession over 16 s, with the speed of transition between these presentations and the transparency of the green hexagons increasing until they were no longer visible.

See Supporting Information Section 6 for screenshots of the experiment as viewed by the Demonstrator and Observer (Supporting Information Fig. S6) and an illustration of the relation Demonstrator and Observer for a given pair of participants (Supporting Information Fig. S7).

As in Experiment 1, we assessed the selection of scoring tiles for all participants (both Demonstrators and Observers), and the selection of redundant selections by the Observers. For the first analysis, note from our preregistrations for Experiments 2–4, we originally planned to assess participant score rather than selection of scoring tiles (recall selection of a scoring tile increases score by 10 points, but that score can also decrease with the repeated selection of unrewarded tiles). To focus solely on participant selections of rewarded tiles, and for ease of comparison with the score analyses of Experiment 1, we present the results using selection of scoring tiles here. Alternative analyses based on score, however, give the same pattern of results throughout Experiments 2–4.

3.1.1. Participants

Forty adult participants (mean age 22.2, range 17–45; 36 females and four males) were recruited at the University of Stirling and took part in exchange for either research participation tokens required for course completion (31 participants) or £4 in cash (nine participants).

3.2. Results

3.2.1. Selection of scoring tiles

The mean number of scoring tiles selected by the Demonstrators was 126.1 ($SD = 47.82$); for the Observers, it was 303.5 ($SD = 76.96$). The mean total scores were 1181.3 ($SD = 471.10$) for the Demonstrators and 2945.4 ($SD = 774.28$) for the Observers.

We constructed a logit-linked mixed model with whether a selected tile was a scoring tile or not as (binary) dependent variable, and participant role (treatment coded), expedition number (centered), and their interaction as fixed effects. Participant identity nested within pair membership and problem number were included as random intercepts.

Observers selected more scoring tiles than Demonstrators ($b = 1.000$, $SE = 0.075$, $z = 13.429$, $p < .001$), and the selection of scoring tiles increased with expedition number ($b = 0.157$, $SE = 0.008$, $z = 19.006$, $p < .001$). The effect of expedition number was more pronounced for the Demonstrators ($b = -0.098$, $SE = 0.010$, $z = -9.834$, $p < .001$).

Comparing scoring tile selections of Experiments 1 and 2, we also find a greater effect of role (Observers selecting more scoring tiles than Demonstrators) in Experiment 2 compared to Experiment 1. See Supporting Information Section 7 for details.

3.2.2. Redundant selections

The mean (of individual Demonstrator means) of potential redundant selections was 118 ($SD = 11.0$; range 90–137). As for Experiment 1, note that if the Observer was perfectly able to make and use valid inferences about the Demonstrator's behavior, they would be able to avoid all potential redundant selections.

For the Observers, the mean (mean over participants of the mean of each participant's performance over the three problems they encountered) number of redundant selections made at least once was 51.6 ($SD = 16.32$), on average 7.1 less than would have been expected if the number of unique Observer selections was assigned randomly: 58.7 ($SD = 6.92$). The comparison of actual redundant selections made at least once compared to the expected value if they were assigned randomly is illustrated in Fig. 3.

We constructed a linear mixed model to investigate the number of redundant selections made by the Observers, comparing them to the number of redundant selections which would be expected if the unique number of Observer selections was made randomly. The fixed effect compared the actual number of redundant selections with the expected values, treatment coded with the expected values as the baseline. Problem number and participant identity nested within pair membership were included as random intercepts.

The actual number of redundant selections made at least once by the Observers was less than would have been expected if the number of unique Observer selections was made randomly ($b = -7.029$, $SE = 3.459$, $t_{119} = -2.032$, $p = .044$).

To investigate the role of expedition number, we constructed an additional model with selection of a redundant selection as (binary) dependent variable, expedition number (centered) as fixed effect, and participant identity as random intercept.

Selections of redundant selections reduced with expedition number ($b = -0.031$, $SE = 0.006$, $z = -4.843$, $p < .001$).

3.3. Discussion

As predicted, and consistent with the results of Experiment 1, Demonstrators and Observers scored higher in later expeditions for each problem, indicating that they were using their own past experience of a problem when making their selections. Observers also selected more scoring tiles than Demonstrators, indicating again that the Observers were making use of the observed social information. In line with our predictions, and in contrast to the results of Experiment 1, Observer participants made fewer redundant selections than would have been expected if they made the same number of unique selections at random.

These results suggest that the Demonstrators did develop non-random search strategies, and that the Observers were able to make and use valid inferences about those search strategies from partial information in this context. In this experiment, therefore, we see Observer behavior which would not only likely preserve the cumulative benefits of repeated cultural transmission in a more naturalistic setting (as in Experiment 1) but which would also provide some protection against performance plateaus.

Comparing these results with those of Experiment 1, it would appear that it is possible for inference to play a role in the adaptation of cultural traits, but it may only do so when the providers of social information behave sufficiently non-randomly, and the quantity of information transmitted contains enough information about the structural regularities of their behavior. There may also have been some effect of the random allocation of rewarded tiles being stressed in Experiment 2: Attempted identification of structural regularities in another's search behavior may be more likely if an Observer is more convinced that there are no structural regularities in the reward positions to discover.

However, as evident from Fig. 3, many participants in the Observer role did not manage to make and use valid inferences about the Demonstrators' past experiences. The mean effect size—Observers on average making seven redundant selections less than would have been expected by chance for each problem—can also be considered small, given that there was the potential within our experimental design for the Observers to avoid *all* potential redundant selections. As noted above, the mean (of individual Demonstrator means) of potential redundant selections was 118, all of which could have been avoided if the Observers had perfectly made and used valid inferences about Demonstrator behavior.

As discussed in Experiment 1, there may have been a greater role for inference if the social information allowed easier and/or more accurate inference of another's behavior. Again, participants in general were much more likely to search some grid locations over others (see Supporting Information Section 5 for illustrative histogram and heat maps). The Demonstrators also suboptimally kept track of their unrewarded selections. The mean proportion of Demonstrator miss selections after the first expedition which were tiles that they had already selected (and found to be non-scoring on at least one previous expedition) was 15% ($SD = 7.0\%$); this was significantly less than the 23% in Experiment 1, however (see Supporting Information Section 8). Even if the Observers were attempting to infer Demonstrator behavior, they likely still had to contend with an egocentric preference for the selection of some tiles over others, and there may have been few transparent structural regularities to detect in any case. And as for Experiment 1, exploratory analyses found no evidence either of Demonstrator task experience leading to more structured search strategies (see Supporting Information Section 4.1), or of Observer task experience affecting the extent to which they made and used inferences about Demonstrator behavior (see Supporting Information Section 4.2).

In two further experiments, we investigated whether in other contexts—specifically those in which there was even greater potential for the behavior of others to be inferred from partial information, due to their having behaved in more predictable or more transparent ways—inference played a greater role in the modification of observed behavior and the adaptation of cultural traits.

4. Experiment 3: Variable exploration costs

We adapted Experiment 2 to assess the effect of there being potentially additional, transparent cues about the Demonstrator's behavior, to investigate whether the Observer would exploit those cues to make and use valid inferences about the Demonstrator's past experience of the task. Specifically, we introduced variability into the cost of selecting unrewarded tiles. This cost was known to the participants in advance of making their selections, and we anticipated that the Demonstrators would search tiles with lower potential costs before those with higher potential costs, akin to, for example, searching for resources in more accessible locations before less accessible ones. If they did behave in this way, then the Demonstrator behavior would be particularly predictable (certainly relative to the Demonstrator behavior in Experiments 1 and 2). If the Observers inferred this behavior and used the inferences, we would expect them not to search the tile with lower potential costs themselves but instead to pick up where the Demonstrators had left off and continue searching tiles with increasing potential costs.

This experiment was registered with the Open Science Framework prior to data collection: <https://osf.io/853kv/>. The data were collected at the same time as that for Experiment 4.

4.1. Methods

The methodology for Experiment 3 followed that of Experiment 2, but with the following changes.

Participants again completed three different search problems, with each problem consisting of a grid of 20×20 tiles. However, rather than all the tiles for a given problem being the same color, the tiles were equally and randomly allocated one of five colors (bright yellow, dull yellow, black, dull blue, and bright blue) so that there are 80 tiles of each color. Once allocated, these tile colors remained fixed for the duration of a problem. A randomly allocated 20 of the 400 tiles were scoring tiles (i.e., hits), and selecting those added 10 points to the participant's score without decreasing the number of "search credits" (cf. "search attempts" in Experiments 1 and 2) remaining. The other 380 tiles were non-scoring (misses), and selecting those did not increase the task score (although they may have decreased it—see below), but used up the number of search credits as determined by the color of the tile as follows: bright yellow, 1 search credit; dull yellow, 2; black, 3; dull blue, 4; bright blue, 5. Note that we refer to *potential* tile costs as the number of search credits only reduced following unrewarded selections. Participants (whether Demonstrators or Observers) had 10 opportunities to search each problem and to earn as high a score as they could within these 10 expeditions. Participants could keep making selections in any given expedition until they have reached (or, by selecting a non-scoring tile with a potential cost greater than the number of search credits remaining, exceeded) their limit of 30 search credits or found all 20 hits. The number of search credits remaining within the current expedition was shown on screen, along with the cumulative score for the entire experiment. Once selected, a tile could not be selected again within the same expedition. When selected, scoring tiles were (as in Experiments 1 and 2) marked with a green hexagon, while selected non-scoring tiles were marked with a red circle. Once selected, a tile could not be selected again within the same expedition.

The motivation for having tiles of different costs, as indicated by their colors, was to increase the likelihood that the Demonstrator would search the grid in a transparently structured way and so increase the likelihood that the Observer would be able to make inferences about the Demonstrator's search history over and above the partial information they received. In Experiment 2, the unrewarded selections always used (the equivalent of) 1 search credit; here, the unrewarded selections use either 1, 2, 3, 4, or 5 search credits (see above). To compensate for the increased costs of selecting the tiles, we increased the number of search credits available from 20 per expedition to 30 (broadly equivalent to the 20 search attempts of Experiment 2). Finally, unlike in Experiment 2 where the transitions between problems were made more salient due to different colored grids (all tiles were either white, gray, or black for each problem), here the appearance of the grids differed from problem to problem due to the different allocations of the tile colors.

See Fig. 4 for an example of how the task would appear to a Demonstrator midway through an expedition.

As for Experiment 2, we analyzed scoring tile selections and redundant selections. We made the same predictions, that is, that participants (both Demonstrators and Observers)

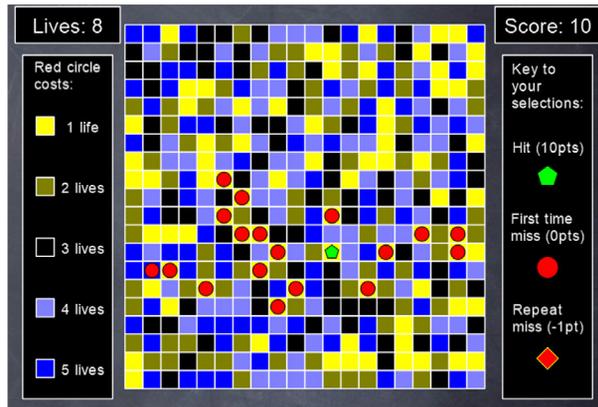


Fig. 4. Example screenshot from Experiment 3. Here, a Demonstrator is midway through an expedition. As indicated at the top left, the participant has eight credits (“lives”) remaining. On the left of the screen are the potential search credit costs of selecting a tile of a given color. If, for example, the participant selected a black tile, they would use up three search credits if it was a miss (and lose 1 point from their score if they had already selected that tile in an earlier expedition for that problem). If it was a hit, however, they did not use up any search credits and gained 10 points. The behavior illustrated here, of primarily searching the “lower cost” tiles, was typical of our Demonstrator participants in earlier expeditions, as intended.

would select more scoring tiles in later expeditions, that Observers would select more scoring tiles than Demonstrators, and that Observers would make fewer Demonstrator-only redundant selections than would have been expected if they were making the same number of unique selections randomly.

4.1.1. Participants

Forty adult participants (mean age 22.4, range 18–49; 27 females and 13 males) were recruited at the University of Stirling and took part in exchange for either research participation tokens required for course completion (32 participants) or £4 in cash (8 participants).

4.2. Results

4.2.1. Selection of scoring tiles

The mean number of scoring tiles selected by the Demonstrators was 93.3 ($SD = 53.73$), compared to 177.1 ($SD = 72.23$) for the Observers. The mean total scores were 871.0 ($SD = 536.64$) for the Demonstrators and 1708.7 ($SD = 726.92$) for the Observers.

We constructed a logit-linked mixed model with whether a selected tile was a scoring tile or not as (binary) dependent variable, and participant role (treatment coded), expedition number (centered), and their interaction as fixed effects. Participant identity nested within pair membership and problem number were included as random intercepts.

Observers selected more scoring tiles than Demonstrators ($b = 0.953$, $SE = 0.166$, $z = 5.741$, $p < .001$), and the selection of scoring tiles increased with expedition number

($b = 0.200$, $SE = 0.010$, $z = 20.796$, $p < .001$). The effect of expedition number was more pronounced for the Demonstrators ($b = -0.095$, $SE = 0.012$, $z = -7.845$, $p < .001$).

4.2.2. Redundant selections

The mean (of individual Demonstrator means) of potential redundant selections was 92 ($SD = 20.7$; range 63–132). As for Experiments 1 and 2, note that if the Observer was perfectly able to make and use valid inferences about the Demonstrator's behavior, they would be able to avoid all potential redundant selections.

For the Observers, the mean (mean over participants of the mean of each participant's performance over the three problems they encountered) number of redundant selections made at least once was 41.9 ($SD = 26.41$), on average 13.5 *more* than would have been expected if the number of unique Observer selections were assigned randomly: 28.4 ($SD = 11.33$). The comparison of actual redundant selections made at least compared to the expected value if they were assigned randomly is illustrated in Fig. 3.

We constructed a linear mixed model to investigate the number of redundant selections made by the Observers, comparing them to the number of redundant selections which would be expected if the unique number of Observer selections were made randomly. The fixed effect compared the actual number of redundant selections with the expected values, treatment coded with the expected values as the baseline. Problem number and pair membership were included as random intercepts. The actual number of redundant selections made at least once by the Observers was *more* than would have been expected if the number of unique Observer selections were made randomly ($b = 13.497$, $SE = 2.775$, $t_{119} = 4.864$, $p < .001$).

To investigate the role of expedition number, we constructed an additional model with selection of a redundant selection as (binary) dependent variable, expedition number (centered) as fixed effect, and participant identity and problem number as random intercepts.

Selections of redundant selections reduced with expedition number ($b = -0.178$, $SE = 0.008$, $z = -12.247$, $p < .001$).

4.3. Discussion

As predicted, and consistent with the results of Experiments 1 and 2, both Demonstrators and Observers scored higher in later expeditions for each problem, indicating that they were using their own past experience of the problem when making their selections. Observers also selected more scoring tiles than Demonstrators, indicating that Observers were making use of the social information available.

Contra our predictions and the results of Experiment 2, Observer participants made *more* redundant selections than would have been expected if they made the same number of unique selections at random. Though the Demonstrators' selected tiles with lower potential cost before those with higher potential cost as we anticipated, the Observers did the same, also starting with the tiles with the lowest potential costs. See Fig. 5 for illustrative figures. So, typically, the Demonstrator searched the tiles with lower costs first, and the social information transmitted indicated that they were only selecting unrewarded tiles with higher costs in their

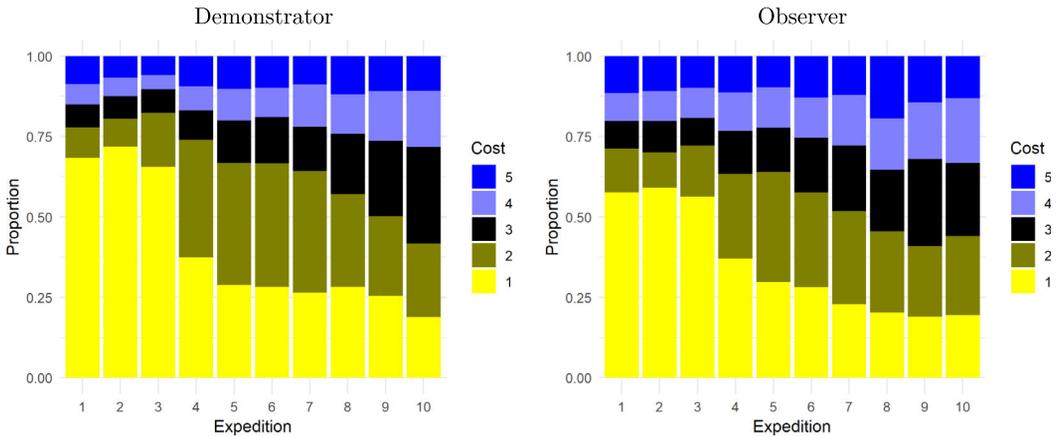


Fig. 5. Proportion of unrewarded selections by cost and expedition for the Demonstrators (left) and Observers (right). Colors shown are those used in the experiment (see Fig. 4). As evident by the greater proportions of lower cost selections in earlier expeditions for both the Demonstrators and Observers, the Observers do not appear to have inferred that the Demonstrators would typically have already discovered rewarded tiles colored as having lower potential costs (note that as the Demonstrators were reselecting the rewarded tiles discovered on earlier expeditions, the rewarded tiles colored as having lower potential costs would be included in the social information available to the Observers). If they had, then we would expect the early Observer expeditions to pick up from where the late Demonstrator expeditions left off, that is, by having increasingly smaller proportions of tiles with lower costs and increasingly larger proportions of tiles with higher costs.

final three expeditions. This would allow the Observer to infer that the rewarded tiles colored as having lower potential costs were likely all discovered and included within the social information available to them. However, the Observer still began by searching the tiles with the lowest potential costs first. Exploratory analyses also indicated that this pattern of behavior *increased* with task experience: Demonstrators search strategies appeared to become more structured (as indicated by reselections of previously selected unrewarded tiles reducing with problem number; Supporting Information Section 9.1), yet Observers increasingly made redundant selections (Supporting Information Section 9.2).

The potential for making and using inferences about the behavior of others here did not lead to participants in receipt of partial information about that behavior to do so. The Observers could have inferred that the Demonstrator's behavior was likely similar to their own, egocentric, preferred strategy (i.e., searching tiles with lower potential costs first), and made use of this by deviating from this strategy. Instead, aside from selecting the rewarded tiles they observed as social information, they broadly replicated the behavior of the Demonstrator.

5. Experiment 4: Interaction

In our final experiment, we investigated the role of inferring behavior from partial information about another individual's experience in a context involving bidirectional interaction

between individuals. We again recruited pairs of participants. But rather than one participant only taking the role of Demonstrator and the other only the role of Observer as in Experiments 1–3, the participants worked as a collaborative dyad aiming to maximize their combined score, guided by partial information about their partner’s behavior.

Interaction and coordination have been shown to lead to successful task outcomes in a variety of contexts, particularly those in which the primary goal of participants directly involves some form of coordination. In these “coordination tasks,” the optimal or target behavior of an individual is dependent on the behavior of another individual or individuals, such as tasks where the primary goal involves successful communication with other participants (e.g., Clark & Wilkes-Gibbs, 1986; Garrod, Fay, Lee, Oberlander, & Macleod, 2007; Kirby, Tamariz, Cornish, & Smith, 2015). These coordination tasks can be contrasted with the “search tasks” we have discussed so far, where the optimal behavior of an individual (such as selecting all the rewarded tiles in a grid) is unaffected by the actions of any other individual.

Interacting pairs in coordination tasks are demonstrably able to, for example, establish linguistic conventions (e.g., Clark & Wilkes-Gibbs, 1986; Garrod & Doherty, 1994) and develop novel communication systems (e.g., Fay, Arbib, & Garrod, 2013; Fay, Garrod, & Roberts, 2008; Galantucci, 2005; Selten & Warglien, 2007) in a range of different contexts. And interaction, in particular having numerous opportunities to observe the behavior of another, including observing how one’s own behavior is responded to, leads to more transparent, predictable behavior which is more likely to be correctly interpreted by those involved in the interaction (e.g., Atkinson, Smith, & Kirby, 2018; Fehér, Wonnacott, & Smith, 2016; Schober & Clark, 1989).

Here, we investigated whether a context which included interaction would lead to a greater role for inferring the behavior of others in a search task. With interaction, pairs of participants might coordinate their search strategies to more successfully infer one another’s behavior and so avoid making the same unrewarded selections. Compared to Experiments 1–3, we expected individual participants to deviate from the search strategies they might have preferred if they were taking part in the task by themselves, with an increased pressure to search the grids in not only a systematic way to allow them to remember their searches of earlier expeditions but also in a way which made the systematicity transparent for their partner.

This experiment was registered with the Open Science Framework prior to data collection: <https://osf.io/kry8g/>.

5.1. *Methods*

The methodology for Experiment 4 followed that of Experiment 2 but with the following changes.

Rather than having two participants who each searched each grid 10 times consecutively, the participants only searched each grid 10 times as a pair, swapping roles after every two expeditions. At the role swap, the individual participant whose turn it was to next attempt the task was told how many points their partner had added to their combined score over the last two expeditions. They then completed two expeditions themselves,

having received information about the second of their partner's two most recent expeditions, that is, the expedition immediately preceding the role swap. As in Experiments 1–3, the received information included the order in which the selections were made, and the selections were visible on screen throughout the two expeditions, again marked by the outlines of green hexagons and red circle to indicate which selections would be rewarded and unrewarded.

As the participants here were not confined to either a Demonstrator or an Observer role throughout, we could not analyze redundant selections as a measure of inferring behavior here. We instead defined “mutually redundant” selections to capture selections which are redundant only from the perspective of the selections made by the participants within the dyad as a whole, that is, non-scoring tiles that were selected at least once by both participants over the 10 searches for each problem, but for which neither received direct information that their partner had made the same selection.

We predicted that the selection of scoring tiles would increase with expedition number, and that there would be fewer mutually redundant selections than would have been expected if the number of unique non-scoring selections made by each participant were made randomly.

5.1.1. Participants

Forty adult participants (mean age 24.7, range 18–49; 26 females, 13 males, and one participant who elected to not provide gender information) were recruited at the University of Stirling and took part in exchange for either research participation tokens required for course completion (16 participants) or £4 in cash (24 participants).

5.2. Results

5.2.1. Selection of scoring tiles

The mean dyad total score was 1016.7 ($SD = 383.76$), with dyads selecting a mean of 106.8 ($SD = 38.73$) scoring tiles.

To assess the effect of expedition number on the selection of scoring tiles, we constructed a logit-linked mixed model with whether a selected tile was a scoring tile or not as (binary) dependent variable and expedition number (centered) as fixed effect. Participant identity nested within pair membership and problem number were included as random intercepts.

The number of scoring tile selections increased with expedition number ($b = 0.150$, $SE = 0.009$, $z = 16.81$, $p < .001$).

5.2.2. Mutually redundant selections

The mean (mean over dyads of the mean of each dyad's performance over the three problems they encountered) number of mutually redundant selections was 6.5 ($SD = 2.62$). The expected number of mutually redundant selections if the number of unique unrewarded selections were made randomly was 7.1 ($SD = 1.32$). The comparison of actual mutually redundant selections compared to the expected value if they were assigned randomly is illustrated in Fig. 3.

We constructed a linear model to investigate the number of mutually redundant selections, comparing them to the number of mutually redundant selections which would be expected if the unique number of selections were made randomly. Note that the random intercepts of dyad identity and problem number specified in our preregistration were dropped to prevent a singular fit.

There was no evidence of a difference between the actual number of mutually redundant selections and that expected if the number of unique unrewarded selections were made randomly ($b = -0.578$, $SE = 0.587$, $t_{118} = -0.984$, $p = .327$).

5.3. Discussion

As predicted, the selection of scoring tiles increased with expedition number, indicating that the participants used their own personal experience of the problem and/or the information they received about their partner's experience to increase their success at the task. Contra our predictions, however, there was no evidence of a difference between the number of mutually redundant selections and that expected if the number of unique unrewarded selections were made randomly. Exploratory analysis also found no evidence of mutually redundant selections decreasing with problem number (see Supporting Information Section 10), and so no evidence of inference developing with task experience. There is therefore no evidence that the participants made and used valid inferences of their partner's experience of the task, and so no support for a context involving interaction having a role for inferring the behavior of another from partial information.

It is possible of course that inferring the behavior of another from partial information would play a role in a different search task context which involved interaction. It is also possible that coordination of search strategies within dyads (where the partial information is transmitted horizontally) would increase with repeated cultural transmission to subsequent "generations" of dyads, that is, where the (partial) behavior of one dyad is (vertically) transmitted to another dyad. In (coordination task) studies where the primary goal for participants has been successful communication, for example, structure in participant behavior increases with cultural transmission to naive participants (e.g., Carr, Smith, Cornish, & Kirby, 2016; Kirby et al., 2015; Silvey, Kirby, & Smith, 2019). If repeated cross-generational transmission leads to increasingly regular, and so potentially increasingly transparent, behavior within dyads, then making and using valid inferences of another individual may become easier, and so possibly more prevalent, in a search task such as ours. However, given the limited evidence for a role of inferring the behavior of others across Experiments 1–4, we do not predict that this would be the case.

6. General discussion

In all four contexts we explored here, participants successfully used both the social information available to them and their own experience of the task. There was also consistent evidence for individual exploration: Participants did not simply replicate observed or previously

produced behaviors, but modified them in ways which had the potential to improve performance further. Such behavior is consistent with that which would lead to increased benefits to users with repeated cultural transmission, allowing them to exploit behaviors which they would only likely be able to discover through the cumulative exploration of others.

However, we saw only limited evidence (Experiment 2), if any (Experiments 1, 3, and 4), of participants making and using valid inferences about another person's experience of the task when modifying the behavior they directly observed, despite the potential task benefits of making and using such inferences. Though novel rewards—those never discovered by the previous participant—could still be discovered (thanks to Observer selections not being completely restricted to the set of all selections made by the Demonstrator), the behavior observed here would have only limited protection against performance plateaus. This behavior is not inconsistent with that which could lead to, for example, the development of increasingly efficient tools, or incremental progress in an area of scientific research. It would not be consistent, however, with also inferring unobservable historical tool designs, or the unpublished studies carried out by other researchers, even where it could advantageous to do so.

So why, given the potential benefits of inferring the behavior of the provider of social information beyond that observed, did we not see greater evidence of it here? It is possible that this is due to some peculiarities of our experimental paradigm, or specific elements in the design of the individual experiments, and that inference does typically play a more substantial role in the adaptation of cultural traits. Our presentation of the task, choice of grid size, or the wording of our instructions, for example, may in some way have discouraged either making inferences and/or using inferred information successfully. Participants may also have made greater use of inference if it involved inferring sampled and rejected behaviors from the exploration of multiple participants, either within a single generation or across multiple generations, rather than the exploration of a single individual. Alternatively, we may have seen such limited use of inference due to our using randomly structured reward spaces: The Observers may have prioritized inference if they believed they could not only infer something about their predecessor's search structure, but also extract information about some structure in the positioning of the rewards (see Mackintosh et al., in prep., for an example of a task involving structured reward spaces within the same experimental paradigm). Our results may also be to some extent unique to the (type of) population we got our participants from, given evidence that responses to social information (Mesoudi, Chang, Murray, & Lu, 2015) and perspective taking (Wu & Keysar, 2007) may be culturally dependent. We may find greater evidence of inference if we sampled participants from a different population, and we would particularly welcome an extension of this line of research which included cross-cultural comparison.

Alternatively, and we suggest more likely, there may be key reasons why inference played a very limited role here, and these may well be applicable to many, or even most, contexts involving the transmission of cultural traits and social information use.

First, the challenges involved both in inferring the behavior of others and successfully using those inferences may be particularly great in contexts involving the types of cultural trait we have considered in these experiments. As discussed above (see Section 5), optimal

behavior in search tasks is not dependent on the behavior of other individuals as it is in coordination tasks. In coordination tasks, a significant degree of behavioral alignment among individuals is essential to task success, and there are two implications of this for a possible role of inference. First, inference is highly likely to be directly and transparently advantageous to task success (though note that effective coordination is possible in the absence of inference, as discussed in, e.g., Barr, 2014; Sulik & Lupyan, 2018). Second, if another's behavior or knowledge is correctly inferred, then the optimal use of the inferences will likely be to largely replicate and incorporate the information inferred into an individual's own behavior. So, for example, if an individual infers the system an interlocutor is using to map a set of signals to a set of referents, then they can use their knowledge of this system to both correctly identify referents from signals they receive, and send signals with a high probability that they will be matched to intended referents. And they do this using the inferred system of signal-referent mappings in the same way as their interlocutor.

By contrast, in search tasks such as those of the experiments we present here, behavioral alignment not being essential for task success leads to different implications for the possible role of inference. First, inferring the behavior of others is not as directly related to task success as it is for coordination goal tasks. Though inferring the behavior of others can aid task success (by providing information about previously sampled and rejected behaviors which an individual would likely be best off avoiding themselves), the inferential step alone cannot lead to an individual performing well at the task. And in the case where the social information they receive comes from someone who has performed poorly at the task in some way (e.g., by forgetting the location of previously discovered rewards), then inferring their behavior may even be detrimental to task success (as avoiding their previously sampled behaviors may lead to the avoidance of those forgotten rewards). An individual in a search task is therefore left with multiple challenges if they do attempt to infer the behavior of another: the primary task itself, and the inferential step. In our experiments, they would have had to infer which selections another participant had made as well as deciding which selections to make themselves to maximize their score, all while also keeping track of their own personal search history. Given the computational challenges of inferring the knowledge and behavior of others, and the low probability of doing so with a high degree of accuracy in many situations, the (perceived) benefits of the inferential step may easily be outweighed by the (perceived) additional costs. And given that the inferential step is optional with regard to completing the task, individuals may be unlikely to attempt it (see Lieder & Griffiths, 2018, for a discussion of how cognitive strategy selection is dependent on such benefit to cost tradeoffs). Furthermore, even if an individual correctly infers the knowledge or behavior of another, then this will only directly provide information about how *not* to behave. A tool-maker, for example, will at best only have inferred which tool designs are *less* efficient than that they can observe directly and wish to improve on. In contrast to the use of inferred information in coordination tasks, the individual is still left with determining which behavior to produce. Making use of inferred information is not as straightforward as simply replicating inferred behavior, for example.

Another reason why there may be a limited role for inference relates to a conflict between making and using inferences about the behaviors of others, and the individual's

own, egocentric, bias for producing a particular, context-dependent, behavior. This conflict can prevent an individual from making inferences, due to the salience, perceived benefits, or ease of a behavior they find preferable. And, if individuals (partially) share the same egocentric biases, not making and using inferences about another's behavior may lead to multiple individuals behaving in similar ways even if they do not observe that behavior in another (as we observed, for example, in Experiments 1 and 3, with the Observers making a large number of redundant selections). Even if the inference-egocentric bias conflict does not prevent the individual from *making* inferences, it may still prevent them from *using* them. Take the Observer behavior in Experiment 3 as an illustration. For all participants, regardless of whether they were assigned to the Demonstrator or Observer role, selecting tiles with lower potential costs was clearly preferred over selecting those with higher potential costs (again, see Fig. 5 for illustration). This bias is understandable, given that a participant could search a greater number of tiles if they selected more with a lower potential cost, and this was part of our experimental design to encourage the Demonstrators to behave in a predictable way (i.e., in accordance with this bias). For the Observers, their egocentric bias to select the tiles with unknown reward values and low potential costs therefore appeared to overcome the more optimal strategy of (a) inferring that the Demonstrator would likely have had the same preference, and so (b) deducing that all the rewarded tiles colored as having lower potential costs would already be included in the transmitted social information visible on screen. And this was in spite of the social information providing additional cues that the Demonstrator had the same bias as the Observer and had made their selections accordingly. Future research could focus on determining the contexts under which inferences about Demonstrator behavior can be *made*, even if participants do not typically go on to use those inferences when given the same task as the Demonstrator. Experiments 1–3 could, for example, be adapted so that the Observers were given the alternative task of predicting Demonstrator selections from partial information.

Despite these general obstacles to inference being involved in the cultural transmission of cultural traits, we do not discount there being some contexts where making and using the inferences of others does play a more substantial role. Even in such cases, however, the challenges to the making and use of inferred information will remain. Egocentric biases can prevent successfully making and using inferences about the behavior of others in contexts where doing so would be more directly beneficial to task success (see Sulik & Lupyan, 2018, for an example of how individuals fail to accurately take the perspective of others in a signaling task). So if, for example, the context was such that an additional goal for the Demonstrators in Experiments 1 and 2 was to produce behavior from which the Observers could accurately infer that behavior from partial information, there is no guarantee that it would increase the likelihood of the Demonstrators doing so (future work could specifically investigate this). Similarly, even if the task for the Observers of Experiment 3 was to infer the behavior of the Demonstrators (behavior likely to be highly predictable by design), this is not to say that they would have done so accurately. The clash of egocentric bias and inference could be reduced by individuals having markedly different egocentric biases, but their having these different biases would likely make any inferred previous behavior less useful to the person who did the inferring, even if the individuals shared the same task goal; an

individual inferring another's previously sampled and rejected behaviors would be of limited use to them if they were unlikely to produce those behaviors themselves anyway.

In evaluating the reasons above, recall that our experimental design made the task of acquiring and beneficially modifying the behavior of another relatively trivial: There was nothing restricting the Observer's access to the social information they were exposed to. The hits and misses of the Demonstrator's preceding (Experiments 1 and 4) or three preceding (Experiments 2 and 3) expeditions remained visible on screen while the Observer was making their selections. They did not, for example, also have to remember where those socially observed hits and misses were. Our results therefore suggest that the role of inferring behavior is limited even in contexts where it would likely be easier to make and use such inferences, that is, when acquiring comparatively transparent traits which can be studied in detail. But for traits which are more opaque or transient, an additional load on the Observer's memory would likely reduce the role of inferring behavior from partial social information even further. Even though the mechanism could be particularly useful in the context of more opaque and/or transient traits, it does not necessarily follow that it is more likely to be employed.

Despite these arguments, it may still be possible that individuals may make and use inferences about previously sampled and rejected behaviors in some contexts (not necessarily involving cumulative cultural traits) where there is a particularly substantial cost to an individual producing a behavior which they had not observed, yet which had already been sampled and rejected by a previous generation. For example, if an individual believes that some food sources may be poisonous but does not know which, they may use inference to avoid readily available foods which they do not observe other individuals eating. Yet even in these cases an individual can avoid maladaptive behavior without making and using inferences about previous behaviors, thanks to, for example, high-fidelity copying and/or social learning strategies which would lead to copying the behavior produced by the majority of observed individuals.

So what, in suggesting at most a minimal role for this mechanism in the adaptation of cultural traits, are the implications for cumulative cultural evolution? How do individuals make specifically beneficial modifications to observed behaviors as they acquire and use them? As discussed in Section 1, some traits may leave physical records as they are repeatedly transmitted, and it may be that in some contexts these are actually examined in some detail before observed traits are modified. On a similar note, in some cases the pay-offs which will be received from specific modifications may be relatively easier to calculate and predict. This may be particularly the case for more transparent traits. However, perhaps most importantly, previously sampled and rejected behaviors can be avoided in cases of cumulative cultural evolution without a need for individuals to infer those behaviors thanks to other tools available to humans in the adaptation of cultural traits, such as language and teaching (see Caldwell, Renner, & Atkinson, 2018; Morgan et al., 2015; Zwirner & Thornton, 2015, for experimental evidence and discussion regarding the roles of language and teaching in cumulative cultural evolution). In Experiments 1–3, for example, if the Demonstrator were able to communicate with the Observer, they could tell them which selections they had already made, or described the general process by which they searched the grid, and the Observer could avoid making redundant selections without

needing to observe those selections directly. Similarly, in Experiment 4, the participants in a dyad would be able to agree on a strategy for searching the grid which would eliminate, or at least greatly reduce the number of, mutually redundant selections.

7. Conclusion

Over four experiments, we tested the hypothesis that inferring behavior from partial social information would have a role in the adaptation of cultural traits, by determining whether individuals would make and use valid inferences of others' behavior over and above the social information directly transmitted. In one experiment, we did find evidence for such a role for inference, but the benefit of participants making and using inferences about the behavior of others was minimal. In the other three experiments, we found no evidence to support the hypothesis, even in contexts where the behavior of another was largely predictable. We suggest that inferring the behavior of others will rarely play a role in the adaptation of cultural traits due to the costs involved in making and using valid inferences about the behaviors of others, and due to the existence of other tools which can facilitate cumulative cultural traits, such as teaching and language.

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program under grant agreement No. 648841 RATCHETCOG ERC-2014-CoG. We thank Bill Thompson for his suggestions in the design of Experiment 3. We also thank Bill Thompson, Juliet Dunstone, Marieke Woensdregt, two anonymous reviewers, and the editor for their helpful comments on earlier versions of this manuscript.

References

- Atkinson, M., Dunstone, J., & Caldwell, C. A. (submitted). Contrasting cultural evolutionary effects of cognitive constraints for copying and analytical social learning.
- Atkinson, M., Smith, K., & Kirby, S. (2018). Adult learning and language simplification. *Cognitive Science*, 42, 2818–2854.
- Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390–412.
- Barr, D. J. (2014). Perspective taking and its impostors in language use: Four patterns of deception. In T. M. Holtgraves (Ed.), *Oxford library of psychology. The Oxford handbook of language and social psychology* (pp. 98–110). Oxford: Oxford University Press.
- Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. *Journal of Memory and Language*, 68(3), 255–278.
- Bates, D., Maechler, M., & Bolker, B. (2013). *lme4: Linear mixed-effects models using S4 classes*. Available at: <https://cran.r-project.org/web/packages/lme4/index.html>

- Boyd, R., & Richerson, P. J. (1996). Why culture is common, but cultural evolution is rare. *Proceedings of the British Academy*, 88, 77–93.
- Caldwell, C. A., Atkinson, M., Blakey, K. H., Dunstone, J., Renner, E., & Wilks, C. E. H. (2020). Experimental assessment of capacities for cumulative culture: Review and evaluation of methods. *Wires Cognitive Science*, 11, e1516.
- Caldwell, C. A., Atkinson, M., & Renner, E. (2016). Experimental approaches to studying cumulative cultural evolution. *Current Directions in Psychological Science*, 25(3), 191–195.
- Caldwell, C. A., Renner, E., & Atkinson, M. (2018). Human teaching and cumulative cultural evolution. *Review of Philosophy and Psychology*, 9(4), 751–770.
- Carr, J. W., Smith, K., Cornish, H., & Kirby, S. (2016). The cultural evolution of structured languages in an open-ended, continuous world. *Cognitive Science*, 12371, 1–32.
- Clark, H. H., & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22, 1–39.
- Dean, L. G., Vale, G. L., Laland, K. N., Flynn, E., & Kendal, R. L. (2014). Human cumulative culture: A comparative perspective. *Biological Reviews*, 89(2), 284–301.
- Fay, N., Arbib, M., & Garrod, S. (2013). How to bootstrap a human communication system. *Cognitive Science*, 37, 1356–1367.
- Fay, N., Garrod, S., & Roberts, L. (2008). The fitness and functionality of culturally evolved communication systems. *Philosophical Transactions of the Royal Society of London Series B, Biological Sciences*, 363 (1509), 3553–3561.
- Fehér, O., Wonnacott, E., & Smith, K. (2016). Structural priming in artificial languages and the regularisation of unpredictable variation. *Journal of Memory and Language*, 91, 158–180.
- Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, 12(6), 499–504.
- Frank, M. C., Tenenbaum, J. B., & Gibson, E. (2013). Learning and long-term retention of large-scale artificial languages. *PLoS ONE*, 8(1), e52500.
- Frost, R., Armstrong, B. C., & Christiansen, M. H. (2019). Statistical learning research: A critical review and possible new directions. *Psychological Bulletin*, 145(12), 1128–1153.
- Galantucci, B. (2005). An experimental study of the emergence of human communication systems. *Cognitive Science*, 29(5), 737–767.
- Garrod, S., & Doherty, G. (1994). Conversation, co-ordination and convention: An empirical investigation of how groups establish linguistic conventions. *Cognition*, 53, 181–215.
- Garrod, S., Fay, N., Lee, J., Oberlander, J., & Macleod, T. (2007). Foundations of representation: Where might graphical symbol systems come from? *Cognitive Science*, 31(6), 961–987.
- Henrich, J., & McElreath, R. (2003). The evolution of cultural evolution. *Evolutionary Anthropology*, 12, 123–135.
- Hills, T. T., Todd, P. M., Lazer, D., Redish, A. D., & Couzin, I. D. (2015). Exploration versus exploitation in space, mind, and society. *Trends in Cognitive Sciences*, 19(1), 46–54.
- Kano, F., Krupenye, C., Hirata, S., Tomonaga, M., & Call, J. (2019). Great apes use self-experience to anticipate an agent's action in a false-belief test. *Proceedings of the National Academy of Sciences of the United States of America*, 116(42), 20904–20909.
- Kirby, S., Tamariz, M., Cornish, H., & Smith, K. (2015). Compression and communication in the cultural evolution of linguistic structure. *Cognition*, 141, 87–102.
- Krupenye, C., Kano, F., Hirata, S., Call, J., & Tomasello, M. (2016). Great apes anticipate that other individuals will act according to false beliefs. *Science*, 354(6308), 110–114.
- Laland, K. N., & Hoppitt, W. (2003). Do animals have culture? *Evolutionary Anthropology*, 12(3), 150–159.
- Lewis, H. M., & Laland, K. N. (2012). Transmission fidelity is the key to the build-up of cumulative culture. *Philosophical Transactions of the Royal Society B*, 367, 2171–2180.
- Lieder, F., & Griffiths, T. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, 43, e1:1–60. <https://doi.org/10.1017/S0140525X1900061X>

- Mackintosh, G., Atkinson, M., Renner, E., & Caldwell, C. A. (in prep). Intentional information sharing promotes cumulative culture relative to inadvertent behavioural cues.
- Mesoudi, A., Chang, L., Murray, K., & Lu, H. J. (2015). Higher frequency of social learning in China than in the West shows cultural variation in the dynamics of cultural evolution. *Proceedings of the Royal Society B: Biological Sciences*, 282(1798), 20142209.
- Mesoudi, A., & Thornton, A. (2018). What is cumulative cultural evolution? *Proceedings of the Royal Society B: Biological Sciences*, 285, 20180712.
- Morgan, T. J. H., Uomini, N. T., Rendell, L. E., Chouinard-Thuly, L., Street, S. E., Lewis, H. M., Cross, C. P., Evans, C., Kearney, R., de la Torre, I., Whiten, A., & Laland, K. N. (2015). Experimental evidence for the co-evolution of hominin tool-making teaching and language. *Nature Communications*, 6, 6029.
- Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K., & Peirce, J. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51, 195–203.
- R Core Team. (2013). *R: A language and environment for statistical computing*. Available at: <http://www.r-project.org/>. Accessed September 4, 2019
- Saffran, J. R., & Kirkham, N. Z. (2018). Infant statistical learning. *Annual Review of Psychology*, 69(1), 181–203.
- Schober, M. F., & Clark, H. H. (1989). Understanding by addressees and overhearers. *Cognitive Psychology*, 21(2), 211–232.
- Schofield, D. P., McGrew, W. C., Takahashi, A., & Hirata, S. (2018). Cumulative culture in nonhumans: Overlooked findings from Japanese monkeys? *Primates*, 59, 113–122.
- Selten, R., & Warglien, M. (2007). The emergence of simple languages in an experimental coordination game. *Proceedings of the National Academy of Sciences of the United States of America*, 104(18), 7361–7366.
- Silvey, C., Kirby, S., & Smith, K. (2019). Communication increases category structure and alignment only when combined with cultural transmission. *Journal of Memory and Language*, 109, 104051.
- Sulik, J., & Lupyan, G. (2018). Perspective taking in a novel signaling task: Effects of world knowledge and contextual constraint. *Journal of Experimental Psychology: General*, 147(11), 1619–1640.
- Tennie, C., Call, J., & Tomasello, M. (2009). Ratcheting up the ratchet: On the evolution of cumulative culture. *Philosophical Transactions of the Royal Society B*, 364, 2405–2415.
- Tomasello, M. (1999). The human adaptation for culture. *Annual Review of Anthropology*, 28, 509–529.
- Wilks, C. E. H., Atkinson, M., & Caldwell, C. A. (in prep). Children's use of social information from multiple models: Cognitive capacities underlying population size effects on cumulative culture.
- Wilks, C. E. H., Rafetseder, E., Renner, E., Atkinson, M., & Caldwell, C. A. (submitted). Cognitive prerequisites for cumulative culture are context-dependent: Children's potential for ratcheting depends on cue longevity.
- Wu, C. M., Schulz, E., Speekenbrink, M., Nelson, J. D., & Meder, B. (2018). Generalization guides human exploration in vast decision spaces. *Nature Human Behaviour*, 2, 915–924.
- Wu, S., & Keysar, B. (2007). The effect of culture on perspective taking. *Psychological Science*, 18(7), 600–606.
- Zwirner, E., & Thornton, A. (2015). Cognitive requirements of cumulative culture: Teaching is useful but not essential. *Scientific Reports*, 5, 16781.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article:

Appendix S1. Inferring Behaviour from Partial Social Information Plays Little or No Role in the Cultural Transmission of Adaptive Traits.