Social Learning Strategies in Networked Groups

Thomas N. Wisdom,[a] Xianfeng Song,[b] Robert L. Goldstone[c]

[a]*Department of Psychological and Brain Sciences, Indiana University*
[b]*Department of Physics, Indiana University*
[c]*Department of Psychological and Brain Sciences, Indiana University*

Abstract

When making decisions, humans can observe many kinds of information about others' activities, but their effects on performance are not well understood. We investigated social learning strategies using a simple problem-solving task in which participants search a complex space, and each can view and imitate others' solutions. Results showed that participants combined multiple sources of information to guide learning, including payoffs of peers' solutions, popularity of solution elements among peers, similarity of peers' solutions to their own, and relative payoffs from individual exploration. Furthermore, performance was positively associated with imitation rates at both the individual and group levels. When peers' payoffs were hidden, popularity and similarity biases reversed, participants searched more broadly and randomly, and both quality and equity of exploration suffered. We conclude that when peers' solutions can be effectively compared, imitation does not simply permit scrounging, but it can also facilitate propagation of good solutions for further cumulative exploration.

Keywords: Social learning; Problem solving; Exploration; Imitation; Conformity bias; Frequency bias; Similarity bias; Innovation diffusion

1. Introduction

The act of learning about the world through others permeates human life. Because of shared patterns in our experiences and preferences, such "social learning" (e.g., about restaurants, schools, or political candidates) allows people to obtain potentially useful information about a large number of options that would be difficult to evaluate directly.

Throughout this article, we will follow Boyd and Richerson's (2005) broad definition of social learning: "the acquisition of behavior by observation or teaching from other conspecifics." This phenomenon has been demonstrated in a wide variety of species and in various contexts, such as choosing mates, foraging for food, and avoiding predators (Galef & Laland, 2005). Models and experiments have also explored the implications of social learning for the development and propagation of communication systems (Galantucci, 2009; Healey, Swoboda, Umata, & King, 2007) and culture (Boyd & Richerson, 2005; Hurley & Chater, 2005). Humans' rare talent among animals for direct and flexible imitation has been called "no-trial learning" (Bandura, 1965), because it is even faster than the one-trial learning observed in animals with a strong built-in tendency to form certain associations (e.g., between the taste of a food and a subsequent stomach ache). This talent allows an imitator to add new behaviors to his or her repertoire without the costs of direct trial-and-error learning.

Modeling has shown that the tendency to imitate others rather than innovate or explore (i.e., "conformity bias") can be adaptive in uncertain environments (Boyd & Richerson, 1985; Kameda & Nakanishi, 2003). Kameda and Tindale (2006) argue that conformity bias evolved in humans and other taxa because of its tendency to promote a net average benefit for individuals, rather than error-free performance. Similarly, socially influenced "informational herding" behavior has been proposed as an explanation for phenomena such as fads and financial panics (Eguíluz & Zimmerman, 2000), but it has been argued that such behavior is the result of otherwise adaptively rational Bayesian reasoning in uncertain conditions (Anderson & Holt, 1997; Banerjee, 1992).

1.1. Factors in social learning strategies

1.1.1. When, who, and which

A simple evolutionary simulation by Rogers (1988) showed that indiscriminate copying confers no net fitness benefit to a population of individual learners in a changing environment, because the effort saved by avoiding direct learning is offset by the costs of using outdated information. Further studies extended this result to show that when social learners can imitate selectively (e.g., imitating only when individual exploration is relatively unreliable and thus more costly), the overall fitness of the population can increase, because both individual and social learning can become more accurate (Boyd & Richerson, 1995; Kameda & Nakanishi, 2002).

Laland (2004) noted the important distinction between social learning strategies that determine *when* imitation should happen and *who* should be imitated. *When* strategies often evaluate the comparative risk or cost of individual learning (e.g., "copy when uncertain" and "copy when information is costly"), while *who* strategies take into account information such as peers' solution payoffs (e.g., "copy the best") and frequency or popularity of solutions among peers (e.g., "copy the majority"). Of course, the two types of strategies are generally combined according to the nature of the problem and the information environment. Baron, Vandello, and Brunsman (1996) found that conformity increased with the difficulty of a perceptual task, and Laughlin and Ellis (1986)

found that conformity to majority judgments about a math problem increased when the solution was not demonstrably clear. McElreath et al. (2008) found that most participants in a multi-armed bandit experiment with peer information used a hierarchical social learning strategy that was primarily payoff dependent but used frequency criteria to choose between options with similar average payoffs. A related bias toward copying options with increasing frequency or "momentum" was found in an observational study of baby names from U.S. Social Security records; American parents were more likely to use names with recent increases, rather than decreases, in popularity (Gureckis & Goldstone, 2009).

If potential solutions to a problem consist of multiple elements (e.g., chemical ingredients, mechanical components, or procedures) that are combined to produce the overall result or payoff, strategies governing *which* elements to imitate become important. When there is more than one possible solution (i.e., the problem is not a search for the combination to a lock or a single needle in a haystack), and the rules governing payoffs in the combinatorial search space are unknown or complex, solutions must be incrementally compared in terms of their elements and payoffs to guide future search efforts. Similar to McElreath et al.'s (2008) hierarchical strategies, the combination of frequency information with payoff information allows for the evaluation of common elements across multiple solutions according to payoff: the more and higher payoff solutions that have a particular element in common, the higher the confidence that the element contributes to a higher payoff. This has the added benefit of being a cognitively simple process, if the solution elements are consistently ordered or otherwise readily comparable. In the absence of payoff information, frequency bias may still allow one to choose good solution elements according to the proportion of others who find them relatively favorable, though absolute payoffs must be evaluated individually. In an open tournament of simulated social learning strategies, Rendell et al. (2010) found that the best performers took advantage of a similar "filtering" of known good solutions by others.

Conversely, a strategy focusing on peer solutions that are similar to one's own (rather than similar to each other) allows one to ignore common elements and potentially deduce differences associated with different payoffs, and thus evaluate potential changes to one's own solution. This is cognitively very similar to a simple individual strategy that makes only small changes to a solution between evaluations. Rubinstein (2003) explored a related strategy in a series of experiments on individual decision making under uncertainty, in which similarities among risky options allowed agents to reduce dimensions of comparison and thus simplify decisions between alternatives—the greater the similarity between two solutions, the simpler the cognitive task of comparison. Similarity bias also allows the preservation of knowledge about one's current solution elements. In sociological studies of innovation diffusion, the use of solution similarity as a criterion for social learning has been shown to help preserve backwards compatibility with previous knowledge, so that earlier benefits are built upon instead of being discarded (Rogers, 2003). This knowledge-preservation effect would accrue for a similarity-biased strategy even without payoff information, though again, each potential change would have to be evaluated individually.

1.1.2. Group characteristics

The size of a social group might be expected to affect individuals' social learning processes, but it is not clear whether N heads are in fact better than $N - 1$. A larger number of co-present learners may lead to a wider variety of shared solutions, and thus potentially improve the performance of the social learning strategies noted above; or it may lead individuals to shirk exploration and rely solely on imitating others' solutions, a phenomenon known more generally as "social loafing" (Latané, Williams, & Harkins, 1979). Kameda, Tsukasaki, Hastie, and Berg (2011) found that in tasks with diminishing marginal returns for individual production inputs to a group endeavor (e.g., in foraging or predator vigilance), a mixed equilibrium arises in which it is individually rational for a subset of group members to consistently contribute their input. Diminishing marginal returns may also affect consumption of social information. Results from a previous experimental puzzle game in our laboratory showed that individual imitation rate was positively correlated with performance only for smaller group sizes, and that this was due to a reduced likelihood of those in larger groups to imitate the best solution among a greater number of peer solutions (Wisdom & Goldstone, 2011). This suggests that cognitive limitations on processing increasing amounts of social information can result in diminishing returns for social learning in larger groups.

1.1.3. Dynamics of learning

The general concept of learning implies changes in behavior over time; even in a static environment, such dynamics manifest in at least two ways. First, as one explores a problem space and accumulates knowledge of it, it is rational to balance exploration for potential improvements with exploitation of existing good solutions, if possible. Second and relatedly, learning strategies are subject to adaptation as performance feedback is received, redirecting learning behavior toward more efficient strategies, or adjusting investment in costly learning processes to optimize or at least "satisfice" performance (Simon, 1956). Previous experiments have shown evidence for such adaptation in the decreasing use of social information over time (in the absence of environmental perturbations) as participants learn more and have less need to rely on knowledge from others (McElreath et al., 2005; Wisdom & Goldstone, 2011).

1.2. Experimental design

1.2.1. Motivations

This research is intended to expand upon and integrate the results of previous social learning studies, to go beyond establishing the existence of learning strategies, and explore the dynamics and performance of their interaction. Like McElreath et al. (2005, 2008), we favor giving participants the opportunity for endogenous social learning in "experimental microsocieties" (Baum, Richerson, Efferson, & Paciotti, 2004), without the use of artificial participants or confederates. This design choice sacrifices some control over participants' choices and behavior, in exchange for a view into more realistic mutual influence. By observing groups of interacting participants, we gain the ability to

document emergent group-level patterns related to the quality, diversity, and transformation of solutions over time (Goldstone & Gureckis, 2009).

We exercised more control in the design of the problem space, to avoid confounding the group dynamics of learning and search with individual "insight" phenomena, because of the difficulty of operationalizing the transfer and use of information in such phenomena. Specifically, we designed the task to have a large number of possible solutions, with a corresponding range of payoffs, to observe processes of improvement and differences in efficiency between different conditions. We also avoided a systematic relation of discernible solution characteristics to the underlying payoff structure of the problem. Instead, we used a combinatorial search space whose payoff function was systematic but partially non-linear, randomized in its relation to visible solution characteristics, and too large to be exhaustively evaluated in the limited time given. Thus, we restricted our investigation to the efficiency of learning and search in a space that was not amenable to solution by conceptual insight. This, of course, only represents a subset of possible search problems, but as we will discuss later, we believe that such spaces are important and fairly common.

Our research is motivated by two basic questions about social learning. First, what strategies are used to dynamically integrate multiple pieces of socially mediated knowledge with one's own knowledge about solutions to a problem? In particular, we wish to confirm and examine in detail biases toward imitation associated with the payoff, popularity, and similarity of solutions, as well as the proportional use of asocial and social learning strategies and their dynamics over time. We expect that strategies will shift to take advantage of the characteristics of different information environments, but experimental investigation is necessary to arbitrate between multiple theoretical predictions. Second, how do individual and aggregate group strategies interact in performance outcomes? In particular, we wish to examine the outcomes associated with different group sizes and different proportions of strategies in a group. For instance, imitation is most likely beneficial to the imitator, but excessive imitation may reduce exploration and the diversity of solutions in a group, and thus impede long-term group performance.

To address these questions, we performed two experiments in which groups of participants explored a large and complex combinatorial problem space while passively sharing information about their solutions with other group members. In the first experiment, we varied the size and complexity of the problem space to gauge the effects of difficulty on social learning. In the second experiment, we manipulated the availability of payoff information about peers' solutions in order to evaluate the resulting changes in their social learning strategies.

1.2.2. Predictions

In a preliminary investigation using a closely related paradigm of social learning in a two-dimensional puzzle game (Wisdom & Goldstone, 2011), we found several results that we expected to reproduce and expand upon here, including retention of high (and increasing) proportions of solution elements from one period to the next, to simplify evaluations of changed elements, as well as decreases in imitation and exploration over time, due to

the adaptation in learning strategies and balancing of exploration and exploitation discussed above. In more difficult conditions, imitation was expected to be more common and exploration rarer, displaying a "copy when information is costly" strategy in the context of a relatively larger and more complex problem space. We also expected larger participant groups to be associated with increases in imitation and decreases in exploration and retention, because participants would be able to rely on a larger pool of good peer solutions to imitate, which would reduce the need to explore individually or rely on their own previous solutions.

Imitation biased toward high-scoring solutions (payoff bias) was predicted based on its immediate benefits. Imitation of solution elements popular among peers (frequency bias) and solutions similar to a participant's own (similarity bias) were predicted based on their utility in, respectively, consolidating and isolating information about solution elements and developing further improvements to solutions. Similar considerations prompted a prediction that when payoff information about peer solutions was unavailable, frequency bias and similarity bias in imitation would be weaker, because of reduced benefits related to the lack of direct evaluative information. These reduced benefits would also reduce scores among participants employing highly imitative strategies, and thus participants would employ less imitation and more exploration, and solution diversity would increase. We expected, however, that rather than improving performance through greater exploration, the impedance of social learning by hiding payoff information would result in lower mean scores because participants would be unable to easily take advantage of good solutions found by others through selective imitation and further improve upon them.

2. Experiment 1: Difficulty manipulation

2.1. Creature game overview

For this experiment, we designed a new computer-based task that incorporated thematic elements of popular games such as virtual pets and fantasy sports leagues (though the task was designed to be far simpler than either of these). In the task, we asked participants to create combinatorial solutions consisting of a small subset ("team") of creature icons from a larger set ("league") over a series of time-limited rounds, and to try to maximize the average payoff of their team using score feedback given after each round. Score feedback was generated according to a stable (within each game) but hidden payoff function, featuring a linear term and pairwise interactions among the icons. In each round, participants could observe each of their fellow players' solutions and associated scores, and imitate them in whole or in part. The size and complexity of the problem space (and thus the task difficulty) were manipulated in two different conditions via the sizes of the overall set of icons and the subset that could be evaluated in one solution, as well as the number of pairwise interactions between icons.

2.2. Methods

Overall, 153 participants were recruited from the Indiana University Psychology Department undergraduate subject pool and were given course credit for taking part in the study. Participants populated each session by signing up at will for scheduled experiments with a maximum capacity of nine persons and were distributed across 39 sessions, as shown in Table 1. The distribution of participants among group sizes was designed to obtain a reasonable number of groups at each size for group-level analysis, but with comparable *N* in all group sizes so that none dominated aggregate analyses.

2.2.1. Participant interface

We implemented the experiment using custom software run in a web browser, and each participant used a mouse to interact with the experimental task. All participants' actions were recorded and synchronized by the game server at the end of each round. The display included an area for the participant's own current solution ("team"), an area that could be toggled to show the participant's team in the previous round or their best-scoring team so far in the game (along with the associated score), an area which showed all of the solution elements (the "league" of potential team members) that were available for selection, and indications of the current round in the game and the amount of time remaining in the current round. In sessions with more than one participant, the display also showed the solution and associated payoff of each other participant from the previous round. The ordering of peers' solutions in each participant's display was kept constant within each condition but not across conditions, to avoid imitation based on past behavior.

Any individual element could be copied from any part of the display to a participant's current solution by dragging and dropping it with the mouse, except for those already in the participant's current solution, which were faded in the display and non-draggable. The current solution could be replaced entirely by another solution by selecting the score box above the latter as a "handle" and dragging it to the current solution area. A short video demonstrating all available actions in the game can be viewed at http://cognitrn. psych.indiana.edu/CreatureGameClip.mov; a screenshot of the interface is shown in Fig. 1.

2.2.2. Instructions

At the beginning of each session, players were given an interactive demonstration of the task, and further informed about the mechanics of the task and what to expect in the

Table 1
Distribution of participants across group sizes in Experiment 1

Group Size	1	2	3	4	5	6	7	8	9
No. of sessions	8	6	5	5	5	2	4	3	1
No. of participants	8	12	15	20	25	12	28	24	9

Fig. 1. Example of experiment interface.

remainder of the experiment session, including the following information. In each game, each creature icon was associated with a certain positive number of points (its own "abilities"), and several unidentified pairs of icons were associated with separate positive point bonuses or negative point penalties (reflecting "how well they got along") when they were both on the same team in the same round. Participants were not given information about the maximum score, the score distribution, or the interacting elements. The icons' display positions and associations with the payoff function were shuffled randomly for each game, so that their appearance and placement in the display did not give clues as to their point values.

Each game consisted of 24 10-s rounds; these parameters were chosen according to the results of pilot studies showing a performance plateau after approximately 24 rounds, and inactivity after the first 10 s of each round. Score feedback (the sum of the individual and pairwise terms described above) was given after each round: If the participant's score had improved from the previous round, the numerical score display counted up to the new score and turned green, and if it had worsened, the display counted down to the new score and turned red. If a participant made no changes to their team before the 10 s of the round were over, the team stayed the same. Participants were informed that the game was intended to be difficult, to elicit strong efforts, and were instructed to do their best to maximize the sum of their teams' scores over all 24 rounds. At the beginning of each game, each participant's team was a random selection of creature icons from the league; thus, participants who made no changes or random changes to their initial teams would have mean scores corresponding approximately to the peaks of the score distributions in Fig. 2.

Each group played eight randomly ordered games, of which half (four games) had league and team sizes of 24 and 5, respectively, and the other half 48 and 6, with more
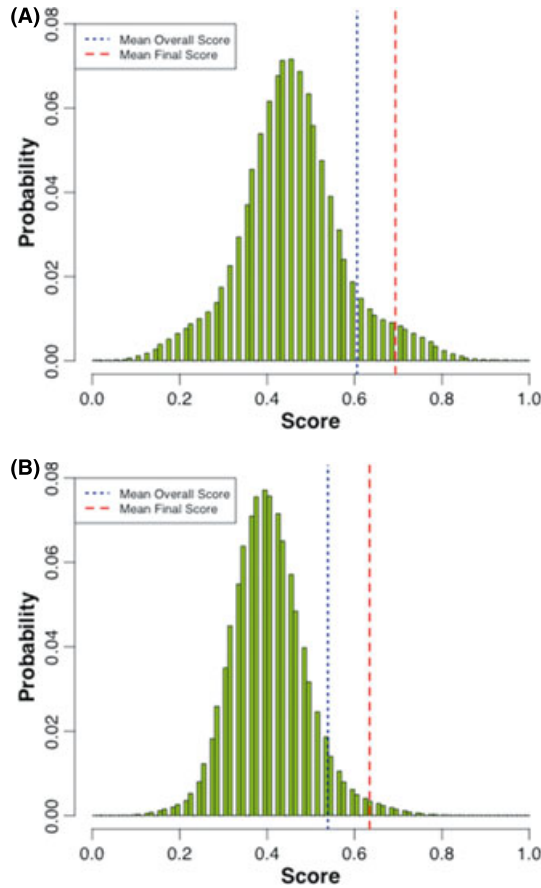
Fig. 2. Distribution of scores for all possible teams in each difficulty condition: (A) league size of 24; (B) league size of 48. Note that the vertical lines show mean participant scores from experimental data, not the means of the distributions. The mean score resulting from random or no participant activity is approximately at the peak of each distribution.

interaction pairs added in the larger league size. These two conditions were intended to vary the level of difficulty of the game, with the larger sizes being more difficult. This was because, although the score distribution and combinatorics made higher absolute scores possible in the larger league size, it also made high-scoring teams comparatively rarer.

2.2.3. Dependent variables

The score for a team was computed by summing the individual point values for each icon, and then adding or subtracting the value of any special pairs present. The pairs did not overlap, and the distribution was designed to be challenging: Pairs that gave large positive bonuses were distributed among icons with small individual point values, and pairs that gave large negative penalties were generally found among icons with large indi-

vidual point values. Possible score ranges for the small and large league and team size combinations were (−6, 51) and (−6, 60), respectively; for ease of comparison and analysis, all scores were normalized to the range (0,1) according to the range of scores possible in each condition. The combinations of individual and pair values resulted in the probability distribution of scores among all possible teams for each condition shown in Fig. 2.

In each round, the following data were automatically recorded for each player: the icons on the current team at the end of the round (or *choices*), the *source* of each icon, and the resulting *score*. The *source* indicated whether each icon was kept unchanged on the team from the round immediately prior (*Retained*), copied back from the player's own best-scoring team so far (*Retrieved*), chosen from the overall league display (*Explored*), or copied from another player's team (*Imitated*). When icons were Imitated from another player, the persistent identifier of that player was recorded to allow further analyses of imitation decisions. Note that source information was not estimated from the contents of players' teams but recorded explicitly according to each player's actions (i.e., the region of the display from which the icon was chosen). In the case of a player replacing the entire team with Imitated icons, only icons which were not already on the player's previous round team were recorded as Imitated; icons which remained on a player's team from one round to the next were always recorded as Retained. The same was true of replacing an entire team with Retrieved icons, or removing an icon and then putting it back on the team via a League choice.

The *similarity* of two solutions was defined as the proportion of elements that the two solutions had in common; for example, solutions with three out of six elements in common had a similarity of 0.5. An *improvement* was defined as an instance of a participant obtaining a score higher than any player's previous scores within a particular condition. A participant's *score rank* in a particular round was defined as the rank of their score (with one being the best) among all scores in the group in that round; individuals with the same score had the same rank. *Guess diversity* for a particular round was defined as the proportion of icons in the league represented on one or more participants' teams during a given round. This value was normalized by the average expected value of this proportion for each participant group size, generated by a Monte Carlo simulation assuming independent random teams. Note that, in general, the words "guess" and "solution" will be used interchangeably throughout this article.

2.3. Experiment 1 results

This section is organized as follows: We present summaries of (a) evidence for specialized social learning strategies; (b) dependent variables in aggregate and across rounds, game order, and participant group size; and (c) examinations of individual and group performance in relation to learning strategies. All analyses were performed on the entire aggregated dataset, except where noted or inappropriate (e.g., Imitation analyses excluded isolated participants).

2.3.1. Specialized social learning strategies

This section presents results regarding social learning biases according to solution payoff, solution element frequency, and solution similarity.

2.3.1.1. Payoff bias: Of all guesses that included Imitated elements, 94.3% imitated only one other participant, 5.1% imitated two participants, and 0.6% imitated more than two participants. Of all instances of single-participant imitation, 82.4% involved imitation of participants whose score rank was 1 (the top score in the group), 10.7% whose score rank was 2, and 7% whose score rank was 3 or below. The score of the imitated participant was greater than that of the imitator in 89.6% of cases, equal to it in 2.6% of cases, and less than that of the imitator in 7.8% of cases. No significant differences across group size, round, similarity, or other factors were observed for these effects.

2.3.1.2. Frequency bias: To measure the bias of participants to choose an icon according to its frequency in peers' choices, we tallied the number of players in the group whose teams included each icon in the previous round (N_{R-1}), as well as the number of the remaining players who added it to their team in the current round via Imitation or Exploration. To convert these figures to normalized frequencies, the first number was divided by the participant group size (N), and the second number was divided by the number of participants who did not possess the icon in the previous round ($N - N_{R-1}$).

The chance probability of imitation (resulting from choosing an icon at random from among all neighbors' teams) scaled with the choice frequency of an icon relative to the team size. The chance level of Exploration (resulting from choosing an icon at random from the league display) is a constant equal to one divided by the league size. Given that league and team size conditions were balanced in all sessions, we used the average value of each to calculate the chance baselines. A linear mixed-effects analysis of imitation probability versus choice frequency showed a positive frequency-dependent bias for Imitation that was significantly greater than chance ($F(1,1128) = 1648$, $p < .0001$, $B = 0.300$; see Fig. 3A), as well as a similar but much smaller frequency-influenced bias for Exploration choices ($F(1,1128) = 268.7$, $p < .0001$, $B = 0.062$; see Fig. 3B). The latter indicates that choosing creatures from the League display did not always strictly equate to Exploration. Notably, the probabilities of Imitation and Exploration only rose above chance when a majority of a participant's neighbors possessed an icon (i.e., when Choice Frequency was greater than 0.5).

2.3.1.3. Frequency-change (momentum) bias: In a similar analysis of "choice momentum," we tallied the change in the number of players whose teams included the icon in the previous two rounds ($N_{R-1} - N_{R-2}$), as well as the number of the remaining players who added it to their team in the current round via Imitation or Exploration. To convert these figures to normalized frequencies, the first number was divided by the participant group size (N), and the second number was divided by the number of participants who did not possess the icon in either of the previous two rounds ($N - \max[N_{R-1}, N_{R-2}]$).
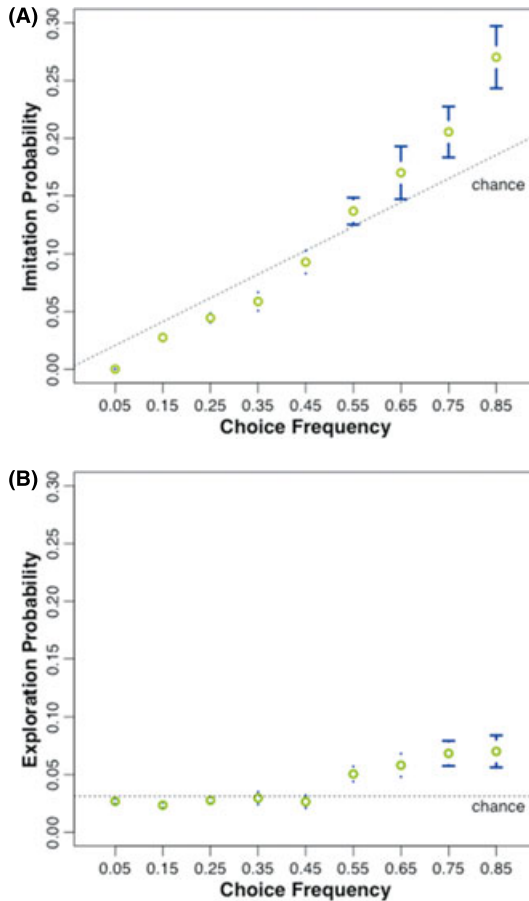
Fig. 3. There were biases toward choosing elements that were more frequently represented on other teams in (A) Imitation and (B) Exploration decisions, showing a *copy the majority* strategy.

The distribution of frequency changes for all icons was very nearly symmetrical around zero, such that an equivalent number of positive and negative proportion changes occurred, with small absolute changes more common than large ones. After log-transforming the Imitation probability data to achieve a normal distribution, a *t*-test of Imitation probability for negative and positive changes in choice frequency showed a significant positive momentum bias ($t(1236) = 18$, $p < .0001$; see Fig. 4A), and a non-significant momentum bias for Exploration (see Fig. 4B).

2.3.1.4. Similarity bias: A comparison between the mean similarity of participants' most recent guesses to those whom they Imitated, and to those whom they did not Imitate, revealed a small but significant difference: 0.550 for imitated versus 0.503 for non-imitated ($t(5029) = -7.10$, $p < .0001$; see Fig. 5). In other words, prior to imitation, the
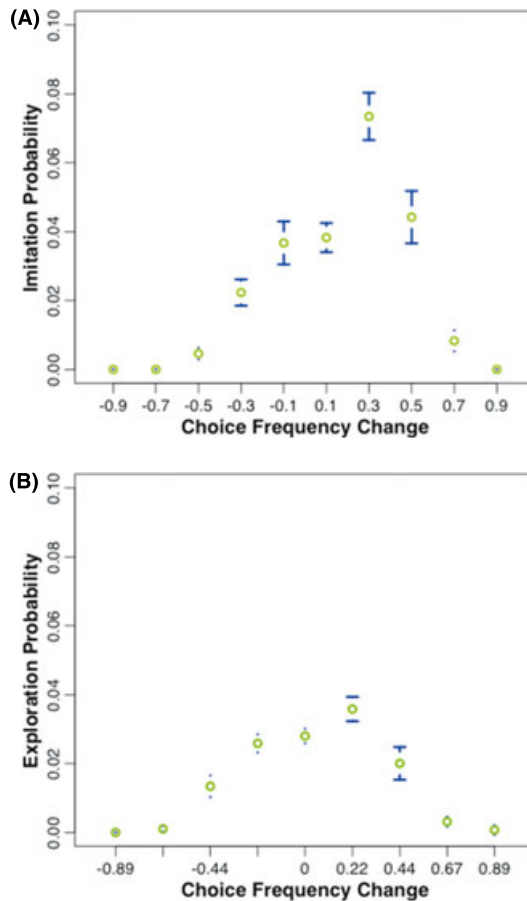
Fig. 4. There were positive momentum biases (toward choosing elements whose representation on other teams was increasing) in (A) Imitation but not in (B) Exploration decisions.

average imitators' guess was more similar to that of the imitated participant(s) than to those of others. This difference remained significant across rounds, even as overall solution similarity increased (i.e., as overall solution diversity decreased; see Fig. 6). No significant trends were observed in linear regressions of guess similarity versus imitated score rank, or similarity versus the score difference between imitator and imitated participants.

2.3.2. Main dependent variable analyses

Average choice source proportions are shown in Table 2. Note that the rows do not add up to 100% because a very small proportion (<1%) of choices were copied back from the player's own previous round team after initially being removed within the same round; these are excluded from subsequent analyses.
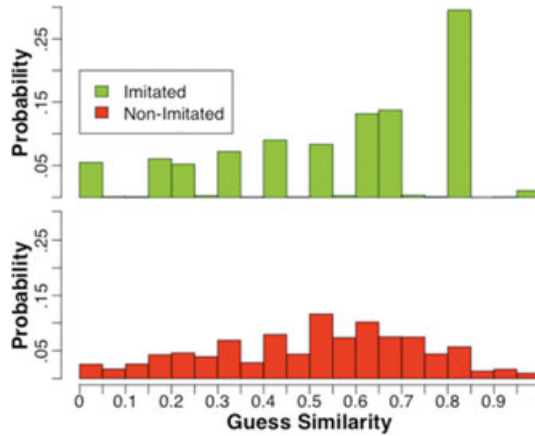
Fig. 5. Imitators' previous teams showed greater similarity to the teams they imitated than to those they did not imitate.
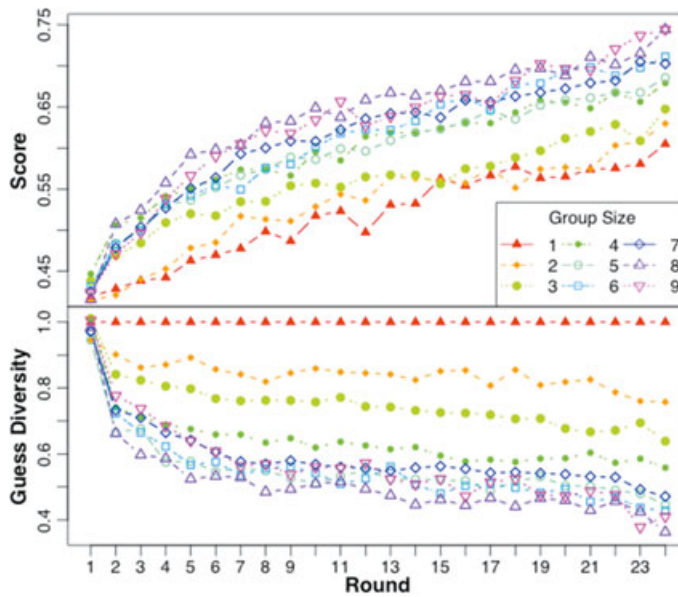

Fig. 6. Mean score increased and mean guess diversity decreased as more rounds were played within a game; stronger effects were observed for larger participant group sizes.

2.3.2.1. League size/difficulty: Participants achieved mean overall scores (averaged across all rounds) and mean final scores for each condition as shown in Fig. 2, with overall differences between conditions shown in Table 3.

2.3.2.2. Rounds, game order, and group size: Linear mixed-effects regression models were used to examine trends across rounds, game order, and group size for each

Table 2
Means (and standard deviations) of choice source proportion for grouped and isolated participants

Choice Source	Imitation	Exploration	Retention	Retrieval
Grouped	8.7% (4.0%)	14.7% (6.0%)	73.8% (6.5%)	2.3% (2.3%)
Isolated	n/a	20.6% (5.6%)	59.2% (18.4%)	19.4% (18.7%)

Table 3
Mean differences between conditions and analyses of all main dependent variables

	Difficulty	Round	Game Order	Group Size
Score	$t(305.8) = 7.6$ diff $= -0.067$***[a]	$F(1,919) = 897.1$ $B = +0.717$***	$F(1,272) = 14.7$ $B = +0.186$**	$F(1,37) = 73.6$ $B = +0.466$***
Guess diversity	$t(233.8) = 2.2$ diff $= -4.2\%$*	$F(1,735) = 188.6$ $B = -0.404$***	$F(1,216) = 20.0$ $B = -0.180$***	$F(1,29) = 38.3$ $B = -0.663$***
Imitation	n.s.	$F(1,681) = 126.0$ $B = -0.453$***	n.s.	$F(1,29) = 22.4$ $B = +0.565$**
Exploration	$t(306.8) = 2.7$ diff $= -1.9\%$**	$F(1,857) = 70.8$ $B = -0.277$***	n.s.	$F(1,37) = 29.0$ $B = -0.563$***
Retention	$t(307.5) = -2.0$ diff $= +2.6\%$*	$F(1,857) = 21.4$ $B = +0.214$***	n.s.	$F(1,37) = 12.1$ $B = +0.433$**
Retrieval	n.s.	$F(1,857) = 9.7$ $B = +0.128$**	n.s.	$F(1,37) = 12.5$ $B = -0.464$**

Note. [a]High – low difficulty. ***$p < .0001$, **$p < .01$, *$p < .05$.

dependent variable, with a random effect of session (see Table 3). Results are summarized briefly below. Scores increased significantly across rounds, game order, and group size, while guess diversity displayed opposite trends (see Figs. 6, 7, 8). As for choice sources, Imitation decreased significantly over rounds but increased in larger group sizes, Exploration decreased over rounds and in larger group sizes, Retention increased over rounds and in larger group sizes, and Retrieval increased over rounds but decreased in larger group sizes (see Figs. 9, 10). No significant trends were found for choice source proportions across game order.

To determine whether the score advantage for larger groups was simply an artifact of the greater chance of observing a better score than one's own given the larger number of guesses being made, we calculated the score difference variance (SDV): the variance of the differences between the top-ranked participant and all other participants within each round, averaged within each game. Using a linear mixed-effects model like the others used for group size analyses above, we confirmed a slight but significant upward trend of SDV across group size ($F(1, 29) = 11.37$, $p = .0021$, $B = 0.262$). However, a similar analysis of Imitation proportion versus SDV did not reveal any significant trend, and controlling for SDV in the Imitation proportion versus group size model above did not alter it significantly. In other words, participants in larger groups imitated each other more often, but this was not simply due to an increased chance of encountering better scores than one's own.
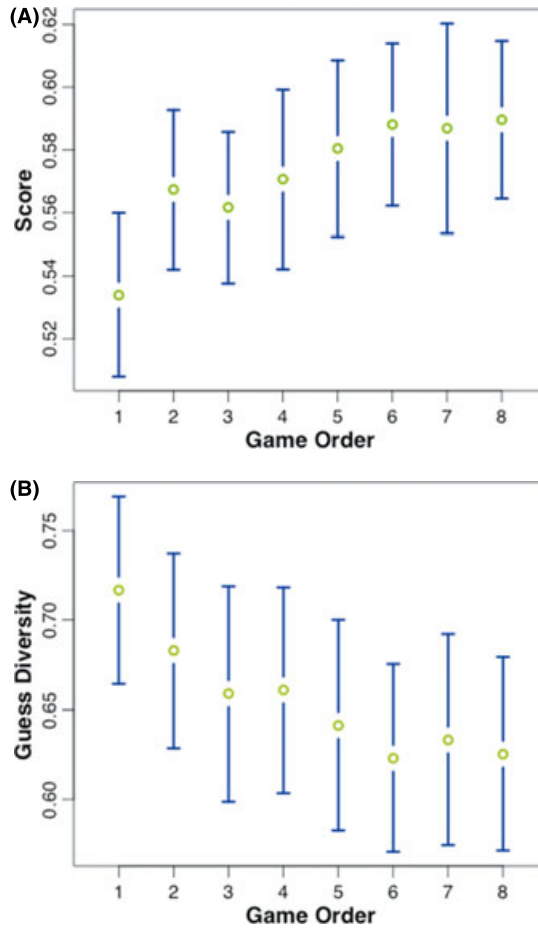
Fig. 7. As more games were played within an experimental session, (A) Mean score increased, and (B) the diversity of solutions decreased.

2.3.3. Learning strategies and performance

2.3.3.1. Choice source strategy: The choice sources of each non-isolated participant over the entire session were analyzed, and each participant's aggregate choice source strategy was categorized according to their proportion of each source. Participants whose choices contained one source in an average proportion greater than the global average for that source plus one standard deviation were labeled with that strategy. Those who fit the above criteria for more than one choice source, or none, were labeled as having a "Mixed" strategy. The score distribution for players in each strategy category is shown in Fig. 11, with the Retain strategy scoring the best, followed by Mixed, Imitate, and Retrieve, with Explore scoring the worst.

2.3.3.2. Individual and group strategy regressions: The above-mentioned figure also summarizes the results of simple regression analyses performed for score versus
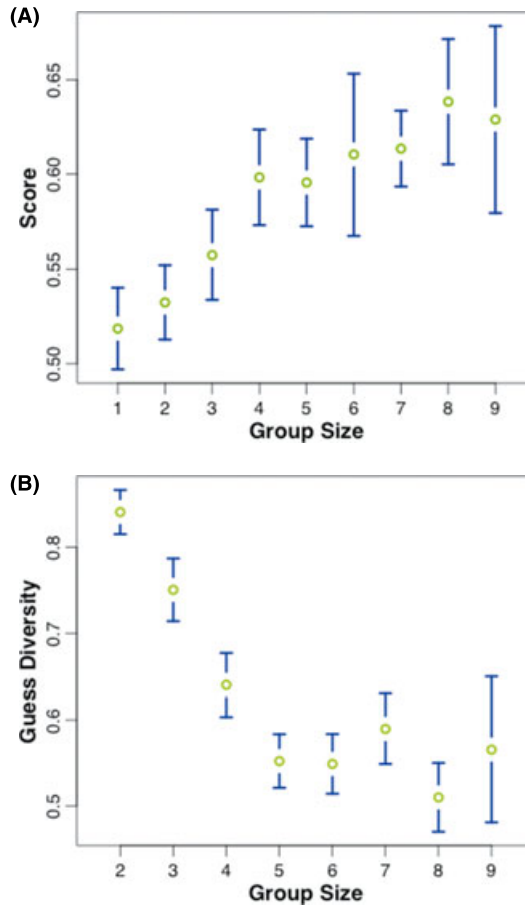
Fig. 8. (A) As participant group size increased, mean scores in a group increased, and (B) the diversity of solutions decreased.

individual and group use of each choice source. A linear regression of mean individual score versus mean individual Imitation proportion showed a significant positive trend—the greater a participant's average proportion of Imitation, the better the participant's score ($F(1,143) = 8.64$, $p = .0038$, $B = 0.239$). A similar positive trend held for Retention ($F(1,143) = 55.72$, $p < .0001$, $B = 0.530$). The opposite was true for individual score versus Exploration, which displayed a significant negative trend ($F(1,143) = 119.8$, $p < .0001$, $B = -0.675$), as did Retrieval ($F(1,143) = 10.93$, $p = .0012$, $B = -0.267$).

Very similar and significant patterns of results (all $p < .001$) were shown in analyses of mean group score versus mean group guess proportion for each choice source, even when each individual's data was excluded from their group's aggregate behavior. That is, an *individual's score* was higher when *the individual's fellow group members* Imitated
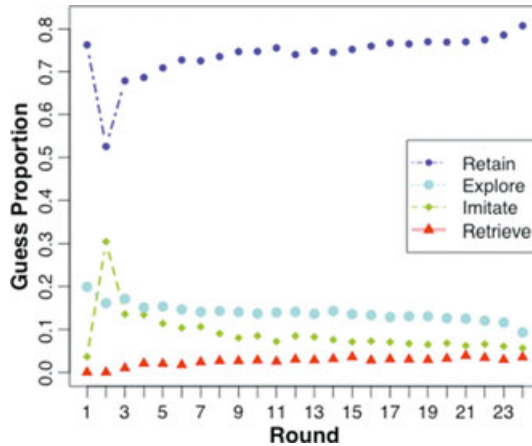
Fig. 9. Mean proportions of Retention and Retrieval increased and Imitation and Exploration decreased as more rounds were played within a game. Imitation jumps in the second round because participants can see their own and peers' scores for the first time.
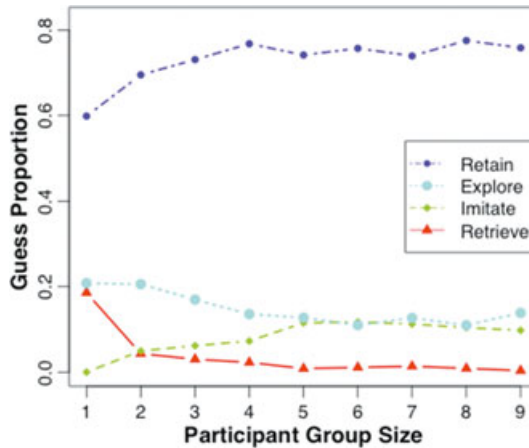

Fig. 10. As participant group size increased, mean proportions of Retention and Imitation increased, and Exploration and Retrieval decreased.

and Retained more, and Explored and Retrieved less. All trends noted above were monotonic; that is, there were no thresholds or inflection points beyond which the relationships changed.

The mean choice source proportions for guesses that resulted in score improvements and those that did not are shown in Table 4. These proportions indicate that Exploration was far more common in improvements than non-improvements. All other choice sources are less common in teams that were improvements, significantly so in the case of Imitation and Retention.
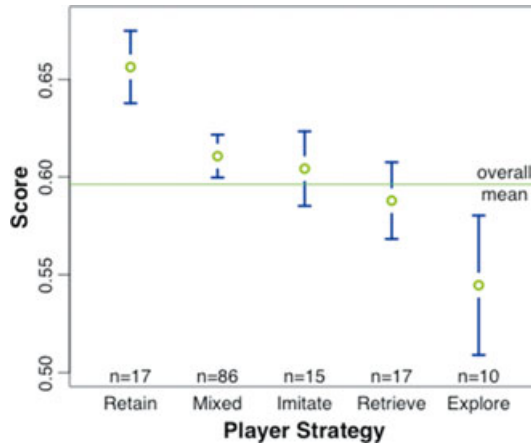
Fig. 11. Mean score versus choice source strategy, showing that cautious and imitative strategies resulted in the best overall performance.

Table 4
Mean (and standard deviation) choice source proportions for improvement and non-improvement guesses

Choice Source	Imitation	Exploration	Retention	Retrieval
Proportion in non-improvement guesses	9.9% (19.4%)	12.9% (14.6%)	74.9% (20.7%)	1.7% (7.4%)
Proportion in improvement guesses	7.9% (16.3%)	18.2% (11.4%)	72.2% (16.4%)	1.5% (6.4%)
Difference	**1.9%****	**−5.3%*****	**2.7%****	0.2%
	$t(504) = 2.72$	$t(549) = -10.32$	$t(547) = 3.68$	n.s.

***p < .0001, **p < .01, *p < .05.

2.4. Experiment 1 discussion

2.4.1. Specialized strategies

Results showed evidence for all expected social learning strategies. Payoff bias was shown in the tendency to imitate peers with higher scores. The small proportion of Imitation of non-top-ranked or even lower scoring peers were not explained by similarity between their guesses or other factors and were likely due to random errors.

Frequency bias was shown in the tendency to imitate solution elements according to their frequency in peers' solutions. The tendency for frequency-biased imitation to rise above the chance level only for solution elements with frequency greater than 50% is consistent with a *copy the majority* strategy, and furthermore with a strict definition of "conformity": not merely a tendency to follow the majority but an exaggerated tendency to do so, sufficiently strong to increase the size of the majority over time (Efferson,

Lalive, Richerson, McElreath, & Lubell, 2008). The greater incidence of imitation for solution elements that were increasing in frequency (relative to those that were decreasing) indicated participants were also sensitive to changes in frequency, not just magnitudes.

Similarity bias was evident in the greater similarity of participants' previous solutions to the peer solutions they imitated (relative to those they did not imitate). This allowed participants to associate differences between their own and similar others' solutions more accurately with payoffs, as well as use information from others' solutions without completely abandoning the knowledge of the problem space accumulated in their own previous solutions.

2.4.2. General strategies

The overall character of participants' learning was fairly cautious, as evidenced by the high mean proportion of Retention (which increased across rounds) and higher mean proportion of Imitation relative to Exploration; this cautious approach was accentuated in the higher difficulty condition in a lower mean proportion of Exploration. Individual guesses became increasingly entrenched over time, as evidenced by the decreasing proportions of Exploration and Imitation, and increasing proportions of Retention and Retrieval, across rounds. This behavior is consistent with the *copy when uncertain* strategy in that more imitation occurred early on in each game when participants had less experience with the current problem space. Guesses became entrenched at the group level across rounds as well (as shown by decreasing group solution diversity) despite decreasing amounts of imitation, because the remaining imitation was increasingly driven by convergent biases toward greater guess similarity, higher choice frequency, and positive choice momentum. These biases may also help explain the decrease in guess diversity in the greater difficulty condition without an accompanying increase in the incidence of imitation. Whereas Baron et al. (1996) found that increasing task difficulty increased the incidence of imitation, in this experiment it appears to have instead changed the focus of the imitation that occurred to favor increased group solution homogeneity.

2.4.3. Group size effects (and lack thereof)

The predicted increase in Imitation with larger group size (after accounting for artifactual score variance explanations), along with decreased Exploration and Retrieval, indicates a bias toward social learning that increased with the number of model solutions available, and the accompanying increase in score indicates that this was a beneficial strategy for this task. Conversely, the reduction in Retrieval with increasing group size indicates a greater dependence by isolated individuals and those in smaller groups on the built-in "memory" of the Best Score option in the game as a source of known good solutions on which to build. The combination of these results implies that in larger groups, this function of memory may be "outsourced" to others who imitate and thus propagate and preserve good solutions within the group. This can be seen as an example of socially distributed cognition (Hutchins, 1995), in which the functions of a cognitive process (such as memory) are enacted by multiple agents interacting dynamically via artifacts (in this case, the task display). A different version of the *copy when uncertain* strategy is

shown here: Imitation is favored when the payoff for Exploration is relatively uncertain, compared to the abundant information available about the content and related payoffs of neighbors' guesses.

2.4.4. Choice strategies and cumulative exploration

The relationship evident between performance and choice strategy, in which above-average Retention and Imitation produce higher scores, while above-average Exploration and Retrieval produce lower scores, indicates that the overall cautious (but not regressive) approach noted above is beneficial for this task. However, a counterpoint for this seemingly simple result is provided by the comparison of choice source proportions between solutions which generated improvements and those that did not, which showed that a substitution of significant amounts of Imitation and Retention with Exploration was required to create new and improved solutions. The fact that substantial amounts of each of the above three choice sources were present in such improved solutions shows that improvements were cumulative, relying on individuals' own past solutions as well as borrowing from others. This, in turn, implies that the adaptive value of Imitation in this context is due to its facilitation of selective learning and the generation of cumulative improvements using smaller amounts of risky Exploration (Boyd & Richerson, 1995; Kameda & Nakanishi, 2003).

2.4.5. Satisficing exploration levels

The decrease in Exploration and increase in Retention in larger groups suggest adaptations by group members to limit risky Exploration to what was required to achieve "good enough" results given the efforts of others. In fact, it may be that the lower end of the distribution of Exploration that actually occurred was nearly optimal for the very thin-tailed distribution of scores in the space of possible solutions (see Fig. 2)—the percentages of possible solutions that have higher scores than the participants' average final score in the lower and higher difficulty conditions are only 4.3% and 1.6%, respectively. This is consistent with the results of Kameda et al. (2011), in which the decreasing marginal returns to individual contribution produce a mixed equilibrium; however, in this experiment, rather than limiting the number of contributors, the mixed equilibrium took the form of limiting the amount of individual contribution of Exploration efforts. Conversely, the increased Exploration and corresponding lower performance of participants in smaller groups was likely due to having fewer peers to observe, thus requiring more costly Exploration per individual to achieve any improvements at all, which resulted in lower average performance.

3. Experiment 2: Score-visibility manipulation

3.1. Overview of changes from Experiment 1

The task used in this experiment was the same as in Experiment 1, with two major changes: (a) the scores associated with peers' solutions were hidden in half of the games

in each experiment session, and shown in the other half; (b) the problem space was changed by adding more positive-scoring bonus interactions between solution elements, which had the effect of making the upper tail of the score distribution longer and fatter, so that there were relatively more solutions with high scores. Modification (a) allowed for the examination of differences in strategies and performance associated with differences in the available social information. Modification (b) allowed us to evaluate potential "funnel effect" explanations of similarity bias and decreasing diversity, by allowing participants to achieve high scores without necessarily converging in the content of their solutions. Because the results of Experiment 1 were qualitatively very similar to the comparable condition in Experiment 2, this manipulation will not be discussed further.

From the research on animal imitation, there are reasons to believe that manipulating score visibility both would, and would not, influence imitation behaviors. On the one hand, the simple presence of an animal at a foraging site can increase the likelihood of other animals joining the animal, irrespective of their knowledge of the first animal's foraging success (Zentall, 2003). As applied to our paradigm, this result would suggest that our participants would continue to imitate others' teams simply based on their observation of these teams. Moreover, the absence of score information might cause our participants to amplify their use of content-related strategies such as frequency and similarity biases. On the other hand, animals do also use information about the foraging success of others to decide where to forage themselves (Smith, Benkman, & Coffey, 1999; Templeton & Giraldeau, 1996). For example, the observed correspondence between observer and demonstrator quails' foraging responses disappears when the observer does not see the reinforcement of demonstrator's responses (Akins & Zentall, 1998). This would suggest that our participants might decrease or eliminate imitation of peers' solutions in the absence of related score information. In contrast to previous research on animal imitation (Donchin, Giraldeau, Valone, & Wagner, 2004), our paradigm allows us to determine whether participants are sensitive to the solution outcomes of their peers in a symbolic and numeric form. Such symbolic and numeric information is vital to the combined payoff- and content-based social learning strategies discussed at the beginning of this article.

3.2. Methods

Overall, 234 participants were recruited from the Indiana University Psychology Department undergraduate subject pool as in Experiment 1 and were distributed across 65 sessions as shown in Table 5.

The task used was nearly identical to that of Experiment 1, with the following changes. To more easily fit the session in the 1-hour time limit required for experiments using our subject pool, there were six games per session instead of eight. In three of these games (the *visible-scores* condition), the scores of other participants were shown along with their solutions from the previous round (as in Experiment 1); in the other three games (the *invisible-scores* condition), other participants' solutions were shown, but not their scores. The games in each session were randomly ordered, and the visibility of peers' scores was the same for all participants and all rounds within each game.

Table 5
Distribution of participants across group sizes in Experiment 2

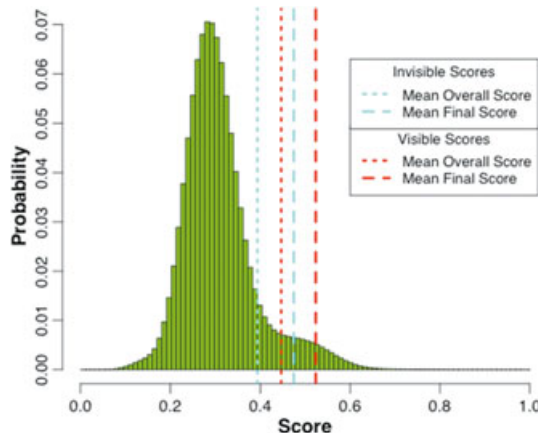| Group Size	1	2	3	4	5	6	7	8	9
No. of sessions	16	8	11	11	7	2	4	4	2
No. of participants	16	16	33	44	35	12	28	32	18

Fig. 12. Distribution of scores for all possible teams in Experiment 2. Note that, as in Fig. 2, the vertical lines show mean scores from the experimental data, not the mean of the distribution.

The distribution of individual point values for the icons was the same as for the larger league size in Experiment 1, but seven new positive bonus interactions were added between icons, and several existing interaction values were shifted to different pairs of icons. These changes had the effect of increasing the complexity of the problem space, as well as increasing the proportion of possible high-scoring teams. As a result, the possible score range changed to (−6,88), but as in Experiment 1, all scores were normalized to the range (0,1) for ease of analysis (but note that due to the above changes, normalized scores cannot be directly compared between Experiments 1 and 2). The combinations of these individual and pair values resulted in the probability distribution of scores among all possible teams shown in Fig. 12 (compare to Fig. 2B).

3.3. Experiment 2 results

The organization of this section is identical to that of Experiment 1: We present summaries of (a) evidence for specialized social learning strategies; (b) dependent variables in aggregate and across rounds, game order, and participant group size; and (c) examinations of learning strategies in relation to individual and group performance.

3.3.1. Specialized social learning strategies
This section presents results regarding social learning biases according to solution payoff, solution element frequency (and changes thereof), and solution similarity.
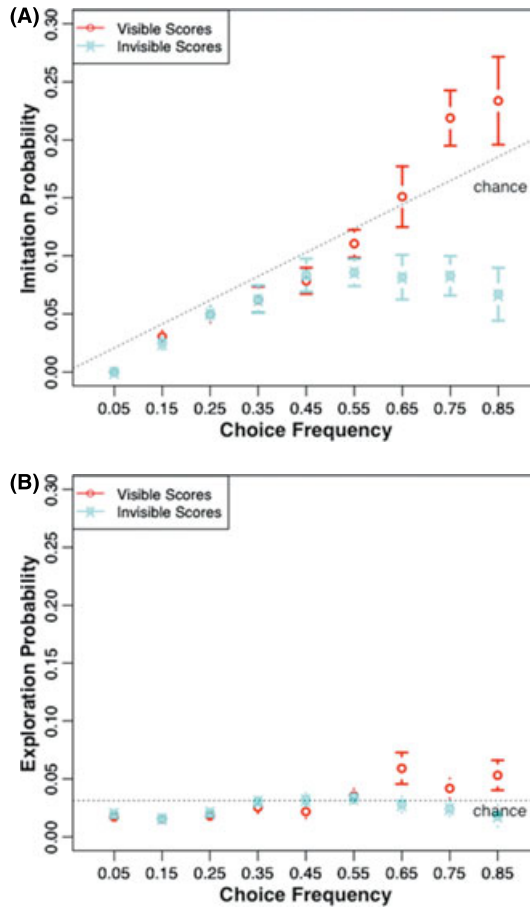
Fig. 13. There were biases toward choosing elements that were more frequently represented on other teams in the visible-scores condition, and less frequently represented on other teams in the invisible-scores condition for (A) Imitation and (B) Exploration decisions.

3.3.1.1. Payoff bias: Of all instances of single-participant imitation, the score of the imitated participant was greater than that of the imitator significantly more often in the visible-scores condition ($t(74) = 16.07$, $p < .0001$); in the invisible-scores condition, the probability was about 54%, or approximately at chance, as would be expected.

3.3.1.2. Frequency bias: As in Experiment 1, we measured the bias of participants to choose an icon according to its frequency in peers' choices. Linear mixed-effects analysis of imitation probability versus choice frequency showed a positive frequency-dependent Imitation bias that was significantly greater than chance in the visible-scores condition ($F(1,604) = 943.25$, $p < .0001$, $B = 0.741$), but significantly lower than chance in the invisible-scores condition ($F(1,604) = 231.67$, $p < .0001$, $B = 0.470$; see Fig. 13A). Likewise, there was a positive frequency-dependent Exploration bias above chance in the
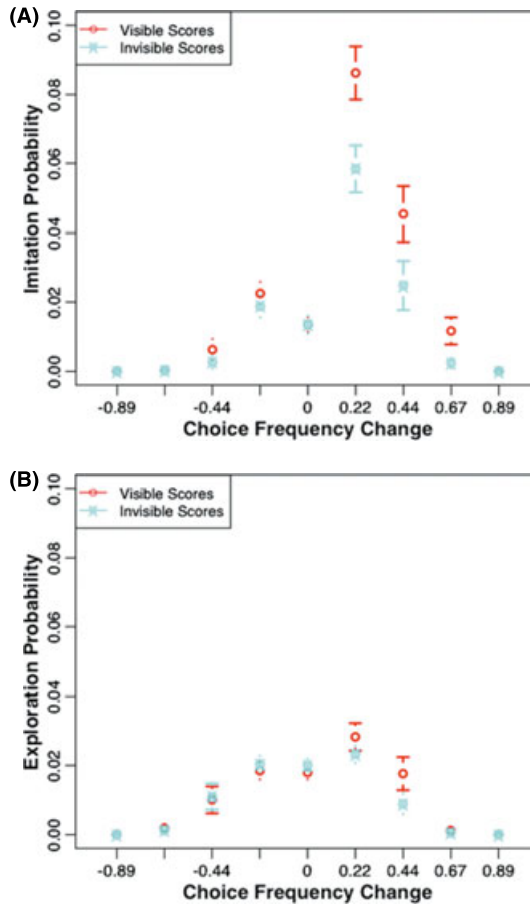
Fig. 14. There were biases toward choosing elements whose representation on other teams was increasing in (A) both visibility conditions for Imitation and (B) only the visible-scores condition for Exploration decisions.

visible-scores condition ($F(1,604) = 181.20$, $p < .0001$, $B = 0.441$) and below chance in the invisible-scores condition ($F(1,604) = 12.78$, $p = .0004$, $B = 0.131$; see Fig. 13B).

3.3.1.3. Frequency-change (momentum) bias: We also repeated the analysis of "choice momentum," by tallying the change in the number of players whose teams included the icon in the previous two rounds, as well as the number of the remaining players who added it to their team in the current round via Imitation or Exploration, and normalizing for group size. After log-transforming the Imitation probability data to achieve an approximately normal distribution, a *t*-test of Imitation probability for negative and positive changes in choice frequency showed a significant positive momentum bias in the visible-scores condition ($t(640) = -14.192$, $p < .0001$), and a smaller positive bias in the invisible-scores condition ($t(661) = -9.98$, $p < .0001$; see Fig. 14A). A slight positive
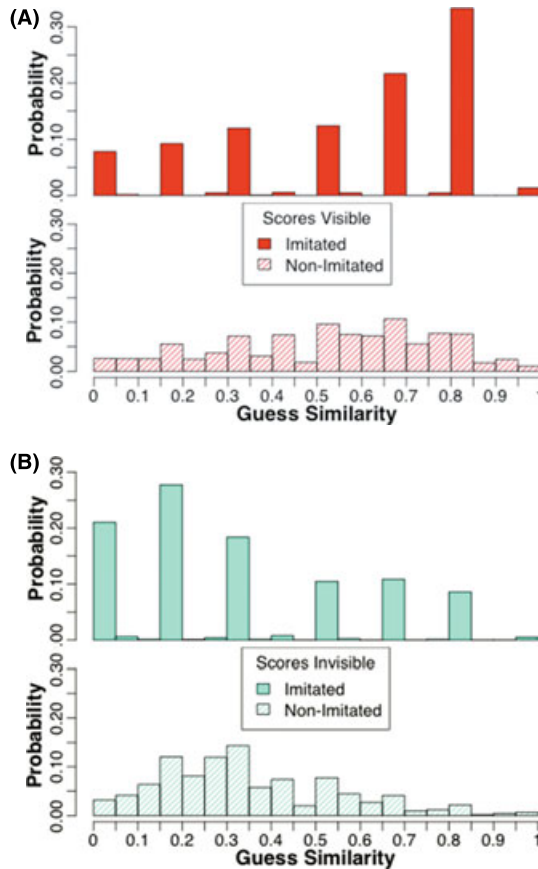
Fig. 15. (A) In the visible-scores condition, imitators' previous guesses showed greater similarity to the guesses they imitated than to those they did not imitate, while (B) in the invisible-scores condition, the opposite effect was observed.

momentum bias was found for Exploration in the visible-scores condition, but no corresponding significant bias was found in the invisible-scores condition (see Fig. 14B).

3.3.1.4. Similarity bias: A comparison between the mean similarity of participants' most recent guesses to those whom they imitated, and to those whom they did not imitate, revealed significant differences in both conditions, but in opposite directions. In the visible-scores condition, there were similarity values of 0.563 for imitated versus 0.524 for non-imitated guesses ($t(5084) = -5.47$, $p < .0001$; see Fig. 15A), replicating the difference found in Experiment 1. In the invisible-scores condition, the opposite was found: There were similarity values of 0.316 for imitated versus 0.346 for non-imitated guesses ($t(4267) = 4.35$, $p < .0001$; see Fig. 15B). In other words, prior to imitation, the average imitators' guess was more similar to that of the imitated participant(s) than to those of others in the visible-scores condition, and less similar in the invisible-scores condition.

Table 6
Mean differences between conditions, and analyses of dependent variables (invisible-scores only)

	Score Visibility	Round	Game Order	Group Size
Score	$t(345.9) = -5.8$ diff $= -0.053$***[a]	$F(1,1126) = 446.5$ $B = +0.727$***	n.s.	$F(1,63) = 15.5$ $B = +0.309$**
Guess diversity	$t(248.6) = 9.8$ diff $= +16.8\%$***	$F(1,1126) = 97.3$ $B = -0.453$***	$F(1,97) = 70.8$ $B = -0.363$***	$F(1,47) = 17.3$ $B = -0.430$**
Imitation	n.s.	$F(1,1077) = 33.3$ $B = -0.219$***	$F(1,97) = 33.0$ $B = +0.289$***	$F(1,47) = 28.0$ $B = +0.500$***
Exploration	n.s.	$F(1,1429) = 134.3$ $B = -0.237$***	$F(1,97) = 41.0$ $B = -0.226$***	$F(1,47) = 14.0$ $B = -0.436$**
Retention	$t(384.3) = -2.5$ diff $= -4.0\%$*	**n.s.**	**n.s.**	**n.s.**
Retrieval	n.s.	n.s.	$\mathbf{F(1,97) = 14.4}$ $\mathbf{B = +0.228}$**	n.s.

Note. Results that differed from the visible-scores condition are in bold.
[a]Invisible-scores – Visible-scores. ***$p < .0001$, **$p < .01$, *$p < .05$.

3.3.2. Main dependent variable analyses

3.3.2.1. Score visibility: Mean overall and final scores in each condition are shown in Fig. 12, with differences between conditions shown in Table 6.

To examine separately how *often* and how *much* participants imitated one another, we measured the mean proportion of solutions in which there was greater than zero Imitation (*Imitation incidence*), as well as the mean Imitation proportion in such cases (*Imitation proportion*). Mean Imitation incidence was significantly higher in the visible-scores condition ($F(1,229) = 31.17$, $p < .0001$), but the distribution of mean Imitation proportions was weighted significantly more heavily toward higher values in the invisible-scores condition, as shown by a Kolmogorov–Smirnoff test of equality of distributions ($D = 0.1893$, $p < .0001$; see Fig. 16). In other words, participants in the invisible-scores condition copied one another less frequently but in larger amounts at a time.

3.3.2.2. Rounds, game order, and group size: As in Experiment 1, linear mixed-effects regression models were used to examine trends across rounds, game order, and group size for each dependent variable, with a random effect of session (see Table 6). Trends for the visible-scores condition were very similar to those in the high-difficulty condition of Experiment 1 and are omitted from Table 6. Trends that changed substantially in significance or direction between the two conditions are shown in bold in Table 6 and are summarized briefly below. Plots of trends for both conditions are shown in Figs. 17–21.

Scores increased over game order in the visible-scores condition, but not in the invisible-scores condition. Imitation decreased over game order in the visible-scores condition but increased in the invisible-scores condition. Retention increased over all three variables in the visible-scores condition but showed no trends in the invisible-scores condition. Finally, Retrieval showed no trends in the visible-scores condition but increased over game order in the invisible-scores condition.
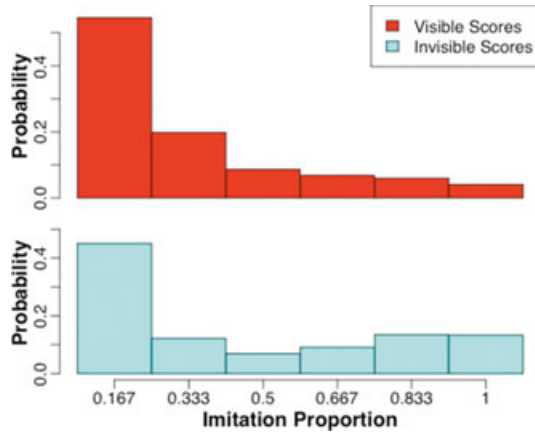
Fig. 16. For guesses that included at least some Imitation, participants in the invisible-scores condition had higher proportions of Imitation in their guesses.
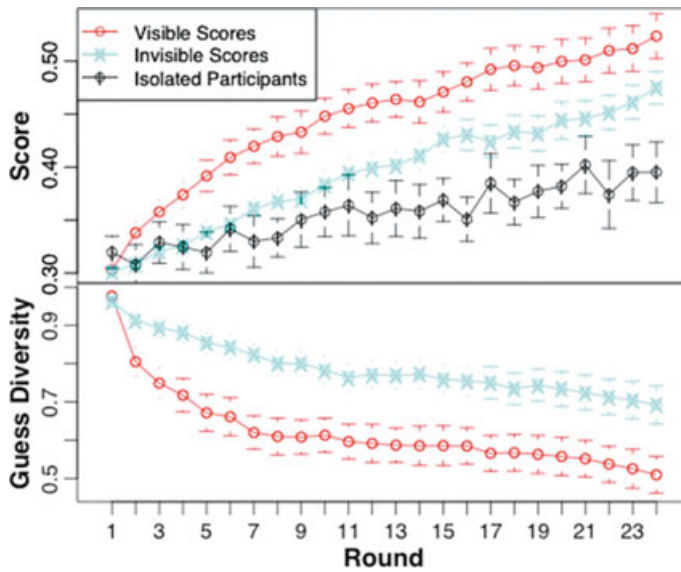

Fig. 17. Scores increased and guess diversity decreased more strongly across rounds in the visible-scores condition than in the invisible-scores condition.

3.3.3. Learning strategies and performance

3.3.3.1. Choice source strategy: As in Experiment 1, the choice sources of each non-isolated participant over the entire session were analyzed, and each participant's choice source *strategy* was categorized according to their proportion of each source. The score distribution for each strategy in each condition is shown in Fig. 22.
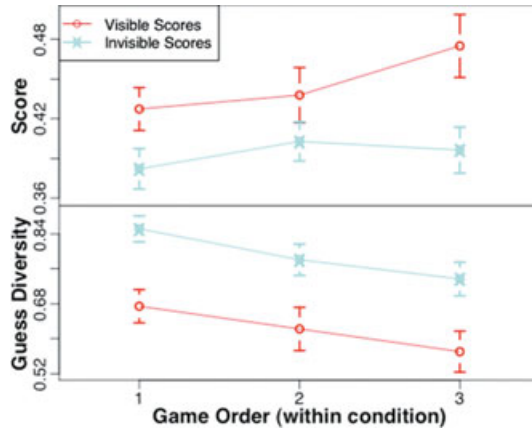
Fig. 18. Score increased significantly only in the visible-scores condition, and guess diversity decreased in both conditions.

3.3.3.2. Improvements: The mean choice source proportions for guesses that resulted in score improvements and those that did not are shown in Table 7. Examining only the solutions that resulted in improvements, 24.2% resulted from guesses that included an Imitation proportion greater than zero in the visible-scores condition, versus 12.2% in the invisible-scores condition. In 52.3% of all improvements in the visible-scores condition, the focal player imitated at least one peer who had previously imitated the focal player, versus 41.5% in the invisible-scores condition. In other words, a player who was imitated by another player often later imitated that same player in the course of creating an improvement, but this happened substantially less often when scores were invisible.

To examine the relative equity of improvement achievement within groups, we defined each participant's *normalized improvement share* as his or her individually achieved proportion of the total improvements achieved by all participants in a session, multiplied by the number of participants in the session. A value of 1 indicated a "fair" share, for example, a participant achieved one third of the improvements in a three-person session. A histogram of normalized improvement share (see Fig. 23) in the visible-scores condition showed a relatively equitable distribution of improvements within groups, with a distribution strongly peaked near a "fair share" of 1 (56% of participants were between 0.4 and 1.2), and only 6.4% of participants having zero improvements (very similar results were found in the data for Experiment 1). In contrast, there was a strongly inequitably skewed distribution in the invisible-scores condition, with only 36.2% of participants having improvement shares between 0.4 and 1.2, and 21.1% having zero improvements. A Kolmogorov–Smirnoff test of equality of distributions indicated that these distributions were significantly different ($D = 0.1789$, $p = .002$). Mean overall score showed a strong positive correlation with improvement share in the invisible-scores condition ($F(1,168) = 64.49$, $p < .0001$, $B = 0.369$), but not in the visible-scores condition.
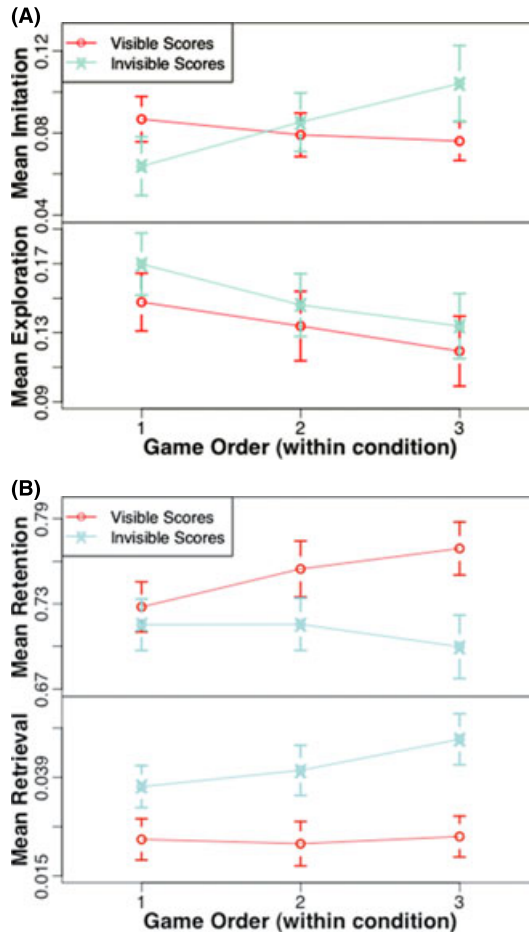
Fig. 19. (A) Imitation and Exploration decreased significantly in the visible-scores condition, while Imitation increased and Exploration decreased in the invisible-scores condition. (B) Retention increased significantly only in the visible-scores condition, and Retrieval increased significantly only in the invisible-scores condition.

3.4. Experiment 2 discussion

3.4.1. Changes in specialized strategies

As predicted, results in the visible-scores condition were very similar to those in Experiment 1, including the Imitation biases toward high-payoff solutions, high-frequency solution elements, and similar solutions. The invisible-scores condition had the predictable effect of eliminating the payoff bias but caused the two content-based strategies to shift in counterintuitive ways. Rather than strengthening these biases toward socially mediated information (as suggested by Abrahamson and Rosenkopf [1997] and Gibbons [2004]), participants actually showed weakened or opposite inclinations. One possible
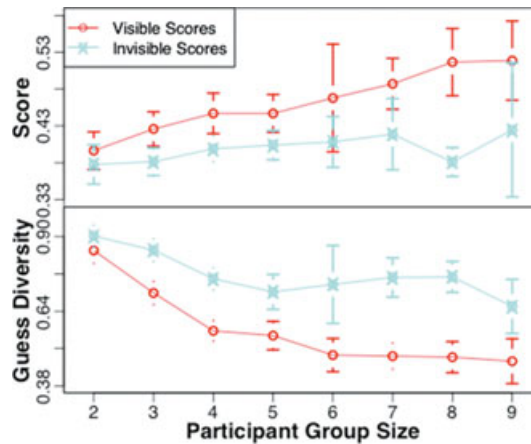
Fig. 20. As participant group size increased, mean scores in a group increased, and the diversity of offered solutions decreased, with both effects weaker in the invisible-scores condition.

explanation for the unpopularity bias and reduced momentum bias we observed is that participants in the invisible-scores condition knew that imitation decisions were not based on reliable performance information, and thus frequency-based biases should be avoided to keep from joining information herds (Banerjee, 1992), as suggested by Giraldeau, Valone, and Templeton (2002).

However, the presence of a dissimilarity bias in combination with the increase in Imitation proportion and decrease in mean Retention suggests an alternate interpretation. Overall, making score information unavailable seemed to shift participants' tactics from *incrementalist* strategies (making small changes to their guesses, informed by frequency- and similarity-based comparisons associated with payoff information) in the visible-scores condition, to *saltationist* strategies (making larger jumps around the problem space, evaluating each individually, and often jumping back to previous known good solutions) in the invisible-scores condition. A reliance on large risky jumps around the problem space would likely pay off about half the time, and those who jumped back to good previous solutions would lose less overall than those who continually jumped around; thus, the invisible-scores conditions showed shallower increases in score over rounds, and an association of higher scores with Retrieval but not Imitation. This saltationist strategy seems to have been more successful than not using Imitation at all, however, as shown by the substantially lower performance of isolated participants. Our results are accordance with the findings of Giraldeau et al. (2002), who found that an inability to combine the use of social and asocial learning simultaneously would result in a lack of benefit for social learning. We are not aware of other results regarding this kind of shift in social learning strategies, but it bears further study.

3.4.2. Imitation and performance

Having demonstrated benefits for Imitation in the previous experiment, the impediment to social learning we introduced in the invisible-scores condition lowered performance as
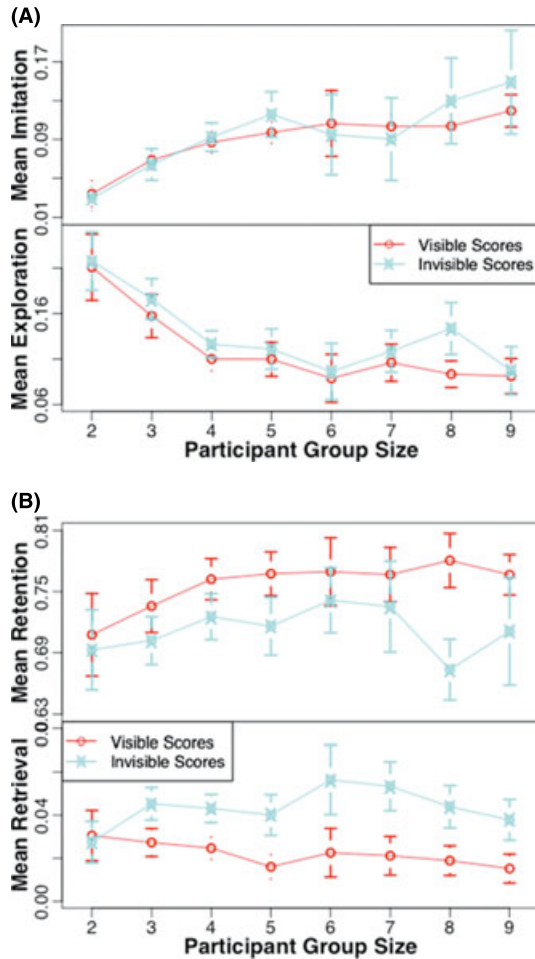
Fig. 21. As participant group size increased, (A) mean proportions of Imitation increased and Exploration decreased in both conditions, and (B) Retention increased only in the visible-scores condition, and Retrieval showed no significant change across group size.

predicted. The reduction in the efficiency of social learning implemented by hiding peers' scores did lead to increased solution diversity but did not improve collective search performance. This is in contrast to the results found in simulations by Lazer and Friedman (2007) and an experiment by Mason, Jones, and Goldstone (2008), in which individuals embedded in various network structures explored a problem space and could view the guesses and feedback of their neighbors. Both studies found that inefficient (less-connected) network structures resulted in better group performance for problems that required broad exploration. This difference between their results and those in the current experiment was likely due to the way that communication efficiency was reduced: Whereas they decreased the connectivity of the social network through which information was
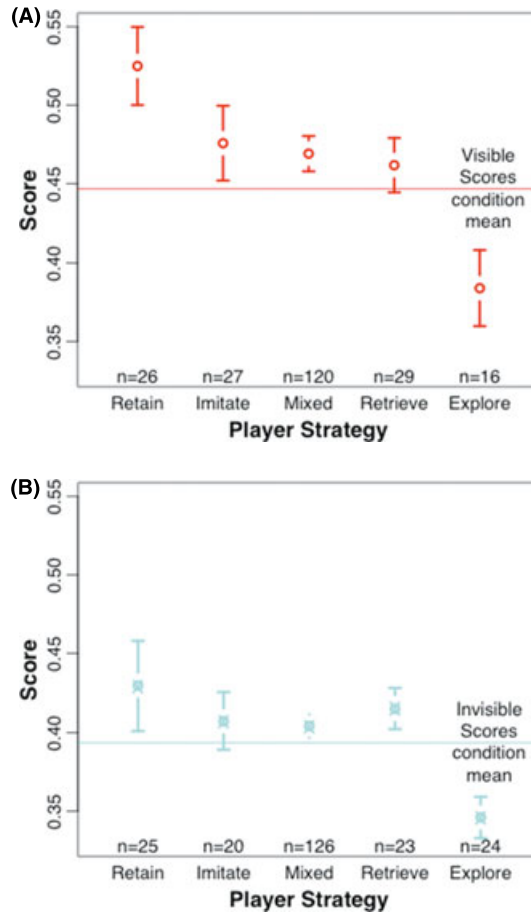
Fig. 22. Score versus choice source strategy in (A) visible-scores and (B) invisible-scores conditions, showing that a cautious high-Retention strategy resulted in the best performance, though a similarly cautious high-Retrieval strategy (returning often to a personal best-so-far) showed good relative performance in the invisible-scores condition.

Table 7
Mean (and standard deviation) choice source proportions for improvement and non-improvement guesses in each condition

Condition	Impr.	% of Guesses	Imitation	Exploration	Retention	Retrieval
Visible scores	No	94.6%	9.1% (19.5%)	11.4%* (13.4%)	76.3%* (21.8%)	2.2% (7.4%)
	Yes	5.4%	8.2% (18.5%)	19.4%* (12.8%)	69.5%* (20.0%)	2.1% (6.7%)
Invisible scores	No	95.6%	10%* (24.2%)	13.3%* (15.1%)	71.2% (25.6%)	4.4% (13.4%)
	Yes	4.4%	3.9%* (12.5%)	21.6%* (14.8%)	70.5% (16.7%)	3.5% (9.0%)

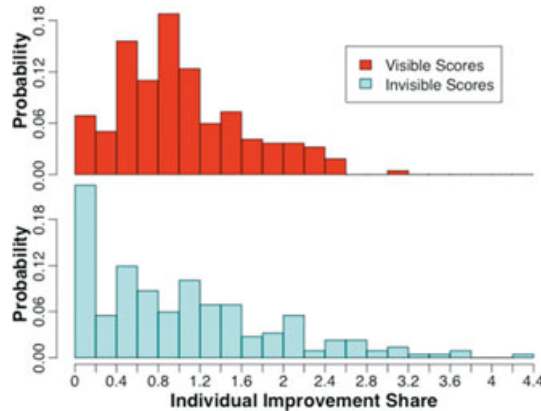*Note*. *Significant differences within condition, $p < .005$.

Fig. 23. Histograms showing relatively equitable achievement of improvements within groups in the visible-scores condition, and an inequitable distribution in the invisible-scores condition.

exchanged, we left the network unchanged but eliminated an important part of the information that participants used to guide imitation decisions—the actual results of peers' solutions. Mason et al. found that limiting participants' ability to see other participants' solutions was effective in avoiding premature group convergence on good, but not great, solutions. In contrast, our participants still had full knowledge of every other participant's solution. However, the participants in our invisible-scores condition could not selectively imitate the best solutions of the peers whose solutions they could see, which is a disadvantage that Mason et al.'s participants did not share.

More recently, Mason and Watts (2012) have found that more efficient networks resulted in better group performance than less efficient networks. They manipulated efficiency by the characteristic length of the path needed to connect pairs of participants. Although their results are consistent with our present results in showing that better communication networks improve group performance, our results are different with respect to the group consequences of individual decisions to imitate. Whereas Mason and Watts found that imitation in individuals was associated with better individual performance (as do we), they found that group performance suffered. In contrast, in our current experiments, individual imitation is associated with better group performance. One likely difference of consequence is the dimensionality of the search problems. Mason and Watts employed a spatial search task similar to earlier collective foraging tasks (Goldstone & Ashpole, 2004; Goldstone, Ashpole, & Roberts, 2005). Although their search space was "rugged" in the sense of having many local maxima, it was only a two-dimensional 100×100 grid. In our studies, the search space was much larger. Even the small search space from Experiment 1 consisted of teams of 5 icons chosen without replacement from a set of 24 candidates, or 42,504 possible teams, while the corresponding number for our larger search space (6-icon teams from a set of 48 candidates) was greater than 12 million. Furthermore, our task was designed so that solutions would naturally build on top of previous solutions. For this reason, imitators are not simply acting as "informational

parasites." They are maintaining an "institutional memory" for previously found good solutions while also tweaking the solutions and thereby possibly improving them. This incremental and ratcheting dynamic is much more pronounced for our relatively complex search space compared to a 2-D spatial foraging task.

3.4.3. Cumulative mutual improvement

The significant correlation of improvement share with mean scores in the invisible-scores conditions shows that individuals who were relatively more skillful (or lucky) were rewarded with proportionately better overall scores compared to others. This was because their fellow players could not easily distinguish their improvements and thus match their payoffs; it is also likely that the unreliability of Imitation made some participants more likely to seek out improvements on their own. In the visible-scores condition this correlation between improvement share and mean score disappeared, but the more equitable distribution of improvements showed that more participants were contributing to their discovery, and mean scores increased significantly such that nearly all participants did better. In other words, when social learning was unimpeded in the visible-scores condition, high and low individual achievers had similar payoffs, but both had higher payoffs compared to the invisible-scores condition, in which high-achievers kept a bigger piece of a smaller pie. Thus, impeding social learning led to relatively greater inefficiency and inequity, and lower long-term performance.

This advantage for more efficient social learning accrued because imitators were not merely scroungers; the substantial proportion of Imitation present in improvements shows that imitated guesses were often the basis for further productive exploration. The cumulative exploration hypothesis is supported by the fact that a larger proportion of improvements were the result of mutual Imitation in the visible-scores condition, in which solution elements were passed between players via copying and built into better solutions in the process. In the invisible-scores condition, the necessity of adopting others' guesses to obtain information about their performance allowed fewer opportunities to evaluate variations on them; it also prevented group members from performing the "filtering" function of copying and consistently retaining only solution elements associated with relatively high scores, so that others would have less chance of copying low-scoring solution elements.

3.5. Experiment 2 conclusions

In Experiment 1, we noted that social loafing strategies did not overwhelmingly dominate participants' behavior, despite the apparent incentive for individuals to under-produce Exploration and free-ride on others via self-interested Imitation. In Experiment 2, more social-loafing-like strategies were found (e.g., larger imitation proportions and fewer cumulative improvements) in the invisible-scores condition by reducing the capacity of individuals to make self-interested imitation decisions, even though they produced more (unproductive) Exploration.

The decreased performance in this task was apparently not due to an underprovision of individual Exploration, but a lack of accompanying evaluative basis for comparing

similarities and differences across solutions, and thus a lack of payoff-based filtering of solution elements. The consistent use of better-performing solutions does not rely on altruistic or publicly minded motives, but such filtering is important for supporting others' successful social learning, as was highlighted in a recent tournament of simulated social learning strategies (Rendell et al., 2010). In summary, we have shown through this experiment that when knowledge is cumulative, efficient, and informed, appropriation is an important step in further provision of the public good of knowledge.

4. General discussion

4.1. Factors that influence imitation—when, whom, and which to imitate

In this study, we found evidence for several social learning strategies that have so far been understudied in humans. The relative reliability of payoffs from social and asocial learning was used in two ways: the higher proportions of Imitation relative to Exploration can be interpreted as a *copy when asocial learning is costly* strategy, and the decrease in Imitation over time can be interpreted as a consequence of a *copy when uncertain* strategy. In Experiment 1 and the visible-scores condition of Experiment 2, participants almost universally showed payoff bias in their social learning, employing *copy when better* and *copy the best* strategies. In these conditions, participants also displayed frequency (i.e., popularity) bias in the form of a particularly strict *copy the majority* strategy, in which solution elements were copied disproportionately often only when greater than 50% of peers possessed them. This shows an ability to identify robust general agreement among multiple peers about element-level contributions to higher payoffs. An opposite *avoid the majority* strategy seemed to be prevalent in the invisible-scores condition of Experiment 2; this may have been simple "herd avoidance" or a consequence of other tactics. Interestingly, however, participants in both score-visibility conditions showed a "momentum bias" toward Imitation of solution elements that were increasing in overall frequency in the group rather than decreasing. This vitiates the "herd avoidance" explanation above and corresponds to similar patterns shown in baby-naming decisions by parents in an examination of 130 years of social security data (Gureckis & Goldstone, 2009).

Unlike previous experiments in which participants choose one option from a set of possibilities, our task design allowed participants to pursue hybrid strategies within a single round, in which they retained some elements of a solution while changing others using both social and asocial learning. One particularly interesting way this was used was similarity-biased imitation. This strategy allows imitators to judge the effect of small changes in solution elements on payoffs, and thus to improve their own payoffs while acquiring new solution elements that are compatible with their previous solution and knowledge of the problem space. This phenomenon is discussed at length by Rogers (2003) in relation to innovation propagation. A bias toward borrowing from similar rather than dissimilar solutions has also been incorporated into general machine learning algorithms featuring multiple agents simultaneously searching for solutions (Goldberg, 1989; Goldstone &

Janssen, 2005). There are two possible drawbacks when agents borrow solution elements from other agents pursuing substantially different solutions: First, they abandon the knowledge of the problem space accumulated in their previous solution; second, there is a strong risk that the resulting blend of solutions will be a suboptimal hybrid not well adapted to the niche of either of the original solutions. Given the complex search landscapes used in the experiments, participants may have been biased to copy solution elements from similar rather than dissimilar solutions to ensure greater solution compatibility.

The fact that the similarity and frequency biases were reversed in the invisible-scores condition of Experiment 2 suggests that any non-strategic (e.g., affiliative) motivations for these behaviors were minor, and that accompanying payoff information was vital to the usefulness of similarity and frequency information in decision making. The alternate *saltationist* behavior in this condition suggests that without peer payoff information, participants may consider incremental moves in the search space less effective than riskier large jumps. It is interesting to note that neither frequency nor similarity information is inherently "social"—that is, alternate solution information used for comparison need not be generated by the efforts of actual peers with social interaction. However, one way to repeatedly generate a variety of solutions with relatively small, evaluable commonalities and differences is to have a group of agents recurrently imitating each other and making changes whose results are visible to others. This has the additional quality of being compatible with individual incentives for finding improvements.

4.2. Group-level effects of imitation and exploration

The results of both experiments show that imitation can be productive for groups as well as individuals, because it enables the preservation of good tentative solutions in "group memory" and their further improvement through cumulative exploration. These results also showed that the pursuit of larger amounts of exploration can result in diminishing returns for both individuals and groups. With full information, as the imitation rate of one's peers increased, one's own score increased, and as the exploration rate of one's peers increased, one's own score decreased. A complete lack of exploration will of course result in a lack of improvements, but this experiment suggests that in a large and complex problem space, productive exploration may be readily incentivized by the potential for generating small improvements based on peers' solutions. This is analogous to the mixed equilibrium for individual contributions to group efforts found by Kameda and Tindale (2006).

These results regarding the group benefits of imitation and collective risks of exploration, taken together with the reductions in diversity over time, imply a view that is at odds with those predicted from a simple producer-scrounger dilemma interpretation of social learning (Kameda & Nakanishi, 2002). Much like "conformity," being a "scrounger" often carries a negative connotation or denotation, such as "social loafing" (Latané et al., 1979). However, such behavior may be appropriate when not all group members' full efforts are required to produce sufficient benefit. In a complex but relatively stable environment, the

best outcome for the group may result from most group members converging on a "good enough" solution quickly to achieve high mean performance, and then introducing productive exploration when necessary. Mesoudi (2008) reached a similar conclusion, finding that despite the possibility for information scrounging behavior to inhibit exploration, participants flexibly switched between individual and social learning to produce good performance. Thus, in some circumstances a suboptimal outcome can result from too much exploration and not enough exploitation, rather than the other way around, because exploration is risky and possibly redundant (and thus wasteful of resources), and imitation helps to concentrate efforts and improve the thoroughness of search in the proximity of known good solutions. Given a baseline inclination to some amount of individual exploration, the limiting factor in improving search performance may be the amount of information sharing and coordination among searchers, which allow them to pool both the benefits and the risks of asocial learning (Hess & Ostrom, 2007).

4.3. The consequences of groups for individual cognition

Much of cognitive science tacitly adopts the assumption of individual people as the unit of cognition. In fact, there are near ubiquitous influences of an individual's social community on the individual's cognition. Structured cognitive behavior can be described at multiple levels, and our thoughts both depend on and determine the social structures that contain us as elements. Often times, an apt unit of analysis is a group of people rather than an individual person (Hutchins, 1995; Louwerse, Dale, Bard, & Jeuniaux, 2012). Even when a thought is possessed by an individual and has a content that is not directly social, the nature of that thought is influenced by others. For example, the complexity of language as understood and spoken by an individual is a function of demographic and historical factors such as the number of speakers of the language (Lupyan & Dale, 2010). Everyday concepts like *computer virus*, *fraction,* and *appetizer* would almost certainly never have been created by an individual living in complete isolation. The emergence of complex symbol systems may in fact be substantially constrained by individuals' opportunities for joint mutual editing of each other's contributions, as found by Healey et al. (2007) using an experimental study of graphical interaction.

Results from the reported experiments suggest specific mechanisms by which an individual's community shapes an individual's cognition, in particular, their search through a high-dimensional space of solutions. First and most obviously, the quality of a person's solutions is positively affected by being in a community of other individuals trying to solve the same problems. The prevalence and positive results of reciprocal imitation point to the beneficial ratchet of improvement that results from individuals mutually benefiting from, and building upon, each others' solutions (Tomasello, 1994). Second, individual imitation biases interact with social context to favor promising search directions. For example, the individual bias to imitate solutions that are increasing in prevalence within the group allows the individual to quickly latch onto promising solutions. Individuals use frequencies and changes of frequencies within a group's solutions to prioritize their own searches, under the appropriate assumption that others are attracted to good solutions.

Third, cognitive decisions to imitate or explore are influenced by group factors such as the diversity, quality, and stability of solutions. Fourth, peers can fill some of the roles usually filled by cognitive components within an individual (Theiner, Allen, & Goldstone, 2010). For example, the need for an individual to personally return to promising previous solutions can be offloaded onto peers because peers will tend to adopt these solutions themselves. Cognitive operations like search, memorization, and choice selection are made less taxing because of the presence of other simultaneous searchers.

4.4. Limitations and implications for future research

Our results regarding the effects of social learning are likely to depend on formal details of the task, parameters of the problem space, and the information environment. The instructions in use likely affect the balance of participant strategies; our emphasis on maximizing average payoffs likely encouraged retention and imitation, whereas maximizing final payoffs or looking for a single best solution may have substantially increased exploration. Similarly, variations in the size and structure (i.e., "ruggedness" or "sparseness") of the problem space could be used to test effects such as the exploration "satisficing" noted earlier (a tradeoff between exploitation and exploration). As mentioned earlier, our deliberate avoidance of task characteristics amenable to "insight" removes such phenomena from the scope of our conclusions; however, these phenomena are often considered to require conceptual or perceptual transformations of the problem space (e.g., Knoblich, Ohlsson, Haider, & Rhenius, 1999), which may not be well served by the comparative and incremental strategies pursued by our participants.

Other limitations relate to our implementations of three categories of phenomena: evaluation, exploration, and information sharing. Our participants received immediate, accurate, universal feedback about their solutions; the use of delayed or noisy information (as in Mason et al., 2008), changes in the environment (as in Kameda & Nakanishi, 2002, 2003), and subjective utility complicate many real-world problems and thus need to be explored in further research on social learning. Exploration in this study was limited to choices from a finite set of options that were equally visible and initially unfamiliar to all participants. The effects of asymmetric information, background knowledge, and the ability to introduce new elements would undoubtedly have altered our results, and exploring the effects of such changes would undoubtedly be useful. Finally, information sharing between participants in our study was passive, accurate, and automatic; allowing obfuscation, deception, and selective sharing according to competitive or cooperative motivations would open up a variety of complexities, as well as connections to related fields such as game theory, economics, finance, political science, management theory, and inter- and intra-group relations.

Beyond these limitations, the differences in dynamics and strategies we found associated with the removal of peer information are (to our knowledge) novel and bear further study. These phenomena could be readily extended to other domains, such as explicit spatial search or gambling tasks. In addition, different types of information (e.g., past payoffs, solution frequency, solution similarity, peer prestige, or other characteristics)

could be systematically "knocked out" to observe individuals' fallback strategies, and the consequences for performance. Such findings have practical implications, for example, in business strategy, where the information disclosed by public firms can lead competitors to behave differently than they would facing private firms, about which less information is disclosed. Relatedly, potential new financial regulations regarding transparency that are intended to increase economic stability may also have unforeseen consequences in the ability of traders to exploit information about others in their trading strategies.

4.5. Conclusions

The most important results we found in this study were the relatively simple strategies of frequency- and similarity-related comparison and imitation, and the consequent benefits to both imitators and those who were imitated, in the form of subsequent improvements available for imitation and further improvement. This synergistic benefit of mutual imitation has been implicated in other animal species as well (Sumpter, 2010). For example, one of the reasons why a cliff swallow engages in costly signaling of insect food sources is that it benefits by having other cliff swallows foraging nearby (Brown & Brown, 1996). The swallows recruited by the signal track the subsequent movements of the insects more effectively than the original swallow could if foraging by itself.

This process of cumulative, incremental, mutual improvement also has intuitive benefits for groups in which it occurs. Tomasello (1994) contends that the capabilities of humans for selective and cumulative social learning constitute a "ratchet effect" that allows culture to develop stably across generations, and that this effect may be unique to humans. Though greatly simplified compared to real social learning environments, our task, in which solutions have multiple components with epistatic relationships, allowed us to examine how such solutions are built cumulatively using simple comparisons and selectively varying proportions of different information sources. This adds realistic complexity beyond that provided by models and experimental settings with simpler problem structures or less flexible learning strategies, and provides insight into how such "ratchet effects" can occur. This study illuminates some details about the dynamic individual and group-level mechanisms by which these benefits can accrue (or not, in some circumstances), and thus contributes substantially to the study of decision making and problem solving in the presence of social interaction. Given imperfect individual memory, the "cultural knowledge pool" (Kameda & Nakanishi, 2003) requires not only the introduction of new information by exploration but also its preservation and amplification through social learning.

Acknowledgments

References

Abrahamson, E., & Rosenkopf, L. (1997). Social network effects on the extent of innovation diffusion: A computer simulation. *Organization Science*, *8*(3), 289–309.

Akins, C. K., & Zentall, T. R. (1998). Imitation in Japanese quail: The role of reinforcement of the demonstrator's response. *Psychonomic Bulletin & Review*, *5*, 694–697.

Anderson, L., & Holt, C. (1997). Information cascades in the laboratory. *American Economic Review*, *87*(5), 847–862.

Bandura, A. (1965). Vicarious processes: A case of no-trial learning. In L. Berkowitz (Ed.), *Advances in experimental social psychology* vol. 2 (pp. 1–55). New York: Academic Press.

Banerjee, A. (1992). A simple model of herd behavior. *Quarterly Journal of Economics*, *107*(3), 797–817.

Baron, R. S., Vandello, J. A., & Brunsman, B. (1996). The forgotten variable in conformity research: The impact of task performance on social influence. *Journal of Personality & Social Psychology*, *71*, 915–927.

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

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

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

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

Brown, C. R., & Brown, M. B. (1996). *Coloniality in the cliff swallows*. Chicago: The University of Chicago Press.

Donchin, E., Giraldeau, L., Valone, T. J., & Wagner, R. H. (2004). Public information: From nosy neighbors to cultural evolution. *Science*, *305*, 487–491.

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

Eguíluz, V. M., & Zimmerman, M. G. (2000). Transmission of information and herd behavior: An application to financial markets. *Physical Review Letters*, *85*, 5659–5662.

Galantucci, B. (2009). Experimental semiotics: A new approach for studying communication as a form of joint action. *Topics in Cognitive Science*, *1*(2), 393–410.

Galef, B. G., & Laland, K. N. (2005). Social learning in animals: Empirical studies and theoretical models. *BioScience*, *55*(6), 489–499.

Gibbons, D. E. (2004). Network structure and innovation ambiguity effects on diffusion in dynamic organizational fields. *Academy of Management Journal*, *47*, 938–951.

Giraldeau, L., Valone, T. J., & Templeton, J. J. (2002). Potential disadvantages of using socially acquired information. *Philosophical Transactions of the Royal Society of London B*, *357*(1427), 1559–1566.

Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Reading, MA: Addison-Wesley.

Goldstone, R. L., & Ashpole, B. C. (2004). Human foraging behavior in a virtual environment. *Psychonomic Bulletin & Review*, *11*, 508–514.

Goldstone, R. L., Ashpole, B. C., & Roberts, M. E. (2005). Knowledge of resources and competitors in human foraging. *Psychonomic Bulletin & Review*, *12*, 81–87.

Goldstone, R. L., & Gureckis, T. M. (2009). Collective behavior. *Topics in Cognitive Science*, *1*(3), 412–438.

Goldstone, R. L., & Janssen, M. A. (2005). Computational models of collective behavior. *Trends in Cognitive Science*, *9*, 424–430.

Gureckis, T. M., & Goldstone, R. L. (2009). How you named your child: Understanding the relationship between individual decision making and collective outcomes. *Topics in Cognitive Science*, *1*(4), 651–674.

Healey, P. G. T., Swoboda, N., Umata, I., & King, J. (2007). Graphical language games: Interactional constraints on representational form. *Cognitive Science*, *31*(2), 285–309.

Hess, C., & Ostrom, E. (Eds.). (2007). *Understanding knowledge as a commons*. Cambridge, MA: MIT Press.

Hurley, S., & Chater, N. (Eds.) (2005). *Perspectives on imitation: From neuroscience to social science*. Cambridge, MA: MIT Press.

Hutchins, E. (1995). *Cognition in the wild*. Cambridge, MA: MIT Press.

Kameda, T., & Nakanishi, D. (2002). Cost–benefit analysis of social/cultural learning in a nonstationary uncertain environment: An evolutionary simulation and an experiment with human subjects. *Evolution and Human Behavior*, *23*(5), 373–393.

Kameda, T., & Nakanishi, D. (2003). Does social/cultural learning increase human adaptability? Rogers' question revisited. *Evolution and Human Behavior*, *24*(4), 242–260.

Kameda, T., & Tindale, R. S. (2006). Groups as adaptive devices: Human docility and group aggregation mechanisms in evolutionary context. In M. Schaller, J. Simpson, & D. T. Kenrick (Eds.), *Evolution and social psychology* (pp. 317–342). New York: Psychology Press.

Kameda, T., Tsukasaki, T., Hastie, R., & Berg, N. (2011). Democracy under uncertainty: The wisdom of crowds and the free-rider problem in group decision making. *Psychological Review*, *118*(1), 76–96.

Knoblich, G., Ohlsson, S., Haider, H., & Rhenius, D. (1999). Constraint relaxation and chunk decomposition in insight problem solving. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *25*(6), 1534–1555.

Laland, K. N. (2004). Social learning strategies. *Learning & Behavior*, *32*(1), 4–14.

Latané, B., Williams, K., & Harkins, S. (1979). Many hands make light the work: The causes and consequences of social loafing. *Journal of Personality and Social Psychology*, *37*(6), 822–832.

Laughlin, P. R., & Ellis, A. L. (1986). Demonstrability and social combination processes on mathematical intellective tasks. *Journal of Experimental Social Psychology*, *22*, 177–189.

Lazer, D., & Friedman, A. (2007). The network structure of exploration and exploitation. *Administrative Science Quarterly*, *52*, 667–694.

Louwerse, M. M., Dale, R., Bard, E. G., & Jeuniaux, P. (2012). Behavior matching in multimodal communication is synchronized. *Cognitive Science*, *36*(8), 1404–1426.

Lupyan, G., & Dale, R. (2010). Language structure is partly determined by social structure. *PLoS ONE*, *5*, e8559. doi:10.1371/journal.pone.0008559.

Mason, W. A., Jones, A., & Goldstone, R. L. (2008). Propagation of innovations in networked groups. *Journal of Experimental Psychology: General*, *137*(3), 422–433.

Mason, W. A., & Watts, D. J. (2012). Collaborative learning in networks. *Proceedings of the National Academy of Sciences*, *109*, 764–769.

McElreath, R., Bell, A. V., Efferson, C., Lubell, M., Richerson, P. J., & Waring, T. (2008). Beyond existence and aiming outside the laboratory: Estimating frequency-dependent and pay-off-biased social learning strategies. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, *363* (1509), 3515–3528.

McElreath, R., Lubell, M., Richerson, P., Waring, T., Baum, W., Edsten, E., Efferson, C., & Paciotti, B. (2005). Applying evolutionary models to the laboratory study of social learning. *Evolution and Human Behavior*, *26*(6), 483–508.

Mesoudi, A. (2008). An experimental simulation of the 'copy-successful-individuals' cultural learning strategy: Adaptive landscapes, producer-scrounger dynamics and informational access costs. *Evolution and Human Behavior*, *29*, 350–363.

Rendell, L., Boyd, R., Cowden, D., Enquist, M., Eriksson, K., Feldman, M. W., Fogarty, L., Ghirlanda, S., Lillicrap, T., & Laland, K. N. (2010). Why copy others? Insights from the social learning strategies tournament. *Science*, *328*(5975), 208–213.

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

Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). New York: Free Press.

Rubinstein, A. (2003). 'Economics and psychology'? The case of hyperbolic discounting. International Economic Review, *44* (4), 1207–1216.

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63* (2), 129–138.

Smith, J. W., Benkman, C. W., & Coffey, K. (1999). The use and misuse of public information. *Behavioral Ecology*, *10*, 54–62.

Sumpter, D. J. T. (2010). *Collective animal behavior*. Princeton, NJ: Princeton University Press.

Templeton, J. J., & Giraldeau, L.-A. (1996). Vicarious sampling: the use of personal and public information by starlings foraging in a simple patchy environment. *Behavorial Ecology and Sociobiology*, *38*, 105–114.

Theiner, G, Allen, C., & Goldstone, R. L. (2010). Recognizing group cognition. *Cognitive Systems Research*, *11*, 378–395.

Tomasello, M. (1994). The question of chimpanzee culture. In R. Wrangham, W. McGrew, F. de Waal & P. Heltne (Eds.), *Chimpanzee cultures* (pp. 301–317). Cambridge, MA: Harvard University Press.

Wisdom, T. N., & Goldstone, R. L. (2011). Imitation, innovation, and problem-solving in a networked group. *Nonlinear Dynamics, Psychology, and Life Sciences*, *15*, 229–252.

Zentall, T. R. (2003). Imitation in animals: How do they do it? *Current Directions in Psychological Science*, *12*, 91–95.