

Social Learning and Cumulative Mutual Improvement in a Networked Group

Thomas N. Wisdom (tnwisdom@indiana.edu)

Department of Psychological and Brain Sciences, Indiana University
1101 East 10th St., Bloomington, Indiana 47405 USA

Robert L. Goldstone (rgoldsto@indiana.edu)

Department of Psychological and Brain Sciences, Indiana University
1101 East 10th St., Bloomington, Indiana 47405 USA

Abstract

We used a simple problem-solving game task to study imitation and innovation in groups of participants. Guesses were composed of multiple elements with linear and interactive effects on score, and score feedback was provided after each of a number of rounds. Participants were allowed to view and imitate the guesses of others during each round, and the score information accompanying others' guesses was either shown or hidden in two conditions. When scores were not visible, social learning was impeded; participants were less efficient in their searching of the problem space and achieved lower performance overall. When scores were visible, higher performance was observed, and results indicated a more equitable sharing of productive exploration among participants within groups as a result of selective imitation and cross-participant cumulative mutual innovations.

Keywords: Social learning; distributed cognition; innovation; imitation; problem solving; innovation diffusion.

Background

The act of learning about the world from others permeates human life. This is evident upon casual reflection about how people gather information and make choices about restaurants or movies, candidates for a job or political office, a new city to live in or a large household purchase, not to mention direct collaboration. Such "social learning" has been defined broadly as "the acquisition of behavior by observation or teaching from other conspecifics" (Boyd & Richerson, 2005). Social learning is a well-studied phenomenon in non-human animals, including foraging choices in starlings (Templeton & Giraldeau, 1996), food preferences in various rodent species (Galef & Giraldeau, 2001), and mate choices in black grouse (Höglund, Alatalo, Gibson & Lundberg, 1995). Humans' rare talent among animals for direct and flexible imitation has been called "non-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 trial-and-error learning.

Social Learning Strategies

Tendencies toward individual and social learning depend on the availability and reliability of information in the environment, including other learners. Laland (2004)

reviews strategies for *when* social learning is chosen, and *who* social learners choose to imitate. The first class of strategies (when to imitate) often uses the relative cost or uncertainty of asocial learning as criteria. For example, learning about predators on one's own can be very dangerous, so many animals have adapted to learn predator responses from others; in at least one instance this learning has occurred across species (Krause, 1993). The second kind of strategy (who to imitate) often relies on absolute or relative performance of candidate solutions (such as *copy the best* or *copy if better* strategies, respectively), or their relative popularity (such as the *copy the majority* strategy); each of these strategies has been shown in several species (Laland, 2004).

Consequences of Social Learning

Rogers (1988) performed simulations showing that in a temporally unstable environment, the extent to which imitation is beneficial depends on how recently the target of imitation has directly sampled the environment. Therefore, the addition of random social learners (information scroungers) to a population of asocial learners (information producers) does not improve the overall fitness of the population, because the costs of learning avoided by imitators will be offset by costs resulting from the increased use of outdated and inaccurate information. Boyd and Richerson (1995) and Kameda and Nakanishi (2003) confirmed and extended these results to show that when social learners can imitate selectively (e.g. imitating 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.

Of course, these models are greatly simplified in several ways, among them the assumption that social learners cannot discriminate between model solutions of varying quality without adopting them first. Even without this capability, the benefits for social learners (and thus average benefits for their group) in temporally stable environments are often assumed to be evident (Kameda & Nakanishi, 2002), but the mechanisms by which these benefits accrue are not necessarily clear. If social learning is essentially scrounging that only benefits imitators, then creating obstacles to social learning will only decrease the average performance of imitators. However, the results of previous experiments (Wisdom, Song & Goldstone, submitted) give

us reason to believe that imitators are often also explorers, and that social learning serves as a vital component of the creation of cumulative improvements. Thus impeding social learning is expected to lead to decreases in the performance of all participants.

Experiment Overview

The following experiment investigates both the causes and consequences of social learning. We employ a task in which participants in groups consisting of between one and nine persons are instructed to individually build solutions, which consist of multiple elements chosen from a larger set of elements over a series of rounds. These solutions are evaluated according to a score function that takes into account both individual element values and interactions between them. Groups of participants play simultaneously, and each can view the tentative solutions of all others. In one condition, participants may view fellow participants' scores alongside their solutions, and in another condition fellow participants' scores are invisible.

Predictions

We made the following predictions. When evaluative information about peer solutions was unavailable, participants would be unable to be sufficiently selective in imitation, and thus participants employing highly imitative strategies would have relatively lower scores. Imitation strategies in both conditions would be biased toward peers with solutions similar to the imitator's, and toward adopting solution elements that were more popular among peers, but these effects would be more pronounced in the invisible-scores condition in order to compensate for the lack of direct evaluative information. Mean scores would be lower for participants (including successful asocial learners) in the invisible-scores condition because they would be unable to easily take advantage of good solutions found by others through selective imitation and further improve upon them.

Methods

Participants were recruited from the Indiana University Psychological and Brain Sciences 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 9 persons. 234 individuals participated in the experiment, distributed across 65 sessions as shown in Table 1.

Table 1. Distribution of participants across group sizes.

Group size	1	2	3	4	5	6	7	8	9
# Sessions	16	8	11	11	5	2	3	3	2
# Participants	16	16	33	44	35	12	28	32	18

Task Details and Instructions

We implemented the experiment using custom software written in Java and Flash and run in a web browser (a version of the task can be run as "Creature League" at <http://groups.psych.indiana.edu/>). Each participant used a mouse to interact with the experimental game. A central game server recorded data and updated participant displays at the end of each round. In the game itself, participants attempted to maximize the scores earned by their chosen subsets ("teams") from a set ("league") of creature icons over 24 rounds. The display included an area for the participant's own current team, another area that could be toggled to show the participant's previous round team or their best-scoring team so far in the game (along with its associated score), a league area which showed all of the icons that were available for selection, and indications of the current round in the game and the amount of time remaining in the current round. If a session included more than one participant, each participant's display also showed the team and, in the visible-scores condition, the associated score for each other participant in the previous round. The ordering of other participants' teams in the display was not kept constant across conditions, to avoid imitation based on past behavior. Icons could be copied from any part of the display to a participant's current team by dragging and dropping them with the mouse, except for those already on the participant's current team. The current team could be replaced entirely by another team by using the score box above it as a "handle" to drag it to the current team area. A screen capture of the task interface is shown in Figure 1.




Figure 1: Example of experiment task display.

At the beginning of each session, players were given a hands-on demo of the game (including the various ways to move creatures to one's current team), and further informed about the mechanics of the game and what to expect in the remainder of the experiment session, including the following information. Each game consisted of 24 rounds, and each round was 10 seconds long. Score feedback 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. At the end of each game, the display showed the player's final score, along with a table of the scores of each player in each round of the game, sorted by average score. The player's own score was highlighted to show their relative performance without placing competitive emphasis on it. Players were instructed to do their best to maximize their teams' scores over all 24 rounds. At the beginning of each game, each player's team was a random selection of creature icons from the league. Each group played 6 games; in 3 of the games, other participants' scores were visible, and in the other 3 they were not. These were called the visible-scores and invisible-scores conditions, respectively, and were played in random order in each session.

In each game, each icon was associated with a certain positive number of points, and several special pairs of icons were associated with separate score bonuses or penalties that captured interactions between icons. 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 which gave large positive bonuses were distributed among icons with small individual point values, and pairs which gave large negative penalties were generally found among icons with large individual point values. There was a greater number of positive interactions than negative ones, to give the score distribution a larger upper tail. For ease of comparison and analysis, all scores were normalized to the range [0,1] according to the minimum and maximum possible scores. The combinations of individual and pair values described above resulted in the probability distribution of scores among all possible teams shown in Fig. 2. Participants were not given explicit information about the maximum score, the score distribution, or the position of the interaction terms. The icons' display position and associations with the point distribution were shuffled randomly for each game, so that their appearance and placement in the display did not give clues as to their point values during the course of an experiment session.

Dependent Variables and Definitions

In each round, the following data were recorded for each player: the icons (*choices*) on the team at the end of the round, the *source* of each icon, and the resulting score. The *source* information indicated whether an icon was unchanged from the previous round (Retained), copied from the player's own best-scoring team so far (Retrieved), chosen from the league display (Innovated), or copied from another player's team (Imitated). When Imitation was chosen, the persistent identifier of the copied player was recorded to allow further analyses of imitation decisions. In the case of a player replacing the entire team with Imitated icons, only the choices that were not already present on the

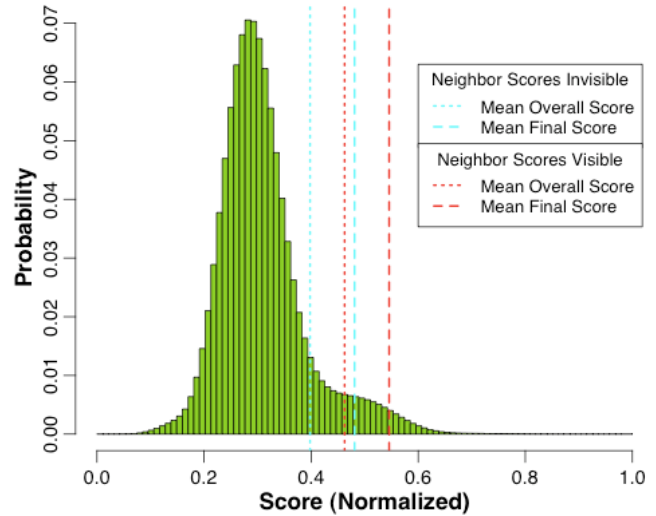


Figure 2: Distribution of scores for all possible teams.

team were counted as Imitated. Similar criteria applied to replacement of an entire team with Retrieved icons, or removing an icon and then returning it to the team via an Innovation choice. *Choice similarity* was defined as the proportion of icons that two teams have in common. An *improvement* was defined as an instance of a participant obtaining a score higher than all prior scores of all players within a particular game. Each participant's *normalized improvement share* was defined as their individually achieved proportion of the total improvements achieved by all participants in a condition, multiplied by the number of participants. A value of 1 indicated a "fair" share, e.g. a participant achieved one-third of the improvements in a three-person session. *Guess diversity* for a group in a particular round was defined as the proportion of icons in the league represented on one or more participants' teams in that round. This value was normalized by the average expected value of this proportion for each participant group size, generated by a Monte Carlo simulation.

Results

Differences in Performance

Grouped participants achieved mean overall (across all rounds) and final normalized scores of .398 and .481 respectively in the invisible-scores condition, and significantly higher scores (.463 and .546) in the visible-scores condition (see Figure 2). Isolated participants achieved mean overall and final scores of .356 and .395. Linear mixed-effects models revealed that score increased significantly with group size in the visible-score condition ($F(1,63)=79.75$, $p<.0001$, $B=0.354$), as well as in the invisible-scores condition ($F(1,63)=14.94$, $p=.0003$, $B=0.129$), though the latter trend was not as strong. Of all grouped participants, 81.7% had higher mean scores in the visible-scores condition than in the invisible-scores condition (see Figure 3).

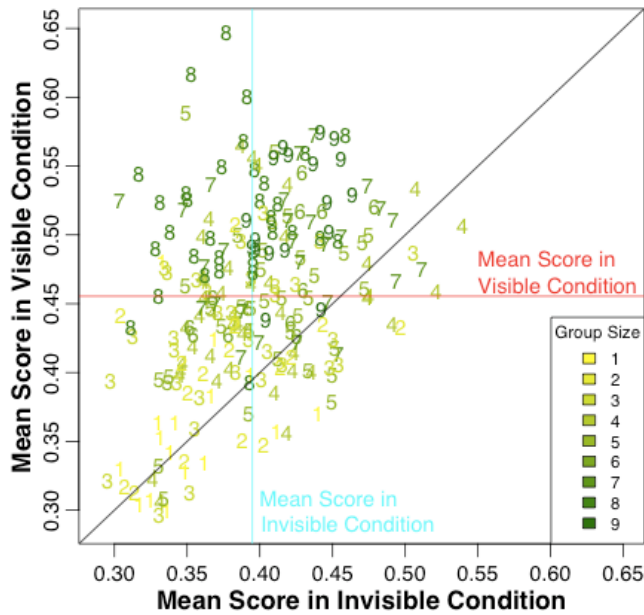


Figure 3: Scattergram of individuals' mean scores in each condition, labeled with their participant group size.

Linear mixed-effects models were used to examine trends across rounds for score and guess diversity, with a random effect of participant group. Analysis of score versus round showed a strong positive trend in the visible-scores condition ($F(1,1494)=295.96$, $p<.0001$, $B=.534$, mean increase=0.188), and a slightly shallower positive trend in the invisible-scores condition ($F(1,1494)=251.93$, $p<.0001$, $B=.615$, mean increase=0.145; see Figure 4). Guess diversity showed a similarly strong decrease across rounds in the visible-scores condition ($F(1,1126)=304.78$, $p<.0001$, $B=-.443$, mean change=-0.468), and a weaker decrease in the invisible-scores condition ($F(1,1126)=97.31$, $p<.0001$, $B=-0.453$, mean change=-0.271; see Figure 4).

Grouped participants achieved an average of 1.21 improvements per person in the visible-scores condition, and 0.95 in the invisible-scores condition. Isolated participants achieved an equivalent average of 2.44 improvements per person. Mean proportions of each choice source for improvement and non-improvement guesses in each condition are shown in Table 2. In both conditions, the proportion of Innovation choices was higher for guesses that yielded improvements relative to non-improvements (invisible-scores: $t(733.20)=-14.03$, $p<.0001$; visible-scores: $t(907.73)=-17.14$, $p<.0001$). In the invisible-scores condition, the proportion of Imitation choices was significantly lower for improvements than non-improvements ($t(916.77)=11.54$, $p<.0001$), while in the visible-scores condition, the proportion of Retention choices was significantly lower for improvements than non-improvements ($t(916.33)=9.34$, $p<.0001$). Overall there was significantly higher Retention in the visible-scores condition ($t(360)=-2.218$, $p=.027$, indicating that guesses changed more slowly).

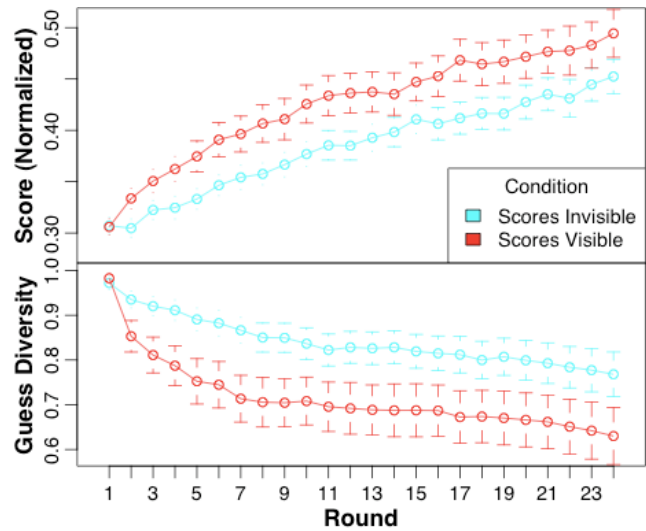


Figure 4: Change in score and guess diversity across rounds in each condition.

Analyses of relationships between mean individual score and mean individual choice source proportions showed a strong negative correlation in both conditions between score and prevalence of Innovation choices (invisible-scores: $F(1,196)=64.16$, $p<.0001$, $B=-0.497$; visible-scores: $F(1,196)=153.5$, $p<.0001$, $B=-0.663$) and a strong positive relationship between score and Retention (invisible-scores: $F(1,196)=15.27$, $p=.0001$, $B=0.269$; visible-scores: $F(1,196)=62.87$, $p<.0001$, $B=0.493$), while a strong positive relationship was shown for Imitation only in the visible-scores condition ($F(1,196)=9.70$, $p=.002$, $B=0.217$), and a strong positive relationship was shown for Retrieval only in the invisible-scores condition ($F(1,196)=14.28$, $p=.0002$, $B=0.261$).

Of all improvements in the invisible-scores condition, 14.5% resulted from guesses that included Imitation, versus 28.4% in the visible-scores condition. In a large majority (>70%) of those cases across both conditions, the focal player imitated at least one peer who had previously imitated the focal player. In other words, a player who was imitated by another player often later imitated the same player in the course of creating an improvement.

Table 2: Mean choice source proportions for (non-) improvement guesses in each condition. (Significant differences within a condition are in **boldface**, and significant differences between conditions are in *italics*.)

Condition	Improvement?	Imit.	Innov.	Retain	Retr.
Invisible Scores	No	.100	.133	<i>.712</i>	<i>.044</i>
	Yes	.039	.216	<i>.705</i>	<i>.035</i>
Visible Scores	No	<i>.091</i>	.114	.763	<i>.022</i>
	Yes	<i>.082</i>	.194	.695	<i>.021</i>

Normalized improvement share showed a relatively equitable distribution of improvements within groups in the visible-scores condition, with the distribution peaked near a "fair" share of 1. In the invisible-scores condition, however, the distribution had a strongly inequitable skew, with a modal share of zero (see Figure 5). A Kolmogorov-Smirnov test of equality of distributions indicated that these distributions were significantly different ($D=0.171$, $p=0.006$). Mean overall score showed a strong positive correlation with improvement share in the invisible-scores condition ($F(1,148)=34.94$, $p<.0001$, $B=0.329$), but this relationship was not evident in the visible-scores condition.

Differences in Strategy

In the visible-scores condition, approximately 79% of imitation events were of the highest-scoring player, while in the invisible-scores condition, all players were approximately equally likely to be imitated with regard to score. A comparison between the mean choice similarity of participants' most recent guesses to those whom they imitated, and to those whom they did not imitate, revealed a slight but significant positive difference in the visible-scores condition: a similarity value of .563 for imitated and .524 for non-imitated guesses ($t(5084.88)=-5.47$, $p<.0001$). The opposite was true in the invisible-scores condition: .317 for imitated and .346 for non-imitated guesses ($t(4041.53)=4.02$, $p<.0001$). In other words, when scores were visible, imitation was biased toward similar guesses, and when scores were invisible, imitation was biased toward dissimilar guesses.

In order to measure the bias of participants to choose an icon according to its frequency in peers' teams, 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. 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 ($N - N_{R-1}$). If a participant had decided to imitate an icon at random from among all neighbors' teams, a certain chance correlation with choice frequency would be expected simply because more high-frequency icons are present. However, a linear mixed-effects analysis of imitation probability versus choice frequency showed a positive frequency bias that was significantly greater than chance in the visible-scores condition ($F(1,604)=943.25$, $p<.0001$, $B=.741$) and significantly below chance in the invisible-scores condition ($F(1,604)=231.67$, $p<.0001$, $B=.470$). This indicates that in the visible-scores condition, participants were biased toward imitating higher-frequency icons at a rate greater than expected by chance, but not in the invisible-scores condition.

Discussion

When scores were visible, participants were heavily biased toward imitating higher-performing peers (displaying the *copy the best* strategy discussed in Laland (2004)), and performance was correlated with the average amount of Imitation in a participant's choices. Participants also showed a bias toward imitating solution elements that were possessed by larger proportions of their fellow participants, similar to the *copy the majority* strategy. Another bias evident in the score-visible condition was toward imitating more similar guesses, which allowed the imitator to make use of social learning while keeping a solution partially compatible with previous solutions and existing knowledge of the problem space, a phenomenon explored in studies of innovation propagation (Rogers, 2003).

As expected, hiding other participants' score information strongly impeded social learning: when others' scores were not visible, the choice of whom to imitate was approximately random with respect to score, and performance was correlated with the average amount of Retrieved information on a participant's team, showing the incentive to focus on previously-acquired information rather than that of others. Surprisingly, participants in the score-invisible condition also seemed to be slightly biased against peer solutions that were similar to their own, as well as icons which were more popular among their peers, perhaps indicating a bias toward novelty, which would help explain the overall decrease in individual Retention in this condition.

However, participants in the invisible-scores condition still showed a slight bias toward imitating more popular icons, indicating that the lack of score information did not cause them to disregard the guesses of their fellow players entirely. Though it conflicts with the finding that imitation in this condition occurred without regard to score, this may explain some of the improvements using Imitation and the




Figure 5: Histograms showing relatively equitable achievement of improvements within groups in the visible-scores condition, and an inequitable distribution in the invisible-scores condition.

positive relationship of score with participant group size in this condition. When players have relatively high incentives to explore for themselves rather than imitate, and yet have some solution elements in common, it is reasonable to conclude that those common solution elements may produce good scores. This is also consistent with many participants' self-reported strategies.

As seen in the increasing score and decreasing guess diversity trends across rounds, average performance increased via the convergence of group members on regions of the problem space that contained high-quality teams. This convergence combined with a small amount of individual exploration caused such regions to be explored more thoroughly and still better solutions to be found. However, in the invisible-scores condition, when imitation was not focused on a small group of better-performing neighbors (because performance information was not available), or similar guesses, this convergence happened much more slowly, search was more diffuse and less efficient, and lower performance resulted.

The significant correlation of improvement share with mean scores in the score-invisible conditions shows that individuals who were relatively more successful at individual exploration were rewarded with proportionately better overall scores compared to others, because their fellow players could not easily copy their improvements and achieve their scores. In the score-visible conditions this relationship disappeared, but 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 approximately the same payoffs, but absolute payoffs were higher for both compared to the invisible-scores condition. This is because imitators were not merely scroungers; the substantial proportion of Imitation present in improvements shows that imitated guesses were often the basis for further cumulative innovations. The cumulative innovation hypothesis is supported by the fact that a large proportion of improvements which used Imitation involved mutual Imitation and improvement, in which solution elements were passed between players via copying and built into better solutions in the process. This enabled a more equitable sharing of the "labor" of producing improvements, and produced more improvements overall.

Gabriel Tarde, one of the founders of social psychology, considered innovation and imitation to be "the fundamental social acts" (Tarde 1903/1969). Cultural conventions can be thought of as a form of large-scale imitation of behaviors that evolve along with their associated populations, subject to accompanying adaptive pressures (Boyd & Richerson, 2005). Innovations are necessary to adapt to the challenges of changing environments, and when members of a group imitate them, adaptive solutions to problems can be effectively preserved within a culture.

The findings of this study point to new avenues for understanding how innovations are generated and spread, as well as how information, incentives and the dynamic

behavioral interactions of individuals create higher-level consequences for the groups to which they belong.

Acknowledgements

The authors would like to thank Xianfeng Song, Zoran Rilak, and Todd Gureckis for their help in designing and programming the experiment, and Frances Kidwell and Bennis Pavisian for their help with running the experiment sessions. This work is funded by National Science Foundation REESE grant 0910218 and a National Science Foundation IGERT traineeship.

References

- Bandura, A. (1965). Vicarious processes: a case of no-trial learning. In L. Berkowitz (Ed.) *Advances in Experimental Social Psychology, Vol. II*. New York: Academic 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.
- Galef, B. G. Jr., & Giraldeau, L. A. (2001). Social influences on foraging in vertebrates: causal mechanisms and adaptive functions. *Animal Behaviour, 61*(1), 3-15.
- Höglund, J., Alatalo, R. V., Gibson, R. M., & Lundberg, A. (1995). Mate-choice copying in black grouse. *Animal Behaviour, 49*(6), 1627-1633.
- Kameda, T., & Nakanishi, D. (2002). Cost-benefit analysis of social/cultural learning in a non-stationary uncertain environment: An evolutionary simulation and an experiment with human subjects. *Evolution and Human Behavior, 23*, 373-393.
- Kameda, T., & Nakanishi, D. (2003). Does social/cultural learning increase human adaptability? Rogers's question revisited. *Evolution and Human Behavior, 24*, 242-260.
- Krause, J. (1993). Transmission of fright reaction between different species of fish. *Animal Behavior, 65*, 595-603.
- Laland, K. N. (2004). Social learning strategies. *Learning & Behavior, 32*(1), 4-14.
- 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.
- Tarde, G. (1969). *The Laws of Imitation*. (Elsie Clews Parsons, Trans.). Chicago: University of Chicago Press. (Original work published 1903.)
- Templeton, J. J., & Giraldeau, L.-A. (1996). Vicarious sampling: the use of personal and public information by starlings foraging in a simple patchy environment. *Behavioral Ecology and Sociobiology, 38*, 105-113.
- Wisdom, T. N., Song, X., & Goldstone, R. L. (manuscript submitted for publication). The effects of peer information on problem-solving in a networked group.