

# Emergent Processes in Group Behavior

Robert L. Goldstone, Michael E. Roberts, and Todd M. Gureckis

*Indiana University*

---

**ABSTRACT**—*Just as neurons interconnect in networks that create structured thoughts beyond the ken of any individual neuron, so people spontaneously organize themselves into groups to create emergent organizations that no individual may intend, comprehend, or even perceive. Recent technological advances have provided us with unprecedented opportunities for conducting controlled laboratory experiments on human collective behavior. We describe two experimental paradigms in which we attempt to build predictive bridges between the beliefs, goals, and cognitive capacities of individuals and patterns of behavior at the group level, showing how the members of a group dynamically allocate themselves to resources and how innovations diffuse through a social network. Agent-based computational models have provided useful explanatory and predictive accounts. Together, the models and experiments point to tradeoffs between exploration and exploitation—that is, compromises between individuals using their own innovations and using innovations obtained from their peers—and the emergence of group-level organizations such as population waves, bandwagon effects, and spontaneous specialization.*

**KEYWORDS**—*collective behavior; imitation; social networks; competition; group behavior; foraging*

---

It is natural for psychologists to focus on the behavior of single individuals, because introspection provides people with motivation and perspective at this level. However, in a literal sense, we are all participating in entities greater than ourselves. Self-organized collectives of people create emergent group-level patterns that are rarely understood or intended by any individual. A business has a style and ethos that transcends its employees. A culture has a nature, integrity, and systematicity that transcends its inhabitants while still being grounded by their

interactions (Atran, Medin, & Ross, 2005). Social phenomena such as the spread of gossip, the World-Wide Web, the popularity of cultural icons, legal systems, and scientific establishments all take on a life of their own, complete with their own self-organized divisions of labor and specialization, dynamics, feedback loops, growth, and adaptations.

A considerable amount of early work on group behavior from social psychology focused on interpersonal relations and the attributes that characterize good leaders or work teams. However, the social patterns that people form are often organized without explicit leaders, chains of command, or fixed communication networks (Ball, 2004). Examples of such spontaneously emerging social patterns include book recommendations on Amazon.com (which evolve based upon similar readers' buying habits), fans at a sport stadium, grassroots political movements, the development of a fully cross-indexed and intricately organized online encyclopedia that any person can edit ([www.wikipedia.org](http://www.wikipedia.org)), and an online venue for media sharing that is freely accessible to both providers and consumers yet still shows striking trends of rich-get-richer popularity ([www.youtube.com](http://www.youtube.com)). In December of 2006, *Time* magazine named "You" as Person of the Year in recognition of the power and sophistication of these grass-roots, decentralized communities.

To understand the structure and dynamics of human collectives like these, we have developed Internet-based experimental platforms that allow groups of 20 to 200 people to interact with each other in real time on networked computers. The experiments use virtual environments in which participants can see the moment-to-moment actions of their peers and immediately respond to their environment by making responses of their own. To understand the results of these experiments, we have developed computational models. Several models of group behavior exist, but rarely are these models tested against detailed data sets obtained from controlled laboratory settings. Often there is a disconcerting mismatch between the simplicity of formal models and the complexities of real-world situations. Our strategy for bridging the gap between computational models and group-behavior phenomena is to create relatively simple laboratory situations involving groups of people interacting in idealized

---

Address correspondence to Robert L. Goldstone, Psychological and Brain Sciences, Indiana University, Bloomington, IN 47401; e-mail: [rgoldsto@indiana.edu](mailto:rgoldsto@indiana.edu).

environments according to easily stated “game rules.” We admittedly sacrifice some external validity in creating idealized experimental scenarios, but this loss is offset by the nearly exact correspondence of the assumptions underlying our psychological experiments to those of the computational models; this allows the models to be aptly applied without sacrificing their concise explanatory value and genuine predictiveness. In what follows, we focus on two scenarios that capture ubiquitous group patterns: the competition of agents for resources and the dissemination of innovations in social networks.

### THE COMPETITIVE SEARCH FOR RESOURCES

A problem faced by all mobile organisms is how to search their environment for resources. Animals forage their environment for food, Web-users surf the Internet for desired data, and industries mine the land for valuable minerals. When an organism forages in an environment that consists, in part, of other organisms that are also foraging, unique complexities arise. The resources available to each individual are affected not just by that individual’s behavior but also by the simultaneous actions of other individuals.


We have been interested in experimentally exploring how human foragers allocate themselves to resources and the time course by which that allocation is achieved. Through the Internet (available to the public at <http://groups.psych.indiana.edu/>), eight groups of 12 to 28 (average = 21) participants competed for food tokens deposited over time in a virtual environment consisting of an  $80 \times 80$  grid of squares. In Goldstone and Ashpole (2004), a virtual environment was created with two resource pools into which valued food tokens could be deposited with different rates of replenishment. Each pool consisted of a compact region of many squares. The participants’ task was to obtain as many resource tokens as possible during the course of a 270-second experiment. A participant obtained a token by being the first to move on top of it. Participants moved square-by-square by pressing arrow keys on their computer keyboard.

Resources were divided between two resource pools in various ways. For example, in one condition, food tokens were split evenly (50/50) between the two pools; other conditions had a 65/35 or 80/20 split. The location of the food within a pool followed a normal (Gaussian) distribution with a mean at the center of the pool and a standard deviation of five horizontal and vertical positions. The locations of the pools were randomized under the constraint that the distance between pools was kept approximately constant. One piece of food was delivered to one of the resource pools every  $4/N$  seconds, where  $N$  is the number of participants. In our “visible” condition, each participant could see each other and the entire food distribution. In our “invisible” condition, they could not see the other participants, and so they gradually acquired knowledge of the resource distributions by virtue of their histories of getting food from each location.

The dynamics of the distribution of agents to resources in each condition are shown in Figure 1, broken down by the three types of resource distribution. Although fast adaptation to the food distributions takes place, the asymptotic distribution of agents systematically undermatches the optimal distribution of agents to resource pools. By undermatching, we mean a distribution of agents that is less uneven than the distribution of resources. For example, in the 65/35 distribution, the 65% pool only attracts an average of 60.6% of the agents. If we were efficiency consultants, we would recommend that some foragers in the less productive pool move to the more productive pool, as the resources there are being relatively underutilized despite the larger crowd. This finding of undermatching has been obtained with other animal groups, including cichlid fish, mallard ducks, and mites (Kennedy & Gray, 1993). We have also observed this undermatching in collectives of citizens of the virtual world Second Life (<http://secondlife.com/>), who were invited to forage for pieces of the world’s currency that we randomly placed in two regions. This undermatching may also explain real-world human foraging behavior, such as the documented inefficiency in sperm whalers’ hunting for whales near the Galapagos Islands in the early 19th century (Whitehead & Hope, 1991).

Our results also reveal periodic fluctuations in resource use. A Fourier analysis was applied to the populations at the resource pools over time to reveal cyclic oscillations of migration. Fourier transformations translate a time-varying signal into a set of sinusoidal components. Each sinusoidal component is characterized by a frequency. The power of a component indicates the strength of a periodic response at that frequency. This analysis revealed significantly greater power in the low-frequency spectra for invisible conditions than for visible conditions. For all three invisible conditions, the peak power was at approximately .02 cycles/second and was particularly high for the most uneven, 80/20 distribution. This means that in the invisible conditions, agents collectively caused waves of relatively dense crowding at one pool that repeated about once every 50 seconds. There was no evidence for population cycles in the visible conditions, presumably because a person who was tempted to leave their dissatisfying pool for greener pastures would be dissuaded if they saw several other people with the same idea already leaving their pool. However, in invisible conditions, agents may become dissatisfied with a pool populated with many other agents, but as they leave such a pool they would not be aware that other agents are also leaving. Thus, the ironic consequence of people’s shared desire to avoid crowds is the emergence of migratory crowds!

In a second experiment (Goldstone, Ashpole, & Roberts, 2005), we ran groups of participants in conditions where food resources, but not fellow foragers, were visible, and vice versa. If people acted like buzzards, using the presence of peers as an indicator of possible food sources, then the presence of a relatively large number of participants at the richer pool would be expected to draw still more participants to the pool. In fact, when



**Fig. 1.** Changes in the number of people at each of two resource pools across 270-second foraging experiments (Goldstone & Ashpole, 2004). Resources were distributed evenly (50/50) or with unequal (65/35 or 80/20) distributions. Participants either were shown the positions of other participants and resources (visible) or not (invisible). The actual distributions of resources (indicated by the straight horizontal lines) are more extreme than the distributions of participants to these resources.

agents can see each other but not the overall pattern of food, their distribution is not uneven enough (e.g., only 62% of the participants were at the 65% pool), but when agents can see food resources but not each other, then their distribution is more uneven than the resource distribution (e.g., 73% of the participants were at the 65% pool). This suggests that people are more like some aphid species than like buzzards, avoiding sites that already have a large crowd of other members of their own species. By this account, overmatching occurred because participants were attracted to the rich, productive pools, and were not dissuaded from approaching these pools by the presence of other participants (because those others were invisible).

To gain greater insight into our results, we chose to model them using an Agent-Based Model. This class of models builds social structures “from the bottom up,” by populating the simulation with many individual virtual agents and allowing emergent organizations to form out of the operation of rules that govern interactions among these agents and their environment. Our EPICURE model (Roberts & Goldstone, 2006, available as an interactive simulation at <http://cognitrn.psych.indiana.edu/Epicure.html>) populates a world with agents that probabilistically decide from moment to moment what spatial grid location

they will approach based on the locations’ values. The first factor that affects a location’s value is its distance: The closer a location is, the more likely it is to be selected as a target destination. Second, once a location has been selected as a target, its value is increased so that it will tend to be selected as a target at the next moment too. This is a way of incorporating consistency in target choices over time. For agents who can see all of the other agents and food, a third factor is that the value of a location increases as the density of food in its vicinity increases, and a fourth one is that the location’s value decreases as the density of other agents increases. In the invisible condition, agents must gradually accumulate a personal history of where they have found food. Every time food is found at a location, the location’s value increases, and this increase diffuses to the nearby locations.

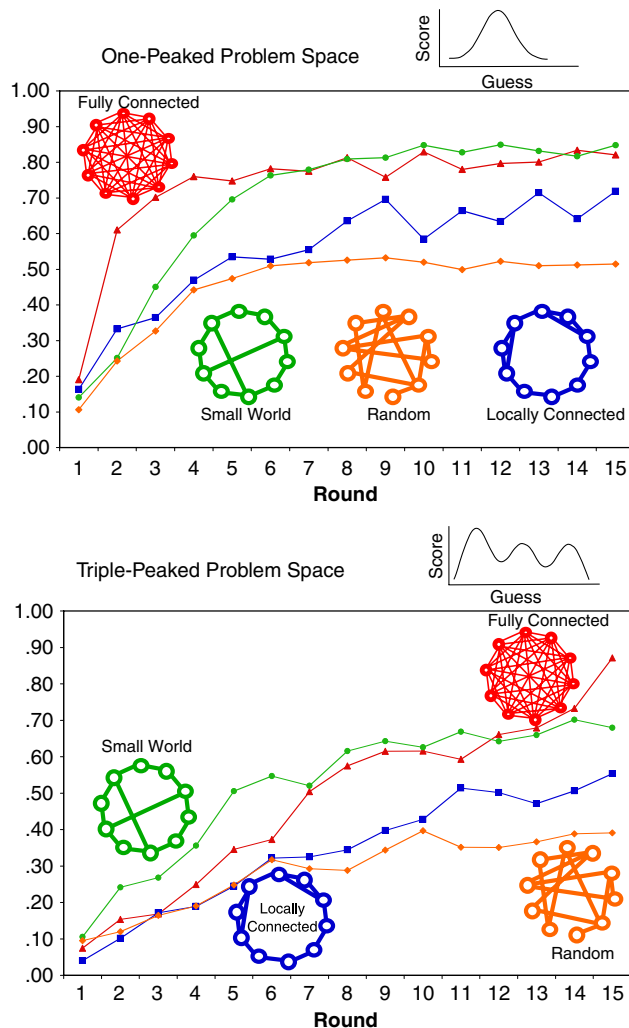
These simple assumptions allowed EPICURE to account for the empirically observed pattern of overmatching and undermatching for the four visibility conditions in which foragers and food could independently be visible or invisible. Why does EPICURE predict undermatching? The critical notion is spatial “turfs.” A single agent can efficiently patrol a compact region of about 10 squares, roughly independent of the food productivity. Two pools that differ in their productivity both have the same

spatial extent and variance, and so can support agents in numbers that are more similar than predicted by the pools' productivities alone. EPICURE also predicted our observed population waves in the invisible conditions, as well as other counterintuitive results found in animal foraging, such as that increasing the distance between two resource pools—and hence increasing travel cost—should decrease undermatching (Baum & Kraft, 1998). One of the best ways to evaluate a model is to see whether behaviors that are not explicitly forced by the rules arise—in other words, are we getting out more than we knew we were putting in the model? By this measure, the model does a good job of explaining collective foraging behavior. Although undermatching and population waves were not explicitly stipulated by the model's assumptions, these behaviors emerge because of agents' tendencies to avoid each other but to be attracted by the same resources.

### DISSEMINATING INNOVATIONS IN SOCIAL NETWORKS

The foraging paradigm involves competitive searching for spatial resources, but we have also studied collective searching for abstract resources. Any organism that is capable of imitating its peers must decide when and how much to imitate others' solutions versus discover its own solutions. To study this in a well-controlled, if somewhat artificial, setting, we had participants guess numbers between 0 and 100 using Internet-connected computers (Mason, Jones, & Goldstone, 2005). Each of the participants' computers then showed them the points earned by their guesses, based upon a hidden scoring function that had either a simple single-peaked or complex triple-peaked form, shown in Figure 2. The triple-peaked form had two local maxima—solutions that were better than their neighboring solutions but not the best possible—and one global maximum. Over 15 rounds, participants received feedback not only on their own guesses but also on their neighbors' guesses. Neighbors were determined by one of four types of network structures: locally connected (connections only to one's immediate neighbors), random, fully connected (everybody connected to everybody else), and small-world (e.g. local connections plus a few long-range connections). Figure 2 shows sample networks for groups with 10 participants.

For the easy, single-peaked function, participants in the fully connected networks converged most quickly on the global maximum, with the random and locally connected networks performing worse. This pattern of results is readily explainable in terms of the propensity of a network to disseminate innovations quickly. Innovations disseminate most quickly in the full network because every individual is informationally connected to every other individual. For the trickier, three-peaked payout function, the small-world network performs better than the fully connected network, particularly for the first half of the trials. The truism of "the more information, the better" is not supported.



**Fig. 2.** Percentage of participants within one standard deviation of the global maximum (best solution) on each round of a problem-solving task for which there were two versions (Mason, Jones, & Goldstone, 2005). Groups of participants guessed numbers between 0 and 100 using Internet-connected computers; each of the participants' computers then showed them the points earned by their own and others' guesses, based upon a hidden scoring function that had either a simple single-peaked (single-peaked problem space) or complex triple-peaked (triple-peaked problem space) form. In the fully connected network, everybody could see each other's guesses and outcomes. In the random network, participants only had access to a set of randomly determined neighbors. In the locally connected network, participants were informationally connected only to their close neighbors. The small-world network also preserved local neighborhoods but additionally had a few distant "short-cut" connections that bridged different local regions. For the one-peaked problem, the best group performance was initially found for the fully connected network. For the triple-peaked problem, the best performance was initially found for the small-world network.

Indeed, problem spaces requiring substantial exploration may benefit from networks with mostly locally connected individuals. The problem with the fully connected network is that everybody ends up knowing the same information, and they thereby become too like-minded, acting like a single explorer rather than like a federation of independent explorers. The small-world structure is an effective compromise between fully exploring a search

space and also quickly disseminating good solutions once they are found (Watts & Strogatz, 1998). When the problem space is even trickier than the three-peaked function, with a single sharp needle peak for the global maximum and a broad local maximum, then the locally connected network performs best, consistent with its preservation of local and quasi-independent communities. Computational modeling work converges on these empirical results in showing that more locally connected social networks are beneficial when the problem the group has to solve is difficult (Hutchins, 1995; Lazer & Friedman, in press). Increasing connectivity among members of real-world cockpit crews has also been shown to hamper group performance by foreclosing exploration (Hutchins, 1995; see also Hinsz, Tindale, & Vollrath, 1997 for a discussion of the danger of groups' over-reliance on shared information).

### FUTURE PROSPECTS FOR RESEARCH ON COLLECTIVE BEHAVIOR

The previously described paradigms are united in exploring group search behavior in both physical and abstract solution spaces. In related work, we have examined the kinds of trail systems that people create when they are motivated to take advantage of the trails left by their predecessors and, in so doing, further reinforce and extend those trails (Goldstone & Roberts, 2006). The resulting trails represent a compromise between going where one wants to go and going where others have gone before. We have begun to apply this work on imitation and exploration to modeling and predicting baby names. Names are interesting because they are roughly neutral in terms of intrinsic value but are culturally meaningful artifacts. "John" is not intrinsically a better name than "Warren" even though it occurs 35 times more frequently in the United States. The distribution of baby names strongly suggests that, as with trails and scholarly citations, the more often a name is used, the more often it will be used in the future.

Other promising areas for experimental research on collective behavior include coalition formation and coordination, social dilemmas, group dynamics, and social specialization. The common principles that repeatedly arise in our group-behavior paradigms include (a) a tradeoff between exploration and exploitation, (b) a compromise between individuals using self- versus other-obtained information, and (c) the emergence of group-level resource usage patterns that result from individual interests but are not always favorable to those interests. Unfavorable manifestations of these principles include inefficient population waves, bandwagon effects (in which people do things because other people do the same), mismatches between agent and resource distributions, disadvantages for highly connected networks, and premature convergence of populations on suboptimal solutions. Despite these pitfalls, collective search continues to be a powerful case of distributed cognition for the simple reason that individual search often fails to provide a

good solution in a limited time, and thus collaborating and sharing solutions with others can dramatically improve search efficiency.

---

### Recommended Reading

- Arrow, H., McGrath, J.E., & Berdahl, J.L. (2000). *Small groups as complex systems: Formation, coordination, development, and adaptation*. Newbury Park, CA: Sage. A highly accessible overview of research on group behavior and dynamics.
- Richerson, P.J., & Boyd, R. (2005). *Not by genes alone: How culture transformed human evolution*. Chicago: University of Chicago Press. An elegant, formal treatment of interactions between cultural and biological evolution.
- Goldstone, R.L., & Janssen, M.A. (2005). Computational models of collective behavior. *Trends in Cognitive Science*, 9, 424–430. Case studies of agent-based computational models of group behavior, and a survey of some of the opportunities and challenges for these models.
- Kennedy, J., Eberhart, R.C., & Shi, Y. (2001). *Swarm intelligence*. San Francisco: Morgan Kaufmann Publishers. Inspired by psychological research on social influence and conformity, the authors propose a novel computational algorithm for having swarms of agents collectively solve problems.
- 

**Acknowledgments**—The authors wish to express thanks to Katy Borner, Rich Shiffrin, William Timberlake, Peter Todd, Winter Mason, and Thomas Wisdom for helpful suggestions on this work. This research was funded by Department of Education, Institute of Education Sciences Grant R305H050116, National Science Foundation Grant 0527920, and National Institute of Humanities-National Institute of Mental Health Training Grant T32 MH019879-12 to the third author. Online playable versions of experiments related to the current work can be found at <http://groups.psych.indiana.edu/>.

### REFERENCES

- Atran, S., Medin, D.L., & Ross, N. (2005). The cultural mind: Environmental decision making and cultural modeling within and across populations. *Psychological Review*, 112, 744–776.
- Ball, P. (2004). *Critical mass*. New York: Farrar, Straus, and Giroux.
- Baum, W.M., & Kraft, J.R. (1998). Group choice: Competition, travel, and the ideal free distribution. *Journal of the Experimental Analysis of Behavior*, 69, 227–245.
- 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., & Roberts, M.E. (2006). Self-organized trail systems in groups of humans. *Complexity*, 15, 43–50.
- Hinsz, V.B., Tindale, R.S., & Vollrath, D.A. (1997). The emerging conceptualization of groups as information processors. *Psychological Bulletin*, 121, 43–64.

- Hutchins, E. (1995). *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Kennedy, M., & Gray, R.D. (1993). Can ecological theory predict the distribution of foraging animals? A critical analysis of experiments on the ideal free distribution. *Oikos*, *68*, 158–166.
- Lazer, D., & Friedman, A. (in press). Parallel problem solving: The social structure of exploration and exploitation. *Administrative Science Quarterly*.
- Mason, W.A., Jones, A., & Goldstone, R.L. (2005). Propagation of innovations in networked groups. *Proceedings of the 27th Annual Conference of the Cognitive Science Society*. Hillsdale, New Jersey: Erlbaum.
- Roberts, M.E., & Goldstone, R.L. (2006). EPICURE: Spatial and knowledge limitations in group foraging. *Adaptive Behavior*, *14*, 291–313.
- Watts, D.J., & Strogatz, S.H. (1998). Collective dynamics of “small-world” networks. *Nature*, *393*, 440–442.
- Whitehead, H., & Hope, P.L. (1991). Sperm whalers off the Galapagos Islands and in the Western North Pacific, 1830–1850: Ideal free whalers? *Ethology & Sociobiology*, *12*, 147–161.