Title

Group Stratification and Coordination Failure in a Continuous N-Player Stag Hunt

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 32(32)

ISSN
1069-7977

Authors
Frey, Seth
Goldstone, Robert L.

Publication Date
2010

Peer reviewed

Group Stratification and Coordination Failure in a Continuous N-Player Stag Hunt

Seth Frey (sethfrey@indiana.edu)
Cognitive Science Program, 819 Eigenmann, 1910 E. 10th St.
Indiana University, Bloomington, IN 47406

Robert L. Goldstone (rgoldsto@indiana.edu)
Psychology Building, Room 338, 1101 E. 10th St.
Indiana University, Bloomington, IN 47405

Abstract

We reveal spontaneous group formation and differentiation in an online dynamic coordination experiment. We observe increased group stratification and attribute it to increases in pairwise cooperative behavior, rather than uncooperative behavior. Our network analyses document the fine scale structure of coordination failure in the face of many established determinants of coordination success. We explore previous work in coordination failure to frame our own findings. Factors that have been previously shown to improve coordination in discrete-time, forced-decision experimental games do not prevent decisive coordination failure in our real-time, asynchronous group decision-making environment.

Keywords: coordination; coordination failure; n-player games; continuous-time games; stag hunt; functional networks; group structure

Introduction

Sixty years of literature have established cooperation as only a special-case outcome in experiments of economic behavior. Where there is a conflict between individual and group interest, the net benefit of all members may suffer in favor of individual gains. Even in coordination games, where individual payoffs are directly related to the group's outcome as a whole, the maximum levels of coordination and personal benefit have been notoriously difficult to attain in the laboratory (Devetag & Ortmann, 2007), either because of uncertainty inherent in the task or strategic uncertainty with regard to the actions of other participants. However, the majority of these coordination situations involve slow, round-based decisions. By contrast, humans and animals often make decisions within the time scales that characterize processing and decision time. In these environments, decision opportunities may be presented with little notice, and part of the decision process is deciding when to decide. Despite the increased cognitive demands, the ability of animal groups to coordinate successfully in these environments is well established in both experimental and observational experiments (Conradt & Roper, 2005; Petit, Gautrais, & Leca, 2009; Couzin, 2009). Recent research on group behavior has worked to understand how these faster-paced decision environments affect coordination.

We build on this work, applying dynamic network methods to estimate the changes in network structure between participants in an n-person stag hunt coordination game. Though we find support for many structural factors that have been proposed (all else being equal) to promote efficient outcomes, participants fail to coordinate on payoff-dominant equilibria, and coordination continues to decay over the course of the experiment. We take advantage of the rich data available from these fine timescale experiments and introduce functional network measures to estimate the emergent structure of participant interactions.

Literature

Experiments in real-time group decision-making document both coordination failure and success. Dyer et al. report a collective navigation experiment in which human groups walked in the induced direction, despite constrained communication, and conflicting individual information (Dyer et al., 2008). Kearns, Suri, and Montfort examined coordination in a network experiment based on a map-coloring problem to investigate the effects of group topology and information differences on cooperation (Kearns et al., 2006). Roberts and Goldstone also looked at the effect of different information treatments in a collective number-guessing game and found that participants spontaneously adopt specific roles without communication(Roberts & Goldstone, 2009). Furthermore, the extent to which group members differentiated themselves predicted how well they solved the coordination problem. Related results have been found in minority and market entry games (Bottazzi & Devetag, 2007; Duffy & Hopkins, 2005). In another network game, continuous-time decision-making is proposed to improve coordination (Berninghaus, Ehrhart, & Ott, 2006). In all of these domains, participants were generally successful at resolving conflicts and finding solutions that brought a net collective benefit.

With continued ties to the study of coarse time-scale political and macroeconomic decision making (as in Schelling, 1980), experimental economists are traditionally interested in synchronous, normal form games wherein all players choose a strategy without knowledge of any other players' choices. However, there is still rich research in games that elicit decisions at finer time scales, as discussed by Berninghaus, Ehrhart, and Keser (1999). This literature has more often found coordination failure than success. Participants in the three games of E. Friedman, Shor, Shenker, & Sopher (2004) frequently failed to converge on any equilibrium, much less those that dominate with respect to payoff or certainty. However, Bottazzi and Devetag find emergent structure that leads to higher coordinative outcomes (2007) and Cheung and

Friedman observe successful coordination in a more complex continuous-time experiment modeling financial speculative attacks (2009).

From a game-theoretic standpoint, the game in this experiment shares several features with so-called stag hunt games, which have been extensively studied in the lab. In the allegorical stag hunt, two neighbors decide whether to hunt hare separately, or to work together to hunt (the more rewarding) stag. Even though both individuals benefit by coordinating for the larger quarry, they may decide that the costs of coordination, and their uncertainty as to the actions of their neighbor, are too great. In its synchronous form, game theory does not predict which outcome is more likely without refinements like risk and payoff dominance (Harsanyi & Selten, 1988). However, in the asynchronous, extended form, in which one neighbor is forced to make a choice that the other can observe beforehand, theory predicts that the rational player will select the high payoff equilibrium (Kuhn, 2009). Similarly, continuous-time games in the laboratory have led to more effective coordinative behavior than comparable synchronous games (D. Friedman & Oprea, 2009).

While real life uncertainty and decision cost dilemmas modeled by the stag hunt are at least as common for groups as for pairs of people, generalized n-player stag hunts have only been investigated in simulation (Pacheco, Santos, Souza, & Skyrms, 2009). However, as pointed out by Resnick (2007), the important coordination experiments of Van Huyck et al. (1990) are relevant. The precedents in both of these investigations seem to predict that adding players to the stag hunt makes coordination failure more likely in our experiment.

We introduce a coordination game building on previous evidence of emergent social conventions in repeated coordination games (Bottazzi & Devetag, 2007; Rankin, Huyck, & Battalio, 2000; Roberts & Goldstone, 2009; Uzzi & Spiro, 2005). Taken together, the above theoretical and experimental results make contrasting predictions about how participants should perform in the asynchronous environment provided by the real-time nature of our group experiment.

Experiment

Paradigm

Groups of two to six participants played an n-player stag hunt game on networked computers. Each player was shown information about the payoff and location of the other uniquely identified participants (Figure 1). Participants were instructed to "harvest" points from any of twelve game tiles. Participants were awarded a tile's points after waiting for a specified amount of time. To introduce the incentives of the stag hunt, tile payoffs increased with each tile's *coordination number*. The coordination number specified the quorum of participants necessary to harvest the points on a tile. When a tile reached its quorum, a visible timer on the tile would start counting down to zero. During this time, any participant on the tile could leave it and other participants could join. If a timer reached zero, each participant on the tile received the

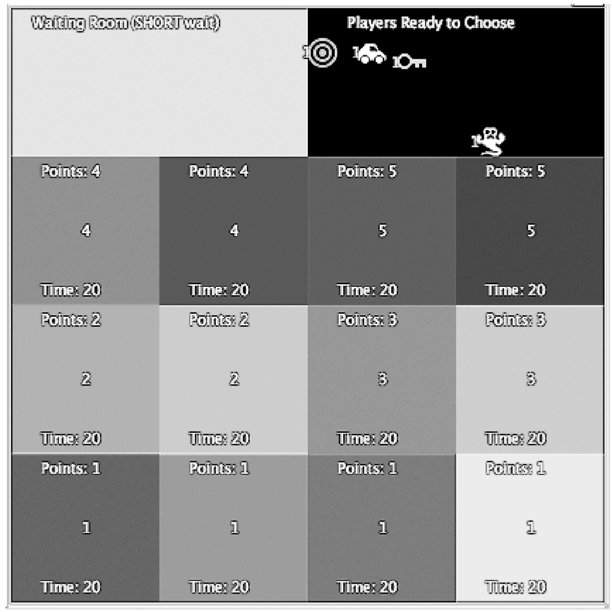

Figure 1: Screenshot of game board. Icons are controlled by participants, each of whom may select any tile. The coordination number in the center of each tile reflects how many participants must wait for the listed Time to each receive the specified number of Points. The number to the left of each player reflects his or her accumulated number of points.

tile's payoff, even if participant presence on a tile was above quorum. If the tile fell below quorum during countdown, the timer was reset and no points were distributed.

The coordination number ranged from one to five, even in cases where there were more than or fewer than five participants in a session. To allow for the possibility of independent action, tiles at each possible coordination level (of five) were distributed redundantly among all possible tiles (twelve). This allowed participants to harvest points alone on tiles with a coordination number equal to one, or to harvest tiles collectively with other participants on tiles with higher coordination numbers. We may have introduced theoretical difficulties by providing multiple tile choices at a given coordination level, but participants showed an ability to coordinate at higher levels, perhaps due to focal point effects or the cheap decision costs of the environment.

In this game, the most efficient strategy was to repeatedly choose one tile with the coordination number equal to the group size. The strategy least vulnerable to strategic uncertainty was to always choose a tile with coordination number equal to one.

Participants started in a "staging area" on the game board from which they selected a tile to harvest from. After distributing its points, a harvested tile was reset and its participants were held for two seconds before being returned to the staging area. Participants were never forced to make any choice, and the ability of any participant to make a choice did not depend on "turns" or any system of rounds that structured when participants made decisions relative to each other.

We thus distinguish between the real-time, asynchronous and non-forced aspects of decision-making in this environment. This less constrained structure allowed us to explore the effects of changes in pay level, risk, and experience on spontaneous group coordination and structure in an environment that evokes real world, short time-scale decision-making.

Manipulations

After a practice round, participants played eight five-minute rounds of the game (permitting up to 75 consecutive harvesting events). We manipulated *pay level* and *harvest time* over rounds in a $2 \times 2 \times 2$ block factorial design ([low versus high payoff]×[long versus short harvest time]×[first versus second block]). The first four rounds (block one) consisted of the 2×2 set of low and high payoffs and short and long harvest times, randomly permuted, and were followed by a second random permutation of the same four conditions. The blocks controlled for order effects and enabled investigation of the effects of experience with the group.

The two levels of pay level were *low* (tile payoff is the same as the tile coordination number) and *high* (tile payoff is the square of tile coordination number). Scaling payoffs with coordination number created an incentive for participants to risk coordinating with other participants for a higher payoff. The two levels of harvest time, the number of seconds necessary to harvest any tile, were *fast* and *slow* (two and ten seconds, respectively). This factor manipulated the amount of time that a tile had to remain at or above quorum to deliver payoffs, and thus how vulnerable participants were to strategic uncertainty—uncertainty as to the beliefs and future actions of their peers.[1]

Participants were introductory psychology undergraduates receiving course credit. Participants did not receive monetary compensation, but observational reports indicate that participants found earning points intrinsically motivating. Participants often spontaneously cheered when they received a large payoff. There were forty-three participants over twelve experiments. Our experiments ultimately had five levels of *group size*, with four groups of 2 participants, two groups of 3, two groups of 4, three groups of 5, and one group of 6. Though it was not controlled, our analysis will treat group size as an independent, between-participant variable. Each experiment lasted one hour.

Measures

The structure of the experiment allowed us to develop and compare a number of measures of coordination. The game was implemented such that the state of each player was recorded approximately twice every second. This enabled investigation of each individual's decisions as a time series, and of participants' decisions together as a multivariate time series.

For each participant, in each condition, we recorded *payoff*, *wait tile*, and *success tile*. Payoff represents how many points a participant earned in a condition. Wait tile represents the mean tile that a participant spent time on. Success tile represents the mean tile that a participant successfully harvested a payoff from. In terms of these measures, highly successful coordination would be reflected by maximum payoffs and wait and success tile values equal to group size, all identical across participants. Conversely, zero-coordination behavior would be reflected in minimal payoffs and wait and success tiles equal to one.

We also used the multivariate nature of the time data to calculate measures of functional proximity between participants in a group. These measures reflect the extent to which any pair of participants coordinated on a tile. We then assembled these dyad weights into fully connected networks representing the pattern of behavioral couplings in the group. Since these networks are fully connected, edge weights replace edge presence in representing heterogeneity within a group. With these graphs of internal group structure, we used a variety of network statistics as additional measures of coordination. These "functional networks" have been used in computational neuroscience to infer how regions of the brain are related, and to determine the extent to which dynamic functional relations correspond to physical anatomical connections (Bullmore & Sporns, 2009; Hagmann et al., 2008). Given the rich data available in the experiment, and the wealth of theoretical issues common to neuroscience and group behavior (Couzin, 2009), the extension of these tools to the study of social behavior was natural. Because functional networks have only recently been applied to the study of group behavior (Nagy, Akos, Biro, & Vicsek, 2010), we implemented three separate measures of proximity: choice distance, mutual information (Cover & Thomas, 2006) and transfer entropy (Schreiber, 2000). Because the three measures gave analogous, consistent results, this investigation will treat only the choice distance, the simplest of the measures.

The *choice distance* between a pair of participants is the number of seconds that they were on different tiles. In the resulting graph, two participants that always made the same choice are *close*, with a choice distance of zero. If their choices were always different, they are *far*, with a distance equal to the number of seconds that they could have been on the same tile. The matrix representing this graph is symmetric, and its diagonal, representing the number of times a participant is on his or her own tile, is always zero.

For every condition, we calculated the mean and variance of each individual's proximity to every other participant in their group. The mean proximity is related to the closeness centrality metric in graph theory and social network theory (Wasserman & Faust, 1994), but it preserves measure units. A change in a participant's mean proximity across some experimental condition entails that the participant coordinated more or less closely, on average, with all other participants in

[1] Strategic uncertainty is strictly uncertainty as to the beliefs and actions of other participants (Huyck et al., 1990), but this experiment provides full information as to the actions of peers in the experiment.

their experiment. The image of a loaf of raisin bread provides a simple visualization of a net change in average choice distance. If yeast caused the loaf to rise, each raisin will be farther from every other. By analogy, an increase in mean choice distance is a net expansion of the graph. A net increase in the variance of this distance, which we interpret as stratification, corresponds to a case where the raisins in the bread all start equidistant from each other, and differently attract and repel each other over time. Together with wait tile and success tile, these graph statistics provide a powerful toolkit for investigating the internal structure of groups in our coordination game.

Table 1: Payoffs

Payoff	Observed		Min:Max[2]		Efficiency	
Harv. time	slow	fast	slow	fast	slow	fast
Low pay	21.5	62.7	25:99.4	75:298	-5%	-3%
High pay	68.1	181	25:430	75:289	9%	9%

Results

Group structure

Groups stratified with time, as seen in an increase by block in the variance of participants' choice distances from each other ($F(1,38) = 10.51, p < 0.05$). However, despite the reliable decrease in other measures of coordination (below), it does not seem that mean choice distance changed by block. The mean choice distance of participants from each other (the number of seconds that they spent on different tiles) was about 33 seconds in blocks one and two. The significant increase in standard deviation about this mean was from 7.5 seconds in the first block to 8.8 seconds in block two.

A change in the variance in a given participant's choice distances suggests an increase in graph heterogeneity: either a selective decrease in the minimum distances, a selective increase in the largest distances, or a combination of the two. An increase in the maximum distance would be consistent with choice refusal, a phenomenon observed in iterated prisoner's dilemma experiments by which participants choose not to reenter a game with a peer who has defected against them in previous iterations (Stanley, Ashlock, & Tesfatsion, 2004). Additionally, an increase in trust between only some peers in the group, and a corresponding net decrease in the minimum distance, could also account for the increase in variance with time. We found only insignificant evidence that the maximum choice distance increases with block ($F(1,38) = 2.34, p = 0.134$ with sample mean maximum distances of 79.8 and 83.1 for blocks one and two) and slightly stronger evidence that the minimum choice distance decreased with time ($F(1,38) = 3.88, p = 0.056$, with sample mean minimum distances of 54.6 and 50.2 by block). This experiment thus provides only minimal support for choice refusal and somewhat stronger support for trust building.

Coordination

In addition to these changes in group structure with time, we also documented decisive coordination failure. We define minimum efficiency as strictly individual, zero-coordination behavior (all players select tiles with coordination number one), and maximum efficiency as group-wide fully cooperative behavior.[2] Normalizing over these extremes, participants performed at negative, or very low efficiencies. For low pay level, observed efficiency was below that of purely individual behavior, presumably due to an excessive time cost of coordination, since waiting on a tile below quorum can yield no payoff. Similarly at the high pay level, the square scaling of payoffs by coordination number brought efficiency up only 9% above minimum (Table 1). Wait tile was on average 47% of maximum for each group size, and success tile was significantly lower at 35% ($t(656) = 6.49, p < .001$) (Table 2). This difference shows that participants consistently took risks on high coordination tiles that were not rewarded with success, even in the face of increasing coordination failure.

This failure of groups to coordinate at higher equilibria increased over the course of the experiment. Wait tile, success tile, and payoff all decreased slightly with block (all $p < 0.001, F(1,38) = 26.2, F(1,38) = 40.3$, and $F(1,42) = 14.4$). results on Table 3).

Only one exception appeared: one of the four 2-person groups had high coordination in the first block and reached near perfect levels of coordination in the second block. Although some groups did not produce significant decreases in coordination with time, this group was the only one to increase coordination, ultimately to the payoff-dominant equilibrium by the end of the experiment. On average, at group size two, wait and success tile stayed the same over block, and payoff did not increase significantly. This serve as limited evidence for the relative stability of two-player stag hunts in continuous time.

Supporting results

Another observed block effect was a significant reduction in participants' information entropy ($F(1,38) = 4.57, p < 0.01$), from 2.67 to 2.46 bits, suggesting that participants had more predictable behavior in the second block. This supports previous observations in a group minority game in which participants tended towards pure strategies over time (Bottazzi & Devetag, 2007), the key difference being that in the minority game experiment participants successfully coordinated towards benefits exceeding those predicted by the mixed-strategy equilibrium for that environment.

Both larger pay levels and shorter harvest times were strongly associated with higher wait tiles and higher success tiles (all $p < 0.001$, Table 3). Thus participants coordinated more when they were receiving higher payoffs, sup-

[2]Coordination number of all harvest tiles equals group size. Maxima in table calculated over four groups of 2, two groups of 3, two groups of 4, three groups of 5, and one group of 6. For the group of six, full cooperation implies that all six participants are on a tile of coordination five, the highest that we implemented.

porting the results of Brandts and Cooper (2006) in a minimum effort game. Participants' average choice distance correlated with wait tile ($\beta = -13.5, p < 0.05$) and success tile ($\beta = -20.0, p < 0.05$).

Table 2: Observed Wait and Success Tile by Group Size

Group size	2	3	4	5	6	ave	max[3]
Wait tile	1.5	1.7	2.6	2.9	3.1	2.5	3.98
% of max	.53	.37	.54	.49	.42	.47	1.00
Succ. tile	1.4	1.5	2.3	2.3	2.6	2.1	3.98
% of max	.40	.27	.45	.33	.31	.35	1.00

Discussion

Group structure

Investigation of participant time series revealed a detailed perspective into the evolution of group structure during the coordination task. Previous investigations into group behavior over time have focused on the differentiation of individuals into less complex roles (Bottazzi & Devetag, 2007; Roberts & Goldstone, 2009). This investigation extends these results by composing specific relationships into a group-level representation of interaction patterns. Groups tended to stratify in time as participants came to preferentially coordinate with and, to a lesser extent avoid, specific peers. In their model of Broadway musical production teams, Guimera et al. (2005) modeled the preference of past team members to work together in the future, but they included no complimentary aversion mechanism corresponding to choice refusal. Our results support this modeling decision.

Coordination failure

To our surprise, participants failed to coordinate on the most efficient outcomes and the extent of their coordination decreased over time. In their review of laboratory coordination failure, Devetag and Ortman propose a number of efficiency-enhancing design principles for coordination success (Devetag & Ortmann, 2007). Within the controlled factors of this experiment, we directly support two of the proposals in their review. We observed that an increase in pay level corresponded to increases in wait tile and payoff, supporting the efficiency-enhancing effect of "lowering the attractiveness of the secure action relative to the risky action required for the efficient equilibrium." Our manipulation of harvest time corresponded to both "lowering the costs of experimentation" and creating "less stringent coordination requirements," all of which are *ceteris paribus* efficiency-enhancing. In addition to this direct support, our design had a majority of the other features observed to be efficiency-enhancing: participants received zero deviation costs, they had repeated encounters in a fixed-group match, and participants had full information as to their peers' states and action choices. Also, although there was no explicit communication, participants' ability to make and change their decisions at low cost may have permitted signaling—some behavioral equivalent of "cheap talk" as Bottazzi and Devetag propose to explain their observations in a group minority game (2007)—thus providing another efficiency-enhancing factor.

In addition to implementing a majority of Devetag and Ortmann's efficiency-enhancing conditions, we also implemented a number of efficiency-enhancing conditions that they did not report. Our groups operated in a continuous time environment in which they made decisions asynchronously, and in which part of their decision was when, or whether, to make a decision. Continuous time environments have been observed to improve coordination in an experimental prisoner's dilemma (D. Friedman & Oprea, 2009) and a number of other studies (Berninghaus et al., 1999; Cheung & Friedman, 2009; E. Friedman et al., 2004). Asynchronous, extended form decision making is hypothesized to improve coordination in stag hunts (Kuhn, 2009). In the face of all of this previous qualitative evidence for high coordination outcomes, our participants still produced net coordination failure and increasing coordination failure with time, in all groups except one of the smallest.

Devetag and Ortmann make no claim for a simple additive relationship between their thirteen design principles, and they repeatedly stress that each factor is only efficiency-enhancing if all else is equal. Therefore, we think it is most reasonable to look to our experiment's major differences to explain the increasing failure of participants to coordinate on payoff-dominant equilibria. Despite the rich evidence supporting continuous time and asynchronous environments as efficiency-enhancing, we suspect that either of these factors or the third, non-forced decision making, may have overwhelmed the other efficiency-enhancing factors built into the experiment.

Conclusion

In an internet-enabled, computer-run coordination experiment, we show spontaneous group formation and differentiation. We also use functional network methods from computational neuroscience to document the fine scale structure of coordination failure in the face of many established determinants of coordination success. Within the experiment's controlled factors, we support previous claims with regard to the efficiency effects of payoff changes and uncertainty, but in the greater context of a coordination failure that increases over time.

Acknowledgments.

The authors wish to acknowledge Charlene Tay and the NSF, Research and Evaluation on Education in Science and Engineering, DRL-0910218

[3]This maximum is calculated over all groups of all group sizes.

Table 3: Main effects on Wait and Success Tile

	Low Pay Level	High Pay Level	Slow Harvest	Fast Harvest	Groups of Three	Groups of Six	Block One	Block Two
Mean wait tile	2.4	2.7	2.3	2.6	1.7	3.1	2.6	2.4
Mean success tile	2.1	2.3	2.0	2.1	1.4	2.6	2.2	2.0

References

Berninghaus, S., Ehrhart, K., & Keser, C. (1999). Continuous-time strategy selection in linear population games. *Experimental Economics*, *2*(1), 41–57.

Berninghaus, S., Ehrhart, K., & Ott, M. (2006). A network experiment in continuous time: The influence of link costs. *Experimental Economics*, *9*(3), 237—251.

Bottazzi, G., & Devetag, G. (2007). Competition and coordination in experimental minority games. *Journal of Evolutionary Economics*, *17*(3), 241–275.

Brandts, J., & Cooper, D. (2006). A change would do you good.... an experimental study on how to overcome coordination failure in organizations. *The American Economic Review*, *96*(3), 669-693.

Bullmore, E., & Sporns, O. (2009, Jan). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat Rev Neurosci*.

Cheung, Y., & Friedman, D. (2009, Jan). Speculative attacks: a laboratory study in continuous time. *Journal of International Money and Finance*.

Conradt, L., & Roper, T. (2005). Consensus decision making in animals. *Trends in Ecology & Evolution*, *20*(8), 449–456.

Couzin, I. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences*, *13*(1), 36–43.

Cover, T. M., & Thomas, J. A. (2006). *Elements of information theory*. Wiley-Interscience.

Devetag, G., & Ortmann, A. (2007). When and why? a critical survey on coordination failure in the laboratory. *Experimental Economics*, *10*(3), 331–344.

Duffy, J., & Hopkins, E. (2005). Learning, information, and sorting in market entry games: theory and evidence. *Games and Economic Behavior*, *51*(1), 31–62.

Dyer, J., Ioannou, C., Morrell, L., Croft, D., Couzin, I., Waters, D., et al. (2008). Consensus decision making in human crowds. *Animal Behaviour*, *75*(2), 461–470.

Friedman, D., & Oprea, R. (2009). A continuous dilemma. *Department of Economics, UCSC*, 657.

Friedman, E., Shor, M., Shenker, S., & Sopher, B. (2004). An experiment on learning with limited information: nonconvergence, experimentation cascades, and the advantage of being slow. *Games and Economic Behavior*, *47*(2), 325–352.

Guimera, R., Uzzi, B., Spiro, J., & Amaral, L. (2005). Team assembly mechanisms determine collaboration network structure and team performance. *Science*, *308*(5722), 697.

Hagmann, P., Cammoun, L., Gigandet, X., Meuli, R., Honey, C., Wedeen, V., et al. (2008). Mapping the structural core of human cerebral cortex. *PLoS Biology*, *6*(7), e159.

Harsanyi, J., & Selten, R. (1988). A general theory of equilibrium selection in games. *MIT Press Books*.

Huyck, J. V., Battalio, R., & Beil, R. (1990). Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*.

Kearns, M., Suri, S., & Montfort, N. (2006). An experimental study of the coloring problem on human subject networks. *Science*, *313*(5788), 824.

Kuhn, S. (2009). Prisoner's dilemma. In *Stanford encyclopedia of philosophy.* Metaphysics Research Lab, Stanford University.

Nagy, M., Akos, Z., Biro, D., & Vicsek, T. (2010). Hierarchical group dynamics in pigeon flocks. *Nature*, *464*(72907290), 890-893.

Pacheco, J., Santos, F., Souza, M., & Skyrms, B. (2009). Evolutionary dynamics of collective action in n-person stag hunt dilemmas. *Proceedings of the Royal Society of London, Series B: Biological Sciences*, *276*(1655), 315–321.

Petit, O., Gautrais, J., & Leca, J. (2009). Collective decision-making in white-faced capuchin monkeys. *Proceedings of the Royal Society B: Biological Sciences*, *276*(1672), 3495.

Rankin, F., Huyck, J. V., & Battalio, R. (2000). Strategic similarity and emergent conventions: Evidence from similar stag hunt games. *Games and Economic Behavior*, *32*(2), 315–337.

Resnick, E. (2007). Cooperation in multi-player stag hunt games.

Roberts, M. E., & Goldstone, R. L. (2009). Adaptive group coordination. In *Proceedings of the thirty-first annual conference of the cognitive science society.* Amsterdam, Netherlands.

Schelling, T. C. (1980). *The strategy of conflict*. Harvard University Press.

Schreiber, T. (2000). Measuring information transfer. *Physical review letters*, *85*(2), 461.

Stanley, E., Ashlock, D., & Tesfatsion, L. (2004). Iterated prisoner's dilemma with choice and refusal of partners. *Staff General Research Papers*.

Uzzi, B., & Spiro, J. (2005). Collaboration and creativity: The small world problem. *American Journal of Sociology*, *111*(2), 447–504.

Wasserman, S., & Faust, K. (1994). *Social network analysis: methods and applications*. Cambridge University Press.