EL FAROL REVISITED: HOW PEOPLE IN LARGE GROUPS LEARN TO COORDINATE THROUGH COMPLEMENTARY SCRIPTS

Corinne Coen (School of Business Administration)
Rick Riolo (Center for the Study of Complex Systems)
University of Michigan

INTRODUCTION

Norms enable and control much of social interaction in groups. Yet, despite calls for systematic research (Hackman, 1976; Feldman, 1984) little work has been done to understand how people jointly form group norms. An exception is a study by Bettenhausen and Murnighan (1991) in which they test how people use scripts based on prior experiences to coordinate actions. The authors observe that when scripts match, norms are formed tacitly. Norms are defined as “regular behavior patterns that are relatively stable and expected by a groups’ members” (p. 21, see also Bettenhausen & Murnighan, 1985). In contrast, when scripts conflict, people must explicitly defend their choice of action and alter the expectations of their partners. Still, the study investigated only a coordination problem in which behaviors had to match. The partners had to agree in order to act.

Another important type of coordination problem is one in which behaviors must be complementary but not identical. For example, in order to avoid conflict or wasted effort, group members commonly adopt different roles. As long as the behavior associated with these roles are compatible, the group can move swiftly toward task completion. Norm formation in this context has not been systematically researched. In this second part of a two-part study, we examine how coordination in the form of regularized behavior may occur tacitly among a relatively large number of actors.

In a now classic problem in the study of complex systems, W. Brian Arthur (1994) built a simulation to model how people might decide whether to attend the Irish music night at the El Farol bar in Santa Fe, New Mexico. Arthur noted that he enjoyed the Irish music except when the bar was overcrowded. He hypothesized that other people with similar preferences to his own must use some sort of inductive reasoning to estimate whether they were likely to enjoy an evening at the bar. In his model, the local newspaper published attendance figures from prior Thursday evenings. Each of 100 potential patrons had a different number of strategies *k* for deciding whether to attend in the upcoming week. These strategies were of the form “I expect the same number of people to attend as attended two weeks ago” or “attendance will be an average of the past four weeks attendance”[1].

All of the patrons felt that more than sixty people at the bar led to overcrowding. Every week each patron assessed whether his or her current strategy had been effective and whether another strategy among the *k* would have been more accurate in estimating levels of attendance over some past number of weeks. In the subsequent week, the patron used the strategy that would have been most effective previously. Arthur found that in a short period of time bar attendance settled into a dynamic equilibrium with a mean close to the desired level of sixty people attending. The patrons were able to coordinate by adopting different but mutually compatible strategies.

---

[1] Examples from John L. Casti’s review of the El Farol problem in Complexity, entitled “Seeing the Light at El Farol”.

Arthur compared these strategies to schemata or scripts (citing, among others, Schank & Abelson, 1977). He described these strategies as internal models built in response to perceived patterns of behavior and representing current hypotheses about how to act. Just as Bettenhausen & Murnighan (1985) defined norms as relatively stable and expected behaviors that are built up from common scripts, Arthur described his bar attendance as a stable behavior pattern that patrons infer from experience and represent as strategies. Bettenhausen & Murnighan (1991) describe how effective scripts are kept in use and ineffective ones altered by bargaining. Arthur's bar patrons keep a strategy in use as a representation of "temporarily fulfilled expectations" (p. 407) and replace it when the expectations were no longer fulfilled. Thus, Arthur's El Farol model maps closely with the assumptions and processes of Bettenhausen & Murnighan's (1991) norm formation experiment.

Recently, another model was built of the El Farol problem that challenges Arthur's results. Fogel, Chellapilla, & Angeline (1999) introduced evolution to their version of the El Farol model. They allowed their agents to make small changes in successful strategies and eliminate poorly performing strategies from the consideration set. While Arthur reported results for a model in which agents choose only from their initial set of k strategies (1994), he also speculated that allowing agents to create new strategies would not qualitatively alter the results he observed. However, Fogel et al. found that the 100 agents rarely achieved the ideal sixty patrons at the bar. In their experiments, the mean attendance was close to 56. What's more, attendance was not normally distributed. They found a bi-modal distribution, with peaks roughly around 40 and 80 patrons attending.

In this study, our goal was to replicate Fogel et al.'s (1999) model, dock it with Arthur's (1994) model[2] and account for the differences in results. We wished to better understand how a population of agents could find complementary strategies and thus coordinate behavior under various conditions. In the following sections, we describe how we built our models, why the two versions differ, and how to improve upon both versions.

THE CURRENT STUDY

Our first step was to replicate Fogel et al.'s model. We built our implementation within our preferred agent-based modeling software, Swarm, in the Objective C language on a Unix platform, sticking to the specifications described by Fogel et al.

Consistent with the Fogel et al. model, each of our agents had 20 strategies. A strategy was a series of weights (from -1.0 - +1.0), each of which was multiplied by a prior week's attendance, to generate an estimated future attendance. In addition, each strategy had a lag weight that was added to the estimated attendance total. Strategies varied in length from two weights (one plus the lag weight) to eleven weights. Every time step, each agent looked at what attendance level each strategy would have predicted for each of the last 12 weeks. Using the sum of squared prediction errors, the agent identified the

[2] See Axtell, Axelrod, Epstein, & Cohen [, 1995 #40] on the docking of computer simulations.

strategy best able to predict past attendance. Then the agent eliminated the ten strategies least able to predict past attendance, and created replicas of the ten best-performing strategies. The agent took these ten "offspring" strategies and slightly altered the weights in each time period (by adding a N(0, 0.1) normally distributed deviation), and either removed one weight, added one weight or left the number of weights the same (with probabilities 1/3). In every time step, the agent allowed ten generations of this mutation before it then selected its preferred strategy.

To replicate Fogel et al's results, we ran our model with all parameters set as they specified. Replicating Arthur's results is more problematic, since he did not fully specify his strategy representation nor his adaptive mechanisms. Thus, we used the Fogel model, with evolution of new strategies turned off, as the conditions analogous to those used by Arthur.

We believe our model is a good replication of Fogel et al., since our version produces the same results on all measures Fogel et al. described in their paper. In particular, our model produced (1) a mean of about 56.3 (s.d. 0.1) patrons attending the bar with the same bi-model distribution reported by Fogel et al, and (2) the variance of attendance within each run was also the same as reported by Fogel et al., 17.6 (s.d. 0.4)[3]. On the other hand, running the model under Arthurian conditions (i.e., with evolution of new strategies turned off), our model generates a mean of 58.7 (s.d. 0.6) patrons attending the bar, with an in-run variance much lower (7.5, s.d. 1.5) than obtained when run under Fogel et al.'s conditions. This mean is just under the threshold of 60, and significantly higher than for Fogel et al.'s evolutionary model.

WHY DOES FOGEL ET AL.'S MODEL LEAD TO LESSER COORDINATION?

We noted that each of three dimensions in Fogel et al.'s model diminished the ability of the group to coordinate its actions. As the evolution rate (the fraction of time steps when strategies are changed) moved from zero to 1, mean attendance dropped from over 58 to nearly 56 patrons. As the agents created new strategies by mutating strategies for more generations, from 0 to 30, the mean attendance slipped from 59 to less than 56. Similarly, as the mean number of strategies per agent move from 0 to 30, mean attendance also dropped below 60, to as little as 58.25. We found that the mean attendance was closest to 60 when the evolution rate went to zero, the generations mutating at each evolution went to zero, and as available strategies slipped to two.

At first glance, these observations seem counter-intuitive. We attribute these results to a number of factors, including (1) the biases introduced by Fogel et al.'s representation of strategies, (2) the pressure toward similar strategies in all agents exerted by Fogel et al.'s evolutionary mechanism, and (3) the non-linear interaction effects between agents in the

---

[3] Note that the in-run variance in attendance is one important measure of how well the public good, the bar, is being utilized. That is, if the variance is high, it means that on some time steps the attendance is way below the threshold and on others it is way above the threshold.

model. We will report on the first two factors in subsequent papers, and focus on the third factor here.

First note that any single agent seems like it would be better off by considering more options (more strategies, more variation among successful strategies), more often. But agents live in a world of other agents. As the number and variety of strategies used by others leads to more erratic behavior in the environment, any agent has a more difficult time assessing its own best choices. A strategy that would have been successful yesterday, is no longer effective in today's changed environment. Stability is more valuable to the system than innovation.

CAN WE IMPROVE ON ARTHUR'S MODEL?

We decided to take this idea about choice to its logical conclusion. We wondered how much Arthur's agents were able to outperform agents behaving randomly. So, we built another model in which agents simply predict attendance randomly (i.e., select a number from a uniform distribution of 0 to 100), and use that random prediction when choosing to attend the Irish music night at the bar or to stay home. These agents used no information from the environment. They did not consider how many people attended in the past. To our surprise, the population of agents in this model were better able to coordinate attendance than the agents in either Fogel et al.'s or Arthur's models! In particular, the mean attendance was 59.4 (s.d. 0.18) and the in-run variance in attendance was 4.9 (s.d. 0.1).

This result led us to think more deeply about the theoretical limits to Arthur's original model. We identified what we considered to be two key shortcomings in Arthur's thinking. First, we observed that all agents in Arthur's world reassess their strategies each time they want to go to the bar. We reasoned that agents who were successful in the past—who attended and did not find the bar overcrowded, or who stayed home and learned that the bar was overcrowded—were unlikely to change their strategies. More technically, we decided that using a model in which agents valued strategies by their ability to exactly predict attendance was excessively demanding. Humans are not good at this type of exacting prediction. Furthermore, the bar attendance problem does not demand this type of predictive ability. It would be sufficient for any agent if it could simply decide whether the bar would or would not be overcrowded.

We built a new model based on our revised view of the problem. Only the least successful fraction of the agents reassessed their strategies. We measured whether the agents were successful by whether they were "happy" (i.e., they stayed home when the bar was overcrowded or went to the bar when less than 60 other people were there), rather than by their strategies' ability to perfectly predict attendance[4]. With these new features, we evaluated the performance of the population as varying fractions of the agents identified themselves as successful. We found that the greatest level of coordination resulted when less than 10% of our 100 agents altered their strategies. This

[4] We decide which agents are "successful" based on attending or not attending, but we still have agents rate strategies based on precise predictions.

small percentage was robust to changes in the evolution rate. In other words, having few agents change strategy was crucial, no matter how often or rarely these strategy changes occurred. In complex systems terms, we found that coordination evolves only when a large fraction of the population sticks to relatively successful strategies while a very small fraction adapts to the fixed behavior of the larger subset.

DISCUSSION

Our research sheds light on issues around norm formation designed to regulate group coordination. Specifically, it helps us to understand how groups of people who need to behave in complementary ways coordinate inductively. By norms, we mean regular behavior patterns that are relatively stable and expected. We know that all groups develop norms and that these norms set in quickly.

Our study of the El Farol problem demonstrates that stabilizing successful behavior is crucial to coordination. When agents who are doing well change their actions, they force everyone else to adapt. This leads to recurrent patterns of change and an inability to settle down around simple issues and move on to the group's task. We saw that it is crucial to the group that most members adopt a single pattern of behavior and stick with it. The rapid adoption of stable behaviors is more important at the system level than optimizing the strategies of each individual member. If adaptation can help the system, only the least successful members should try new strategies. Through doing these tests, the least successful members might improve the group's performance but without de-stabilizing their teammates.

Too rapid change upsets all members of the group. The effectiveness of strategies is a relative measure. Thus, when everyone is reassessing and changing their own behaviors, they destroy the value of their teammates' assessment measures. The effective strategy can only be measured by the past and with too much change, the past does not map to the future. By fixing their own pattern of choices, members of the group create a space in which improvements can be made through slight adjustments in the areas of lowest performance. Coordination demands stability. Stability is best built in these conditions with approximations to success and allowing the system to settle before adjustments are made.

FUTURE RESEARCH

We have begun examining the learning experience of individual agents within poorly coordinated and the well-coordinated groups in the El Farol model. Our preliminary results suggest that while innovation by a large fraction of agents in a population generates a low overall success rate, success is distributed normally among the individual agents. In contrast, in conditions of random choice and when a small fraction of the population makes adaptive changes, approximately 70% of the agents are successful much of the time and the remaining 30% are unsuccessful much of the time. In confirming our preliminary results, we are seeking to explain why and how this divergence in the success of learning by individuals occurs in overall well-coordinated groups.

Bibliography

Arthur, W. B. (1994) Complexity in economic theory: inductive reasoning and bounded rationality. American Economic Association Papers and Proceedings, Vol. 84: 406-411.

Axtell, R., Axelrod, R., Epstein, J. M., & Cohen, M. D. (1995). Aligning simulation models: a case study and results. Computational and mathematical organization theory, 1(2), 123-142.

Casti, J. L. (1996) Seeing the Light at El Farol. Complexity, 7-10.

Bettenhausen, K. L., & Murnighan, J. K. (1985). The emergence of norms in competitive decision making groups. Administrative Science Quarterly, 30, 350-372.

Bettenhausen, K. L., & Murnighan, J. K. (1991). The development of an intragroup norm and the effects of interpersonal and structural challenges. Administrative Science Quarterly, 36, 20-35.

Feldman, D. C. (1984). The development and enforcement of group norms. Academy of Management Review, 9, 47-53.

Fogel, D.B., Chellapilla, K., & Angeline, P. J. (1999) Inductive reasoning and bounded rationality reconsidered. IEEE Transactions on Evolutionary Computation, 3, 2, 142-146.

Hackman, J. R. (1976). Group influences on individuals. In M. Dunnette (Ed.), Handbook of Industrial Organizational Psychology (pp. 1455-1526). Chicago: Rand McNally.

Shank, R., & Abelson, R. P. (1977) Scripts, plans, goals, and understanding: An inquiry into human knowledge structures. Hillside, NJ: Erlbaum.