Behavioral Experiments on Competitive Contagion

LILI DWORKIN and MICHAEL KEARNS, University of Pennsylvania

1. INTRODUCTION

We describe behavioral experiments based on the model of networked competitive contagion developed by Goyal and Kearns in [2012], which built upon the work of [Bharathi et al. 2007; Borodin et al. 2010; Chasparis and Shamma 2010]. In this model, there are two distinct "infections" competing for maximum spread through a social network. For instance, these infections might represent the products of two rival firms, or the ideas of two opposing political parties. Each individual in the network must choose only one of the two options, and this decision is stochastically influenced by the choices of his neighbors. We imagine that each competing firm can "seed" the initial adoption of its product by selecting certain individuals to receive promotions or give-aways. The goal of each firm is to choose these seeds to maximize eventual adoption throughout the network at the expense of its competitor.

In our experiments, a player takes the perspective of a firm, and chooses seed vertices for a series of networks. We evaluate a player on the performance of his seed choice against the distribution over the seed choices of all other players. Our goal is to study the collective performance of the population as a whole, which is measured by the population's distance to equilibrium on each network.

Our use of behavioral experiments to study strategic interaction in networks is inspired by the long line of work summarized in [Kearns 2012]. In these previous experiments, each participant was given control of a single vertex in a network, and provided with a local view of his immediate neighbors only. The experiment was performed simultaneously, with each player present in the same room. Our setting is notably different. We allow each subject a global view of the network when running simulations and making seed choices. Additionally, participation is asynchronous, with each participant playing on a web-based simulator at his convenience. The new setup allows for greater flexibility in experiment design, and makes participation easier for the subjects. Our motivation for shifting from local to global information stems from the increased awareness of networks in everyday life; individuals no longer simply act *within* networks, but have begun to reason *about* them as well. Thus, there is great value in discovering how people understand these mathematical models.

2. EXPERIMENT DESIGN

The precise game mechanics we consider are as follows. There are two competing players, denoted "Red" and "Blue." Each player is allowed to choose two seed vertices, which are immediately infected with his color. If a vertex is chosen by both players, its color is chosen at random. After the initial seeding, stochastic adoption dynamics determine the spread of each infection. We use a discrete time model. On each step, we consider all uninfected vertices that are adjacent to an infected vertex. If such a vertex has more Red neighbors than Blue neighbors, it becomes Red, and similarly for Blue. If a vertex's neighbors are evenly split, the color of the vertex is chosen at random. Thus, infection spreads via a stochastic breadth-first traversal. At the end of the process, every vertex will be infected. The score of each player depends on the fraction of vertices that are eventually infected with its color.

Rather than having two players compete directly, we use an asynchronous design in which each subject essentially competes against all other players simultaneously. Our population consists of 101 students taking the undergraduate course "Networked Life" at the University of Pennsylvania. Each

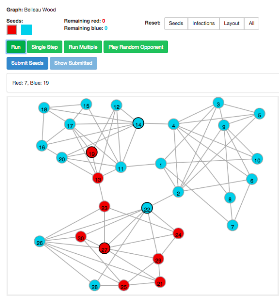

Fig. 1. A screenshot of our simulator.

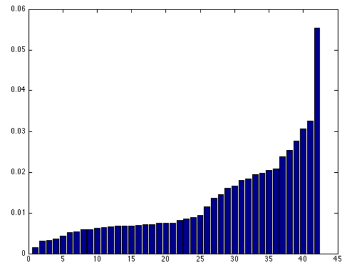

Fig. 2. Sorted histogram of the population's regrets on each of the 42 graphs.

student was required (for a class assignment) to choose the Red seeds for a series of 42 networks. The students were given access to a web-based simulator that allows them to experiment with different choices of Red and Blue seeds, see the randomized outcomes, and compute average adoptions over many simulations. A screenshot is shown in Figure 1. Importantly, the simulator also allows a student to play against the seed choices already submitted by his classmates (in the Blue role). That is, by pressing the "Play Random Opponent" button, a player can (repeatedly) sample a Blue seed pair from the distribution of submissions, and then run simulations to determine which Red seeds compete favorably. This permits a player to optimize his strategy against the current population of opponents. The students were also allowed to update their choices for the Red seeds as often as desired for the week-long duration of the experiment. Thus, the population evolved over time, as students returned to change their choices in response to the updates made by their classmates.

We next describe our scoring mechanism. The rules are designed to emphasize that, from an individual player's perspective, this is a two-player game in which the opponent is the *distribution* over all other players' choices. Let $G = (V, E)$ be a graph on which the game is played, and let $I_G(x, y)$ denote the expected fraction of Red infections when $x, y \in V^2$ are the initial seed pairs chosen by the Red and Blue player, respectively.[1] Let P_G denote the population distribution over seed pairs, and let U_G be the support of P_G. The *score* of seed pair x on G is then defined as $S_G(x) = \sum_{y \in U_G} P_G(y) I_G(x, y)$. In other words, the score of x is the expected fraction of infections that x wins against the population distribution P_G. A player's final score is the average of his seed scores over all graphs.

3. RESULTS

Given the incentives described above, the right measure of collective performance is the population's distance from *equilibrium*, the state in which no player can unilaterally improve his score by changing his seed choice.[2] To quantify this distance, we introduce the concept of *regret*, which measures how close each player's score is to the best player's score. More formally, the regret R_G of the population P_G on graph G is $\sum_{x \in U_G} P_G(x)(\max_{y \in U_G} S_G(y) - S_G(x))$. It is easily verified that equilibrium is reached if and only if 1) $R_G = 0$, and 2) there exists no seed pair z such that $S_G(z) > \max_{y \in U_G} S_G(y)$. Checking the satisfaction of condition 2 requires enumerating through all pairs of vertices, and is therefore intractable. Thus, we use R_G as a one-sided approximation of how far the population is from equilibrium.

[1]We estimate $I_G(x, y)$ using 1000 offline simulations.
[2]Since players are scored against the population distribution, the equilibrium concept here is actually that of an Evolutionary Stable Strategy (ESS).

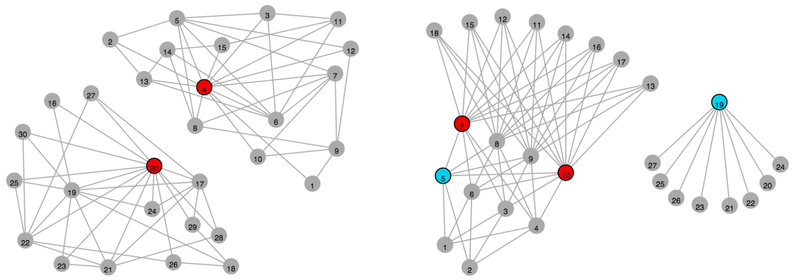

(a) Spotsylvania (b) Kasserine Pass

Fig. 3. Spotsylvania and Kasserine Pass had the lowest and highest regret of all 42 graphs, respectively. On each graph, we plot the majority seed choice in red. On Kasserine Pass, we plot a less frequent but better alternative in blue.

The average regret across all 42 graphs is 0.0129, less than 3% of the average player score of 0.4911, indicating strong collective performance by this measure. Yet if we analyze the graphs individually, we find significant variation. See Figure 2 for a sorted histogram of the population's regret on each graph. Broadly speaking, there seem to be two types of graphs. The first type consists of graphs in which there exists a *dominant strategy*, i.e. a seed pair that always infects at least as many vertices as the seed pair it plays against, regardless of the identity of this competing seed pair. In this case, the distribution over opponents' choices is irrelevant, and the game becomes a single-player optimization problem in which the goal is to identify the dominant choice. See Figure 3(a) for an illustration of such a graph, where the dominant strategy is the seed pair shown in red. Note that equilibrium can be achieved on this type of graph when everyone plays the dominant strategy. In general, the population performed quite well on such graphs, with almost everyone converging on the dominant seed choice.

Graphs of the second type have no dominant strategy, and so the best seed choices depend strongly on the population distribution. Here a coordination problem arises: in order to reach equilibrium, different portions of the population must choose different seeds. The task now involves game-theoretic strategizing rather than pure optimization. See Figure 3(b) for an example. If all players choose seeds (7,10), one should choose (5, 19). But if all players choose (5, 19), one should choose (4, 7). In our study, the majority chose seeds (7,10), while a minority chose (5, 19) and thus achieved a better score, causing high regret. As a further example, note that all three-component graphs lack a dominant strategy because a player is only allowed two seeds, and so the choice of which components in which to play depends on the population distribution. In order to reach equilibrium, certain fractions of the population have to play in each component pair, but we never observed such behavior. As a result, the average regret on three-component graphs is 0.0297, more than twice the overall average.

The discrepancy in performance on graphs of each type suggests that when equilibrium requires diversity among strategies, subjects perform more poorly. To quantify this further, we calculate the entropy of the population distribution of each graph. Note that on graphs of the first type, equilibrium is reached when everyone makes the same seed choice, and therefore the entropy of the distribution is zero. But on graphs of the second type, equilibrium requires a distribution with non-zero entropy. We found a strong positive correlation between the population's regret on a graph and the entropy of the distribution (namely a Pearson coefficient of 0.6659 with insignificant p-value), which supports the observation that the population struggles with the coordination of diverse strategies.

We conclude by reiterating that, on average, human subjects demonstrate good performance on this problem, and are capable of approaching equilibrium even on difficult graphs. While we have restricted our attention to graph-level analyses, we note that there are interesting subject-level findings as well, which suggest that certain players are much better than the average at this task.

REFERENCES

Shishir Bharathi, David Kempe, and Mahyar Salek. 2007. Competitive Influence Maximization in Social Networks. In *Proceedings of the 3rd International Conference on Internet and Network Economics (WINE'07)*. Springer-Verlag, Berlin, Heidelberg, 306–311.

Allan Borodin, Yuval Filmus, and Joel Oren. 2010. Threshold Models for Competitive Influence in Social Networks. In *Proceedings of the 6th International Conference on Internet and Network Economics (WINE'10)*. Springer-Verlag, Berlin, Heidelberg, 539–550.

Georgios C. Chasparis and J.S. Shamma. 2010. Control of preferences in social networks. In *Proceedins of the 49th IEEE Conference on Decision and Control*. 6651–6656.

Sanjeev Goyal and Michael Kearns. 2012. Competitive Contagion in Networks. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing (STOC '12)*. ACM, New York, NY, USA, 759–774.

Michael Kearns. 2012. Experiments in Social Computation. *Commun. ACM* 55, 10 (Oct. 2012), 56–67.