

Bayesian Security Games for Controlling Contagion

Jason Tsai¹, Yundi Qian¹, Yevgeniy Vorobeychik², Christopher Kiekintveld³, Milind Tambe¹

¹University of Southern California, Los Angeles, CA
{jasontts, yundi.qian, tambe}@usc.edu

²Sandia National Laboratories, Livermore, CA
yvorobe@sandia.gov

³University of Texas at El Paso, El Paso, TX
cdkiekintveld@utep.edu

ABSTRACT

Influence blocking games have been used to model adversarial domains with a social component, such as counterinsurgency. In these games, a mitigator attempts to minimize the efforts of an influencer to spread his agenda across a social network which is modeled as a graph. Previous work has assumed that the influence graph structure is known with certainty by both players. However, in reality, there is often significant information asymmetry between the mitigator and the influencer. We introduce a model of this information asymmetry as a two-player zero-sum Bayesian game. Nearly all past work in influence maximization and social network analysis suggests that graph structure is fundamental in strategy generation, leading to an expectation that solving the Bayesian game exactly would be vastly superior to any technique that does not account for uncertainty about the network structure. Surprisingly, we show through extensive experimentation on synthetic and real-world social networks that many common forms of uncertainty can be addressed near-optimally by ignoring the vast majority of it and simply solving an abstracted game with a few randomly chosen types. This suggests that optimal strategies of games that do not model the full range of uncertainty in influence blocking games are in many cases robust to uncertainty about the structure of the influence graph.

1. INTRODUCTION

Social contagion has long been of great interest in the literature on marketing, the spread of rumors, and, recently, in the context of Arab Spring [11, 13, 19]. Our specific focus is on counterinsurgency, which we view as a competition for the support of local leaders. Counterinsurgency can be modeled as a game with two strategic players, the insurgents and the peacekeepers, in which the insurgents aim to spread their views, unrest, etc. among the local population, while the peacekeepers wish to minimize the resulting contagion by engaging in their own influence campaign [8, 7, 20]. The key *computational* question we address is: given limited resources, how to select which of the local leaders to influence to minimize the global impact of the insurgency.

These ‘influence blocking’ games have received recent attention in the security games literature [20], where they have been modeled using graphs with nodes representing the tribal leaders and edges

representing possible transmission of influence. However, this line of work has assumed that full information about network structure is available to both players and that accurate network knowledge is crucial to generating high-quality strategies. In practice, informational challenges abound in counterinsurgency, where the insurgents are typically an indigenous group that has an informational advantage, and the mitigators often have uncertainty about their knowledge of the social network [8].

In this work, we model counterinsurgency as an influence blocking game with asymmetric information. Specifically, we assume that the influencer (an insurgent group) has perfect knowledge of the graph structure, while the mitigator is uncertain about the influence network. In the resulting Bayesian game, a type of the influencer corresponds to a particular instantiation of the influence graph, and the mitigator must reason over the distribution over these graphs (i.e., influencer types) in order to compute an optimal strategy.

Past work in influence maximization and social network analysis highlight the importance of graph structure in strategy generation [11, 3, 6]. In addition, previous work on Bayesian security games has shown that not accounting for even small degrees of payoff uncertainty can lead to large drops in solution quality [12]. Thus, we expect strategies generated without modeling most of the uncertainty about graph structure to do far worse than the optimal solution to the Bayesian game. Supporting this intuition, we show that there are cases where a mitigator who has incorrect information about a single edge can suffer unbounded loss and that quantifying the impact of changing a single edge in a given graph is #P-Hard. We also show empirically that, indeed, under our models of uncertainty, optimal mitigator strategies for different influencer types are vastly different.

However, while past work has focused on sophisticated algorithms for Bayesian security games [9, 12, 22], we showcase the opposite approach that runs directly counter to what intuition and our initial experiments suggest: ignoring the vast majority of uncertainty has minimal impact. Specifically, we show through extensive experiments that computing a mitigation strategy based on a game with only a few *randomly* sampled influencer types yields near-optimal rewards for widely varied models of uncertainty. We experiment on 3 different synthetic graph models with and without resource imbalances on both sides, 5 models of uncertainty, weighted/unweighted counting of nodes, varied edge weight distributions, varied graph sizes, varied degrees of uncertainty, and varied degrees of sampling. We also conduct experiments on two real-world social networks using two different models graph construction. In all, we studied over 200 experimental settings and consistently observe the same result: simple sampling techniques

perform near-optimally. This suggests that even in domains as challenging as ours, models which ignore uncertainty may nevertheless be robust to it.

2. MODEL

2.1 Asymmetric Information Game

We model counterinsurgency as a two-player Bayesian zero-sum game situated on a graph in which two players, the influencer (denoted by I) and the mitigator (denoted by M) compete to maximize influence over the nodes. Formally, let $G = (V, E)$ be a graph with weighted nodes V and edges E , and for each edge $(i, j) \in E$, let p_{ij} be the probability that node i 's opinion will influence node j . Suppose that the influencer initially attempts to influence a subset of nodes $S_I \subseteq V$ to his cause, and the mitigator's initial influence is aimed at a subset of nodes $S_M \subseteq V$. We model propagation of influence in the graph as a synchronized independent cascade process [11, 20] as follows. For nodes $v \in S_I \cap S_M$ which both players initially try to influence, initial "activation" (e.g., actual opinion adoption) happens in either player's favor with equal probability, while all the remaining nodes adopt the view of (are activated by) the player who directly targets them. Next, we activate *all* edges (i, j) in the graph with the corresponding probability of influence, p_{ij} , yielding a subgraph upon which influence can spread. At that point, the influence process proceeds through a sequence of iterations. In each iteration, if a node j has not yet adopted an opinion but has active edges to neighbors who have, j either adopts the opinion of these neighbors when it is unanimous, or adopts each opinion with equal probability if j 's active neighbors disagree. Viewing now the initial target nodes S_I and S_M as the strategies of the players I and M respectively, let $\sigma(S_I, S_M)$ be the expected value of nodes that adopt the influencer's opinion following the independent cascade process described above. We define the utility of the influencer to be $U_I(S_I, S_M) = \sigma(S_I, S_M)$.

We now depart from the model of Tsai et al. [20] by relaxing the complete/symmetric information assumption. Specifically, we assume that the influencer knows the actual influence graph G exactly, while the mitigator is uncertain about its true structure, and only knows the probability distribution over possible graphs. Let λ be an index identifying a particular graph G_λ , and let us make explicit the dependence of the expected influence on the graph, denoting it by $\sigma(S_I, S_M, \lambda)$. Finally, we denote by P the probability distribution over λ , with P_λ being the probability that the true graph is the one identified by λ . From the mitigator's perspective, the influencer's decision will depend on his type; that is, on the true graph which the influencer observes. Thus, we view the influencer's strategy S_I as a function of λ , with S_I^λ denoting the influencer's strategy when his type is λ . The mitigator's utility is then $U_M(S_I, S_M) = -E_{\lambda \sim P}[\sigma(S_I^\lambda, S_M, \lambda)]$.

2.2 Models of Networks and Uncertainty

Numerous stochastic generative models for graphs have been proposed to generate synthetic instances of graphs that resemble real social networks [14]; some of the best known examples are the preferential attachment process, which generates scale-free graphs [2], and the process of generating small-world networks pioneered by Watts and Strogatz [21]. Recently, a new generative model, BTER, has been developed, and the authors convincingly demonstrated that this model matches the important properties of real-world networks, such as the distribution of degrees and clustering coefficients, far better than previously proposed methods [17]. BTER graphs feature a scale-free collection of densely clustered community structures (dense Erdős-Rényi subgraphs), which are

sparsely interconnected by 'inter-community' edges. We conducted experiments on BTER graphs (including variations in community density and interconnectedness), small-world graphs (Watts-Strogatz), preferential attachment graphs, and real-world networks from two villages in India. Due to space limitations, we show results for BTER graphs and two of the Indian villages here and post the remainder in an online appendix: <http://main2013.webs.com>.

We consider several ways to model the mitigator's uncertainty about the graph. *Influential Node* uncertainty models uncertainty about which nodes are most connected, motivated by the fact the identity of the most socially connected and influential individuals is a function of the local culture which is more familiar to the influencer than the mitigator. Specifically, we start with a baseline graph, then, for each type, choose a set of j nodes and add k new randomly chosen edges from each of these nodes to others. It is important to note that in BTER graph, these j nodes are the *only* nodes that can potentially have inter-community edges under this uncertainty. These inter-community edges are particularly important in contagion games because they enable the spread of influence across groups. The second model, *Random Edge* uncertainty, is the simplest: the mitigator has perfect information about the nodes in the graph, and is uncertain about which edges out of a given set exist. The third model of uncertainty, *Inter-community Edge* uncertainty, models the mitigator's uncertainty about a subset of the inter-community edges (i.e., which edges out of a given set of inter-community edges exist). The fourth model of uncertainty, *Inter/Intra-Community Edge* uncertainty, models uncertainty about a combination of inter-community *and* intra-community edges and addresses the concern that Inter-community Edge Uncertainty may provide additional information by being restricted to the critical inter-community edges. Note that in Inter-community Edge, Inter/Intra-Community Edge, and Random Edge uncertainty, we have a type λ for each possible subset of uncertain edges in the graph so the number of types could be as large as $2^{|E|}$. The fifth model of uncertainty, *Inter-community Edge Set* uncertainty, models uncertainty over which *set* of inter-community edges exists (i.e., which set of 8 edges exists). The final three uncertainties, which highlight inter-community edges, apply only to BTER graphs. Here we only show results for Influential Node and Inter-community Edge uncertainty and post the rest on the online appendix. The results omitted are extremely similar to those for Inter-community Edge uncertainty.

The counterinsurgency literature [8] makes clear that military intelligence explicitly performs 'intelligence preparation of the battlefield (IPB)' to ascertain the structure and dynamics of a local population with high fidelity. Therefore, we are not interested in cases in which the entire social network is largely unknown or misunderstood. Instead, we focus on situations with a generally correct social network in which uncertainty is about the details of the network structure.

3. THE CHALLENGES OF UNCERTAINTY

The first question to ask is whether we can bound the impact of a small amount of uncertainty, because that may help us bound the total loss of solution quality given some uncertainty. We show that, in general, ignoring uncertainty can yield a solution that is arbitrarily poor for the mitigator. Consider the graph shown in Figure 1 in which the edge from A to B is uncertain, $N > M$, and both players have a single resource. Suppose that the influencer chooses to influence node A with probability 1. If the mitigator mistakenly assumes the edge does not exist, then his best response is to influence node C with probability 1, but his actual loss amounts to $\frac{N}{2}$ as compared to the true best-response of playing B ($\frac{M}{2}$). A similar situation arises when the mitigator assumes the opposite. Thus, since




Figure 1: Unbounded loss

M is arbitrary, by ignoring the uncertainty of just a single edge the mitigator can suffer unbounded loss.

The network in the above example is rather artificial, so it is natural to wonder what happens under a more realistic model of a network and uncertainty. To this end, we investigate the following empirical question: under our models of uncertainty, if we were to compute an optimal strategy assuming a single influencer type, how much would that strategy vary for different types? To answer this, we take a Bayesian game with 40 types and compute an optimal mitigation strategy for each possible influencer type λ under the assumption of complete information. This yields a mixed strategy, S_M^λ , for each possible influencer type. Next, we select a type b uniformly at random and measure the fraction of pure strategies in the support of S_M^b that is different from the pure strategies in the support of each S_M^λ for $\lambda \neq b$. In Figure 2 we report the *average* fractional difference over 20 independent instances of 40-type Bayesian games on 40-node BTER graphs (edges vary from 130 to 200) with Influential Node uncertainty (more details on our standard setup will be presented shortly). Note that 1 in this case indicates that the mixed strategy for a randomly chosen type does not share a *single* pure strategy with the mixed strategy computed for *any* other type. As can be seen here for this instance, and is generally true under this uncertainty, nearly all instances show minimal overlap in the pure strategies used by each of the type-specific optimal strategies.




Figure 2: Comparison of mixed strategies

Finally, we turn to the question of complexity, where the result is very clear and very negative. At a high level, the challenge of efficiently reducing the runtime of computing equilibria in our setting lies in quantifying the impact of even small changes in the graph structure. If this could quickly and accurately be determined, then types could be efficiently clustered and bounds could be placed on the quality loss. The fact that computing the expected influence is #P-Hard [4] should already give us pause. Indeed, a simple corollary of this result reveals that such quantification is intractable in general.

PROPOSITION 1. *Computing the difference in expected influence for a given seed set even when a single edge is added to a graph is #P-Hard. (Proof available in appendix)*

4. DOUBLE ORACLE ALGORITHM

Even though we formulated influence blocking as a zero-sum game, which can in principle be solved using linear programming, computing an equilibrium of this game in our case is challenging for three reasons. First, payoff estimation requires determining the value of $\sigma(S_I^\lambda, S_M, \lambda)$, which has been shown to be #P-Hard [4]. Therefore, even constructing the payoff matrix for this linear program is non-trivial. Second, the strategy sets for both players are exponentially large, making it impractical to store the entire payoff matrix even if we could compute payoffs efficiently for a pair of player strategies. Third, because we model uncertainty over graph instances, the number of influencer types can be exponentially large.

The first problem was addressed in prior research [20] by introducing the LSMI heuristic for faster estimation of $\sigma(\cdot)$, which we also use here. The Bayesian double-oracle algorithm introduced by Halvorson et al. [5] provides a possible solution to the second problem. This algorithm begins with a small subset of pure strategies for each player and iteratively adds best-response strategies to the existing subgame. The algorithm ends when no new best-responses need to be added, at which point it has provably converged to the equilibrium of the full game. In the context of Bayesian games, Halvorson et al. propose computing the best response for every player type, which in our case means that we compute the influencer’s best response for each type (graph), and add all of these pure strategies in each iteration. This approach runs into our third and final problem: the exponential number of types. Since computing a best response for a given type requires a non-negligible amount of computation, having to do this for every type will simply not scale. To address this, we now show empirically that simple heuristics actually produce near-optimal solutions.

5. THE POWER OF SIMPLE

The results presented thus far, as well as the intuition from the vast literature on influence maximization [13, 4, 3], suggest that carefully accounting for our uncertainty about graph structure is crucial to obtaining high quality solutions. Next, we present a small, representative subset of an extensive collection of experiments, all showing precisely the opposite: we need only to randomly sample a few types from the type distribution and solve the resulting game as if no other types exist, to obtain solutions that are nearly optimal. This is quite surprising, particularly since we have already shown, via the example in Figure 1, that ignoring even a single influencer type can yield arbitrarily poor solutions even with only two types.

All the results below are an average of 20 game instances and were run on machines with CPLEX 12.2, 2.8 GHz CPU, and 4GB of RAM. Unless otherwise stated, experiments were run on 40-node graphs (130 to 200 edges), contagion probabilities on edges drawn from a $\mathcal{N}(0.4, 0.2)$ distribution, node values varying uniformly from 1-10, each player having two seed nodes ($|S_I| = |S_M| = 2$), and payoffs estimated using the LSMI heuristic introduced by Tsai et al. [20]. Monte Carlo payoff estimations produced similar results but could not be meaningfully scaled. Since an optimal benchmark is necessary, the best-response oracles iteratively evaluate each available action to determine the best response, rather than using greedy hill-climbing common in the influence maximization literature. Unless otherwise stated, Influential Node uncertainty selects 3 nodes and gives each 4 additional edges. Moreover, *only* these 12 edges could potentially connect communities, making the chosen nodes not only more connected (average degree, excluding uncertain edges, varies from 3-5 with maximums

of 9), but also incident to the more consequential edges. For Inter-Community Edge uncertainty we varied the number of uncertain edges between 1 and 6 (the optimal technique could not scale to more edges). We focus throughout on the mitigator strategy obtained by drawing a random subset of the influencer’s types and solving the game assuming no other types exist (referred to as *Random Sampling*).

5.1 Experiments




Figure 3: Reward comparison, BTER graphs

In our first set of results, shown in Figure 3, we consider the impact of the number of randomly sampled types on solution quality (only a combination of BTER and two models of uncertainty are shown here, as these results exhibit the *greatest* approximation error from random sampling; extensive other studies, included in the online appendix, offer even more dramatic support of our argument). These experiments use the *same* 40-node games that were featured in Figure 2 that showed pure strategies used by individual types have minimal overlap. The x -axis shows the number of sampled types, while the mitigator utility is plotted on the y -axis. The key point is that with only about 2-5 randomly sampled types we obtain a solution that is very nearly optimal, despite the fact that only using a single influencer type yields a relatively poor mitigator reward (Figure 3b). While results in the optimization literature such as sample average approximation theory [18] show that random sampling can converge exponentially fast to optimal solutions, our “convergence” is uncannily quick.




Figure 4: Inter-community Edge uncertainty

Next, we fix the number of randomly sampled types used to generate a solution at 2, and increase the number of *actual* types (increasing the degree of uncertainty). The graph sizes were fixed to 40 nodes. Intuitively, we would expect that the performance of *Random Sampling* should degrade significantly as we increase uncertainty by adding types. In addition, we compare the random sampling strategy to an even simpler heuristic which uses only a single type with the highest probability; we call this *Max Prob*. Figures 4 and 5 are representative of a broad array of experiments we ran in this space (see online appendix). In addition to considering




Figure 5: Influential Node uncertainty

several types of uncertainty, we also varied the density of connections among communities (low density uses $\rho = 0.5$ as the probability of inter-community edges, while high density uses $\rho = 0.9$). Perhaps the most surprising finding in these experiments is that the quality of *Random Sampling* relative to optimal degrades very little as we increase the number of types. While we could not compute optimal solutions for games involving more types, this finding suggests that we may need to sample a decreasing (rather than a constant) fraction of all possible types as the number of total types increases.

In our final set of results using synthetic graphs, we study the impact of the size of the underlying network. The number of edges varied from 28 (20 nodes) to 188 (40 nodes) with up to 6 edges differing between types for Inter-community Edge uncertainty and up to 24 edges for Influential Node uncertainty (12 new edges per type). Here, we keep the number of nodes/edges about which we are uncertain fixed, and increase the network size. Consequently, our expectation is that smaller networks would exhibit significantly greater difference between random sampling and optimal, since uncertainty involves a greater fraction of the graph. Figure 6 shows




Figure 6: Scale-up of graph size

little evidence of this: the quality of simple heuristics relative to optimal is little affected by the fraction of the graph that is uncertain.

In all, we studied variations involving BTER, preferential attachment, and small-world generative models of networks, all five models of uncertainty described above, and, moreover, considered numerous variations in the parameter space of all the generative models of graphs and uncertainty about them, the number of resources that players had, etc. For example, we studied games in which the mitigator was allowed to initially impact 3 or 4 nodes, while the influencer was restricted to 2, and vice versa; we varied the distribution of contagion probabilities between 0.4 and 0.7; for BTER we additionally examined different degree distributions. All these results (see online appendix) exhibit essentially the same trends that we showed here.

Finally, we conducted a set of experiments on a real-world social network dataset released in 2012 that was obtained via survey data

in 75 Indian villages.¹ The survey asked the inhabitants of the villages a series of questions to ascertain their relationship with other people in the village (e.g., would you invite him in for tea, do you go to temple with him, would you loan him money, etc.). From this data, a social network can be constructed by beginning with a complete graph with edge weights of 0.0, increasing the weight of an edge corresponding to a positive answer to a survey question by 0.1, and then normalizing all weights. For our experiments, we use the household-level data for two of the smaller villages (8 and 10), because even the double-oracle optimization does not scale to larger networks. The results in Figure 7 use Influential Node uncertainty, and each type now chooses 8 random nodes and gives each 10 new edges to maintain the same fraction of uncertainty, since the India data sets have 77 or 94 nodes and an average degree of 7.7 or 7.4. As the figure testifies, our results are not an artifact of synthetic graph models that we generate, but can be observed on graphs based on actual social network data as well. Moreover, we studied others variations, such as changing the way we weigh the edges (increasing the relative weight of stronger relationships), with the results virtually identical to what we show here (see appendix: <http://main2013.webs.com>).




Figure 7: Influential Node uncertainty, 0.1-Weight-Scheme

5.2 Initial Analysis

The results shown are surprising in their extremity, especially in light of the result presented previously demonstrating minimal overlap of pure strategies in optimal strategies for individual types. We now explore why this might be occurring in these games. We begin by again comparing the type-by-type overlap in mixed strategies, but this time we focus on the *nodes* used instead of the actions (i.e., sets of nodes) used as a more granular metric. In addition to comparing against a randomly chosen base type as before, we also compare each type’s individual optimal strategy against the optimal strategy for the full Bayesian game.

Figure 8 shows the results for 20 trials where each bar represents the average percentage of node overlap for a single trial (averaged over all types). We see in Figure 8a that there is a 60-80% difference (average of 74%) in the nodes used by individual types when compared against a randomly chosen type’s optimal strategy. Thus, while each type’s strategy may differ, they may all contain a set of core nodes that overlap more with the optimal strategy, which may cause the type-specific strategies to perform well overall. In Figure 8b we show the results when comparing the overlap in nodes used between the optimal strategy and each type’s individual optimal strategy. The difference drops to the 40-70% range (average




Figure 8: Overlap in nodes used

of 62%), suggesting that each individual type’s optimal strategy uses nearly half of the nodes used by the optimal Bayesian strategy. The existence of such a core of nodes that are part of the optimal strategy for many types can partially explain the success of simple sampling techniques.

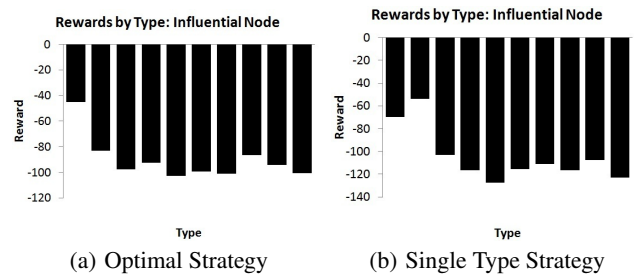


Figure 9: Performance distribution, Influential Node uncertainty

Digging deeper, we examine the type-by-type reward obtained by the sampling strategy versus the reward obtained by the optimal Bayesian strategy under Influential Node uncertainty. These experiments were examined instance by instance, and we show a typical example to illustrate the trends. In Figure 9a, we show results for the optimal strategy’s performance on each type, where each bar shows the mitigator’s reward on the y -axis for a particular type. Note that there is one extremely high reward type and the others are all in the -80 to -100 range. In Figure 9b, we show the results for a randomly chosen single type’s optimal strategy. Here the majority of rewards range from -100 to -120, suggesting that the overlap in nodes showed previously may stem some of the losses (maximum loss is -200 in expectation), but does not explain the whole story. An additional clue, however, is that the single type’s strategy actually performed extremely well for two types here, leading to a comparable average reward despite the poorer performance in most cases. While this was not universally true, in most cases a similar phenomenon occurred. This suggests that within the forms of uncertainty explored here, an optimal strategy for one type tends to be near-optimal for a handful of other types (despite minimal overlap in optimal support sets). Under what formal circumstances this reliably occurs, however, remains an open question.

6. DISCUSSION

The phenomenon of simple techniques providing highly effective solutions has also been observed elsewhere [18, 10, 16, 1]. In addition to novelty of our influence-driven, network-based model, our work differs from these in other important ways. As noted earlier, our results differ from previous work using sampling techniques in their extremity (i.e., only one or two samples needed),

¹Abhijit Banerjee; Arun Chandrasekar; Esther Duflo; Matthew Jackson, 2011-08, "Social Networks and Microfinance", <http://hdl.handle.net/1902.1/16559> UNF:5:4EmgOYAQGaoQugFowckNfA== Jameel Poverty Action Lab [Distributor] V5 [Version]

and unlike research in heuristic techniques for equilibrium computation, our work focuses on the power of extremely few samples instead of general heuristics.

A closer examination of previous literature in security games that addresses uncertainty reveals that similar phenomena may have been true elsewhere but went unexplored. In Yin et al. [22], the authors provide a novel algorithm (HUNTER) for optimally handling Bayesian Stackelberg games with many types. While the algorithm is orders of magnitude faster than previously proposed optimal algorithms, the authors report that BRASS, a far less complex solution method [15], achieves an average loss of 0.7 in a game where the range of rewards for optimal solutions ranged from -26 to 17 compared against their algorithm. One again wonders whether a sampling approach would have worked extremely well here too.

Our work does not dispute the fact that extremely large Bayesian zero-sum games remain very challenging to solve well in general and there are certainly problem classes that are not amenable to simple heuristics. In Kiekintveld et al. [12], for example, the authors introduce several techniques for handling large numbers of Bayesian types to address payoff uncertainty and they show that simple techniques do not perform near-optimally. Our work stresses the need to verify whether or not simple techniques work before embarking on extensive algorithmic gymnastics to achieve minimal gains in solution quality. Although we have provided some analysis of why this occurs in our domain, this is only the beginning of research in this direction and clearly more work is needed. Finally, our findings give hope that many very challenging problems in computational game theory may actually be very effectively addressed by simple techniques: the power of simple.

7. REFERENCES

- [1] O. Armantier, J.-P. Florens, and J.-F. Richard. Approximation of nash equilibria in bayesian games. *Journal of Applied Econometrics*, 23(7):965–981, 2008.
- [2] A.-L. Barabási and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [3] C. Budak, D. Agrawal, and A. E. Abbadi. Limiting the spread of misinformation in social networks. In *WWW*, pages 665–674, 2011.
- [4] W. Chen, C. Wang, and Y. Wang. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *KDD*, pages 1029–1038, 2010.
- [5] E. Halvorson, V. Conitzer, and R. Parr. Multi-step multi-sensor hide-seeker games. In *IJCAI*, pages 159–166, 2009.
- [6] X. He, G. Song, W. Chen, and Q. Jiang. Influence blocking maximization in social networks under the competitive linear threshold model. In *SDM*, pages 463–474, 2012.
- [7] N. J. Howard. *Finding optimal strategies for influencing social networks in two player games*. Masters thesis, MIT, Sloan School of Management, June 2011.
- [8] B. W. K. Hung. *Optimization-Based Selection of Influential Agents in a Rural Afghan Social Network*. Masters thesis, MIT, Sloan School of Management, June 2010.
- [9] M. Jain, M. Tambe, and C. Kiekintveld. Quality-bounded solutions for finite bayesian stackelberg games: Scaling up. In *In AAMAS*, pages 997–1004, 2011.
- [10] M. Kearns, Y. Mansour, and A. Y. Ng. A sparse sampling algorithm for near-optimal planning in large markov decision processes. *Machine Learning*, 49:193–208, 2002.
- [11] D. Kempe, J. M. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *KDD*, pages 137–146, 2003.
- [12] C. Kiekintveld, J. Marecki, and M. Tambe. Approximation methods for infinite bayesian stackelberg games: Modeling distributional payoff uncertainty. In *International Conference on Autonomous Agents and Multiagent Systems*, 2011.
- [13] J. Leskovec, A. Krause, C. Guestrin, C. Faloutsos, J. M. VanBriesen, and N. S. Glance. Cost-effective outbreak detection in networks. In *KDD*, pages 420–429, 2007.
- [14] M. Newman. *Networks: An Introduction*. Oxford University Press, 2010.
- [15] J. Pita, M. Jain, F. Ordóñez, M. Tambe, S. Kraus, and R. Magori-Cohen. Effective solutions for real-world stackelberg games: When agents must deal with human uncertainties. In *The Eighth International Conference on Autonomous Agents and Multiagent Systems*, 2009.
- [16] R. Porter, E. Nudelman, and Y. Shoham. Simple search methods for finding a nash equilibrium. In *Games and Economic Behavior*, pages 664–669, 2004.
- [17] C. Seshadhri, T. G. Kolda, and A. Pinar. Community structure and scale-free collections of Erdős-Rényi graphs. *Phys. Rev. E*, 85:056109, May 2012.
- [18] A. Shapiro and T. Homem-de Mello. On the rate of convergence of optimal solutions of monte carlo approximations of stochastic programs. *SIAM Journal on Optimization*, 11(1):70–86, 2000.
- [19] M. Trusov, R. E. Bucklin, and K. Pauwels. Effects of word-of-mouth versus traditional marketing: Findings from an internet social networking site. *Journal of Marketing*, 73, September 2009.
- [20] J. Tsai, T. H. Nguyen, and M. Tambe. Security games for controlling contagion. In *AAAI*, 2012.
- [21] D. J. Watts and S. H. Strogatz. Collective dynamics of small-world networks. *Nature*, 393(6684):440–442, June 1998.
- [22] Z. Yin and M. Tambe. A unified method for handling discrete and continuous uncertainty in bayesian stackelberg games. In *In AAMAS*, pages 855–862, 2012.