Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence

New Models for Competitive Contagion

Moez Draief
Electrical and Electronic Engineering
Imperial College
m.draief@imperial.ac.uk

Hoda Heidari
Computer and Information Science
University of Pennsylvania
hoda@seas.upenn.edu

Michael Kearns
Computer and Information Science
University of Pennsylvania
mkearns@cis.upenn.edu

Abstract

In this paper, we introduce and examine two new models for *competitive contagion* in networks, a game-theoretic generalization of the viral marketing problem. In our setting, firms compete to maximize their market share in a network of consumers whose adoption decisions are stochastically determined by the choices of their neighbors. Building on the switching-selecting framework introduced by Goyal and Kearns, we first introduce a new model in which the payoff to firms comprises not only the number of vertices who adopt their (competing) technologies, but also the network *connectivity* among those nodes. For a general class of stochastic dynamics driving the local adoption process, we derive upper bounds on (1) the (pure strategy) Price of Anarchy (PoA), which measures the inefficiency of resource use at equilibrium, and (2) the Budget Multiplier, which captures the extent to which the network amplifies the imbalances in the firms' initial budgets. These bounds depend on the firm budgets and the maximum degree of the network, but no other structural properties. In addition, we give general conditions under which the PoA and the Budget Multiplier can be unbounded. We also introduce a model in which budgeting decisions are *endogenous*, rather than externally given as is typical in the viral marketing problem. In this setting, the firms are allowed to choose the number of seeds to initially infect (at a fixed cost per seed), as well as which nodes to select as seeds. In sharp contrast to the results of Goyal and Kearns, we show that for almost any local adoption dynamics, there exists a family of graphs for which the PoA and Budget Multiplier are unbounded.

Introduction

In the traditional viral marketing problem, firms attempt to maximize the adoption of their product or service in a large social network. To this end, each of them *seeds* a set of initial "infections" in the network via product give-aways, marketing campaigns targeting the seed individuals, and so on. The product may then spread through the network via local stochastic dynamics accounting for local recommendations or influence between neighbors, known as "word of mouth" effects. Previous papers on this subject mainly focus on designing (Kempe, Kleinberg, and Tardos 2003; 2005; Mossel and Roch 2010) or improving (Chen, Wang, and Yang 2009; Chen, Yuan, and Zhang 2010; Borgs et al. 2014; Goyal, Lu, and Lakshmanan 2011) algorithms for finding a seed set that (approximately) maximizes the total number of vertices that ultimately adopt the product. More recently, a number of papers (Goyal and Kearns 2012; Bharathi, Kempe, and Salek 2007; Borodin, Filmus, and Oren 2010; Clark and Poovendran 2011; Carnes et al. 2007; Dubey, Garg, and Meyer 2006; Vetta 2002; Borodin et al. 2013; Tzoumas, Amanatidis, and Markakis 2012; Alon et al. 2010) have examined models of *competitive contagion* that take a game-theoretic perspective on the traditional problem: two or more players or firms compete in a large social network, each with the goal of maximizing their individual market share, possibly at the expense of others.

In this paper, we introduce and examine two new natural models for competitive contagion in a network. Existing models of influence maximization and competitive contagion assume that firms benefit merely according to the *number* of nodes that eventually adopt their product. We introduce a framework where the payoffs to firms capture both their market share and the *connectivity* within that market share in the network. In many natural settings, the goal is not simply that many nodes adopt a product, but also subsequently use a networked service requiring the product. For instance, Skype users are more valuable if they are in a densely connected subnetwork of other Skype users with whom they use the service to interact. We thus consider utility functions that combine both the size of the market share and the connectivity within that share. For a broad family of stochastic dynamics—concave switching function and linear selection function— we prove upper bounds on both the pure strategy PoA and the Budget Multiplier which depend on the firm budgets and the maximum degree of the network, but no other structural properties. We also find broad conditions under which the PoA and the Budget Multiplier can be unbounded.

Previous works on the subject of influence maximization (e.g. see (Goyal and Kearns 2012; Kempe, Kleinberg, and Tardos 2003)) assume that the budgets available to firms to seed initial infections in the network are fixed and exogenously determined. In many settings this might not be realistic, as firms are free to adjust their budgets in order to capture a larger market share. We therefore also consider a model in which budgeting decisions are *endogenous*. While

Model \ Measure	Price of Anarchy	Budget Multiplier
Goyal and Kearns model	$\Theta(1)$	$\Theta(1)$
The Connectivity model	$\Theta(K_{max}^2, d_{max})$	$\Theta(K_{max})$
The Endogenous model	unbounded	unbounded

Table 1: Summary of our upper bound results compared to the work of Goyal and Kearns (Goyal and Kearns 2012) for the case where the adoption dynamics exhibit decreasing returns to local adoption. Here K_{max} denotes the maximum number of "seeds" firms can spend to initialize the adoption of their product in the network, and d_{max} is the maximum degree of a vertex in the network.

the results of (Goyal and Kearns 2012) establish fairly general conditions on local dynamics yielding bounded PoA and Budget Multiplier independent of network structure, we show that such robustness vanishes in the case of endogenous budgets: for almost *any* choice of dynamics, the PoA and the Budget Multiplier may be unbounded for certain network structures. The informal intuition is that firms may engage in "bidding wars" for sub-optimal seed infections that essentially eradicate subsequent market share gains.

Table 1 shows a summary of our results compared to (Goyal and Kearns 2012).

Related Work

We contribute to the study of influence maximization in a networked setting (see (Kempe, Kleinberg, and Tardos 2003; 2005; Mossel and Roch 2010; Chen, Wang, and Yang 2009; Chen, Yuan, and Zhang 2010; Borgs et al. 2014; Goyal, Lu, and Lakshmanan 2011)), where the goal is to find a small set of influential nodes in the network whose aggregated influence is maximized. We also contribute to the study of competition in networked environments (see (Butters 1977; Grossman and Shapiro 1984)). We mainly build on the game-theoretic framework introduced in (Goyal and Kearns 2012) for studying the competitive influence maximization in a social network. In this work, the authors identify broad conditions on the adoption dynamics — namely, decreasing returns to local adoption — under which the PoA is uniformly bounded above, across all networks. They also find sufficient conditions on the adoption dynamics — namely, proportional local adoption between competitors — for bounded pure strategy Budget Multiplier. In our work we investigate similar problems in more general settings.

To the best of our knowledge, our work is the first to take the connectivity among adopters into account in the payoff function of players. There are however several previous papers that look at a similar quantity, but with different goals. For example, (Quan et al. 2012) and (Chaoji et al. 2012) investigate the relationship between connection features of individuals and the popularity of a content in a social network; and (Katona, Zubcsek, and Sarvary 2011) investigates the effect of the connectivity among current adopters on a potential consumer's behavior. Our connectivity model is also remotely related to the large body of work on clustering (see (Aggarwal and Wang 2010) for a survey).

The impact of endogenous budgets have been investigated on various economic problems, including auctions (Kotowski 2013; Burkett 2011; Ausubel, Burkett, and Filiz-Ozbay 2013) and housing markets (Pereyra 2012).

The Framework

For the underlying diffusion dynamics, we consider the switching-selection framework introduced in (Goyal and Kearns 2012). Before introducing our new models, we first review this framework.

Consider a 2-player [1] game of competitive adoption on a (possibly directed) graph G over n vertices. G is known to the players, R(ed) and B(lue). The two players simultaneously choose some number of vertices to initially seed; after this seeding, the stochastic dynamics of local adoption determines how each player's seeds spread throughout G to create adoptions by new nodes. Each player seeks to maximize her payoff which is a function of her eventual adopters.

More precisely, suppose that player $p \in \{R, B\}$ has $K_p \in \mathbb{N}$ seed infections; Each player p then chooses an allocation $a_p = (a_{p1}, a_{p2}, ..., a_{pn})$ of budget across the n vertices, where $a_{pj} \in \mathbb{N}$ and $\sum_{j=1}^{n} a_{pj} = K_p$. In (Goyal and Kearns 2012) the authors assume that K_p is exogenously given.We will see that this assumption is crucial for obtaining their upper bounds on the Price of Anarchy and the Budget Multiplier.

Let A_p be the set of allocations for player p, which is her pure strategy space. A mixed strategy for player p is a probability distribution σ_p on A_p. Let $\mathcal{A}_p$ denote the set of probability distributions for player p. The two players simultaneously choose their strategies (σ_R, σ_B). Consider any realized initial allocation (a_R, a_B) for the two players. Let $V(a_R) = \{v | a_{Rv} > 0\}$, $V(a_B) = \{v | a_{Bv} > 0\}$ and let $V(a_R, a_B) = V(a_R) \cup V(a_B)$. A vertex v becomes initially infected if one or more players assigns a seed to infect v. If both players assign seeds to the same vertex, then the probability of initial infection by a player is proportional to the seeds allocated by the player (relative to the other player).

Following the allocation of seeds, the stochastic contagion process on G determines how these Red and Blue infections generate new adoptions in the network. Time is considered to be discrete for this process, and the state of a vertex v

[1] One can also define the single-player setting by letting the opponent's strategy be the empty set. In the single-player setting the interesting problem is that of maximizing the profit. One can easily show that as is the case for other contagion models, including (Kempe, Kleinberg, and Tardos 2003), profit maximization is hard in our models. In fact it may even be more difficult, since in the connectivity model the payoff function is not submodular, and in the endogenous model it is not monotone.

at time t is denoted $s_{vt} \in \{U, R, B\}$, where U stands for Uninfected, R stands for infection by Red, and B stands for infection by Blue. We assume there is an *update schedule* which determines the order in which vertices are considered for state updates. Note that this schedule does not need to be deterministic. We also assume once a vertex is infected, it is never a candidate for updating again.

For the stochastic update of an uninfected vertex v, we will consider the *switching-selection* model. In this model, updating is determined by the application of two functions to v's local neighborhood: $f(x)$ (the *switching* function), and $g(y)$ (the *selection* function). More precisely, let α_R and α_B be the fraction of v's neighbors infected by R and B, respectively, at the time of the update, and let $\alpha = \alpha_R + \alpha_B$ be the total fraction of infected neighbors. The function f maps α to the interval $[0, 1]$ and g maps $\alpha_R/(\alpha_R + \alpha_B)$ (the relative fraction of infections that are R) to $[0, 1]$. These two functions determine the stochastic update in the following fashion:

1. With probability $f(\alpha)$, v becomes infected by *either* R or B; with probability $1 - f(\alpha)$, v remains in state U(ninfected), and the update ends.
2. If it is determined that v becomes infected, it becomes infected by R with probability $g(\alpha_R/(\alpha_R + \alpha_B))$, and infected by B with probability $g(\alpha_B/(\alpha_R + \alpha_B))$.

We assume $f(0) = 0$ (infection requires exposure), $f(1) = 1$ (full neighborhood infection forces infection), and f is increasing (more exposure yields more infection); and $g(0) = 0$ (players need some local market share to win an infection), $g(1) = 1$. Note that since the selection step above requires that an infection takes place, we also have $g(y) + g(1-y) = 1$, which implies $g(1/2) = 1/2$.

Given a graph G and an initial allocation of seeds, the dynamics described above — determined by f, g, and the update schedule — yield a number of adopters for the two players, and as mentioned earlier, the payoff to player $p \in \{R, B\}$, which we denote by $\Pi_p(\sigma_R, \sigma_B)$, is a function of her eventual adopters. Here is one possible choice for the payoff function: let the random variable χ_p denote the *number* of infections p gets at the termination of the dynamics for the strategy profile (σ_R, σ_B); in (Goyal and Kearns 2012) authors assume $\Pi_p(\sigma_R, \sigma_B) = \mathbf{E}[\chi_p|(\sigma_R, \sigma_B)]$, where the expectation is over any randomization in the player strategies, in the choice of initial allocations, and the randomization in the stochastic updating dynamics. Shortly in the connectivity and the endogenous budgets model we will see more general choices for the payoff function.

Given any payoff function, each player seeks to maximize her own expected payoff, and this results in competition among the players. In the resulting game a *Pure Nash Equilibrium* is a profile of pure strategies (a_R, a_B) such that a_p maximizes player p's payoff given the strategy a_{-p} of the other player.

A *Mixed Nash Equilibrium* is a pair $\sigma = (\sigma_R, \sigma_B)$ of independent probability distributions that satisfies

$$\mathbf{E}_{a\sim\sigma}[\Pi_p(a)] \geq \mathbf{E}_{a_{-p}\sim\sigma_{-p}}[\Pi_p(a'_p, a_{-p})]$$

for every p and $a'_p \in A_p$. In the above $a = (a_p, a_{-p})$.

For a fixed graph G and stochastic update dynamics, the *maximum social welfare* allocation is the (deterministic) allocation (a^*_R, a^*_B) (obeying the budget constraints if any exists) that maximizes $\Pi_R(a_R, a_B) + \Pi_B(a_R, a_B)$. For the same fixed graph and update dynamics, let (σ_R, σ_B) be the equilibrium strategies that *minimize* $\Pi_R(\sigma_R, \sigma_B) + \Pi_B(\sigma_R, \sigma_B)$ among all equilibria[2]. Then the *Price of Anarchy* (or PoA) is defined to be

$$\frac{\Pi_R(a^*_R, a^*_B) + \Pi_B(a^*_R, a^*_B)}{\Pi_R(\sigma_R, \sigma_B) + \Pi_B(\sigma_R, \sigma_B)}.$$

The Price of Anarchy is a measure of the inefficiency in resource use created due to decentralized, non-cooperative behavior by the two players.

We also study the *Budget Multiplier*, which measures the extent to which network structure and dynamics can amplify initial *resource* inequality across the players. For example, when we have external budget constraints K_R, K_B, with $K_R \geq K_B$, we define the Budget Multiplier as follows: for any fixed graph G and stochastic update dynamics, let (σ_R, σ_B) be the equilibrium that *maximizes* the ratio

$$\frac{\Pi_R(\sigma_R, \sigma_B)}{\Pi_B(\sigma_R, \sigma_B)} \times \frac{K_B}{K_R}$$

among all equilibria[3]. The resulting maximized ratio is the Budget Multiplier, and it measures the extent to which the larger budget player can obtain a final market share that exceeds her share of the initial budgets[4].

Finally, we will restate some of the results in (Goyal and Kearns 2012) which we will make use of later in the paper.

Lemma 1 *Let a_R and a_B be any sets of seed vertices for the two players. Then if f is concave and g is linear,*

$$\mathbf{E}[\chi_R|(a_R, \emptyset)] \geq \mathbf{E}[\chi_R|(a_R, a_B)]$$

and

$$\mathbf{E}[\chi_B|(\emptyset, a_B)] \geq \mathbf{E}[\chi_B|(a_R, a_B)].$$

Lemma 2 *Let a_R and a_B be any sets of seeded nodes for the two players. If f is increasing,*

$$\mathbf{E}[\chi_R + \chi_B|(a_R, a_B)] \geq \mathbf{E}[\chi_R|(a_R, \emptyset)].$$

The Connectivity Model

In some cases, such as video conferencing and messaging applications like WhatsApp, ooVoo or Skype, not only the number of adopters matters; it is also important to the firms how well those adopters are connected to each other, as the

[2]If we restrict attention to pure Nash equilibria only, the resulting ratio is called the *pure strategy* Price of Anarchy.

[3]If we restrict attention to pure Nash equilibria only, the resulting ratio is called the *pure strategy* Budget Multiplier.

[4]We will later introduce the endogenous budgets model in which budgets are not externally constrained. In that model we can either use the same definition for the Budget Multiplier, or use a similar one in which we replace the number of seeds players spend with the cost per seed for each.

use of these products takes place along the edges of the social network rather than at a vertex alone. Motivated by the above examples, we introduce the *connectivity* model in which a firm's goal is to maximize the number of her adopters *plus* the number of edges among those adopters; the second term here models the strength of the connectivity among the adopters.

More precisely, consider a *pure* strategy profile (a_R, a_B) where a_p denotes the strategy of player $p \in \{R, B\}$. We define the random variable γ_p to be the eventual number of edges among adopters of product p. Player p then seeks to maximize her payoff which is equal to

$$\mathbf{E}[\gamma_p + \chi_p | (a_R, a_B)].$$

Note that due to linearity of expectation $\mathbf{E}[\gamma_p + \chi_p | (a_R, a_B)] = \mathbf{E}[\gamma_p | (a_R, a_B)] + \mathbf{E}[\chi_p | (a_R, a_B)]$. To simplify the statement of our results we denote $\mathbf{E}[\gamma_p | (a_R, a_B)]$ by δ_p and $\mathbf{E}[\chi_p | (a_R, a_B)]$ by θ_p. In addition, we denote the strategy of player p in the maximum social welfare solution by a_p^* and her payoff by $\theta_p^* + \delta_p^*$. Also let $\theta = \theta_R + \theta_B$, $\delta = \delta_R + \delta_B$, $\theta^* = \theta_R^* + \theta_B^*$, and $\delta^* = \delta_R^* + \delta_B^*$.

We will see that in this new model, when f is concave and g is linear, upper bounds on the pure PoA and Budget Multiplier still exist, but they can depend on the budget constraints and the structure of the network. In addition we will see that if f is convex and g is linear, then the PoA and Budget Multiplier can be unbounded. Note that the results and techniques presented in this section can be easily extended to the case where $\Pi_p = \mathcal{B}\delta_p + \mathcal{D}\theta_p$ with $\mathcal{B}, \mathcal{D}$ being positive constants. Detailed discussion is omitted due to space constraints.

We first illustrate the connectivity model with an example for which the equilibrium looks quite different compared to the original model of (Goyal and Kearns 2012).

Example 1 *Consider a graph G consisting of 3 components C_1, C_2 and C_3, where C_1, C_2 are star networks of size $(N+1)$ with central nodes v_1, v_2 pointing to N peripheral vertices, and C_3 is a complete undirected network of size $3\sqrt{N}$. Suppose that both players have a budget equal to 1. Let $f(x) = x^\alpha$ for some $\alpha > 0$, and let g be linear. Suppose that the update schedule is relatively long so that any connected component containing a non-zero number of seeds eventually becomes entirely infected.*

It can be shown that in the original model, the equilibrium of the game on G is where Red and Blue put their seeds on v_1, v_2 and each get an expected number of infections equal to $(N+1)$. However, in the connectivity model, the equilibrium is where both players put their seeds on C_3 and each get an expected payoff approximately equal to $3(3N/2 + \sqrt{N})/2$. Thus the connectivity objective causes the players to prefer to compete for the highly connected vertices, as opposed to each taking an isolated low-connectivity component.

In what follows, we will use the following notations: K_{max} denotes $\max\{K_R, K_B\}$, K_{min} denotes $\min\{K_R, K_B\}$, and d_{max} is the maximum degree of any vertex in the graph G.

Theorem 1 *If f is concave and g is linear, then in the connectivity model, the pure strategy Budget Multiplier is bounded from above by $8(K_{max}+1)$.*

Proof: Without loss of generality suppose $K_R \geq K_B$. Let (S_R, S_B) denote the pure Nash equilibrium that maximizes the ratio $\frac{\Pi_R}{\Pi_B} \times \frac{K_B}{K_R}$. We assume that in this equilibrium, Red has the higher payoff, that is $\theta_R + \delta_R \geq \theta_B + \delta_B$ (if this is not the case the budget multiplier is bounded from above by 1). We will assume that $\delta_R \geq \theta_R$. Since if $\delta_R < \theta_R$, from Theorem 4 in (Goyal and Kearns 2012), we know that in the equilibrium the Blue player gets at least $\frac{K_B}{2K_R}\theta_R$ infections (the Blue player can achieve this by imitating the K_B Red seeds that result in the highest *number* of infections). Therefore, we have

$$\Pi_B \geq \theta_B \geq \frac{K_B}{2K_R}\theta_R = \frac{K_B}{4K_R}(\theta_R + \theta_R) \geq \frac{K_B}{4K_R}(\theta_R + \delta_R)$$

and so we can conclude that the Budget Multiplier is at most 4.

Given the above assumptions, next we assign a unique label to each of the vertices in S_R and attribute subsequent Red infections to exactly one of these seeds. More precisely, let the solo Red process be the stochastic dynamical process on the network when only Red seeds S_R are present. Suppose $S_R = \{v_1, \ldots, v_{K_R}\}$. At time 0, label each vertex $v_i \in S_R$ by a different color C_i; also label all other vertices by U. In the subsequent steps, whenever a new vertex gets infected in the process we assign to it one of the K_R labels $\{C_i\}_{i=1}^{K_R}$ in the following manner: when updating a vertex v, we first compute the fraction α_v^R of neighbors whose current label is one of $C_1, \ldots, C_{K_R}$, and with probability $f(\alpha_v^R)$ we decide that an infection will occur (otherwise the label of v is updated to U). If an infection occurs, we simply choose an infected neighbor of v uniformly at random, and update v to have the same label (which will be one of the C_i's). It can easily be observed that at every step, the dynamics of the $(S_R, \emptyset)$ process are faithfully implemented if we drop label subscripts and simply view any label C_i as a generic Red infection R. Furthermore, at all times every infected vertex has only one of the labels C_i. Thus if we denote the expected number of edges with endpoints labeled C_i, C_j by $\Omega_{i,j}^R$, we have $\mathbf{E}[\gamma_R | (S_R, \emptyset)] = \sum_{i \leq j} \Omega_{i,j}^R$.

Next we claim that the blue player can choose K_B of the Red seeds as her strategy such that in expectation she gets at least $\frac{K_B}{4K_R(K_R+1)}$ share of the Red edges, which as we will see, results in the desired bound on the Budget Multiplier.

To prove the above claim, observe that the Blue player can consider the $\frac{K_B}{2}$ color pairs with highest Ω^R values. If pair (i, j) is in that set, the Blue player adds both i and j to her seed set. Since there are $K_R(K_R+1)/2$ color pairs in total, the expected number of edges Blue gets by choosing this strategy, is at least $\frac{1}{4}\frac{K_B/2}{K_R(K_R+1)/2}\delta_R$ (the factor $\frac{1}{4}$ is present due to the fact that when Blue seeds v_i, v_j, the two vertices both become initially infected by Blue with probability $1/4$). Therefore we have

$$\frac{\theta_R + \delta_R}{\theta_B + \delta_B} \leq \frac{2\delta_R}{\frac{K_B}{4K_R(K_R+1)}\delta_R} \leq 8(K_R+1)\frac{K_R}{K_B}$$

Next, we simply multiply the left and the the right hand side of the above inequality by $\frac{K_B}{K_R}$ and replace K_R with K_{max} to obtain the claimed bound on the Budget Multiplier. This finishes the proof. ∎

One can easily find families of examples in which the Budget Multiplier does actually grow with K_{max}.

Example 2 *Suppose $K_R = K > 2$ and $K_B = 2$. Consider the network G which is a complete graph of size K. One can easily see that in any equilibrium Red gets $\Theta(K)$ vertices while Blue gets $\Theta(1)$ vertices. Since all the Red vertices are connected to each other, the number of edges among them is $\Theta(K^2)$. Similarly the number of edges among Blue adopters is $\Theta(1)$. Therefore the Budget Multiplier is $\Theta(K)$.*

Theorem 2 *If f is concave and g is linear, then in the connectivity model, the pure strategy Price of Anarchy is bounded from above by $2+2(1+d_{max})(1+8(K_{max}+1)\frac{K_{max}}{K_{min}})$.*

Proof: Let (S_R, S_B) denote the pure Nash equilibrium and (a^*_R, a^*_B) denote the maximum social welfare solution. Without loss of generality, suppose the Red player gets the higher payoff in (a^*_R, a^*_B). Since (S_R, S_B) is a Nash equilibrium, the deviation of Red player to a^*_R should not be profitable, i.e. $\theta_R + \delta_R$ must be larger than the payoff Red gets by deviating to a^*_R. Let's denote that payoff by $(\theta'_R + \delta'_R)$. Next we find a lower bound on $(\theta'_R + \delta'_R)$. We claim the following holds:

$$\theta'_R + \delta'_R \geq (\theta^*_R - \theta) + (\delta^*_R - \theta d_{max}) \tag{1}$$

To prove the above, we first note that when Blue is not present, the number of infection R gets by switching to a^*_R is at least θ^*_R (Lemma 1). Also when Red is not present, the number of infection B gets by seeding S_B is at most θ (Lemma 2). Now by adding S_B to a^*_R, the total number of infections remains at least θ^*_R. In addition the number of Blue infections will decrease and become less than or equal to θ. So the number of Red infections will be at least $(\theta^*_R - \theta)$, i.e.

$$\theta'_R \geq (\theta^*_R - \theta) \tag{2}$$

Next, we observe that when Blue is not present, the number of edges Red gets using strategy a^*_R is at least δ^*_R. To see this, just note that by the departure of a^*_B from (a^*_R, a^*_B) all the vertices become just more likely to adopt Red. So the Red edges remain Red.

Now if the Blue player comes back with strategy S_B, this can result in at most θ Blue infections, and therefore number of Red edges decreases by at most θd_{max}, as each Blue vertex can take at most d_{max} edges away from Red, meaning that

$$\delta'_R \geq (\delta^*_R - \theta d_{max}). \tag{3}$$

Combining (2) and (3), we get

$$\theta'_R + \delta'_R \geq (\theta^*_R - \theta) + (\delta^*_R - \theta d_{max})$$

Combining the above with $\theta_R + \delta_R \geq \theta'_R + \delta'_R$ (which must hold because of the property of Nash equilibrium), we obtain the following:

$$\theta_R + \delta_R \geq (\theta^*_R - \theta) + (\delta^*_R - \theta d_{max}) \tag{4}$$

Now to prove the desired bound on PoA we can write

$$\begin{aligned}
& \theta_R + \delta_R \geq (\theta^*_R - \theta) + (\delta^*_R - \theta d_{max}) \\
\Rightarrow \quad & 1 + (1 + d_{max})\frac{\theta}{\theta_R + \delta_R} \geq \frac{\theta^*_R + \delta^*_R}{\theta_R + \delta_R} \\
\Rightarrow \quad & 1 + (1 + d_{max})\frac{\theta + \delta}{\theta_R + \delta_R} \geq \frac{\theta^*_R + \delta^*_R}{\theta + \delta} \\
\Rightarrow \quad & 2 + 2(1 + d_{max})\frac{\theta + \delta}{\theta_R + \delta_R} \geq \frac{\theta^* + \delta^*}{\theta + \delta}.
\end{aligned} \tag{5}$$

In the third line we used the fact that $\delta \geq 0$ and $\theta_B + \delta_B \geq 0$, and in the fourth line we used the fact that in the maximum social welfare solution the Red player gets the higher payoff, or equivalently $2(\theta^*_R + \delta^*_R) \geq \theta^* + \delta^*$. Next we note that

$$\frac{\theta + \delta}{\theta_R + \delta_R} \leq \begin{cases} 1 + 1 & \textit{if } \theta_R + \delta_R \geq \theta_B + \delta_B \\ 1 + 8 & \textit{if } K_R \geq K_B \\ 1 + \mathcal{M}\frac{K_B}{K_R} & \textit{otherwise.} \end{cases}$$

where $\mathcal{M}$ is the Budget Multiplier[5]. Using the bound we obtained in Theorem 1 for Budget Multiplier, we get

$$\frac{\theta + \delta}{\theta_R + \delta_R} \leq 1 + 8(K_B + 1)\frac{K_B}{K_R}.$$

Combining the above with (5) and replacing K_B with K_{max} and K_R with K_{min}, we get

$$2 + 2(1 + d_{max})(1 + 8(K_{max} + 1)\tfrac{K_{max}}{K_{min}}) \geq \tfrac{\theta^* + \delta^*}{\theta + \delta}$$

and that finishes the proof. ∎

Finally, using constructions similar to the ones in (Goyal and Kearns 2012), we show that the concavity of f and the linearity of g are necessary for obtaining the upper bounds.

Proposition 1 *If f is linear and $g(y) = y^s/(y^s + ((1-y)^s)$ for some $s > 1$, then in the connectivity model, for any $V > 0$, there exists a graph G for which the Budget Multiplier is greater than V.*

Proof:(sketch) The idea, closely following Theorem 5 in (Goyal and Kearns 2012), is to construct a layered graph that amplifies the punishment in the selection function, and as a result makes the Budget Multiplier arbitrarily large: The top layer of this amplifying graph is where in the equilibrium solution players put their seeds; as we go down the layers the share of vertices that adopt the product of the higher budget player, becomes larger and larger; and therefore in the last layer, which is a huge one that makes up for the majority of the payoff each player receives, the larger budget player receives a payoff much higher than what the opponent receives. As a result the Budget Multiplier becomes very large. By adjusting the parameters of this construction, we can make the Budget Multiplier arbitrarily large and this proves the theorem. ∎

Proposition 2 *If $f(x) = x^r$ for some $r > 1$, and g is linear, then in the connectivity model, for any $V > 0$, there exists a graph G for which the Price of Anarchy is greater than V.*

[5]The discussion for the second item where $K_R \geq K_B$ is very similar to the proof of Theorem 1, and due to space constraints we do not repeat it here.

Proof:(sketch) The idea, closely following Theorem 2 in (Goyal and Kearns 2012), is to create a layered directed graph whose dynamics rapidly amplify the convexity of f. When we take two such amplification components of differing sizes, one equilibrium is the case where players coordinate on the smaller component, while the maximum social welfare solution lies in the larger component. As a result we have an example in which the PoA is large. By adjusting the parameters of this construction, we can make the PoA arbitrarily large and this proves the theorem. ∎

The Endogenous Budgets Model

As we have mentioned, previous works assume that there are external constraints on the maximum number of seeds players can spend to initialize the adoption of their product in the network. We argue that this assumption is not necessarily realistic, since in some settings (such as the case of product give-aways), if firms feel they can offset higher marketing expenditures by winning greater market share, they are free to do so. This motivates us to relax this assumption and investigate the case in which firms are allowed to choose the number of seeds they want to allocate, given the constraint that each seed has a cost associated with it.

More precisely, we propose the *endogenous budgets model*, which is the following modification of the original framework: Each firm $p \in \{R, B\}$ has a cost per seed denoted by $c_p \geq 0$, and for each *new* (non-seed) vertex that adopts her product, firm p benefits b_p[6]. So if θ_p denotes the (expected) eventual number of non-seed infections that firm p obtains by initially spending K_p seeds, her payoff is equal to $b_p \times \theta_p - c_p \times K_p$. Without loss of generality we can assume $b_p = 1$, and write the payoff as:

$$\theta_p - c_p \times K_p$$

In this section we show that in the endogenous budgets model, for a broad setting of f, g and c, there are examples in which the PoA is unbounded[7].

Beside the PoA we also look at the quantity $\frac{\theta^*}{\theta}$ (recall that this quantity represented the PoA in the original model of (Goyal and Kearns 2012)).We will see that similar to the PoA in our model, this quantity is also unbounded, showing that the unbounded PoA in the endogenous model is not merely due to the introduction of costs (i.e. the term $-c_p \times K_p$) to the payoff function.

We first illustrate the endogenous budgets model with an example where the equilibrium can look completely different compared to the original model of (Goyal and Kearns 2012).

Example 3 *Consider a graph G consisting of a central node v pointing to N vertices $v_1, v_2, ..., v_N$, each of which points to 3 different (unimportant) vertices. Suppose $c_p = K_p = 1$ for $p \in \{R, B\}$. Let $f(x) = x^\alpha$ for some $\alpha > 0$, and let g be linear. Suppose that the update schedule is relatively long so that a connected component containing a nonzero number of seeds eventually becomes entirely infected.*

It can be shown that in the original model of (Goyal and Kearns 2012), the equilibrium is where both players put their only seed on v and each get an expected payoff equal to $(4N + 1)/2$, while in the endogenous budgets model, the equilibrium is the case where there is exactly one Red and one Blue seed on each of $v_1, v_2, ..., v_N$, resulting in a payoff equal to $N/2$ for each player.

[6]We assume seed vertices do not provide any payoff, only cost.

[7]Similar results hold for the Budget Multiplier. The construction—which is basically a properly tuned star-like network—is straightforward, and therefore omitted due to space constraints.

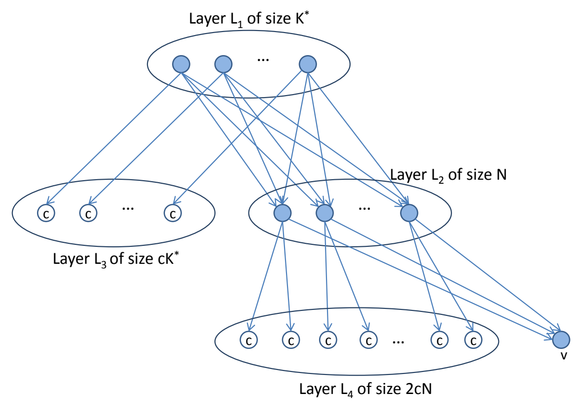

Figure 1: A graph with PoA $\geq (2c + 1)N$.

Theorem 3 *Suppose $c_p = c$ for all $p \in \{R, B\}$ and $c \in \mathbb{N}$. Then regardless of f, g, in the endogenous budgets model the PoA can be unbounded.*

Proof: We first prove the theorem for the case of $c = 1$. Consider the graph H represented in Figure 1. We claim that in this graph the maximum social welfare solution is the case where there is a single seed on each node in layer L_1. Also we claim that the equilibrium solution is the case where players both spend exactly one seed on each node of layer L_2. If this holds then it is easy to see that the PoA is equal to $\frac{3N+1+K^*-K^*}{2N+1-2N} = \frac{3N+1}{1} = 3N + 1$; And since N can be arbitrarily large, we can conclude that the PoA is unbounded. Also we have $\frac{\theta^*}{\theta} = \frac{3N+1+K^*}{2N+1}$, and since K^* can be as large as we want, $\frac{\theta^*}{\theta}$ is unbounded as well.

It only remains to prove the above claims. First note that no node in layers L_3, L_4 is ever selected as a seed in the maximum social welfare solution, because selecting those nodes does not result is any new infections (as all the edges are directed from top to bottom). Therefore since layer L_3 is empty of seeds, putting one seed on a node of layer L_1 always pays off immediately as it has a neighbor in L_3 which becomes subsequently infected with probability 1 ($f(1) = 1$). This means that if a node in L_1 is not already seeded in the maximum social welfare solution, one can seed it without decreasing the payoff. So we can assume that in the maximum social welfare solution all the nodes in layer L_1 are seeded (it is easy to see that one seed per node suffices). Now note that once all the vertices in L_1 are seeded, the rest of the network will become infected with probability 1 ($f(1) = 1$). So we can conclude that this case is indeed the maximum social welfare solution.

Now we prove our claim about the equilibrium solution. Suppose the Red player has a seed on every node of layer L_2; we compute the best response of the Blue player to this strategy. Note that adding a seed to layers L_1, L_3, L_4 does not increase Blue's payoff, so her only choice is to allocate seeds to vertices in layer L_2. Suppose Blue has $k < N$ seeds on L_2 each on a different node. This means that the expected number of Blue seeds in L_2 is equal to $k/2$. Now note that for each new seed that Blue adds to a vertex u in L_2:

1. If Blue does not have any other seed on u, by adding one, she increases her payoff by a positive amount: with probability $\frac{1}{2}$, u becomes a Blue seed. If that happens, the 2 neighbors of u in layer L_4 certainly adapt Blue; also the probability that v adopts Blue increases (as g is increasing). Since the cost per seed is equal to 1 and the change in the expected number of Blue infections is larger than $\frac{1}{2} \times 2 = 1$, this action increases Blue's payoff.
2. If Blue does already have one seed on u, by adding another, at best she changes her payoff by $\frac{1}{6}(2+1) - 1 \leq 0$ (the change in the probability of u becoming initially Blue as a result of this new seed is $\frac{1}{6}$. If this happens, in the best case scenario, both v and the two neighbors of u in L_4 become subsequently infected). Since the change is negative this action decreases Blue's payoff.

The above shows that the best response of the Blue player is to spend exactly one seed on each vertex of L_2. Similarly, the best response of the Red player to Blue player's strategy is to put exactly one seed on every vertex of L_2, showing that this allocation is indeed an equilibrium. This finishes the proof for the case of $c = 1$.

It is now easy to see that the above example can be easily generalized to the case where the cost per seed is a positive integer c: It suffices to replace every node in layers L_3, L_4 with c vertices with the same neighbor set in L_1, L_2 as the original vertex. With a similar type of reasoning one can show that PoA is equal to $\frac{(2c+1)N+1+K^*c-K^*c}{2cN+1-2cN} = \frac{(2c+1)N+1}{1} = (2c+1)N+1$ which can be arbitrarily large by respectively choosing N sufficiently large. ∎

Corollary 1 *Suppose $c_p = c$ for all $p \in \{R, B\}$ and $c \geq 1$ ($c \in \mathbb{N}$). Then regardless of f, g, the quantity $\frac{\theta^*}{\theta}$ can be unbounded in the endogenous budgets model.*

Proof: The same construction in the previous proof works here too. We have that $\frac{\theta^*}{\theta} = \frac{(2c+1)N+1+K^*c}{2cN+1}$, so it can be arbitrarily large by choosing K^* big enough.

Finally we investigate the case where $c < 1$.

Theorem 4 *Suppose $c_p = c < 1$ for $p \in \{R, B\}$ and g is linear. Then regardless of f, the PoA can be unbounded in the endogenous budgets model.*

Proof: Consider the graph H in Figure 2. We choose K sufficiently large. The maximum social welfare solution is obviously the case where a single seed is put on v, resulting in payoff $K + [2cK] + 1 - c$ which is larger than K.

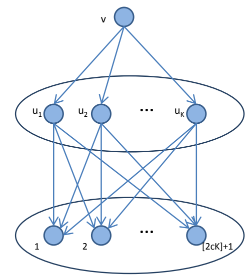

Figure 2: A graph with PoA $\geq K$.

Next we claim that one equilibrium is the case where players put one seed on each of $u_1, ..., u_K$ resulting in total payoff of $[2cK]+1-2cK < 1$, and therefore a PoA larger than K. To prove our claim, we just need to show that if the strategy of one player (say Red) is to put exactly one seed on each vertex of the second layer, then the best response of the Blue player will be to do the same thing. Suppose Blue chooses x vertices among $u_1, ..., u_K$ as her seed set. Then her expected payoff would be at most $(x/2K)([2cK]+1) - cx$ which is increasing in x (one can easily see this by taking the derivative: $([2cK]+1)/2K - c \geq 0$). This means that the best response is to choose $x = K$ which proves our claim.

Finally, we note that since K in the above construction can be chosen arbitrarily large, the PoA is unbounded and this finishes the proof. ∎

Concluding Remarks

We relaxed two of the main restricting assumption in previous papers on competitive influence maximization. We saw that exogenous budget constraints are crucial for obtaining upper bounds on the PoA and the Budget Multiplier. We also saw that if in addition to the number of adopters, firms take the connectivity among those adopters into account, upper bounds on the PoA and the Budget Multiplier still exist, but they can depend weakly on the network structure and the budget constraints. Our work suggests a number of additional interesting open problems, including the following:

1. In the endogenous model, we assumed that the cost per seed is a fixed value. It would be interesting to investigate the case where the cost is in fact a more complex function of the seed set.
2. In the connectivity model, we considered a simple notion of connectivity (i.e. the number of edges). Another interesting direction is to investigate more complex notions of connectivity.

Acknowledgements

We give warm thanks to Sanjeev Goyal for extensive discussions on the models and results presented here. This research was partially supported by MURI award W911NF-12-1-0509 of the United States Army Research Office (ARO).

References

Aggarwal, C., and Wang, H. 2010. A survey of clustering algorithms for graph data. *Advances in Database Systems* 40.

Alon, N.; Feldman, M.; Procaccia, A.; and Tennenholtz, M. 2010. A note on competitive diffusion through social networks. *Information Processing Letters* 110(6).

Ausubel, L.; Burkett, J.; and Filiz-Ozbay, E. 2013. An experiment on auctions with endogenous budget constraints. *Manuscript*.

Bharathi, S.; Kempe, D.; and Salek, M. 2007. Competitive influence maximization in social networks. In *Proceedings of the Conference on Web and Internet Economics*.

Borgs, C.; Brautbar, M.; Chayes, J.; and Lucier, B. 2014. Maximizing social influence in nearly optimal time. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*.

Borodin, A.; Braverman, M.; Lucier, B.; ; and Oren, J. 2013. Strategyproof mechanisms for competitive influence in networks. In *Proceedings of the International World Wide Web Conference*.

Borodin, A.; Filmus, Y.; and Oren, J. 2010. Threshold models for competitive influence in social networks. In *Proceedings of the Conference on Web and Internet Economics*.

Burkett, J. 2011. Endogenous budget constraints in auctions. *Manuscript*.

Butters, G. 1977. Equilibrium distribution of prices and advertising. *Review of Economic Studies* 44(3).

Carnes, T.; Nagarajan, C.; Wild, S.; and van Zuylen, A. 2007. Maximizing influence in a competitive social network: A follower's perspective. In *Proceedings of the ACM Conference on Economics and Computation*.

Chaoji, V.; Ranu, S.; Rastogi, R.; and Bhatt, R. 2012. Recommendations to boost content spread in social networks. In *Proceedings of the International World Wide Web Conference*.

Chen, W.; Wang, Y.; and Yang, S. 2009. Efficient influence maximization in social networks. In *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.

Chen, W.; Yuan, Y.; and Zhang, L. 2010. Scalable influence maximization in social networks under the linear threshold model. In *Proceedings of the IEEE International Conference on Data Mining*.

Clark, A., and Poovendran, R. 2011. Maximizing influence in competitive environments: A game-theoretic approach. In *Decision and Game Theory for Security*, volume 7037 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg.

Dubey, P.; Garg, R.; and Meyer, B. D. 2006. Competing for customers in a social network: The quasi-linear case. In *Proceedings of the Conference on Web and Internet Economics*.

Goyal, S., and Kearns, M. 2012. Competitive contagion in networks. In *Proceedings of the ACM Symposium on Theory of Computing*.

Goyal, A.; Lu, W.; and Lakshmanan, L. V. 2011. Simpath: An efficient algorithm for influence maximization under the linear threshold model. In *Proceedings of the IEEE International Conference on Data Mining*.

Grossman, G., and Shapiro, C. 1984. Informative advertising with differentiated products. *Review of Economic Studies* 51(1).

Katona, Z.; Zubcsek, P. P.; and Sarvary, M. 2011. Network effects and personal influences: The diffusion of an online social network. *Journal of Marketing Research* 48(3).

Kempe, D.; Kleinberg, J.; and Tardos, E. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the ACM SIGKDD Conference onKnowledge Discovery and Data Mining*.

Kempe, D.; Kleinberg, J.; and Tardos, E. 2005. Influential nodes in a diffusion model for social networks. In *Proceedings of the International Colloquium on Automata, Languages, and Programming*.

Kotowski, M. 2013. First-price auctions with budget constraints. *Manuscript*.

Mossel, E., and Roch, S. 2010. Submodularity of influence in social networks: From local to global. *SIAM Journal of Computing* 39(6).

Pereyra, J. S. 2012. Housing markets with endogenous budget constraints. *Manuscript*.

Quan, H.; Milicic, A.; Vucetic, S.; and Wu, J. 2012. A connectivity-based popularity prediction approach for social networks. In *Proceedings of the IEEE International Conference on Communications*.

Tzoumas, V.; Amanatidis, C.; and Markakis, E. 2012. A game-theoretic analysis of a competitive diffusion process over social networks. In *Proceedings of the Conference on Web and Internet Economics*.

Vetta, A. 2002. Nash equilibria in competitive societies with applications to facility location, traffic routing and auctions. In *Proceedings of the IEEE Symposium on Foundations of Computer Science*.