

Competitive Contagion in Networks

Sanjeev Goyal, Hoda Heidari, Michael Kearns

Abstract

We develop a game-theoretic framework for the study of competition between firms who have budgets to “seed” the initial adoption of their products by consumers located in a social network. We identify a general property of the adoption dynamics — namely, decreasing returns to local adoption — for which the inefficiency of resource use at equilibrium (the *Price of Anarchy*) is uniformly bounded above, across all networks. We also show that if this property is violated the Price of Anarchy can be unbounded, thus yielding sharp threshold behavior for a broad class of dynamics.

We provide similar results a new notion, the *Budget Multiplier*, that measures the extent that imbalances in player budgets can be amplified at equilibrium.

Keywords:

1. Introduction

The role of social networks in shaping individual choices has been brought out in a number of studies over the years.¹ In the past, the deliberate use of such social influences by external agents was hampered by the lack of good data on social networks. In recent years, data from on-line social networking sites along with other advances in information technology have created interest in ways that firms and governments can use social networks to further their goals.²

In this work, we study competition between firms who use their resources to maximize product adoption by consumers located in a social network.³ The social network may transmit information about products, and adoption of products by neighbors may have direct consumption benefits. The firms, denoted *Red* and *Blue*, know the graph which defines the social network and offer similar or interchangeable products or services. The two firms simultaneously choose to allocate their resources on subsets of consumers, i.e.,

Email addresses: sg472@cam.ac.uk (Sanjeev Goyal), hoda@seas.upenn.edu (Hoda Heidari), mkearns@cis.upenn.edu (Michael Kearns)

¹See e.g., Coleman [1] on doctors’ prescription of new drugs, Conley and Udry [2] and Foster and Rosenzweig [3] on farmers’ decisions on crop and input choice, and Feick and Price [4], Reingen et al. [5], and Godes and Mayzlin [6] on brand choice by consumers.

²The popularity of terms such as *word of mouth marketing*, *viral marketing*, *seeding the network* and *peer-leading intervention* is an indicator of this interest.

³Our model may apply to other settings of competitive contagion, such as between two fatal viruses in a population.

to *seed* the network with initial adoptions. The stochastic dynamics of local adoption determine how the influence of each player’s seeds spreads throughout the graph to create new adoptions. Our work thus builds upon recent interest in models of competitive contagion [7, 8, 9].

A distinctive feature of our framework is that we allow for a broad class of local influence processes. We decompose the dynamics into two parts: a *switching function* f , which specifies the probability of a consumer switching from non-adoption to adoption as a function of the fraction of his neighbors who have adopted *either* of the two products Red and Blue; and a *selection function* g , which specifies, conditional on switching, the probability that the consumer adopts (say) Red as function of the fraction of adopting neighbors who have adopted Red (the decomposition into switching and selection has some precursors in the literature, see [8] for example). Each firm seeks to maximize the total number of consumers who adopt its product. Broadly speaking, the switching function captures “stickiness” of the (interchangeable) products based on their local prevalence, and the selection function captures preference for firms based on local market share.

This framework yields a rich class of competitive strategies, which depend in subtle ways on the dynamics, the relative budgets of the players, and the structure of the social network. Section 5 illustrates this point: We will see that even in a simple fixed network, the equilibrium solution depends greatly on the parameters of the model. Also we will see how these parameters affect other properties of the equilibrium, such as the Price of Anarchy and the Budget Multiplier.

We then continue with understanding two general features of equilibrium: first, the efficiency of resource use by the players (Price of Anarchy, and Price of Stability) and second, the role of the network and dynamics in amplifying ex-ante resource differences between the players (Budget Multiplier).

Our *first* set of results concern efficiency of resource use by the players. For a fixed graph and fixed local dynamics (given by f and g), and budgets of K_R and K_B seed infections for the players, let (S_R, S_B) be the sets of seed infections that maximize the joint expected infections (payoffs) $\Pi_R(S_R, S_B) + \Pi_B(S_R, S_B)$ subject to $|S_R| = K_R$ and $|S_B| = K_B$, and let σ'_R and σ'_B be Nash equilibrium strategies obeying the budget constraints that *minimize* the joint payoff $\Pi_R(\sigma_R, \sigma_B) + \Pi_B(\sigma_R, \sigma_B)$ across all Nash equilibria (σ_R, σ_B) . The *Price of Anarchy* (or PoA)⁴ is then defined as:

$$\frac{\Pi_R(S_R, S_B) + \Pi_B(S_R, S_B)}{\Pi_R(\sigma'_R, \sigma'_B) + \Pi_B(\sigma'_R, \sigma'_B)}.$$

Similarly, if σ''_R and σ''_B are Nash equilibrium strategies obeying the budget constraints that *maximize* the joint payoff $\Pi_R(\sigma_R, \sigma_B) + \Pi_B(\sigma_R, \sigma_B)$ across all Nash equilibria (σ_R, σ_B) , the *Price of Stability* (or PoS)⁵ is then defined as:

$$\frac{\Pi_R(S_R, S_B) + \Pi_B(S_R, S_B)}{\Pi_R(\sigma''_R, \sigma''_B) + \Pi_B(\sigma''_R, \sigma''_B)}.$$

⁴The PoA is a measure of the maximum potential inefficiency created by non-cooperative/decentralized activity. In our context, if we suppose that consumers get positive utility from consumption of firms’ products then the PoA also reflects losses in consumer welfare.

⁵The Price of Stability (PoS) compares outcomes in the ‘best’ equilibrium with socially optimal outcomes; one may interpret the PoS as a measure of the minimum inefficiency created by non-cooperative, as opposed to merely decentralized, activity.

Our first main result, Theorem 1, shows that if the switching function f is concave and the selection function g is linear, then the PoA (and as a result, the PoS) is uniformly bounded above by 4, across all networks. The main proof technique we employ is the construction of certain coupled stochastic dynamical processes that allow us to demonstrate that, under the assumptions on f and g , the departure of one player can only benefit the other player, even though the total number of joint infections can only decrease. This in turn lets us argue that players can successfully defect to the maximum social welfare solution and realize a significant fraction of its payoff, thus implying they must also do so at equilibrium.

Our next main result, Theorem 2, shows that even a small amount of convexity in the switching function f can lead to arbitrarily high PoS (and as a result PoA). This result is obtained by constructing a family of layered networks whose dynamics effectively compose f many times, thus amplifying its convexity. Equilibrium and large PoS are enforced by the fact that despite this amplification, the players are better off playing near each other: this means that if one player locates in one part of the network, the other player has an incentive to locate close by, even if they would jointly be better off locating in a different part of the network. Taken together, our upper and lower bounds yield sharp threshold behavior in a parametric classes of dynamics. For example, if the switching function is $f(x) = x^r$ for $r > 0$ and the selection function g is linear, then for all $r \leq 1$ the PoA is at most 4, while for any $r > 1$ it can be unbounded.

Our *second* set of results are about the effects of networks and dynamics on budget differences across the players. We introduce and study a new quantity called the *Budget Multiplier*. For any fixed graph, local dynamics, and initial budgets, with $K_R \geq K_B$, let (σ_R, σ_B) be the Nash equilibrium for which the quantity

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

is *maximized*; this quantity is just the ratio of the final payoffs divided by the ratio of the initial budgets. The resulting maximized quantity is the Budget Multiplier, and it measures the extent to which the larger budget player can obtain a final market share that exceeds their share of the initial budgets.

Theorem 3 shows that if the switching function is concave and the selection function is linear, then the (pure strategy) Budget Multiplier is bounded above by 2, uniformly across all networks. The proof imports elements of the proof for the PoA upper bound, and additionally employs a method for attributing adoptions back to the initial seeds that generated them.

Our next result, Theorem 4, shows that even a slight departure from linearity in the selection function can yield unbounded Budget Multiplier. The proof again appeals to network structures that amplify the nonlinearity of g by self-composition, which has the effect of “squeezing out” the player with smaller budget. Combining the Budget Multiplier upper and lower bounds again allows us to exhibit simple parametric forms yielding threshold behavior: for instance, if f is linear and g is from the well-studied Tullock contest function family (discussed shortly), which includes linear g and therefore bounded Budget Multiplier, even an infinitesimal departure from linearity can result in unbounded Budget Multiplier.

2. Related Literature

Our paper contributes to the study of competitive strategy in network environments. We build a framework which combines ideas from economics (contests, competitive seeding and advertising) and computer science – uniform bounds on properties of equilibria, as in the Price of Anarchy – to address a topical and natural question. The Tullock contest function was introduced in Tullock [10]; for an axiomatic development see Skaperdas [11]. For early and influential studies of competitive advertising, see Butters [12] and Grossman and Shapiro [13]. The Price of Anarchy (PoA) was introduced by Koutsoupias and Papadimitriou [14], and important early results bounding the PoA in networked settings regardless of network structure were given by Roughgarden and Tardos [15]. The tension between equilibrium and Nash efficiency is a recurring theme in economics; for a general result on the inefficiency of Nash equilibria, see Dubey [16].

More specifically, we contribute to the study of influence in networks. This has been an active field of study in the last decade, see e.g., Ballester, Calvo-Armengol and Zenou [17]; Bharathi, Kempe and Salek [7]; Galeotti and Goyal [18]; Kempe, Kleinberg, and Tardos [19, 20]; Mossel and Roch [21]; Borodin, Filmus, and Oren [8]; Chasparis and Shamma [9]; Carnes et al [22]; Dubey, Garg and De Meyer [23]; Vetta [24]. There are three elements in our framework which appear to be novel: one, we consider a fairly general class of adoption rules at the individual consumer level which correspond to different roles which social interaction can potentially play (existing work often considers specific local dynamics); two, we study competition for influence in a network (existing work has often focused on the case of a single player seeking to maximize influence), and three, we introduce and study the notion of Budget Multiplier as a measure of how networks amplify budget differences. To the best of our knowledge, our results on the relationship between the dynamics and qualitative features of the strategic equilibrium are novel. Nevertheless, there are definite points of contact between our results and proof techniques and earlier research in (single-player and competitive) contagion in networks that we shall elaborate on where appropriate.

3. Model

3.1. Graph, Allocations, and Seeds

We consider a 2-player game of competitive adoption on a (possibly directed) graph G over n vertices. G is known to the two players, whom we shall refer to as R (ed) and B (lue).⁶ We shall also use R, B and U (ninfected) to denote the state of a vertex in G , according to whether it is currently infected by one of the two players or uninfected (Note that we are overloading the notation by giving the same name to the nodes' states and to the players, however throughout it will be clear from the context which one we are referring to). The two players simultaneously choose some number of vertices to initially seed; after this seeding, the stochastic dynamics of local adoption (discussed below) determine how each player's seeds spread throughout G to create adoptions by

⁶The restriction to 2 players is primarily for simplicity; our main results on PoA can be generalized to a game with 2 or more players, see Appendix B.

new nodes. Each player seeks to maximize their (expected) total number of eventual adoptions.⁷

More precisely, suppose that player $p = R, B$ has *budget* $K_p \in \mathbf{N}_+$; Each player p chooses an allocation of budget across the n vertices, $a_p = (a_{p1}, a_{p2}, \dots, a_{pn})$, where $a_{pj} \in \mathbf{N}_+$ and $\sum_{j=1}^n a_{pj} = K_p$. Let L_p be the set of allocations for player p , which is their pure strategy space. A mixed strategy for player p is a probability distribution σ_p on L_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 $V(a_R) = \{v | a_{vR} > 0\}$, $V(a_B) = \{v | a_{vB} > 0\}$ 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). More precisely, fix any allocation (a_R, a_B) . For any vertex v , the initial state s_v of v is in $\{R, B\}$ if and only if $v \in V(a_R, a_B)$. Moreover, $s_v = R$ with probability $a_{vR}/(a_{vR} + a_{vB})$, and $s_v = B$ with probability $a_{vB}/(a_{vR} + a_{vB})$.

Following the allocation of seeds, the stochastic contagion process on G determines how these R and B infections generate new adoptions in the network. We consider a discrete time model for this process. The state of a vertex v at time t is denoted $s_{vt} \in \{U, R, B\}$, where U stands for Uninfected, R stands for infection by R , and B stands for infection by B .

3.2. The Switching-Selection Model

We assume there is an *update schedule* which determines the order in which vertices are considered for state updates. The primary simplifying assumption we shall make about this schedule is that once a vertex is infected, it is never a candidate for updating again.

Within this constraint, we allow for a variety of behaviors, such as randomly choosing an uninfected vertex to update at each time step (a form of *sequential* updating), or updating all uninfected vertices simultaneously at each time step (a form of *parallel* updating). We can also allow for an *immunity* property — if a vertex is exposed once to infection and remains uninfected after updating, it is never updated again. Update schedules may also have finite termination times or conditions — for instance, if the firms primarily care about the number of adoptions in the coming fiscal year. We can also allow schedules that update each uninfected vertex only a fixed number of times. In our framework, a schedule which perpetually updates uninfected vertices will eventually cause any connected G to become entirely infected, thus trivializing the PoA (though not necessarily the Budget Multiplier), but we allow for considerably more general schedules.⁸

For the stochastic update of an uninfected vertex v , we will consider what we shall call 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

⁷Throughout the paper, we shall use the terms *infection* and *adoption* interchangeably.

⁸The proof of Lemma 1 specifies the technical property we need of the update schedule, which is consistent with the examples mentioned here and many others.

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 (uninfected), 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 take place, we also have $g(y) + g(1 - y) = 1$, which implies $g(1/2) = 1/2$. We assume that the switching and selection functions are the same across vertices.⁹

We think of the switching function as specifying how rapidly adoption increases with the fraction of neighbors who have adopted (i.e. the stickiness of the interchangeable products or services), regardless of their R or B value; while the selection function specifies the probability of infection by each firm in terms of the local relative market share split.¹⁰ In addition to being a natural decomposition of the dynamics, our results will show that we can articulate properties of f and g which sharply characterize the PoA and Budget Multiplier. In Section 5, we shall provide economic motivation for this formulation and also illustrate with specific parametric families of functions f and g . We also discuss more general models for the local dynamics at a number of places in the paper. The Appendix also illustrates how these switching and selection functions f - g may arise out of optimal decisions made by consumers located in social networks.

Relationship to Other Models. It is natural to consider both general and specific relationships between our models and others in the literature, especially the widely studied *general threshold* model [19, 21, 7]. One primary difference is our allowance of rather general choices for the switching and selection functions f and g , and our study of how these choices influence equilibrium properties. When considering concave f — which is a special case of sub-modularity — the relationship becomes closer, and our proof techniques bear similarity to those in the general threshold model (particularly the extensive use of coupling arguments). Nevertheless there seem to be elements of our model not

⁹This is for expositional simplicity only; our main results on PoA and Budget Multiplier carry over to a setting with heterogeneity across vertices (so long as the selection function remains symmetric across the two players).

¹⁰In the threshold model a consumer switches to an action once a certain fraction of society/neighborhood adopts that action (Granovetter, 1978). In our model, heterogeneous thresholds can be captured in terms of different switching function f .

easily captured in the general threshold model, including our allowance of rather general update schedules that may depend on the state of a vertex; the general threshold model asks that all randomization (in the form of the selection of a random threshold for each vertex) occur prior to the updating process, whereas our model permits repeated randomization in subsequent updates, in a possibly state-dependent fashion. We shall make related technical comments where appropriate.

3.3. Payoffs and Equilibrium

Given a graph G and an initial allocation of seeds (a_R, a_B) , the dynamics described above — determined by f , g , and the update schedule — yield a stochastic number of eventual infections for the two players. For $p = R, B$, let χ_p denote this random variable at the termination of the dynamics. Given strategy profile (σ_R, σ_B) , the payoff to player $p = R, B$ is $\Pi_p(\sigma_R, \sigma_B) = \mathbf{E}[\chi_p | (\sigma_R, \sigma_B)]$. Here the expectation is over any randomization in the player strategies in the choice of initial allocations, and the randomization in the stochastic updating dynamics. A Nash equilibrium is a profile of strategies (σ_R, σ_B) such that σ_p maximizes player p 's payoff given the strategy σ_{-p} of the other player.

3.4. Price of Anarchy, Price of Stability and Budget Multiplier

For a fixed graph G , stochastic update dynamics, and budgets K_R, K_B , the *maximum payoff* allocation is the (deterministic) allocation (a_R^*, a_B^*) obeying the budget constraints that maximizes $\mathbf{E}[\chi_R + \chi_B | (a_R, a_B)]$. For the same fixed graph, update dynamics and budgets, let (σ'_R, σ'_B) be the Nash equilibrium strategies that *minimize* $\mathbf{E}[\chi_R + \chi_B | (\sigma_R, \sigma_B)]$ among all Nash equilibria (σ_R, σ_B) — that is, the Nash equilibrium with the smallest joint payoff, then the *Price of Anarchy* (or PoA) is defined to be

$$\frac{\mathbf{E}[\chi_R + \chi_B | (a_R^*, a_B^*)]}{\mathbf{E}[\chi_R + \chi_B | (\sigma'_R, \sigma'_B)]}$$

The Price of Anarchy is a measure of the inefficiency in resource use created due to decentralized/ non-cooperative behavior by the two players. In the context of competition between firms, one interpretation of the PoA is as a measure of the relative improvement in efficiency effected by a hypothetical merger of the firms.

For the same setting if (σ''_R, σ''_B) is the Nash equilibrium strategies that *maximize* $\mathbf{E}[\chi_R + \chi_B | (\sigma_R, \sigma_B)]$ among all Nash equilibria (σ_R, σ_B) — that is, the Nash equilibrium with the largest joint payoff, then the *Price of Stability* (or PoS) is defined to be

$$\frac{\mathbf{E}[\chi_R + \chi_B | (a_R^*, a_B^*)]}{\mathbf{E}[\chi_R + \chi_B | (\sigma''_R, \sigma''_B)]}$$

The Price of Stability is a measure of the inefficiency in resource use created due to non-cooperative behavior by the two players. In the context of competition between firms, one interpretation of the PoS is as a measure of the relative improvement in efficiency effected by a hypothetical objective authority that can help players reach a good Nash equilibrium.

We also introduce and study a new quantity called the *Budget Multiplier*. The Budget Multiplier measures the extent to which network structure and dynamics can amplify initial resource inequality across the players. Thus for any fixed graph G and stochastic

update dynamics, and initial budgets K_R, K_B , with $K_R \geq K_B$, let (σ_R, σ_B) be the Nash equilibrium for which the ratio

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

is *maximized*. 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 their share of the initial budgets.

4. Local Dynamics: Motivation

In this section, we provide some examples of the decomposition of the local update dynamics into a switching function f and a selection function g . As discussed above, we view the switching function as representing how contagious a product or service is, regardless of which competing party provides it; and we view the selection function as representing the extent to which a firm having majority local market share favors its selection in the case of adoption. We illustrate the richness of this model by examining a variety of different mathematical choices for the functions f and g , and discuss examples from the domain of technology adoption that might (qualitatively) match these forms. Finally, to illustrate the scope of this formulation, we also discuss examples of natural update dynamics that *cannot* be decomposed in this way.

A fairly broad class of dynamics is captured by the following parametric family of functions. The switching function

$$f(x) = x^r \quad r \geq 0$$

and the selection function

$$g(y) = y^s / (y^s + (1 - y)^s) \quad s \geq 0.$$

Regarding this form for f , for $r = 1$ we have linear adoption. For $r < 1$ we have f concave, corresponding to cases in which the probability of adoption rises quickly with only a small fraction of adopting neighbors, but then saturates or levels off with larger fractions of adopting neighbors. In contrast, for $r > 1$ we have f convex, which at very large values of r can begin to approximate threshold adoption behavior — the probability of adoption remains small until some fraction of neighbors has adopted, then rises rapidly. See Figure 1.

Regarding this form for g , which is known as the *Tullock contest function* (Tullock (1980)), for $s = 1$ we have a (linear) *voter* model in which the probability of selection is proportional to local market share. For $s < 1$ we have what we shall call an *equalizing* g , by which we mean that selection of the minority party in the neighborhood is favored relative to the linear voter model $g(y) = y$; and for $s > 1$ we have a *polarizing* g , meaning that the minority party is disfavored relative to the linear model. As s approaches 0, we approach the completely equalizing choice $g \equiv 1/2$, and as s approaches infinity, we approach the completely polarizing *winner-take-all* g ; see Figure 1.

These parametric families of switching and selection functions will play an important role in illustrating our general results. The appendix discussed a few technology adoption examples which are (qualitatively) covered by these families of functions. In the

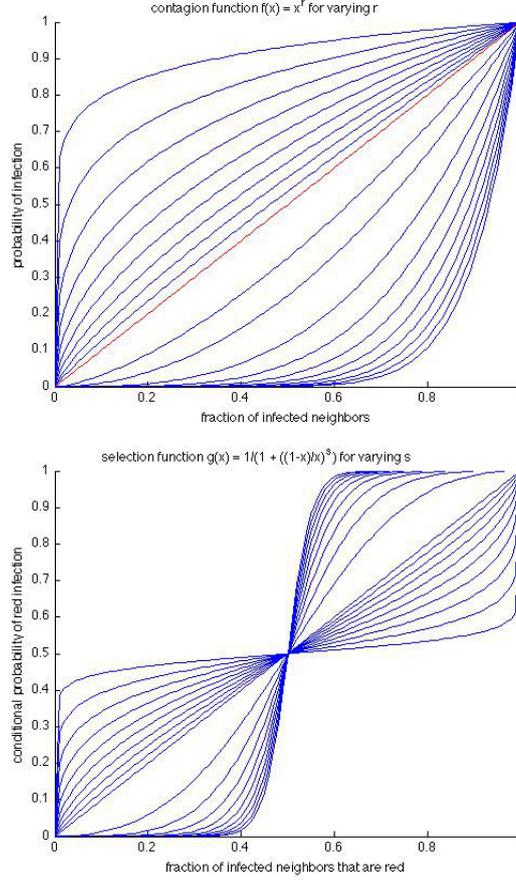


Figure 1: Left: Plots of $f(x) = x^r$ for varying choices of r , including $r = 1$ (linear, red line), $r < 1$ (concave), and $r > 1$ (convex). Right: Plots of $g(y) = y^s / (y^s + (1 - y)^s)$ for varying choices of s , including $s = 1$ (linear, red line), $s < 1$ (equalizing), and $s > 1$ (polarizing).

proofs of some of our results, it will sometimes be convenient to use a more general adoption function formulation with some additional technical conditions that are met by our switching-selection formulation. We will refer to this general, single-step model as the *generalized adoption function* model. In this model, if the local fractions of Red and Blue neighbors are α_R and α_B , the probability that we update the vertex with an R infection is $h(\alpha_R, \alpha_B)$ for some *adoption function* h with range $[0, 1]$, and symmetrically the probability of B infection is thus $h(\alpha_B, \alpha_R)$. Let us use $H(\alpha_R, \alpha_B) = h(\alpha_R, \alpha_B) + h(\alpha_B, \alpha_R)$ to denote the total infection probability under h . Note that we can still always decompose h into a two-step process by defining the switching function to be $f(\alpha_R, \alpha_B) = H(\alpha_R, \alpha_B)$ and defining the selection function to be

$$g(\alpha_R, \alpha_B) = h(\alpha_R, \alpha_B) / (h(\alpha_R, \alpha_B) + h(\alpha_B, \alpha_R))$$

which is the infection-conditional probability that R wins the infection. The switching-

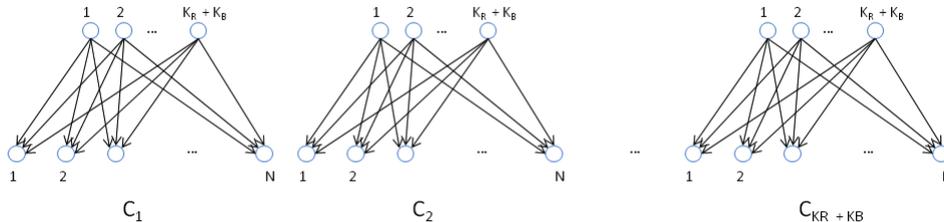


Figure 2: A sample graph G .

selection model is thus the special case of the generalized adoption function model in which $H(\alpha_R, \alpha_B) = f(\alpha_R + \alpha_B)$ is a function of only $\alpha_R + \alpha_B$, and $g(\alpha_R, \alpha_B)$ is a function of only $\alpha_R/(\alpha_R + \alpha_B)$.¹¹

5. Examples

We provide a simple example showing the effect of the switching and selection functions, and the relative budgets of the players, on the equilibrium solution.

Consider the graph G illustrated in figure 2. Here N is chosen to be much larger than $(K_R + K_B)$. The update schedule we consider for this example is a single parallel update, however all the claims hold as well for any other update schedule in which every node is updated exactly once. We also assume $f(x) = x^\alpha$, therefore f is convex when $\alpha > 1$, it is linear when $\alpha = 1$, and it is concave when $0 \leq \alpha \leq 1$.

Observe that in an equilibrium on G , no seed is ever placed on the bottom layer of a component, therefore in order to simplify the statement of our claims, we say a player *chooses* one component, when she puts some of her seeds on the top layer of that component.

We claim that the following propositions hold:

- If f is concave, regardless of the form of g (whether linear, equalizing or punishing), the only equilibrium is the case where players put each of their seeds in a separate component of G , i.e. there is no more than one seed in each component.

Here is the reason: Suppose not and in the equilibrium, $r + b > 1$ seeds are on the top layer of one of the components, say C_i , where r is the number of Red seeds, and b is the number of blue seeds on C_i . We know that the expected number of infections in that component is $T = N \times f(\frac{r+b}{K_R+K_B})$. Let us define the “fair share”

¹¹While the decomposition in terms of a switching function and a selection function accommodates a fairly wide range of adoption dynamics there are some cases which are ruled out. Consider the choice $h(x, y) = x(1 - y^2)$; it is easily verified that the total probability of adoption $H(x, y)$ is increasing in x and y . But $H(x, y)$ clearly cannot be expressed as a function of the form $f(x + y)$. Similarly, it is easy to construct an adoption function that is not only not decomposable, but violates monotonicity. Imagine consumers that prefer to adopt the majority choice in their neighborhood, but will only adopt once their local neighborhood market is sufficiently settled in favor of one or the other product. The probability of total adoption may then be higher with $x = 0.2, y = 0$ as compared to $x = y = 0.4$.

of Red and Blue in C_i to be $T \times \frac{r}{r+b}$ and $T \times \frac{b}{r+b}$ respectively. Observe that regardless of g , there exist a player (say Red) whose actual share is less than or equal to her fair share. Now note that if the Red player moves her seeds from C_i and put each of them in a separate empty component, her expected payoff would be $N \times f(\frac{1}{K_R+K_B}) \times r$ which we claim is larger than her current payoff. Note that Red's current payoff is at most $N \times f(\frac{r+b}{K_R+K_B}) \times \frac{r}{r+b}$, so it suffices to show that:

$$N \times f(\frac{1}{K_R+K_B}) \times r > N \times f(\frac{r+b}{K_R+K_B}) \times \frac{r}{r+b}$$

To prove the above holds, we note that it is equivalent to:

$$f(\frac{1}{K_R+K_B}) > f(\frac{r+b}{K_R+K_B}) \times \frac{1}{r+b}$$

and the last equation holds due to concavity of f . This means that the Red player can increase her payoff by moving her seeds from C_i to empty components of G , which contradicts the assumption of being in an equilibrium.

- If f is convex and g is linear, the only equilibrium is the case where both players put all of their seeds in the same component.

The reason the above holds is as follows: When g is linear every player gets her fair share (as defined earlier) of the total number of infections in each component. Now note that every time a player moves her seeds from one component to a component with larger number of seeds on, due to the convexity of f , she is increasing the total number of her infections in G . To be more precise, suppose there are $r+b$ seeds on component C_i and $r'+b'$ seeds on component C_j where $r+b \geq r'+b'$. The payoff per seed for Red in C_i and C_j are:

$$N \times f(\frac{r+b}{K_R+K_B})/(r+b)$$

and

$$N \times f(\frac{r'+b'}{K_R+K_B})/(r'+b')$$

respectively. By moving her seeds from C_j to C_i , the payoff per seed for Red will become:

$$N \times f(\frac{r+b+r'}{K_R+K_B})/(r+b+r')$$

Which is larger than the previous two due to the convexity of f . This means that the best response to a given opponent strategy is always to find the component with the largest number of seeds, and put all the seeds there. Therefore the only equilibrium will be the case where all the seeds are on the same component.

- If f is convex, g is equalizing, and $\frac{K_R}{K_B} = 1$, depending on whether f or g has a stronger impact, the equilibrium solution can vary. In other words the convexity of f makes players put their seeds close to one another to increase the total number of infections, whereas the equalizing effect of g makes them avoid having the majority of seeds in any component. Therefore players face a trade-off that results in various equilibrium solutions.

If we fix g and make f very convex then the only equilibrium will be the case where both players put all of their seeds in the same component. The reason is that adding a seed to a larger group of seeds increases the total number of infections by a relatively large amount that makes the unfair share acceptable.

On the other hand, if we fix g and make f less and less convex, not being the majority becomes a more important condition to satisfy. If f is linear for example, one equilibrium is the case where Red and Blue seeds are paired with each other on K_R different components (it is very easy to verify this: first note that moving seeds does not change the total number of infections in G ; second, moving a seed to an empty component clearly does not add to the number of infections one can get –although the assignment remains stable, and third and last, moving a seed to a component that already has a few seeds on, just makes the corresponding player the majority on that component, which is not desirable).

- If f is convex, g is punishing, and $\frac{K_R}{K_B} = 1$, one equilibrium is the case where both players put all of their seeds in the same component. Note that in this case the convexity of f makes players put their seeds close to each other. In addition the punishing effect of g makes being the majority desirable. Now it is easy to see that the above case is an equilibrium: given the opponent puts all her seeds on one component, the best one can do is to add her seeds to the same component.

If f is too convex, then the above will be the only equilibrium. As f gets closer to linear, more stable solutions emerge as well.

- If f is convex, g is equalizing, and $\frac{K_R}{K_B} > 1$, again various solutions are possible depending on f and g .

If f is very convex the only equilibrium is the case where both players put all of their seeds in the same component. As f becomes closer to linear, the red player plays as diverse as she can by putting each of her seeds on different components of G to avoid becoming the majority in any component. The best response of Blue then would be to play each of her seeds in a different component (not necessarily empty of Red seeds).

- If f is convex, g is punishing, and $\frac{K_R}{K_B} > 1$, then the Red (larger budget) player tries to gather her seeds close to each other and the Blue player's seeds, as it makes both the total infections and her share, larger at the same time. The blue player on the other hand avoids playing close to Red seeds, unless f is very convex and he gets very few infections in case of not playing close to Red seeds. Thus depending on which function has a stronger impact, the equilibrium solution may be different.

If f is very convex the only equilibrium is the case where both players put all of their seeds in the same component. As f becomes closer to linear, the blue player can

play farther from Red. If f is linear for example, there is no pure Nash equilibrium, but the high-level strategies of the players would be to place their seeds in a way that they become the majority in as many components as they can.

The general points the above example makes are the following:

1. Fixing g , As f becomes more and more concave, each player tends to keep their seeds far from one another and also from the opponent's seeds.
2. Fixing g , the more convex f gets, the more the players tendency to keep their seeds close to each other and also close to the opponent's seeds.
3. Fixing f , the more equalizing g gets, the player with smaller budget tends to play near the large-budget player, while the large budget player avoids being the majority. The quantity $\frac{K_R}{K_B}$ indicates how effective this is: if it is small, e.g. 1, g does not affect the equilibrium greatly.
4. Fixing f , the more punishing g gets, the player with larger budget tends to play near the small-budget player, while the small budget player avoids playing close to her opponent. The quantity $\frac{K_R}{K_B}$ indicates how effective this is: if it is small, e.g. 1, g does not affect the equilibrium greatly.

Here we chose a symmetric network to be able to see the pure effect of f , g and $\frac{K_R}{K_B}$ on the equilibrium. When the underlying network is not symmetric players face another trade-off that rises due to the effect of the network structure. For example although it is in general true that the more concave f gets, the more far apart players put their seeds, this does not hold in all networks, e.g. if the network has only a few important vertices, then players still play near each other to compete for those high impact vertices.

6. Results: Price of Anarchy and Price of Stability

We first develop a number of examples to illustrate the trade-offs and the issues with regard to costs of decentralization. We then state and prove a theorem providing general conditions in the switching-selection model under which the Price of Anarchy is bounded by a constant that is independent of the size and structure of the graph G . The simplest characterization is that f being any concave function (satisfying $f(0) = 0$, $f(1) = 1$ and f increasing), and g being the linear voter function $g(y) = y$ leads to bounded PoS; but we shall see the conditions allow for certain combinations of concave f and nonlinear g as well. We then prove a lower bound showing that the concavity of f is required for bounded PoS in a very strong sense — a small amount of convexity can lead to unbounded PoS.

6.1. PoA: Examples

Suppose that budgets of the firms are $K_R = K_B = 1$, and the update rule is such that all vertices are updated only once. The network contains two connected components with 10 vertices and 100 vertices, respectively. In each component there are 2 influential vertices, each of which is connected to the other 8 and 98 vertices, respectively. So in component 1, there are 16 directed links while in component 2 there are 196 directed links in all.

- Suppose that the switching function and the selection function are both linear, $f(x) = x$ and $g(y) = y$. Then there is a unique equilibrium in which players place their seeds on distinct influential vertices of component 2. The total infection is then 100 and this is the maximum number of infections possible with 2 seeds. So here the PoA is 1.
- Let us now alter the switching function such that $f(1/2) = \epsilon$ for some $\epsilon < 1/25$ (keeping $f(1) = 1$, as always), but retain the selection function to be $g(y) = y$. Now there also exists an equilibrium in which the firms locate on the influential vertices of component 1. In this equilibrium payoffs to each player are equal to 5. Observe that for $\epsilon < 1/25$, a deviation to the other component is not profitable: it yields an expected payoff equal to $1 + \epsilon \times 98$, and this is strictly smaller than 5. Since it is still possible to infect component 2 with 2 seeds, the PoA is 10. Here inefficiency is created by a coordination failure of the players.
- Finally, suppose there is only one component with 110 vertices, with 2 influential vertices and 108 vertices receiving directed links. Then equilibrium under both switching functions considered above will involve firms locating at the 2 influential vertices and this will lead to infection of all vertices. So the PoA is 1, irrespective of whether the switching function is linear $f(x) = x$ or whether $f(1/2) < 1/25$.

Thus for a fixed network, updating rule and selection function, variations in the switching function can generate large variations in the PoA. Similarly, for fixed update rule and switching and selection functions, a change in the network structure yields very different PoA.

Theorem 1 provides a set of sufficient conditions on switching and selection function, under which the PoA is uniformly bounded from above. Theorem 2 shows how even small violations of these conditions can lead to arbitrarily high PoA.

6.2. PoA: Upper Bound

We find it useful to state and prove our theorems using the generalized adoption model formulation described in section 4, but with some additional conditions on h that we now discuss. If $h(\alpha_R, \alpha_B)$ (respectively, $h(\alpha_B, \alpha_R)$) is the probability that a vertex with fractions α_R and α_B of R and B neighbors is infected by R (respectively, B), we say that the total infection probability $H(\alpha_R, \alpha_B) = h(\alpha_R, \alpha_B) + h(\alpha_B, \alpha_R)$ is *additive in its arguments* (or simply *additive*) if H can be written $H(\alpha_R, \alpha_B) = f(\alpha_R + \alpha_B)$ for some increasing function f — in other words, h permits interpretation as a switching function. We shall say that h is *competitive* if $h(\alpha_R, \alpha_B) \leq h(\alpha_R, 0)$ for all $\alpha_R, \alpha_B \in [0, 1]$. In other words, a player always has equal or higher infection probability in the absence of the other player.

Concave f and linear g . Observe that the switching-selection formulation always satisfies the additivity property by definition. Moreover, in the switching-selection formulation, if g is linear, the competitiveness condition becomes

$$h(x, y) = f(x + y)(x/(x + y)) \leq f(x) = h(x, 0)$$

or

$$f(x + y)/(x + y) \leq f(x)/x.$$

This condition is satisfied by the concavity of f . We will later see that the following theorem also applies to certain combinations of concave f and nonlinear g . The first theorem can now be stated.

Theorem 1. *If the adoption function $h(\alpha_R, \alpha_B)$ is competitive and H is additive in its arguments, then Price of Anarchy is at most $\frac{1}{4}$ for any graph G .*

Proof We establish the theorem via a series of lemmas and inequalities that can be summarized as follows. Let (S_R^*, S_B^*) be an initial allocation of infections that gives the maximum joint payoff, and let (S_R, S_B) be a pure¹² Nash equilibrium with S_R being the larger set of seeds, so $K_R = |S_R^*| = |S_R| \geq K_B = |S_B^*| = |S_B|$. We first establish a general lemma (Lemma 1) that implies that the set S_R^* alone (without S_B^* present) must yield payoffs close to the maximum joint payoff (Corollary 1). The proof involves the construction of a coupled stochastic process technique we employ repeatedly in the paper.¹³ We then contemplate a deviation by the Red player to (S_R^*, S_B) . Another coupling argument (Lemma 2) establishes that the total payoffs for both players under (S_R^*, S_B) must be at least those for the Red player alone under (S_R^*, \emptyset) . This means that under (S_R^*, S_B) , one of the two players must be approaching the maximum joint infections. If it is Red, we are done, since Red's equilibrium payoff must also be this large. If it is Blue, Lemma 1 implies that Blue could still get this large payoff even after the departure of Red. Next we invoke Lemma 2 to show that total eventual payoff to both players under (S_R, S_B) must exceed this large payoff accruing to Blue, proving the theorem.

Lemma 1. *Let A_R and A_B be any sets of seed vertices for the two players. Then if h is competitive,*

$$\mathbf{E}[\chi_R | (A_R, \emptyset)] \geq \mathbf{E}[\chi_R | (A_R, A_B)]$$

and

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

Proof We provide the proof for the first statement involving χ_R ; the proof for χ_B is identical. We introduce a simple *coupled simulation* technique that we shall appeal to several times throughout the paper. Consider the stochastic dynamical process on G under two different initial conditions: both A_R and A_B are present (the *joint* process, denoted (A_R, A_B) in the conditioning in the statement of the lemma); and only the set A_R is present (the *solo Red* process, denoted (A_R, \emptyset)). Our goal is to define a new stochastic process on G , called the *coupled process*, in which the state of each vertex v will be a pair $\langle X_v, Y_v \rangle$. We shall arrange that X_v faithfully represents the state of

¹²The extension to mixed strategies is straightforward and omitted.

¹³The theorem includes concave (and therefore sub-modular) f and makes extensive use of coupling arguments to prove local-to-global effects (of which Lemmas 1 and 2 are examples); this bears a similarity to the work of Mossel and Roch [21], and it has been suggested that our proofs might be simplified by appeal to their results. However, we have not been able to apply their results in our context. Two features of our framework seem to make direct application difficult: first, the important role of competitive effects, which is explicit in Lemma 1; and second, the variety of updating schedules we consider appear not be covered by the general threshold model which underlies the Mossel and Roch analysis. While we suspect more direct relationships might be possible in special cases, here we provide proofs specific to our model.

a vertex in the joint process, and Y_v the state in the solo Red process. However, these state components will be correlated or coupled in a deliberate manner. More precisely, we wish to arrange the coupled process to have the following properties:

1. At each step, and for any vertex state $\langle X_v, Y_v \rangle$, $X_v \in \{U, R, B\}$ and $Y_v \in \{U, R\}$.
2. Projecting the states of the coupled process onto either component faithfully yields the respective process. Thus, if $\langle X_v, Y_v \rangle$ represents the state of vertex v in the coupled process, then the $\{X_v\}$ are stochastically identical to the joint process, and the $\{Y_v\}$ are stochastically identical to the solo Red process.
3. At each step, and for any vertex state $\langle X_v, Y_v \rangle$, $X_v = R$ implies $Y_v = R$.

Note that the first two properties are easily achieved by simply running *independent* joint and solo Red processes. But this will violate the third property, which yields the lemma, and thus we introduce the coupling.

For any vertex v , we define its initial coupled process state $\langle X_v, Y_v \rangle$ as follows: $X_v = R$ if $v \in A_R$, $X_v = B$ if $v \in A_B$, and $X_v = U$ otherwise; and $Y_v = R$ if $v \in A_R$, and $Y_v = U$ otherwise. It is easily verified that these initial states satisfy Properties 1 and 3 above, thus encoding the initial states of the two separate processes.

Assume for now that the first vertex¹⁴ v to be updated in the X and Y processes are the same — i.e. the same vertices are updated in both the joint and solo update schedules, which may in general depend on the state of the network in each. We now describe the coupled updates of v . Let α_v^R denote the fraction of v 's neighbors w such that $X_w = R$, and α_v^B the fraction such that $X_w = B$. Note that by the initialization of the coupled process, α_v^R is also equal to the fraction of $Y_w = R$ (which we denote $\tilde{\alpha}_v^R$).

In the joint process, the probability that v is updated to R is $h(\alpha_v^R, \alpha_v^B)$, and to B is $h(\alpha_v^B, \alpha_v^R)$. In the solo Red process, the probability that v is updated to R is $h(\alpha_v^R, 0)$, which by competitiveness is greater than or equal to $h(\alpha_v^R, \alpha_v^B)$.

We can thus define the update dynamics of the coupled process as follows: pick a real value z uniformly at random from $[0, 1]$. Update the state $\langle X_v, Y_v \rangle$ of v as follows:

- X_v update: If $z \in [0, h(\alpha_v^R, \alpha_v^B))$, update X_v to R ; if $z \in [h(\alpha_v^R, \alpha_v^B), h(\alpha_v^R, \alpha_v^B) + h(\alpha_v^B, \alpha_v^R)]$, update X_v to B ; otherwise, update X_v to U . Note that the probabilities X_v are updated to R and B exactly match those of the joint process, as required by Property 2 above. See Figure 3.
- Y_v update: If $z \in [0, h(\alpha_v^R, 0)]$, update Y_v to R ; otherwise, update Y_v to U . The probability Y_v is updated to R is thus exactly $h(\alpha_v^R, 0)$, matching that in a solo Red process. See Figure 3.

Since by competitiveness, $z \in [0, h(\alpha_v^R, \alpha_v^B))$ implies $z \in [0, h^R(\alpha_v^R, 0)]$, we ensure Property 3. Thus in subsequent updates we shall have $\alpha^R \leq \tilde{\alpha}^R$. Thus as long as $h(\alpha^R, \alpha^B) \leq h(\tilde{\alpha}^R, 0)$ we can continue to maintain the invariant. These inequalities follow from competitiveness.

So far we have assumed the same vertices were candidates for updating in both the joint and solo processes; while this may be true for some update schedules, including

¹⁴Or vertices; the argument will be similar if more than one vertex are updated at the same time.

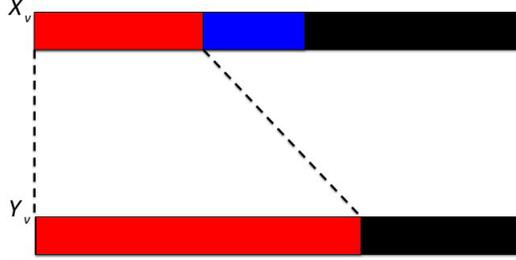


Figure 3: Illustration of the coupled dynamics defined in the proof of Lemma 1. In the update dynamic for X_v (top line), the probabilities of Red and Blue updates are represented by disjoint line segments of length $h(\alpha_v^R, \alpha_v^B)$ and $h(\alpha_v^B, \alpha_v^R)$ respectively. By competitiveness, the Red segment has length less than $h(\alpha_v^R, 0)$, which is the probability of Red update of Y_v (bottom line). The dashed red lines indicate this inequality. Thus by the arrangement of the line segments we enforce the invariant that $X_v = R$ implies $Y_v = R$.

fixed or predetermined schedules, in general it will not be. So we discuss the case of random, and parallel schedules separately. For a random schedule in which the order of state updates is determined randomly, our coupling argument still works, we just need to fix one specific realization of the random schedule at the very beginning and then repeat our coupling argument. Finally we handle parallel schedules by considering three cases. Case 1: Assuming Property 3 holds, if a vertex is a candidate for updating in both processes, we can maintain this property by performing the coupled updates described above. Case 2: If a vertex v is a candidate for updating only in the solo red process, then by Property 3 X_v cannot be R , so Property 3 will still hold after the update of Y_v . Case 3: Finally, if v is a candidate for updating only in the joint process, then if $Y_v = R$, Property 3 will still hold after the update of X_v , and if $Y_v = U$ and all neighbors of v in the joint process are B , Property 3 will remain true after the update. The only case remaining is that $Y_v = U$ and v has R neighbors in the joint process. This is impossible for the parallel update schedule: v should have also been a candidate for updating in the solo red process, since by Property 3 v has weakly more R neighbors in the solo process.

Since Properties 2 and 3 hold on an update-by-update basis in any run or sample path of the coupled dynamics, they also hold in expectation over runs, yielding the statement of the lemma. \square (Lemma 1)

Corollary 1. *Let A_R and A_B be any sets of seeded nodes for the two players. Then if the adoption function $h(\alpha^R, \alpha^B)$ is competitive,*

$$\mathbf{E}[\chi_R + \chi_B | (A_R, A_B)] \leq \mathbf{E}[\chi_R | (A_R, \emptyset)] + \mathbf{E}[\chi_B | (\emptyset, A_B)].$$

Proof Follows from linearity of expectation applied to the left hand side of the inequality, and two applications of Lemma 1. \square (Corollary 1)

Let (S_R^*, S_B^*) be the maximum joint payoff seed sets. Let (S_R, S_B) be any (pure) Nash equilibrium, with S_R having the larger budget. Corollary 1 implies either $\mathbf{E}[\chi_R | (S_R^*, \emptyset)]$

or $\mathbf{E}[\chi_B | (\emptyset, S_B^*)]$ is at least as great as $\mathbf{E}[\chi_R + \chi_B | (S_R^*, S_B^*)]/2$; so assume without loss of generality¹⁵ that $\mathbf{E}[\chi_R | (S_R^*, \emptyset)] \geq \mathbf{E}[\chi_R + \chi_B | (S_R^*, S_B^*)]/2$. Let us now contemplate a unilateral deviation of the Red player from S_R to S_R^* , in which case the strategies are (S_R^*, S_B) . In the following lemma we show that the *total* number of eventual adoptions for the two players is larger than adoptions accruing to a single player under solo seeding.

Lemma 2. *Let A_R and A_B be any sets of seeded nodes for the two players. If H is additive,*

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

Proof We employ a coupling argument similar to that in the proof of Lemma 1. We define a stochastic process in which the state of a vertex v is a pair $\langle X_v, Y_v \rangle$ in which the following properties are obeyed:

1. At each step, and for any vertex state $\langle X_v, Y_v \rangle$, $X_v \in \{R, B, U\}$ and $Y_v \in \{R, U\}$.
2. Projecting the state of the coupled process onto either component faithfully yields the respective process. Thus, if $\langle X_v, Y_v \rangle$ represents the state of vertex v in the coupled process, then the $\{X_v\}$ are stochastically identical to the joint process (A_R, A_B) , and the $\{Y_v\}$ are stochastically identical to the solo Red process (A_R, \emptyset) .
3. At each step, and for any vertex state $\langle X_v, Y_v \rangle$, $Y_v = R$ implies $X_v = R$ or $X_v = B$.

We initialize the coupled process in the obvious way: if $v \in A_R$ then $X_v = R$, if $v \in A_B$ then $X_v = B$, and $X_v = U$ otherwise; and if $v \in A_R$ then $Y_v = R$, and $Y_v = U$ otherwise. Let us fix a vertex v to update, and let α_v^R, α_v^B denote the fraction of neighbors w of v with $X_w = R$ and $X_w = B$ respectively, and let $\tilde{\alpha}_v^R$ denote the fraction with $Y_w = R$. Initially we have $\alpha_v^R = \tilde{\alpha}_v^R$.

We assume the vertex or vertices v to be updated in the X and Y processes are the same; the fact that the update schedules may cause these sets to differ is handled in the same way as in the proof of Lemma 1. On the first update of v in the joint process (A_R, A_B) , the total probability infection by either R or B is

$$H(\alpha_v^R, \alpha_v^B) = h(\alpha_v^R, \alpha_v^B) + h(\alpha_v^B, \alpha_v^R).$$

In the solo process (A_R, \emptyset) , the probability of infection by R is $h(\alpha_v^R, 0) \leq h(\alpha_v^R, 0) + h(0, \alpha_v^R) = H(\alpha_v^R, 0) \leq H(\alpha_v^R, \alpha_v^B)$ where the last inequality follows by the additivity of H .

We thus define the update dynamics in the coupled process as follows: pick a real value z uniformly at random from $[0, 1]$. Update $\langle X_v, Y_v \rangle$ as follows:

¹⁵If this does not hold, instead of (S_R^*, S_B^*) , consider the seed set (T_R^*, T_B^*) in which $T_R^* = S_B^* \cup (S_R^* - S)$ and $T_B^* = S$ where S is any subset of S_R^* of size K_B for which $S \cup S_B^* \cup (S_R^* - S) = S_R^* \cup S_B^*$. Note that this new seed set also maximizes the joint payoff.

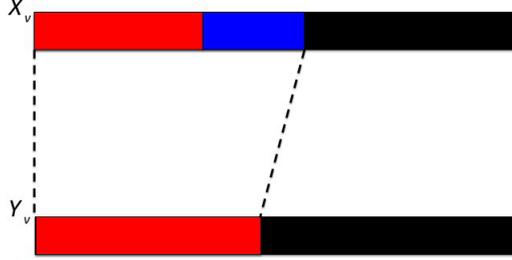


Figure 4: Illustration of the coupled dynamics defined in the proof of Lemma 2. In the update dynamic for X_v (top line), the probabilities of Red and Blue updates are represented by line segments of length $h(\alpha_v^R, \alpha_v^B)$ and $h(\alpha_v^B, \alpha_v^R)$ respectively. By additivity of H , together these two segments are greater than $h(\alpha_v^R, 0)$ which is the probability of Red update of Y_v (bottom line). This inequality is represented by the dashed black lines.

- X_v update:
 If $z \in [0, h(\alpha_v^R, \alpha_v^B))$, update X_v to R ;
 if $z \in [h(\alpha_v^R, \alpha_v^B), h(\alpha_v^R, \alpha_v^B) + h(\alpha_v^B, \alpha_v^R)]$
 $\equiv [h(\alpha_v^R, \alpha_v^B), H(\alpha_v^R, \alpha_v^B)]$, update X_v to B ;
 otherwise update X_v to U . See Figure 4.
- Y_v update: If $r \in [0, h(\alpha_v^R, 0))$, update Y_v to R ;
 otherwise update Y_v to U . See Figure 4.

It is easily verified that at each such update, the probabilities of R and B updates of X_v are exactly as in the joint (A_R, A_B) process, and the probability of an R update of Y_v is exactly as in the solo (A_R, \emptyset) process, thus maintaining Property 2 above. Property 3 follows from the previously established fact that $h(\alpha_v^R, 0) \leq H(\alpha_v^R, \alpha_v^B)$, so whenever Y_v is updated to R , X_v is updated to either R or B .

Notice that since $h(\alpha_v^R, 0) \geq h(\alpha_v^R, \alpha_v^B)$ by competitiveness, for the overall theorem (which requires competitiveness of h) we *cannot* ensure that $Y_v = R$ is always accompanied by $X_v = R$. Thus the Red infections in the solo process may exceed those in the joint process, yielding $\tilde{\alpha}_v^R > \alpha_v^R$ for subsequent updates. To maintain Property 3 in subsequent updates we thus require that $\tilde{\alpha}_v^R \leq \alpha_v^R + \alpha_v^B$ implies $h(\tilde{\alpha}_v^R, 0) \leq H(\tilde{\alpha}_v^R, 0) \leq H(\alpha_v^R, \alpha_v^B)$ which follows from the additivity of H . Also, notice that since the lemma holds for every fixed A_R and A_B , it also holds in expectation for mixed strategies. \square (Lemma 2)

Continuing the analysis of a unilateral deviation by the Red player from S_R to S_R^* , we have thus established

$$\begin{aligned}
& \mathbf{E}[\chi_R + \chi_B | (S_R^*, S_B)] \\
&= \mathbf{E}[\chi_R | (S_R^*, S_B)] + \mathbf{E}[\chi_B | (S_R^*, S_B)] \\
&\geq \mathbf{E}[\chi_R | (S_R^*, \emptyset)] \\
&\geq \mathbf{E}[\chi_R + \chi_B | (S_R^*, S_B^*)] / 2
\end{aligned}$$

where the equality is by linearity of expectation, the first inequality follows from Lemma 2, and the second inequality from Corollary 1. Thus at least one of $\mathbf{E}[\chi_R|(S_R^*, S_B)]$ and $\mathbf{E}[\chi_B|(S_R^*, S_B)]$ must be at least $\mathbf{E}[\chi_R + \chi_B|(S_R^*, S_B^*)]/4$.

If $\mathbf{E}[\chi_R|(S_R^*, S_B)] \geq \mathbf{E}[\chi_R + \chi_B|(S_R^*, S_B^*)]/4$, then since (S_R, S_B) is Nash, $\mathbf{E}[\chi_R|(S_R, S_B)] \geq \mathbf{E}[\chi_R + \chi_B|(S_R^*, S_B^*)]/4$, and the theorem is proved. The only remaining case is where $\mathbf{E}[\chi_B|(S_R^*, S_B)] \geq \mathbf{E}[\chi_R + \chi_B|(S_R^*, S_B^*)]/4$. But Lemma 1 has already established that $\mathbf{E}[\chi_B|(\emptyset, S_B)] \geq \mathbf{E}[\chi_B|(S_R^*, S_B)]$, and we have $\mathbf{E}[\chi_R + \chi_B|(S_R, S_B)] \geq \mathbf{E}[\chi_B|(\emptyset, S_B)]$ from Lemma 2. Combining, we have the following chain of inequalities:

$$\begin{aligned} \mathbf{E}[\chi_R + \chi_B|(S_R, S_B)] &\geq \mathbf{E}[\chi_B|(\emptyset, S_B)] \\ &\geq \mathbf{E}[\chi_B|(S_R^*, S_B)] \\ &\geq \mathbf{E}[\chi_R + \chi_B|(S_R^*, S_B^*)]/4 \end{aligned}$$

thus establishing the theorem. □ (Theorem 1)

Concave f , non-linear g . Recall that the switching-selection formulation in which f is concave and g is linear satisfies the hypothesis of the Theorem above. But Theorem 1 also provides more general conditions for bounded PoA in the switching-selection model. For example, suppose we consider switching functions of the form $f(x) = x^r$ for $r \leq 1$ (thus yielding concavity) and selection functions of the Tullock contest form $g(y) = y^s/(y^s + (1 - y)^s)$, as discussed in Section 3.2. Letting a and b denote the local fraction of Red and Blue neighbors for notational convenience, this leads to an adoption function of the form $h(a, b) = (a + b)^r/(1 + (b/a)^s)$. The condition for competitiveness is

$$h(a, 0) - h(a, b) = a^r - (a + b)^r/(1 + (b/a)^s) \geq 0.$$

Dividing through by $(a + b)^r$ yields

$$\begin{aligned} (a/(a + b))^r - 1/(1 + (b/a)^s) \\ = 1/(1 + (b/a))^r - 1/(1 + (b/a)^s) \geq 0. \end{aligned}$$

Making the substitution $z = b/a$ and moving the second term to the right-hand side gives

$$1/(1 + z)^r \geq 1/(1 + z^s).$$

Thus competitiveness is equivalent to the condition $1 + z^s \geq (1 + z)^r$ for all $z \geq 0$. It is not difficult to show that any $s \in [r, 1]$ will satisfy this condition. In other words, the more concave f is (i.e. the smaller r is), the more equalizing g can be (i.e. the smaller s can be) while maintaining competitiveness. By Theorem 1 we have thus shown:

Corollary 2. *Let the switching function be $f(x) = x^r$ for $r \leq 1$ and let $g(y) = y^s/(y^s + (1 - y)^s)$ be the selection function. Then as long as $s \in [r, 1]$, the Price of Anarchy is at most 4 for any graph.*

6.3. PoS: Lower Bound

We now show that concavity of the switching function is required in a very strong sense — essentially, even a slight convexity leads to unbounded PoS.

Theorem 2. *Let the switching function be $f(x) = x^r$ for any $r > 1$, and let the selection function be linear $g(y) = y$. Then for any $V > 0$, there exists a graph G for which the Price of Stability is greater than V .*

Proof The idea is to create a layered, directed graph whose dynamics rapidly amplify the convexity of f . We show that for any arbitrary natural number V , there is a graph in which the price of stability is larger than V .

Suppose $K_R = 1$ and $K_B = K$ where $K > V$. Let $\epsilon = \frac{1}{(K+1)^4}$. Since $r > 1$ and $\frac{K}{K+1} < 1$ we can choose a constant l such that $(\frac{K}{K+1})^{r^l} < \epsilon$.

Consider the graph G consisting of two components A, B depicted in figure 5. Component B is a $(l+2)$ -layered graph, where the first layer has $(K+1)$ vertices, the second, third, and the $(l+1)$ th layer all have M vertices with $M = \frac{K+1}{\epsilon} = (K+1)^5$, and the last layer has N vertices where N is much larger than K (say $N = O(10^K)$). A node in layer i ($1 \leq i \leq l$) of component B has an outgoing edge to every node in layer $(i+1)$. Component $A = K_{1,Q}$ is a complete bipartite graph with $Q = \frac{Ml+N}{K+1} + 1$.

When the first layer of B is entirely covered with seeds i.e. all the $(K+1)$ seeds are placed on u_1, \dots, u_{K+1} , the eventual number of infected nodes in that component is equal to $K+1+lM+N$. We claim that this case maximizes the eventual number of infections: First observe that if there are less than $(K+1)$ seeds in component B , it is not possible for infections in its last layer to be larger than ϵN : If the seeds are on the first layer, the number of infections in the last layer would be $(\frac{K}{K+1})^{r^l} \times N$ which is less than ϵN due to the way we chose l . Also if the seeds are on a layer other than the first layer of B , then no more than $\frac{K}{M} < \epsilon$ fraction of that layer can be infected which results in a number of infections less than ϵN in the last layer. Therefore it is not possible for number of infections in the last layer of B to be larger than ϵN when there are less than $(K+1)$ initial seeds put on component B . (We note that the claims with regard to infections are ‘approximate’: the actual infections will converge to these approximations, as the layers get large).

Since the total number of infections in component A cannot be greater than the number of nodes in it i.e. $\frac{Ml+N}{K+1} + 2$, the expected total number of infections in G when the first layer of B is not entirely covered with seeds would be less than $\epsilon N + lM + K + \frac{Ml+N}{K+1} + 2$. The later is less than $K+1+lM+N$ once N is chosen sufficiently large. Therefore we can conclude that the expected number of infections is maximized when all the seeds are on the first layer of component B .

Note that the assignment that maximizes the expected number of infections is not an equilibrium since in that case the Red player gets (in expectation) $\frac{Ml+N}{K+1} + 1$ infections, and therefore she is better off by moving her only seed to u where she gets slightly more, i.e. $\frac{Ml+N}{K+1} + 2$ infections in expectation.

Next we proceed to find the equilibrium: we claim that regardless of the strategy of the Blue player, the dominant strategy for the Red player is to put her only seed on u . To prove this we consider all possible cases: Above we discussed that if the Blue player plays entirely on component B , the Red player prefers to place her seed on u . In other cases where the Blue player has a few seeds in A and a few seeds in B , if the Red player puts her seed on u , in the worst case (which is the case where all the Blue seeds are also put on u), she gets in expectation $\frac{Ml+N+K+1}{(K+1)^2}$ infections. Whereas if she plays in component B , what she gets is bounded by the maximum infections that can

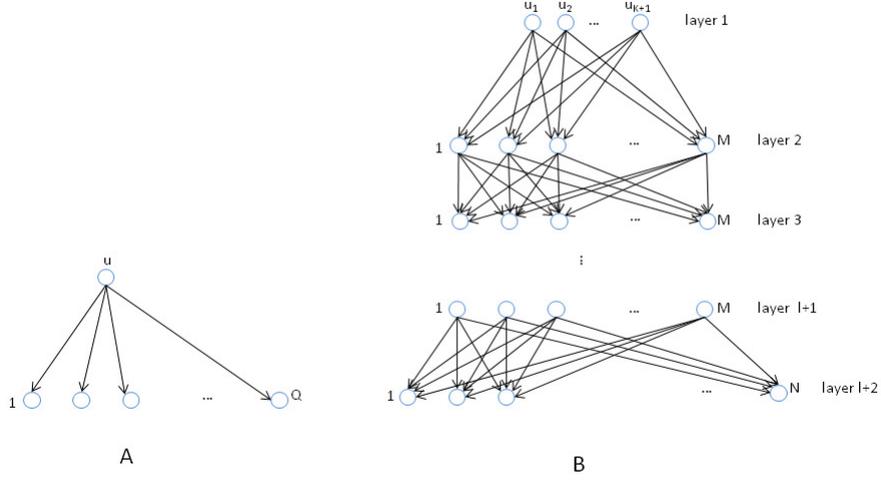


Figure 5: A graph with PoS greater than K

possibly be achieved in B with K seeds which is, as argued before, at most $\epsilon N + lM + K$ (The later is equal to $\frac{N}{(k+1)^4} + lM + K$). Thus when N is much larger than K , we have $\frac{Ml+N+K+1}{(K+1)^2} > \frac{N}{(k+1)^4} + lM + K$, which means the better strategy for the Red player is to stay on u , even in the worst case where there is a heavy competition for getting that node.

Now that we showed selecting u as the seed set is the dominant strategy for the Red player, it is easy to see the best move for the Blue player is to put all her seeds on u as well. Suppose not and the Blue player has $z < K$ seeds on u and $K - z > 0$ seeds in B . For each seed that Blue moves from component B to u , she increase her expected infection in component A by at least $(\frac{K}{K+1} - \frac{K-1}{K}) \times \frac{Ml+N}{(K+1)} \geq \frac{1}{(K+1)^2} \times \frac{Ml+N}{(K+1)}$ (which is equal to $\frac{Ml+N}{(K+1)^3}$), while in the component B the expected number of infections she is losing at most $\epsilon N + lM + K$ (which is equal to $\frac{N}{(k+1)^4} + lM + K$). Since $N \gg K$ it is always beneficial for the Blue player to move her seeds from component B to u . This means that the only equilibrium is the case where all the seeds are placed on u .

As the result of above, the price of stability for this network is at least:

$$\frac{Ml + N + (K + 1)}{\frac{Ml+N+(K+1)}{K+1} + 1}$$

It is very easy to see that by choosing N sufficiently large, the above can get arbitrarily close to $(K + 1)$ which is larger than V . This completes the proof. \square

Combining Theorem 1 and Theorem 2, we note that for $f(x) = x^r$ and linear g we obtain the following sharp threshold result:

Corollary 3. *Let the switching function be $f(x) = x^r$, and let the selection function be linear, $g(y) = y$. Then:*

- *For any $r \leq 1$, the Price of Anarchy (and as a result, Price of Stability) is at most 4 for any graph G ;*
- *For any $r > 1$ and any V , there exists a graph G for which the Price of Stability (and as a result, Price of Anarchy) is greater than V .*

7. Results: Budget Multiplier

We start with some examples to illustrate the issues that arise. We then derive sufficient conditions for bounded Budget Multiplier, and show that violations of these conditions can lead to unbounded Budget Multiplier.

8. Budget Multiplier: Examples

Suppose that budgets of the firms are $K_R = 1$, $K_B = 2$ and the update rule is such that all vertices are updated only once. The network contains 3 influential vertices, each of which has a directed link to all the other $n - 3$ vertices, respectively. So there are $3(n - 3)$ links in all. Let $n \gg 3$.

- Suppose the switching function and selection function are both linear, i.e., $f(x) = x$ and $g(y) = y$. There is a unique equilibrium and in this equilibrium, players will place their resources on distinct influential vertices. The (expected) payoffs to player R are $n/3$, while the payoffs to player B are $2n/3$. So the Budget Multiplier is equal to 1.
- Next, suppose the switching function is convex with $f(2/3) = 1/25$, and the selection function $g(y)$ is as in Tullock (1980). Suppose the two players place their resources on the three influential vertices. The payoffs to R are $g(1/3)n$, while firm B earns $g(2/3)n$. Clearly this is optimal for firm B as any deviation can only lower payoffs. And, it can be checked that a deviation by firm R to one of the influential vertices occupied by player B will yield a payoff of $n/100$ (approximately). So the configuration specified is an equilibrium so long as $g(1/3) \geq 1/100$. The Budget Multiplier is now (approximately) 50.
- Finally, suppose the network consists of ℓ equally-sized connected components. In each component, there is 1 influential vertex which has a directed link to each of the $(n/\ell) - 1$ other vertices. In equilibrium each player locates on distinct influential vertices, *irrespective* of whether the switching function is convex or concave and whether the Tullock selection function is linear ($s = 1$) or whether it is polarizing ($s > 1$). The Budget Multiplier is now equal to 1.

These examples show that for fixed network and updating rule, variations in the switching and selection functions generate large variations in Budget Multiplier. Moreover, for fixed switching and selection functions the payoffs depend crucially on the network.

Theorem 3 provides a set of sufficient conditions on the switching and selection function, under which the Budget Multiplier is uniformly bounded. Theorem 4 shows how even small violations of these conditions can lead to arbitrarily high Budget Multiplier. Theorem 5 illustrates the role of concavity of the switching function in shaping the Budget Multiplier.

8.1. Budget Multiplier: Upper Bound

As in the PoA analysis, it will be technically convenient to return to the generalized adoption function model. Recall that for PoA, competitiveness of h and additivity of H were needed to prove upper bounds, but we didn't require that the implied selection function be linear. Here we introduce that additional requirement, and prove that the (pure strategy) Budget Multiplier is bounded.

Theorem 3. *Suppose the adoption function $h(\alpha^R, \alpha^B)$ is competitive, that H is additive in its arguments, and that the implied selection function is linear:*

$$g(\alpha_R, \alpha_B) = \frac{h(\alpha_R, \alpha_B)}{h(\alpha_R, \alpha_B) + h(\alpha_B, \alpha_B)} = \alpha_R / (\alpha_R + \alpha_B)$$

Then the pure strategy Budget Multiplier is at most 2 for any graph G .¹⁶

Proof The proof borrows elements from the proof of Theorem 1, and introduces the additional notion of tracking or attributing indirect infections generated by the dynamics to specific seeds.

Consider any pure Nash equilibrium given by seed sets S_R and S_B in which $|S_R| = K > |S_B| = L$. For our purposes the interesting case is one in which

$$\mathbf{E}[\chi_R | (S_R, S_B)] \geq \mathbf{E}[\chi_B | (S_R, S_B)]$$

and so

$$\mathbf{E}[\chi_R | (S_R, S_B)] \geq \mathbf{E}[\chi_R + \chi_B | (S_R, S_B)] / 2.$$

Since the adoption function is competitive and additive, Lemma 1 implies that

$$\mathbf{E}[\chi_R | (S_R, \emptyset)] \geq \mathbf{E}[\chi_R | (S_R, S_B)]$$

That is, the Red player only benefits from the departure of the Blue player.

Let us consider the dynamics of the solo Red process given by (S_R, \emptyset) . We first introduce a faithful simulation of these dynamics that also allows us to attribute subsequent infections to exactly one of the seeds in S_R ; we shall call this process the *attribution simulation* of (S_R, \emptyset) . Thus, let $S_R = \{v_1, \dots, v_K\}$ be the initial Red infections, and let us label v_i by R_i , and label all other vertices U . All infections in the process will also be assigned one of the K labels R_i 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 R_1, \dots, R_K , and

¹⁶The theorem actually holds for any equilibrium in which the player with the larger budget plays a pure strategy; the player with smaller budget may always play mixed. It is easy to find cases with such equilibria. The theorem also holds for general mixed strategies under certain conditions — for instance, when both f and g are linear and the larger budget is an integer multiple of the smaller.

with probability $H(\alpha_v^R, 0) = h(\alpha_v^R, 0) + h(0, \alpha_v^R)$ 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 R_i). It is easily seen 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 R_i as a generic Red infection R . Furthermore, at all times every infected vertex has only one of the labels R_i . Thus if we denote the expected number of vertices with label R_i by $\mathbf{E}[\chi_{R_i} | (S_R, \emptyset)]$, we have $\mathbf{E}[\chi_R | (S_R, \emptyset)] = \sum_{i=1}^K \mathbf{E}[\chi_{R_i} | (S_R, \emptyset)]$. Let us assume without loss of generality that the labels R_i are sorted in order of decreasing $\mathbf{E}[\chi_{R_i} | (S_R, \emptyset)]$.

We now consider the payoff to Blue under a deviation from S_B to the set $\hat{S}_B = \{v_1, \dots, v_L\} \subset S_R$ — that is, the L “most profitable” initial infections in S_R . Our goal is to show that the Blue player must enjoy roughly the same payoff from these L seeds as the Red player did in the solo attribution simulation.

Lemma 3.

$$\begin{aligned} \mathbf{E}[\chi_B | (S_R, \hat{S}_B)] &\geq \frac{1}{2} \sum_{i=1}^L \mathbf{E}[\chi_{R_i} | (S_R, \emptyset)] \\ &\geq \frac{L}{2K} \mathbf{E}[\chi_R | (S_R, \emptyset)] \end{aligned}$$

Proof The second inequality follows from

$$\mathbf{E}[\chi_R | (S_R, \emptyset)] = \sum_{i=1}^K \mathbf{E}[\chi_{R_i} | (S_R, \emptyset)],$$

established above, and fact that the vertices in S_R are ordered in decreasing profitability. For the first inequality, we introduce coupled attribution simulations for the two processes (S_R, \emptyset) (the solo Red process) and (S_R, \hat{S}_B) . For simplicity, let us actually examine (S_R, \emptyset) and $(S_R - \hat{S}_B, \hat{S}_B)$; the latter joint process is simply the process (S_R, \hat{S}_B) , but in which the contested seeded nodes in \hat{S}_B are all won by the Blue player. (The proof for the general (S_R, \hat{S}_B) case is the same but causes the factor of $1/2$ in the lemma.)

The coupled attribution dynamics are as follows: as above, in the solo Red process, for $1 \leq i \leq L$, the vertex v_i in S_R is initially labeled R_i , and all other vertices are labeled U . In the joint process, the vertex v_i is labeled B_i for $i \leq L$ (corresponding to the Blue invasions of S_R), while for $L < i \leq K$ the vertex v_i is labeled R_i as before. Now at the first update vertex v , let α_v^R be the fraction of Red neighbors in the solo process, and let $\tilde{\alpha}_v^R$ and $\tilde{\alpha}_v^B$ be the fraction of Red and Blue neighbors, respectively, in the joint process.

Note that initially we have $\alpha_v^R = \tilde{\alpha}_v^R + \tilde{\alpha}_v^B$. Thus by additivity H , the total probabilities of infection $H(\alpha_v^R, 0)$ and $H(\tilde{\alpha}_v^R, \tilde{\alpha}_v^B)$ in the two processes must be identical. We thus flip a common coin with this shared infection probability to determine whether infections occur in the coupled process. If not, v is updated to U in both processes. If so, we now use a coupled attribution step in which we pick an infected neighbor of v at random and copy its label to v in *both* processes. Thus if a label with index $i \leq L$ is chosen, v will be updated to R_i in the solo process, and to B_i in the joint process; whereas if $L < i \leq K$ is chosen, the update will be to R_i in both processes. It is easily verified that each of the two processes faithfully implement the dynamics of the solo and joint attribution processes, respectively.

This coupled update dynamic maintains two invariants: infections are always matched in the two processes, thus maintaining $\alpha_v^R = \tilde{\alpha}_v^R + \tilde{\alpha}_v^B$ for all v and every step; and for all $i \leq L$, every R_i attribution in the solo Red process is matched by a B_i attribution in the joint process, thus establishing the lemma. \square (Lemma 3)

Thus, by simply imitating the strategy of the Red player in the L most profitable resources, the Blue player can expect to infect $(1/2)(L/K)$ proportion of infections accruing to Red in isolation. Since (S_R, S_B) is an equilibrium, the payoffs of Blue in equilibrium must also respect this inequality. \square (Theorem 3)

8.2. Budget Multiplier: Lower Bound

We have already seen that concavity of f and linearity of g lead to bounded PoA and Budget Multiplier, and that even slight deviations from concavity can lead to unbounded PoA. We now show that fixing f to be linear (which is concave), slight deviations from linearity of g towards polarizing g can lead to unbounded Budget Multiplier, for similar reasons as in the PoA case: graph structure can amplify a slightly polarizing g towards arbitrarily high punishment of the minority player.

Theorem 4. *Let the switching function be $f(x) = x$, and let the selection function be of Tullock contest form, $g(y) = y^s / (y^s + ((1-y)^s))$, where $s > 1$. Then for any $V > 0$, there exists a graph G for which the Budget Multiplier is greater than V . More precisely, there is a family of graphs for which the Budget Multiplier grows linearly with the population size (number of vertices).*

Proof As in the PoA lower bound, the proof relies on a layered amplification graph, this time amplifying punishment in the selection function rather than convexity in the switching function. The graph will consist of two components, C_1 and C_2 .

Let us fix the budget of the Red player to be 3, and that of the Blue player to be 1 (the proof generalizes to other unequal values). C_1 is a directed, layered graph with $k + 1$ layers. The first layer has 4 vertices, and layers 2 through k have $n \gg 4$ vertices, while layer $k + 1$ has n_1 vertices, where we shall choose $n_1 \gg n$, meaning that payoffs in C_1 are dominated by infections in the final layer.

The second component C_2 is a 2-layer directed graph, with 1 vertex in the first layer and n_2 in the final layer, and all directed edges from layer 1 to 2. We will eventually choose $n_2 \ll n_1$, so that C_1 is the much bigger component. We choose an update rule in which each layer is updated in succession and only once.

Consider the configuration in which Red places its 3 infections in the first layer of C_1 , and Blue places its 1 infection in the first layer of C_2 . We shall later show that this configuration is a Nash equilibrium. In this configuration, the expected payoff to Red is approximately $\sum_{i=2}^k (3/4)n + (3/4)n_1$ by linearity of f ; notice that the selection function does not enter since the players are in disjoint components. Similarly, the expected payoff to Blue is $(n_2 + 1)$. In this configuration, the ratio of Red and Blue expected payoffs is thus at least $(3/4)n_1 / (n_2 + 1) \simeq (3/4)n_1 / n_2$ (as n_1, n_2 are chosen to be very large numbers), whereas the initial budget ratio is $1/3$. So the Budget Multiplier for this configuration is at least $n_1 / (4n_2)$.

We now develop conditions under which this configuration is an equilibrium. It is easy to verify that red is playing a best response. Moving vertices to later layers of C_1 lowers Red's payoff, since $n \gg 4$ and f is linear. Finally, moving infections to invade the first layer of C_2 will lower Red's payoff as long as, say, $(1/4)n_1$ (Red's current payoff per initial infection in the final layer of C_1) exceeds n_2 (the maximum amount Red could get in C_2 by full deviation), or $n_1 \gg 4n_2$.

We now turn to deviations by Blue. Moving the solo Blue initial infection to the second layer of C_2 is clearly a losing proposition. So consider deviations in which Blue moves to vertices in component 1. If he moves to the lone unoccupied vertex in layer 1 of C_1 , his payoff is approximately:

$$\begin{aligned} & \sum_{i=2}^k g^{(i)}(1/4)n + g^{(k+1)}(1/4)n_1 \\ &= \sum_{i=2}^k \frac{(1/4)^{s^i}}{(1/4)^{s^i} + (3/4)^{s^i}} n \\ & \quad + \frac{(1/4)^{s^{k+1}}}{(1/4)^{s^{k+1}} + (3/4)^{s^{k+1}}} n_1 \end{aligned}$$

Similarly, if Blue directly invades a Red vertex, Blue's payoff is approximately

$$\chi = \frac{1}{2} \times \left(\sum_{i=2}^k \frac{(1/3)^{s^i}}{(1/3)^{s^i} + (2/3)^{s^i}} n + \frac{(1/3)^{s^{k+1}}}{(1/3)^{s^{k+1}} + (2/3)^{s^{k+1}}} n_1 \right)$$

Since in both cases Blue's payoff is being exponentially dampened at each successive layer, it is easy to see that the second deviation is more profitable. Finally, Blue may choose a vertex in a later layer of C_1 , but again by $n \gg 4$ and the linearity of f , this will be suboptimal.

Thus as long as we arrange that n_2 — Blue's payoff without deviation — exceeds χ above, we will have ensured that no player has an incentive to deviate from the specified strategy configuration. Let us scale n_1 up as large as necessary to have χ dominated by the term involving n_1 , and now set n_2 to equal that term:

$$n_2 = \frac{(1/3)^{s^{k+1}}}{(1/3)^{s^{k+1}} + (2/3)^{s^{k+1}}} n_1$$

in order to satisfy the equilibrium condition. The ratio n_1/n_2 , which we have already shown above lower bounds the Budget Multiplier, is thus a function that is increasing exponentially in k for any fixed $s > 1$. Thus by choosing k sufficiently large, we can force the Budget Multiplier larger than any chosen value. \square (Theorem 4)

Combining Theorem 3 and Theorem 4, we note that for linear f and Tullock g , we obtain the following sharp threshold result, which is analogous to the PoA result in Corollary 3.

Corollary 4. *Let the switching function f be linear, and let the selection function g be Tullock, $g(y) = y^s/(y^s + (1 - y)^s)$. Then:*

- For $s = 1$, the Budget Multiplier is at most 2 for any graph G ;
- For any $s > 1$ and any V , there exists a graph G for which the Budget Multiplier is greater than V .

In fact, if we permit a slight generalization of our model, in which certain vertices in the graph are “hard-wired” to adopt only one or the other color (so there is no use for the opposing player to seed them), unbounded Budget Multiplier also holds in the Tullock case for $s < 1$ (equalizing). So in this generalization, linearity of g is required for bounded Budget Multiplier.

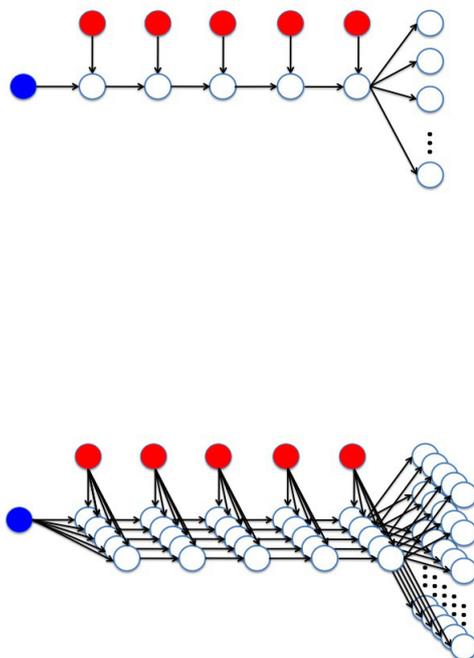


Figure 6: Illustration of the construction in the proof of Theorem 5. Top: Basic gadget. Bottom: Equilibrium construction.

We have thus shown that even when the switching function is “nice” (linear), even slight punishment in the selection function can lead to unbounded Budget Multiplier. Recall that we require switching and selection functions to be 0 (1, respectively) on input 0 (respectively) and increasing, and additionally that $g(1/2) = 1/2$. The following theorem shows that if f is allowed to be a sufficiently convex function, then the Budget Multiplier is again unbounded for *any* selection function. This establishes the importance of concavity of f for both the PoA and Budget Multiplier.

Theorem 5. *Let the switching function f satisfy $f(1/2) = 0$ and $f(1) = 1$. Then for any value $V > 0$, there exists G such that the Budget Multiplier is greater than V .*

Proof Let the Blue player have 1 initial infection and the Red player have $K \geq 2$ (the proof can be generalized to any unequal initial budgets, which we comment on below). Consider the directed graph shown in the top panel of Figure 6, where we have arranged the 1 Blue and K Red seeded nodes in a particular configuration. Aside from the initially infected vertices, this graph consists of a directed chain of K vertices, whose final vertex then connects to a large number $N \gg K$ of terminal vertices. Let us update each vertex in the chain from left to right, followed by the terminal vertices.

Let us first compute the expected payoffs for the two players in this configuration. First, note that since $f(1) = 1$, it is certain that every vertex in the chain will be infected in sequence, followed by all of the terminal vertices; the only question is which player will win the most. By choosing $N \gg K$ we can ignore the infections in the chain and just focus on the terminal vertices, which will be won by whichever player infects the final chain vertex. It is easy to see that the probability this vertex is won by Blue is $1/2^K$, since Blue must “beat” a competing Red infection at every vertex in the chain. Thus the expected payoffs are approximately $N/2^K$ for Blue and $N(1 - 1/2^K)$ for Red. If this configuration were an equilibrium, the Budget Multiplier would thus be $2^K/K$, which can be made as large as desired by choosing K large enough.

However, this configuration is not an equilibrium — clearly, either player would be better off by simplifying initially infecting the final vertex of the chain, thus winning all the terminal vertices. This is fixed by the construction shown in the bottom panel of Figure 6, where we have replicated the chain and terminal vertices M times, but have only the original $K + 1$ seeded nodes as common “inputs” to all of these replications. Notice that now if either player defects to an uninfected vertex, neither player will receive *any* infections in any of the other replications, since now there is a missing “input infection” and reaching the terminal vertices requires all $K + 1$ input infections since $f(1/2) = 0$ (each chain vertex has two inputs, and if either is uninfected, the chain of infections halts). Similarly, if either player attempts to defect by invading the seeded nodes of the other player, there will be no payoff for either player in any of the replications. Thus the most Blue can obtain by deviation is N (moving its one infection to the final chain vertex of a single replication), while the most Red can obtain is KN (moving all of their infections to the final chain vertices of K replications). The equilibrium requirements are thus $M(N/2^K) > N$ for Blue, and $MN(1 - 1/2^K) > KN$ for Red. The Blue requirement is the stronger one, and yields $M > 2^K$. The Budget Multiplier for this configuration is the same as for the single replication case, and thus if we let K be as large as desired and choose $M > 2^K$, we can make the Budget Multiplier exceed any value. \square (Theorem 5)

It is worth noting that even if the Blue player has $L > 1$ seeded nodes, and we repeat the construction above with chain length $K + L - 1$, but with Blue forced to play at the beginning of the chain, followed by all the Red infections, the argument and calculations above are unchanged: effectively, Blues L seeded nodes are no better than 1 infection, because they are simply causing a chain of $L - 1$ Blue infections before then facing the chain of K Red inputs. In fact, even if we let $L \gg K$, Blue’s payoff will still be a factor of $1/2^K$ smaller than Red’s. Thus in some sense the theorem shows that if f is

sufficiently convex, not only is the Budget Multiplier unbounded, but the much *smaller* initial budget may yield arbitrarily *higher* payoffs!

9. Concluding Remarks

We have developed a general framework for the study of competition between firms who use their resources to maximize adoption of their product by consumers located in a social network. This framework yields a very rich class of competitive strategies, which depend in subtle ways on the dynamics, the relative budgets of the players and the structure of the social network. We identified properties of the dynamics of local adoption under which resource use by players is efficient or unboundedly inefficient. Similarly, we identified adoption dynamics for which networks neutralize or dramatically accentuate ex-ante resource difference across players.

There are a number of other questions which can be fruitfully investigated within our framework. One obvious direction is to understand the structure of equilibria in greater detail, and in particular how it is related to network structure. While our results on the PoA and Budget Multiplier demonstrate that network structure can interact in dramatic ways with the switching and selection functions at equilibrium, a more general and detailed understanding would be of interest.

Other interesting directions include algorithmic issues such as computing equilibria and best responses in our framework, and how their difficulty depends on the switching and selection functions; and the multi-stage version of our game, in which the two firms may gradually spend their seed budgets in a way that depends on the evolving state of the network.

References

- [1] J. Coleman, *Medical Innovation: A Diffusion Study*, 2nd Edition, Bobbs-Merrill, 1966, new York.
- [2] T. Conley, C. Udry, Learning about a new technology: Pineapple in ghana, *American Economic Review* (2010) 35–69.
- [3] A. Foster, M. Rosenzweig, Learning by doing and learning from others: Human capital and technical change in agriculture, *Journal of Political Economy* 103 (6) (1995) 1176–1209.
- [4] L. Feick, L. Price, The market maven: A diffuser of marketplace information, *Journal of Marketing* 51 (1) (1987) 83–97.
- [5] P. Reingen, B. Foster, J. Brown, S. Seidman, Brand congruence in interpersonal relations: A social network analysis, *The Journal of Consumer Research* 11 (3) (1984) 771–783.
- [6] D. Godes, D. Mayzlin, Using online conversations to study word of mouth communication, *Marketing Science* 23 (4) (2004) 545–560.
- [7] S. Bharathi, D. Kempe, M. Salek, Competitive influence maximization in social networks, *Internet and Network Economics Lecture Notes in Computer Science* 4858 (2007) 306–311.
- [8] A. Borodin, Y. Filmus, J. Oren, Threshold models for competitive influence in social networks, *Proc. Workshop on Internet and Network Economics (WINE)* (2010) 539–550.
- [9] G. Chasparis, J. Shamma, Control of preferences in social networks, *49th IEEE Conference on Decision and Control* (2010) 6651–6656.
- [10] G. Tullock, *Towards a theory of the rent-seeking society*, Texas A&M University Press, 1980.
- [11] S. Skaperdas, Contest success functions, *Economic Theory* 7 (2) (1996) 283–290.
- [12] G. Butters, Equilibrium distribution of prices and advertising, *Review of Economic Studies* 44 (1977) 465–492.
- [13] G. Grossman, C. Shapiro, Informative advertising with differentiated products, *Review of Economic Studies* 51 (1984) 63–82.
- [14] E. Koutsoupias, C. Papadimitriou, Worst-case equilibria, *STACS* (1999) 404–413.

- [15] T. Roughgarden, E. Tardos, How bad is selfish routing?, Proc. 41st Annual IEEE Symposium on Foundations of Computer Science (FOCS) (2000) 236–259.
- [16] P. Dubey, Inefficiency of nash equilibria, Mathematics of Operations Research 11 (1) (1986) 1–8.
- [17] C. Ballester, A. Calvo-Armengol, Y. Zenou, Who’s who in networks. wanted: The key player, Econometrica 74 (5) (2006) 1403–1417.
- [18] A. Galeotti, S. Goyal, Influencing the influencers: a theory of strategic diffusion, Rand Journal of Economics 40 (3) (2009) 509–532.
- [19] D. Kempe, J. Kleinberg, E. Tardos., Maximizing the spread of influence through a social network, Proc. 9th ACM SIGKDD Intl. Conf. on Knowledge Discovery and Data Mining (2003) 137–146.
- [20] D. Kempe, J. Kleinberg, E. Tardos, Influential nodes in a diffusion model for social networks, Proc. 32nd International Colloquium on Automata, Languages and Programming (ICALP) (2005) 1127–1138.
- [21] E. Mossel, S. Roch, Submodularity of influence in social networks: From local to global, SIAM J. Computing 39 (6) (2007) 2176–2188.
- [22] T. Carnes, C. Nagarajan, S. Wild, A. van Zuylen, Maximizing influence in a competitive social network: A follower’s perspective, Proc. Ninth ACM Conference on Electronic Commerce (EC) (2007) 351–360.
- [23] P. Dubey, R. Garg, B. D. Meyer, Competing for customers in a social network: The quasi-linear case, Proc. Workshop on Internet and Network Economics (WINE) (2006) 162–173.
- [24] A. Vetta, Nash equilibria in competitive societies with applications to facility location, traffic routing and auctions, Proceedings of the 43rd Symposium on the Foundations of Computer Science (FOCS) (2002) 416–425.

Appendix A. Technology adoption examples

This section illustrates a range of applications for our model.

- *Social Network Services (Facebook, Google+, etc.):* Here adoption probabilities might grow slowly with a small fraction of adopting neighbors, since there is little value in using (any) social networking services if none of your friends are using them; thus a convex switching function f ($r > 1$) might be a reasonable model. However, given that it is currently difficult or impossible to export friends and other settings from one service to another, there are strong platform effects in service selection, so a polarizing or even winner-take-all selection function g ($s > 1$) might be most appropriate.
- *Televisions (Sharp, Sony, etc.):* Televisions were immediately useful upon their introduction, without the need for adoption by neighbors, since they allowed immediate access to broadcast programming; the adoption by neighbors serves mainly as a route for information sharing about value of the product. The information value of more neighbors adopting a product is falling with adoption and so a concave f might be appropriate. Compared to social networking services, the platform effects are lower here, and so a linear or equalizing g is appropriate.
- *Mobile Phone Service (Verizon, T-Mobile, etc.):* Mobile phone service was immediately useful upon its introduction without adoption by neighbors, since one could always call land lines, thus arguing for a concave f . Since telephony systems need to be interoperable, platform effects derive mainly from marketing efforts such as “Friends and Family” programs, and thus are extant but perhaps weak, suggesting an equalizing g .

Appendix B. The Multiple Player Version of the Game

Here we generalize the results in Sections 6 and 7 to the multiplayer version where there are $m > 2$ firms/players marketing their products on the network.

We first generalize the dynamics as follows: Let the players be denoted by $1, 2, \dots, m$. When updating an uninfected vertex v , if the fraction of its neighbors who have adopted the product of the i th player is x_i ($1 \leq i \leq m$), then $f(x_1 + x_2 + \dots + x_m)$ specifies the probability with which v becomes infected. In addition $g(\frac{x_i}{x_1 + \dots + x_m})$ specifies the probability that v adopts the i th product.

Suppose S_i is the seed set chosen by the i th player, and χ_i denotes the number of people who eventually adopt product i ($1 \leq i \leq m$). We need the following notations:

Definition 1. $\Pi(S_1, S_2, \dots, S_m)$ denotes the expected number of people who eventually adopt some product, i.e. $\Pi(S_1, S_2, \dots, S_m) = \mathbf{E}[\chi_1 + \dots + \chi_m | (S_1, \dots, S_m)]$.

Definition 2. For each $1 \leq i \leq m$, $\Pi_i(S_1, S_2, \dots, S_m)$ denotes the expected number of people who eventually adopt the product of i 'th player, i.e. $\Pi_i(S_1, S_2, \dots, S_m) = \mathbf{E}[\chi_i | (S_1, \dots, S_m)]$.

Definition 3. For each $1 \leq i \leq m$, $\Pi_{-i}(S_1, S_2, \dots, S_m)$ denotes the expected number of people who eventually adopt some product other than the i 'th product, i.e.

$$\Pi_{-i}(S_1, S_2, \dots, S_m) = \mathbf{E}[\sum_{j \neq i} \chi_j | (S_1, \dots, S_m)]$$

For the m -player version of the game, if f is concave and g is linear, then the following inequalities hold¹⁷:

Lemma 4. $\Pi_1(S_1, \emptyset, \dots, \emptyset) \geq \Pi_1(S_1, S_2, \dots, S_m)$.

Proof This is equivalent to Lemma 1. To prove the claim, we just map the spread process corresponding to (S_1, S_2, \dots, S_m) to the spread process corresponding to (S_R, S_B) in the two-player version of the game, in which we assume $S_R = S_1$ and $S_B = S_1 \cup \dots \cup S_m$. It is straightforward to show that at each time step, $\Pi_R(S_R, S_B) = \Pi_1(S_1, S_2, \dots, S_m)$ and $\Pi_B(S_R, S_B) = \Pi_{-1}(S_1, S_2, \dots, S_m)$. Then by following the same steps as in proof of Lemma 1, we can conclude the desired inequality. \square

Lemma 5. $\Pi_{-1}(S_1, S_2, \dots, S_m) \leq \Pi_{-1}(\emptyset, S_2, \dots, S_m)$.

Proof Similar to the proof of Lemma 4. \square

Lemma 6. $\Pi(T_1, S_2, \dots, S_m) \leq \Pi(S_1, S_2, \dots, S_m)$, where $T_1 \subseteq S_1$.

Proof This follows directly from monotonicity of the switching function. \square

Applying (a series of) the above inequality, we can easily obtain the following two inequalities:

¹⁷We present the inequalities for the first player but similar equalities hold for every other player as well.

Corollary 5. $\Pi(S_1, S_2, \dots, S_m) \geq \Pi(S_1, \emptyset, \dots, \emptyset)$.

Corollary 6. $\Pi(S_1, S_2, \dots, S_m) \geq \Pi(\emptyset, S_2, \dots, S_m)$.

We are now ready to show that for the multiplayer version of the game, there is an $O(m)$ upper bound on the PoA, when the switching function is concave and the selection function is linear:

Theorem 6. *When there are m players ($m > 2$) competing, if f is concave and g is linear, then the PoA is less than or equal to $2m$.*

Proof Suppose not, and there is a graph G in which PoA is greater than $2m$. Let (S_1, S_2, \dots, S_m) denote the strategy profile of players in the worst equilibrium, and $(S_1^*, S_2^*, \dots, S_m^*)$ denote the strategy that maximizes the total number of infections. Since PoA is greater than $2m$, we have:

$$\frac{1}{2m} \Pi(S_1^*, S_2^*, \dots, S_m^*) > \Pi(S_1, S_2, \dots, S_m) \quad (\text{B.1})$$

which is equivalent to:

$$\frac{1}{m} \Pi(S_1^*, S_2^*, \dots, S_m^*) > 2\Pi(S_1, S_2, \dots, S_m) \quad (\text{B.2})$$

In addition, according to pigeon hole principle, there is a player i such that

$$\Pi_i(S_1^*, S_2^*, \dots, S_m^*) \geq \frac{1}{m} \Pi(S_1^*, S_2^*, \dots, S_m^*)$$

Without loss of generality, assume $i = 1$ and we have:

$$\Pi_1(S_1^*, S_2^*, \dots, S_m^*) \geq \frac{1}{m} \Pi(S_1^*, S_2^*, \dots, S_m^*) \quad (\text{B.3})$$

According to lemma 4:

$$\Pi_1(S_1^*, \emptyset, \dots, \emptyset) \geq \Pi_1(S_1^*, S_2^*, \dots, S_m^*) \quad (\text{B.4})$$

So combining the above with inequality B.3 we have:

$$\Pi_1(S_1^*, \emptyset, \dots, \emptyset) \geq \frac{1}{m} \Pi(S_1^*, S_2^*, \dots, S_m^*) \quad (\text{B.5})$$

Now we replace the strategy of the first player in the worst equilibrium, i.e. S_1 , with S_1^* and show that $\Pi_1(S_1^*, S_2, \dots, S_m) > \Pi_1(S_1, S_2, \dots, S_m)$. This means that S_1 is not the best response to S_{-1} , which contradicts with the assumption of being in an equilibrium, and it finishes the proof.

We just need to prove that the above inequality holds.

According to corollary 5, we have $\Pi(S_1^*, S_2, \dots, S_m) \geq \Pi(S_1^*, \emptyset, \dots, \emptyset)$. Combining this with inequality B.5 we get $\Pi(S_1^*, S_2, \dots, S_m) \geq \frac{1}{m} \Pi(S_1^*, S_2^*, \dots, S_m^*)$. Then combining this last inequality with inequality B.2, we have:

$$\Pi(S_1^*, S_2, \dots, S_m) > 2\Pi(S_1, S_2, \dots, S_m) \quad (\text{B.6})$$

or equivalently:

$$\frac{1}{2}\Pi(S_1^*, S_2, \dots, S_m) > \Pi(S_1, S_2, \dots, S_m) \quad (\text{B.7})$$

On the other hand, according to corollary 6, $\Pi(\emptyset, S_2, \dots, S_m) < \Pi(S_1, S_2, \dots, S_m)$, and therefore combining with equation B.7 we get:

$$\frac{1}{2}\Pi(S_1^*, S_2, \dots, S_m) > \Pi(\emptyset, S_2, \dots, S_m) \quad (\text{B.8})$$

According to lemma 5, we know that $\Pi_{-1}(S_1^*, S_2, \dots, S_m) \leq \Pi_{-1}(\emptyset, S_2, \dots, S_m)$. Based on the fact that $\Pi(\emptyset, S_2, \dots, S_m) = \Pi_{-1}(\emptyset, S_2, \dots, S_m)$, we can write it as

$$\Pi_{-1}(S_1^*, S_2, \dots, S_m) \leq \Pi(\emptyset, S_2, \dots, S_m)$$

Now combining this last inequality with B.8 we get:

$$\frac{1}{2}\Pi(S_1^*, S_2, \dots, S_m) > \Pi_{-1}(S_1^*, S_2, \dots, S_m) \quad (\text{B.9})$$

Since by definition $\Pi(S_1^*, S_2, \dots, S_m) = \Pi_{-1}(S_1^*, S_2, \dots, S_m) + \Pi_1(S_1^*, S_2, \dots, S_m)$, and by inequality B.8 we have $\frac{1}{2}\Pi(S_1^*, S_2, \dots, S_m) > \Pi_{-1}(S_1^*, S_2, \dots, S_m)$, the following holds:

$$\Pi_1(S_1^*, S_2, \dots, S_m) > \frac{1}{2}\Pi(S_1^*, S_2, \dots, S_m). \quad (\text{B.10})$$

According to inequality B.7 we have that $\Pi_1(S_1^*, S_2, \dots, S_m) > \Pi(S_1, S_2, \dots, S_m)$. Since it is obvious that $\Pi(S_1, S_2, \dots, S_m) \geq \Pi_1(S_1, S_2, \dots, S_m)$, we can conclude that $\Pi_1(S_1^*, S_2, \dots, S_m) > \Pi_1(S_1, S_2, \dots, S_m)$ as desired. \square

It is easy to see that the rest of the results can also be generalized to the multiplayer version. Since the proofs are straightforward, we just state the results here:

Theorem 7. *When there are m players ($m > 2$) competing, if f is slightly convex and g is linear, then the PoS can be unbounded.*

To accommodate the definition of the Budget Multiplier for the case of more than 2 players, we change the definition as follows: For any fixed graph, local dynamics, and initial budgets, with $K_1 \geq K_2 \geq \dots \geq K_m$, let $(\sigma_1, \sigma_2, \dots, \sigma_m)$ be the Nash equilibrium that *maximizes* the quantity

$$\frac{\Pi_i(\sigma_1, \sigma_2, \dots, \sigma_m)}{\Pi_j(\sigma_1, \sigma_2, \dots, \sigma_m)} \times \frac{K_j}{K_i}$$

among all Nash equilibria $(\sigma_1, \sigma_2, \dots, \sigma_m)$ and all choices of $1 \leq i < j \leq m$; The resulting maximized quantity is the Budget Multiplier.

Theorem 8. *When there are m players ($m > 2$) competing, if f is concave and g is linear, then the pure strategy Budget Multiplier is bounded by 2 across all networks.*

Theorem 9. *When there are m players ($m > 2$) competing, if f is slightly convex and g is linear, then the Budget Multiplier can be unbounded.*