Competing for Customers in a Social Network: The Quasi-linear Case

Pradeep Dubey[1,2], Rahul Garg[3], and Bernard De Meyer[4,2]

[1] Center for Game Theory, Stony Brook University, USA
[2] Cowles Foundation, Yale University, USA
[3] IBM India Research Lab, New Delhi, India
[4] Cermsem, Univesité Paris 1, Paris, France

Abstract. There are many situations in which a customer's proclivity to buy the product of any firm depends not only on the classical attributes of the product such as its price and quality, but also on who else is buying the same product. We model these situations as games in which firms compete for customers located in a "social network". Nash Equilibrium (NE) in pure strategies exist and are unique. Indeed there are closed-form formulae for the NE in terms of the exogenous parameters of the model, which enables us to compute NE in polynomial time.

An important structural feature of NE is that, if there are no a priori biases between customers and firms, then there is a cut-off level above which high cost firms are blockaded at an NE, while the rest compete *uniformly* throughout the network.

We finally explore the relation between the connectivity of a customer and the money firms spend on him. This relation becomes particularly transparent when externalities are dominant: NE can be characterized in terms of the invariant measures on the recurrent classes of the Markov chain underlying the social network.

1 Introduction

Consider a situation in which firms compete for customers located in a "social network". Any customer i has, of course, a higher proclivity to buy from firm α, if α lowers its price relative to those quoted by its rivals. But another, quite independent, consideration also influences i's decision. He is keen to conform to his neighbors in the network. If the bulk of them purchase firm β's product, then he is tempted to do likewise, even though β may be charging a higher price than α. Customer i's behavior thus involves a delicate balance between the "externality" exerted by his neighbors and the more classical constituents of demand — the price and the intrinsic quality of the product itself. Such externalities arise naturally in several contexts (see, e.g., [1],[5],[6],[3],[8],[7]).

The externality in demand clearly has significant impact on the strategic interaction between the firms. Firm α may spend resources marketing its product to i, not because α cares about i per se as a client, but because i enjoys the position of a "hub" in the social network and so wields influence on other potential clients that are of value to α. This in turn might instigate rival firms to

P. Spirakis et al. (Eds.): WINE 2006, LNCS 4286, pp. 162–173, 2006.

spend further on i, since they wish to wean i away from an excessive tilt toward α; causing α to increase its outlay on i even more, unleashing yet another round of incremental expenditures on i.

The scenario invites us to model it as a non-cooperative game between the firms[1]. We take our cue from [1],[5] which explore the optimal marketing strategy of a *single* firm, based on the "network value" of the customers. Our innovation is to introduce competition between *several* firms in this setting. The model we present is more general than that of [1],[5], though inspired by it. As in [1],[5], the social network, specifying the field of influence of each customer, is taken to be exogenous. Rival firms choose how much money to spend on each customer. For any profile of firms' strategies, we show that the externality effect stabilizes over the social network and leads to unambiguous customer-purchases. A particular instance of our game arises when firms compete for advertisement space on different web-pages in the Internet (see Section 2.1).

Our main interest is in understanding the structure of the Nash Equilibria (NE) of the game between the firms. Will they end up as regional monopolies, operating in separate parts of the network? Or will they compete fiercely throughout? Which firms will enter the fray, and which will be blockaded? And how will the money spent on a customer depend on his connectivity in the social network?

For ease of presentation, the focus of this paper is on the quasi-linear[2] case (which includes the model in [1], by setting # firms = 1). We show that NE are unique and can be computed in polynomial time via closed-form expressions involving matrix inverses. It turns out that, provided that there are no a priori biases between firms and customers, any NE has a cut-off cost: all firms whose costs are above the cut-off are blockaded, and the rest enter the fray. Moreover there is no "regionalization" of firms in an NE: each active firm spends money on *every* customer-node of the social network. The money spent on node i is related to the connectivity of i, but the relation is somewhat subtle, though expressible in precise algebraic form. When externalities are dominant, however, this relation becomes more transparent: NE can be characterized in terms of the invariant measures on the recurrent classes of the Markov chain underlying the social network (see Section 4). In particular suppose that the graph representing the social network is undirected and connected, all the neighbors of any customer-node exert equal influence on him, and each company values all the nodes equally. Then, at the NE, the money spent by a company on a node is proportional to the degree of the node.

2 The Model

There is a finite set $\mathcal{A}$ of firms and $\mathcal{I}$ of customers. We shall define a strategic game Γ among the firms. The customers themselves are non-strategic in our model and described in behavioristic terms.

[1] Customers are not strategic in our model. As in [1],[5], they are described in behavioristic terms.

[2] For generalizations to the non-linear case, see Section 4.

Firm $\alpha \in \mathcal{A}$ can spend m_i^α dollars on customer $i \in \mathcal{I}$ by way of marketing its product to him. This could represent the discounts or special warranties offered by α to i (in effect lowering, for i, the fixed price that α has quoted for its product), or free add-ons of supplementary products, or simply the money spent on advertising to i, etc. The strategy set of firm α may thus be viewed as[3] $R_+^{\mathcal{I}}$, with elements $m^\alpha \equiv (m_i^\alpha)_{i\in\mathcal{I}}$.

Consider a profile of firms' strategies $m \equiv (m^\alpha)_{\alpha\in\mathcal{A}} \in R_+^{\mathcal{I}\times\mathcal{A}}$. The *proclivity* of customer i to buy from any particular firm α clearly depends on the profile m, i.e., not just the expenditure of α but also that of its rivals. We denote this proclivity by $p_i^\alpha(m)$. One can think of $p_i^\alpha(m)$ as the quantity of α's product purchased by i. Or, interpreting i to be a mass of customers such as those who visit a web page i, one can think of $p_i^\alpha(m)$ as the fraction of mass i that goes to α (or, equivalently, as the probability of i going to α). In either setting, we take $p_i(m) \equiv (p_i^\alpha(m))_{\alpha\in\mathcal{A}} \in [0,1]^{\mathcal{A}}$. (When $p_i^\alpha(m)$ is a quantity, there is a physical upper bound on customer i's capacity to consume which, w.l.o.g., is normalized to be 1).

The benefit to any particular firm α from its *clientele* $p^\alpha(m) \equiv (p_i^\alpha(m))_{i\in\mathcal{I}}$ is $\sum_{i\in\mathcal{I}} u_i^\alpha p_i^\alpha(m)$ and the cost of its expenditures m^α is $\sum_{i\in\mathcal{I}} c_i^\alpha m_i^\alpha$.

Thus α's *payoff* in the game is given by

$$\Pi^\alpha(m) = \sum_{i\in\mathcal{I}} u_i^\alpha p_i^\alpha(m) - \sum_{i\in\mathcal{I}} c_i^\alpha m_i^\alpha$$

It remains to define the map from m to $p(m)$.

Customer i's proclivity p_i^α to purchase from firm α is clearly positively correlated with α's expenditure m_i^α on i, and negatively correlated with the expenditures $m_i^{-\alpha} \equiv (m_i^\beta)_{\beta\in\mathcal{I}\setminus\{\alpha\}}$, of α's rivals.

In addition we suppose that there is a positive externality exerted on i by the choice of any neighbor j: increases in p_j^α may boost p_i^α. Negative cross-effects of p_j^β on p_i^α, for $\beta \neq \alpha$, can be incorporated under certain assumptions (which we make precise in [2]), but here we suppose that they are absent.

By way of an example of such an externality, think of firms' products as specialized software. Then if the users with whom i frequently interfaces (i.e., i's "neighbors") have opted for α's software, it will suit i to also purchase predominantly from α in order to more smoothly interact with them. Or else suppose the firms are in an industry focused on some fashion product. Denote by i's neighbors the members of i's peer group with whom i is eager to conform. Once again, p_i^α is positively correlated with p_j^α where j is a neighbor of i. Another typical instance comes from telephony: if most of the people, who i calls, subscribe to service provider α and if α-to-α calls have superior connectivity compared with α-to-β calls, then i may have incentive to subscribe to α even if α is costlier than β.

To define the map from m to $p(m)$, we must turn to the social network. It is represented by a directed, weighted graph $G = (\mathcal{I}, E, w)$. The nodes of G are

[3] Budget constraints on expenditures can be incorporated via cost functions (see Section 4).

identified with the set of customers $\mathcal{I}$. Each directed edge $(i, j) \in E \equiv \mathcal{I} \times \mathcal{I}$ has weights $(w^\alpha_{ij})_{\alpha \in \mathcal{A}}$, where $w^\alpha_{ij} \geq 0$ is a measure of the influence j has on i, with regard to purchases from α. Precisely, if $p^\alpha = (p^\alpha_j)_{j \in \mathcal{I}}$ denotes the proclivities of purchases, then the externality impact of p^α on i is $\sum_{j \in \mathcal{I}} w^\alpha_{ij} p^\alpha_j$. We assume that $\sum_{j \in \mathcal{I}} w^\alpha_{ij} \leq 1$, for all $i \in \mathcal{I}$ and $\alpha \in \mathcal{A}$. (One may view $(\mathcal{I}, E^\alpha, w^\alpha)$ as the social network relevant for firm α, with $E^\alpha = \{(i, j) \in E : w^\alpha_{ij} > 0\}$).

Let us now make explicit how firms' expenditures, in conjunction with the externality effect, determine purchases in the social network.

Fix a profile $m \equiv (m^\beta)_{\beta \in \mathcal{A}} \equiv ((m^\beta_j)_{j \in \mathcal{I}})_{\beta \in \mathcal{A}}$ of firms' strategies.

For any firm α and customer i, let $\gamma^\alpha_i(m_i) \in [0, 1]$ denote the proclivity with which i is *initially* impelled to buy from firm α on account of the direct "marketing impact", where (recall) $m_i \equiv (m^\beta_i)_{\beta \in \mathcal{A}}$ gives the expenditures induced on i by m.

Denoting $(m^\beta_i)_{\beta \in \mathcal{A} \setminus \{\alpha\}}$ by $m^{-\alpha}_i$, it stands to reason that the impact $\gamma^\alpha_i(m^\alpha_i, m^{-\alpha}_i)$ be strictly increasing in m^α_i for any fixed $m^{-\alpha}_i$. We assume this and a little bit more: γ^α_i is also concave in m^α_i for fixed $m^{-\alpha}_i$, reflecting the diminishing returns to α of incremental dollars spent on i.

A canonical example we have in mind is $\gamma^\alpha_i(m_i) = m^\alpha_i / \overline{m}_i$ where $\overline{m}_i \equiv (\sum_{\beta \in \mathcal{I}} m^\beta_i)$ (with $\gamma^\alpha_i(0) \equiv 0$). In short, i's probability of purchase from different firms is simply set proportional to the money they spend on him[4].

Customer i weights the two factors (i.e., the externality impact and the marketing impact) by θ^α_i and $1 - \theta^\alpha_i$, where $0 \leq \theta^\alpha_i < 1$. Thus, given a strategy profile m, the final steady-state proclivities of purchase $p(m) \equiv (p^\alpha(m))_{\alpha \in \mathcal{A}} \in [0, 1]^{\mathcal{I} \times \mathcal{A}}$, where $p^\alpha \equiv (p^\alpha_j(m))_{j \in \mathcal{I}}$, must satisfy.

$$p^\alpha_i(m) = (1 - \theta^\alpha_i)\gamma^\alpha_i(m_i) + \theta^\alpha_i \sum_{j \in \mathcal{I}} w^\alpha_{ij} p^\alpha_j(m) \tag{1}$$

for all $\alpha \in \mathcal{A}$ and $i \in \mathcal{I}$.

Define the $|\mathcal{I}| \times |\mathcal{I}|$-matrices: $I \equiv$ identity, $\Theta^\alpha \equiv$ the diagonal matrix with $\Theta^\alpha_{ii} = \theta^\alpha_i$ and $W^\alpha \equiv$ the matrix with entries w^α_{ij}. Then equation (1) reads

$$p^\alpha(m) = (I - \Theta^\alpha)\gamma^\alpha(m) + \Theta^\alpha W^\alpha p^\alpha(m).$$

Since $I - \Theta^\alpha W^\alpha$ is invertible (its row sums being less than 1), we obtain

$$p^\alpha(m) = (I - \Theta^\alpha W^\alpha)^{-1}(I - \Theta^\alpha)\gamma^\alpha(m).$$

This gives

$$\Pi^\alpha(m) = [u^\alpha]^\top (I - \Theta^\alpha W^\alpha)^{-1}(I - \Theta^\alpha)\gamma^\alpha(m) - [c^\alpha]^\top m^\alpha \tag{2}$$

---

[4] More generally, $\gamma^\alpha_i(m_i) = (m^\alpha_i / \overline{m}_i)(\overline{m}_i)^r$ where $0 \leq r < 1$. We may think of $(\overline{m}_i)^r$ as the "market penetration", which rises with the total money spent. (If $\gamma^\alpha_i(m_i)$ is to be a probability, one must amend $(\overline{m}_i)^r$ to $\max\{(\overline{m}^r_i), 1\}$ or a suitably smoothed version of this function.)

where $u^\alpha \equiv (u_j^\alpha)_{j\in\mathcal{I}} \in R_+^{\mathcal{I}}$ and $c^\alpha \equiv (c_j^\alpha)_{j\in\mathcal{I}} \in R_{++}^{\mathcal{I}}$ are column vectors and $\top$ stands for the transpose operation. Denote

$$v^\alpha \equiv [u^\alpha]^\top (I - \Theta^\alpha W^\alpha)^{-1}(I - \Theta^\alpha) \tag{3}$$

Then (2) may be rewritten:

$$\Pi^\alpha(m) = \sum_{i\in\mathcal{I}} (v_i^\alpha \gamma_i^\alpha(m_i) - c_i^\alpha m_i^\alpha) \tag{4}$$

Our key assumption on $\gamma_i^\alpha(m_i)$ is that it depends only on the variables m_i^α and $\overline{m}_i^{-\alpha} \equiv \sum_{\beta\in\mathcal{A}\setminus\{\alpha\}} m_i^\beta$, i.e., firm α is affected only by the aggregate[5] expenditure of its rivals.

Assume $\gamma_i^\alpha(m_i^\alpha, \overline{m}_i^{-\alpha})$ is continuous; and, furthermore, increasing and differentiable w.r.t. m_i^α whenever $\overline{m}_i \equiv \sum_{\beta\in\mathcal{A}} m_i^\beta = m_i^\alpha + \overline{m}_i^{-\alpha} > 0$. Let

$$\phi_i^\alpha(m_i^\alpha, \overline{m}_i^{-\alpha}) \equiv \frac{\partial}{\partial m_i^\alpha} \gamma_i^\alpha(m_i^\alpha, \overline{m}_i^{-\alpha})$$

and next define

$$\lambda_i^\alpha(r_i^\alpha, \overline{m}_i) \equiv \phi_i^\alpha(r_i^\alpha \overline{m}_i, (1 - r_i^\alpha)\overline{m}_i)$$

(Thus $r_i^\alpha \equiv m_i^\alpha / \overline{m}_i$.) We suppose that

$$\lambda_i^\alpha \text{ is strictly decreasing in } r_i^\alpha \text{ and in } \overline{m}_i \tag{5}$$

for fixed $\overline{m}_i$ and r_i^α respectively. This condition reflects the diminishing returns on incremental dollars spent by α; it also states that an incremental dollar of α counts for less when α's rivals have put in more money.

We also assume that

$$\lim_{\delta\to 0} \frac{\gamma_i^\alpha(\delta, 0)}{\delta} = \infty. \tag{6}$$

Note that both conditions (5) and (6) are satisfied by our canonical example and its variants in footnote 4.

Finally we assume that for each customer there exist at least two firms that value him:

$$\forall i \in \mathcal{I}, \exists \alpha, \alpha' \in \mathcal{A} \text{ such that } : \alpha \neq \alpha' \text{ and } u_i^\alpha > 0 \text{ and } u_i^{\alpha'} > 0. \tag{7}$$

This will create enough competition in an NE to ensure that positive money is bid on each client, enabling us to steer clear of possible discontinuity[6] of γ_i^α at 0.

[5] Aggregation is a form of anonymity that is common to many markets. It says, in essence, that if a firm pretends to be two entities and splits its expenditure between them, this has no effect on *other* firms. This form of "anonymity toward numbers" is tantamount to aggregation.

[6] As occurs in our canonical example.

2.1 An Example: Competition for Advertisement on the Web

Think of the web as a set $\mathcal{I}$ of pages, each of which corresponds to a distinct node of a graph. A directed arc (i, j) means that there is a link from page j to page i.

At the beginning of any period, two kind of "surfers" visit page i. There are those who transit to i from other pages j in the web. Furthermore, there are "fresh arrivals", entering the web for the first time, via page i at rate ψ_i.

At the end of the period, a fraction $(1 - \theta_i)$ of the population on the page i exits the web, while the remaining fraction θ_i continues surfing (where $0 \leq \theta_i < 1$). The weight on (i, j), which we denote ω_{ij}, gives the probability that a representative surfer, who is on page j and who continues surfing, moves on to page i (or, alternatively, the fraction of surfers on page j who transit to page i). Thus $\sum_{i\in\mathcal{I}} \omega_{ij} = 1$ for all $j \in \mathcal{I}$.

Companies $\alpha \in \mathcal{A}$ compete for advertisement on the web pages. If they spend $m_i \equiv (m_i^\alpha)_{\alpha\in\mathcal{A}}$ dollars to place their ads on page i, they get "visibility" (time, space) on page i in proportion to the money spent. Thus the probability that a surfer views company α's ad on page i is $m_i^\alpha/\overline{m}_i = \gamma_i^\alpha(m_i^\alpha, \overline{m}_i)$.

The payoff of a company is the aggregate "eyeballs" of its advertisement obtained, in the long run (i.e., in the steady state).

To compute the payoff, let us first examine the population distribution of surfers across nodes in the unique steady state of the system.

Denote by ϕ_i denote the arrival rate of surfers (of both kinds) to page i. Then, in a steady state, we must have

$$\phi_i = \psi_i + \sum_{j\in\mathcal{I}} \omega_{ij}\theta_j\phi_j$$

for all $i \in \mathcal{I}$. In matrix notation, this is

$$\phi = \psi + \Omega\Theta\phi$$

where $\phi \equiv (\phi_i)_{i\in\mathcal{I}}$ and $\psi \equiv (\psi_i)_{i\in\mathcal{I}}$ are column vectors, Θ is the diagonal $\mathcal{I} \times \mathcal{I}$ matrix with entries $\theta_{ii} = \theta_i$, and Ω is the $\mathcal{I} \times \mathcal{I}$ matrix with entries ω_{ij}. Hence

$$\phi = (I - \Omega\Theta)^{-1}\psi$$

The total eyeballs (per period) obtained by company α is then

$$\sum_{i\in\mathcal{I}} \phi_i\gamma_i^\alpha(m)$$

which fits the format of (4).

More generally, suppose surfers have bounded recall of length k. Then firm α will only care about any surfer's eyeballs in the last k periods prior to the surfer's exit. When $k = 1$, α's payoff is

$$\sum_{i\in\mathcal{I}} (1 - \theta_i)\phi_i\gamma_i^\alpha(m)$$

The expression for v_i^α will become complicated when the recall $k > 1$ (more so, if discounting of past memory is incorporated). But the payoffs in all these cases still fit the format of (4).

Generalizing in a different direction, suppose that surfers at page i, who have spent t periods in the web, exit at rate θ_i^t for $t = 1, 2 \ldots$. Denote by Θ^t the diagonal matrix whose ii^{th} entry is θ_i^t. Then $\phi = (I + \Omega\Theta^1 + \Omega\Theta^2\Omega\Theta^1 + \ldots)\psi$, which is well-defined provided we assume $\theta_i^t \leq \Delta < 1$ for some Δ (for all t, i). This retains the format of (4) though the expression for v_i^α becomes even more complicated. One could also incorporate bounded recall in this setting, without departing from (4).

Notice that the "externality" in the above examples is reflected in the movement of traffic across pages in the web. Also notice that the games derived are *anonymous* i.e. $v_i^\alpha = v_i$ for all α. Such games will be singled out for special attention later.

2.2 Uniqueness of Nash Equilibrium

Recall that a strategy profile m is called a **Nash Equilibrium**[7] (NE) of the game Γ if

$$\Pi^\alpha(m) \geq \Pi^\alpha(\tilde{m}^\alpha, m^{-\alpha}) \ \ \forall\, \tilde{m}^\alpha \in R_+^{\mathcal{I}}$$

for all $\alpha \in \mathcal{A}$ (where $m^{-\alpha} \equiv (m^\beta)_{\beta \in \mathcal{I}\setminus\{\alpha\}}$).

Theorem 1. *Under hypotheses (5), (6), (7), there exists a unique Nash Equilibrium in the quasi-linear model.*

Proof: See [2].

2.3 Characterization of Nash Equilibrium

Theorem 2. *Consider our canonical case: $\gamma_i^\alpha(m_i) = m_i^\alpha/\overline{m}_i$ (other closed-form expressions for the γ_i^α will lead to analogous characterizations). Fix customer i and rank all the firms in $\mathcal{A} \equiv \{1, 2, \ldots, n\}$ in order of increasing $\kappa_i^\alpha \equiv c_i^\alpha/v_i^\alpha$ (see (3) for the definition of v_i^α). For convenience denote this order $\kappa_i^1 \leq \kappa_i^2 \leq \ldots \leq \kappa_i^n$. Let*

$$k_i = \max\left\{ l \in \{2, \ldots, n\} : (l-2)\kappa_i^l < \sum_{\alpha=1}^{l-1} \kappa_i^\alpha \right\} \tag{8}$$

In the unique NE, firms $1, \ldots, k_i$ will spend money on customer i as follows:

$$m_i^\alpha = \left(\frac{k_i - 1}{\sum_{\beta=1}^{k_i} \kappa_i^\beta} \right) \left(1 - \frac{(k_i - 1)\kappa_i^\alpha}{\sum_{\beta=1}^{k_i} \kappa_i^\beta} \right) \tag{9}$$

Firms $k_i + 1, \ldots, n$ put no money on customer i.

---

[7] Throughout we confine attention to "pure" strategies.

Proof: See [2].

According to Theorem 2, companies α can be ranked, at each customer-node i, according to their "effective costs" κ_i^α. The money m_i^α, spent by α on i, is a strictly decreasing function of κ_i^α upto some threshold, after which it becomes zero.

Theorem 2 confirms the obvious intuition that $m_i^\alpha = 0$ if $v_i^\alpha = 0$ (i.e., $\kappa_i^\alpha = \infty$, recalling that $c_i^\alpha > 0$ by assumption). It also brings to light a different, and more important, feature of NE. First recall that, by (3), v_i^α may well be highly positive even though the direct value u_i^α of customer i to company α is zero. This is because v_i^α incorporates the *network value* of i, stemming from the possibility that i may be exerting a big externality on other customers whom α does directly value. Now, since κ_i^α falls with v_i^α, (9) reveals that α may be spending a huge m_i^α on i even when u_i^α is zero, *purely* on account of the network value of i.

2.4 Impact of the Social Network on Nash Equilibrium

To get a better feel for Theorem 2, it might help to consider some examples.

Suppose there are five customers $\{1, 2, \ldots, 5\}$ and four firms $\{\alpha_1, \alpha_2, \beta_1, \beta_2\}$. The customers are arranged in a linear network, with i connected to $i+1$ via an undirected (i.e., directed both ways) edge, for $i = 1, 2, 3, 4$. Suppose each node is equally influenced by its neighbors in the purchase of any firm's product. Thus $(w_{11}^\gamma, w_{12}^\gamma, w_{13}^\gamma, w_{14}^\gamma, w_{15}^\gamma) = (0,\ 1,\ 0,\ 0,\ 0)$, $(w_{21}^\gamma, w_{22}^\gamma, w_{23}^\gamma, w_{24}^\gamma, w_{25}^\gamma) = (0.5,\ 0,\ 0.5,\ 0,\ 0)$ etc., for any company γ. Further suppose $\theta_i^\gamma = 0.1$ and $c_i^\gamma = 1$ for all γ and i. Finally let $u^{\alpha_1} = u^{\alpha_2} = (1,\ 1,\ 0,\ 0.1,\ 0.1)$ and $u^{\beta_1} = u^{\beta_2} = (0.1,\ 0.1,\ 0,\ 1,\ 1)$. Formula (3) yields $v^{\alpha_1} = v^{\alpha_2} = (0.950,\ 0.998,\ 0.055,\ 0.102,\ 0.095)$ and $v^{\beta_1} = v^{\beta_2} = (0.095,\ 0.102,\ 0.055,\ 0.998,\ 0.950)$ and hence $\kappa^{\alpha_1} = \kappa^{\alpha_2} = (1.053,\ 1.002,\ 18.182,\ 9.779,\ 10.514)$ and $\kappa^{\beta_1} = \kappa^{\beta_2} = (10.514,\ 9.779,\ 18.182,\ 1.002,\ 1.053)$. It follows from Theorem 2 that firms α_1 and α_2 will put no money on customers $4, 5$ and positive money on the rest; while firms β_1 and β_2 will put no money on customers $1, 2$ and positive money on the rest. In effect, there will "regionalization" of customers into α-territory $\{1, 2, 3\}$ and β-territory $\{3, 4, 5\}$. The only overlap is customer 3, who is of zero direct value u_3^γ to all firms γ and yet is being equally targeted by them, purely on account of his network value.

The situation dramatically changes when the game is anonymous i.e., $v_i^\alpha = v_i$ and $c_i^\alpha = c^\alpha$ for all α and i. (The first identity holds in particular — see (3) — when $w_{ij}^\alpha = w_{ij}$, $\theta_i^\alpha = \theta_i$, and $u_i^\alpha = u_i$, for all α, i and j, i.e., there are no a priori biases between firms and customers.) Our analysis in Section 2.3 immediately implies that we can rank the firms, independently of i, by their costs; say (after relabeling)

$$c^1 \leq c^2 \leq \ldots \leq c^n$$

At the Nash Equilibrium a subset of low-cost firms $\{1, \ldots, k\}$ will be active (see (8), while all the higher-cost firms $\{k+1, \ldots, n\}$ will be blockaded, where

$$k = \max \left\{ l \in \{2, \ldots, n\} : (l-2)c^l < \sum_{\beta=1}^{l-1} c^\beta \right\}$$

Each active firm $\alpha \in \{1, \ldots, k\}$ will spend an amount $m_i^\alpha > 0$ on all the nodes $i \in \mathcal{I}$ that is proportional to v_i. Indeed, by (9), we have

$$m_i^\alpha = \frac{v_i(k-1)}{\sum_{\beta=1}^k c^\beta}\left(1 - \frac{(k-1)c^\alpha}{\sum_{\beta=1}^k c^\beta}\right)$$

which also shows that $\overline{m}^\alpha \geq \overline{m}^\beta$ if $\alpha < \beta$, i.e., lower cost firms spend more money than their higher-cost rivals. Finally, adding across α we obtain

$$\overline{m}_i = \frac{v_i(k-1)}{\sum_{\beta=1}^k c^\beta}$$

Thus there is no regionalization of customer territory at NE, with firms operating in disjoint pieces of the social network. Instead, firms that are not blockaded, compete *uniformly* throughout the social network.

3 When Externalities Become Dominant

3.1 A Markov Chain Perspective

It is often is too expensive for a firm α to provide meaningful subsidies m_i^α to each customer i. Indeed the marketing division of firm α is typically allocated a fixed budget M^α and, if there is a large population of customers, then the individual expenditures m_i^α must perforce be small. In this event, customers' behavior is predominantly driven by the externality effect of their neighbors. We can capture the situation in our model by supposing that all the θ_i^α are close to 1.

Thus we are led to inquire about the limit of the NE as the $\theta_i^\alpha \longrightarrow 1$ for all α and i. (In this scenario we will also obtain a more transparent relation between NE and the graphical structure of the social network.)

To this end — and even otherwise— it is useful to recast our model in probabilistic terms. Assume, for simplicity, that $\sum_{j \in \mathcal{I}} w_{ij}^\alpha = 1$ for all i and α. Let us consider a Markov chain with $\mathcal{I}$ as the state space and W^α as the transition matrix (i.e., w_{ij}^α is the probability of going from i to j.). Let i_t denote the (random) state of the chain at date $t = 0, 1, 2, \ldots$. Suppose that, upon arrival in state i_t, a choice $L_t \in \{Stop, Move\}$ is made with $Prob(L_t = Move) = \theta_{i_t}^\alpha$. Let T be the first time $L_t = Stop$ and consider the random variable $\gamma_{i_T}^\alpha(m)$. If $\phi^\alpha(i)$ denotes the conditional expectation $E[\gamma_{i_T}^\alpha(m) | i_0 = i]$, then clearly the I-dimensional vector ϕ^α, substituted for $p^\alpha(m)$, satisfies equation (1). Since this equation has a unique solution, it must be the case that $p^\alpha(m) = \phi^\alpha$.

Recall that each vector u^α is positive, and so we may write $u^\alpha = y^\alpha \xi^\alpha$, where $y^\alpha > 0$ is a scaler and ξ^α is a probability distribution on $\mathcal{I}$. The weighted sum $[u^\alpha]^\top p(m)$ is then equal to $y^\alpha \sum_{i \in \mathcal{I}} \xi_i^\alpha \phi^\alpha(i)$ which in turn can be expressed as $y^\alpha E[\gamma_{i_T}^\alpha(m)]$, provided we assume that the probability distribution of the initial state i_0 is ξ^α. Therefore the vector v^α / y^α is just the probability distribution of i_T initializing the Markov chain at ξ^α.

We want to analyze the asymptotics of v^α as the θ_i^α converge to 1 (since the unique NE of our games are determined by v^α). Let us first consider the simple case when $\theta_i^\alpha = \theta^\alpha$ for all i. Then the random time T becomes independent of the Markov chain and we get easily that $prob(T = t) = (1 - \theta^\alpha)(\theta^\alpha)^t$.

Therefore

$$\begin{aligned} v_i^\alpha / y^\alpha &= prob(i_T = i) \\ &= \textstyle\sum_{t=0}^\infty prob(T = t) prob(i_t = i | T = t) \\ &= \textstyle\sum_{t=0}^\infty prob(T = t) prob(i_t = i) \\ &= \textstyle\sum_{t=0}^\infty prob(T = t) E[\mathbb{1}_i(i_t)] \\ &= E[\textstyle\sum_{t=0}^\infty (1 - \theta^\alpha)(\theta^\alpha)^t \mathbb{1}_i(i_t)] \end{aligned}$$

where $\mathbb{1}_i$ is the indicator function of i: $\mathbb{1}_i(j) = 0$ if $j \neq i$ and $\mathbb{1}_i(i) = 1$.

Recall that a sequence $\{a_t\}_{t \in \mathbb{N}}$ of real numbers is said to

i) Abel -converge to a if $\lim_{\theta \to 1} \sum_{t=0}^\infty (1 - \theta)(\theta)^t a_t = a$.

ii) Cesaro-converge to a if $\lim_{N \to \infty} N^{-1} \sum_{t=0}^{N-1} a_t = a$.

The Frobenius theorem (see, e.g., line 11 on page 65 of [4]) states that a Cesaro-convergent sequence is Abel-convergent to the same limit. So, to analyse the limit behavior of v_i^α, it is sufficient to consider the Cesaro-convergence of $\{\mathbb{1}_i(i_t)\}_{t \in \mathbb{N}}$.

The finite state-set $\mathcal{I}$ of our Markov chain can be partitioned into recurrent classes $I_1^\alpha, \ldots, I_{k(\alpha)}^\alpha$ and a set of transient states I_0^α. Each recurrent class I_s^α is the support of a unique invariant probability measure μ_s^α.

If the Markov process starts within a recurrent class I_s^α (i.e., $i_0 \in I_s^\alpha$), then the ergodic theorem states that, for an arbitrary function f on $\mathcal{I}$, $N^{-1} \sum_{t=0}^{N-1} f(i_t)$ converges almost surely to $E_{\mu_s^\alpha}[f]$.

If it starts at a transient state $i \in I_0^\alpha$, then we may define the first time τ that it enters $\cup_{s \geq 1} I_s^\alpha$. Let S be the index of the recurrence class i_τ belongs to. The ergodic theorem also tells us in this case that $N^{-1} \sum_{t=0}^{N-1} f(i_t)$ converges almost surely to the random variable $E_{\mu_S^\alpha}[f]$.

Let us define $\hat{\mu}^{\alpha,i}$ as the expectation $E[\mu_S^\alpha]$, if $i \in I_0^\alpha$ and as μ_s^α if $i \in I_s^\alpha$ ($s \geq 1$). Then we clearly get $E[N^{-1} \sum_{t=0}^{N-1} f(i_t) | i_0 = i] \longrightarrow E_{\hat{\mu}^{\alpha,i}}[f]$. Therefore, denoting $\hat{\mu}^\alpha \equiv \sum_{i \in \mathcal{I}} \xi_i^\alpha \hat{\mu}^{\alpha,i}$, the Frobenius theorem implies

Theorem 3. *As θ^α tends to 1, v_i^α converges to $y^\alpha E_{\hat{\mu}^\alpha}[\mathbb{1}_i] = y^\alpha \hat{\mu}_i^\alpha$.*

Corollary 1. *Suppose that the graph of the underlying social network is undirected and connected. Further suppose*

$$\theta_i^\alpha = \theta, \ w_{ik'} = w_{ik} \text{ and } \sum_{j \in \mathcal{I}} w_{ij} = 1$$

for all $\alpha \in \mathcal{A}, i \in \mathcal{I}$ and k, k' such that $w_{ik} > 0$ and $w_{ik'} > 0$ (i.e., all the nodes connected to i have the same influence on i). Finally suppose that u_i^α is invariant of i for all α (i.e., each company values all clients equally), w.lo.g. $u_i^\alpha = 1/|\mathcal{I}|$ for all α and i. Then as θ tends to 1, the money spent at NE by a company on any node is proportional to the degree[8] of the node.

[8] Recall that the degree of a node in an undirected graph is the number of edges incident on it.

To verify the corollary note that the invariant measure is (obviously) proportional to the degree. By Theorem 3, $v_i^\alpha = v_i$ converges to the degree of i as θ tends to 1. But, by Section 2.4, m_i^α is proportional to v_i.

Let us now deal with the general case where θ_i^α are not all the same. We will analyze the situation where θ_i^α is a function of a parameter θ going to 1 with the following hypotheses:

$$\lim_{\theta \to 1} \theta_i^\alpha(\theta) = 1, \text{ for all } i \tag{10}$$

$$\theta_i^\alpha(\theta) < 1, \text{ for all } i \text{ and } \theta < 1. \tag{11}$$

$$0 < \lim_{\theta \to 1} \frac{1 - \theta_i^\alpha(\theta)}{1 - \theta_1^\alpha(\theta)} = \delta_i^\alpha < \infty \tag{12}$$

For simplicity, we will also assume that $\mathcal{I} = I_1^\alpha$, i.e., there is just one recurrent class comprising all the vertices.

Theorem 4. *Under (10), (11), (12), v_i^α converges to $y^\alpha \frac{\delta_i^\alpha \mu_i^\alpha}{\sum_{j \in \mathcal{I}} \delta_j^\alpha \mu_j^\alpha}$ as θ tends to 1.*

Proof: See [2].

4 Generalizations

We have reported on some of the key results in [2]. But as shown in [2], much of the analysis can be extended to the case when externalities, utilities and costs are not necessarily linear but satisfy certain concavity/convexity conditions. In particular it can be shown that, if externalities form a "contraction", the strategic game between the firms is well-defined. Furthermore, under standard convexity hypothesis, NE continue to exist in pure strategies (see Theorem 1 of [2]). The important fixed-budget case

$$C^\alpha(m) = \begin{cases} 0 \text{ if } \sum_{i \in \mathcal{I}} m_i^\alpha \le M^\alpha \\ -\infty \text{ otherwise} \end{cases}$$

is admitted by us, as C^α is convex. (One may imagine here that the marketing division of each company α has been allocated a budget M^α to spend freely as it likes.)

It is no longer true that NE are unique (see the simple example in [2]). But if there is "enough competition" between firms, in the sense that each firm has "sufficiently many" rivals whose characteristics are "nearby", uniqueness of NE is restored. Uniqueness also holds if firms' valuations of clients are anonymous (i.e., there are no a priori biases between firms and clients), no matter how heterogenous the costs of the firms (for details see Section 5 of [2]).

Finally in [2], we also show that cross-effects (of p_j^β on p_i^α) can be incorporated, under some constraints, in our model without endangering the existence of NE.

Acknowledgements

We are grateful to Sachindra Joshi and Chandan Deep Nath for introducing us to the social network literature (in particular to the references [1,5]), and for stimulating initial discussions.

References

1. Pedro Domingos and Matt Richardson. Mining the network value of customers. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 57–66, San Francisco, USA, August 2001.
2. Pradeep Dubey, Rahul Garg, and Bernard De Meyer. Competing for customers in a social network. Stony Brook Working Paper (WP 06-1), August 2006. http://www.sunysb.edu/economics/research/working_paper/.
3. Matthew Jackson. The economics of social networks. In *Proceedings of the 9th World Congress of the Econometric Society (to appear), edited by Richard Blundell, Whitney Newey, and Torsten Persson*, Cambridge University Press, July 2005.
4. C. N. Moore. Summability of series. *The American Mathematical Monthly*, 39(2):62–71, February 1932.
5. Matthew Richardson and Pedro Domingos. Mining knowledge-sharing sites for viral marketing. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 61–70, Edmonton, Alberta, Canada, July 2002.
6. J. Scott. *Social Network Analysis: A Handbook 2nd Ed.* Sage Publications, London, 2000.
7. Carl Shapiro and Hal R. Varian. *Information Rules: A Strategic Guide to the Network Economy.* Harvard Business School Press, November 1998.
8. Oz Shy. *The Economics of Network Industries.* Cambridge University Press, 2001.