COMPETING FOR CUSTOMERS IN A SOCIAL NETWORK

By

Pradeep Dubey, Rahul Garg and Bernard De Meyer

August 2006

COWLES FOUNDATION DISCUSSION PAPER NO. 1591

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
YALE UNIVERSITY
Box 208281
New Haven, Connecticut 06520-8281

http://cowles.econ.yale.edu/

Competing for Customers in a Social Network

Pradeep Dubey* Rahul Garg† Bernard De Meyer‡

11 August 2006§

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 in general. In the quasi-linear version of the model, NE turn out to be unique and can be precisely characterized. 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 also 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.

Finally we consider convex (instead of linear) cost functions for the firms. Here NE need not be unique as we show via an example. But uniqueness is restored if there is enough competition between firms or if their valuations of clients are anonymous.

JEL Classification Numbers: A14, C72, D11, D21, L1, L2
Key words: Social network, game theory, Nash equilibrium, competition game on a 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., [2],[7],[8],[4],[10],[9]).

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 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 [2],[7] 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

*pradeepkdubey@yahoo.com, Center for Game Theory, Stony Brook University and Cowles Foundation, Yale University, USA
†grahul@in.ibm.com, IBM India Research Lab, New Delhi, INDIA
‡demeyer@univ-paris1.fr, PSE, Univesité Paris 1, FRANCE and Cowles Foundation, Yale University, USA
§This paper incorporates and supersedes [3]
[1]Customers are not strategic in our model. As in [2],[7], they are described in behavioristic terms.

present is more general than that of [2],[7], though inspired by it. As in [2],[7], 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 4.2).

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?

In Section 2 we describe a general non-linear model. So long as the externalities are a contraction, the strategy-to-outcome map (and thus the game) is well-defined. We show in Section 3 that NE exist in pure strategies under the standard convexity assumptions.

Section 4 specializes to the quasi-linear case (and includes the model in [2], by setting # firms = 1). Here we prove that NE are unique and can be easily 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 6). 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.

In Section 5 we consider convex (rather than linear) cost functions, which include the fixed-budget case where each firm can spend freely up to an exogenously specified limit. NE need not be unique as we show via an example in Section 5.1. But if there is "sufficient competition", in that each firm has enough rivals whose characteristics are nearby, then the uniqueness of NE is restored (see Section 5.3). Uniqueness also holds if firms' valuation of clients are anonymous (see Section 5.4), no matter how heterogenous their costs.

2 The General 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.

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[2] $R_+^{\mathcal{I}}$, 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 given by a function $U^\alpha : [0,1]^{\mathcal{I}} \to R$.

There is also a cost $C^\alpha(m^\alpha)$ to α from incurring the expenditures m^α. A natural candidate is $C^\alpha(m^\alpha) = \sum_{i\in\mathcal{I}} m_i^\alpha$, which simply totals the money spent by α. If the money is borrowed at interest rates that rise with subsequent tranches, $C^\alpha(m^\alpha)$ is a piecewise linear, convex function of $\sum_{i\in\mathcal{I}} m_i^\alpha$. Or, if the firm can freely spend up to some budget limit M, $C^\alpha(m^\alpha)$ is 0 if $\sum_{i\in\mathcal{I}} m_i^\alpha \le M$ and ∞ otherwise. (This is still a convex function). Our formulation of cost is general and includes these as special cases.

[2]Budget constraints on expenditures can be incorporated via cost functions (see Section 5 and Remark 1 in Section 3).

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

$$\Pi^\alpha(m) = U^\alpha(p^\alpha(m)) - C^\alpha(m^\alpha)$$

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 Section 8.3), but for the bulk of the paper 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.

This externality gives rise to a natural dynamic: if, at some time $t \geq 0$, others' proclivities to purchase are given by $q_{-i}^\alpha(m,t) \equiv (q_j^\alpha(m_j,t))_{j\in\mathcal{I}\setminus\{i\}}$, we will have $q_i^\alpha(m,t+1) = F_i^\alpha(m, q_{-i}^\alpha(m,t))$ for some function F_i^α. (Here $F_i^\alpha(m,0)$ may be viewed as the initial proclivity at time 0, which is created by the marketing expenditures m and does not take the externality into account.)

We shall suppose that the influence on i of his neighbors, albeit positive, is only partial, i.e., i puts positive weight on the money m_i^α that firm α offers to him and is not *solely* guided by the externality effect[3]. Then increasing $\sum_{j\in\mathcal{I}\setminus\{i\}} q_j^\alpha$ by Δ will no doubt boost q_i^α, but by strictly less than Δ. We make the somewhat weaker assumption that the function F_i^α is a *contraction*, i.e.,

$$\|F_i^\alpha(m, q_{-i}^\alpha) - F_i^\alpha(m, \tilde{q}_{-i}^\alpha)\| \leq K\|q_{-i}^\alpha - \tilde{q}_{-i}^\alpha\|, \text{for all } m, q_{-i}^\alpha, \tilde{q}_{-i}^\alpha$$

where $K < 1$ and $\|\cdot\|$ denotes the maximum norm.

Since the F_i^α are contractions, this dynamic process settles very quickly (geometrically) to a steady state $p^\alpha(m)$ (the unique fixed point fixed point of $F^\alpha \equiv (F_i^\alpha)_{i\in\mathcal{I}}$):

$$p_i^\alpha(m) = F_i^\alpha(m, p_{-i}^\alpha(m)), \text{for all } i \in \mathcal{I}, \alpha \in \mathcal{A}$$

We shall ignore in this paper the transient phase of the dynamic because if $q^\alpha(m,t)$ is viewed as a proclivity to purchase, then it will only be put into effect once it becomes stable. Would a customer buy a new car of a particular company when he is still in the process of revising his mind based on the feedback from his neighbors? On the other hand, if $q^\alpha(m,t)$ represents actual purchases that are occurring repeatedly in small quantities, then the aggregate purchase in the steady state overwhelms the small volume traded during the very short transient phase.

In either scenario a firm need only worry about the steady state behavior of customers in evaluating its payoff. It thus seems natural to suppose that the outcome engendered by a strategy profile m is the unique fixed point $p^\alpha(m)$ of $F^\alpha(m,\cdot)$. This fully defines the map from m to $p(m)$, and thereby the strategic game Γ between the firms.

However, at this level of abstraction, it is hard to imagine that firms can come to know the functions $(F^\alpha)_{\alpha\in\mathcal{A}}$. The social interaction between customers tends to be quite subtle and it is not easy for firms to generally predict the outcome with any degree of accuracy. But there are scenarios in which the interaction gets channelled through networks that are common knowledge. In particular this is possible in the wired world where the interaction may be tracked online and made explicit. (See for example [8],[1] as well as Section 4.2.) Then $(F^\alpha)_{\alpha\in\mathcal{A}}$ can become "manifest" to the companies, enabling them to compute the effect of the interaction, and thus to participate in the kind of game we are describing. Indeed we will focus on networks in most of our analysis.

[3]In Section 6 we examine the scenario when externalities become dominant

3 Existence of Nash Equilibrium

Recall that a strategy profile m is called a **Nash Equilibrium**[4] (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 $m^{-\alpha} \equiv (m^\beta)_{\beta \in \mathcal{I}\setminus\{\alpha\}}$).

It turns out that NE exist in our model under quite general conditions which we list below.

AI: The cost function $C^\alpha : R_+^{\mathcal{I}} \to R_+$ is continuous, convex and strictly increasing[5].

AII: The benefit function $U^\alpha : [0,1]^{\mathcal{I}} \to R$ is continuous, concave and increasing.

AIII: The externality function $F_i^\alpha(m,q) = F_i^\alpha(m^\alpha, m^{-\alpha}, q_{-i}^\alpha)$ is continuous if $m >> 0$[6]; and is concave and increasing in m^α for every fixed $m^{-\alpha}, q_{-i}^\alpha$. Furthermore F_i^α is a contraction in q for every fixed m.

Our last assumption has to do with the possible discontinuity of the function $F_i^\alpha(m,q)$ as $m_i \equiv (m_i^\alpha)_{\alpha\in\mathcal{A}} \to 0$. We require that, for each customer i, there be at least two distinct firms who value i, so that the competition between them will ensure that the total money spent on i is positive in any NE. The intuition is that, if m_i is too small, either of the two firms could spend a "sliver" on i, which costs very little, but is nevertheless overwhelmingly more than other firms' expenditures on i, and thus is able to "buy out" i, contradicting that it has optimized. Formally, denoting by $m_{-i}^{-\tau}$ the vector m with the component m_i^τ suppressed, we have

AIV For each customer i, there exist two distinct firms α and β such that:

(i) Both firms value i, i.e., U^α and U^β are strictly increasing in the variables p_i^α and p_i^β respectively.

(ii) Customer i responds to the marketing of both firms, i.e., for $\tau = \alpha$ or $\tau = \beta$,

$$\liminf [F_i^\tau((m_i^\tau + \delta, m_{-i}^{-\tau}), q_{-i}^\tau) - F_i^\tau(m, q_{-i}^\tau)]/\delta = \infty$$

where the lim inf is taken over sequences $\{m, \delta\}$ that satisfy the conditions: $(m_i, \delta) \to 0$, $(m_i^\tau + \delta)/m_i^{\tau'} \to \infty$ for all $\tau' \in \mathcal{A}\setminus\{\tau\}$, and $m_i^\tau \leq m_i^{\tau'}$ for some $\tau' \in \mathcal{A}\setminus\{\tau\}$.

To interpret the second part of AIV, take $\tau = \alpha$ and consider a unilateral deviation by α wherein α increases m_i^α to $m_i^\alpha + \delta$. Since all $\beta \in \mathcal{A}\setminus\{\alpha\}$ have expenditures m_i^β on i that are vanishingly small compared to the expenditure $m_i^\alpha + \delta$ made by α, firm α must have 100% of the "marketing impact" on i in the limit, on account of its deviation. On the other hand, it has less than 50% of the impact, prior to its deviation, since its expenditure is over-matched by at least one rival firm. But the jump from 50% to 100% is non-negligible since — as was said — i is not guided solely by the externality effect of his neighbors, and since the marketing impact affects his proclivities by (say) at least $\theta > 0$. The bracketed term $[\ldots]$ is thus of the order of $\theta/2$ and so the whole term goes to infinity like $\theta/2\delta$ as $\delta \to 0$. Our assumption is weaker, allowing for the total probability of purchase across all firms by customer i to go to zero (sufficiently slowly) as the aggregate expenditure $m_i \to 0$.

We are ready to state our main existence result.

Theorem 1. *Assume AI, AII, AIII, AIV. Then a Nash Equilibrium (NE) exists in the game* Γ. *Moreover, if* m *is a NE,* $\overline{m}_i > 0$ *for all* $i \in \mathcal{I}$.

Proof: See the Appendix.

Remarks

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

[5] i.e., if $x \geq y, x \neq y$ imply $g(x) \geq g(y)$, we say that g is increasing. If the last inequality is strict, we say that g is strictly increasing.

[6] i.e., each component of m is strictly positive

(1) Theorem 1 remains intact (with obvious amendments in the proof) if we drop the strictly increasing property of C^α and replace it by the requirement that α's expenditures must lie in a compact, convex subset S^α of $R_+^{\mathcal{I}}$. (Strictness is only used to bound the expenditures of α.) This is tantamount to taking C^α to be convex and continuous on S^α and $-\infty$ on $R_+^{\mathcal{I}} \backslash S^\alpha$.

(2) If F_i^α is continuous in m (even when $m \to 0$) then AIV can be dropped. In need only be postulated for those i where continuity fails. Existence of NE remains intact, but now the total money spent on a client may be zero.

4 The Quasi-Linear Model

4.1 The Data of the Economy

We turn to a quasi-linear version of our model, which is particularly transparent, and in which NE are not only unique but can be precisely characterized. The social network now has a concrete representation in terms of a directed, weighted graph $G = (\mathcal{I}, E, w)$. The nodes of G are identified with the set of customers $\mathcal{I}$. Each directed edge $(i,j) \in E \equiv \mathcal{I} \times \mathcal{I}$ has weights $(w_{ij}^\alpha)_{\alpha \in \mathcal{A}}$, where $w_{ij}^\alpha \geq 0$ is a measure of the influence j has on i, with regard to purchases from α. Precisely, if $p^\alpha = (p_j^\alpha)_{j \in \mathcal{I}}$ denotes the proclivities of purchases, then the externality impact of p^α on i is $\sum_{j \in \mathcal{I}} w_{ij}^\alpha p_j^\alpha$. We assume that $\sum_{j \in \mathcal{I}} w_{ij}^\alpha \leq 1$, 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 $E^\alpha = \{(i,j) \in E : w_{ij}^\alpha > 0\}$).

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

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

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

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

A canonical example we have in mind is $\gamma_i^\alpha(m_i) = m_i^\alpha / \overline{m}_i$ where $\overline{m}_i \equiv (\sum_{\beta \in \mathcal{I}} m_i^\beta)$ (with $\gamma_i^\alpha(0) \equiv 0$). In short, i's probability of purchase from different firms is simply set proportional to the money they spend on him[7].

Customer i weights the two factors (i.e., the externality impact and the marketing impact) by θ_i^α and $1 - \theta_i^\alpha$, where $0 \leq \theta_i^\alpha < 1$. 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 $p^\alpha \equiv (p_j^\alpha(m))_{j \in \mathcal{I}}$, must satisfy

$$p_i^\alpha(m) = (1 - \theta_i^\alpha)\gamma_i^\alpha(m_i) + \theta_i^\alpha \sum_{j \in \mathcal{I}} w_{ij}^\alpha p_j^\alpha(m) \tag{1}$$

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

The fact that (1) has a unique solution follows, of course from our analysis of the general model, once one observes that the map $(p_i^\alpha)_{i \in \mathcal{I}} \longmapsto (\theta_i^\alpha \sum_{j \in \mathcal{I}} w_{ij}^\alpha p_j^\alpha)_{i \in \mathcal{I}}$ is a contraction since $\theta_i^\alpha < 1$ and $\sum_{j \in \mathcal{I}} w_{ij}^\alpha \leq 1$.

Define the $\mathcal{I} \times \mathcal{I}$-matrices: $I \equiv$ identity, $\Theta^\alpha \equiv$ the diagonal matrix with $\Theta_{ii}^\alpha = \theta_i^\alpha$ and $W^\alpha \equiv$ the matrix with entries w_{ij}^α. 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

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

[7]More generally, $\gamma_i^\alpha(m_i) = (m_i^\alpha / \overline{m}_i)(\overline{m}_i)^r$ where $0 \leq r < 1$. We may think of $(\overline{m}_i)^r$ as the "market penetration", which rises with the total money spent. Notice $p_i^\alpha(m)$ is effectively bounded. This is so because the derivative of γ_i^α (w.r.t. m_i^α) goes to zero as $m_i^\alpha \to \infty$, while the cost of m_i^α is fixed — see later — at $c_i^\alpha > 0$, bounding m_i^α (and so γ_i^α). (If $\gamma_i^\alpha(m_i)$ is to be a probability, one must amend $(\overline{m}_i)^r$ to $\max\{(\overline{m}_i^r), 1\}$ or a suitably smoothed version of this function.)

It still remains to specify U^α, C^α and γ_i^α. We take U^α and C^α to be linear:

$$U^\alpha(p^\alpha) = \sum_{j\in\mathcal{I}} u_j^\alpha p_j^\alpha$$
$$C^\alpha(m^\alpha) = \sum_{j\in\mathcal{I}} c_j^\alpha m_j^\alpha$$

with $u_j^\alpha \geq 0$ and $c_j^\alpha > 0$ for all $j \in \mathcal{I}$. This gives

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

where u^α and c^α are the column vectors $(u_j^\alpha)_{j\in\mathcal{I}}$, $(c_j^\alpha)_{j\in\mathcal{I}}$ 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[8] expenditure of its rivals.

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

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

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

which is the analogue of AIV in our quasi-linear world.

Note that both conditions (5) and (6) are satisfied by our canonical example and its variants in footnote 8.

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

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

[8]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.

4.2 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.

4.3 Uniqueness of Nash Equilibrium

Theorem 2. *Under hypotheses (5), (6), (7), there exists a unique Nash Equilibrium in the quasi-linear model.*

Proof: First observe that condition (5) implies that $(\partial/\partial m_i^\alpha)\gamma_i^\alpha(m_i^\alpha, m_i^{-\alpha})$ is decreasing in m_i^α (for any fixed $m_i^{-\alpha}$), i.e., γ_i^α is concave in its first variable. Thus, in conjunction with (6), and (7), all the requirements of Theorem 1 are met and an NE exists, with $\overline{m}_i \equiv \sum_{\beta\in\mathcal{I}} m_i^\beta > 0$ for all $i \in \mathcal{I}$[9].

Suppose $m \equiv (m^\alpha)_{\alpha\in\mathcal{I}}$ and $\eta \equiv (\eta^\alpha)_{\alpha\in\mathcal{I}}$ are two NE's. Denote $r_i^\alpha \equiv m_i^\alpha/\overline{m}_i$ and $s_i^\alpha \equiv \eta_i^\alpha/\overline{\eta}_i$ (where, recall, $\overline{m}_i \equiv \sum_{\alpha\in\mathcal{A}} m_i^\alpha$ etc.). It suffices to show that $\overline{m}_i = \overline{\eta}_i$ and $r_i^\alpha = s_i^\alpha$ for all $\alpha \in \mathcal{A}$ and all $i \in \mathcal{I}$.

The first-order conditions[10] for maximizing payoffs imply

$$
\begin{aligned}
v_i^\alpha \lambda_i^\alpha(r_i^\alpha, \overline{m}_i) &= c_i^\alpha \text{ if } m_i^\alpha > 0 && (8)\\
v_i^\alpha \lambda_i^\alpha(r_i^\alpha, \overline{m}_i) &\le c_i^\alpha \text{ if } m_i^\alpha = 0 && (9)\\
v_i^\alpha \lambda_i^\alpha(s_i^\alpha, \overline{\eta}_i) &= c_i^\alpha \text{ if } \eta_i^\alpha > 0 && (10)\\
v_i^\alpha \lambda_i^\alpha(s_i^\alpha, \overline{\eta}_i) &\le c_i^\alpha \text{ if } \eta_i^\alpha = 0 && (11)
\end{aligned}
$$

Fix $i \in \mathcal{I}$ and suppose w.l.o.g. that $\overline{m}_i \le \overline{\eta}_i$.

Step 1: $s_i^\alpha \le r_i^\alpha$ for all $\alpha \in \mathcal{A}$.

Proof: First note that, by (4), $v_i^\alpha = 0$ implies $m_i^\alpha = 0$ in any NE. Let $s_i^\alpha > 0$ (otherwise the claim is vacuously true) and so we must have $v_i^\alpha > 0$. Suppose, to the contrary, that $s_i^\alpha > r_i^\alpha$. Since $\overline{\eta}_i \ge \overline{m}_i$, condition (5) on λ_i^α imply $\lambda_i^\alpha(s_i^\alpha, \overline{\eta}_i) < \lambda_i^\alpha(r_i^\alpha, \overline{m}_i)$. But (8), (9) and (10) and the fact that $v_i^\alpha > 0$ together yield $\lambda_i^\alpha(s_i^\alpha, \overline{\eta}_i) \ge \lambda_i^\alpha(r_i^\alpha, \overline{m}_i)$, a contradiction.

Step 2: $s_i^\alpha = r_i^\alpha$ for all $\alpha \in \mathcal{A}$.

Proof: Immediate from step 1, since $\sum_{\alpha\in\mathcal{A}} s_i^\alpha = 1 = \sum_{\alpha\in\mathcal{A}} r_i^\alpha$.

Step 3: $\overline{m}_i = \overline{\eta}_i$

Proof: Suppose $\overline{\eta}_i > \overline{m}_i$ (by assumption we already have $\ge$). By step 2, and condition (ii) on λ_i^α, we have LHS of (10) $<$ LHS of (8). Since $\sum_{\beta\in\tau} r_i^\beta = 1$ there exists β' such that $r_i^{\beta'} > 0$. By step 2, $s_i^{\beta'} = r_i^{\beta'}$, so both (10) and (8) hold, hence LHS of (10) = LHS of (8), a contradiction. This proves step 3.

Since the choice of i was arbitrary, we have shown that $\overline{\eta}_i = \overline{m}_i$ and $r_i^\alpha = s_i^\alpha$ for all $\alpha \in \mathcal{A}$ and all $i \in \mathcal{I}$. Thus $m = \eta$, establishing the uniqueness of NE. ∎

4.4 Characterization of Nash Equilibrium

Theorem 3. *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 $\kappa_i^1 \le \kappa_i^2 \le \ldots \le \kappa_i^n$. Let*

$$
k_i = \max\left\{ l \in \{2, \ldots, n\} : (l-2)\kappa_i^l < \sum_{\alpha=1}^{l-1} \kappa_i^\alpha \right\} \tag{12}
$$

[9]For better perspective, here is an alternative proof that $\overline{m}_i > 0$. Suppose, to the contrary, that $\overline{m}_i = 0$ for some i. By assumption, there exists α such that $u_i^\alpha > 0$. By (3),

$$
\begin{aligned}
v^\alpha &= [u^\alpha]^\top (I - \Theta^\alpha W^\alpha)^{-1}(I - \Theta^\alpha)\\
&= [u^\alpha]^\top \sum_{n=0}^{\infty} (\Theta^\alpha W^\alpha)^n (I - \Theta^\alpha)
\end{aligned}
$$

from which it follows that $[u]^\top$ is being multiplied by a matrix with non negative entries and strictly positive diagonal entries. Hence $v_i^\alpha > 0$.

Let firm α unilaterally deviate from m by spending a small δ on customer i. By (4), his change in payoff is

$$
v_i^\alpha \gamma_i^\alpha(\delta, 0) - c_i^\alpha \delta
$$

which, using (7), becomes positive for small enough δ, contradicting that m is an NE. We conclude that $\overline{m}_i > 0$ for all $i \in \mathcal{I}$.

[10]Since $\overline{m}_i > 0$ and $\overline{\eta}_i > 0$, and the γ_i^α are differentiable away from zero, these conditions can be invoked.

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

$$m_i^\alpha = \left(\frac{k_i - 1}{\sum_{\beta=1}^{k_i} \kappa_i^\beta}\right)\left(1 - \frac{(k_i - 1)\kappa_i^\alpha}{\sum_{\beta=1}^{k_i} \kappa_i^\beta}\right) \tag{13}$$

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

Proof: Note that $\lambda_i^\alpha(r_i^\alpha, \overline{m}_i) = (1 - r_i^\alpha)/\overline{m}_i$ in our canonical case. Thus the first-order conditions (8) and (9) become

$$\frac{1 - r_i^\alpha}{\overline{m}_i} = \kappa_i^\alpha \text{ if } r_i^\alpha > 0 \tag{14}$$

$$\frac{1 - r_i^\alpha}{\overline{m}_i} \leq \kappa_i^\alpha \text{ if } r_i^\alpha = 0 \tag{15}$$

It follows at once that, if $r_i^\alpha > 0$, then $r_i^\beta > 0$ whenever $\kappa_i^\beta \leq \kappa_i^\alpha$. Hence we only need check that: (i) condition (14) holds for $1 \leq \alpha \leq k_i$; (ii) condition (15) holds for $k_i + 1 \leq \alpha \leq n$; and (iii) $r_i^\alpha > 0$ for $1 \leq \alpha \leq k_i$.

Note that (13) implies

$$\overline{m}_i = \sum_{\beta=1}^{k_i} m_i^\beta = \frac{k_i - 1}{\sum_{\beta=1}^{k_i} \kappa_i^\beta} \tag{16}$$

And (13) and (16) imply

$$\begin{aligned} 1 - r_i^\alpha &= 1 - \frac{m_i^\alpha}{\overline{m}_i} \\ &= \frac{(k_i - 1)\kappa_i^\alpha}{\sum_{\beta=1}^{k_i} \kappa_i^\beta} \end{aligned}$$

Then (i) follows from the above equation and (16).

It suffices to show (ii) for $\alpha = k_i + 1$, since LHS of (15) $= 1/\overline{m}_i$ for all $\alpha \geq k_i + 1$ (on account of $r_i^\alpha = 0$) and since RHS of (15) rises with α.

Taking $l = k_i + 1$, and violating the inequality in (12), we obtain

$$(k_i - 1)\kappa_i^{k_i+1} \geq \sum_{\beta=1}^{k_i} \kappa_i^\beta$$

i.e.,

$$\frac{\sum_{\beta=1}^{k_i} \kappa_i^\beta}{k_i - 1} \leq \kappa_i^{k_i+1}$$

which, together with (16), implies (ii).

Finally, taking $l = k_i$ in (12), we have

$$(k_i - 2)\kappa_i^k \leq \sum_{\beta=1}^{k_i-1} \kappa_i^\beta$$

Adding κ_i^k to both sides gives

$$(k_i - 1)\kappa_i^k < \sum_{\beta=1}^{k_i} \kappa_i^\beta$$

i.e.,

$$1 - \frac{(k_i - 1)\kappa_i^k}{\sum_{\beta=1}^{k_i} \kappa_i^\beta} > 0$$

But, since $\kappa_i^\alpha \leq \kappa_i^k$ for $\alpha \leq k_i$, this yields

$$1 - \frac{(k_i - 1)\kappa_i^\alpha}{\sum_{\beta=1}^{k_i} \kappa_i^\beta} > 0$$

for $1 \leq \alpha \leq k_i$. Then, by (13), $m_i^\alpha > 0$, i.e., $r_i^\alpha > 0$ for $1 \leq \alpha \leq k_i$ verifying (iii). ∎

According to Theorem 3, 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 3 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, (13) 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.

4.5 Impact of the Social Network on Nash Equilibrium

To get a better feel for Theorem 3, 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 3 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. Assume that there are no a priori biases between firms and customers: $w_{ij}^\alpha = w_{ij}$ and $\theta_i^\alpha = \theta_i$ for all $\alpha \in \mathcal{A}$ and $i, j \in \mathcal{I}$. It then follows from (3) that the game is anonymous, i.e., $v_i^\alpha = v_i$ for all α and i. To simplify the analysis, further assume: $c_i^\alpha = c^\alpha$. Our analysis in Section 4.4 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 (12), 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 (13), 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, by (16), we have

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

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.

5 Uniqueness with Convex Costs

Consider the quasi-linear model but with convex, instead of linear, cost functions C^α. We will no longer assume that the C^α are continuous. In particular, the *fixed-budget case*

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

is admitted by us, as C^α is still 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.)

5.1 Multiple Nash Equilibria

Unfortunately it is no longer true that NE are unique. Consider the following simple fixed-budget example. There are two customers and two firms. The budgets of the two firms are identical: $M^1 = M^2 = 1$. Suppose $v_1^2 = v_1^1 = 1$, $v_2^2 = 0.02$ and $v_2^1 = 500$. (The u_i^α can be adjusted, given any $w_{ij}^\alpha > 0$ and $0 \leq \theta^\alpha < 1$, to guarantee that the v_i^α take on these values.) Finally take our canonical marketing function, i.e., $\gamma_i^\alpha(m) = m_i^\alpha/\overline{m}_i$ if $\overline{m}_i > 0$ and 0 otherwise. By Theorem 1 (see Remark 1) there exists NE; and, if m is any NE, we must have $\overline{m}_1 > 0$ and $\overline{m}_2 > 0$. Now if $m_i^\alpha = 0$ for any α and i then the rival firm β can reduce m_i^β and shift money to the other client $j \neq i$, improving its payoff. We conclude that $m_i^\alpha > 0$ for $i = 1, 2$ and $\alpha = 1, 2$. Therefore m is an NE if, and only if, the first order conditions below hold

$$\frac{v_1^1 m_1^2}{(m_1^1 + m_1^2)^2} = \frac{v_2^1 m_2^2}{(m_2^1 + m_2^2)^2} \equiv c^1, \tag{17}$$

$$\frac{v_1^2 m_1^1}{(m_1^1 + m_1^2)^2} = \frac{v_2^2 m_2^1}{(m_2^1 + m_2^2)^2} \equiv c^2 \tag{18}$$

along with $m >> 0$, $m_1^1 + m_2^1 = M^1$ and $m_1^2 + m_2^2 = M^1$ (Here c^α can be interpreted as the marginal utility of a dollar for company α at the NE).

By straightforward algebra, we obtain

$$m_i^1 = \frac{c_2 (v_i^1)^2 v_i^2}{(c_1 v_i^2 + c_2 v_i^1)^2}, \tag{19}$$

$$m_i^2 = \frac{c_1 (v_i^2)^2 v_i^1}{(c_1 v_i^2 + c_2 v_i^1)^2}, \tag{20}$$

$$M^1 + M^2 = \sum_{i=1}^{2} v_i^1 v_i^2 \frac{1}{(c_1 v_i^2 + c_2 v_i^1)} \tag{21}$$

and

$$\sum_{i=1}^{2} \frac{v_i^1 v_i^2}{(c_1 v_i^2 + c_2 v_i^1)^2} (M^2 c_2 v_i^1 - M^1 c_1 v_i^2) = 0 \tag{22}$$

Clearly if (c_1, c_2) solves(22), then so does $(\lambda c_1, \lambda c_2)$ for any $\lambda > 0$. So consider (22) with $c_2 = 1$, which yields (substituting our values for v_i^α and M^α) a cubic equation in c_1 with three positive roots, whose approximate values are $\tilde{c_1} = 1.087109$, $\tilde{c_1} = 47.1973$, and $\tilde{c_1} = 24800.020967164826$. But in order to satisfy (21) we must have $c_2 = \lambda$, $c_1 = \lambda\tilde{c_1}$, where λ satisfies

$$\lambda = \frac{1}{2}(\frac{1}{1 + \tilde{c_1}} + \frac{10}{500 + 0.02\tilde{c_1}})$$

Thus we get three distinct pairs $(c_1, c_2) \approx (0.271305, 1.808711)$, $(c_1, c_2) \approx (0.96071, 0.0203552)$ and $(c_1, c_2) \approx (125, 0.0050402)$ which give (via (19) and (20)) three distinct η, $\tilde{\eta}$ and $\overline{\eta}$ as NE:

$$\eta_1^1 \approx 0.919868, \quad \eta_2^1 = 1 - \eta_1^1$$

$$\eta_1^2 \approx 0.999996, \quad \eta_2^2 = 1 - \eta_1^2$$

and

$$\tilde{\eta}_1^1 \approx 0.0211485, \quad \tilde{\eta}_2^1 = 1 - \tilde{\eta}_1^1$$
$$\tilde{\eta}_1^2 \approx 0.998152, \quad \tilde{\eta}_2^2 = 1 - \tilde{\eta}_1^2$$

and

$$\overline{\eta}_1^1 \approx 3.2256 \times 10^{-7}, \quad \overline{\eta}_2^1 = 1 - \overline{\eta}_1^1$$
$$\overline{\eta}_1^2 \approx 0.0079995, \quad \overline{\eta}_2^2 = 1 - \overline{\eta}_1^2$$

(The reader may numerically check that η, $\tilde{\eta}$ and $\overline{\eta}$ are indeed approximate solutions to (17) and (18).)

Notice that the two companies have widely disparate valuations of client 2 in our counter example: $v_2^1 = 500$ and $v_2^2 = 0.02$. Curiously, if we replicate each company, the counterexample disappears and uniqueness of NE is restored. More generally uniqueness holds if, for each company there are "sufficiently many" other companies whose characteristics are "nearby"[11]. Of course the words in quotes must be made precise (which we shall do in Section 5.3).

But the counter example does show the need to impose additional constraints on the marketing functions γ_i^α to guarantee uniqueness of NE, a matter to which we now turn.

5.2 Uniqueness with Multi-Concavity

We shall first present an abstract result and later bring it to bear on our model. Let $S^\alpha \subset R_+^{\mathcal{I}}$ be the (closed, convex) strategy-set of $\alpha \in \mathcal{A}$. Given a strategy profile $(m^\beta)_{\beta\in\mathcal{A}} \in \times_{\beta\in\mathcal{A}} S^\beta$, suppose the payoff to any firm α depends on his action s^α and the aggregate $\sum_{\beta\in\mathcal{A}\setminus\{\alpha\}} m^\beta$ of others' actions. So we may take α's payoff Π^α to be defined on $S^\alpha \times R_+^{\mathcal{I}}$. It will be convenient to extend the domain of Π^α to $R^{\mathcal{I}} \times R_+^{\mathcal{I}}$ by putting $\Pi^\alpha(m^\alpha, .) = -\infty$ if $m^\alpha \notin S^\alpha$. Assume $\Pi^\alpha(m_i^\alpha, m_{-i}^{-\alpha})$ is concave in m_i^α for any fixed $m_{-i}^{-\alpha}$, and let $\partial_{m_i^\alpha}\Pi^\alpha(m)$ denote its superdifferential w.r.t. m_i^α.

Theorem 4. *Suppose for all α and i, the superdifferential $\partial_{m_i^\alpha}\Pi^\alpha(m)$ is a correspondence $h_i^\alpha(m_i^\alpha, \overline{m}_i, \overline{m}^\alpha)$ that is strictly decreasing in m_i^α and decreasing in both $\overline{m}_i$ and $\overline{m}^\alpha$*[12]*. Then there is at most one Nash equilibrium.*

The proof is based on the following lemma which may be of independent interest. Let B be an $\mathcal{I} \times \mathcal{A}$ matrix with entries B_i^α for $i \in \mathcal{I}$ and $\alpha \in \mathcal{A}$. Define $\overline{B}^\alpha \equiv \sum_i B_i^\alpha$ and $\overline{B}_i \equiv \sum_\alpha B_i^\alpha$. A cell (i, α) is said to be positive for B if $\overline{B}_i \geq 0$, $\overline{B}^\alpha \geq 0$ and $B_i^\alpha > 0$. It is said to be negative for B if it is positive for $-B$. Finally it is said to be signed if it is either positive or negative for B.

Lemma 1. *Any non zero matrix B has a signed cell.*

Proof: Suppose A has no signed cell. Let $\mathcal{A}^+ \equiv \{\alpha \in \mathcal{A} : \overline{B}^\alpha \geq 0\}$, $\mathcal{A}^- \equiv \mathcal{A} - \mathcal{A}^+$,; and similarly $\mathcal{I}^+ \equiv \{i \in \mathcal{I} : \overline{B}_i \geq 0\}$, $\mathcal{I}^- \equiv \mathcal{I} - \mathcal{I}^+$.

Since B has no positive cell,

$$B_i^\alpha \leq 0 \text{ for all } \alpha \in \mathcal{A}^+ \text{ and } i \in \mathcal{I}^+ \tag{23}$$

Similarly, since B has no negative cell,

$$B_i^\alpha \geq 0 \text{ for all } \alpha \in \mathcal{A}^- \text{ and } i \in \mathcal{I}^- \tag{24}$$

Now, for all $\alpha \in \mathcal{A}^-$, $0 > \overline{B}^\alpha = \sum_{\mathcal{I}^+} B_i^\alpha + \sum_{\mathcal{I}^-} B_i^\alpha$. By (24) this is only possible if $\sum_{\mathcal{I}^+} B_i^\alpha < 0$.

We will prove that $\mathcal{A}^- = \emptyset$. Indeed, if $\mathcal{A}^- \neq \emptyset$, then

$$\sum_{\mathcal{A}^-}\sum_{\mathcal{I}^+} B_i^\alpha < 0.$$

[11] As we expand the neighborhood of characteristics that defines "nearby", we will need to put in more companies in that neighborhood.

[12] A correspondence $\Lambda(x, y, z)$ is said to be decreasing in x if $v \in \Lambda(x, y, z)$, $w \in \Lambda(x', y, z)$, $x' > x$ together imply $v \geq w$. If the last inequality is strict, we say that Λ is strictly decreasing in x.

On the other hand, for all $i \in \mathcal{I}^+$ we have $0 \leq \overline{B}_i = \sum_{\mathcal{A}^+} B_i^\alpha + \sum_{\mathcal{A}^-} B_i^\alpha$. By (23) this is only possible if $\sum_{\mathcal{A}^-} B_i^\alpha \geq 0$ and therefore $\sum_{\mathcal{A}^-} \sum_{\mathcal{I}^+} B_i^\alpha \geq 0$, a contradiction.

A symmetric argument shows that $\mathcal{I}^- = \emptyset$. So, by (23), all elements B_i^α are non-positive and their row sums $\overline{B}_i$ are non-negative. This is only possible if $B = 0$. ∎

Proof of Theorem 4: Let $m \in R^{\mathcal{I}\times\mathcal{A}}$ and $\eta \in R^{\mathcal{I}\times\mathcal{A}}$ be two NE of the game. Define the matrix $B = \eta - m$. We shall show that it has no signed cell, so that by the lemma, $B = 0$, proving $m = \eta$.

Suppose i, α is a positive cell for B, then we have $m_i^\alpha < \eta_i^\alpha$, $\overline{m}_i \leq \overline{\eta}_i$, $\overline{m}^\alpha \leq \overline{\eta}^\alpha$. By the strictly decreasing property of h_i^α in its first variable, if $x \in h_i^\alpha(m_i^\alpha, \overline{m}_i, \overline{m}^\alpha)$ and $y \in h_i^\alpha(\eta_i^\alpha, \overline{\eta}_i, \overline{\eta}^\alpha)$, we must have $x > y$. But this contradicts the first order conditions of NE, according to which $0 \in h_i^\alpha(m_i^\alpha, \overline{m}_i, \overline{m}^\alpha)$ and $0 \in h_i^\alpha(\eta_i^\alpha, \overline{\eta}_i, \overline{\eta}^\alpha)$, $\forall i, \alpha$.

By a symmetric argument, B has no negative cell. ∎

To apply Theorem 4 to our model, we focus on the case when costs are convex in total expenditure and benefits are linear, so that the payoff function may be written

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

(We shall refer to this in brief as *"the quasi-linear model with convex costs"*.)

The superdifferential of $-C^\alpha$ is clearly decreasing in $\overline{m}^\alpha$ since C^α is convex. It therefore suffices to check that the superdifferential of γ_i^α can be expressed as a correspondence of two variables m_i^α and $\overline{m}_i$, which is strictly decreasing in m_i^α and decreasing in $\overline{m}_i$.

Consider, for concreteness, our canonical marketing function $\gamma_i^\alpha = m_i^\alpha / \overline{m}_i$. Note

$$\frac{\partial \gamma_i^\alpha}{\partial m_i^\alpha} = \frac{\overline{m}_i - m_i^\alpha}{(\overline{m}_i)^2} \equiv h(m_i^\alpha, \overline{m}_i)$$

Clearly h is strictly decreasing in m_i^α. But

$$\frac{\partial h}{\partial \overline{m}_i} = \frac{-\overline{m}_i + 2m_i^\alpha}{(\overline{m}_i)^3}$$

is non-positive if, and only if, $m_i^\alpha \leq \overline{m}_i/2$.

In the light of this, Theorem 4 (or, rather, its proof) immediately yields

Corollary 1. *In the quasi-linear model with convex costs and canonical marketing, there is at most one NE in the region*

$$\Omega \equiv \{m : m_i^\alpha \leq \overline{m}_i/2 \textit{ for all } \alpha \in \mathcal{A} \textit{ and } i \in \mathcal{I}\}$$

No wonder that, in the counterexample of Section 5.1, the NE were not contained in Ω.

5.3 Competition Restores Uniqueness

Throughout this section we confine ourselves to the fixed-budget model with the canonical marketing function, which was also the context of the counterexample.

We shall show that, with "enough competition" no NE can be outside Ω. This will guarantee uniqueness of NE (by Corollary 1).

First consider the time-honored device of creating competition by replicating the companies, i.e., for any α, there is a replica (twin) $\tilde{\alpha}$ with identical characteristics ($\theta^\alpha = \theta^{\tilde{\alpha}}$, $u^\alpha = u^{\tilde{\alpha}}$, $W^\alpha = W^{\tilde{\alpha}}$, $M^\alpha = M^{\tilde{\alpha}}$). It suffices to show that replicas act identically in any NE, for then obviously $m \in \Omega$.

We shall prove this by contradiction. Suppose m is an NE with $m^\alpha \neq m^{\tilde{\alpha}}$. Since [13] $\sum_{i\in\mathcal{I}} m_i^\alpha = M^\alpha = M^{\tilde{\alpha}} = \sum_{i\in\mathcal{I}} m_i^{\tilde{\alpha}}$, there exist clients i and j such that $r_i^\alpha > r_i^{\tilde{\alpha}}$ and $r_j^\alpha < r_j^{\tilde{\alpha}}$ (where, recall, $r_i^\alpha \equiv m_i^\alpha/\overline{m}_i$ etc.).

[13] Clearly both the companies will spend all their money at any NE, since each puts positive value on at least one customer-node.

The first order conditions of NE are

$$\frac{v_i^\alpha(1-r_i^\alpha)}{\overline{m}_i} \geq \frac{v_j^\alpha(1-r_j^\alpha)}{\overline{m}_j} \tag{25}$$

$$\frac{v_i^{\tilde{\alpha}}(1-r_i^{\tilde{\alpha}})}{\overline{m}_i} \leq \frac{v_j^{\tilde{\alpha}}(1-r_j^{\tilde{\alpha}})}{\overline{m}_j} \tag{26}$$

(Since $r_i^\alpha > 0$, we have $m_i^\alpha > 0$ and so the LHS of (25) must equal $c^\alpha \equiv$ the marginal utility of a dollar to α at the NE. But RHS of (25) is at most c^α, proving (25). A similar argument can be made for (26).)

But $v_i^\alpha = v_i^{\tilde{\alpha}}$ and $v_j^\alpha = v_j^{\tilde{\alpha}}$ by (3). So LHS of (26)> LHS of (25) $\geq$ RHS of (25) > RHS of (26), contradicting (26).

This establishes uniqueness of NE under replication.

But it is not necessary to have exact replicas. It suffices to assume that, for each company α, there are sufficiently many rivals whose characteristics are "close enough" to those of α. (As we relax the notion of "closeness", we will need to put in more rivals.) Precisely, we have:

Theorem 5. *Consider the quasi-linear model with fixed positive budgets*[14] *and canonical marketing. For any $\beta \in \mathcal{A}$, denote*[15] $v_{min}^\beta = \min\{v_j^\beta : j \in \mathcal{I}, v_j^\beta > 0\}$ *and* $v_{max}^\beta = \max\{v_j^\beta : j \in \mathcal{I}, v_j^\beta > 0\}$. *Fix an integer $n \geq 2$. Assume that, for each $\alpha \in \mathcal{A}$ there exists $\mathcal{A}^\alpha(n) \subset \mathcal{A}\backslash\{\alpha\}$ such that $v_i^\beta = 0$ if and only if $v_i^\alpha = 0$ for all $\beta \in \mathcal{A}^\alpha(n)$ and all $i \in \mathcal{I}$. Furthermore assume, for every $\alpha \in \mathcal{A}$, that*

(i) $|\mathcal{A}^\alpha(n)| \geq n$
(ii)(a) $\exists k_1 > 0, k_2 > 0$ *s.t.* $k_1k_2 < (2n-1)/n$; *and*
(b) $k_1^{-1} \leq v_i^\alpha/v_i^\beta \leq k_2$ *for all* $\beta \in \mathcal{A}^\alpha(n)$ *and all* $i \in \{j \in \mathcal{I} : v_j^\alpha > 0\}$.
(iii) $M^\alpha - M^\beta \leq \frac{3}{16|\mathcal{I}|}(v_{min}^\beta/v_{max}^\beta)M^\beta$ *for all* $\beta \in \mathcal{A}^\alpha(n)$

Then Γ has a unique NE.

Proof. By Corollary 1, we need only verify that, if m is an NE, then $m \in \Omega$, i.e., $r_i^\alpha \leq 1/2$ for all α and i. Suppose some $r_i^\alpha > 1/2$. Then, clearly, there exists $\beta \in \mathcal{A}^\alpha(n)$ such that $r_i^\beta < 1/2n$

Claim 1. *There exists $j \in \mathcal{I}\backslash\{i\}$ such that $r_j^\beta > r_j^\alpha$.*

Proof. It suffices to show that β has more money left to spend[16] on $\mathcal{I}\backslash\{i\}$ than does α:

$$M^\alpha - m_i^\alpha < M^\beta - m_i^\beta$$

i.e.,

$$m_i^\alpha - m_i^\beta > M^\alpha - M^\beta \tag{27}$$

Now

$$\begin{aligned} m_i^\alpha - m_i^\beta &= (r_i^\alpha - r_i^\beta)\overline{m}_i \\ &> \left(\frac{1}{2} - \frac{1}{2n}\right)\overline{m}_i \\ &= \left(\frac{n-1}{2n}\right)\overline{m}_i \\ &\geq \frac{1}{4}\overline{m}_i \end{aligned} \tag{28}$$

[14] Leading to convex costs (see the beginning of Section 5).

[15] Recall that $(u_j^\beta)_{j\in\mathcal{I}} \neq 0$ by assumption, hence (see (3)) we have $(v_j^\beta)_{j\in\mathcal{I}} \neq 0$.

[16] Since, $v_j^\beta > 0$ for some $j \in \mathcal{I}$, it follows that $\sum_{j\in\mathcal{I}} m_j^\beta = M^\beta$.

Consider any $k \in \mathcal{I} \backslash \{i\}$ such that $r_k^\beta > 0$ (we shall deal shortly with the case that no such k exists). The first order conditions of NE imply

$$v_i^\beta \left(\frac{1 - r_i^\beta}{\overline{m}_i} \right) \leq v_k^\beta \left(\frac{1 - r_k^\beta}{\overline{m}_k} \right)$$

Hence

$$\begin{aligned} \overline{m}_i &\geq \frac{v_i^\beta}{v_k^\beta} \left(\frac{1 - r_i^\beta}{1 - r_k^\beta} \right) \overline{m}_k \\ &\geq (v_{min}^\beta / v_{max}^\beta) \left(\frac{1 - 1/2n}{1 - 0} \right) \overline{m}_k \\ &\geq (v_{min}^\beta / v_{max}^\beta) \frac{3}{4} \overline{m}_k \end{aligned}$$

Summing over all k such that $r_k^\beta > 0$, we get

$$|\mathcal{I}| \; \overline{m}_i \geq \frac{3}{4} (v_{min}^\beta / v_{max}^\beta) M^\beta$$

i.e.,

$$\overline{m}_i \geq \frac{3}{4|\mathcal{I}|} (v_{min}^\beta / v_{max}^\beta) M^\beta$$

(When $r_k^\beta = 0$ for all $k \in \mathcal{I} \backslash \{i\}$, we must have $\overline{m}_i \geq M^\beta$ and therefore the above inequality still holds.)

This inequality along with (28), implies

$$\begin{aligned} m_i^\alpha - m_i^\beta &> \frac{1}{4} \frac{3}{4|\mathcal{I}|} (v_{min}^\beta / v_{max}^\beta) M^\beta \\ &= \frac{3}{16|\mathcal{I}|} (v_{min}^\beta / v_{max}^\beta) M^\beta \end{aligned}$$

On the other hand, condition (iii) says

$$M^\alpha - M^\beta \leq \frac{3}{16|\mathcal{I}|} (v_{min}^\beta / v_{max}^\beta) M^\beta$$

Therefore

$$m_i^\alpha - m_i^\beta > M^\alpha - M^\beta$$

establishing the claim. □

The first-order conditions of NE imply that (25) and (26) hold, with β substituted for $\tilde{\alpha}$. This yields[17]

$$\frac{v_i^\alpha v_j^\beta}{v_i^\beta v_j^\alpha} \geq \frac{(1 - r_j^\alpha)(1 - r_i^\beta)}{(1 - r_j^\beta)(1 - r_i^\alpha)} \geq \frac{1 - 1/2n}{1/2} = \frac{2n - 1}{n}$$

Using (ii)(b), we get

$$k_1 k_2 \geq \frac{2n - 1}{n}$$

which contradicts (ii)(a). We conclude that $r_i^\alpha \leq 1/2$ for every α and i, proving the theorem. □

[17] Recalling that (a) $r_j^\beta > r_j^\alpha$, $r_i^\alpha > 1/2$, $r_i^\beta < 1/2n$; (b) $r_j^\beta < 1$ since at least two companies bid on any customer-node in an NE; (c) by (ii) (and the obvious fact that a company bids money on a node only if it places positive value on it) we have $v_i^\alpha > 0, v_i^\beta > 0, v_j^\alpha > 0, v_j^\beta > 0$; (d) both $\overline{m}_i$ and $\overline{m}_j$ are positive by Theorem 1. These conditions together imply that the LHS and RHS of (25) and (26) are all positive.

5.4 Anonymous Valuations Restore Uniqueness

NE are unique even with heterogeneous convex costs and marketing functions that are more general than our canonical example, provided that companies' valuations of clients are identical: $v_i^\alpha = v_i$ for all α and i. This is not an unnatural assumption. It holds when w_{ij}^α and θ_i^α are invariant of α, as in our examples in Section 4.2.

Suppose that costs are given by differentiable and convex functions of total expenditure: $C^\alpha(m^\alpha) \equiv C^\alpha(\overline{m}^\alpha)$ for all $\alpha \in \mathcal{A}$. Further suppose that the marketing impact of α on i can be factored in terms of i's expenditure m_i^α and the total expenditure $\overline{m}_i$.

AV For all $\alpha \in \mathcal{A}$ and $i \in \mathcal{I}$

$$\lambda_i^\alpha(r_i^\alpha, \overline{m}_i) = \frac{\lambda^\alpha(r_i^\alpha)}{f_i(\overline{m}_i)}$$

(when $\overline{m}_i > 0$), where f_i is strictly increasing and λ^α is strictly decreasing. (Recall $r_i^\alpha \equiv m_i^\alpha / \overline{m}_i$.)

Note that, in our canonical case, $\lambda_i^\alpha(r_i^\alpha, \overline{m}_i) = (1 - r_i^\alpha)/\overline{m}_i$ and so (AV) is satisfied. The related examples in footnote 7 also satisfy (AV).

Theorem 6. *Suppose $v_i^\alpha = v_i$ for all α and i, and that assumption AV holds. Then there exists a unique Nash Equilibrium.*

Proof: Let m be an NE. As argued in the proof of Theorem 3, $\overline{m}_i > 0$ for all $i \in \mathcal{I}$, so that the derivatives $\lambda_i^\alpha(r_i^\alpha, \overline{m}_i)$ are well defined.
Step 1: $r_i^\alpha = r_j^\alpha \equiv r^\alpha$ for all $i \in \mathcal{I}, j \in \mathcal{I}$ and $\alpha \in \mathcal{A}$.
Proof: Suppose $r_i^\alpha > r_j^\alpha$ for some α, i, j. Since $\sum_{\beta \in \mathcal{A}} r_i^\beta = 1 = \sum_{\beta \in \mathcal{A}} r_j^\beta$, there exists β such that $r_i^\beta < r_j^\beta$. Since $r_i^\alpha > 0$, the first-order conditions for α at i and j are (where $\xi^\alpha(\overline{m}^\alpha) = (d/d\overline{m}^\alpha) C^\alpha(\overline{m}^\alpha)$)

$$v_i \frac{\lambda^\alpha(r_i^\alpha)}{f_i(\overline{m}_i)} = \xi^\alpha(\overline{m}^\alpha)$$

and

$$v_j \frac{\lambda^\alpha(r_j^\alpha)}{f_j(\overline{m}_j)} \leq \xi^\alpha(\overline{m}^\alpha)$$

which gives

$$\frac{v_i \lambda^\alpha(r_i^\alpha)}{f_i(\overline{m}_i)} \geq \frac{v_j \lambda^\alpha(r_j^\alpha)}{f_j(\overline{m}_j)} \tag{29}$$

Similarly

$$\frac{v_i \lambda^\beta(r_i^\beta)}{f_i(\overline{m}_i)} \leq \frac{v_j \lambda^\beta(r_j^\beta)}{f_j(\overline{m}_j)} \tag{30}$$

From (29) and (30) we obtain

$$\frac{\lambda^\alpha(r_i^\alpha)}{\lambda^\alpha(r_j^\alpha)} \geq \frac{\lambda^\beta(r_i^\beta)}{\lambda^\beta(r_j^\beta)} \tag{31}$$

But, by (AV), LHS of (31) < 1 and RHS of (31) > 1, a contradiction. ■

Let η be another NE and define $s_i^\alpha \equiv \eta_i^\alpha / \overline{\eta}_i$. As shown in Step 1, $s_i^\alpha = s_j^\alpha \equiv s^\alpha$ for all $\alpha \in \mathcal{A}$ and $i, j \in \mathcal{I}$.
Step 2: $r^\alpha = s^\alpha$ for all $\alpha \in \mathcal{A}$.
Proof: Suppose not. W.l.o.g. let $\overline{m} \geq \overline{\eta}$. Clearly there exists $i \in \mathcal{I}$ such that $\overline{m}_i \geq \overline{\eta}_i$ and (since $\sum_{\beta \in \mathcal{A}} r^\beta = 1 = \sum_{\beta \in \mathcal{A}} s^\beta$) there exists $\alpha \in \mathcal{A}$ such that $r^\alpha > s^\alpha \geq 0$. The first order conditions give the following:

$$\frac{v_i \lambda^\alpha(r^\alpha)}{f_i(\overline{m}_i)} = \xi^\alpha(r^\alpha \overline{m}) \tag{32}$$

$$\frac{v_i \lambda^\alpha(s^\alpha)}{f_i(\overline{\eta}_i)} \leq \xi^\alpha(s^\alpha \overline{\eta}) \tag{33}$$

By the convexity of the cost function and the fact that $r^\alpha > s^\alpha$ and $\overline{m} \geq \overline{\eta}$, RHS of (32) $\geq$ RHS of (33), so

$$\frac{\lambda^\alpha(r^\alpha)}{f_i(\overline{m}_i)} \geq \frac{\lambda^\alpha(s^\alpha)}{f_i(\overline{\eta}_i)}$$

But, since $r^\alpha > s^\alpha$ and $\overline{m}_i \geq \overline{\eta}_i$, (AV) implies $\lambda^\alpha(r^\alpha) < \lambda^\alpha(s^\alpha)$ and $f_i(\overline{m}_i) \geq f_i(\overline{\eta}_i)$, contradicting the last displayed inequality. ■

Step 3: $\overline{m} = \overline{\eta}$.

Proof: Suppose w.l.o.g. $\overline{m} > \overline{\eta}$. Then there exists i such that $\overline{m}_i > \overline{\eta}_i$. Also there clearly exists α such that $r^\alpha > 0$. Now, consider the first order conditions (32) and (33). Since $r^\alpha = s^\alpha$ by Step 2, RHS of (32) $\geq$ RHS of (33), which implies $f_i(\overline{\eta}_i) \geq f_i(\overline{m}_i)$. Thus $\overline{\eta}_i \geq \overline{m}_i$ since f_i is strictly increasing by (AV). This is a contradiction. ■

Steps 1, 2 and 3 imply that $m = \eta$, proving Theorem 4. ■

5.4.1 Structure of Nash Equilibrium

It is natural to consider the case where the marketing impact is an anonymous function of expenditures, as in our canonical example.

AVI $\lambda_i^\alpha = \lambda_i^\beta \equiv \lambda_i$ for all α, β and i.

Also assume that λ_i is concave and increasing in m_i^α, as in Section 4.1. In this event, even without the factorization of AV, we can describe an interesting structural feature of NE (though we do not know if they are unique).

Recall that $\xi^\alpha(\overline{m}^\alpha) = (d/d\overline{m}^\alpha)C^\alpha(\overline{m}^\alpha)$.

Theorem 7. *Consider the model stated above. Let m be any NE. Denote $\mathcal{I}(\alpha) = \{i \in \mathcal{I} : m_i^\alpha > 0\}$. There is an ordering $\alpha_1 \leq \alpha_2 \leq \ldots \leq \alpha_n$ of the firms such that $\mathcal{I}(\alpha_1) \subset \mathcal{I}(\alpha_2) \subset \ldots \subset \mathcal{I}(\alpha_n)$. In other words* **the clientele of active firms are always nested**.

Proof:

First we shall show

$$m_i^\alpha > 0 \Rightarrow m_i^\beta > 0$$

for all β, α such that $\xi^\beta(0) \leq \xi^\alpha(0)$. Suppose, to the contrary, that $m_i^\alpha > 0$ and $m_i^\beta = 0$ for some β such that $\xi^\beta(0) \leq \xi^\alpha(0)$. The first order conditions for α and β become

$$v_i \lambda_i(r_i^\alpha, \overline{m}_i) = \xi^\alpha(\overline{m}^\alpha)$$

$$v_i \lambda_i(0, \overline{m}_i) \leq \xi^\beta(0)$$

where $r_i^\alpha \equiv m_i^\alpha/\overline{m}_i > 0$. Since $\xi^\alpha(\overline{m}^\alpha) > \xi^\alpha(0) \geq \xi^\beta(0)$, we have $\lambda_i(r_i^\alpha, \overline{m}_i) > \lambda_i(0, \overline{m}_i)$ contradicting that λ_i is increasing as in Section 4.1. The theorem now easily follows. ■

6 When Externalities become Dominant

6.1 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 $Prob(L_t = Move) = \theta^\alpha_{i_t}$. Let T be the first time $L_t = Stop$ and consider the random variable $\gamma^\alpha_{i_T}(m)$. If $\phi^\alpha(i)$ denotes the conditional expectation $E[\gamma^\alpha_{i_T}(m)|i_0 = i]$, 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 $y^\alpha \sum_{i \in \mathcal{I}} \xi^\alpha_i \phi^\alpha(i)$ which in turn can be expressed as $y^\alpha E[\gamma^\alpha_{i_T}(m)]$, 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^α_i converge to 1 (since the unique NE of our games are determined by v^α). Let us first consider the simple case when $\theta^\alpha_i = \theta^\alpha$ 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

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

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 [5]) states that a Cesaro-convergent sequence is Abel-convergent to the same limit. So, to analyse the limit behavior of v^α_i, 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 $I^\alpha_1, \ldots, I^\alpha_{k(\alpha)}$ and a set of transient states I^α_0. Each recurrent class I^α_s is the support of a unique invariant probability measure μ^α_s.

If the Markov process starts within a recurrent class I^α_s (i.e., $i_0 \in I^\alpha_s$), 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 $E_{\mu^\alpha_s}[f]$.

If it starts at a transient state $i \in I^\alpha_0$, then we may define the first time τ that it enters $\cup_{s \geq 1} I^\alpha_s$. 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 $E_{\mu^\alpha_S}[f]$.

Let us define $\hat{\mu}^{\alpha,i}$ as the expectation $E[\mu^\alpha_S]$, if $i \in I^\alpha_0$ and as μ^α_s if $i \in I^\alpha_s$ ($s \geq 1$). 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 $\hat{\mu}^\alpha \equiv \sum_{i \in \mathcal{I}} \xi^\alpha_i \hat{\mu}^{\alpha,i}$, the Frobenius theorem implies

Theorem 8. *As θ^α tends to 1, v^α_i converges to $y^\alpha E_{\hat{\mu}^\alpha}[\mathbb{1}_i] = y^\alpha \hat{\mu}^\alpha_i$.*

Corollary 2. *Suppose that the graph of the underlying social network is undirected and connected. Further suppose*

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

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^α_i is invariant of i for all α (i.e., each company values all clients equally), w.lo.g. $u^\alpha_i = 1/|\mathcal{I}|$ 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 of the node.

Proof: It is evident that the invariant measure is proportional to the degree. By Theorem 8, $v^\alpha_i = v_i$ converges to the degree of i as θ tends to 1. But, by Section 4.5, m^α_i is proportional to v_i. ∎

It might be useful to illustrate Theorem 8 with a simple example. Consider the network with four nodes, corresponding to the following matrix W^α for firm α:

$$W^\alpha = \begin{pmatrix} 0 & .5 & .5 & 0 \\ .5 & 0 & 0 & .5 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Also let $u_1^\alpha = 7, u_2^\alpha = 8, u_3^\alpha = 10, u_4^\alpha = 11$. Nodes 3 and 4 are clearly absorbent: once reached by the process, they become permanent. Nodes 1 and 2 are transient. Hence $I_0^\alpha = \{1,2\}$, $I_1^\alpha = \{3\}$ and $I_2^\alpha = \{4\}$.

The invariant probability measure μ_1^α (resp. μ_2^α) places all the weight on node 3 (resp. 4). Therefore,

$$\hat{\mu}^{\alpha,3} = \mu_1^\alpha = (0,0,1,0) \text{ and } \hat{\mu}^{\alpha,4} = \mu_2^\alpha = (0,0,0,1)$$

Observing that $prob(i_{t+1} \in \{3,4\} | i_t = i) = 1/2$, when $i \in I_0$, we conclude that, if the process starts in I_0 at time 0, then the first time τ it will reach $\{3,4\}$ is a geometric random variable with parameter $1/2$: $prob(\tau = t) = 1/2^t$, $t = 1,2,3,\ldots$.

As above, let S denote the index of the recurrence class of i_τ. When $i_0 = 1$, then $S = 1$ whenever τ is an odd number and $S = 2$ otherwise. Clearly,

$$prob(S = 2 | i_0 = 1) = \sum_{k=1}^{\infty} prob(\tau = 2k) = \sum_{k=1}^{\infty} 1/4^k = 1/3$$

and thus $\hat{\mu}^{\alpha,1} = 2/3\mu_1^\alpha + 1/3\mu_2^\alpha = (0,0,2/3,1/3)$.

A similar argument shows that $\hat{\mu}^{\alpha,2} = 1/3\mu_1^\alpha + 2/3\mu_2^\alpha = (0,0,1/3,2/3)$, since, if $i_0 = 2$, then the event $\{S = 1\}$ corresponds to the even values of τ.

Clearly

$$\xi^\alpha = \frac{1}{36}(7,8,10,11)$$

which yields

$$\hat{\mu}^\alpha = (0,0,\frac{13}{27},\frac{14}{27})$$

and hence $v_1^\alpha, v_2^\alpha, v_3^\alpha$ and v_4^α converge to $0, 0, 52/3$ and $56/3$.

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:

$$\lim_{\theta \to 1} \theta_i^\alpha(\theta) = 1, \text{ for all } i \tag{34}$$

$$\theta_i^\alpha(\theta) < 1, \text{ for all } i \text{ and } \theta < 1. \tag{35}$$

$$0 < \lim_{\theta \to 1} \frac{1 - \theta_i^\alpha(\theta)}{1 - \theta_1^\alpha(\theta)} = \delta_i^\alpha < \infty \tag{36}$$

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

Theorem 9. *Under (34), (35), (36), 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 the Appendix.

7 Acknowledgements

We are grateful to Sachindra Joshi and Chandan Deep Nath for introducing us to the social network literature (in particular to the references [2, 7]), and for stimulating initial discussions.

8 Appendix

8.1 Proof of Theorem 1

Step 1: Since C^α is strictly increasing and convex, $C^\alpha(x) \to \infty$ as $||x|| \to \infty$. Thus there exists a scalar b such that $C^\alpha(m^\alpha) > U^\alpha(1,1,\dots,1) - U^\alpha(0,0,\dots,0)$ whenever $||m^\alpha|| \geq b$. Define $S^\alpha = \{m^\alpha \in R_+^{\mathcal{I}} : ||m^\alpha|| \leq b\}$. Clearly, no firm α would spend more than b, for it then could be better off spending zero on every customer. W.l.o.g. we may confine α's strategies to the compact convex set S^α.

Step 2: For any $\alpha \in \mathcal{A}$ and $j \in \mathcal{I}$, if $m >> 0$ then $p_i^\alpha(m^\alpha, m^{-\alpha})$ is concave and increasing in m^α for every fixed $m^{-\alpha}$.

Proof of Step 2: Recall

$$\begin{aligned} p_i^\alpha(m) &= F_i^\alpha(m, p_{-i}^\alpha(m)) \\ &= \lim_{t\to\infty} F_i^\alpha(m, p_{-i}^\alpha(m,t)) \end{aligned} \tag{37}$$

where

$$p^\alpha(m, t+1) = F_i^\alpha(m, p_{-i}^\alpha(m,t)) \tag{38}$$

and $p_{-i}^\alpha(m, 0) = 0$.

For brevity, say that a real-valued function $h(m^\alpha, m^{-\alpha}, y)$ defined on a vector space "has property (*)" if it is concave and increasing in m^α for every fixed choice of $m^{-\alpha}$ and y.

Note $F_i^\alpha(m, 0) = F_i^\alpha(m^\alpha, m^{-\alpha}, 0)$ satisfies (*) by assumption AII. Assume $F_i^\alpha(m, p_{-i}^\alpha(m,t))$ satisfies (*). Then one may check that $F_i^\alpha(m, p_{-i}^\alpha(m, t+1))$ also satisfies (*). Indeed this follows from (38) and the obvious fact that the function $G(z, g(z))$ is concave and increasing in z whenever both g and G are concave and increasing.

Step 2 now follows from (37) and an obvious limiting argument.

Step 3: For $\epsilon > 0$, define the game Γ^ϵ by truncating the strategy sets to $S^{\epsilon,\alpha} = S^\alpha \cap \{m^\alpha \in R_+^{\mathcal{I}} : m_j^\alpha \geq \epsilon \ \forall j \in \mathcal{I}\}$. Then Γ^ϵ has an NE.

Proof of Step 3: Obviously $p_i^\alpha(m^\alpha, m^{-\alpha})$ is continuous in $m \equiv (m^\alpha, m^{-\alpha})$ and (by Step 2) concave in m^α. Moreover U^α is continuous, concave and increasing in all its variables by assumption AII. It follows that $\Pi^\alpha(m) = U^\alpha(p^\alpha(m)) - C^\alpha(m^\alpha)$ is continuous in m and concave in m^α. The existence of NE now follows from the standard Nash argument [6].

Step 4: Let $m(\epsilon)$ be an NE of Γ^ϵ and select a subsequence $\epsilon_n \to 0$ so that $m(\epsilon_n) \to m$ as $n \to \infty$. Then m is an NE of Γ.

Proof of Step 4: We need only verify that m is a point of continuity of the payoff functions. This will follow if $m_j \neq 0$ for all $j \in \mathcal{I}$. Suppose, to the contrary, $m_i = 0$ for some i, i.e., $m_i^\tau(\epsilon_n) \to 0$ for all $\tau \in \mathcal{A}$. Let α and β be as in assumption (AIV). By going to a subsequence if necessary, assume $m_i^\alpha(\epsilon_n) \leq m_i^\beta(\epsilon_n)$ for all n. Choose $\delta_n \to 0$ such that $\delta_n / m_i^\tau(\epsilon_n) \to \infty$ as $n \to \infty$, for all $\tau \in \mathcal{A} \backslash \{\alpha\}$ (e.g. take $\delta_n = \max\{\sqrt{m_i^\tau(\epsilon_n)} : \tau \in F \backslash \{\alpha\}\}$). Let α spend δ_n more on i. (This deviation is feasible for large enough n, since $m_i^\alpha < b$ and $\delta_n \to 0$.) The incremental cost of the deviation to α is at most $C_+^\alpha \delta_n$ where C_+^α is the maximum of the right hand derivative of C^α evaluated at (see Step 1) the point $(b, b, \dots, b)$. We will show that α's gain in benefit is strictly more for small enough δ_n. Let $p^\alpha(-), p^\alpha(+) \in R_+^{\mathcal{I}}$ be the probabilities achieved before and after α's unilateral deviation to the extra expenditure δ_n. As shown in Step 2, $p^\alpha(+) \geq p^\alpha(-)$ component-wise. But, since U^α is increasing, the gain in benefit is at least $B[p_i^\alpha(+) - p_i^\alpha(-)]$ where $B = \min\{\partial U^\alpha(p^\alpha)/\partial p_i^\alpha : p^\alpha \in [0,1]^{\mathcal{I}}\} > 0$ with $\partial U^\alpha / \partial p_i^\alpha$ denoting the right-hand derivative of the concave function U^α. Now, denoting by $\overline{\delta}_n \in R_+^{\mathcal{I} \times \mathcal{A}}$ the vector whose $i\alpha^{th}$ component is δ_n and all other components are 0, we have

$$\begin{aligned} p_j^\alpha(+) - p_j^\alpha(-) &= F_i^\alpha(m + \overline{\delta}_n, p_{-i}^\alpha(m+\delta)) - F_i^\alpha(m, p_{-i}^\alpha(m)) \\ &= F_i^\alpha(m + \overline{\delta}_n, p_{-i}^\alpha(m)) - F_i^\alpha(m, p_{-i}^\alpha(m)) \\ &\quad + F_i^\alpha(m + \overline{\delta}_n, p_{-i}^\alpha(m+\delta)) - F_i^\alpha(m + \overline{\delta}_n, p_{-i}^\alpha(m)) \end{aligned}$$

The term $F_i^\alpha(m + \overline{\delta}_n, p_{-i}^\alpha(m+\delta)) - F_i^\alpha(m + \overline{\delta}_n, p_{-i}^\alpha(m))$ is non-negative by the assumption that F_i^α is increasing and the fact that (see Step 2) $p_{-i}^\alpha(m + \overline{\delta}_n) \geq p_{-i}^\alpha(m)$. By AIV the term $F_i^\alpha(m + \overline{\delta}_n, p_{-i}^\alpha(m)) -$

$F_i^\alpha(m, p_{-i}^\alpha(m)$, is at least $K\delta_n$ where K can be chosen arbitrarily large for small enough δ_n; in particular, to ensure that $BK > C_+^\alpha$. But then the gain in payoff is at least $BK\delta_n$ which exceeds the loss $C_+^\alpha\delta_n$, for small enough δ_n. This shows that α can benefit from unilateral deviation at $m(\epsilon_n)$, for small enough ϵ_n, contradicting that $m(\epsilon_n)$ is an NE of Γ^{ϵ_n}. We conclude that $m_i \neq 0$ as was to be shown. ■

8.2 Proof of Theorem 9

It will be convenient to create a micro-model of how the decision $L_t \in \{Stop, Move\}$ is taken in our Markov chain. Before starting the Markov chain, one can, for each state i, consider an infinite sequence of independent decisions $\{L_k^i\}_{k=0,1,2,\ldots}$, with $prob(L_k^i = Move) = \theta_i^\alpha(\theta)$. Each time the process comes to state i the decision to Stop or to Move is taken according to the first unused decision L_k^i. In other words, if N_t^i denotes the number of visits of state i up to time t, we get $L_t = L^{i_t}_{N_t^{i_t}}$.

Let then K_i denote the smallest k such that $L_k^i = Stop$. Clearly K_i is a geometric random variable with parameter $\theta_i^\alpha(\theta)$, so that, for an integer k, $prob(K_i = k) = (1-\theta_i^\alpha)(\theta_i^\alpha)^k$ and $P(K_i > k) = (\theta_i^\alpha)^{k+1}$.

The event $\{i_T = i\}$ coincides then with $\{\exists t | N_t^i \geq K_i \ \& \ N_t^j < K_j, \forall j \neq i\}$.

The ergodic theorem tells us that, as t goes to ∞, $n_t^i \equiv N_t^i/t$ converges almost surely to the random variable $E_{\mu^\alpha}[\mathbb{1}_i] = \mu_i^\alpha$, where μ^α is the unique invariant measure ($\mathcal{I} = I_1$).

Therefore, for all $\epsilon > 0$, there exists N such that $prob(A) > 1-\epsilon$, where $A \equiv \{\forall t > N, \forall i \in \mathcal{I} : |\mu_i^\alpha - n_t^i| < \epsilon\}$. Define also $B \equiv \{K_i > N\}$. Then

$$\begin{aligned}
prob(i_T = i) &\geq prob(\{i_T = i\} \cap A \cap B) \\
&= prob(\{\exists t | tn_t^i \geq K_i \ \& \ tn_t^j < K_j, \forall j \neq i\} \cap A \cap B) \\
&\geq prob(\{\exists t | t(\mu_i^\alpha - \epsilon) \geq K_i \ \& \ t(\mu_j^\alpha + \epsilon) < K_j, \forall j \neq i\} \cap A \cap B) \\
&\geq prob(\{\tfrac{K_i}{\mu_i^\alpha - \epsilon} < \tfrac{K_j}{\mu_j^\alpha + \epsilon}, \forall j \neq i\}) - prob(A^c) - prob(B^c)
\end{aligned}$$

Since $\frac{\mu_i^\alpha - \epsilon}{\mu_j^\alpha + \epsilon}$ is decreasing in ϵ,

$$\lim_{\epsilon \to 0} prob(\{\frac{K_i}{\mu_i^\alpha - \epsilon} < \frac{K_j}{\mu_j^\alpha + \epsilon}, \forall j \neq i\}) = prob(\{\frac{K_i}{\mu_i^\alpha} < \frac{K_j}{\mu_j^\alpha}, \forall j \neq i\}).$$

Therefore

$$prob(i_T = i) \geq prob(\{\frac{K_i}{\mu_i^\alpha} < \frac{K_j}{\mu_j^\alpha}, \forall j \neq i\}) - prob(B^c)$$

Next, with $[x]$ being the integer part of the real number x, we get

$$\begin{aligned}
prob(\{\tfrac{K_i}{\mu_i^\alpha} < \tfrac{K_j}{\mu_j^\alpha}, \forall j \neq i\}) &= \textstyle\sum_{k=0}^\infty (1-\theta_i^\alpha)(\theta_i^\alpha)^k \prod_{j \neq i} (\theta_j^\alpha)^{\left[\frac{\mu_j^\alpha}{\mu_i^\alpha} k\right]+1} \\
&\geq \left(\textstyle\prod_{j \neq i} \theta_j^\alpha\right)(1-\theta_i^\alpha)\textstyle\sum_{k=0}^\infty \left(\prod_{j \in \mathcal{I}} (\theta_j^\alpha)^{\frac{\mu_j^\alpha}{\mu_i^\alpha}}\right)^k \\
&= \left(\textstyle\prod_{j \neq i} \theta_j^\alpha\right) \frac{1-\theta_i^\alpha}{1-\prod_{j \in \mathcal{I}} (\theta_j^\alpha)^{\frac{\mu_j^\alpha}{\mu_i^\alpha}}}
\end{aligned}$$

Using (34) and (36), the limit of this RHS as $\theta \to 1$ is

$$\frac{\delta_i^\alpha \mu_i^\alpha}{\sum_{j \in \mathcal{I}} \delta_j^\alpha \mu_j^\alpha} \tag{39}$$

Similarly,

$$\begin{aligned}
prob(i_T = i) &\leq prob(\{i_T = i\} \cap A \cap B) + prob(A^c) + prob(B^c) \\
&= prob(\{\exists t | tn_t^i \geq K_i \ \& \ tn_t^j < K_j, \forall j \neq i\} \cap A \cap B) + prob(A^c) + prob(B^c) \\
&\leq prob(\{\exists t | t(\mu_i^\alpha + \epsilon) \geq K_i \ \& \ t(\mu_j^\alpha - \epsilon) < K_j, \forall j \neq i\}) + prob(A^c) + prob(B^c) \\
&\leq prob(\{\tfrac{K_i}{\mu_i^\alpha + \epsilon} \leq \tfrac{K_j}{\mu_j^\alpha - \epsilon}, \forall j \neq i\}) + prob(A^c) + prob(B^c)
\end{aligned}$$

Letting ϵ go to 0 yields

$$prob(i_T = i) \le prob(\{\frac{K_i}{\mu_i^\alpha} \le \frac{K_j}{\mu_j^\alpha}, \forall j \ne i\}) + prob(B^c)$$

But

$$\begin{aligned} prob(\{\tfrac{K_i}{\mu_i^\alpha} \le \tfrac{K_j}{\mu_j^\alpha}, \forall j \ne i\}) &= \textstyle\sum_{k=0}^\infty (1-\theta_i^\alpha)(\theta_i^\alpha)^k \prod_{j\ne i} (\theta_j^\alpha)^{\left[\frac{\mu_j^\alpha}{\mu_i^\alpha} k\right]} \\ &\le \left(\textstyle\prod_{j\ne i} \theta_j^\alpha\right)^{-1} (1-\theta_i^\alpha) \textstyle\sum_{k=0}^\infty \left(\textstyle\prod_{j\in\mathcal{I}} (\theta_j^\alpha)^{\frac{\mu_j^\alpha}{\mu_i^\alpha}}\right)^k \\ &= \left(\textstyle\prod_{j\ne i} \theta_j^\alpha\right)^{-1} \frac{1-\theta_i^\alpha}{1-\prod_{j\in\mathcal{I}} (\theta_j^\alpha)^{\frac{\mu_j^\alpha}{\mu_i^\alpha}}} \end{aligned}$$

As θ goes to 1, this also converges to (39) , which therefore is also the limit of $prob(i_T = i)$. ∎

8.3 Cross effects

We show that cross-effects (of p_j^β on p_i^α) can be incorporated, under some constraints, in our general model without endangering the existence of NE.

For $\alpha \in \mathcal{A}$, define the partial order $\overset{\alpha}{\succeq}$ on $R^{\mathcal{I}\times\mathcal{A}}$ by:

$$x \overset{\alpha}{\succeq} y \text{ if and only if } \forall i \in \mathcal{I}, x_i^\alpha \ge y_i^\alpha \text{ and } \forall \beta \ne \alpha, \forall i \in \mathcal{I}, x_i^\beta \le y_i^\beta$$

Assume that the contraction mapping $F(m,p)$ can be written as $F(m,p) = (1-\theta)\gamma(m) + \theta G(p)$ where G is non-expansive (this includes our quasi-linear model). Our assumptions on γ, G are:

AVII γ_i^α is concave in m_i^α, fixing $m_{-i}^{-\alpha}$; and is convex in m_i^β, for $\beta \in \mathcal{A}\backslash\{\alpha\}$, fixing $m_{-i}^{-\beta}$.

AVIII G is affine and *α-increasing* (i.e., $p \overset{\alpha}{\succeq} p'$ implies $G(p) \overset{\alpha}{\succeq} G(p')$).

(It can easily be checked that our canonical example satisfies AVII.)

Lemma 2. *Assume that the utilities U^α are concave and increasing and that AVII, AVIII hold. Then $U^\alpha(p(m^\alpha, m^{-\alpha}))$ is concave in m^α for fixed $m^{-\alpha}$.*

Proof: Indeed, if $\lambda \in [0,1]$, if $m^\alpha = \lambda m'^\alpha + (1-\lambda)m''^\alpha$, then AVII implies

$$\gamma(m^\alpha, m^{-\alpha}) \overset{\alpha}{\succeq} \lambda\gamma(m'^\alpha, m^{-\alpha}) + (1-\lambda)\gamma(m''^\alpha, m^{-\alpha}) \tag{40}$$

Define inductively $p_0 = p'_0 = p''_0 = 0$; $p_{n+1} = (1-\theta)\gamma(m^\alpha, m^{-\alpha}) + \theta G(p_n)$, $p'_{n+1} = (1-\theta)\gamma(m'^\alpha, m^{-\alpha}) + \theta G(p'_n)$ and $p''_{n+1} = (1-\theta)\gamma(m''^\alpha, m^{-\alpha}) + \theta G(p''_n)$.

Clearly $p_0 \overset{\alpha}{\succeq} \lambda p'_0 + (1-\lambda)p''_0$. Now, suppose by induction that $p_n \overset{\alpha}{\succeq} \lambda p'_n + (1-\lambda)p''_n$. Then, since G is affine and satisfies AVIII, $G(p_n) \overset{\alpha}{\succeq} G(\lambda p'_n + (1-\lambda)p''_n) = \lambda G(p'_n) + (1-\lambda)G(p''_n)$. Adding this to (40) yields $p_{n+1} \overset{\alpha}{\succeq} \lambda p'_{n+1} + (1-\lambda)p''_{n+1}$.

Now observe that, as n goes to ∞, $p_n \to p(m^\alpha, m^{-\alpha})$, $p'_n \to p(m'^\alpha, m^{-\alpha})$ and $p''_n \to p(m''^\alpha, m^{-\alpha})$. Therefore

$$p(m^\alpha, m^{-\alpha}) \overset{\alpha}{\succeq} \lambda p(m'^\alpha, m^{-\alpha}) + (1-\lambda)p(m''^\alpha, m^{-\alpha})$$

In particular, $\forall i \in \mathcal{I}$: $p_i^\alpha(m^\alpha, m^{-\alpha}) \ge \lambda p_i^\alpha(m'^\alpha, m^{-\alpha}) + (1-\lambda)p_i^\alpha(m''^\alpha, m^{-\alpha})$. Since U^α just depends on p^α and is increasing and concave, we get

$$\begin{aligned} U^\alpha(p^\alpha(m^\alpha, m^{-\alpha})) &\ge U^\alpha(\lambda p^\alpha(m'^\alpha, m^{-\alpha}) + (1-\lambda)p^\alpha(m''^\alpha, m^{-\alpha})) \\ &\ge \lambda U^\alpha(p^\alpha(m'^\alpha, m^{-\alpha})) + (1-\lambda)U^\alpha(p^\alpha(m''^\alpha, m^{-\alpha})) \end{aligned}$$

∎

Lemma 2 implies the existence of NE in the standard manner (see proof of Theorem 1).

The key property invoked is that G be affine and α-increasing. Thus writing $p_i^\alpha = k_i^\alpha + \sum_{\beta\in\mathcal{A}, j\in\mathcal{I}\setminus\{i\}} w_{ij}^{\alpha\beta} p_j^\beta$ we must have

$$w_{ij}^{\alpha\beta} \geq 0 \text{ if } \beta = \alpha$$

$$w_{ij}^{\alpha\beta} \leq 0 \text{ if } \beta \neq \alpha$$

Of course additional constraints need to be imposed on the $w_{ij}^{\alpha\beta}$ to ensure that G is non-expansive (e.g., $0 \leq \sum_{\beta\in\mathcal{A}, j\in\mathcal{I}\setminus\{i\}} w_{ij}^{\alpha\beta} \leq 1$ will suffice).

References

[1] Social network. http://en.wikipedia.org/wiki/Social_network.

[2] 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.

[3] Pradeep Dubey and Rahul Garg. Competing for customers in a social network. Technical Report RI05009, IBM Research Report, September 2005.

[4] 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.

[5] C. N. Moore. Summability of series. *The American Mathematical Monthly*, 39(2):62–71, February 1932.

[6] John Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Science*, 36:48–49, 1950.

[7] 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.

[8] J. Scott. *Social Network Analysis: A Handbook 2nd Ed.* Sage Publications, London, 2000.

[9] Carl Shapiro and Hal R. Varian. *Information Rules: A Strategic Guide to the Network Economy.* Harvard Business School Press, November 1998.

[10] Oz Shy. *The Economics of Network Industries.* Cambridge University Press, 2001.