Control of Preferences in Social Networks

Chasparis, Georgios; Shamma, Jeff S.

2012

Control of Preferences in Social Networks*

Georgios C. Chasparis† Jeff S. Shamma‡

January 30, 2012

Abstract

We consider the problem of deriving optimal advertising policies for the spread of innovations in a social network. We seek to compute policies that account for i) endogenous network influences, ii) the presence of competitive firms, that also wish to influence the network, and iii) possible uncertainties in the network model. Contrary to prior work in optimal advertising, which also accounts for network influences, we assume a dynamical model of preferences and we compute optimal policies for either a finite or infinite horizon. These optimal policies are related to and extend prior introduced notions of centrality measures usually considered in sociology. We also compute robust optimal policies in the case where the evolution of preferences is affected by misspecified dynamics or uncertainties which can be modeled as external disturbances of the nominal dynamics. Under these perturbed dynamics, we formulate a max-min optimization to compute an optimal policy which is robust to a class of norm-bounded uncertainties. We also show that the optimization exhibits a certainty equivalence property, i.e., the optimal values of the control variables are the same as if there were no uncertainty. Finally, we investigate the scenario where a competitive firm also tries to influence the network. In this case, robust optimal solutions are computed in the form of i) Nash and Stackelberg equilibria, and ii) max-min solutions.

1 Introduction

This paper is concerned with the derivation of optimal advertising strategies in a network of customers whose preferences are affected by both their neighbors and the incentives provided through advertising. The contribution of this paper lies in the inclusion of three important factors in the derivation of an optimal advertising strategy: i) dynamic network effects in the formation of preferences, ii) possible misspecifications/uncertainties in the assumed model of evolution of preferences, and iii) uncertainty in the intentions of a competitive firm that also tries to influence the network.

*This work was supported by AFOSR projects #FA9550-05-1-0321 and #FA9550-09-1-0420.

†G.C. Chasparis is with the Department of Automatic Control, Lund University, Lund 221 00-SE, Sweden. E-mail: georgios.chasparis@control.lth.se. URL: http://www.control.lth.se/chasparis.

‡J.S. Shamma is with the School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332. E-mail: shamma@gatech.edu. URL: http://www.prism.gatech.edu/∼jshamma3.

The literature on optimal advertising starts with the pioneering work of [1] in a monopoly framework and it has been extended to differential games in oligopolies, a detailed survey of which can be found in [2]. The main objective of this line of work, as very well stated in [3], is to *set up an optimal control problem to determine the optimal rate of advertising expenditures over time in a way that maximizes the net profit of the firm.* To this end, prior work has focused on i) the derivation of dynamic models which capture the sales response to advertising, and ii) the computation of an optimal policy of advertising as a function of the sales.

Those models which capture the effect of advertising on sales are usually described by means of a differential or difference equation which describe the evolution of the state (usually the sales rate or the market share) as a function of the state and the advertising expenditures. We will generally assume that firms have some way of knowing or estimating the dynamics of sales response to advertising. The estimation of these dynamics will not be part of this work. Moreover, several sales-to-advertising models are also a function of other properties of the product, such as its price or quality, which will not be considered here.

Prior sales-to-advertising models usually capture the following phenomena: i) advertising effects persist over the current period but diminish with time [1], ii) marginal advertising effects diminish or remain constant with the size of advertising [4], iii) advertising effects diminish with the size of sales [1, 5, 6], iv) advertising effects diminish with the size of competitive advertising [7, 5, 8, 9, 10, 11], and v) advertising effects are affected by word-of-mouth communication (or excess advertising) [12].

Depending on the formulation of sales response to advertising, models have also been categorized in: i) sales response models (where the state is the rate of sales) [1], ii) market share models (where the state is the share of the market) [5], iii) diffusion models (which capture the market growth) [13], and iv) goodwill models (which capture the evolution of advertising capital) [14].

Our model is also related to those models. It exhibits diminishing returns with time in the absence of advertising effort, constant marginal returns with the size of advertising, and diminishing returns with the size of competitive advertising. It extends traditional advertising models by also considering the effect of word-of-mouth communication through a network of interactions similarly to [15, 16]. We model network effects similarly to the model of [16]. However, the analysis here is not restricted to the equilibrium state of the evolution of preferences. Instead, the dynamics of network effects become part of the optimization. Using this model, we are able to derive analytically optimal advertising strategies which are related to and extend prior introduced notions of centrality measures usually considered in sociology [17].

At the same time, we would also like to consider the possibility that we are uncertain about the accuracy of the preferences update, instead of assuming a deterministic update. Usually stochastic extensions of existing models have been considered, e.g., [18, 19]. In this paper, we would like to consider uncertainties that can incorporate possible unmodeled dynamics. Under these perturbed dynamics, we formulate a max-min optimization to compute an optimal policy which is robust to a class of norm-bounded uncertainties. We show, as probably expected, that the optimization exhibits a certainty equivalence property, that is, the optimal values of the control variables are the same as if there were no uncertainty.

Finally, we also investigate the possibility that a competitive firm also tries to influence the network, introducing a second form of uncertainty. In this case, and when the objective of the competitive firm is to maximize its sales, the strategy of the competitive firm may not be known. We will either assume that i) the competitive firm has the form of a competitive fridge which tries to enter the market, introducing a notion of sequential optimization (expressed by a Stackelberg solution), or ii) both firms have the ability of simultaneous play (expressed by a Nash solution). Under these scenarios, we provide a complete characterization of Nash solutions (also Stackelberg solutions) within the set of open-loop strategies. These solutions are also a subset of closed-loop (or Markovian) Nash solutions. A complete characterization of the set of closed-loop Nash solutions is going beyond the scope of this paper, since it is highly case-dependent, i.e., it depends on the class of policies which will be considered reasonable for the scenario of interest. However, the proposed framework can be easily utilized to provide closed-loop Nash solutions when the class of policies, over which we are optimizing, is specified. Finally, we investigate the scenario where firms are also uncertain about the objectives of the competitor, which can be formulated as a max-min optimization.

The remainder of the paper is organized as follows. Section 2 describes the problem under consideration. Section 3 discusses some necessary background on dynamic programming. Section 4 derives finite- and infinite-horizon optimal policies in a monopoly under unperturbed and perturbed preferences update. Section 5 computes Stackelberg and Nash solutions in a duopoly. Finally, Section 6 presents concluding remarks.

Notation: For any vector $x \in \mathbb{R}^n$, where x_i is its ith entry,

- $|x|$ denotes its Euclidean norm,
- $|x|_\infty \triangleq \max\{|x_1|, ..., |x_n|\}$,
- $\max_1^+(x) \triangleq \max\{0, x_1, x_2, ..., x_n\}$,
- $\max_i^+(x) \triangleq \max\left\{\{0, x_1, x_2, ..., x_n\} \setminus \bigcup_{k=1}^{i-1} \max_k^+(x)\right\}$, for $i > 1$,
- for some $\alpha > 0$, $\mathrm{sat}(x; \alpha) \triangleq (y_1, y_2, ..., y_n)$ such that

$$y_i = \begin{cases} \alpha & x_i \geq \alpha \\ x_i & 0 < x_i < \alpha \\ 0 & x_i \leq 0 \end{cases}, \quad i = 1, 2, ..., n.$$

2 Problem Description

2.1 Evolution of preferences

The problem considers a pair of firms $\mathcal{L} = \{a, b\}$ and a finite set of customers or nodes $\mathcal{I} = \{1, 2, ..., n\}$.[1] We will denote a firm by $\ell \in \mathcal{L}$ and a customer by $i \in \mathcal{I}$. We assume that the customers are nodes in a given directed network, which is described by a row stochastic matrix W.[2] The matrix W captures how nodes' proclivities towards the product of either firm a or firm b are affected by its neighbors.

Let $x^\ell_{i,k} \geq 0$ denote the proclivity of node i towards buying the product of firm $\ell \in \{a, b\}$ at time k, and

$$x^\ell_k \triangleq (x^\ell_{1,k}, x^\ell_{2,k}, ..., x^\ell_{n,k}) \in \mathbb{R}^n_+$$

be the vector of proclivities over the whole network. We will refer to this vector as the *state* of firm ℓ and we will denote by $\mathcal{S}^\ell \subset \mathbb{R}^n_+$ the corresponding set of states.

Firm $\ell \in \mathcal{L}$ is able to influence the proclivity of node $i \in \mathcal{I}$ towards its product by marketing its product to node i, e.g., by offering discounts or warranties. Let $u^\ell_{i,k} \geq 0$ denote the amount of funds that firm ℓ spends on marketing its product to node i at time k, and

$$u^\ell_k \triangleq (u^\ell_{1,k}, u^\ell_{2,k}, ..., u^\ell_{n,k}) \in \mathbb{R}^n_+$$

be the vector of funds firm ℓ spends over the set of nodes $\mathcal{I}$. We will refer to this quantity as the *control* of firm ℓ. We will also assume that the amount of funds each firm can spend at any given time cannot be larger than M^ℓ, i.e.,

$$\sum_{i \in \mathcal{I}} u^\ell_{i,k} \leq M^\ell \quad \text{for all } k = 0, 1, ... \tag{1}$$

Let $\mathcal{C}^\ell \subset \mathbb{R}^n_+$ denote the resulting constraint set of controls.

The specific relation between the controls and the states is motivated by the work of [16, 20] on social influence network theory and it is described by the following difference equation:

$$x^\ell_{k+1} = \Theta W x^\ell_k + (I - \Theta)\varphi(u^\ell_k, u^{-\ell}_k) \tag{2}$$

which provides the proclivity of node i at time $k + 1$ as a convex combination of i) a weighted average of the proclivities of the neighbors and ii) the external influence caused by both own and competitive advertising. The notation $-\ell$ denotes the complementary set $\mathcal{L} \backslash \ell$. The matrix Θ is

[1] An extension of the forthcoming analysis to multiple number of firms will be straightforward.

[2] A row stochastic matrix W is a nonnegative matrix which also satisfies $W\mathbf{1} = \mathbf{1}$, i.e., the sum of its entries in any row is equal to 1.

assumed diagonal such that

$$\Theta = \mathrm{diag}\{\theta_1, \theta_2, ..., \theta_n\},$$

with diagonal entries satisfying

$$0 \leq \theta_i < 1, \quad \forall i \in \mathcal{I}. \tag{3}$$

The constraint (3) has a natural interpretation since it implies that there is no node that completely ignores external influence. Furthermore, in the absence of external influence, it also models diminishing returns with time. We will simplify notation by rewriting the dynamics in the form:

$$x_{k+1}^{\ell} = Ax_k^{\ell} + B\varphi(u_k^{\ell}, u_k^{-\ell}), \tag{4}$$

where $A \triangleq \Theta W$ and $B \triangleq I - \Theta$. Variations of this nominal model will also be considered later on in this paper when firms are uncertain about the accuracy of the model.

The function $\varphi : \mathcal{C}^{\ell} \times \mathcal{C}^{-\ell} \to [0, \alpha_1] \times ... \times [0, \alpha_n]$, for some $\alpha_i > 0$, $i \in \mathcal{I}$, maps the control vectors of both firms to a vector of influences over the set of nodes $\mathcal{I}$. It is assumed to be nonnegative and bounded above, i.e., the amount of external influence is finite. We will refer to this function as the *influence function.* We would like function φ to also satisfy the following property:

Property 2.1 *The influence function $\varphi : \mathcal{C}^{\ell} \times \mathcal{C}^{-\ell} \to [0, \alpha_1] \times ... \times [0, \alpha_n]$, for some $\alpha_i > 0$, $i \in \mathcal{I}$, is such that:*

1. $\varphi_i(u_k^{\ell}, u_k^{-\ell}) \geq 0$, *if* $u_{i,k}^{\ell} \geq u_{i,k}^{-\ell}$;
2. $\varphi_i(u_k^{\ell}, u_k^{-\ell}) = 0$, *if* $u_{i,k}^{\ell} < u_{i,k}^{-\ell}$.

In other words, we would like the influence of firm ℓ's advertising to be nonnegative when firm ℓ is investing more on advertising than its competitor, and zero, otherwise. That is, a customer would be influenced towards either one of the firms depending on the relative size of their advertising.

One candidate function which satisfies the above property is the following:

$$\varphi_i(u_k^{\ell}, u_k^{-\ell}) \triangleq \mathrm{sat}(u_{i,k}^{\ell} - u_{i,k}^{-\ell}; \alpha_i) \tag{5}$$

for some $\alpha_i > 0$, $i = 1, 2, ..., n$.

We will refer to the above model as ***duopoly***. When, instead, $u_{i,k}^{-\ell} \equiv 0$ for all $i \in \mathcal{I}$ and $k = 0, 1, ...,$ we will refer to this model as ***monopoly***.

The proposed update of preferences exhibits diminishing returns to both *own* and *competitive* advertising, which is due to the definition of the influence function. It also exhibits diminishing returns with time, due to the definition of the matrix Θ. Finally, it models the effect of word-of-mouth (or excess) advertising due to the assumed network of connections.

2.2 Objective

The *utility* of firm $\ell \in \mathcal{L}$ at time k is defined as:

$$g(x_k^\ell, u_k^\ell) = V(x_k^\ell) - C(u_k^\ell) \tag{6}$$

where we assume that the *reward* is linear with the proclivities of the nodes, i.e.,

$$V(x_k^\ell) = v^{\mathrm{T}} x_k^\ell,$$

for some vector $v \in \mathbb{R}_+^n$, and the *cost* is linear with the funds spent on advertising, i.e.,

$$C(u_k^\ell) = c^{\mathrm{T}} u_k^\ell,$$

for some $c \in \mathbb{R}_+^n$.

For some discount factor $\beta \in (0, 1)$, the *objective* of firm ℓ has the following form

$$\max_{\pi^\ell \in \Pi^\ell} \left\{ J_{\pi^\ell}(x) \triangleq \lim_{N \to \infty} \sum_{k=0}^{N-1} \beta^k g(x_k^\ell, \mu_k^\ell(x_k^\ell)) \right\} \tag{7}$$

over the set of infinite sequences of policies Π^ℓ with elements $\pi^\ell = (\mu_0^\ell, \mu_1^\ell, ...)$ where μ_k^ℓ is a function from the set of states $\mathcal{S}$ to the set of controls $\mathcal{C}$. The above optimization is subject to the dynamics (4). Later on, we are also going to consider variations of this optimization, especially when dynamics (4) are perturbed and robust optimal policies need to be derived.

For the remainder of the paper, the proposed advertising model characterized by the dynamics of (4) and the utility function (6) will be denoted by $\mathscr{M}$.

2.3 Assumptions and preliminaries

For the remainder of the paper, we are also going to consider the following assumptions:

Assumption 2.1 $\beta v^{\mathrm{T}} B - c^{\mathrm{T}} > 0$.

This assumption can be written equivalently as

$$\beta v_i (1 - \theta_i) - c_i > 0, \quad i = 1, 2, ..., n.$$

It implies that, for every unit of advertising effort, *the discounted return received from each node is strictly greater than the corresponding cost.* This is a reasonable assumption and it is also related to the existence of a non-degenerate solution to the optimization problems considered here.

Assumption 2.2 $\alpha_i^\ell \geq M_i^\ell$ *for all* $i \in \mathcal{I}$ *and* $\ell \in \mathcal{L}$.

This assumption implies that each customer's capacity of getting influenced through advertising is larger than the advertising power of each firm. This is not a necessary assumption for the existence of solutions, however, it simplifies the following analysis. The derivation of the corresponding solutions in case Assumption 2.2 does not hold is also straightforward and qualitatively remains identical.

In the presentation of the model, we have implicitly assumed that the evolution of preferences is governed by identical dynamics for both firms. This assumption allows for a cleaner presentation of the analysis, however, as it will become obvious later, it does not change qualitatively the solutions.

We also assume that the utility functions of both firms are identical. This implies that benefits and costs are materialized as a function of the proclivities and investments similarly for both firms. This is a reasonable assumption, however, the following analysis can be easily modified to include the case of different utility functions.

Note, finally, that the proposed preferences update (4) constitutes a linear time-invariant system with bounded inputs. It is straightforward to show that the above system is *input-output stable* in the sense that there exists nonnegative constants ζ, θ such that the solution to the difference equation, denoted $x(k, x_0, u)$, satisfies $|x(k, x_0, u)| \leq \zeta + \theta \|u\|_\infty$, where $\|u\|_\infty \triangleq \sup\{|u_k| : k \in \mathbb{Z}_+\}$. This is due to the fact that W is a row stochastic matrix and Θ satisfies the constraint (3). The constraint (3) on matrix Θ also implies the *controllability* (cf., [21]) of the system (A, B), simply because $\operatorname{rank}(B) = \operatorname{rank}(I - \Theta) = n$.

2.4 Alternative models and discussion

The dynamics of preferences (4) is based on the assumption that agents are bounded rational, since their preferences are a weighted average of neighbors' preferences. Full rationality instead may not necessarily lead to better models due to the resulting computational complexity. A similar model in the context of evolution of preferences without external influence has also been considered by [22, 23] to study the diffusion of innovations and norms in a social network. This model has also been related to alternative measures of centrality as discussed in [17, 24].

In this paper, we modified the model used by [22, 23] to include the possibility of an external control influence (4), e.g., due to advertising effects. The proposed model bears similarities with several previously introduced advertising models, e.g., the goodwill models of [14], new product diffusion models [13] or extensions of the Vidale-Wolfe model [1]. In the following subsections we discuss some of the similarities and differences between these models with the proposed $\mathscr{M}$.

2.4.1 Comparison with goodwill models

Advertising goodwill models (see, e.g., [2, Section 3.5]) capture the evolution of the advertising capital. For example, the advertising goodwill model introduced in the seminal paper [14] assumes the following dynamics

$$\dot{G}(t) = u(t) - \delta G(t), \tag{8a}$$

where $G(t)$ here represents the advertising capital. The main difference with the proposed model $\mathscr{M}$ is that the latter includes directly the interpersonal influences through the assumed communication network modeling a form of word-of-mouth communication. Note also that the control input or advertising effort u influences directly the advertising capital. Similar is the assumption in $\mathscr{M}$, where the advertising effort directly influences the preferences of all nodes. As we will see later, this is not necessarily the case in other advertising models, where the advertising effort only applies to the *undecided* part of the population. In other words, both $\mathscr{M}$ and the goodwill models investigate situations where the product is recently launched in the market and all customers are willing to revise their preferences regardless of prior preferences.

The dynamics (8a) can also be modified to include the possibility of multiple firms, e.g., the models in [25, 26]. For example, the model considered in [26] assumes

$$\dot{G}_i(t) = \sqrt{u_i(t)} - \delta G_i(t), \quad G_i(0) = G_{i0} > 0, \quad i \in \{1, 2\}, \tag{8b}$$

and the sales rate x_i (similarly to the proposed vector of proclivities) depends on the advertising capital of both firms, i.e., $x_i = x_i(G_1, G_2)$, where $\partial x_i / \partial G_i > 0$ and $\partial x_i / \partial G_j < 0$ for $i \neq j$.

Note that the square root of the control input in (8b), which has also been used in other advertising models (see, e.g., [5]), captures diminishing returns with the size of advertising effort. Alternatively, diminishing returns can also be modeled indirectly by considering a squared cost in the utility function. For example, in [6] the term u_i^2 is considered instead in the cost function, or in [27] more general non-linear functions of u_i are considered which are convex increasing. In the proposed model $\mathscr{M}$, diminishing returns with the advertising effort are modeled indirectly by assuming the saturation effect of the influence function.

A squared cost term in the utility functional could also be included in the proposed model in comparison to the initially proposed model $\mathscr{M}$ of Section 2.1. For example, an alternative utility functional that incorporates diminishing returns with the size of advertising could be:

$$g(x_k^\ell, u_k^\ell) = v^{\mathrm{T}} x_k^\ell - \left(u_k^\ell\right)^{\mathrm{T}} C u_k^\ell \tag{9}$$

where $C \triangleq \operatorname{diag}(c)$, i.e., C is a diagonal matrix where the diagonal entries coincide with the entries of the vector c. Some of the nice analytical properties of $\mathscr{M}$ are also shared by the above quadratic cost function (9), such as the forthcoming analytical solution of the monopoly optimization problem.

2.4.2 Comparison with market-share response models

The previously described goodwill advertising models and the proposed model $\mathscr{M}$ differ from the class of market-share response models emanating from the model of Vidale-Wolfe [1]. An extension of this model to a duopoly has been considered by [7] and is described by:

$$\dot{x}_i = (1 - x_i - x_j)u_i - \delta_i x_i, \quad x_i(0) = x_{i0}, \quad i, j \in \{1, 2\}, \quad i \neq j. \tag{10}$$

A small modification can also account for *excess advertising* effects due to word-of-mouth influences in the population, such as the model in [28] described by

$$\dot{x}_i = (1 - x_i - x_j)u_i - \delta_i x_i + e_i(u_i - u_j)(x_i + x_j), \quad x_i(0) = x_{i0}, \quad i, j \in \{1, 2\}, \quad i \neq j, \tag{11}$$

where the last term represents the persons switching from firm j to i as a result of the word-of-mouth processes.

Contrary to both $\mathscr{M}$ and the goodwill advertising models, where the advertising effort applies directly to the whole population, in the market-share response generalizations of Vidale-Wolfe's model [1], the control applies only to the *undecided* part of the population. The last term of the dynamics (11), which models excess advertising, applies to the *decided* part of the market and models transfers due to excess of advertising. This term also resembles the influence function φ considered in $\mathscr{M}$, where the influence on a node depends only on the excess part of the advertising efforts at that node.

Note, however, that a small modification of the proposed model $\mathscr{M}$ can account for behaviors that are present in the market-share models [1]. For example, if we instead consider the following influence function:

$$\varphi_i(u_k^\ell, u_k^{-\ell}) \triangleq \operatorname{diag}\left(\alpha^\ell \mathbf{1} - x_k^{-\ell}\right) u^\ell - \operatorname{diag}\left(\alpha^\ell \mathbf{1} - x_k^\ell\right) u^{-\ell}. \tag{12}$$

then the advertising efforts of either firm applies only on the part of the market which is either undecided or has different preferences. When we assume the alternative dynamics with the influence function (12), then an analytical derivation of a closed-form solution, even for the monopoly framework, is not feasible any more. In the forthcoming analysis, we will only consider the initially proposed influence function which provides closed-form solutions, however future work may include alternative forms of the influence function that accept only computational solutions.

Similar remarks also hold for the models emanating from the Lanchaster model of combat, such as the models of [29, 8, 9, 10, 11]. The main difference of Lanchester models with the Vidale-Wolfe models is that in the latter ones the effect of competitive advertising onto the market share is indirectly included (through the undecided portion of the market). Instead, in the Lanchester models, the effect of competitive advertising is directly included in the dynamics of market share.

This discussion reveals the flexibility of the proposed model $\mathscr{M}$ to incorporate alternative behaviors or modeling ideas which have already been discussed in prior literature. In several cases though, it is also desirable that a sales-to-advertising model also provides closed-form optimal solutions which avoids computational burdens. The proposed model $\mathscr{M}$ and its extensions herein exhibit most of the observed phenomena of sales-to-advertising models and, as we will discuss later, it provides attractive closed-form expressions of optimal strategies under several scenarios.

3 Dynamic Programming Background

The notation and part of the analysis in this section follows [30].

3.1 The dynamic programming algorithm

Denote by $\mathcal{J}$ the set of all extended real-valued functions of the form $J : \mathcal{S} \to \mathbb{R}^*$, defined on the state space $\mathcal{S}$ and taking values on the extended real line $\mathbb{R}^* = [-\infty, +\infty]$.

For some time horizon $N \in \mathbb{N}$, consider the generic finite-horizon optimization problem:

$$\max_{\pi \in \Pi} \left\{ J_{N,\pi}(x_0) \triangleq E\left\{ g(x_N) + \sum_{k=0}^{N-1} \beta^k g(x_k, \mu_k, w_k) \right\} \right\} \tag{13}$$

over any admissible policy $\pi = \{\mu_0, \mu_1, ..., \mu_{N-1}\} \in \Pi$, where $\mu_k \in \mathcal{M}$ for all k, and $\mathcal{M}$ is the set of functions from the set of states $\mathcal{S}$ to the set of controls $\mathcal{C}$. Furthermore, $g(x_N)$ defines the cost at the final stage, which depends only on the final state x_N.

The above optimization is subject to the system dynamics

$$x_{k+1} = f(x_k, u_k, w_k),$$

where $\{w_k\}$ denotes a noise sequence taking values in a measurable space $(\mathcal{W}, \mathfrak{F})$. Denote $J_N^*(x)$ the optimal value of the N-stage objective function. Finally, assume that

$$|g(x, u, w)| < \infty, \quad \text{for all } x \in \mathcal{S}, u \in \mathcal{C}, w \in \mathcal{W}.$$

For any function $J \in \mathcal{J}$, define the following function

$$(TJ)(x) \triangleq \max_{u \in \mathcal{C}(x)} E\{g(x, u, w) + \beta J(f(x, u, w))\}, \quad x \in \mathcal{S}.$$

Note that $(TJ)(\cdot)$ is the *optimal value function* for the one stage problem that has stage cost g and terminal cost βJ.

Also, we will denote by T^k the composition of the mapping T with itself k times; i.e., for all

$k = 1, 2, ...$, we write

$$(T^k J)(x) = (T(T^{k-1}J))(x), \quad x \in \mathcal{S}.$$

For convenience, we also write $(T^0 J)(x) = J(x)$.

Similarly, for any function $J \in \mathcal{J}$ and any policy $\mu : \mathcal{S} \to \mathcal{C}$, we denote:

$$(T_\mu J)(x) \triangleq E\{g(x, \mu(x), w) + \beta J(f(x, \mu(x), w))\}. \tag{14}$$

Again, $T_\mu J$ may be viewed as the cost function associated with the policy μ for the one-stage problem that has stage cost g and terminal cost βJ.

The **dynamic programming algorithm** (DP) is the following algorithm; for any $k = 1, ..., N$ compute

$$J_k(x) = (TJ_{k-1})(x), \tag{15}$$

with initial condition $J_0(x) = g(x)$. The last step of the DP algorithm provides the N-stage value, $J_N(x)$, $x \in \mathcal{S}$.

Define

$$H(x, u, J) \triangleq E\left\{g(x, u, w) + \beta J(f(x, u, w))\right\}. \tag{16}$$

Assumption 3.1 *The above sequence $\{J_k\} \subset \mathcal{J}$ is a non-decreasing sequence satisfying $H(x, u, J_1) < \infty$, and*

$$\lim_{k\to\infty} H(x, u, J_k) = H(x, u, \lim_{k\to\infty} J_k),$$

for all $x \in \mathcal{S}$ and $u \in \mathcal{C}$.

The above assumption excludes problems where exchangeability of expectation with the limit is not possible. This assumption is satisfied when we consider a monotonously increasing sequence of functions $\{J_k\}_k$ in $\mathcal{J}$ and also the functions J_k are measurable with respect to the probability measure under consideration. This will be due to the Lebesgue's Increasing Convergence Theorem (cf., [31]).

Proposition 3.1 (Optimality of DP) *Let Assumption 3.1 hold, and assume that $J_{k,\pi}(x) < \infty$ for all $x \in \mathcal{S}$, $\pi \in \Pi$, and $k = 1, 2, ..., N$. Then*

$$J_N^* = T^N(J_0).$$

Proof. See Proposition 3.1 in [30]. □

3.2 Infinite horizon problems

Consider now the infinite horizon optimization problem:

$$\max_{\pi \in \Pi} \left\{ J_\pi(x_0) = \lim_{N \to \infty} E \left\{ \sum_{k=0}^{N-1} \beta^k g(x_k, \mu_k(x_k), w_k) \right\} \right\}, \tag{17}$$

over any admissible infinite policy $\pi = \{\mu_0, \mu_1, ...\}$ and subject to the system dynamics

$$x_{k+1} = f(x, u, w).$$

Let us also define the optimal value of this problem as

$$J^*(x) \triangleq \sup_{\pi \in \Pi} J_\pi(x). \tag{18}$$

The following is a condition on the optimal stationary policy.

Proposition 3.2 (Optimal stationary policy) *Consider the infinite horizon optimization problem of (17) and assume that*

$$J_0(x) \leq H(x, u, J_0), \quad \forall x \in \mathcal{S}, \quad \forall u \in \mathcal{C}$$

where $J_0(x) = g(x)$. *Then, the optimal value of the infinite horizon optimization problem is*

$$J^*(x) = \lim_{N \to \infty} J_N(x). \tag{19}$$

where $J_N(x)$ *is the* N*-th stage value of the dynamic programming algorithm. Let also Assumption 3.1 hold. Then, a stationary policy* $\pi^* = (\mu^*, \mu^*, ...) \in \Pi$ *is optimal if and only if*

$$T_{\mu^*}(J_{\pi^*}) = T(J_{\pi^*}). \tag{20}$$

Proof. See Proposition 5.5 in [30]. □

4 Optimal Policy in Monopoly

In this section, we compute the optimal policy of a firm when there is no competitive firm, and also the dynamics are either a) unperturbed, or b) perturbed. Since we consider a single firm, we will skip the superscript ℓ for the remainder of this section.

4.1 Unperturbed dynamics

The dynamics we consider in this section are described by (4) with $u_k^{-\ell} \equiv 0$, i.e.,

$$x_{k+1} = Ax_k + B\varphi(u_k) \triangleq f(x_k, u_k). \tag{21}$$

In the remainder of the section we compute the optimal policy for the 1) finite-horizon, and 2) infinite-horizon optimization problem.

First, define:

$$\tilde{A}_k \triangleq \sum_{j=0}^{k} \beta^j A^j$$

and

$$h_{k+1}^{\mathrm{T}} \triangleq \beta v^{\mathrm{T}} \tilde{A}_k B - c^{\mathrm{T}},$$

for $k = 0, 1, \ldots.$ Note that $\tilde{A}_0 = I$ and $h_1^{\mathrm{T}} = \beta v^{\mathrm{T}} B - c^{\mathrm{T}}$.

Before computing the solutions to the finite- and infinite-horizon optimization problems, note that:

Claim 4.1 *$v^{\mathrm{T}} \tilde{A}_{k+1} \geq v^{\mathrm{T}} \tilde{A}_k$ for all $k = 0, 1, \ldots.$*

Proof. First note that

$$\begin{aligned} v^{\mathrm{T}} \tilde{A}_{k+1} &= v^{\mathrm{T}} \sum_{j=0}^{k+1} \beta^j A^j \\ &= v^{\mathrm{T}} \sum_{j=0}^{k} \beta^j A^j + v^{\mathrm{T}} \beta^{k+1} A^{k+1} \geq v^{\mathrm{T}} \tilde{A}_k. \end{aligned}$$

where the last inequality results from the fact that all the entries of matrix A are nonnegative. □

4.1.1 Finite-horizon optimization

We first consider the finite-horizon optimization

$$\max_{\pi \in \Pi} \left\{ J_\pi(x_0) \triangleq g(x_N) + \sum_{k=0}^{N-1} \beta^k g(x_k, \mu_k(x_k)) \right\}. \tag{22}$$

where $g(x) \triangleq v^{\mathrm{T}} x$ defines the utility at the last stage.

Proposition 4.1 (Nth stage optimal policy for monopoly) *Consider the finite horizon optimization problem (22) under the dynamics (21). The Nth stage optimal value of the dynamic* *programming iteration, is*

$$J_N^*(x) = v^{\mathrm{T}} \tilde{A}_N x + \sum_{k=0}^{N-1} \beta^k h_{N-k}^{\mathrm{T}} u_{N-k}^*. \tag{23}$$

The optimal control at time k, for $k = 0, 1, ..., N-1$, is $u_{N-k}^* = (u_{1,N-k}^*, ..., u_{n,N-k}^*)$, where

$$u_{i,N-k}^* = \begin{cases} M & i = \arg\max_1^+ (h_{N-k}) \\ 0 & \textit{otherwise.} \end{cases} \tag{24}$$

Proof. We are going to show the statement by induction. According to the dynamic programming algorithm, the k-th stage optimal value is

$$J_k(x) = \max_{u_k \in \mathcal{C}(x)} \{g(x, u_k) + \beta J_{k-1}(f(x, u_k))\}$$

where $J_0(x) = g(x) = v^{\mathrm{T}} x$. By applying the operator T to J_0, we get the optimal value for the first stage, which is

$$\begin{aligned} J_1(x) &= (TJ_0)(x) \\ &= \max_{u_1 \in \mathcal{C}(x)} \{g(x, u_1) + \beta J_0(f(x, u_1))\} \\ &= \max_{u_1 \in \mathcal{C}(x)} \left\{v^{\mathrm{T}} x - c^{\mathrm{T}} u_1 + \beta v^{\mathrm{T}}(Ax + Bu_1)\right\} \\ &= \max_{u_1 \in \mathcal{C}(x)} \left\{(v^{\mathrm{T}} + \beta v^{\mathrm{T}} A)x + (\beta v^{\mathrm{T}} B - c^{\mathrm{T}})u_1\right\} \\ &= v^{\mathrm{T}} \tilde{A}_1 x + h_1^{\mathrm{T}} u_1^*. \end{aligned}$$

where the optimal stage control is $u_1^* = (u_{1,1}^*, ..., u_{n,1}^*)$ such that

$$u_{i,1}^* = \begin{cases} M & i = \arg\max_1^+ (h_1) \\ 0 & \text{otherwise.} \end{cases} \tag{25}$$

Note that the value $J_1(\cdot)$ is given by expression (23) if we set $N = 1$ and the optimal stage control u_1^* is given by expression (24) if we set $N = 1$ and $k = 0$.

Assume that the value iteration for the N-step optimization horizon gives (23), i.e.,

$$J_N(x) = v^{\mathrm{T}} \tilde{A}_N x + \sum_{k=0}^{N-1} \beta^k h_{N-k}^{\mathrm{T}} u_{N-k}^* \tag{26}$$

where $u^*_{N-k} = (u^*_{1,N-k}, ..., u^*_{n,N-k})$ is such that

$$u^*_{i,N-k} = \begin{cases} M & i = \arg\max^+_1 (h_{N-k}) \\ 0 & \text{otherwise,} \end{cases}$$

for $k = 0, 1, ..., N-1$.

Consider now an $(N+1)$-step optimization horizon. The value at $N+1$ is:

$$\begin{aligned} J_{N+1}(x) &= (TJ_N)(x) \\ &= \max_{u_{N+1}\in\mathcal{C}} \{g(x, u_{N+1}) + \beta J_N(f(x, u_{N+1}))\} \\ &= v^{\mathrm{T}} \left(I + \beta \tilde{A}_N A\right) x + \max_{u_{N+1}\in\mathcal{C}} h^{\mathrm{T}}_{N+1} u_{N+1} + \beta \sum_{k=0}^{N-1} \beta^k h^{\mathrm{T}}_{N-k} u^*_{N-k} \\ &= v^{\mathrm{T}} \tilde{A}_{N+1} x + \sum_{k=0}^{N} \beta^k h^{\mathrm{T}}_{N+1-k} u^*_{N+1-k} \\ &= v^{\mathrm{T}} \tilde{A}_{k+1} x + \sum_{i=1}^{k+1} \beta^i \left(\beta v^{\mathrm{T}} B \tilde{A}_{k-i+1} - c^{\mathrm{T}}\right) u^*_{k-i+1} \end{aligned} \tag{27}$$

where $u^*_{N+1} = (u^*_{1,N+1}, ..., u^*_{n,N+1})$ is such that

$$u^*_{i,N+1} = \begin{cases} M & i = \arg\max^+_1 (h_{N+1}) \\ 0 & \text{otherwise,} \end{cases} \tag{28}$$

for $i = 1, 2, ..., n$.

Thus, we showed that the values of the dynamic programming iteration are provided by equation (23).

Finally, to show optimality of the dynamic programming iteration, subtract equations (26) from (27) to get:

$$J_{N+1}(x) - J_N(x) = v^{\mathrm{T}} \left(\tilde{A}_{N+1} - \tilde{A}_N\right) x + \sum_{k=0}^{N-1} \beta^k \left(h^{\mathrm{T}}_{N+1-k} u^*_{N+1-k} - h^{\mathrm{T}}_{N-k} u^*_{N-k}\right) + \beta^N h^{\mathrm{T}}_1 u^*_1.$$

By Claim 4.1, we have that

$$v^{\mathrm{T}} \left(\tilde{A}_{N+1} - \tilde{A}_N\right) x \geq 0 \quad \text{for all } x \in \mathcal{S}.$$

Given also Assumption 2.1 and the form of optimal control (28), we get that

$$h^{\mathrm{T}}_{N+1} u^*_{N+1} \geq h^{\mathrm{T}}_N u^*_N \geq ... \geq h^{\mathrm{T}}_1 u^*_1 > 0.$$

Therefore, $J_{N+1}(x) \geq J_N(x)$ for all $x \in \mathcal{S}$ and Assumption 3.1 is satisfied. Then, by Proposition 3.1, the dynamic programming iteration provides the optimal value of the finite-horizon optimization (22). $\Box$

4.1.2 Infinite-horizon optimization

We would like to solve the following optimization problem:

$$\max_{\pi \in \Pi} \left\{ J_\pi(x_0) \triangleq \lim_{N \to \infty} \sum_{k=0}^{N-1} \beta^k g(x_k, \mu_k(x_k)) \right\} \tag{29}$$

subject to the discrete-time dynamics (21).

Before we compute the solution to the infinite horizon optimization problem, recall the definition of $H(x, u, J)$ from (16). Given also that $J_0(x) = v^{\mathrm{T}} x$, it is straightforward to show under Assumption 2.1 that:

Claim 4.2 *$J_0(x) \leq H(x, u, J_0)$, for all $x \in \mathcal{S}$ and $u \in \mathcal{C}(x)$.*

Note also that:

Lemma 4.1 *The matrix $(I - \beta A)$ is non-singular for any $\beta \in (0, 1)$.*

Proof. Note that by construction, $(I - \beta A)$, is strictly diagonally dominant,[3] since the magnitude of its i-th diagonal entry $1 - \beta\theta_i w_{ii}$ satisfies

$$\begin{aligned} 1 - \beta\theta_i w_{ii} &= 1 - \beta\theta_i (1 - \sum_{j \neq i} w_{ij}) \\ &= 1 - \beta\theta_i + \beta \sum_{j \neq i} \theta_i w_{ij} > \beta \sum_{j \neq i} \theta_i w_{ij}, \end{aligned}$$

i.e., it is strictly larger than the sum of magnitudes of all non-diagonal entries of the ith row. By Levy-Desplanques theorem (cf., [32]) the matrix $(I - \beta A)$ is non-singular. $\Box$

Lemma 4.2 *Let $\beta \in (0, 1)$ and $A \in \mathbb{R}^{n \times n}$ such that $(I - \beta A)$ is non-singular. Then*

$$\tilde{A}_k = \sum_{j=0}^{k} \beta^j A^j = (I - \beta A)^{-1} (I - \beta^{k+1} A^{k+1}), \tag{30}$$

[3] A matrix is *strictly diagonally dominant* if in every row of the matrix, the magnitude of the diagonal entry in that row is larger than the sum of the magnitudes of all the other (non-diagonal) entries in that row.

$k = 0, 1, \ldots$. *Furthermore, if* $\lim_{k\to\infty} A^k$ *exists, then*

$$\tilde{A}_\infty \triangleq \sum_{j=0}^{\infty} \beta^j A^j = (I - \beta A)^{-1}.$$

Proof. To show the first statement, simply multiply from the left with $(I - \beta A)$. The second statement is a direct consequence of (30) if we take the limit as $k \to \infty$. □

Define also:

$$h_\infty^{\mathrm{T}} \triangleq \beta v^{\mathrm{T}} \tilde{A}_\infty B - c^{\mathrm{T}}.$$

Proposition 4.2 (Optimal Stationary Policy in Monopoly) *Consider the infinite horizon optimization problem (29) under the deterministic and unperturbed dynamics (21). Then, the stationary policy* $\pi^* = (\mu^*, \mu^*, \ldots) \in \Pi$*, such that* $\mu^*(x) = (\mu_1^*, \mu_2^*, \ldots, \mu_n^*)$ *with*

$$\mu_i^* = \begin{cases} M & i = \arg\max_1^+ (h_\infty) \\ 0 & \textit{otherwise} \end{cases} \tag{31}$$

for $i \in \mathcal{I}$*, is an optimal policy for the infinite horizon optimization problem. Furthermore, the optimal infinite value is*

$$J^* = v^{\mathrm{T}} \tilde{A}_\infty x + \frac{M}{1-\beta} \max_1^+ (h_\infty). \tag{32}$$

Proof. Due to Claim 4.2, we have

$$J_0(x) \leq H(x, u, J_0)$$

for all $x \in \mathcal{S}$ and $u \in \mathcal{C}(x)$. Also, as we showed in the proof of Proposition 4.1, due to Claim 4.1 and Assumption 2.1, $J_{k+1}(x) \geq J_k(x)$ for every $x \in \mathcal{S}$. Thus, Assumption 3.1 is satisfied and, according to Proposition 3.2, in order to show that the stationary policy $\pi^* = (\mu^*, \mu^*, \ldots)$ is optimal, it suffices to show that

$$T_{\mu^*}(J_{\pi^*}) = T(J_{\pi^*}).$$

First, we compute $J_{\pi^*}(x)$: Similarly to Proposition 4.1 and taking into account (30), the stationary policy π^* establishes the following sequence of values

$$\begin{aligned} J_{N,\pi^*} &= v^{\mathrm{T}} \tilde{A}_N x + \sum_{k=0}^{N-1} \beta^k h_{N-k}^{\mathrm{T}} \mu^* \\ &= v^{\mathrm{T}} \tilde{A}_\infty (I - \beta^{N+1} A^{N+1}) x + \\ &\quad \sum_{k=0}^{N-1} \beta^k \left(\beta v^{\mathrm{T}} \tilde{A}_\infty (I - \beta^{N-k} A^{N-k}) B - c^{\mathrm{T}} \right) \mu^* \end{aligned}$$

$$= v^{\mathrm{T}}\tilde{A}_{\infty}x + \sum_{k=0}^{N-1}\beta^k h_{\infty}^{\mathrm{T}}\mu^* - \beta^{N+1}v^{\mathrm{T}}\tilde{A}_{\infty}A^{N+1}x - \beta^{N+1}v^{\mathrm{T}}\tilde{A}_{\infty}\sum_{k=0}^{N-1}A^{N-k}B\mu^*.$$

Note that

$$\sum_{k=0}^{N-1}A^{N-k}B\mu^* = \sum_{k=1}^{N}A^kB\mu^* = \sum_{k=1}^{N}W^k\Theta^k(I-\Theta)\mu^*.$$

Since the diagonal entries of Θ satisfy $0 \le \theta_i < 1$ for every $i \in \mathcal{I}$ and μ^* is bounded, the above series is convergent. Therefore, we have

$$J_{\pi^*} \triangleq \lim_{k\to\infty} J_{k,\pi^*} = v^{\mathrm{T}}\tilde{A}_{\infty}x + \frac{1}{1-\beta}h_{\infty}^{\mathrm{T}}\mu^*.$$

Given that $\mu^* = (\mu_1^*, \mu_2^*, ..., \mu_n^*)$ where μ_i^* is given by (31), we have

$$h_{\infty}^{\mathrm{T}}\mu^* = M \cdot \max_1^+(h_{\infty}). \tag{33}$$

Thus,

$$J_{\pi^*} = v^{\mathrm{T}}\tilde{A}_{\infty}x + \frac{M}{1-\beta}\max_1^+(h_{\infty}).$$

We are ready now to compute $T_{\mu^*}(J_{\pi^*})$ and $T(J_{\pi^*})$. In particular,

$$\begin{aligned} T_{\mu^*}(J_{\pi^*}) &= g(x,\mu^*) + \beta J_{\pi^*}(f(x,\mu^*)) \\ &= v^{\mathrm{T}}\left(I + \beta\tilde{A}_{\infty}A\right)x + h_{\infty}^{\mathrm{T}}\mu^* + \frac{\beta M}{1-\beta}\max_1^+(h_{\infty}). \end{aligned}$$

Due to condition (33) and the fact that $I + \beta\tilde{A}_{\infty}A \equiv \tilde{A}_{\infty}$, we have

$$T_{\mu^*}(J_{\pi^*}) = v^{\mathrm{T}}\tilde{A}_{\infty}x + \frac{M}{1-\beta}\max_1^+(h_{\infty}).$$

Finally,

$$\begin{aligned} T(J_{\pi^*})(x) &= \max_{u\in\mathcal{C}(x)}\{g(x,u) + \beta J_{\pi^*}(f(x,u))\} \\ &= v^{\mathrm{T}}(I + \beta\tilde{A}_{\infty}A)x + \max_{u\in\mathcal{C}(x)}\left\{h_{\infty}^{\mathrm{T}}u\right\} + \frac{\beta M}{1-\beta}\max_1^+(h_{\infty}) \\ &= v^{\mathrm{T}}\tilde{A}_{\infty}x + M\max_1^+(h_{\infty}) + \frac{\beta M}{1-\beta}\max_1^+(h_{\infty}) \\ &= v^{\mathrm{T}}\tilde{A}_{\infty}x + \frac{M}{1-\beta}\max_1^+(h_{\infty}). \end{aligned}$$

Hence, we showed that

$$T_{\mu^*}(J_{\pi^*}) = T(J_{\pi^*}),$$

which implies that π^* is an optimal stationary policy. Also, J_{π^*} provides the optimal value of the infinite-horizon optimization. $\square$

Trying to interpret the optimal stationary policy (31), the firm is going to invest the largest possible amount M to the node i which corresponds to the maximum entry of

$$h_\infty^{\mathrm{T}} = \beta v^{\mathrm{T}} \tilde{A}_\infty B - c^{\mathrm{T}} = \beta v^{\mathrm{T}} (I - \beta A)^{-1} (I - \Theta) - c^{\mathrm{T}}.$$

Note that this decision is affected by the following factors:

1. how easily node i can be influenced by the firm's advertising policy, which is measured by $1 - \theta_i$,

2. how large is the "*network value*" of node i throughout the optimization horizon, expressed by the ith entry of $\beta v^{\mathrm{T}} (I - \beta A)^{-1}$, which measures the effect of every unit of advertising effort spent in i on the proclivities of all nodes that are connected directly or indirectly to i,

3. how small is the cost of every unit of advertising effort in node i, expressed by c_i.

Note also that the matrix $(I - \beta A)^{-1}$, which influences the optimal decision, can be interpreted as a measure of the centrality of the nodes. In fact, Bonacich in his work on measures of centrality [17], introduced the following centrality measure:

$$c(\gamma, \beta) \triangleq \gamma (I - \beta A)^{-1} A \mathbf{1}, \tag{34}$$

where γ is a scaling factor. When $\gamma = 1$, $c(1, \beta)$ has several nice interpretations. To see this, note that the centrality measure is equivalently written as:

$$c(1, \beta) = \left(\sum_{k=0}^{\infty} \beta^k A^k \right) A \mathbf{1} = (I + \beta A + \beta^2 A^2 + ...) A \mathbf{1}. \tag{35}$$

Therefore, the centrality $c(1, \beta)$ is a measure of closeness, and it is high for a node which is connected to other nodes with short and highly weighted paths. The parameter β represents the degree of information (benefits in our model) that is transmitted from one node to another node. In our case, where A is a row stochastic matrix, the above centrality measure takes on the following form

$$c(1, \beta) = (I + \beta A + \beta^2 A^2 + ...) \mathbf{1} = (I - \beta A)^{-1} \mathbf{1}.$$

Trying to translate this centrality measure in the language of our dynamic model of the evolution of preferences, we can say that it represents a measure of the relative importance of nodes (in terms of benefits) when the initial condition is $x_0 = \mathbf{1}$ and there is no external influence (i.e., there is no control input).

Note that in our dynamic model both the initial condition and the control input affect the returns of the advertising firm. Since, though, we are only interested in the computation of the optimal advertising policy, an appropriate centrality (or network value) measure would be $\beta v^{\mathrm{T}} \tilde{A}_\infty B - c^{\mathrm{T}}$. The highest entry of this vector will provide the highest benefits over time. Note that when $\beta = 0$, the control input does not have any implication to the returns. In that case, centrality could be measured by $v^{\mathrm{T}} \tilde{A}_\infty$, since it is only the initial condition that affects the returns.

4.2 Perturbed Dynamics

In this section, we are going to consider a family of perturbations of the nominal model (21), described by

$$x_{k+1} = Ax_k + B\varphi(u_k) + Fq_k, \tag{36}$$

where we have neglected the effect of the second firm. The term q_k corresponds to an unknown signal caused possibly by misspecified system dynamics. The sequence $\{q_k\}$ may feed back in a possibly nonlinear way on the history of x. We will impose the following constraint on the size of any instance of this perturbation sequence:

$$|q_k| \leq \eta, \quad \text{for all } k = 0, 1, ..., \tag{37}$$

where $\eta > 0$ is a measure of the firm's confidence of the accuracy of the nominal model. Let $\mathcal{Q}$ denote the resulting constraint set of disturbances.

Note that due to the presence of the unknown (but bounded) signal q_k our initial assumption that $\mathcal{S} \subset \mathbb{R}^n_+$ may be violated. As we noted though in Section 2.3, the system is input-output stable, therefore an appropriate shift of the state can always guarantee that the dynamics will evolve within the positive cone. In particular, consider $\bar{x} \in \mathbb{R}^n_+$, such that

$$Fq_k + \bar{x} \geq 0, \tag{38}$$

for all q_k satisfying (37), and define instead the dynamics:

$$x_{k+1} = Ax_k + B\varphi(u_k) + Fq_k + \bar{x} \triangleq f(x_k, u_k, q_k). \tag{39}$$

Note that shifting the dynamics by $\bar{x}$ does not change qualitatively the model, since the state x still describes propensities, but relative to $\bar{x}$.

For some $F \in \mathbb{R}^{n \times n}$ let us also define the vector

$$r^{\mathrm{T}}_{k+1} \triangleq \beta v^{\mathrm{T}} \tilde{A}_k F, \quad k = 0, 1, ...,$$

with $r_1^{\mathrm{T}} = \beta v^{\mathrm{T}} F$. Let also:

$$r_\infty^{\mathrm{T}} \triangleq \beta v^{\mathrm{T}} \tilde{A}_\infty F.$$

We would like to solve the following optimization for the computation of a robust solution:

$$\max_{\pi \in \Pi} \min_{\sigma \in \Sigma} \left\{ J_{(\pi,\sigma)}(x_0) \triangleq \lim_{N \to \infty} \sum_{k=0}^{N-1} \beta^k g(x_k, \mu_k(x_k)) \right\}, \tag{40}$$

subject to the perturbed dynamics (39) and the constraints (37)–(38). Here Σ denotes the set of sequences of policies $\sigma = (\nu_0, \nu_1, ...)$ of the uncertainty, where ν_k is a function from the set of states $\mathcal{S}$ to $\mathcal{Q}$. Note also that due to the new shifted dynamics, a utility function of the form

$$g(x, u) = v^{\mathrm{T}} x - c^{\mathrm{T}} u - \lambda(\bar{x})$$

would have been more appropriate. However, in that case, and since the last term is a constant, the optimal policy of (40) would have been identical.

Proposition 4.3 (Optimal policy under uncertainty) *Consider the infinite horizon optimization of (40) under the perturbed dynamics (39) and the constraint (37)–(38). The optimal stationary policy is $\mu^* = (\mu_1^*, ..., \mu_n^*)$, such that*

$$\mu_i^* = \begin{cases} M & i = \arg\max_1^+ (h_\infty) \\ 0 & \textit{otherwise} \end{cases}, \quad i \in \mathcal{I}. \tag{41}$$

Proof. To solve this optimization problem, we implement the dynamic programming iteration. In fact, we recursively implement the operator $T(\cdot)$ defined as

$$(TJ)(x) \triangleq \max_{u \in \mathcal{C}} \min_{q \in \mathcal{Q}} \{ g(x, u) + \beta J(f(x, u, q)) \}, \tag{42}$$

for any $x \in \mathcal{S}$. The dynamic programming iteration successively gives:

$$J_N(x) \quad = \quad v^{\mathrm{T}} \tilde{A}_N x + \sum_{k=0}^{N-1} \beta^k h_{N-k}^{\mathrm{T}} u_{N-k}^* + \sum_{k=0}^{N-1} \beta^k r_{N-k}^{\mathrm{T}} q_{N-k}^* + \sum_{k=0}^{N-1} \beta^{k+1} v^{\mathrm{T}} \tilde{A}_{N-k} \bar{x},$$

for all $N = 1, 2, ...$, where u_k^* and q_k^* denote the sequences of optimal investments and disturbances, respectively. In particular, $u_k^* = (u_{1,k}^*, ..., u_{n,k}^*)$ and $q_k^* = (q_{1,k}^*, ..., q_{n,k}^*)$, are such that

$$u_{i,k}^* = \begin{cases} M & i = \arg\max_1^+ (h_k) \\ 0 & \text{otherwise} \end{cases}, \quad i \in \mathcal{I},$$

and

$$r_k^{\mathrm{T}} q_k^* = -\eta \left| r_k \right|_\infty .$$

In other words, the disturbance places all its weight on the maximum (in absolute value) entry of r_k, or

$$q_{i,k}^* = \begin{cases} -\eta & i = \arg\max_1^+ (r_k) \\ 0 & \text{otherwise} \end{cases}, \quad i \in \mathcal{I}.$$

The order of max and min in the definition of the operator $T(\cdot)$ does not change the optimal policies. Note also that:

$$\begin{aligned} H(x,u,q,J_0) &= g(x,u) + \beta J_0(f(x,u,q)) \\ &= J_0(x) + \beta v^{\mathrm{T}} A x + \beta v^{\mathrm{T}}(Fq + \bar{x}) + (\beta v^{\mathrm{T}} B - c^{\mathrm{T}}) u \\ &\geq J_0(x) \end{aligned}$$

for all $x \in \mathcal{S}$, $u \in \mathcal{C}^*$, $q \in \mathcal{Q}^*$ and under condition (38). Thus, from Proposition 3.2, the dynamic programming iteration provides the optimal infinite value.

Consider the stationary policy (41) for the monopolistic firm and the stationary policy $\sigma^* = (\nu^*, ..., \nu^*)$ for the disturbance such that

$$r_\infty^{\mathrm{T}} \nu^* = -\eta \left| r_\infty \right|_\infty .$$

Similarly to the proof of Proposition 4.2, the corresponding infinite value is

$$\begin{aligned} J_{(\pi^*,\sigma^*)}(x) &= v^{\mathrm{T}} \tilde{A}_\infty x + h_\infty^{\mathrm{T}} \lim_{N\to\infty} \sum_{k=0}^{N-1} \beta^k \mu^* + r_\infty^{\mathrm{T}} \lim_{N\to\infty} \sum_{k=0}^{N-1} \beta^k \nu^* + \beta v^{\mathrm{T}} \tilde{A}_\infty \sum_{k=0}^{N-1} \beta^k \bar{x} \\ &= v^{\mathrm{T}} \tilde{A}_\infty x + \frac{M}{1-\beta} \max_1^+ (h_\infty) - \frac{\eta}{1-\beta} \left| r_\infty \right|_\infty + \frac{\beta}{1-\beta} v^{\mathrm{T}} \tilde{A}_\infty \bar{x}. \end{aligned}$$

By following similar reasoning to the proof of Proposition 4.2, we can show that

$$T_{(\mu^*,\nu^*)}(J_{(\pi^*,\sigma^*)}) = T(J_{(\pi^*,\sigma^*)}).$$

Therefore, according to Proposition 3.2, (π^*, σ^*) provides the optimal lower value. It is also straightforward to show that the sequence of policies (π^*, σ^*) also provides the optimal upper value, defining this way a solution to the max-min optimization problem. □

Note that the robust optimal policy for the perturbed model coincides with the optimal policy for the unperturbed or riskless model, i.e., exhibits a certainty equivalence property. Such property was expected due to the linearity of the perturbed model (36) and the linearity of the utility function

(6).

5 Optimal Policy in Duopoly

5.1 Preliminaries

The previous section computed the optimal robust policy for the problem of monopoly under norm-bounded model uncertainty. In this section, we would also like to include the possibility that a competitive firm tries to influence the preferences of the customers towards buying its own product as described by the more general duopoly model (4).

The presence of a competitive firm introduces a new source of uncertainty. We will either assume that i) the competitive firm has the form of a competitive fringe which tries to enter the market, introducing a notion of sequential optimization (expressed by a Stackelberg solution), or ii) both firms have the ability of simultaneous play (expressed by a Nash solution).

Each firm $\ell \in \mathcal{L}$ solves the following optimization problem:

$$\max_{\pi^\ell \in \Pi^\ell} \left\{ J_{(\pi^\ell, \pi^{-\ell})}(x_0^\ell) \triangleq \lim_{N \to \infty} \sum_{k=0}^{N-1} \beta^k g\left(x_k^\ell, \mu_k^\ell(x_k^\ell)\right) \right\} \tag{43}$$

subject to the system dynamics

$$x_{k+1}^\ell = A x_k^\ell + B \varphi(\mu_k^\ell, \mu_k^{-\ell}) \tag{44}$$

where $\pi^\ell = (\mu_1^\ell, \mu_2^\ell, ...)$ and $\pi^{-\ell} = (\mu_1^{-\ell}, \mu_2^{-\ell}, ...)$ are the infinite sequences of policies of the firms ℓ and $-\ell$, respectively.

Definition 5.1 (Stackelberg solution) *A Stackelberg solution is a pair of policies* $(\pi^{\ell *}, \pi^{-\ell *}) \in \Pi^\ell \times \Pi^{-\ell}$ *such that*

$$\pi^{-\ell *} \in \mathrm{BR}_{-\ell}(\pi^{\ell *}) \triangleq \arg\max_{\pi^{-\ell}} \left\{ J_{(\pi^{-\ell}, \pi^\ell)}(x_0^{-\ell}) \,\middle|\, \pi^{\ell *} \right\}$$

and, furthermore,

$$\pi^{\ell *} \in \arg\max_{\pi^\ell \in \Pi^\ell} \left\{ J_{(\pi^\ell, \pi^{-\ell})}(x_0^\ell) \,\middle|\, \pi^{-\ell} \in \mathrm{BR}_{-\ell}(\pi^\ell) \right\}.$$

In the above definition of a Stackelberg solution, we will refer to firm ℓ as the *leader* and firm $-\ell$ as the *follower*. Note that the definition implies that firm ℓ (or *leader*) announces first its policy, while firm $-\ell$ (or *follower*) reacts to that policy.

Definition 5.2 (Nash solution) *A Nash solution is a pair of policies* $(\pi^{\ell *}, \pi^{-\ell *}) \in \Pi^\ell \times \Pi^{-\ell}$ *such that*

$$\pi^{-\ell *} \in \mathrm{BR}_{-\ell}(\pi^{\ell *}) \triangleq \arg\max_{\pi^{-\ell} \in \Pi^{-\ell}} \left\{ J_{\pi^{-\ell}}(x_0^{-\ell}) \,\middle|\, \pi^{\ell *} \right\}$$

and, furthermore,

$$\pi^{\ell*} \in \mathrm{BR}_\ell(\pi^{-\ell*}) \triangleq \arg\max_{\pi^\ell \in \Pi^\ell} \left\{ J_{\pi^\ell}(x_0^\ell) \,\middle|\, \pi^{-\ell*} \right\}.$$

We will also refer to these solutions as *Markovian* or *closed-loop Nash solutions*. If, instead, the maximization in the definition of the Nash solution is restricted to the set of sequences of control inputs in $\mathcal{C}^\ell$, then the corresponding solutions will be referred to as *open-loop Nash solutions.* Note that these definitions of Nash solutions implicitly assumes a simultaneous announcement of policies from both firms.

A straightforward implication of the above definitions is the following claim.

Claim 5.1 *Any Stackelberg solution is also a Nash solution.*

5.2 Open-loop stationary Nash equilibria

In this section, we will restrict our attention to open-loop Nash equilibria that are also stationary, i.e., time-independent. Before characterizing this family of Nash solutions, define the set of actions $\mathcal{A}^\ell \triangleq \{\alpha_1, \alpha_2, ..., \alpha_n\}$, $\ell \in \mathcal{L}$, such that for each $i \in \{1, 2, ..., n\}$, $\alpha_i = (\alpha_{i,1}, \alpha_{i,2}, ..., \alpha_{i,n})$ where

$$\alpha_{i,j} \triangleq \begin{cases} M & j = \arg\max_i^+(h_\infty), \\ 0 & \text{otherwise,} \end{cases} \qquad j = 1, 2, ..., n.$$

In other words, the action α_i corresponds to investing all available funds to the ith largest non-negative entry of h_∞. Note that the set of actions define an isomorphic set of stationary policies, i.e., for each action α_i there is a stationary policy $(\alpha_i, \alpha_i, ...)$. Let us also denote by $J_{(i,j)}(x)$ the corresponding infinite horizon value for initial condition x when one firm applies stationary policy $(\alpha_i, \alpha_i, ...)$ and the other firm applies stationary policy $(\alpha_j, \alpha_j, ...)$. Any other open-loop stationary policy μ^ℓ can be represented as a mixture of actions in $\mathcal{A}^\ell$, i.e.,

$$\mu^\ell = \begin{cases} \alpha_1, & \text{with probability } p_1^\ell \\ \alpha_2, & \text{with probability } p_2^\ell \\ ... & \\ \alpha_n, & \text{with probability } p_n^\ell \end{cases}, \quad \ell \in \mathcal{L}, \tag{45}$$

where $p_i^\ell \geq 0$, $i \in \mathcal{I}$, and $\sum_i p_i^\ell = 1$. The corresponding value of the objective function (43) for any open-loop stationary policy is characterized by the following proposition.

Proposition 5.1 (Payoffs under stationary open-loop policies) *When both firms $\ell \in \mathcal{L}$ apply an open-loop stationary policy $\pi^\ell = (\mu^\ell, \mu^\ell, ...)$ satisfying (45), the infinite value of the objective* *function $J_{(\pi^\ell,\pi^{-\ell})}$ defined by (43), is*

$$J_{(\pi^\ell,\pi^{-\ell})} = \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{I}} J_{(i,j)} p_i^\ell p_j^{-\ell},$$

where

$$J_{(i,j)}(x) = \begin{cases} v^{\mathrm{T}}\tilde{A}_\infty x + \frac{1}{1-\beta}[-c^{\mathrm{T}}\alpha_i], & i=j, \\ v^{\mathrm{T}}\tilde{A}_\infty x + \frac{1}{1-\beta}[h_\infty^{\mathrm{T}}\alpha_i], & i\neq j, \end{cases} \quad x\in\mathcal{S}^\ell, \quad \ell\in\mathcal{L}. \tag{46}$$

Proof. When the pair of stationary policies $(\pi^\ell,\pi^{-\ell})$ is applied, where $\pi^\ell = (\mu^\ell,\mu^\ell,...)$ and $\pi^{-\ell} = (\mu^{-\ell},\mu^{-\ell},...)$, the corresponding value of the objective function of firm ℓ will be:

$$\begin{aligned} J_{(\pi^\ell,\pi^{-\ell})}(x) &= v^{\mathrm{T}}\tilde{A}_\infty x + \lim_{N\to\infty}\sum_{k=0}^{N-1}\beta^k\left[(h_\infty+c)^{\mathrm{T}}\varphi\left(\mu^\ell(x),\mu^{-\ell}(x)\right) - c^{\mathrm{T}}\mu^\ell(x)\right] \\ &= v^{\mathrm{T}}\tilde{A}_\infty x + \frac{1}{1-\beta}\left[(h_\infty+c)^{\mathrm{T}}\varphi\left(\mu^\ell(x),\mu^{-\ell}(x)\right) - c^{\mathrm{T}}\mu^\ell(x)\right] \end{aligned}$$

for some initial state $x\in\mathcal{S}^\ell$. If $\mu^\ell = \mu^{-\ell} = \alpha_i$, then the corresponding infinite value of the objective function of ℓ, denoted $J_{(i,i)}$, is:

$$J_{(i,i)}(x) = v^{\mathrm{T}}\tilde{A}_\infty x + \frac{1}{1-\beta}[-c^{\mathrm{T}}\alpha_i].$$

If, instead, $\mu^\ell = \alpha_i$ and $\mu^{-\ell} = \alpha_j$ with $i\neq j$, the corresponding infinite value of the objective function ℓ, denoted $J_{(i,j)}$, is:

$$J_{(i,j)}(x) = v^{\mathrm{T}}\tilde{A}_\infty x + \frac{1}{1-\beta}[h_\infty^{\mathrm{T}}\alpha_i].$$

Then, the corresponding expected return of firm $\ell\in\mathcal{L}$ is:

$$\begin{aligned} J_{(\pi^\ell,\pi^{-\ell})}(x) &= v^{\mathrm{T}}\tilde{A}_\infty x + \frac{1}{1-\beta}\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{I}}\left[(h_\infty+c)^{\mathrm{T}}\varphi(\alpha_i,\alpha_j) - c^{\mathrm{T}}\alpha_i\right] p_i^\ell p_j^{-\ell} \\ &= \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{I}}\left[v^{\mathrm{T}}\tilde{A}_\infty x + \frac{1}{1-\beta}\left[(h_\infty+c)^{\mathrm{T}}\varphi(\alpha_i,\alpha_j) - c^{\mathrm{T}}\alpha_i\right]\right] p_i^\ell p_j^{-\ell} \\ &= \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{I}} J_{(i,j)} p_i^\ell p_j^{-\ell}, \end{aligned}$$

which concludes the proof. □

Thus, we may define an equivalent one-stage symmetric game of two players, finite set of actions $\mathcal{A}^\ell = \{\alpha_1,\alpha_2,...,\alpha_n\}$ for each player $\ell\in\mathcal{L}$, and payoff matrix of the row player which is given by Table 1.

	α_1	α_2		α_n
α_1	$J_{(1,1)}$	$J_{(1,2)}$	...	$J_{(1,n)}$
α_2	$J_{(2,1)}$	$J_{(2,2)}$	...	$J_{(2,n)}$
	...	...	...	...
α_n	$J_{(n,1)}$	$J_{(n,2)}$	...	$J_{(n,n)}$

Table 1: Equivalent one-shot symmetric game in open-loop stationary policies.

A direct consequence of Claim 5.1 is the following:

Claim 5.2 *The following hold:*

1. $J_{(i,j)}(x) \geq J_{(i,i)}(x)$ *for all* $i, j \in \mathcal{I}$ *with* $i \neq j$;
2. $J_{(i,j)}(x) = J_{(i,j')}(x)$ *for all* $i, j, j' \in \mathcal{I}$ *such that* $j \neq i$ *and* $j' \neq i$;
3. $J_{(i,j)}(x) \geq J_{(j,i)}(x)$ *for all* $i, j \in \mathcal{I}$ *with* $i > j$.

Proposition 5.2 (Stackelberg & Nash solutions) *Consider the optimization problem (43) under the dynamics (44) and the constraints (1) with* $M^{\ell} = M^{-\ell}$, *i.e., both firms have identical advertising power. For any* $\ell \in \mathcal{L}$, *the pair of open-loop stationary policies* $\pi^* = (\pi^{\ell *}, \pi^{-\ell *})$ *where* $\pi^{\ell *} = (\mu^{\ell *}, \mu^{\ell *}, ...)$ *and* μ^{ℓ} *is defined by (45) satisfying either*

1. $p_1^{\ell} = p_2^{-\ell} = 1$, *or*
2. $p_1^{\ell} = p_2^{-\ell} = \frac{J_{(1,2)} - J_{(2,2)}}{J_{(1,2)} - J_{(1,1)} + J_{(2,1)} - J_{(2,2)}}$,

defines an open-loop Nash solution. Furthermore, when $\ell \in \mathcal{L}$ *has the opportunity to announce first its policy, any one of the above pairs of open-loop stationary policies also defines an open-loop Stackelberg solution.*

Proof. The first claim is a direct consequence of Claim 5.2 and the fact that any one of the policies corresponding to the cases (1) and (2) defines a Nash equilibrium for the equivalent one-shot symmetric game of Table 1.

Assume now that ℓ has the opportunity to announce its strategy first. In order to show that $(\pi^{\ell *}, \pi^{-\ell *})$ defines a Stackelberg solution, we need to verify that the leader's policy $\pi^{\ell *}$ guarantees maximum return over all possible announced policies. It is straightforward to show that any announced policy that does not allocate all available funds to $\arg\max_1^+(h_\infty)$ will result to a best response of the follower that can only decrease leader's optimal value. $\square$

The conclusions of Proposition 5.2 do not necessarily hold when we consider different spending powers for the firms, i.e., when $M^{\ell} \neq M^{-\ell}$. However, extending the conclusions of Proposition 5.2, to that case is straightforward.

Another straightforward implication of Proposition 5.2 is summarized in the following corollary.

Corollary 5.1 *The open-loop stationary Nash solutions characterized by Proposition 5.2 are also closed-loop Nash solutions.*

This is due to the fact that open-loop strategies are a subset of Markovian or state-dependent strategies.

A complete characterization of the set of closed-loop Nash solutions is going beyond the scope of this paper, since it is highly case-dependent, i.e., it depends on the class of policies which will be considered reasonable for the application of interest. For example, if we assume that the class of strategies over which the optimization is executed are affine functions of the state, then a new class of closed loop Nash solutions can easily be computed using the framework proposed in this paper.

5.3 Max-min solutions

Computation of Nash equilibria by either firm requires that firms are aware of the performance indices of the competitors. Such an assumption may be quite strong especially in a competitive environment.

When firms are not able to compute Nash solutions, computing an optimal strategy which is robust to any possible policy of the competitor might be the only possibility. Such an optimization problem can be formulated as a max-min optimization.

We are going to consider two firms with different expenditure capabilities. In particular, we consider the following two scenarios: a) $M^\ell > M^{-\ell}$, and b) $M^\ell \leq M^{-\ell}$ for any $\ell \in \mathcal{L}$.

The firm $\ell \in \{a, b\}$ solves the following max-min optimization problem:

$$\max_{\pi \in \Pi} \min_{\sigma \in \Sigma} \left\{ J_{(\pi,\sigma)}(x_0) \triangleq \lim_{N \to \infty} \sum_{k=0}^{N-1} \beta^k g\left(x_k, \mu_k(x_k)\right) \right\} \tag{47}$$

over the set Π of infinite sequences of policies $(\mu_0, \mu_1, ...)$ and subject to the system dynamics

$$x_{k+1} = Ax_k + B\varphi(\mu_k, \nu_k). \tag{48}$$

The set Σ denotes the collection of infinite sequences of policies $(\nu_0, \nu_1, ...)$ of the competitor. In words, the above optimization reflects the situation at which the firm wishes to announce a strategy which will provide the optimal returns assuming that the competitor acts to minimize these returns. To simplify notation, we have removed the superscript ℓ from the above optimization variables. It is straightforward to show that the following holds:

Proposition 5.3 *Consider the optimization problem (47) under the dynamics (48) and the constraints (1). If $M^\ell > M^{-\ell}$, i.e., the advertising power of the firm is larger than the competitor's* *one, then the optimal strategy of the firm will be a stationary policy* $(\mu^*, \mu^*, ...)$ *such that*

$$\mu_i^* = \begin{cases} M & i = \arg\max^+ (h_\infty) \\ 0 & \textit{otherwise} \end{cases}, \quad i \in \mathcal{I}. \tag{49}$$

Note that this is not necessarily the case when the advertising power of the firm is less than the corresponding one of the competitor's. It is straightforward to see that in that case, any strategy will be optimal, since the competitor has the power to counteract any announced strategy of the firm.

6 Conclusions

We discussed the problem of deriving optimal advertising strategies in a network of customers. Contrary to prior work, the dynamics of preferences were also affected by an underlying network of connections which introduces a form of word-of-mouth effects between nodes. The derived optimal policies are related to and extend prior introduced notions of centrality measures usually considered in sociology. Although the assumed model of the evolution of preferences might be the outcome of an identification process, it is likely that we are uncertain about its accuracy. To this end, we also considered a perturbed model which models possible misspecifications or uncertainties of the nominal model, and we derived robust optimal strategies. It was shown that the monopoly model exhibits a certainty equivalence property, i.e., the optimal strategies for the perturbed model coincide with the optimal strategies for the unperturbed or riskless model. Such behavior can be attributed to both the linearity of the monopoly dynamics and the linearity of the utilities. Finally, we investigated robust policies in a duopoly framework. In particular, we characterized the set of open-loop Nash solutions, which also happens to be closed-loop Nash solutions. The model can easily be utilized to accommodate scenarios at which more complicated forms of strategies are of interest, leading to new forms of closed-loop Nash solutions. We finally characterized the set of max-min solutions in a duopoly framework, when the firm makes no assumptions about the utilities of the competitive firm.

References

[1] M. L. Vidale and H. B. Wolfe, "An operations research study of sales response to advertising," *Operations Res.*, vol. 5, pp. 370–381, 1957.

[2] S. Jørgensen and G. Zaccour, *Differential Games in Marketing.* Kluwer Academic Publishers, 2004.

[3] S. P. Sethi, "Dynamic optimal control models in advertising: A survey," *SIAM Review*, vol. 19, no. 4, pp. 685–725, Oct. 1977.

[4] G. Leitmann and W. E. Schmitendorf, "Profit maximization through advertising: A nonzero sum differential game approach," *IEEE Transactions on Automatic Control*, vol. 23, no. 4, pp. 646–650, 1978.

[5] J. Case, *Economics and the competitive process.* New York, NY: New York University Press, 1979.

[6] K. Deal, "Optimizing advertising expenditures in a dynamic duopoly," *Operations Research*, vol. 27, pp. 682–692, 1979.

[7] K. Deal and S. Zionts, "A differential games solution to the problem of determining the optimal timing of advertising expenditures," in *Proc. Second Annual Northeast Regional AIDS Conference*, Atlanta, GA, 1973.

[8] G. M. Erickson, "A model of advertising competition," *Journal of Marketing Research*, vol. 22, pp. 297–304, 1985.

[9] ——, "Empirical analysis of closed-loop duopoly advertising strategies," *Management Science*, vol. 38, no. 12, pp. 1732–1749, 1992.

[10] P. K. Chintagunta and N. J. Vilcassim, "An empirical investigation of advertising strategies in a dynamic duopoly," *Management Science*, vol. 38, no. 9, pp. 1230–1244, 1992.

[11] G. E. Fruchter and S. Kalish, "Closed-loop advertising strategies in duopoly," *Management Science*, vol. 43, no. 1, pp. 54–63, 1997.

[12] S. Jørgensen, "A survey of some differential games in advertising," *Journal of Economic Dynamics and Control*, vol. 4, pp. 341–369, 1982.

[13] F. M. Bass, "New product growth model for consumer durables," *Management Science*, vol. 15, pp. 215–227, 1969.

[14] M. Nerlove and K. J. Arrow, "Optimal advertising policy under dynamic conditions," *Economica*, vol. 39, pp. 129–142, 1962.

[15] F. Alkemade and C. Castaldi, "Strategies for the diffusion of innovations on social networks," *Computational Economics*, vol. 25, pp. 3–23, 2005.

[16] P. K. Dubey, B. D. Meyer, and R. Garg, "Competing for customers in a social network," Cowles Foundation, Tech. Rep. 1591, Nov. 2006.

[17] P. Bonacich, "Power of centrality: A family of measures," *American Journal of Sociology*, vol. 92, no. 5, pp. 1170–1182, 1987.

[18] S. Sethi, "Deterministic and stochastic optimization of a dynamic advertising model," *Optimal Control Applications & Methods*, vol. 4, no. 2, pp. 179–184, 1983.

[19] A. Prasad and S. P. Sethi, "Competitive advertising under uncertainty: A stochastic differential game approach," *Journal of Optimization Theory and Applications*, vol. 123, no. 1, pp. 163–185, Oct. 2004.

[20] N. E. Friedkin, "Norm formation in social influence networks," *Social Networks*, vol. 23, pp. 167–189, 2001.

[21] T. Kailath, *Linear Systems.* Prentice-Hall, 1980.

[22] N. E. Friedkin and E. Johnsen, "Social influence networks and opinion change," *Advances in Group Processes*, vol. 16, pp. 1–29, 1999.

[23] B. Golub and M. O. Jackson, "Naive learning in social networks: Convergence, influence, and the wisdom of crowds, Tech. Rep. FEEM Working Paper No. 64, 2007.

[24] N. E. Friedkin, "Theoretical foundations for centrality measures," *American Journal of Sociology*, vol. 96, pp. 1478–1504, 1991.

[25] C. Fershtman, "Goodwill and market shares in oligopoly," *Economica*, vol. 51, pp. 271–281, 1984.

[26] P. K. Chintagunta, "Investigating the sensitivity of equilibrium profits to advertising dynamics and competitive effects," *Management Science*, vol. 39, pp. 1146–1162, 1993.

[27] J. P. Gould, "Diffusion processes and optimal advertising policy," in *Microeconomic foundations of employment and inflation theory*, E. P. et al., Ed. W.W. Norton, New York, 1970, pp. 338–368.

[28] K. Deal, S. Sethi, and G. Thompson, "A bilinear-quadratic differential game in advertising," in *Control theory in mathematical economics, Proceedings of the third Kingston Conference, Part B*, ser. Lecture notes in pure and applied mathematics, P. Liu and J. Sutinen, Eds. Marcel Dekker, New York, 1979, vol. 1200.

[29] G. C. Kimball, "Some industrial applications of military operations research methods," *Operations Research*, vol. 5, pp. 201–204, 1957.

[30] D. Bertsekas and S. Shreve, *Stochastic Optimal Control: The Discrete Time Case.* London: Academic Press, Inc. (London) Ltd., 1978.

[31] F. Jones, *Lebesgue integration on Euclidean space.* Jones and Bartlett Publishers, Inc., 1993.

[32] R. A. Horn and C. R. Johnson, *Matrix Analysis.* Cambridge, UK: Cambridge University Press, 1985.