A Marketing Game

a rigorous model for strategic resource allocation

Matthew G. Reyes
Independent Researcher and Consultant
Ann Arbor, Michigan
amarketinggame.com
matthewgreyes@yahoo.com

ABSTRACT

In this paper we introduce a model of marketing-influenced socially-contingent decision-making with respect to a market of alternative Products. The model is based on the foundation of socially-contingent random utility theory [17], [3], and as such, is *data-driven* in the sense of viewing consumers' perception of Product utility as a parametrization of the observed frequencies with which consumers choose among the alternative Products. This paper introduces a *marketing strength* into this parametrization, and a *marketing response* indicating the marketing strength applied to a consumer as a function of investment by a Company. We analytically and numerically illustrate the utility of this model by showing in a market of two alternatives with no inherent biases and uniform social biases on a cycle network, for any dollar allocation between two consumers, marketing share is optimized by targeting consumers equidistant on the network; and that the optimal dollar allocation between two consumers depends on the size of the budget.

KEYWORDS

random utility, interaction games, product adoption, Glauber dynamics, Gibbs fields, marketing, utility, graphical models

ACM Reference Format:

Matthew G. Reyes. 2018. A Marketing Game: a rigorous model for strategic resource allocation. In *Proceedings of ACM Machine Learning in Graphs conference (MLG'18)*. ACM, New York, NY, USA, 8 pages. https://doi.org/10.475/123_4

1 INTRODUCTION

Consider a *market* consisting of two alternatives, Product A and Product B, between which consumers choose according to their perception of the value of these two alternatives. The Products may be commercial products or political candidates, for example. Company A and Company B *market* their respective Products in an effort to enhance consumers' perception of the utility of these Products. This paper is concerned with the problem of how Companies A and B should allocate marketing resources to optimize their respective *market share* [2], the fraction of consumers that choose their Product.

MLG'18, August 2018, London, UK

ACM ISBN 123-4567-24-567/08/06.
https://doi.org/10.475/123_4

We model consumer choice within the framework of random utility [4], [17], in particular its extension to socially-contingent decision-making [3]. A tenet of individual choice theory [26], [14], [17], is that consumers seek to maximize perceived utility among alternatives A and B. The theory of *random* utility [4], [17], states that such utility maximization, *with respect to the market* consisting of alternatives A and B, will result in choices that *appear* random due to the fact that decisions made with respect to this market will nevertheless be influenced by considerations external to this market. Randomness in individual choice can also be derived by hypothesizing *bounded rationality* [25] on the part of consumers. A strength of the random utility model is that utilities can be viewed as a parametrization of the probabilities with which consumers make decisions within a particular market, and thus can be estimated as those that predict *observed* frequencies of choice by consumers.

We assume consumers belong to a *social network*, in which pairs of consumers whose choices are contingent upon one another's are referred to as *neighbors*. Such socially-contingent decision-making is referred to as a *game* [20]. The simplest such socially-contingent parametrization includes *inherent bias* of individual consumers towards one or the other Product, and *social bias* reflecting the degrees to which neighboring consumers influence each others' decisions. In this paper we add to this parameterization action taken by a *marketer* for Companies A and B to enhance consumers' perception of the utility of Products A and B, respectively. In particular, we introduce a *marketing response* for individual consumers that determines the effectiveness of marketing in biasing their choices.

There has been a great deal of interest in modeling how products are adopted on social networks. Some of this interest stems from Malcolm Gladwell's *The Tipping Point* [11]. In this work he compared the fast, widespread adoption of a product to the contagious nature of epidemics, and he can be credited, at least in part, with our current jargon in referring to a product *going viral*. Much of the early technical work on product adoption within a network pursued the epidemic approach [29], modeling and analyzing the spread of product preference within a social network *as a virus*. The epidemic analogy is rather misleading, however, with respect to the influences dictating *choices* made on a network. For example, to use the vernacular of epidemics in the context of a market with Products A and B, we might say that consumer i has been infected with virus A, a neighbor j possibly with virus B. Then, if i infects j with virus A, j would no longer be infected with virus B, perhaps initiating an *epidemic* in which Product A is widely adopted. But this is *not* how viruses spread, and therefore, product adoption should *not* be modeled as viruses.

The framework of socially-contingent decision-making through the concept of games has been pursued by many researchers [3], [31], [30], [19], [8]. Most of this work has modeled consumer choice with best-response dynamics, which involve deterministic maximization between utility assigned to the different Products. When stochastic choice dynamics have been considered [3], [19], it has typically been in an effort to facilitate convergence of the network dynamics to a particular Nash equilibrium, for example the payoff- as opposed to the risk-dominant choice [13], [31]. Considered in the framework of socially-contingent random utility, the problem of Product adoption depends on the particular Product market. (Near) best-response dynamics and Nash equilibria make sense for markets in which utility with respect to the particular market, for example *fitting in* with respect to social conventions, outweighs considerations external to the market. However, for Products that are either commercial products or political candidates, external factors *do* influence decision-making. More importantly, a Company need not concern itself with what these external factors are, and how much utility they may have to particular consumers. Rather, the beauty of the random utility framework is in viewing utility as a parametrization of *actual frequencies* of consumer choice, so that a Company can choose a parametrization corresponding to measurable data, then simply estimate the values of the parametrization that predict the choices actually being made by consumers.

Within the framework of socially-contingent decision-making, many have considered the problem of which consumers to target in order to facilitate adoption of a Product, for example [30], [19], [8]. Invariably, these papers address the problem by considering an initial *seeding* of a Product, followed by (best-response) choice dynamics determined solely by inherent and social biases. But Companies continue to market their Products long after the Products are released. Moreover, the reason Companies market their Products is the belief that such efforts will influence the decision-making that takes place within the social network. As such, the marketer needs to be included in a parametrization of (stochastic) consumer choice.

A Company's *marketing allocation* consists of a subset of consumers, together with respective amounts and types of marketing for consumers within this subset. In general, the marketing allocation will vary over time, as a Company collects data and infers from this data where further gains can be made or losses need to be cut. However, what is important to note is that with respect to the problem of resource allocation, one must account for the fact that marketing is applied (more or less) continuously and is not simply a matter of inducing an initial configuration of sorts.

Marketing is designed to enhance *perception* of utility. The extent to which additional marketing biases a consumer's perception of utility will vary according to the amount of marketing to which the consumer has already been exposed. Therefore it is natural to invoke the well-known Weber-Fechner law regarding changes in perception as a function of stimulus intensity [5]. The scale and shift of the resulting sigmoid function will in general depend on the type of marketing, and this dependence will likely be consumer-specific. In Section 3.2 we propose a *marketing response* to capture the relationship between dollar investment in a particular type of marketing to a particular consumer, and the resulting *applied bias* exerted upon the consumer's choice response.

In this paper we consider a very popular type of choice response dynamics, the *logit* model. The logit model is frequently used in modeling the product choices of consumers [12], [16], due to its parametrization of choice response in terms of varying influences. It is well known that the logit model resulting from random utility theory is a special case of market share [2] in expressing the fraction of a market that a company gets due to various influences. In Section 2.2 we briefly discuss assumptions and implications of this model as well as alternative models that may be more appropriate for different markets.

In general, optimization of a Company's marketing allocation will involve simulation of choice response dynamics in order to assess expected market share resulting from particular candidate marketing allocations. However, a nice property of logit choice response dynamics is that, under conditions which we discuss briefly in Section 2.3, they converge to an equilibrium Gibbs distribution [10], which makes it relatively simple to illustrate the broader concepts of this approach. For example, in Section 3.3 we imagine a Company optimizing total bias resulting from a two unit marketing allocation with respect to such an equilibrium Gibbs distribution on a cycle. We show analytically that for any distribution between the two marketing allocations, total bias is increased by spacing the allocations evenly around the cycle. We show numerically that if the marketing budget is large enough, then distributing the marketing budget evenly between the two allocations is optimal, for any allocation placement, while if the budget is small, then it may be better to skew the budget distribution towards one consumer.

This paper is organized as follows. In Section 2 we provide background on random utility and the resulting choice response dynamics and equilibria. In Section 3.1 we discuss performance criterion for our game, in Section 3.2 we describe the marketing response function, and in Section 3.3 we illustrate analysis of our model.

2 BACKGROUND: UTILITY, CHOICE RESPONSE DYNAMICS, AND EQUILIBRIA OF COORDINATED DECISIONS

Consider a set V of *consumers* who interact on a *social network* $G = (V, E)$, where E is the set of *neighbors*, or pairs of individuals who *interact* on the network. We say that two consumers i and j *interact* on G if there is direct communication between the two in the sense that each observes and is potentially influenced by the decisions of the other. We say that the interaction among neighbors is *pairwise* in the sense of being *additive* in the sense that the *absolute influence* exerted on consumer i's decision by a neighbor j is independent of the presence of other neighbors. That being said, the *relative influence* of a neighbor on a given consumer *will* depend on the presence of other neighbors and the absolute influences that those other neighbors respectively exert on the consumer. The set of neighbors of consumer i is denoted ∂i. Let A and B be two *Companies*, each providing a single *Product*. To distinguish between a Company and its Product, we may refer to Company A and Product A, for example, though if no distinction is required in a given context, we will simply refer to A. We are interested in the *market* consisting of the exchange of value between consumers and Companies A and B. The value provided by consumers can be money or a vote, for example, exchanged for the perceived utilities derived from Products A and B. We will make use of the *variable* x_i, where

$$x_i = \begin{cases} 1 & \text{consumer } i \text{ chooses } A \\ -1 & \text{consumer } i \text{ chooses } B \end{cases}, \quad (1)$$

and we will abuse notation somewhat and let x_i refer *both* to the numerical value (1 or -1) *and* the choice (A or B).

At time t, the configuration of choices $\mathbf{x}^{(t)} = (x_1^{(t)}, \ldots, x_{|V|}^{(t)})$ on the network represent preferences for the two Products. Consumer i observes the choices $\{x_j^{(t)} : j \in \partial i\}$ of his neighbors, for example through posts on social media. Each consumer will *update* his choice with a given frequency, which we will assume to be a Poisson distribution, and we will further assume that this frequency is the same for all consumers. A well-known consequence of the Poisson assumption is that the probability that two consumers update their choices simultaneously is essentially zero, which means the dynamic process of choice updates can be modeled as an appropriately-defined *Glauber dynamics* [9]. We will discretize time according to the times at which consumers update their choices. This model of choice updates is standard [3].

2.1 Random Rational Utility

We adopt the model of random utility [4], [17]. Assume consumers are *rational*, that is, that they seek to *maximize* utility

$$U = \begin{bmatrix} u_A + \epsilon_A \\ u_B + \epsilon_B \end{bmatrix}, \quad (2)$$

where u_A is a *known utility* derived from choosing Product A, and ϵ_A is an *unknown utility* derived from choosing Product A. Likewise for Product B. The significance of the unknown utilities ϵ_A and ϵ_B is that, with respect to the market of alternatives A and B, a consumer's perception of the utility of these alternatives will be influenced by factors that a modeler is unable to account for.

We decompose the known utility that a consumer assigns to a Product into *inherent bias*, *social bias* from neighbors, and *applied bias* from marketing. For each pair of neighbors i and j, there is a parameter θ_{ij} representing the *social bias* or *interaction strength* between i and j. For each consumer i, there is parameter α_i indicating an *inherent bias* toward Product A, where $\alpha_i < 0$ indicates an inherent bias towards Product B. Consumer i receives an allocation of *marketing strength* m_A^i from Company A and an allocation of *marketing strength* m_B^i from Company B. It is important to note that the marketing strengths m_A^i and m_B^i represent perceived utility that consumer i assigns to Products A and B, respectively, as a result of marketing, rather than investment of resources by Companies A and B. That is, a company could spend a lot of money on ineffective marketing or very little money on rather effective marketing, and it is natural to assume that such a *marketing response* would be consumer-specific. We discuss this further in Section 3.2.

Given choices $\mathbf{x}_{\partial i}$ of his neighbors, the contingent utilities for consumer i conditioned on the choices $\mathbf{x}_{\partial i}$ of his neighbors are

$$U = \begin{bmatrix} \alpha_i + m_A^i - m_B^i + \sum_{j \in \partial i} \theta_{ij} x_j + \epsilon_A^i \\ -\alpha_i - m_A^i + m_B^i - \sum_{j \in \partial i} \theta_{ij} x_j + \epsilon_B^i \end{bmatrix}. \quad (3)$$

The main assumption in this parametrization is the symmetric influence between neighbors, i.e., the utility derived by a consumer depends on his choice and the choice of his neighbor, i.e., the contingent utility for consumer i from choosing A given that his neighbor j chooses B, is the same as the utility that j receives from choosing A given that i chooses B. See [3] for a decomposition of *games* corresponding to symmetric influence.

2.2 Choice Response Dynamics

Assume that consumer i updates his choice at time t. He observes the choices $\{x_j : j \in \partial i\}$ of his neighbors. Modeling the unknown sources of utility ϵ_A and ϵ_B as independent and identically distributed *extreme values*, it can be shown [17], [27], that in maximizing the decomposition of utility in (3), consumer i chooses between Products A and B with respective conditional probabilities $p(A|\mathbf{x}_{\partial i}^{(t)})$ and $p(B|\mathbf{x}_{\partial i}^{(t)})$, summarized as

$$p_{i|\mathbf{x}_{\partial i}^{(t)}} = \begin{bmatrix} \frac{1}{Z_{i|\mathbf{x}_{\partial i}^{(t)}}} \exp\{\alpha_i + m_A^i - m_B^i + \sum_{j \in \partial i} \theta_{ij} x_j^{(t)}\} \\ \frac{1}{Z_{i|\mathbf{x}_{\partial i}^{(t)}}} \exp\{-\alpha_i - m_A^i + m_B^i - \sum_{j \in \partial i} \theta_{ij} x_j^{(t)}\} \end{bmatrix}, \quad (4)$$

where $Z_{i|\mathbf{x}_{\partial i}^{(t)}} = \exp\{\alpha_i + m_A^i - m_B^i + \sum_{j \in \partial i} \theta_{ij} x_j^{(t)}\} + \exp\{-\alpha_i - m_A^i + m_B^i - \sum_{j \in \partial i} \theta_{ij} x_j^{(t)}\}$ is referred to as the *local partition function* for i conditioned on $\mathbf{x}_{\partial j}^{(t)}$. Using (1), this can be expressed succinctly as

$$p(x_i|\mathbf{x}_{\partial i}^{(t)}) = \frac{1}{Z_{i|\mathbf{x}_{\partial i}^{(t)}}} \exp\{\sum_{j \in \partial i} \theta_{ij} x_i x_j^{(t)} + x_i(\alpha_i + m_A^i - m_B^i)\}. \quad (5)$$

One can show [14] that the logit response model of (5) satisfies the *independence from irrelevant alternatives* (IIA) property, which states that the relative likelihood of choosing one product over another does not change when an additional product is available as an alternative. This property is due to the assumption that the unknown sources of utility for the different choices are independent of one another. Luce [14] derived the logit choice model viewing IIA as an axiom. While IIA is often violated in practice, in markets where the alternatives do not possess the IIA property, dependence in unknown sources of utility among the alternatives leads to either a *nested* or *cross-nested* [27] logit model, which can be formulated as sequential logit models where, for example, a logit choice is made between *types* of alternatives, followed by a type-specific logit model among alternatives of the chosen type. In other words, even in situations where the logit model is not appropriate *per se*, it still serves as the basic building block for more realistic choice models. The reader is referred to [27] for an excellent discussion of ((cross) nested) logit models. If one assumes that the unknown sources of utility ϵ_A and ϵ_B are instead *sums* of i.i.d. random variables, then one would derive a *probit* model for the choice dynamics [15].

2.3 Steady-State Behavior

The collection of choice responses (5) constitute a *specification* [10], [3] in that it *specifies* a set of conditional distributions with respect to which individual consumers make their choices. If the specification satisfies the Dobrushin-Langford-Ruelle (DLR) consistency condition [10], then there is *at least one* equilibrium Gibbs probability measure, on the set of choice configurations on the network, whose conditional distributions for individual consumers conditioned on the choices of their neighbors, are precisely those given by the specification. In this paper we do not consider the issue of phase transition [10] in which there are multiple equilibria.

The contingent influence of two neighbors is referred to as a *game*. A game is said to be a *potential game* [18] if, rather than a separate utility function for each neighbor, there is a single function of the bivariate choices, called a *potential* function, such that if either participant to the game changes his choice, the resulting difference in the value of the potential function equals the difference in utility that the consumer who changed his choice receives. Blume [3] shows that 2×2 symmetric games always have a potential, so the response dynamics of (5) converge to the *equilibrium Gibbs distribution* given by

$$p(\mathbf{x};\theta) = \frac{1}{Z(\theta)} \exp\{\sum_{\{i,j\}\in E} \theta_{ij} x_i x_j + \sum_{i\in V} \theta_i x_i\}, \tag{6}$$

where $\theta_i = \alpha_i + m^i_A - m^i_B$, and

$$Z(\theta) = \sum_{\mathbf{x}} \exp\{\sum_{\{i,j\}\in E} \theta_{ij} x_i x_j + \sum_{i\in V} \theta_i x_i\} \tag{7}$$

is the (global) partition function.

For consumer i, we will refer to θ_i as the *direct bias* on i; and for a neighbor $j \in \partial i$, to θ_{ij} as the *social bias* on i from j. The direct bias is the sum of the *inherent bias* α_i and the *applied bias* $m^i_A - m^i_B$.

2.4 Belief Propagation

Belief Propagation [23] is a standard algorithm for computing probabilities within Gibbs fields. In this section we provide a brief overview of relevant concepts used in Section 3.3 to optimize marketing allocation on a cycle.

For consumer i, we define the *self-potential*[1]

$$\Phi_i \triangleq \begin{bmatrix} e^{\alpha_i + m^i_A - m^i_B} \\ e^{-\alpha_i - m^i_A + m^i_B} \end{bmatrix},$$

and likewise for a pair of neighboring consumers i and j with interaction strength θ_{ij}, we define the *edge-potential*

$$\Psi_{ij} \triangleq \begin{bmatrix} e^{\theta_{ij}} & e^{-\theta_{ij}} \\ e^{-\theta_{ij}} & e^{\theta_{ij}} \end{bmatrix}. \tag{8}$$

The *belief* Z_i for consumer i is a vector with components $[Z_i(x_i)]$ defined as

$$Z_i(x_i) = \Phi_i(x_i) \sum_{\mathbf{x}_{V\setminus i}} \prod_{j,k} \Psi_{jk}(x_j, x_k) \prod_j \Phi_j(x_j). \tag{9}$$

Our interest in the belief Z_i is due to the fact that normalizing it gives the probabilities that consumer i chooses A or B. That is,

$$p_i(x) = \frac{Z_i(x)}{\sum_{x\in\{A,B\}} Z_i(x)}.$$

Belief Propagation (BP) is an algorithm for computing beliefs. It was designed, and is both optimal and efficient, for acyclic networks, and involves the recursive computation of *messages* between neighboring consumers. It is helpful to think of the message from consumer j to a neighbor i as being computed by the *fusion* of incoming message vectors $m_{k\to j}$ from neighbors other than i, and the vector Φ_j for consumer j, followed by *propagation* via multiplication with the matrix Ψ_{ij}. The belief at node i can be computed from incoming messages as

$$Z_i = \Phi_i \prod_{j\in\partial i} m_{j\to i}, \tag{10}$$

where $\prod$ is component-wise multiplication. The messages are computed recursively as

$$m_{j\to i} = \Psi_{ij}\Phi_j \prod_{k\in\partial j\setminus i} m_{k\to j}. \tag{11}$$

It will be useful in the following analysis to keep in mind the identity[2]

$$\begin{aligned}\Psi_{ij} &= \begin{bmatrix} e^{\theta} & e^{-\theta} \\ e^{-\theta} & e^{\theta} \end{bmatrix} \\ &= \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} C_{\theta_{ij}} & 0 \\ 0 & S_{\theta_{ij}} \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix},\end{aligned}$$

where $C_\theta = \cosh\theta$ and $S_\theta = \sinh\theta$. This can be seen from the spectral decomposition of Ψ_{ij} [1]. With this decomposition, the product of Ψ matrices corresponding to a sequence of consumers with degree 2 and no direct biases, simplifies to

$$\begin{aligned}&\Psi_{k+1,k}\cdots\Psi_{i-1,i} \\ &= \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \begin{bmatrix} \prod_{j=i}^{k+1} C_{\theta_{j,j-1}} & 0 \\ 0 & \prod_{j=i}^{k+1} S_{\theta_{j,j-1}} \end{bmatrix} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}.\end{aligned} \tag{12}$$

3 A MARKETING GAME

We now introduce the main idea of this paper: that the choice response models rigorously developed in economics and other social sciences [26], [14], [17]; their natural extension to interdependent decision-making [3], [6]; and the plethora of theory and algorithms developed for such interconnected systems [28], provide a strong foundation for now analyzing and optimizing the influence of marketing on the coordinated decision-making of those in the network. Moreover, the models derived within this framework, i.e., choice responses (5) and equilibria (6), are *data-driven* in the sense of agreeing with observed data but making no additional assumptions [7]. The importance of a data-driven approach to marketing has long proved successful [21] in the targeting of individual consumers. The framework of socially-contingent random utility enables a systematic approach wherein a model of a social network's collective decision-making can be learned from data. By including marketing allocation within this model, a Company can optimize expected market share over candidate market allocations.

[1] This is a rather common abuse of notation, in which both the exponent and the exponential can be referred to as a potential.

[2] Beliefs can be correctly computed by eliminating common factors between components of potentials and messages. We do so here for simplicity.

3.1 Performance Criterion

The natural performance criterion for a Company to consider in optimizing its marketing allocation is *market share* [2]. The market share for Company A with respect to consumer k is the probability $p_k(A)$ that consumer k chooses Product A. Likewise for Product B. With respect to the entire network, the market share of a Company is the sum, over all consumers in the network, of the probabilities that each consumer chooses their Product. The *bias* of consumer k is the difference in the probabilities of selecting A and B, i.e., the moment for consumer k:

$$\mu_k = p_k(A) - p_k(B) = \frac{Z_k(A) - Z_k(B)}{Z_k(A) + Z_k(B)} . \tag{13}$$

The *total bias* on the network is the sum of the biases of all consumers,

$$\sum_{k \in V} \mu_k = \sum_{k \in V} [p_k(A) - p_k(B)] = \sum_{k \in V} p_k(A) - \sum_{k \in V} p_k(B) , \tag{14}$$

which is the difference in market share of the two companies.

3.2 Marketing Allocation

The utilities m_A^i and m_B^i due to marketing indicate *marketing strengths* applied to a consumer through effective marketing. The *art* of marketing has long been a *science* [21], driven by research and the articulation of precise rules. Through market research, a Company learns effective means of biasing consumer choice through enhanced perception of Product utility. This suggests the need for a *marketing response* function that captures the relationship between expenditure on a form of marketing and the resulting *applied bias*, or perception of utility, that such marketing induces in a consumer.

The respective *marketing responses* R_A^i and R_B^i for Companies A and B will be functions of appropriately dollar denominated units of investment, denoted d_A^i and d_B^i, in marketing to consumer i. The resulting applied biases m_A^i and m_B^i can be expressed as

$$m_A^i = R_A^i(d_A^i) \qquad m_B^i = R_B^i(d_B^i) .$$

The precise shape of these functions would, in practice, be determined by market research such as surveys and focus groups. However, we can assume certain characteristics that such functions would possess.

There would be a saturation effect or diminishing returns where additional marketing would not result in any appreciable increased likelihood of a consumer choosing a Product. Moreover, the response of a consumer to marketing by a Company would likely depend on any inherent bias the consumer has towards Products A or B. If a consumer has a bias towards one or the other product, we would expect that he will be less responsive to marketing from *both* Companies than if he has no bias. For instance, if a consumer is biased in favor of Product A, marketing by Company A will only incrementally add to the consumer's perception of Product A's utility. On the other hand, if a consumer is biased in favor of Product B, marketing by Company A can only do so much.

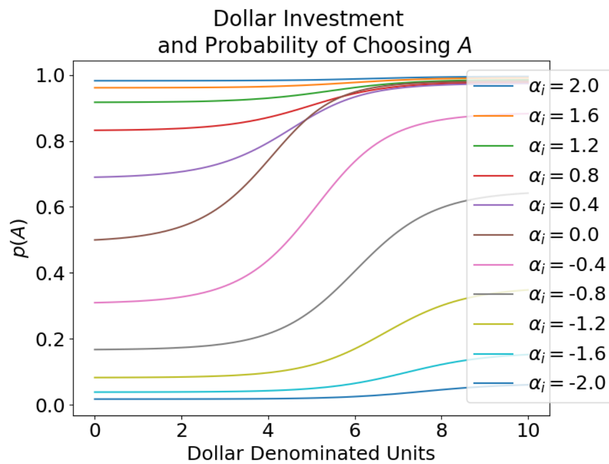

Figure 1: Probability of choosing Product A as a function of investment d_A^i for different values of α_i.

Consider a consumer i whose choice response probability distribution depends only on his inherent bias α_i. That is, we assume that consumer i does not receive any marketing and we are not taking into account any social bias influencing his decision. The probability that consumer i chooses Product A, then, can be computed from (5) by setting $m_A^i = m_B^i = 0$ and assuming that either there are no neighbors or setting $\theta_{ij} = 0$ for each. Now consider the following marketing responses:

$$m_A^i = \frac{1}{1+|\alpha_i|}\left[\frac{2}{1+e^{5+|\alpha_i|-d_A^i}} - \frac{2}{1+e^{5+|\alpha_i|}}\right] , \tag{15}$$

$$m_B^i = \frac{1}{1+|\alpha_i|}\left[\frac{2}{1+e^{5+|\alpha_i|-d_B^i}} - \frac{2}{1+e^{5+|\alpha_i|}}\right] ,$$

which are sigmoid functions shifted and scaled by the inherent bias α_i. Sigmoid functions express the Weber-Fechner law which can model perceived utility of a Product resulting from varying degrees of exposure [5]. In general, the effect of inherent bias on the marketing strength will depend on the product in favor of which the consumer is biased. However, for now, to keep things simple and get the ball rolling, let us assume that the only difference is the expenditure d_A^i or d_B^i. Figure 1 shows the resulting probability of choosing A as a function of investment d_A^i, for different values of the inherent bias α_i, again assuming no social biases. Recall that $\alpha_i > 0$ indicates a bias in favor of A, $\alpha_i < 0$ in favor of B. Incorporating the marketing response functions, the direct bias on consumer i is now

$$\theta_i = \alpha_i + \frac{1}{1+|\alpha_i|}\frac{2}{1+e^{5+|\alpha_i|-d_A^i}} - \frac{1}{1+|\alpha_i|}\frac{2}{1+e^{5+|\alpha_i|-d_B^i}} = \alpha_i + \frac{1}{1+|\alpha_i|}\left[\frac{2}{1+e^{5+|\alpha_i|-d_A^i}} - \frac{2}{1+e^{5+|\alpha_i|-d_B^i}}\right] .$$

3.3 Optimization of Marketing Allocation

We now consider the heart of *A Marketing Game*, optimizing the sets of consumers that Companies A and B target for marketing using the optimization criterion of *total bias* (14) discussed in Section 3.1.

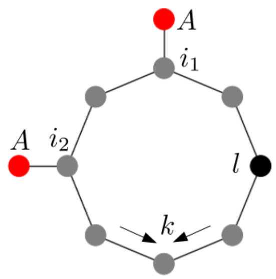

Figure 2: A cycle with consumer l indicated for conditioning, and two units of marketing applied to consumers i_1 and i_2 by Company A.

To provide a template for the kinds of questions one can address within the framework of this model, and motivate further research into this problem, we focus our analysis on an optimal allocation for a single Company under the assumption of uniform social biases on a cycle of N consumers, for example as illustrated in Figure 2. Optimal allocations will be considered with respect to the *true* direct and social bias parameters $\{\theta_i, \theta_{ij}\}$ on the network, implicitly assuming perfect estimation.

As shown in the following proposition, if there are no direct biases on the network, consumers are equally likely to choose Products A and B.

Proposition 3.1. *Let G be a cycle network with uniform social biases. Then for each consumer k, the bias of k is $\mu_k = 0$.*

Proof. One can compute beliefs on a cycle using Conditioning [22], [24], a variant of BP that operates by opening up a cyclic network at a set of nodes constituting a *loop cutset*, a subset of nodes whose removal eliminates all cycles from the network, then performing standard acyclic BP on the opened network *for each configuration of the loop cutset nodes*. Then, the beliefs computed under different hypotheses of the loop cutset nodes are summed together to get overall beliefs (9). See Figure 2 for an illustration. Opening up the cycle at a node l and using (11), (10), and (12), one can show that the belief at consumer k, conditioned on $x_l = A$ is

$$Z_k^{(x_l=A)} = \begin{bmatrix} C_\theta^N + C_\theta^{l-k} S_\theta^{N-l+k} + C_\theta^{N-l+k} S_\theta^{l-k} + S_\theta^N \\ C_\theta^N - C_\theta^{l-k} S_\theta^{N-l+k} - C_\theta^{N-l+k} S_\theta^{l-k} + S_\theta^N \end{bmatrix},$$

and likewise, that the belief at consumer k conditioned on $x_l = B$ is

$$Z_k^{(x_l=B)} = \begin{bmatrix} C_\theta^N - C_\theta^{l-k} S_\theta^{N-l+k} - C_\theta^{N-l+k} S_\theta^{l-k} + S_\theta^N \\ C_\theta^N + C_\theta^{l-k} S_\theta^{N-l+k} + C_\theta^{N-l+k} S_\theta^{l-k} + S_\theta^N \end{bmatrix}.$$

From this we can compute the belief for consumer k as

$$\begin{aligned} Z_k &= Z_k^{(x_l=A)} + Z_k^{(x_l=B)} \\ &= \begin{bmatrix} C_\theta^N + S_\theta^N \\ C_\theta^N + S_\theta^N \end{bmatrix}, \end{aligned}$$

and using (13), we see that $\mu_k = 0$. □

We now consider the case that a single unit of marketing allocation is placed at a consumer i on the cycle. As we will see, the beliefs, and therefore bias, at a consumer k will be a function of the distance between i and k. Let Δ_i^k and $\bar{\Delta}_i^k$ be the distances between consumers i and k going either way around the cycle.

Proposition 3.2. *Let G be a cycle network with uniform social biases, and suppose i is the consumer receiving the single unit marketing allocation from Company A. Then, the bias of consumer k is*

$$\mu_k = \frac{S_{m_A}\left[C_\theta^{\Delta_i^k} S_\theta^{\bar{\Delta}_i^k} + C_\theta^{\bar{\Delta}_i^k} S_\theta^{\Delta_i^k}\right]}{C_{m_A}\left[C_\theta^N + S_\theta^N\right]}.$$

Proof. We again using the idea of Conditioning to compute beliefs on the cycle. While the node l at which one opens up the cycle need not coincide with the consumer i receiving a marketing allocation, it does simplify the calculation. Doing this, one can show that conditioned on $x_i = A$, the belief at consumer k is

$$Z_k^{(x_i=A)} = e^{m_A}\begin{bmatrix} C_\theta^N + C_\theta^{\Delta_i^k} S_\theta^{\bar{\Delta}_i^k} + C_\theta^{\bar{\Delta}_i^k} S_\theta^{\Delta_i^k} + S_\theta^N \\ C_\theta^N - C_\theta^{\Delta_i^k} S_\theta^{\bar{\Delta}_i^k} - C_\theta^{\bar{\Delta}_i^k} S_\theta^{\Delta_i^k} + S_\theta^N \end{bmatrix},$$

and likewise, that the belief at consumer k conditioned on $x_i = B$ is

$$Z_k^{(x_i=B)} = e^{-m_A}\begin{bmatrix} C_\theta^N - C_\theta^{\Delta_i^k} S_\theta^{\bar{\Delta}_i^k} - C_\theta^{\bar{\Delta}_i^k} S_\theta^{\Delta_i^k} + S_\theta^N \\ C_\theta^N + C_\theta^{\Delta_i^k} S_\theta^{\bar{\Delta}_i^k} + C_\theta^{\bar{\Delta}_i^k} S_\theta^{\Delta_i^k} + S_\theta^N \end{bmatrix}.$$

The conclusion is attained by adding $Z_k^{(x_i=A)}$ and $Z_k^{(x_i=B)}$. □

Figure 3 (a) shows the probability that consumers choose Product A as a function of their location on a cycle for a given placement of marketing allocation. Using (15) for the marketing response, Figure 3 (b) shows the total bias of the network as a function of dollar denominated investment. While we assume uniform social biases in the derivations of this section, extension to non-uniform social biases is straightforward, for example by replacing $C_\theta^{\Delta_i^j}$ with $\prod_{k=1}^{j-1} C_{\theta_{k,k+1}}$. In Figure 3 (c) we show a pattern of non-uniform social biases that decrease linearly in both directions from a given site, and in (d) we look at total bias as a function investment when the single unit marketing allocation is placed at the consumer with largest or smallest social biases, for different inherent biases of this consumer. The takeaway from (d) is that at lower levels of investment, it is better for a Company to target consumers who are loyal to their Product, while at higher levels of investment, it may be better to target consumers who are biased towards another Product but in a region of the network with large social biases.

Noting how rapidly the probabilities in Figure 3 (a) decrease as a function of the distance from the single marketing allocation, a natural question is whether, for a given amount of marketing investment, it is better to distribute the marketing among multiple consumers. We now address the question of allocating two units of marketing on a cycle. To further economize notation, we define

$$\lambda(i,j) \triangleq C_\theta^{\Delta_i^j} S_\theta^{\bar{\Delta}_i^j} + C_\theta^{\bar{\Delta}_i^j} S_\theta^{\Delta_i^j}.$$

Using this notation we now address the optimal allocation of two units of marketing given a fixed investment in each.

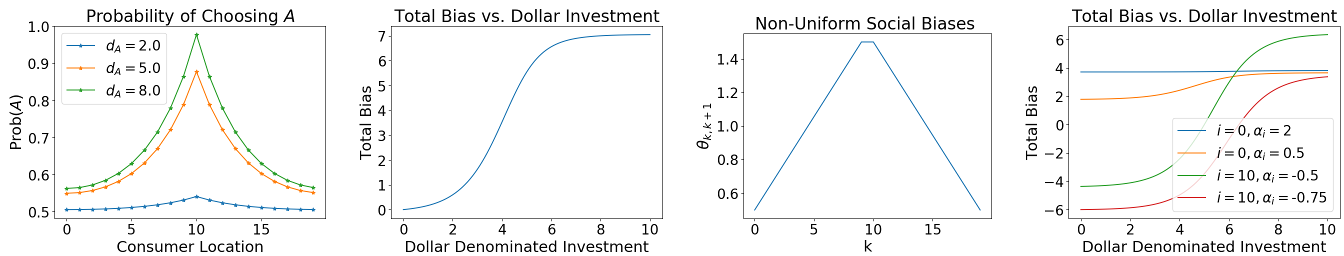

(a) (b) (c) (d)

Figure 3: (a) and (b) assume uniform social biases, and respectively show: the probability of choosing A for each consumer on the cycle, and the total bias as a function of dollar denominated investment; (c) non-uniform social biases, and (d) total bias as a function of investment for different inherent biases at two different consumer locations, with the social biases of (c).

Proposition 3.3. *Assume a cycle with no inherent biases and uniform social biases. Then, the optimal placement of two units of marketing, $m_A^{i_1}$ and $m_A^{i_2}$, is at evenly spaced consumers on the cycle.*

Proof. Let $\Delta_{i_1}^{i_2}$ and $\bar{\Delta}_{i_1}^{i_2}$ be the distances between consumers i_1 and i_2 going either way around the cycle. Likewise $\Delta_{i_1}^{k}$ and $\bar{\Delta}_{i_1}^{k}$ for the distances between consumers i_1 and k; and $\Delta_{i_2}^{k}$ and $\bar{\Delta}_{i_2}^{k}$ for consumers i_2 and k. The total bias on the cycle with respective allocation strengths $m_A^{i_1}$ and $m_A^{i_2}$ at consumers i_1 and i_2 is

$$\sum_k \mu_k \quad = \quad \sum_k \frac{C_{m_A^{i_2}} S_{m_A^{i_1}} \lambda(i_2,k) + C_{m_A^{i_1}} S_{m_A^{i_2}} \lambda(i_1,k)}{C_{m_A^{i_1}} C_{m_A^{i_2}} \lambda(i_1,i_1) + S_{m_A^{i_1}} S_{m_A^{i_2}} \lambda(i_1,i_2)} \,. \tag{16}$$

Note that for a given allocation $m_A^{i_1}$ and $m_A^{i_2}$, the denominator

$$D(i_1,i_2,m_A^{i_1},m_A^{i_2}) \triangleq C_{m_A^{i_1}} C_{m_A^{i_2}} \lambda(i_1,i_1) + S_{m_A^{i_1}} S_{m_A^{i_2}} \lambda(i_1,i_2)$$

is constant for all consumers k on the cycle. As for the numerator, note that

$$\begin{aligned}
&\sum_k C_{m_A^{i_2}} S_{m_A^{i_1}} \lambda(i_2,k) + C_{m_A^{i_1}} S_{m_A^{i_2}} \lambda(i_1,k) \\
&\qquad = \sum_k C_{m_A^{i_2}} S_{m_A^{i_1}} \lambda(i_2,k) + \sum_k C_{m_A^{i_1}} S_{m_A^{i_2}} \lambda(i_1,k) \\
&\qquad = C_{m_A^{i_2}} S_{m_A^{i_1}} \left[\sum_k C_\theta^{\bar{\Delta}_{i_2}^k} S_\theta^{\Delta_{i_2}^k} + \sum_k C_\theta^{\Delta_{i_2}^k} S_\theta^{\bar{\Delta}_{i_2}^k} \right] \\
&\qquad\quad + C_{m_A^{i_1}} S_{m_A^{i_2}} \left[\sum_k C_\theta^{\bar{\Delta}_{i_1}^k} S_\theta^{\Delta_{i_1}^k} + \sum_k C_\theta^{\Delta_{i_1}^k} S_\theta^{\bar{\Delta}_{i_1}^k} \right] \\
&\qquad = C_{m_A^{i_2}} S_{m_A^{i_1}} \left[\sum_{d=0}^{N} C_\theta^{d} S_\theta^{N-d} + \sum_{d=0}^{N} C_\theta^{N-d} S_\theta^{d} \right] \\
&\qquad\quad + C_{m_A^{i_1}} S_{m_A^{i_2}} \left[\sum_{d=0}^{N} C_\theta^{d} S_\theta^{N-d} + \sum_{d=0}^{N} C_\theta^{N-d} S_\theta^{d} \right] \\
&\qquad \triangleq N(m_A^{i_1}, m_A^{i_2})
\end{aligned}$$

is constant for all i_1 and i_2, depending only on the marketing strengths $m_A^{i_1}$ and $m_A^{i_2}$. Thus the total bias (16) is

$$\sum_k \mu_k \quad = \quad \frac{N(m_A^{i_1}, m_A^{i_2})}{D(i_1,i_2,m_A^{i_1},m_A^{i_2})}.$$

It is straightforward to show that $D(i_1,i_2,m_A^{i_1},m_A^{i_2})$ is minimized at $\Delta_{i_1}^{i_2} = N/2$, or in other words, when i_1 and i_2 are evenly spaced along the cycle. □

Extrapolating from these results, we conjecture that Proposition 3.3 can be extended to the case of a fixed investment distribution among k units of allocation, in that the optimal placement of the k units would be equally placed around the cycle. Such a conclusion would seem to agree with recent criticism of the so-called influentials hypothesis [30], in that it is better to distribute marketing throughout the network.

In Figure 4 we assume uniform social biases, a fixed investment $d = d_A^1 + d_A^2$ among two allocations, and consider the total bias as a function of the investment d_A^1 in the first allocation at site $i_1 = 0$, for different placements of the second allocation i_2. That is, we consider the distribution of marketing investment in terms of both the locations of the consumers within the network and the monetary expenditure for each consumer. In (a), where $d = 10$, we can see that for all candidate placements of the second allocation, if the investment by Company A is skewed towards either allocation, Company A will fail to achieve the total bias that is possible with a more equitable distribution of its marketing investment. Furthermore, as shown in the above proposition, for any distribution of the marketing investment d between i_1 and i_2, total bias is maximized with i_1 and i_2 evenly spaced around the cycle. The maximum possible total bias is achieved with i_1 and i_2 evenly spaced around the cycle, with an equitable distribution of the total marketing investment.

On the other hand, in (b), where $d = 8$, for any placement i_2 of the second allocation, distributing the budget evenly actually results in the smallest possible total bias. That is, when the budget is small, it is better to distribute resources unevenly. Making the budget d even smaller results in even more skewed resource allocations, such that if the budget is small enough, total bias is maximized by allocating the entire resource budget to i_1. Nevertheless, Proposition

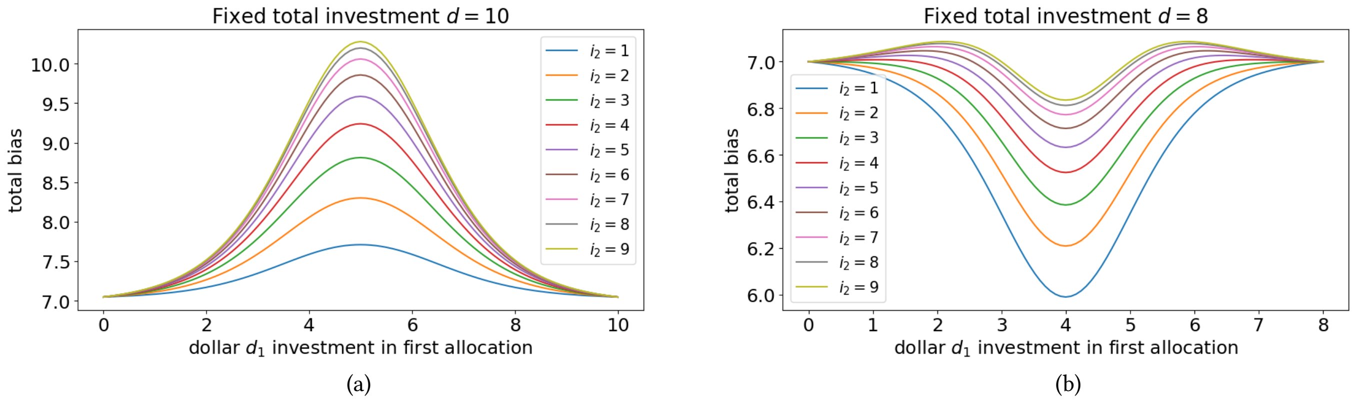

Figure 4: Total bias as a function of dollar investment d_1 in consumer $i_1 = 0$ for different locations i_2 of the second allocation, with total dollar investment (a) $d = d_1 + d_2 = 10$; (b) $d = d_1 + d_2 = 8$.

3.3 still holds, that for any distribution d_A^1, d_A^2 of resources, evenly spaced i_1 and i_2 is optimal. Note by comparing (a) and (b) that when the budget is smaller ($d = 8$), the optimal d_A^1 is different from the optimal d_A^1 when the budget is larger ($d = 10$). That is, it is not simply a matter of allocating a certain amount to consumer i_1 and the remainder to the second allocation. With regards to the influentials hypothesis, and criticism thereof, whether it is better to distribute resources throughout the network or concentrate them on a small number of consumers depends upon the marketing budget. And the only way to determine whether it is "small" or "large" is to run the analysis in the particular scenario. It is a strength of this model that it provides a framework for doing such analysis.

4 CONCLUDING REMARKS

In this paper we have introduced a model for consumer decision-making on a social network that builds upon the rigorous foundation of random utility theory and its extension to socially-contingent decision-making. In particular, we have included the *marketer* within the parametrization of random consumer choice, and have proposed a *marketing response* characterizing the *applied bias* resulting from appropriately denominated units of investment. We have illustrated the ability of our model to address questions relevant to the strategic distribution of marketing resources. Extensions of this approach to markets with more than two alternatives is straightforward. As a final note, one can use the derivations of the previous section to show that if Companies A and B each get a single unit of marketing allocation of equal strength, the Nash equilibria, between Companies A and B, consist of allocating to the same consumer, or two antipodally located consumers.

REFERENCES

[1] R.J. Baxter, "Exactly Solved Models in Statistical Mechanics," Dover, 2007.
[2] D.E. Bell, R.L. Keeney, and J.D.C. Little, "A Market Share Theorem," *Journal of Marketing Research*, **12**(2), May 1975.
[3] L.E. Blume, "Statistical Mechanics of Strategic Interaction," *Games and Economic Behavior*, **5**(3), 1993.
[4] H.D. Block and J. Marschak, "Random orderings and stochastic theories of responses," *Contributions to Probability and Statistics*, Stanford University Press, 1960.
[5] S.H. Britt, "How Weber's Law Can Be Applied to Marketing," *Business Horizons*, **18**(1), 1975.
[6] W.A. Brock and Durflauf, "Discrete Choice with Social Interactions," *Review of Economic Studies*, **68**, 2001.
[7] T.M. Cover and J.A. Thomas, "Elements of Information Theory," Wiley, 1991.
[8] A. Fazeli and A. Jadbabaie, "Game Theoretic Analysis of a Strategic Model of Competitive Contagion and Product Adoption in Social Networks," http://repository.upenn.edu/ese_papers/618, December 2012.
[9] R.J. Glauber, "Time-Dependent Statistics of the Ising Model," *Journal of Mathematical Physics*, **4**, 1963.
[10] H.O. Georgii, "Gibbs Measures and Phase Transitions," De Grutyer, 1988.
[11] M. Gladwell, "The Tipping Point," Little, Brown, and Company, 2000.
[12] P.E. Green, F.J. Carmone, and D.P. Wachspress, "On the Analysis of Qualitative Data in Marketing Research," *Journal of Marketing Research*, February 1977.
[13] J.C. Harsanyi and R. Selton, "A General Theory of Equilibrium Selection in Games", MIT Press, 1988.
[14] R.D. Luce, "Individual Choice Behavior," Dover, 1959.
[15] R.D. Luce, "Thurstone's Discriminal Process Fifty Years later," *Psychometrika*, **42**(4), 1977.
[16] N.K. Malhotra, "The Use of Linear Logit Models in Marketing Research," *Journal of Marketing Research*, February 1984.
[17] D. McFadden, "Conditional Logit Analysis of Qualitative Choice Behavior," *In Frontiers in Econometrics*, New York: Academic Press, 1974.
[18] D. Monderer and L.S. Shapley, "Potential Games," *Games and Economic Behavior*, **14**, 1996.
[19] A. Montanari and A. Saberi, "The spread of innovations in social networks," *PNAS*, **107**(47), 2010.
[20] J. Von Neumann and O. Morgernstern, "Theory of Games and Economic Behavior," Princeton University Press, 1944.
[21] D. Ogilvy, "Ogilvy On Advertising," Vintage, 1985.
[22] J. Pearl, "A Constraint Propagation Approach To Probabilistic Reasoning," *Uncertainty in Artificial Intelligence*, Elsevier, New York, pps. 357-369, 1986.
[23] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufman, June 2014.
[24] M.G. Reyes and D.L. Neuhoff, "Local Conditioning for Undirected Graphs," *Information Theory and Applications* workshop, February 2017.
[25] H.A. Simon, "A Behavioral Model of Rational Choice," *The Quarterly Journal of Economics*, **69**(1), 1955.
[26] L.L. Thurstone, "Psychological Analysis," *The Amer. Jrnl. Psych.*, Vol. 38, No. 3, July 1927.
[27] K. Train, "Discrete Choice Models with Simulation," Cambridge University Press, 2002.
[28] M. Wainwright and M. Jordan, "Graphical Models, Exponential Families, and Variational Inference," Technical Report 649, UC Berkeley, 2003.
[29] D.J. Watts, "A simple model of global cascades on random networks," *PNAS*, Vol. 99, No. 9, 2002.
[30] D.J. Watts and P.S. Dodds , "Influentials, Networks, and Public Opinion Formation," *Jrnl. Consumer Research*, Vol. 34, December 2007.
[31] H.P. Young,"The Evolution of Conventions," *Econometrica*, Vol. 61, No. 1, January 1993.