

Likes, Budgets, and Equilibria: Designing Contests for Socially Optimal Advertising

Sayantika Mandal¹, Harman Agrawal¹, and Swaprava Nath¹

¹IIT Bombay

{22d0379,22b3012,swaprava}@iitb.ac.in

Abstract

Firms (businesses, service providers, entertainment organizations, political parties, etc.) advertise on social networks to draw people’s attention and improve their awareness of the brands of the firms. In all such cases, the competitive nature of their engagements gives rise to a game where the firms need to decide how to distribute their budget over the agents on a network to maximize their brand’s awareness. The firms (players) therefore need to optimize how much budget they should put on the vertices of the network so that the spread improves via direct (via advertisements or free promotional offers) and indirect marketing (words-of-mouth). We propose a *two-timescale* model of decisions where the communication between the vertices happen in a faster timescale and the strategy update of the firms happen in a slower timescale. We show that under fairly standard conditions, the best response dynamics of the firms converge to a *pure strategy Nash equilibrium*. However, such equilibria can be away from a socially optimal one. We provide a characterization of the *contest success functions* and provide examples for the designers of such contests (e.g., regulators, social network providers, etc.) such that the Nash equilibrium becomes unique and *social welfare maximizing*. Our experiments show that for realistic scenarios, such contest success functions perform fairly well.

1 Introduction

In today’s digitally connected world, the landscape of marketing has undergone a profound transformation. Traditional advertising channels, such as newspapers and television, still have relevance, but the burgeoning influence of social media platforms cannot be overstated. This shift is not merely a matter of preference; it’s driven by the fundamental human inclination to trust recommendations from within our social circles. Using this inherent trust, firms seek to maximize the reach of their products and services through strategic dissemination within social networks.

In a competitive market setting, potential customers (vertex in the social graph) have varying level of information available on different firms. In this paper, we measure this information disparity a vertex has about firms through its *awareness level*. It captures how prominently a firm occupies an agent’s attention or belief relative to its competitors. The awareness of the customers therefore consists of two components: (a) a direct marketing impact coming from the firms through advertisements, promotional offers, etc., and (b) an indirect impact coming from recommendations from their neighbors in the social graph. The consolidated impact on the awareness therefore depends on the network centralities of the vertex w.r.t. its neighbors and from the competing budget allocated to every node by all the

firms. The competing budget component is modeled through the widely used *contest success function* (CSF) (Skaperdas, 1996). Our analysis considers that the firms play best response to each other’s committed budgets to reach equilibrium, and analyzes the welfare at such equilibria.

1.1 Our Contributions

The contributions of this paper is summarized as follows.

- We introduce a fast and slow timescale model of awareness spread in a social network where on the firms update their budget strategies on the slow timescale and the individuals update their awareness on the faster timescale (Section 2.1).
- We generalize the awareness level convergence on the faster timescale for the individuals in a network from two firms in Bimpikis et al. (2016) to m firms (Section 2.1.2).
- On the slower timescale, we show that the budget strategy of the firms converge to a Nash equilibrium (NE) (Theorem 1). However, the *social welfare* at the Nash equilibrium may not be optimal (Example 1).
- Our results show how a network policymaker or a regulatory authority can set the rules of the contest, by choosing an appropriate *contest success function*, such that the NE is unique and the social welfare is maximized (Theorem 2).
- Though the theoretical guarantees depend on certain assumptions, experiments show that the convergence to NE and its desirable social welfare properties are achievable on realistic synthetic data (Section 5).

Due to space limitations, some details of the proofs and experiments have been deferred to the supplementary material.

1.2 Related Work

The related literature of this paper can be broadly classified into three strands.

Modeling and structural guarantees of competitive influence. A significant line of research has focused on formalizing competitive diffusion processes and analyzing the properties of equilibria. Foundational work by Bharathi et al. (2007) study the diffusion of competing innovations in social networks, proposing a tractable game-theoretic model, deriving a $(1 - 1/e)$ -approximation algorithm for best responses, and bounding the *price of competition*. Early works (Alon et al., 2010; Goyal and Kearns, 2012) formalize the game theoretic competitive influence maximizing model and emphasize the role of topology in the existence of Nash equilibrium and influence amplification. Zuo et al. (2022) introduce the Online Competitive Influence Maximization (OCIM) problem, and design bandit-based algorithms with sublinear regret. Tzoumas et al. (2012) models the diffusion process of CIM in a simultaneous two-player non-cooperative game as a linear threshold model. Fazeli et al. (2012a,b); Fazeli et al. (2016) develop a series of game-theoretic models for competitive product diffusion in social networks, where firms simultaneously allocate fixed budgets between seeding strategies and product quality, aiming to maximize the lower bound of their product adoptions.

Marketing resource allocation. A parallel stream of work was investigating the budget allocation perspective, where firms allocate varying resources over nodes to compete for attention. In this scenario, nodes tend to adopt the product of the company that invests the most in them. Notable studies in this area include (Masucci and Silva, 2014, 2017; Varma et al., 2018, 2019; Ansari et al., 2019). The work by Bimpikis et al. (2016) is particularly relevant, as it models targeted advertising in networks, characterizes Nash equilibria, and highlights the impact of network centrality on strategic allocation. Maehara et al. (2015) extend the problem with a

bipartite influence model and establish it as a potential game with a pure Nash equilibrium and a price of anarchy of 2.

Fairness, welfare considerations Several other studies have examined fairness and welfare optimization in competitive influence or marketing games on social networks. Wang et al. (2022); Lu et al. (2013) focus on fair seed or budget allocation by a central host to ensure equitable influence spread among advertisers. Rahmattalabi et al. (2021); Becker et al. (2022) design fairness-driven and randomized mechanisms to ensure balanced ad reach and representation across communities. Banerjee et al. (2019, 2020) emphasize maximizing social welfare at the node level under diffusion-based utility models. These works primarily focus on fairness or social welfare optimization rather than equilibrium outcomes.

Our work complements and extends these strands by providing theoretical guarantees for the convergence of best-response dynamics to Nash equilibria in a two-timescale competitive advertising model. We generalize the awareness model of Bimpikis et al. (2016) to accommodate competition among m firms rather than just two, thereby significantly broadening the scope. Additionally, we study the design of contest-success functions that ensure unique and socially optimal equilibria, bridging the gap between strategic resource allocation and mechanism design in competitive influence networks. We validate these results empirically on synthetic networks.

2 Preliminaries

We consider a market with m competing firms denoted by $\mathcal{M} = \{1, 2, \dots, m\}$, each having an equal and fixed marketing budget that they strategically allocate across users in a social network. Firms compete to attract the attention of these users by distributing their budgets to maximize brand visibility and influence. We consider *awareness* as a cardinal metric of a user's attention towards a particular firm, which lies within $[0, 1]$. This marketing competition is modeled as a finite non-cooperative game on a social network $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N} = \{1, 2, \dots, n\}$ denotes the set of users in the network and $\mathcal{E} = (e_{ij})_{i,j \in \mathcal{N}}$ represents their influence structure through an adjacency matrix, with $e_{ii} = 0$ for all $i \in \mathcal{N}$. Each firm aims to maximize the overall awareness of its brand across the network, subject to its budget constraints.

2.1 A Two-Timescale Communication Model

Figure 1 below illustrates the concept of communication flow on two-time scales: a faster one and a slower one.




Figure 1: *The two-timescale concept*

The slower timescale, $k = 0, 1, 2, \dots, \infty$, governs firms' budget allocations across each time period k of length T , while the faster timescale, $t = 0, T, 2T, \dots, \infty$, captures how information spreads and stabilizes between users within each period. On the onset of the slower timescale, firms offer varying budget shares to different users (*direct communication*) with probability $(1 - \alpha)$, who then disseminate the information received to their neighbors. This information propagation or awareness flow among the nodes (*peer-to-peer communication*) happens in the faster timescale with probability α and converges to its steady state value before the beginning of the next period of the slower timescale. We assume that T is large such that the fast timescale process converges between two periods. This is often observed in practice where

the advertisement budgets (and strategies) are updated by the firms once in every month or even more infrequently, while the viewers of the advertisements update their actions very frequently (perhaps every hour). In this paper, we will overload t to be the running variable within two slower timescale periods k and $k + 1$, where t runs from 0 to T .

2.1.1 Awareness levels of each firm

Each firm $s \in M$ chooses a strategy vector $\mathbf{b}_s = (b_{s1}, \dots, b_{sn})$ representing the allocation of its marketing budget across the n nodes in the network. The set of feasible budget vectors is denoted by $\mathcal{B} = \{\mathbf{b} \in \mathbb{R}^{m \times n} : \mathbf{b}_{si} \geq 0, \forall s \in \mathcal{M}, \forall i \in \mathcal{N}, \mathbf{1}^\top \mathbf{b}_s \leq C\}$ where matrix $\mathbf{b}$, represented as $(\mathbf{b}_s, \mathbf{b}_{-s})$ denotes the joint strategy profile of all firms. Here $\mathbf{b}_{-s}$ denotes the allocation vectors of all competing firms other than s . Each firm is assumed to have an equal total marketing budget, C , which is normalized to 1. We will call this game an *awareness competition game* (ACG) among the firms in the rest of the paper where each player (firm) s chooses their budget vector $\mathbf{b}_s(k)$ at every slow timescale period k .

Let $\mathbf{x}_s(t) = (x_{s1}(t), \dots, x_{sn}(t))$ denote the awareness level vector for firm s at time t , where $x_{si}(t) \in [0, 1]$ represents the awareness of node i toward firm s . The overall awareness level of a node evolves as a linear combination of awareness gained directly through firms' marketing efforts and indirectly through influence from neighboring nodes in the social network. To capture the temporal dynamics, awareness is modeled as a running average of updates across all preceding periods, ensuring both recent and past interactions contribute to the current state.

The basic update rule follows the framework of Bimpikis et al. (2016), which we extend to m players (firms) in a two-timescale setting to separate the faster awareness dynamics from the slower strategic budget updates. Formally, the awareness of node i toward firm s at time t of the faster timescale during the period k to $k + 1$ of the slower timescale is given by:

$$x_{si}(t) = \frac{1}{t} \left[\alpha \sum_{j \in \mathcal{N}_i} e_{ji} x_{sj}(t-1) + (1 - \alpha) h(b_{si}(k), \mathbf{b}_{-si}(k)) \right] + \left(1 - \frac{1}{t}\right) x_{si}(t-1) \quad (1)$$

Here, $e_{ji} \in [0, 1]$ is the (j, i) -th entry of the adjacency matrix $\mathcal{E}$, that represents the influence of agent j on i . The adjacency matrix $\mathcal{E}$ is assumed to be sub-stochastic, i.e., $\sum_{j \in \mathcal{N}_i} e_{ji} \leq 1, \forall i \in \mathcal{N}$, where $\mathcal{N}_i$ is the neighbors of i . The function $h(b_{si}(k), \mathbf{b}_{-si}(k))$, called the *Contest Success Function* (CSF), defines the probability that firm s successfully influences node i given the competing budget allocations. Note that the function is identical for all (firm, agent) pairs and depends only on the budget invested on node i by all the firms. We will use the shorthand $h_{si}(\mathbf{b})$ to refer to $h(b_{si}(k), \mathbf{b}_{-si}(k))$ when the dependency is clear.

2.1.2 Limiting Value of Awareness Vector

The awareness for each node updates according to Equation (1). The steady-state awareness vector for a given firm across the network is given by the following expression that follows an analysis similar to Bimpikis et al. (2016). The convergence happens between two consecutive time periods of the slow timescale k as we discussed earlier. The steady-state awareness vector of firm s at the end of each period k is given by

$$\mathbf{x}_s(k) \approx (1 - \alpha)(I - \alpha \mathcal{E})^{-1} \mathbf{h}_s(\mathbf{b}_s(k), \mathbf{b}_{-s}(k)), \quad (2)$$

where $\mathbf{h}_s(\mathbf{b}_s(k), \mathbf{b}_{-s}(k))$ is the vector $(h(b_{si}(k), \mathbf{b}_{-si}(k)), i \in \mathcal{N})$. The approximation arises due to potential convergence issues over a finite time horizon. However, for analytical tractability, we treat this relationship as an equality in the remainder of the paper. Refer to the supplementary material (Section A.1) for the detailed proof of this derivation, where we extend the two-firm case presented by Bimpikis et al. (2016) to the general m -firm setting.

2.1.3 Payoff of a firm

We define the utility (payoff) of a firm as the difference between its consolidated awareness over the network and the cost incurred through advertising. This formulation captures the trade-off between maximizing awareness and minimizing expenditure. The utility of firm s is given by

$$\begin{aligned} u_s(\mathbf{b}_s(k), \mathbf{b}_{-s}(k)) &= \sum_{i \in \mathcal{N}} x_{si}(k) - \sum_{i \in \mathcal{N}} b_{si}(k) \\ &= \sum_{i \in \mathcal{N}} (1 - \alpha) [(I - \alpha \mathcal{E})^{-1}]_i h_{si}(\mathbf{b}(k)) - \sum_{i \in \mathcal{N}} b_{si}(k). \end{aligned} \quad (3)$$

We will denote $M_i = (1 - \alpha) [\mathbf{1}^T (I - \alpha \mathcal{E})^{-1}]_i = (1 - \alpha) c_i$ as the weighted centrality of each node i where c_i is the non-negative *absorption centrality* (Bimpikis et al., 2016). In the matrix form, this can be expressed as

$$u_s(\mathbf{b}(k)) = \mathbf{1}^T \left[(1 - \alpha) (I - \alpha \mathcal{E})^{-1} \mathbf{h}_s(\mathbf{b}(k)) \right] - \mathbf{1}^T \mathbf{b}_s(k), \quad (4)$$

where firm s strategically allocates its budget across nodes to maximize this utility by balancing awareness gains against advertising costs in a competitive environment.

2.2 Design Desiderata

In this setting, we expect the contest to reach an equilibrium where we want to satisfy certain desirable properties. In the following, we define the equilibrium concept, the goals of design, and a performance metric.

Nash Equilibrium

In this setting, the firms are the players and they choose their budgets as actions over the slower timescale. A natural property that we would look for is that of *pure strategy Nash equilibrium*.

Definition 1 (Nash equilibrium). A budget profile $\mathbf{b}^* = (\mathbf{b}_s^*, \mathbf{b}_{-s}^*) \in \mathcal{B}$ is a *pure strategy Nash equilibrium* (PSNE) if

$$u_s(\mathbf{b}_s^*, \mathbf{b}_{-s}^*) \geq u_s(\mathbf{b}_s, \mathbf{b}_{-s}^*), \forall \mathbf{b}_s, \forall s \in \mathcal{N}.$$

Social Welfare

While the competing agents reach an NE, the goal of a planner (network policymaker or regulator) is to ensure some *social welfare*. We consider the generalized p -mean welfare function, which provides a unifying framework for different welfare criteria (Barman et al., 2020).

Definition 2 (p -mean social welfare). For an aggregated budget allocation profile $\mathbf{b} \in \mathcal{B}$ among m firms, the p -mean social welfare is defined as

$$W_p(\mathbf{b}) = \left(\frac{1}{m} \sum_{s \in \mathcal{M}} u_s(\mathbf{b})^p \right)^{\frac{1}{p}}, \quad (5)$$

where the parameter $p \in (-\infty, 1]$ controls the trade-off between equity and efficiency in welfare aggregation.

Special cases of this measure include the following: (a) Utilitarian (average) welfare ($p = 1$): $W_{\text{avg}}(\mathbf{b}) = \frac{1}{m} \sum_{s \in \mathcal{M}} u_s(\mathbf{b})$, (b) Egalitarian welfare ($p = -\infty$): $W_{\text{egal}}(\mathbf{b}) = \min_{s \in \mathcal{M}} u_s(\mathbf{b})$, and (c) Nash social welfare ($p \rightarrow 0$): $W_{\text{NSW}}(\mathbf{b}) = (\prod_{s \in \mathcal{M}} u_s(\mathbf{b}))^{1/m}$.

Definition 3 (Price of Anarchy). Given the p -mean social welfare measure $W_p(\mathbf{b})$ defined over all feasible budget allocation profiles $\mathbf{b} \in \mathcal{B}$, let $\mathcal{B}^*$ denote the set of Nash equilibria of the ACG. The *Price of Anarchy (PoA)* is defined as

$$\text{PoA} = \frac{\max_{\mathbf{b} \in \mathcal{B}} W_p(\mathbf{b})}{\min_{\mathbf{b} \in \mathcal{B}^*} W_p(\mathbf{b})}, \quad (6)$$

where the numerator corresponds to the welfare under the socially optimal allocation, and the denominator represents the welfare at the worst Nash equilibrium. A higher PoA indicates greater inefficiency arising from competitive behavior.

2.3 Existence of Nash Equilibrium

Existence of a PSNE in this context is guaranteed due to Nash's theorem (Nash Jr, 1950). To see this, given the current game Γ , construct another game Γ' where the pure strategy of every firm $s \in \mathcal{M}$ is to pick one user $i \in \mathcal{N}$ on which the firm invests the entire budget of 1. Every mixed strategy of Γ' is a pure strategy of Γ . Since by Nash's theorem, a mixed strategy Nash equilibrium exists for any finite game, PSNE exists in our game. We will refer to the term PSNE as Nash equilibrium (NE) in the rest of the paper.

3 Best Response Dynamics

Even though the existence of an NE is guaranteed in a game, it is not always easy for the players to discover it and play at the equilibrium. However, there are types of games where the *best response dynamics* (BRD) of the players lead to an NE. BRD is an iterative method where each player sequentially updates their strategy to maximize their utility, given the current strategies of the other players. We show in this section that under certain assumptions, our awareness competition game (ACG) converges to an NE when players play BRD.

At the beginning of each slow-timescale period k , each firm $s \in \mathcal{M}$ choose their budget $\mathbf{b}_s(k)$ while awareness levels $\mathbf{x}_s(t)$ evolve until time $t = T$ on that strategy profile, that is assumed to converge to a steady-state value before the next slow-timescale period $k + 1$. At $k + 1$, each firm updates its budget as a best response (to maximize the utility function in Equation (4)) to competitors' previous budgets $\mathbf{b}_{-s}(k)$. Algorithm 1 illustrates the implementation of BRD for the ACG where every agent updates their budget in the direction of the gradient of their utility as given in line 4. The gradient of the utility function of firm s with respect to its own budget vector $\mathbf{b}_s$ is defined as

$$\begin{aligned} \mathbf{g}_s(\mathbf{b}(k), \boldsymbol{\mu}(k)) &= \nabla_{\mathbf{b}_s} u_s(\mathbf{b}(k)) + \mu_s(k) \in \mathbb{R}^{n \times 1} \\ \text{where } \mu_s(k) &= \begin{cases} -\nabla_{\mathbf{b}_s} u_s(\mathbf{b}(k)), & \text{if } \mathbf{b}(k) + \gamma \mathbf{g}_s(\mathbf{b}(k), \boldsymbol{\mu}(k)) \notin \mathcal{B} \\ 0, & \text{otherwise} \end{cases} \end{aligned} \quad (7)$$

We vectorize the joint budget matrix $\mathbf{b}$ for all m firms across n agents into a flattened column vector $\mathbf{b} \in \mathbb{R}^{mn \times 1}$. Consequently, the corresponding joint gradient for all firms can be written concisely by stacking the individual firm gradients as $\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k)) = [\mathbf{g}_s(\mathbf{b}(k), \boldsymbol{\mu}(k))]_{s \in \mathcal{M}} \in \mathbb{R}^{mn \times 1}$.

Algorithm 1 Best Response Dynamics

- 1: **Initialize** budget vectors $\mathbf{b}_s(0) \in \mathcal{B}$, $\forall s \in \mathcal{M}$, number of iterations K
 - 2: **Choose** step size $\gamma > 0$
 - 3: **for** each iteration $k = 0, 1, 2, \dots, K - 1$ **do**
 - 4: Compute $\mathbf{g}_s(\mathbf{b}(k), \boldsymbol{\mu}(k))$ for all $s \in \mathcal{M}$ (Equation (7))
 - 5: For each firm $s \in \mathcal{M}$, update: $\mathbf{b}_s(k+1) = \mathbf{b}_s(k) + \gamma \mathbf{g}_s(\mathbf{b}(k), \boldsymbol{\mu}(k))$
 - 6: **end for**
 - 7: Set $\mathbf{b}^* \leftarrow \mathbf{b}(K)$ $\triangleright$ Nash equilibrium
 - 8: **return** $\mathbf{b}^*$
-

We define a few matrices as follows. The gradient of $u_{s_1}(\mathbf{b})$ w.r.t. $\mathbf{b}_{s_1}$ is denoted as $\nabla_{\mathbf{b}_{s_1}} u_{s_1}(\mathbf{b})$ for an arbitrary $s_1 \in \mathcal{M}$. The second derivative of this expression w.r.t. $\mathbf{b}_{s_2}$ for any $s_2 \in \mathcal{M}$ is defined as the $n \times n$ matrix G_{s_1, s_2} where $[G_{s_1, s_2}]_{i,j} = \frac{\partial^2}{\partial b_{s_2, j} \partial b_{s_1, i}} u_{s_1}(\mathbf{b})$. The matrix $G = [G_{s_1, s_2}]_{s_1 \in \mathcal{M}, s_2 \in \mathcal{M}}$ is a block matrix of dimension $mn \times mn$ where every block consists of a G_{s_1, s_2} matrix.

3.1 Convergence of Best Response Dynamics

In this section, we present the conditions on the utility functions and analyze the convergence of the BRD in the ACG to a Nash equilibrium.

Assumption 1 (Utility Properties). The agents' utility functions $u_s(\mathbf{b}), s \in \mathcal{M}$ satisfy the following properties:

1. *Strong concavity*: The matrix $G(\mathbf{b})$ is λ -negative semidefinite, i.e.,

$$G(\mathbf{b}) + \lambda I_{mn \times mn} \preceq 0,$$

2. *Bounded operator norm*: The operator norm of G , defined as $\|G\|_{\text{op}} := \inf \{c > 0 : \|G\mathbf{v}\| \leq c\|\mathbf{v}\|, \forall \mathbf{v}\}$ is bounded, i.e., $\|G(\mathbf{b})\|_{\text{op}} \leq B$.

for some constants $\lambda, B > 0$.

With Assumption 1, we now prove that the BRD converges to a Nash equilibrium, formalized in the following theorem.

Theorem 1 (Convergence to a Nash Equilibrium). *Under Assumption 1, the best response dynamics (Algorithm 1) converges to a Nash Equilibrium $\mathbf{b}^*$. More precisely, for any given $\epsilon > 0$ and any initial point $(\mathbf{b}(0), \boldsymbol{\mu}(0))$ in Algorithm 1, the gradients for $K > \frac{2B^2}{\lambda^2} \log \left(\frac{\|\mathbf{g}(\mathbf{b}(0), \boldsymbol{\mu}(0))\|_2}{\epsilon} \right)$ iterations satisfy $\|(\mathbf{g}(\mathbf{b}(K)), \boldsymbol{\mu}(K))\|_2 < \epsilon$, provided that the step size is selected as $\gamma = \frac{\lambda}{B^2}$.*

Proof. The convergence of the budget update rule in line 5 of Algorithm 1 is analyzed through the L^2 norm of the update direction. Applying a first-order Taylor expansion to the gradient term yields

$$\mathbf{g}(\mathbf{b}(k+1), \boldsymbol{\mu}(k)) = \mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k)) + G(\mathbf{b}')(k)(\mathbf{b}(k+1) - \mathbf{b}(k)) \quad (8)$$

where $G(\mathbf{b}')$ is the Hessian-like matrix (as previously defined) evaluated at some intermediate point $\mathbf{b}'$ between $\mathbf{b}(k)$ and $\mathbf{b}(k+1)$. Substituting $\mathbf{b}(k+1) - \mathbf{b}(k)$, we get:

$$\mathbf{g}(\mathbf{b}(k+1), \boldsymbol{\mu}(k)) = \mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k)) + \gamma G(\mathbf{b}')(k) \mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))$$

Taking the squared norm of both sides and expanding, we have

$$\|\mathbf{g}(\mathbf{b}(k+1), \boldsymbol{\mu}(k))\|_2^2 = \|\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k)) + \gamma G(\mathbf{b}')(k) \mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2^2$$

$$\begin{aligned}
&= \|\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2^2 + 2\gamma \mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))^T G(\mathbf{b}') \mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k)) \\
&\quad + \gamma^2 \|G(\mathbf{b}') \mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2^2
\end{aligned} \tag{9}$$

Under the assumption of strong concavity of utility, we obtain the following inequality

$$\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))^T G(\mathbf{b}') \mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k)) \leq -\lambda \|\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2^2 \tag{10}$$

Additionally, assuming that the operator norm of the Hessian-like matrix is bounded, $\|G(\mathbf{b}')\|_{\text{op}} \leq B$ and using the standard property of operator norms $\|A\mathbf{v}\|_2 \leq \|A\|_{\text{op}} \|\mathbf{v}\|_2$, we obtain the following inequalities

$$\begin{aligned}
\|G(\mathbf{b}') \mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2 &\leq B \|\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2 \\
\|G(\mathbf{b}') \mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2^2 &\leq B^2 \|\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2^2
\end{aligned} \tag{11}$$

Substituting the bounds from (10) and (11) in (9):

$$\|\mathbf{g}(\mathbf{b}(k+1), \boldsymbol{\mu}(k))\|_2^2 \leq \|\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2^2 (1 - 2\gamma\lambda + \gamma^2 B^2) \tag{12}$$

For convergence, the coefficient of the right-hand side must be less than one, leading to the step-size condition

$$1 - 2\gamma\lambda + \gamma^2 B^2 < 1 \quad \Rightarrow \quad \gamma < \frac{2\lambda}{B^2}$$

Choosing the step size as $\gamma = \frac{\lambda}{B^2}$ and assuming the dual variables $\boldsymbol{\mu}(k)$ are updated such that

$$\|\mathbf{g}(\mathbf{b}(k+1), \boldsymbol{\mu}(k+1))\|_2 \leq \|\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2$$

Eq. (12) simplifies to

$$\|\mathbf{g}(\mathbf{b}(k+1), \boldsymbol{\mu}(k+1))\|_2^2 \leq \|\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2^2 \left(1 - \frac{\lambda^2}{B^2}\right)$$

Applying this recursively over k iterations yields

$$\|\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2^2 \leq \left(1 - \frac{\lambda^2}{B^2}\right)^k \|\mathbf{g}(\mathbf{b}(0), \boldsymbol{\mu}(0))\|_2^2$$

Under the inequality $(1-x)^k \leq e^{-xk}$ and taking the square roots, we obtain an exponential rate of convergence

$$\|\mathbf{g}(\mathbf{b}(k), \boldsymbol{\mu}(k))\|_2 \leq e^{-\frac{\lambda^2}{2B^2}k} \|\mathbf{g}(\mathbf{b}(0), \boldsymbol{\mu}(0))\|_2 \tag{13}$$

The number of iterations $k = K$ required to achieve a desired error tolerance $\epsilon > 0$ is

$$\begin{aligned}
e^{-\frac{\lambda^2}{2B^2}K} \|\mathbf{g}(\mathbf{b}(0), \boldsymbol{\mu}(0))\|_2 &= \epsilon \\
\Rightarrow K &= \frac{2B^2}{\lambda^2} \log \left(\frac{\|\mathbf{g}(\mathbf{b}(0), \boldsymbol{\mu}(0))\|_2}{\epsilon} \right)
\end{aligned} \tag{14}$$

Thus, the best response dynamics converges to a Nash Equilibrium in $K = \mathcal{O}(\log \frac{1}{\epsilon})$. $\square$

The result above ensures that the BRD will lead to a NE. However, there is a natural tension between the equilibrium and the welfare properties. In the following section, we show how this tension can be mitigated using an appropriate contest success function (CSF).

4 Welfare and Nash Equilibrium

In this section, we first present an example that shows the need of designing appropriate CSFs. We construct the following counterexample for three widely-used social welfare functions, namely, Nash, utilitarian, and egalitarian social welfare functions (see the paragraph below Definition 2). However, this example can be adapted for any $p \in (-\infty, 1]$.

Example 1 (A Counterexample). Consider two firms A and B competing over a fully connected network of five agents, with the asymmetric influence matrix $\mathcal{E}$ as shown in Table 1. The optimal budgets that maximize the p -mean social welfare (Definition 2) is denoted as $(\mathbf{b}_A^{\text{opt}}, \mathbf{b}_B^{\text{opt}}) = \arg\max_{\mathbf{b}} W_p(\mathbf{b})$.

Nodes	1	2	3	4	5
1	0	0.3901	0.3003	0.2456	0.0640
2	0.0669	0	0.3715	0.2578	0.3037
3	0.0149	0.7005	0	0.1534	0.1313
4	0.1407	0.2334	0.4025	0	0.2234
5	0.4340	0.0989	0.2072	0.2599	0

Table 1: Influence Matrix $\mathcal{E}$

We choose the following CSF as given in (Skaperdas, 1996) with the parameters $k > 0, \delta \in (0, 1]$:

$$h_{si}(\mathbf{b}) = \frac{e^{kb_{si}}}{\sum_{r \in \mathcal{M}} e^{kb_{ri}} + \delta}$$

The total budget for both the firms and the peer-to-peer communication factor (see Section 2.1) are $C = 1$ and $\alpha = 0.5$ respectively. The CSF parameters are chosen as $k = 1$ and $\delta = 0.5$. The welfares corresponding to utilitarian, egalitarian, and Nash social welfare functions are given by $W_{\text{avg}}, W_{\text{egal}},$ and W_{NSW} respectively.

The NE is given by $\mathbf{b}_A^* = (0.1496, 0.0, 0.5263, 0.1907, 0.1334), \mathbf{b}_B^* = (0.0337, 0.0, 0.9663, 0.0, 0.0)$. The utilities are given by $u_A(\mathbf{b}^*) = 1.0431, u_B(\mathbf{b}^*) = 1.1021$, and therefore $W_{\text{avg}}(\mathbf{b}^*) = 1.0431, W_{\text{egal}}(\mathbf{b}^*) = 1.0726, W_{\text{NSW}}(\mathbf{b}^*) = 1.0722$.

On the other hand, the welfare maximizing solutions for these three social welfare functions are given by the trivial budget vector $\mathbf{b}_A^{\text{opt}} = \mathbf{b}_B^{\text{opt}} = (0.0, 0.0, 0.0, 0.0, 0.0)$. The optimal social welfares are $u_A(\mathbf{b}^{\text{opt}}) = u_B(\mathbf{b}^{\text{opt}}) = 2.0$ and $W_{\text{avg}}(\mathbf{b}^{\text{opt}}) = W_{\text{egal}}(\mathbf{b}^{\text{opt}}) = W_{\text{NSW}}(\mathbf{b}^{\text{opt}}) = 2.0$. $\square$

We can observe that the optimal strategy of each firm is to invest nothing on the nodes and this gives a higher social welfare than their equilibrium strategies. It can be deduced that competition induces wasteful spending, while collective welfare is maximized when all firms refrain from investing. This inefficiency stems from the structure of the existing contest success function, which encourages firms to spend resources that do not contribute to higher total welfare. To overcome this, we next explore how an appropriately designed CSF can align equilibrium behavior with collective social welfare.

4.1 The Optimal Contest Success Function

We consider a planner who can design the CSF in our awareness competition game. The planner's goal is to choose a CSF that shapes firms' incentives so that their strategic behavior leads to a welfare-maximizing outcome. Here, we identify the properties a CSF must satisfy for the resulting Nash equilibrium to also maximize social welfare, and we provide an example of a functional form that meets these conditions.

Assumption 2. The CSF $h_{si}(\mathbf{b}) = h(b_{si}, \mathbf{b}_{-si})$ for each player $s \in \mathcal{M}$ on node $i \in \mathcal{N}$, satisfies the following properties:

1. (*Strict concavity*)

$$\frac{\partial^2 h_{si}}{\partial b_{si}^2} < 0 \quad \forall s \in \mathcal{M}. \quad (15)$$

2. (*Strategic substitutability*)

$$\frac{\partial^2 h_{si}}{\partial b_{si} \partial b_{ri}} < 0 \quad \forall s, r \in \mathcal{M}, r \neq s. \quad (16)$$

3. (*Dominance of diminishing returns*) For any pair of distinct players $s, r \in \mathcal{M}$ and for all $0 \leq t \leq 1 - \alpha, \alpha > 0$:

$$\left. \frac{\partial^2 h_{si}}{\partial b_{si}^2} \right|_{\substack{b_{si}=t+\alpha \\ b_{ri}=t, \forall r \neq s}} < \left. \frac{\partial^2 h_{si}}{\partial b_{si} \partial b_{ri}} \right|_{\substack{b_{si}=t \\ b_{ri}=t+\alpha \\ b_{ki}=t, \forall k \neq r, s}} < 0. \quad (17)$$

The above conditions in the assumption constitute the m -player extensions of the conditions introduced by Bimpikis et al. (2016) in *Proposition 5* to establish the existence and uniqueness of a symmetric Nash equilibrium in their two-player game. Intuitively, the *strict concavity* condition in (15) ensures that each player's CSF exhibits diminishing returns with respect to its own investment, implying that the marginal benefit of additional spending decreases as the player allocates more budget to a node. The *strategic substitutability* condition in (16) captures the negative externalities among competing players, meaning that an increase in one player's spending reduces the marginal effectiveness of another player's investment on the same node.

The *dominance of diminishing returns* condition in (17) contrasts these two opposing effects on firm s 's marginal success. The left-hand side captures the *diminishing returns* effect on a firm s when it increases its own spending from t to $t + \alpha$, while all other $m - 1$ firms maintain their spending at t . The right-hand side represents the *strategic substitutability* effect on firm s when a competitor r raises its spending to $t + \alpha$, while firm s and the remaining $m - 2$ firms keep their spending fixed at t . The inequality thus states that the adverse effect of a firm's own diminishing returns outweighs the negative strategic impact exerted by any single competitor.

Theorem 2. Under Assumption 2, the best response dynamics (Algorithm 1) converges to a unique and symmetric NE that maximizes the p -mean social welfare for all $p \in (-\infty, 1]$.

Proof. Bimpikis et al. (2016) (*Proposition 5*) prove that under assumption 2 (for 2 players), the Nash equilibrium is always unique and symmetric. The proof can be generalized straightforward to $m \geq 2$ players as detailed in Section B.

Since the equilibrium solution is guaranteed to be symmetric always, this means the budget vectors of each firm s are the same, i.e. $b_{si} = b_i, \forall s \in \mathcal{M}$. This implies that the outcome of the contest is same for all the players: $h_{si}(\mathbf{b}) = h_i(\mathbf{b}), \forall s \in \mathcal{M}$, and hence the utilities (3) for all the players are also the same.

$$u_s(\mathbf{b}) = \sum_{i=1}^n M_i h_i(\mathbf{b}) - \sum_{i=1}^n b_i = u(\mathbf{b})$$

Since NE is symmetric only, let's restrict the strategy space to be symmetric profiles only. This reduces the p -mean welfare expression (Definition 2) to be the same as the utility expression:

$$W_p(\mathbf{b}) = \left(\frac{1}{m} \sum_{s=1}^m u(\mathbf{b})^p \right)^{\frac{1}{p}} = \left(\frac{1}{m} \cdot m u(\mathbf{b})^p \right)^{\frac{1}{p}} = u(\mathbf{b}).$$

The strict concavity (15) of the CSF $h_{si}(\mathbf{b})$ in b_{si} implies that each firm's utility $u_s(\mathbf{b})$ is concave in its own decision vector $\mathbf{b}_s$, since it combines a concave total awareness term $\sum_{i=1}^n M_i h_{si}(\mathbf{b})$ and the negative of a convex (linear) cost term $\sum_{i=1}^n b_{si}$. The p -mean welfare $W_p(\mathbf{b})$, being a concave aggregation of the concave utilities $\{u_s(\mathbf{b})\}_{s \in \mathcal{M}}$, is therefore concave in the joint budget profile $\mathbf{b}$, and remains concave in $\mathbf{b}$ when restricted to the symmetric manifold ($b_{si} = b_i, \forall s$).

Since both the planner's welfare-maximization problem and each firm's utility-maximization problem are concave programs with convex feasible sets, the KKT conditions are both necessary and sufficient for optimality. Hence, solving the KKT systems yields the optimal welfare-maximizing allocation and the firms' Nash equilibrium strategies, respectively.

Consider the planner's constrained maximization game over the symmetric manifold:

$$\max_{\mathbf{b} \in \mathbb{R}_+^n, \sum_i b_i \leq C} W_p(\mathbf{b}),$$

where C is the total budget of each firm. Let $(\mathbf{b}, \hat{\lambda}, \hat{\mu})$ be a KKT triple for this planner problem. The KKT conditions, for each node $i \in \mathcal{N}$ are:

$$\frac{\partial W_p}{\partial b_i} - \hat{\lambda} + \hat{\mu}_i = 0, \quad \hat{\mu}_i \geq 0, \hat{\lambda} \geq 0, \quad \hat{\mu}_i b_i = 0, \hat{\lambda}(C - \sum_i b_i) = 0.$$

Now consider any single firm s solving its own constrained problem. Its KKT system is identical to the planner's:

$$\frac{\partial u_s}{\partial b_{si}} - \lambda_s + \mu_{si} = 0, \quad \mu_{si} \geq 0, \lambda_s \geq 0, \quad \mu_{si} b_{si} = 0, \lambda_s(C - \sum_i b_{si}) = 0.$$

On symmetric profiles, one has $\partial_{b_i} W_p(\mathbf{b}) = \partial_{b_{si}} u_s(\mathbf{b})$ for every $s \in \mathcal{M}, i \in \mathcal{N}$, since the derivative of W_p equals that of each firm's payoff evaluated on the symmetric manifold. Hence, setting $\lambda_s = \hat{\lambda}$ and $\mu_{si} = \hat{\mu}_i$ shows that $(\mathbf{b}, \lambda_s, \mu_{si})$ satisfies each firm's KKT conditions. Therefore, $\mathbf{b}$ is a best response for every firm, and consequently a symmetric Nash equilibrium of the planner's welfare-maximizing game. □

Since the Nash equilibrium obtained in Theorem 2 is unique and coincides with the welfare-maximizing allocation, the equilibrium outcome achieves the same social welfare as the optimal solution. Therefore, the ratio defining the Price of Anarchy equals one.

Corollary 1. *Under the conditions of Theorem 2, the Price of Anarchy (PoA) equals one, i.e., $PoA = 1$.*

The next proposition illustrates a class of contest success functions that meet the requirements of Assumption 2, through which Theorem 2 can be implemented.

Proposition 1. The contest success function h of the following form

$$h_{si}(\mathbf{b}) = \frac{f(b_{si})}{\sum_s f(b_{si}) + \delta} \tag{18}$$

with $f(\cdot)$ concave and increasing, $f(0) = 0$, and $\delta > 0$ satisfies the properties stated in Assumption 2, and thus ensures that the resulting unique and symmetric Nash equilibrium maximizes the p-mean social welfare.

Proof. We will prove that this general form of CSF satisfies the 3 conditions of the Assumption 2. Consider the case of 2 players with strategies x and y . The CSF can be written as:

$$h(x, y) = \frac{f(x)}{f(x) + f(y) + \delta}$$

Let $D(x, y) = f(x) + f(y) + \delta$. Assume that the $f(\cdot) > 0$ is increasing, i.e., $f'(\cdot) > 0$. Calculating the partial derivative:

$$\frac{\partial h(x, y)}{\partial x} = \frac{f'(x)(f(y) + \delta)}{(f(x) + f(y) + \delta)^2}$$

Calculating the second derivative:

$$h_{xx}(x, y) = \frac{\partial^2 h(x, y)}{\partial x^2} = \frac{(f(y) + \delta)[f''(x)D(x, y) - 2f'(x)^2]}{D(x, y)^3}$$

Since f is concave and increasing, $f'' < 0$ and $f' > 0$. Thus, $f''(x)D(x, y) - 2f'(x)^2 < 0$ and hence $h_{xx}(x, y) < 0$ which satisfies the strict concavity condition (15).

Calculating the mixed partial derivative:

$$h_{xy}(x, y) = \frac{\partial^2 h(x, y)}{\partial x \partial y} = -f'(x)f'(y) \frac{\delta + f(y) - f(x)}{D(x, y)^3}$$

The dominance of diminishing returns condition (17) for 2 players can be written as:

$$\frac{\partial^2 h(t + \alpha, t)}{\partial x^2} < \frac{\partial^2 h(t, t + \alpha)}{\partial x \partial y} < 0, \quad \forall 0 \leq x \leq 1 - \alpha \text{ and } \alpha > 0 \quad (19)$$

The right-hand inequality is the strategic substitutability condition. In the mixed partial derivative expression, substitute $x = t$, $y = t + \alpha$:

$$h_{xy}(t, t + \alpha) = -f'(t)f'(t + \alpha) \frac{\delta + f(t + \alpha) - f(t)}{D(t, t + \alpha)^3}$$

Since $\delta > 0$ and $f(t + \alpha) > f(t)$ (as f is increasing), $\delta + f(t + \alpha) - f(t) \geq 0$. Hence, $h_{xy}(x, y) < 0$ which satisfies the strategic substitutability condition (16).

The left-hand inequality in (19) also implies:

$$\frac{\partial h(t_1, t_2)}{\partial x} < \frac{\partial h(t_2, t_1)}{\partial x} \quad \text{for } 0 \leq t_2 < t_1 \leq 1$$

Now we'll check if the general concave Tullock function satisfies this simpler implied inequality.

$$\begin{aligned} \frac{\partial h(t_1, t_2)}{\partial x} &= \frac{f'(t_1)(f(t_2) + \delta)}{(f(t_1) + f(t_2) + \delta)^2} \\ \frac{\partial h(t_2, t_1)}{\partial x} &= \frac{f'(t_2)(f(t_1) + \delta)}{(f(t_1) + f(t_2) + \delta)^2} \end{aligned}$$

Since the denominators are identical and positive, this simplifies to checking if:

$$\frac{f'(t_1)}{f(t_1) + \delta} < \frac{f'(t_2)}{f(t_2) + \delta} \quad (20)$$

Let's define a new function $G(t) = \frac{f'(t)}{f(t) + \delta}$. The inequality holds if $G(t)$ is a strictly decreasing function, i.e., $G'(t) < 0$.

$$G'(t) = \frac{f''(t)(f(t) + \delta) - f'(t)^2}{(f(t) + \delta)^2}$$

Since f is concave and increasing, $f'' < 0$ and $f' > 0$. Thus, $G'(t) < 0$. Therefore for $t_2 < t_1$, $G(t_2) > G(t_1)$ which is the inequality to be proved in (20). Thus we can conclude that the general concave Tullock function satisfies (19).

We have shown that for the two-player case, all three conditions of Assumption 2 are satisfied. The proof naturally extends to the m -player case as well where $f(x) = f(b_{si})$ and $f(y)$ is replaced by $\sum_{r \neq s} f(b_{ri})$. The strict concavity (15) and the negative mixed partial derivative of the m -player CSF (16) follow directly due to similar derivative expressions. For the dominance of diminishing returns condition (17), the key idea is to verify that the inequality holds when comparing any two firms while keeping all others' expenditures fixed, which can be established analogously to the two-player case. $\square$

5 Experiments

The theoretical results established in the previous sections depend on certain assumptions. In this section, we conduct numerical experiments on a simulated social network to see how the algorithms perform even when such assumptions cannot be enforced. In particular, we aim to verify the convergence of the Best-Response Dynamics (BRD) to a unique and symmetric Nash Equilibrium (NE) that maximizes the p -mean social welfare, as stated in Theorem 2. We test this via the Price of Anarchy (PoA) which converges to one for large number of iterations, confirming that the welfare-maximizing Nash equilibrium is indeed attainable in practice.

5.1 Data and Network Generation

To implement our theoretical model, we first require a realistic social network that captures how individuals interact and influence one another. Since no such network or customer data is publicly available, both the network and underlying data must be generated synthetically to resemble real-world marketing environments.

5.1.1 Data Generation

We synthesize a population of customers with basic demographic attributes such as age, gender, and geographic location (represented as spatial grids). Using these demographics, we simulate an auxiliary social network that serves solely to generate realistic product adoption data in a principled way following the literature that generates the social network and their influence probabilities from users' demographic data. Connections between users are formed probabilistically via Bernoulli trials based on demographic similarity and spatial proximity, with the connection probability computed following the formulation in Guarino et al. (2021). To capture directional influence, each edge is assigned a random influence weight drawn from a Beta distribution.

Using this auxiliary network and the Independent Cascade (IC) diffusion model (Kempe et al., 2003), we simulate product adoption process for multiple firms' products. For each customer, we record whether and when a product was adopted, resulting in a dataset that contains both demographic details and synthetic product holding histories (adoption dates). This dataset is then treated as our synthetic *observed market data*.

5.1.2 Social Network Construction

From the synthetic dataset, using the same Bernoulli trials procedure as in the auxiliary network to determine edge existence, we reconstruct the main social influence network to be used in our equilibrium experiments. Here, the influence probabilities are not assigned randomly but are estimated using a Bernoulli model with *maximum likelihood estimation* as proposed in Goyal et al. (2010), where each interaction is treated as a Bernoulli trial. Specifically, the probability of node j influencing node i is estimated as $e_{ji} = \frac{A_{j \rightarrow i}}{A_j}$, which constitutes entries of the network adjacency matrix $\mathcal{E}$. Here $A_{j \rightarrow i}$ denotes the number of products propagated from j to i and A_j to be the total products adopted by j , capturing the actual influence dynamics in the network. The resulting directed network forms the adjacency matrix $\mathcal{E}$ used in the contest simulations.

Finally, we validate the reconstructed network by testing whether it can accurately predict future product adoptions from past diffusion data. The model demonstrates strong predictive performance, suggesting that the generated network meaningfully captures underlying influence dynamics and is suitable for evaluating Nash equilibrium convergence. The detailed procedures for data and network generation are provided in the supplementary material (Section C).

5.2 BRD Simulations

We begin by generating a synthetic dataset comprising 100 customers and 10 product holdings. From this dataset, we construct a social network with 100 nodes and estimate its adjacency matrix $\mathcal{E}$, which satisfies our initial assumption of being substochastic. We consider a market with three firms competing over this connected network. The probability of peer-to-peer communication α is set to 0.5. The CSF follows the general functional form given in Proposition 1 (Eq. 18). Specifically, we consider the Tullock (Tullock et al., 1980) CSF $f(b_{si}) = b_{si}^q$, which is concave for $q \leq 1$ and becomes non-concave for $q > 1$. In addition, we analyze two other commonly used functional forms: the concave logarithmic form $f(b_{si}) = \log(1 + b_{si})$ and the non-concave exponential form $f(b_{si}) = \exp(b_{si})$.

The BRD (Algorithm 1) is implemented on this network to analyze convergence to a NE under both concave and non-concave CSF conditions. We experiment with three types of budget initializations: (i) *uniform*, where all nodes start with equal budgets ($1/n$ for every firm), (ii) *random*, where initial budgets are assigned randomly, and (iii) *biased*, where a large share of the budget is allocated to a single random node and the remainder distributed uniformly across others. Since uniform is deterministic, we initialize that once, and the other two in equal proportions in our total number of simulations. For the BRD implementation, we use a step size of 0.0001 and a maximum of 5000 iterations. In the concave CSF setting (including the Tullock CSF with $q = 1$ and the logarithmic form), and in the non-concave setting (including the Tullock CSF with $q = 2$ and the exponential form), each simulation is repeated 50 times with different budget initializations. To assess the efficiency of the equilibria found by the BRD algorithm, we compute the welfare-maximizing budget matrix $\mathbf{b}^{\text{opt}}$ and track the evolution of the ratio of welfare at each iteration to the optimal welfare. Specifically, for iteration k , we define the welfare ratio as

$$R(k) = \frac{W_p(\mathbf{b}(k))}{W_p(\mathbf{b}^{\text{opt}})},$$

where $\mathbf{b}(k)$ is the budget vector at iteration k of BRD, and $\mathbf{b}^{\text{opt}}$ is the budget vector that maximizes the p -mean social welfare.




Figure 2: Convergence behavior of the welfare ratio $R(k)$ across all three types of p -mean welfare for (a) concave Tullock CSF ($q = 1$), (b) non-concave Tullock CSF ($q = 2$), (c) log CSF and (d) exp CSF.

5.3 Results

Figure 2 presents the evolution of the mean welfare ratio $R(k)$ over BRD iterations for three types of p -mean welfare objectives: (i) Nash social welfare (W_{NSW}), (ii) Average welfare (W_{avg}), and (iii) Egalitarian welfare (W_{egal}). In the concave CSF setting (Figures 2a and 2c), as the algorithm converges, $R(k)$ approaches 1 for all three welfare measures, indicating that the equilibrium achieves welfare close to the optimal and hence the $PoA = 1$. The shaded regions indicate variability, which is large initially and shrinks as BRD converges, reflecting reduced fluctuations in welfare outcomes over iterations.

In the non-concave CSF setting (Figures 2b and 2d), mean welfare ratio $R(k)$ for all three types of social welfare is substantially below the optimum value of 1. Moreover, the variance across iterations remains high, with fluctuations persisting even in later stages, indicating instability in the best-response dynamics. This inefficiency arises from the lack of concavity in the underlying utility functions, which prevents the BRD algorithm from converging to a stable equilibrium. Consequently, the equilibria obtained are suboptimal, reflecting both poor welfare performance and inconsistent convergence behavior.

6 Conclusion

In this paper, we developed a two-timescale framework to study how competing firms allocate advertising budgets across social networks to maximize brand awareness. We generalized existing two-firm results to an arbitrary number of firms and showed that, under standard conditions, the firms' best-response dynamics converge to a pure strategy Nash equilibrium, which may not coincide with the social optimum. To address this, we characterized a class of contest success functions that network regulators can design to ensure that the Nash equilib-

rium is unique and welfare-maximizing. Our experiments confirm these theoretical insights, demonstrating stable convergence to optimal welfare across different contest success functions on realistic synthetic networks, achieving the theoretical guarantee of $\text{PoA} = 1$.

This study thus offers both theoretical and practical guidance for designing fair and efficient competitive environments for brand promotion and information dissemination in social networks. Future directions include extending the current symmetric framework to asymmetric settings reflecting firms’ heterogeneous budgets and influence, and incorporating node-level incentives to better model user engagement and design socially beneficial diffusion policies.

References

- Noga Alon, Michal Feldman, Ariel D Procaccia, and Moshe Tennenholtz. A note on competitive diffusion through social networks. *Information Processing Letters*, 110(6):221–225, 2010. (Cited on page 2)
- Amirhossein Ansari, Masoud Dadgar, Ali Hamzeh, Jörg Schlötterer, and Michael Granitzer. Competitive influence maximization: integrating budget allocation and seed selection. *arXiv preprint arXiv:1912.12283*, 2019. (Cited on page 2)
- Prithu Banerjee, Wei Chen, and Laks VS Lakshmanan. Maximizing welfare in social networks under a utility driven influence diffusion model. In *Proceedings of the 2019 international conference on management of data*, pages 1078–1095, 2019. (Cited on page 3)
- Prithu Banerjee, Wei Chen, and Laks VS Lakshmanan. Maximizing social welfare in a competitive diffusion model. *arXiv preprint arXiv:2012.03354*, 2020. (Cited on page 3)
- Siddharth Barman, Umang Bhaskar, Anand Krishna, and Ranjani G Sundaram. Tight approximation algorithms for p-mean welfare under subadditive valuations. *arXiv preprint arXiv:2005.07370*, 2020. (Cited on page 5)
- Ruben Becker, Gianlorenzo D’angelo, Sajjad Ghobadi, and Hugo Gilbert. Fairness in influence maximization through randomization. *Journal of Artificial Intelligence Research*, 73:1251–1283, 2022. (Cited on page 3)
- Shishir Bharathi, David Kempe, and Mahyar Salek. Competitive influence maximization in social networks. In *Internet and Network Economics: Third International Workshop, WINE 2007, San Diego, CA, USA, December 12-14, 2007. Proceedings 3*, pages 306–311. Springer, 2007. (Cited on page 2)
- Kostas Bimpikis, Asuman Ozdaglar, and Ercan Yildiz. Competitive targeted advertising over networks. *Operations Research*, 64(3):705–720, 2016. (Cited on pages 2, 3, 4, 5, 10, 19, and 21)
- VS Borkar. *Stochastic Approximation: A Dynamical Systems Viewpoint*. Cambridge University Press, Cambridge, UK, 2008. (Cited on page 20)
- Shih-Fen Cheng, Daniel M. Reeves, Yevgeniy Vorobeychik, and Michael P. Wellman. Notes on equilibria in symmetric games. In *Proc. 6th Workshop Decision Theoretic Game Theoretic Agent*, pages 23 – 28, 2004. URL <https://api.semanticscholar.org/CorpusID:10284620>. (Cited on page 21)
- Arastoo Fazeli and Ali Jadbabaie. Game theoretic analysis of a strategic model of competitive contagion and product adoption in social networks. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 74–79. IEEE, 2012a. (Cited on page 2)

- Arastoo Fazeli and Ali Jadbabaie. Targeted marketing and seeding products with positive externality. In *2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 1111–1117. IEEE, 2012b. (Cited on page 2)
- Arastoo Fazeli, Amir Ajorlou, and Ali Jadbabaie. Competitive diffusion in social networks: Quality or seeding? *IEEE Transactions on Control of Network Systems*, 4(3):665–675, 2016. (Cited on page 2)
- Amit Goyal, Francesco Bonchi, and Laks VS Lakshmanan. Learning influence probabilities in social networks. In *Proceedings of the third ACM international conference on Web search and data mining*, pages 241–250, 2010. (Cited on pages 14, 24, and 25)
- Sanjeev Goyal and Michael Kearns. Competitive contagion in networks. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*, pages 759–774, 2012. (Cited on page 2)
- Stefano Guarino, Enrico Mastrostefano, Massimo Bernaschi, Alessandro Celestini, Marco Cianfriglia, Davide Torre, and Lena Rebecca Zastrow. Inferring urban social networks from publicly available data. *Future Internet*, 13(5):108, 2021. (Cited on pages 13 and 22)
- David Kempe, Jon Kleinberg, and Éva Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 137–146, 2003. (Cited on pages 13 and 23)
- Wei Lu, Francesco Bonchi, Amit Goyal, and Laks VS Lakshmanan. The bang for the buck: fair competitive viral marketing from the host perspective. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 928–936, 2013. (Cited on page 3)
- Takanori Maehara, Akihiro Yabe, and Ken-ichi Kawarabayashi. Budget allocation problem with multiple advertisers: A game theoretic view. pages 428–437, 2015. (Cited on page 2)
- Antonia Maria Masucci and Alonso Silva. Strategic resource allocation for competitive influence in social networks. In *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, pages 951–958. IEEE, 2014. (Cited on page 2)
- Antonia Maria Masucci and Alonso Silva. Advertising competitions in social networks. In *2017 American Control Conference (ACC)*, pages 4619–4624. IEEE, 2017. (Cited on page 2)
- John F Nash Jr. Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950. (Cited on page 6)
- J.R. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, UK, 1998. (Cited on page 20)
- Aida Rahmattalabi, Shahin Jabbari, Himabindu Lakkaraju, Phebe Vayanos, Max Izenberg, Ryan Brown, Eric Rice, and Milind Tambe. Fair influence maximization: A welfare optimization approach. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11630–11638, 2021. (Cited on page 3)
- Stergios Skaperdas. Contest success functions. *Economic theory*, 7:283–290, 1996. (Cited on pages 2 and 9)
- Gordon Tullock et al. Efficient rent seeking. *Toward a theory of the rent-seeking society*, 97:112, 1980. (Cited on page 14)

- Vasileios Tzoumas, Christos Amanatidis, and Evangelos Markakis. A game-theoretic analysis of a competitive diffusion process over social networks. In *International Workshop on Internet and Network Economics*, pages 1–14. Springer, 2012. (Cited on page 2)
- Vineeth S Varma, Irinel-Constantin Morărescu, Samson Lasaulce, and Samuel Martin. Marketing resource allocation in duopolies over social networks. *IEEE control systems letters*, 2(4):593–598, 2018. (Cited on page 2)
- Vineeth S Varma, Samson Lasaulce, Julien Mounthanyvong, and Irinel-Constantin Morărescu. Allocating marketing resources over social networks: A long-term analysis. *IEEE Control Systems Letters*, 3(4):1002–1007, 2019. (Cited on page 2)
- Pengzi Wang, Yiming Zhu, Kai Han, Zhizhuo Yin, Qing Xiu, and Pan Hui. Fairness-aware algorithms for seed allocation in social advertising. In *2022 IEEE 24th Int Conf on High Performance Computing & Communications; 8th Int Conf on Data Science & Systems; 20th Int Conf on Smart City; 8th Int Conf on Dependability in Sensor, Cloud & Big Data Systems & Application (HPCC/DSS/SmartCity/DependSys)*, pages 464–473. IEEE, 2022. (Cited on page 3)
- Lander Willem, Thang Van Hoang, Sebastian Funk, Pietro Coletti, Philippe Beutels, and Niel Hens. Socrates: an online tool leveraging a social contact data sharing initiative to assess mitigation strategies for covid-19. *BMC research notes*, 13(1):293, 2020. (Cited on page 22)
- Jinhang Zuo, Xutong Liu, Carlee Joe-Wong, John CS Lui, and Wei Chen. Online competitive influence maximization. In *International conference on artificial intelligence and statistics*, pages 11472–11502. PMLR, 2022. (Cited on page 2)

A Appendix to Section 2.1.2

A.1 Limiting Value of Awareness

Consider the awareness update equation (1):

$$x_{si}(t) = \frac{1}{t} \left[\alpha \sum_{j \in \mathcal{N}_i} e_{ji} x_{sj}(t-1) + (1-\alpha) h(b_{si}(k), b_{-si}(k)) \right] + \left(1 - \frac{1}{t}\right) x_{si}(t-1)$$

Let us define the vector $\mathbf{y}_s(t)$ of size $n+m$, where the first n entries correspond to the vector $\mathbf{x}_s(t)$ and for $[\mathbf{y}_s(t)]_{n+i} = 0$ if $s \neq i$ and $[\mathbf{y}_s(t)]_{n+s} = 1$ if $s = i$. Now consider the matrix W , whose element W_{ij} represents the rate of transitioning from state i to j :

$$W = \begin{bmatrix} -1 & \alpha e_{12} & \dots & \alpha e_{1n} & (1-\alpha)h_{11} & \dots & (1-\alpha)h_{m1} \\ \alpha e_{21} & -1 & \dots & \alpha e_{2n} & (1-\alpha)h_{12} & \dots & (1-\alpha)h_{m2} \\ \alpha e_{31} & \alpha e_{32} & \dots & \alpha e_{3n} & (1-\alpha)h_{13} & \dots & (1-\alpha)h_{m3} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \alpha e_{n1} & \alpha e_{n2} & \dots & -1 & (1-\alpha)h_{1n} & \dots & (1-\alpha)h_{mn} \\ 0 & 0 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 0 & 0 & \dots & 0 \end{bmatrix}$$

We also define the following diagonal matrix D :

$$[D(t-1)]_{ii} = \begin{cases} \frac{1}{t}, & \text{if } i \in \mathcal{N}, \\ 1, & \text{if } i \in \{n+1, \dots, n+m\}. \end{cases}$$

The expected value of the awareness vector $\mathbf{y}_s$, representing the awareness levels of firm s at time t given the history of information exchange up to time t , denoted by $\{m(l)\}_{l < t}$, can be expressed as

$$\mathbb{E}[\mathbf{y}_s(t) \mid \{m(l)\}_{l < t}] = (I_{n+m} + D(t-1)W) \mathbf{y}_s(t-1), \quad (21)$$

where I is the identity matrix, and W and D are the corresponding interaction and diagonal weighting matrices, respectively. Equation (21) follows directly from rewriting the awareness update equation (1) in matrix form. Since the expected awareness levels can be expressed analogously to the formulation in Bimpikis et al. (2016) for the proof of *Proposition 1*, we can employ a similar analytical approach to derive the limiting awareness levels for each firm. Thus, the evolution of the awareness vector $\mathbf{y}_s$ for firm s can be expressed from Eq. (21) as follows:

$$\mathbf{y}_s(t) = (I + D(t-1)W) \mathbf{y}_s(t-1) + (\text{random noise})$$

Define the noise term as

$$\begin{aligned} \eta(t) &= D^{-1}(t-1) [\mathbf{y}_s(t) - (I + D(t-1)W) \mathbf{y}_s(t-1)] \\ \implies \mathbf{y}_s(t) &= \mathbf{y}_s(t-1) + D(t-1) (W \mathbf{y}_s(t-1) + \eta(t)) \end{aligned} \quad (22)$$

Given that the random noise has zero mean and bounded variance, it can be derived $\forall t \geq 1$:

$$\mathbb{E}[\eta(t) \mid \mathcal{F}(t-1)] = 0 \quad \text{a.s.}, \quad (23)$$

$$\mathbb{E} [\|\eta(t)\|^2 \mid \mathcal{F}(t-1)] \leq \sup_{l < t} K(1 + \|\mathbf{y}(l)\|^2) \quad \text{a.s.} \quad (24)$$

where $\{\mathcal{F}(l)\}_{l \geq 0}$ is the family of σ -fields with $\mathcal{F}(l) = \sigma(\{m(l)\}_{l < t})$ and

$$K = \max(3\|W\|^2 + 2\|I\|^2, \|I\|^2)$$

The above equations and Borkar (2008) imply that Eq. (22) tracks a time-independent ordinary differential equation (ODE) represented as

$$\frac{dz_s(\hat{t})}{d\hat{t}} = Wz_s(\hat{t}) \quad \forall \hat{t} > 0 \quad (25)$$

where the initial state is $z_s(0) = \mathbf{y}_s(0)$ and the awareness levels converge to the limit of the above ODE, $z_{s,\text{lim}}$ almost surely. Hence,

$$\mathbf{y}_{s,\text{lim}} = \lim_{t \rightarrow \infty} \mathbf{y}_s(t) = z_{s,\text{lim}} = \lim_{\hat{t} \rightarrow \infty} e^{W\hat{t}} z_s(0)$$

We define the transition probability matrix V corresponding to a Markov chain with m absorbing states based on the matrix W , following the formulation in (Norris, 1998, Chapter 2). Each entry V_{ij} denotes the probability that, given a transition from state i , the next state will be j .

$$V = \begin{bmatrix} 0 & \frac{\alpha e_{12}}{W_1} & \cdots & \frac{\alpha e_{1n}}{W_1} & \frac{(1-\alpha)h_{11}}{W_1} & \cdots & \frac{(1-\alpha)h_{m1}}{W_1} \\ \frac{\alpha e_{21}}{W_2} & 0 & \cdots & \frac{\alpha e_{2n}}{W_2} & \frac{(1-\alpha)h_{12}}{W_2} & \cdots & \frac{(1-\alpha)h_{m2}}{W_2} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ \frac{\alpha e_{n1}}{W_n} & \frac{\alpha e_{n2}}{W_n} & \cdots & 0 & \frac{(1-\alpha)h_{1n}}{W_n} & \cdots & \frac{(1-\alpha)h_{mn}}{W_n} \\ 0 & 0 & \cdots & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 & 0 & \cdots & 1 \end{bmatrix}$$

where $W_i = \sum_{j \neq i} W_{ij}$ is the total rate of leaving state i which sums up to 1. It is a normalization factor that turns the rates of interaction from W matrix into probabilities.

Since Eq. (25) also defines a continuous time Markov chain, using results from Norris (1998), Chapter 3:

$$\lim_{\hat{t} \rightarrow \infty} e^{W\hat{t}} = \lim_{t \rightarrow \infty} V^t$$

Therefore, we arrive at the following *unique* limiting value of the awareness level vector for each firm s

$$\mathbf{y}_{s,\text{lim}} = \lim_{t \rightarrow \infty} \mathbf{y}_s(t) = \lim_{\hat{t} \rightarrow \infty} e^{W\hat{t}} \mathbf{y}_s(0) = \lim_{t \rightarrow \infty} V^t \mathbf{y}_s(0) \quad (26)$$

Representing the matrix V into four block matrices as follows:

$$\begin{bmatrix} \alpha \mathcal{E} & E \\ 0 & I_m \end{bmatrix}$$

where $\alpha \mathcal{E}$ denotes the $n \times n$ fractional adjacency matrix, 0 represents an $m \times n$ zero matrix, I_m is the $m \times m$ identity matrix, and E is the $n \times m$ matrix containing the $(1 - \alpha)$ -weighted contest success function values for each node-firm pair (as it can be visualized from the elements of V matrix).

$$V^t = \begin{bmatrix} \alpha^t \mathcal{E}^t & \sum_{l=0}^t (\alpha^{\mathcal{E}})^l E \\ 0 & I_m \end{bmatrix}$$

$$\lim_{t \rightarrow \infty} V^t = \begin{bmatrix} 0 & (I - \alpha^{\mathcal{E}})^{-1} E \\ 0 & I_m \end{bmatrix}$$

Here, $\lim_{t \rightarrow \infty} (\alpha^{\mathcal{E}})^t \rightarrow 0$ as $\rho(\alpha^{\mathcal{E}}) < 1$ which ensures matrix powers decay to zero and $\lim_{t \rightarrow \infty} \sum_{l=0}^t (\alpha^{\mathcal{E}})^l E$ is an infinite Neumann series, hence equals $(I - \alpha^{\mathcal{E}})^{-1} E$. For firm s , $[\mathbf{y}_s]_{n+s} = 1$ and $[\mathbf{y}_s]_{n+i} = 0$ for $i \neq s$. Hence, in (26),

$$[\mathbf{y}_{s,\text{lim}}]_{1:n} = (I - \alpha^{\mathcal{E}})^{-1} [E]_s$$

where $[E]_s$ is the s^{th} column of the E matrix or the $(1 - \alpha)$ -weighted contest success function values for firm s . Thus, at the end of each time period T , the awareness levels of a firm s converge to a limiting value $\mathbf{x}_s(k) = [\mathbf{y}_{\text{lim}}^s]_{1:n}$ as defined in (2):

$$\mathbf{x}_s(k) = (I - \alpha^{\mathcal{E}})^{-1} (1 - \alpha) \mathbf{h}_s(\mathbf{b})$$

B Appendix to Section 4.1

We restate and extend *Proposition 5* of Bimpikis et al. (2016), which establishes that under Assumption 2, the Nash equilibrium in the two-player setting is unique and symmetric. We provide a formal argument demonstrating that the uniqueness and symmetry of the equilibrium continue to hold for any number of players $m \geq 2$.

Following Cheng et al. (2004), the game admits a pure strategy symmetric equilibrium because it is a symmetric game with compact and convex strategy spaces and continuous and quasi-concave utility functions. Consequently, each player's best-response correspondence is single-valued, and the first-order optimality conditions cannot be simultaneously satisfied by two distinct symmetric profiles, ensuring uniqueness of the equilibrium.

Let us suppose that an asymmetric equilibrium exists. Then there exist 2 distinct firms $p \in \mathcal{M}$ and $q \in \mathcal{M}$ such that $\mathbf{b}_p \neq \mathbf{b}_q$. Because each firm has the same budget, there exist nodes i and j such that $b_{pi} > b_{qi}$ and $b_{qj} > b_{pj}$. For each firm, the KKT conditions require equality of marginal utilities across all active agents:

$$c_i \frac{\partial h(b_{pi}, \mathbf{b}_{-pi})}{\partial x} = c_j \frac{\partial h(b_{pj}, \mathbf{b}_{-pj})}{\partial x}$$

But because h is strictly concave and $\frac{\partial^2 h(x,y)}{\partial x \partial y} < 0$, condition (17) of Assumption 2 cannot hold simultaneously for both p and q when allocations are unequal across agents. This leads to a contradiction, ruling out asymmetric allocations.

C Appendix to Section 5.1

C.1 Data Generation

The data generation pipeline in Figure 3 describes the data generation procedure as summarized in Section 5.1. Here we shall provide all the details regarding the procedure.

C.1.1 Customer Demographics Data

Customer demographics are generated to reflect realistic diversity in age, gender, and location. Ages (18–65+) and genders follow real population statistics for India, while customer locations are created using a grid-based system, where the entire region is divided into small square tiles. Each tile represents a geographic area and, on average, contains four customers.




Figure 3: Phase 1: Synthetic data generation pipeline.

The tiles near the center have higher population density, meaning more customers are placed there, while those toward the edges have fewer customers. The distance between tiles is computed using standard distance formulas, giving a realistic sense of how far customers are from one another geographically.

C.1.2 Social Network Generation

The social network is generated by combining demographic and geographic data to create realistic interactions. Age group interactions are modeled using an interaction matrix, and social contact values between age groups are incorporated from the SOCRATES tool (Willem et al., 2020), which provides empirical estimates of contact frequencies between distinct age demographics. Spatial proximity is incorporated based on tile locations. Interaction probabilities are then computed, adding Gaussian noise for variability, and edges are formed probabilistically through *Bernoulli trials*. Specifically, the probability that an edge (u, v) exists in the graph G is determined by the following equation (Guarino et al., 2021):

$$\Pr[u, v] = \frac{\mu \cdot N}{2} \cdot \frac{m_{g_u, g_v} \cdot s_{g_u, g_v}}{\sum_{i \leq j} (m_{i, j} \cdot s_{i, j})} \cdot \frac{D(u, v)}{\sum_{u' \in V_{g_u}, v' \in V_{g_v}} (D(u', v'))} \quad (27)$$

where:

- μ : Average degree of graph G ,
- N : Number of users,
- g_u : The age label of node u ,
- $D(u, v)$: Euclidean distance between u and v ,
- V_i : The set of nodes having age label i ,
- $m_{i, j}$: Count of individuals between age groups i and j ,
- $s_{i, j}$: Frequency of social ties between age groups i and j , derived from a social contact matrix.

Figure 4 depicts a social network generated for 100 users with 25 tiles. The colours of the nodes depict the various age groups. As it can be seen, the average number of users in each tile is more or less 4.




Figure 4: Social network generated for 100 users (25 tiles).

C.1.3 Influence Probabilities Initialization

In initializing influence probabilities within the social network, we use the *Beta distribution*, which flexibly models probabilities in the finite range $[0, 1]$. It is defined as

$$f(x; \alpha, \beta) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha, \beta)}, \quad (28)$$

where α and β are the shape parameters and $B(\alpha, \beta)$ is the beta function. Random initialization of these probabilities introduces variability that reflects the heterogeneity of real-world influence dynamics and serves as the basis for generating product holdings in subsequent stages.

C.1.4 Product Holdings Generation

We consider a market with m competing firms, each introducing a total of 10 products simultaneously on a specific date: January 1st, 2020. The adoption status of each user for these products is recorded, along with the date of adoption where applicable. Initially, a random subset of users is selected to adopt certain products through Bernoulli trials, thereby initializing the adoption states on the start date. Subsequently, product (or awareness) propagation occurs via the *Independent Cascade (IC) diffusion model* (Kempe et al., 2003), a probabilistic diffusion model in which, at each time step, an active node has a single chance to activate each of its inactive neighbors with a given probability. It is implemented as follows: for a user u who has adopted product P_i , $i = 1, \dots, 10$, and each of its neighbors v in the network, a random number $r \in [0, 1]$ is drawn to determine whether v becomes aware of or adopts the product. If $r \leq p_{uv}$, where p_{uv} denotes the influence probability from user u to v , then v adopts product P_i . The adoption (or activation) time of v is then sampled to occur after the activation time of u . This procedure is repeated iteratively for all products $P_1, \dots, P_{10}$. Through this process, the diffusion of product awareness unfolds dynamically across the network, capturing realistic influence and adoption patterns that emerge from both initial random adoptions and subsequent cascade effects.

C.2 Network Generation

The network generation process illustrated in Figure 5 outlines the social network construction described in Section 5.1. In the following, we provide a detailed explanation of this procedure.




Figure 5: Phase 2: Social network generation pipeline.

C.2.1 Influence Probabilities Estimation

In this section, we describe the methodology used to estimate the influence probabilities within the social network generated in Section C.1.2, which collectively form the influence matrix $\mathcal{E}$. Our approach builds on Goyal et al. (2010), who estimate influence probabilities from action logs, or the past history of product adoptions. The method relies on some assumptions for simplifying the estimation process and ensuring practical applicability. The assumptions are:

1. *Independence of Influence:* Probability of various friends influencing a node i are independent of each other.
2. *Static Influence Probabilities:* The influence probabilities are static and do not change with time.
3. *Bernoulli Distribution:* The influence propagation between users is modeled using a binary outcome, following a *static model*.

The Bernoulli distribution is utilized to model the binary outcome of influence propagation between users. Under this framework, each interaction between a user and their neighbor is treated as a Bernoulli trial, where success denotes influence and failure denotes no influence. The element e_{ji} of the network adjacency matrix $\mathcal{E}$ which denotes the influence probability of node j on node i , is estimated using the Maximum Likelihood Estimator (MLE), calculated as the ratio of successful attempts to total attempts.

$$e_{ji} = \frac{A_{j \rightarrow i}}{A_j} \quad (29)$$

where $A_{j \rightarrow i}$ denotes the number of products propagated from node j to i and A_j is the total products adopted by j . This estimation method ensures that the influence probabilities reflect the actual propagation dynamics observed within the network, allowing for accurate modeling of influence dynamics and subsequent analysis.

C.2.2 Product Holdings Estimation

In this section, we employ the *General Threshold model* (Goyal et al., 2010) to estimate product holdings within the network. The model uses threshold activation, where a node becomes activated if the cumulative influence from its activated neighbors exceeds a node-specific threshold. For each node i , a threshold θ_i represents the cumulative influence required for activation.

The *joint influence probability*, denoted as $p_i(\mathcal{A}_i)$, captures the likelihood of node i becoming activated given a set $\mathcal{A}_i \subseteq \mathcal{N}_i$ of activated neighbors. Assuming independence of influences from different neighbors, this probability is

$$p_i(\mathcal{A}_i) = 1 - \prod_{j \in \mathcal{A}_i} (1 - e_{ji}), \quad (30)$$

where e_{ji} is the influence probability from node j to node i . Node i activates if $p_i(\mathcal{A}_i) \geq \theta_i$. Using these principles, we estimate product holdings by considering node-specific thresholds and joint influence probabilities derived from the network structure.

C.3 Network Validation Experiments

Table 2: Network Accuracy Results on Synthetic Data

Beta Distribution Parameters			Accuracy (%)		
Network	α	β	n=2000	n=3000	n=5000
Sparse	0.5	2	95.05 $\pm$ 0.11	95.91 $\pm$ 0.08	95.80 $\pm$ 0.07
Medium	1	1	95.38 $\pm$ 0.11	95.67 $\pm$ 0.09	95.85 $\pm$ 0.07
Medium	3	3	96.03 $\pm$ 0.08	96.04 $\pm$ 0.09	95.87 $\pm$ 0.07
Dense	2	0.5	95.73 $\pm$ 0.11	95.78 $\pm$ 0.09	95.61 $\pm$ 0.07

C.3.1 Experimental Setup

The experimental setup aims to validate the effectiveness of the network generation technique. Initially, customer demographic data is generated for varying population sizes, ranging from 100 to 10,000 individuals. The number of tiles in the grid-based region is determined as the square root of the number of agents, keeping the average number of users per tile as 4. This demographic data serves as the basis for generating the social network, utilizing a probabilistic approach, similar to the one described in Section C.1.2. To account for stochasticity, the network generation process is repeated 100 times. The influence probabilities are initialized as described in Section C.1.3 using the *Beta distribution*. Different datasets are created by varying the density of the networks during the data generation pipeline. This is done by tweaking the beta distribution parameters α and β , producing networks of varying densities, such as sparse, medium, and dense.

Following network generation, the final dataset from the data generation module is truncated at January 1st, 2022. The training set is then constructed using data only up to this date. Influence probabilities are estimated for the training set utilizing the *Bernoulli model* as described in Section C.2.1. Subsequently, product holdings are estimated for the years from 2022 to 2024 from the directed social graphs obtained earlier using the *General threshold Model* as described in Section C.2.2. The test data comprises the complete product holdings from 2020 to the present, allowing for comparisons between predicted and actual product holdings after 2022. Through these experiments, we aim to validate the robustness and efficacy

of the proposed network generation approach in accurately capturing real-world influence dynamics within the social network.

C.3.2 Results

Table 2 presents the mean accuracy and standard deviation of predicted product holdings across different network densities (sparse to dense) and population sizes ($n = 2000$, $n = 3000$, and $n = 5000$). The results show that predicted values closely align with observed ones, with accuracies consistently between 95%–96%. Although smaller networks ($n = 2000$) exhibit slightly higher variability, larger populations yield more stable and consistent results, demonstrating the robustness and reliability of the proposed network generation method.