Niche-Seeking in Influence Maximization with Adversary

Long-Foong Liow
School of Information Systems
Singapore Management University
lfliow.2009@msis.smu.edu.sg

Shih-Fen Cheng
School of Information Systems
Singapore Management University
sfcheng@smu.edu.sg

Hoong Chuin Lau
School of Information Systems
Singapore Management University
hclau@smu.edu.sg

ABSTRACT

In hotly contested product categories dominated by a few powerful firms, it is quite common for weaker or late entrants to focus only on particular segments of the whole market. The rationale for such strategy is intuitive: to avoid direct confrontation with heavy-weight firms, and to concentrate in segments where these weaker firms have comparative advantages. In marketing, this is what people called "go niche or go home". The niche-building strategy may rely on "homophily", which implies that consumers in a particular market segment might possess certain set of attributes that cause them to appreciate certain products better (in other words, weaker firms would customize their products to target some particular market segments and not the mass market). On the other hand, the niche-building strategy may also rely on the network effect, which implies that consumers having social relationship would reinforce each other via their respective adoptions. In this case, weaker firms should recognize such inter-customer network and concentrate only on customers belonging to certain set of strategic clusters. In this paper, we present the model for building effective niche-seeking strategies. For simplicity, we assume that the adoption choice depends only on the network effects (in other words, a customer will choose the product that is chosen by the majority of her neighbor). The social network is directed, and there will be two firms, one with significantly more marketing budget than the other firm. Firms take turns making investment choices on which customer to convert. For both firms, their budgets are fixed over time and unused budget will not carry over to future time periods. With this model, we manage to show that a simple strategy based on the evaluation of individual customer's "value" can effectively identify and secure niches within randomly generated scale-free networks. We also show that such niche-building strategy indeed performs better in the long run than a myopic strategy that only cares about immediate market gains.

International Conference on Electronic Commerce '12, August 6-8, 2012, Singapore Management University, Singapore.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Theory, Social Network, Marketing

Keywords

influence network, duopoly influencer marketing

1. INTRODUCTION

Social systems have been shown to be an important factor in affecting consumption behavioral patterns since the 1960s. The seminal work by Bass [3] marks the dawn of an era where researchers begin to explore the significance of networks in explaining or predicting product adoptions or innovation diffusions. The Bass model is closely related to the work on network externality in economics (e.g., see [4]) in that it adopts a macroscopic view, investigating adoptions or diffusions at the industry level. In the Bass model, the impact of networks is aggregated as the count on *previous adopters*, and the future adoption is then a function of this aggregated count. The simple and elegant Bass model was later expanded to model adoptions of products with *successive generations* (e.g., high-tech products like DRAM or consumer electronics) [12] and diffusion process with decision variables (e.g., price) [13] as well.

With the prevalence of technologies and devices that can accurately capture the digital traces of an individual (e.g., smartphones or social network sites like Facebook) recently, it becomes increasingly plausible to investigate adoption or diffusion processes at microscopic level. Now researchers are able to investigate and infer the micro-structures that are behind these macro-outcomes instead of fitting the observed statistics at macro-levels. Such micro-structural insights can be utilized to explicitly describe ripple effects of adoption or diffusion among interconnected individuals, and this lead to the intensive study of cascading phenomenon. In particular, researchers are studying how one could maximize the influence/ diffusion/ adoption in a given network through a targeted set of individuals which is well defined as the *influence maximization problem* by [9]. More concretely, given a directed weighted graph, $G = (V, E, W)$ with vertices V as users, edges E as relationships with weight function W: E→[0,1] which denotes the influence probabilities, the goal is to select a subset $S \subseteq V$ for initiating the diffusion process so as to maximize $\sigma(S)$, the number of vertices influenced by S at the end of diffusion process. The dynamics of influence propagation can be represented by one of many existing models, such as the linear threshold model, general threshold model and utility based model. Most of the propagation models in literature assume progressive activation in which an activated node cannot revert back to inactive mode. This assumption implies that a consumer is unable to change his choice after an initial purchase, which is rarely practical in the context of business marketing.

In this paper, we extend the classical influence maximization problem by introducing an adversary to the model and relaxing the assumption on progressive activation. Our research is motivated by an emerging e-Commerce practice known as *influencer marketing*, which channels marketing investment to specific key individuals, known as *influencers*, instead of the mass market. The main idea is to generate substantial awareness and subsequently possible sales from potential buyers who are strongly connected to these influencers. The *influencers* serve as conduit to the entire buyer segment, and are perceived as individuals who shape the purchasing decisions of true potential buyers. Since all marketers are aware of such phenomena and may deploy similar practice to compete in the same market, it might eventually lead marketers to engage the same group of key individuals in a repeated manner. It is therefore essential from each marketer's perspective to design their marketing strategies (to perform influence maximization) strategically under the presence of adversary.

In this paper, we narrow our study to a market where two players (the *incumbent/adversary* and the *entrant*) compete for their respective market shares on a single product. Players are to take turns making investment decisions (which particular customer to convert) in their respective decision epochs. Each player is endowed with a fixed (yet different) marketing budget at the beginning of each epoch (unused budget will not carry over), and the decision process will go on forever. Of particular interest in this paper is the assumption that the adversary is endowed with a higher budget. Due to the fact that the competition goes on forever, it's important for players to properly account for their opponent's future moves. This is where our model departs from the classical influence maximization problem. An interesting question we would like to answer is whether such consideration would cause the weaker player (the entrant) to adopt a niche-seeking strategy.

We model our problem as a two-player influence maximization problem with infinite horizon. Both the incumbent and entrant are assumed to have the information on the budget of their opponent. Each customer's decision is governed by her utility, which comprises network influence effect and monetary incentives she receives (if any). Furthermore, we introduce a *contractual lock-in* constraint which prevents a customer from changing product choice too often (she cannot change her product choice for a fixed period after she adopts a new choice). We propose a minimax algorithm that allows players to reason strategically on which customer to invest in at each time period. The above process will continue until some form of steady state is reached. Instead of seeking immediate gain in market share (what classical influence maximizer would do), we define a simple value function to measure the *long-term value* of each customer, and ask entrant to make decision based on this value. We illustrate empirically that, under certain conditions, our proposed *value-based* approach is niche-seeking, which can indeed secure a larger pool of customers as opposed to conventional count-based (myopic) approach, when competing against a stronger opponent.

2. BACKGROUND AND RELATED WORK

2.1 Background

The rules of game, players' role, their decision model and information sets are define as follow. Our game model is comprised of non zero-sum, pure strategies repeated sequential two-players competition game with network influence. The objectives of both players are to maximize the number of consumer adopting their respective product/ service within a network. The structure of consumer markets must conform to *scale-free* network phenomenon which is conjectured to be good representation for most social network structure. Both players have complete information of the ply depth used but incomplete information on the type of node selection policy used by each other. The *incumbent* player perceived his opponent to deploy a simplistic count-based approach. The role of *entrant* player is to decide which node selection policy to deploy for forming the list of consumers to influence in sequential manner. The game is initialized with consumers adopting neither product. Then the *entrant* will take the lead play to select a consumer (from a pool of eligible consumers) to instigate persuasive incentive based on his decision model, and the *incumbent* will perform the same subsequently after observing the new game state. The game will repeat in this manner until either one of the following termination criteria is met. The termination criteria are, 1) repetition of game state and 2) four hundred game stage is played with no repetition of game state observed. Both players are forbade from targeting the same consumer within a specified game stages. We termed this constraint as *contractual lock-in* in this paper.

2.2 Related Work

Network influences are observable dispersion effects which normally originated from a small local group to larger interconnected structure through various means of diffusion. A phenomena which resulted from network behavioral interaction. Research in this area can be broadly classified into two categories as the influence propagation models, and algorithms to estimate the propagation effects. Here we will cover only the former which is more relevant to our research.

Granovetter and Schelling [6] were among the first to propose models that capture the progressive change of nodes in networks. The concept of *linear threshold* model proposed by [6] was based on node-specific thresholds. An *inactive* node v at time *t-1* will become active (and remain active) at time t under the following condition:

$$\sum_{u \in N(v)} w_{v,u} X_{u,t-1} \geq Threshold(v) \quad (1)$$

The variable $X_{u,t-1}$ is 1 if u was active at time *t-1* and *0* otherwise. The variable $w_{v,u}$ denotes the degree of influence on v by u (the level of which u being active will contribute to v being active). Intuitively, if a predetermined fraction (*threshold*) of v is less than the sum of degree of influence from its *active* neighbors. This model is also known to be the foundation for a large body of work in the Sociology domain. Subsequently the Granovetter's Linear Threshold model is generalized as *general threshold* [10] model as it can be deduced into a re-parameterization form of node v with monotone *activation function* $f_v : 2^v \rightarrow [0,1]$, and activation threshold θ_v which is chosen independently, uniformly and randomly from the interval of (0,1]. A node v will become *active* at time *t+1* if and only if $f_v(S) \geq \theta_v$, where S denotes the set of active nodes at time t. This model is different from the Linear Threshold model as it focuses on the *cumulative influence* of all nodes from a set S instead of the individual attempts of nodes $u \in S$.

Most relevant to our problem context was the *utility function-based* propagation model commonly used to evaluate consumer product selection. Janssen and Jager [7] proposed this model which incorporated cognitive behavioral theories, to study the consumer purchasing decision from a psychological perspective. The authors concluded that the behavioral processes which drive the consumers decision are mainly based upon their needs such as low prices, high social comparison, and type of cognitive processing that the consumers utilize. [5] generalized the notion to formulate an utility model to represent the consumer personal pleasure in consuming a product. The value in the utility model is derived based on the consumer's perception of a product quality, and their tendency to follow the trends within a localized community. Here the decision making process of consumers is assumed to be adversarial and responsive to network externalities, so a product with the highest utility value will be selected. A general utility model proposed by [5] is,

$$U_i = (1 - ft)(Q_i - Q_{des}) + ft(N_i/N_p). \tag{2}$$

U_i denotes the amount of pleasure a consumer derives from product i. The follower tendency of a consumer, ft has a value in the range of [0,1]. Q_i is the consumer's experienced quality of product i and has a range value of [0,1]. Here the consumers are assumed to have a minimum quality requirement, Q_{des} of 0.5 to be satisfied. The number of consumers who select product i and the total number of consumers are denoted as N_i and N_p respectively.

3. THE MODEL

3.1 Influence Propagation with Investment

We assume that there are two players competing for their respective market shares in a social network modeled as an acyclic directed graph. The nodes in the directed graph represent individuals and links represent relationships among individuals (since the graph is directed, the impacts are directional). The two players are to take turns making decisions with infinite decision horizon. At each decision epoch, the player will be endowed a fixed player-specific budget and could spend this budget on one of the nodes. As stated earlier, unspent budget cannot be carried over.

To focus on the network effect, we assume that for each node (customer), her decision depends solely on two factors: 1) the influence from her neighbors, and 2) the direct investment from any player. Expressed formally, customer n's utility value for the product owned by player i in time period t is:

$$u_i^t(n) = \sum_{m:a_{m,n}=1} \frac{\mathrm{I}_{\{c^{t-1}(m)=i\}}}{|\{m' : a_{m',n} = 1\}|} + M_i^t(n), \tag{3}$$

where $c^{t-1}(m)$ represents the choice of node m in time period $t-1$, $a_{m,n}$ denotes the linkage from m to n (1 if such link exists, 0 otherwise), and this term evaluates the network influence effects (influence from all neighbors which link to him) on node n. $M_i^t(n)$ represents the investment by player i on node n in time period t. With the above utility function definition, node n's choice in time period t is then simply:

$$c^t(n) = \begin{cases} \arg\max_i u_i^t(n), & T^t(n) \geq \tau; \\ c^{t-1}(n), & \text{otherwise.} \end{cases} \tag{4}$$

Note that in (4), customer n is only allowed to change her decision if she has maintained a particular choice for more than τ time periods ($T^t(n)$ is the number of time periods customer n has maintained its current choice). This particular feature is to emulate the minimum length of contract one has to sign on when a new product is chosen. This design also helps to eliminate simple cycles among players (players keep selecting the most crucial node).

When a new investment is made, or the time period has progressed (thus changing the value of $T^t(n)$), some nodes might end up with new product choices, and these changes will create ripple effects that need to be properly accounted for. Considering the potential interactions among connected nodes, we have to propagate these updates using proper order. The procedure is described as follows:

1. Let $\mathcal{S}$ be the set containing all nodes.
2. Let $\mathcal{R} \equiv \{n | n \in \mathcal{S},\ a_{m,n} = 0, \forall m \in \mathcal{S}\}$ (in other words, $\mathcal{R}$ is the set of *root nodes* in $\mathcal{S}$). Let $\mathcal{S} \leftarrow \mathcal{S} \setminus \mathcal{R}$.
3. For $n \in \mathcal{R}$, compute $u_i^t(n)$ for $i = 1$ and 2 following (3). $c^t(n)$ can then be found from (4).
4. If $\mathcal{S} = \emptyset$, stop, otherwise, go to step 2.

The above procedure always terminates if the graph is acyclic.

3.2 The Node-Selection Problem

With the above influence propagation model, players are allowed to take turns making decisions on which node to invest in. Given the complexity of the influence propagation model described in 3.1, even with perfect information on $\{c^t(n)\}$ and $\{T^t(n)\}$, a player has to rely on pure enumeration to find the *best* node to invest in. Note that the above problem is only with one time period and not considering adversary. Although it might be possible to formulate player's decision making problem with infinite horizon and adversary, it will be computationally intractable. As such, when we design player's strategy, we explicitly define number of future time periods to be included in the evaluation, and treat that as a strategy parameter (we call it the *look-ahead time periods*).

Although we can make single-player's strategy tractable by setting a small-enough look-ahead time periods, such limitation would create some unexpected issue in how we evaluate the importance of each node. In most influence maximization problems, the importance of a node can be characterized by the number of converts it can bring in through influence propagation. When adversaries are present and horizon is infinite, we can still estimate the importance of a node by using average or discounted measure (commonly used techniques in infinite horizon decision making problems) and having appropriate opponent model. Unfortunately, if we artificially limit the number of periods that we look ahead, these classical approach will not work anymore. To see this, assume the look-ahead period is just 1, implying that this player would be myopic. In this case, the strategy is essentially an influence maximizer that simply chooses the node that would result in maximum immediate gain; however, as one would expect, given that there is an adversary, such gain might be short-lived, and the choice might turn out to be a *short-term gain, long-term loss.*

One way to deal with such undesirable side effect of limiting planning horizon is to properly define a *value function* that would approximate a node's true value suppose we are able to reason with infinite horizon. In our initial study, we defined one such estimation function, and to distinguish it from the conventional ways of estimating a node's value, we call the strategy that relies on conventional measure the *count-based* approach, and the strategy that relies on the value function the *value-based* approach.

Formally speaking, the count-based approach is the strategy where player i use the following function to evaluate the total value for all nodes under his control in time period t:

$$v_i^t = \sum_n \mathrm{I}_{\{c^t(n)=i\}}. \tag{5}$$

With (5), the importance value of a particular node m not owned by i is simply $v_i^t(c^t(m) \text{ set to } 1) - v_i^t$.

On the other hand, the value-based approach can take many different forms. In our study, we use a simple function that focuses not just on the quantity of nodes under his control, but also on the strength of the control. Such strength can be quantified by summing up utilities for nodes that are controlled by the player i in time t:

$$v_i^t = \sum_{n:c^t(n)=i} u_i^t(n)^2. \tag{6}$$

Similarly, individual node's value can be computed as in the count-based case.

To account for the adversary who is effectively competing for the market share in a zero-sum fashion, we introduce a minimax procedure to enable players to reason strategically. Given a state tuple $(\{c^t(n)\}, \{T^t(n)\})$, the current player (the maximizing player) will explore all feasible choices, and for each choice, compute the objective value v_i^t by using either Equation (5) or (6) from the state space at number of look-ahead moves. Now it's adversary's turn to make choice, the assumption is that he will make choice that minimizes the maximizing player's objective value. In general both players are maximizing their own payoff value calculated according to their objective function at their number of look-ahead moves. The number of look-ahead moves allowed will be a player-specific parameter in our model. This process will continue until we reach one of the termination criteria.

To improve the performance of the above minimax search procedure, we apply a standard α-β pruning on the search tree. The α and β values refer to the lower bound for the maximizing player and the upper bound for the minimizing player. All nodes in the search tree that have values lower than α (for the maximizing player) or higher than β (for the minimizing player) will be pruned.

4. MODEL EVALUATION

Social influence networks are commonly modeled using a class of graph structures known as the scale-free networks, which exhibit high clustering coefficient, small mean shortest path length properties and power-law degree distribution. Scale-free networks exhibit higher fraction of nodes with large (larger than average) number of in-degree edges connected to them in a network. So networks of any topology that comply to the three properties above can be classified as scale-free model (also known as power-law degree distribution networks). Kempe et. al. [8] used the *Kleinberg's Small World* network structure as the basis to study gossip protocols for spreading information in a communication network. In our experiments, we employed the *JGraphT* Java graph library which contain mathematical graph theory objects and algorithm for generating the synthetic scale-free networks. These networks contain 100 nodes which are a reasonable size for representing influential social network with strong ties. According to Adam et. el., procurement decision of a consumer is influenced mainly by her neighbors with *strong ties*, compares to all others [1].

Given the size of our experimental networks, minimax algorithm is employed to evaluate the decision game tree. α-β pruning heuristic is incorporated in the minimax algorithm to reduce the size of state space need by pruning decision branches that prove to be less promising. The computation effort for α-β approach is upper bounded by the brute force approach in a complete tree search for each game stage. Our adversarial search problem is emulated using a simulation model written in Java. When the size of influence network is scaled up, the dynamics of our problem will result in an exponential growth in number of state space. However when such situation occurs, the sampling-based approaches can be used to improve the computational efficiency by sacrificing the comprehensiveness of search.

The following design is used to address our conjecture in this research. A total of 96 instances are evaluated using 6 sets of distinct network structure. The consumers in the influence networks are initialized under two conditions for the experiments. When the market players enter an unknown market space untainted by competition, there is ample amount of opportunity for growth given that all the consumers are yet to be explored. This is synonymous to a consumer influence network with no initial obligation (*contractual lock-in*) and every consumer can be targeted by the players. All the consumers adopt neither product choice at the initialization of game. Here we refer this condition as *null* initial state. On the other hand, the market can already be hotly contested, and consumers in the network are all already committed to certain existing obligations with different remaining lengths. Under this experimental condition, the game is initialized with each consumer assigned a random *contractual lock-in* period.

Table 1: The design of experiments.

Parameter	Incumbent	Entrant
Node Selection	count-based, value-based	
Budget Ratio	1	1/3, 1/2
Initial Lock-in	null, randomized	
Play Sequence	first, second	
Network Structure	1, 2, 3, 4, 5, 6	

The stronger player's budget is set to 1, and the weaker player is assigned either 1/3 or 1/2. Since the utility function of a customer for a particular product is determined by the ratio of her incoming neighbors owning that product, a budget of 1 implies the stronger player can convert *any* customer, while the weaker player can only convert a much smaller fraction of all customers. To remove first-mover's advantage, each player will get to play first for each network instance.

4.1 Results and Discussions

Given the budget limitation constraint on the *entrant* player, we foresee his market share to be lower than the adversary. However we conjectured that the *entrant* player could garner a larger market share simply by changing her node selection policy to consider long-term node values. Using the synthetic consumer influence networks data, we attested the effectiveness of the proposed *value-based* node selection policy. Without going into details, we would like to illustrate visually the niche-seeking behavior by our value-based strategy. As illustrated in Figure 1, a cluster is quickly identified and captured by the entrant player, and such behavior is consistently observed in the steady state and also for other network instances. For the rest of the section, we will define quantitative measures that allow us to quantify the niche-seeking behaviors.

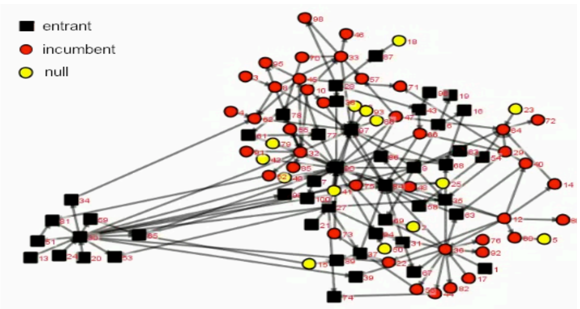

Figure 1: The effects of value-based approach.

4.1.1 *Identifying Set of Steady States*

Recall that our model is with infinite horizon, thus the adoption status might change radically from epoch to epoch. Therefore, unless this dynamic adoption process comes to a complete stop (e.g., one player completely dominates the market), it would be unfair to take any snapshot and conclude the performances of player's strategies.

If complete stop is not observe in a particular experiment instance, what we can do instead is to identify the *repeating states*. According to Section 3.1, the state at time t can be described by two vectors: $\{c^t(n)\}$ and $\{T^t(n)\}$. For the purpose of performance evaluation, we only need $\{c^t(n)\}$. Based on the above description, at least one state will eventually appear for the second time, since the set of feasible values of $c^t(n)$ is finite and t is unbounded. Suppose the state in time t is denoted as S_t and $S_t = S'$, if S' is observed again in time $t+\delta$, the experiment can be terminated early and the set $\{S_t, S_{t+1}, \ldots, S_{t+\delta-1}\}$ will contain all the repeating states of this experiment instance. We can make this claim because all strategies we proposed are deterministic and not dependent on $\{T^t(n)\}$.

This is an important observation, as we can now use the set of repeating states to quantify player's performances. When we report the experiment results, we focus on two measures: the average and the volatility of market shares over the set of repeating states. The analysis of experiment results is presented in the next subsection.

4.1.2 *Analysis*

The set of steady states defined in the previous subsection will appear repeatedly and infinitely. Because of this, a natural way to characterize a player's performance is to compute the average market share for all states in the set. Besides examining averages, we also compute the volatility of market shares from state to state (in the same repeating set); such measure allows us to judge the stability and robustness of a particular player's strategy as well.

The computation of volatility is inspired by and borrowed from the financial literature. To illustrate how it's computed, let $\{m_1, m_2, \ldots, m_\delta\}$ be the sequence of market shares for a particular player over the set of steady states. We can first compute the log market share relatives as: $r_i = \ln(m_{i+1}/m_i)$. The volatility can then be computed by computing the standard deviation of all r_i's.

The above two performance measures under different experiment settings (combination of *null* and *random* initialization and budget ratio of 1/3 and 1/2) are presented in Table 2.

Table 2: Summary of experiment results (CB and VB stand for count- and value-based respectively).

			# of Adopters (std. dev.) Adversary Strategy CB	VB	Volatility (against CB)
null	$\frac{1}{3}$	CB	45.2 (0.015)	24.3 (0.036)	0.0625
		VB	67.8 (0.031)	45.1 (0.023)	0.0335
	$\frac{1}{2}$	CB	46.3 (0.018)	25.5 (0.039)	0.0761
		VB	69.1 (0.027)	45.7 (0.021)	0.0337
random	$\frac{1}{3}$	CB	30.8 (0.027)	17.8 (0.028)	0.0170
		VB	40.9 (0.019)	28.0 (0.031)	0.0232
	$\frac{1}{2}$	CB	32.0 (0.021)	23.1 (0.023)	0.0130
		VB	44.4 (0.030)	30.8 (0.037)	0.0191

There are two important insights we can draw from the results presented in Table 2:

- For the entrant player, it's always better to adopt the value-based strategy. By simply adopting value-based strategy (which consider a node's long-term value), the entrant can perform considerably better than the myopic count-based strategy by more than 50% in all cases.

- For the incumbent player, it's also beneficial to adopt the value-based strategy. For all setups, the incumbent performs significantly better if he uses value-based strategy. Moreover, when the network is initialized with *null* status, we can see that the incumbent player performs worse than the value-based entrant player if he adopts count-based strategy. This is so for both entrant player's budget levels (1/2 and 1/3).

The above two observations from the experiment results strongly support our conjecture that to succeed in a competitive influence maximization game with infinite horizon, it's very important to evaluate a node's long-term value correctly. Although our value function is extremely simple, it's still significantly better than the conventional approach that just myopically maximize immediate market gains. In terms of volatility, we can see that against a count-based adversary in a null initialization, the entrant will enter a much stable steady state if he chooses the value-based approach. However, such difference diminishes when the initialization becomes random.

4.1.3 Niche Performance Metric

In the previous subsection, we show that value-based strategy performs better than the count-based approach under all circumstances and for both entrant and incumbent players. To find out whether such value-based strategy would result in players building niche customer base, we define a *niche performance metric* to measure the degree of niche seeking.

This metric calculates the proportion of edges where both the source and target nodes adopt the same product type over total number of edges. When the source and target nodes are adopting the same product type, it implies an influence propagation effect from the source to target node. On the contrarily, when the source and target node's product type is different from each other, it symbolizes a non-continuity in influence spread.

In our experiments, we observed that regardless of setups, the value-based approach always ends up with higher niche performance metric in steady states. For the null initialization, the comparison is summarized in Table 3. For the random initialization, the comparison is summarized in Table 4.

These experiment results confirm our second conjecture: it's indeed more advantageous for weaker player to concentrate on smaller niche in the scale-free networks.

Table 3: Niche performance metric under null initialization.

	Fraction of Adversary's Budget			
	1/3		1/2	
	CB	VB	CB	VB
Average	0.425	0.711	0.443	0.718
Std. dev.	0.019	0.033	0.024	0.038

Table 4: Niche performance metric under random initialization.

	Fraction of Adversary's Budget			
	1/3		1/2	
	CB	VB	CB	VB
Average	0.247	0.365	0.262	0.388
Std. dev.	0.038	0.054	0.023	0.046

5. CONCLUSIONS

Influence maximization is a well-studied problem in the literature. It captures researchers' attention recently due to the fact that societal-scale social networks are increasingly ubiquitous. However, the classical research on influence maximization lacks either the explicit modeling of adversaries or the consideration of time.

In this paper, we construct a simple two-player, infinite horizon influence maximization model to illustrate the importance of considering both adversary and time. The count-based approach, which represents the conventional influence maximization approach, is shown to perform much worse than the value-based approach, which approximate the long-term values of individual nodes. Furthermore, we show that the value-based approach achieves the better performance via niche-seeking, which is measured using our niche performance metric.

Our model provides a simple and clean framework for future study on extending well-known results on influence maximization to real-world marketing problems.

6. REFERENCES

[1] P. Adams. How your customers' social circles influence what they buy, what they do and where they go. Presented Web 2.0 Expo 2011, San Francisco, California, USA, 2011.

[2] R. Aumann. Advances in game theory. Number 52 in Annals of mathematics studies, Princeton, NJ, 1964. Princeton Univ. Press.

[3] F. M. Bass. A new product growth for model consumer durables. *Management Science*, 15(5):215–227, 1969.

[4] N. Economides. The economics of networks. *International Journal of Industrial Organization*, 14(6):673–699, 1996.

[5] M. Garlick and M. Chli. An agent-based simulation of lock-in dynamics in a duopoly. In *the 9th International Conference on Autonomous Agents and Multiagent Systems, Toronto, Canada, May 10 - May 14, 2010*, pages 1545–1546, 2010.

[6] M. Granovetter. Threshold model of collective behavoir. *The American Journal of Sociology*, 83(6), 1978.

[7] M. Janssen and W. Jager. An integrated approach to simulating behavioral processes: A case study of the lock-in of consumption patterns. *Journal of Artificial Societies and Social Simulation*, 2(2), 1999.

[8] Kempe, Kleinberg, and Demers. Spatial gossip and resource location protocols. *JACM: Journal of the ACM*, 51, 2004.

[9] D. Kempe, J. Kleinberg, and E. 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*, KDD '03, pages 137–146, New York, NY, USA, 2003. ACM.

[10] D. Kempe, J. Kleinberg, and E. Tardos. Influential nodes in a diffusion model for social networks. *Computer*, 3580:1127–1138, 2005.

[11] D. Kreps and R. Wilson. Sequential equilibria. *Econometrica*, 50:863–894, 1982.

[12] J. A. Norton and F. M. Bass. A diffusion theory model of adoption and substitution for successive generations of high-technology products. *Management Science*, 33(9):1069–1086, 1987.

[13] B. Robinson and C. Lakhani. Dynamic price models for new-product planning. *Management Science*, 21(10):1113–1122, 1975.