TiSA: Time-Dependent Social Network Advertising

Wei Wang[†], Linlin Yang[‡], Qing Liao[†], Xiang Zhu[*], Qian Zhang[†]
[†]Department of Computer Science and Engineering,
[‡]Fok Ying Tung Graduate School,
Hong Kong University of Science and Technology
[*]College of Computer, National University of Defense Technology
Email: {gswwang, lyangah, qnature, qianzh}@ust.hk
zhuxiang@nudt.edu.cn

***Abstract*—Much of today's online social network (OSN) system relies on advertising for financial support. To improve the effectiveness of advertising, online advertisers tend to leverage influential users to deliver ads. Most of existing efforts on online advertising have focused on single-shot scenarios or assume static OSN models, while they overlook the fact that actions of advertising affect users' behaviors. In this paper, we investigate the behaviors of Sina Weibo users over three months, and make the observation that advertising affects the behaviors of the user's followers, which in turn has an impact on the effectiveness of future advertising. Based on this observation, we propose TiSA, a time-dependent advertising framework, which considers the future impact of advertising. Under this framework, the advertiser and the user make their decisions based on their instant utilities as well as future utilities. We also devise a learning algorithm with provable convergence to obtain the optimal policies. Evaluations using three month traces of 975 Sina Weibo users have been conducted, and the results validate the effectiveness of the proposed framework by showing that the utilities of all entities are significantly improved compared with traditional systems.**

I. Introduction

Social network advertising provides financial support for a large portion of today's online social network (OSN) ecosystem, and is displayed in a variety of forms embedded in web sites. As the behavior of a user is affected by others via perceived relationships and received information from online friends, organizations, and societies, online advertisers tend to leverage the power of influential people in OSNs to increase brand awareness and influence people's purchasing actions [1], [2]. Many companies like IBM, Microsoft, and Walmart have adopted this advertising strategy, and distributed advertisements (ads) among influential people who are willing to post the ads.

To improve the effectiveness of social network advertising, many efforts have been devoted to identifying influential people [3] and targeted advertising [4], [5]. However, most of these studies focus on "single-shot" scenarios or assume static OSN models, while the correlations and interactions between users' influence and advertising in a long-term perspective have not yet been investigated. The motivation behind this is the profound impact of advertising on users' future behaviors in OSNs. For example, a user continuing posting ads within a short time may annoy some of its followers and even drive them away, which in turn affects the effectiveness of advertising in future. Without considering this future impact, the advertising strategies are myopic and fail to maximize advertisers profits in a long-term perspective.

The target of this paper is to fill the gaps in long-term social network advertising. In this paper, we argue that the impact of advertising on users' future behaviors should be considered. We first characterize the impact of advertising in OSNs by investigating Sina Weibo, which is the largest microblogging social network in China. We observe that frequent advertising has a negative impact on the user's number of followers, which undermines the advertising effectiveness in future. Therefore, the advertiser and the users should carefully choose their policies to strike a balance between the instant utilities and future utilities.

To obtain long-term advertising policies, we propose *TiSA*, a **Ti**me-dependent **S**ocial network **A**dvertising framework that considers the dynamics of OSNs and the impact of advertising. In TiSA, we model the strategic and dynamic interactions between a user and an advertiser as a stochstic game, where both players concern about their long-term utilities. In each time slot, both the user and the advertiser observe the current number of the user's followers, as well as the numbers of followers in the past, and then take their actions based on their policies. Specifically, the user decides the number of ads to post in each time slot, while the advertiser determines the price it pays the user for posting each ad. To obtain the optimal policies of the user and the advertiser, we design a learning algorithm for TiSA with provable convergence.

The main contributions of this paper are summarized as follows.

- We investigate the behaviors of 975 Sina Weibo users over three months, and characterize the impact of advertising in OSNs. These observations have profound implications on the design of long-term advertising systems in OSNs.
- We propose a time-dependent advertising framework that maximizes the long-term utilities of the advertiser and the user. To the best of our knowledge, this is the first advertising framework that considers the impact of advertising and the dynamics of followers.
- We model the dynamic interactions between the advertiser and the user as a stochastic game, and derive the

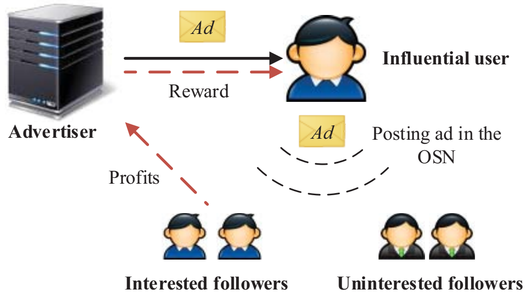

Fig. 1. Social network advertising system.

optimal policies via a learning algorithm with provable convergence.

- We evaluate our framework using three-month traces of 975 users. Promisingly, the results show that TiSA outperforms baseline schemes in terms of long-term utilities.

The reminder of this paper is structured as follows. We begin in Section II with the advertising system model in social networks, and then characterize the impact of advertising by investigating Sina Weibo. Section III describes the detailed time-dependent advertising problem formulation, and presents a learning algorithm to obtain the optimal policies. Evaluation of TiSA is presented in Section IV. Section V reviews related work, followed by conclusion in Section VI.

II. Characterizing Advertising in Online Social Networks

In this section, we first describe the advertising system in online social networks by introducing the user model, the advertiser model, and their interactions. Then, we investigate the impact of advertising in social networks by analyzing the user behaviors in Sina Weibo.

A. Advertising System in Online Social Networks

An advertising system in an OSN consists of an advertiser, influential users who post ads in the OSN for the advertiser, and the followers of the influential users. The influential users refer to the users connect directly or indirectly with a massive number of users. In the OSN, a user can post messages which can be seen by its followers. In this paper, the relationship of following and being followed can be either reciprocal (e.g., friend-relationship in Facebook) or non-reciprocal (e.g., relationship in Twitter and Sina Weibo).

In the OSN, the advertiser delivers its ads to users via influential users, who post ads forwarded by the advertiser. In turn, the advertiser pays a certain amount of reward to the influential users for posting ads. The ads posted by the influential users can be seen by their followers. Among these followers, there may be interested users who are potential buyers, from whom the advertiser earns profits. However, there may exist followers who are not interested in the ads, or may even regard them as spams. In this case, the uninterested users can be annoyed by frequent appearance of ads in their incoming messages, and unfollow the users posting lots of

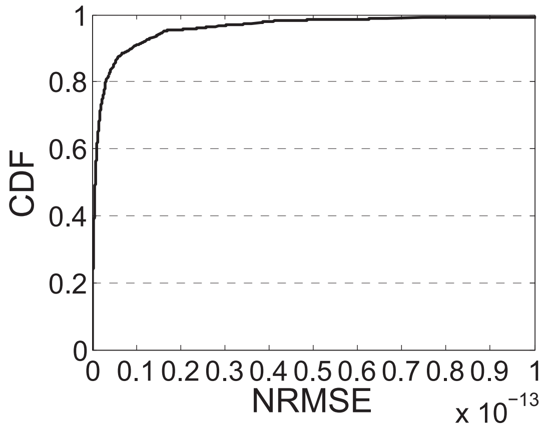

Fig. 2. CDF of NRMSE for linear regression.

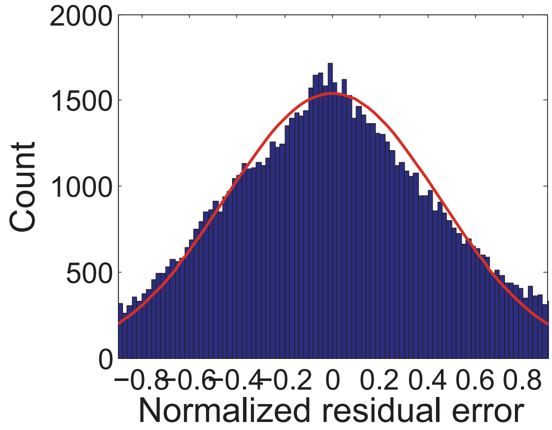

Fig. 3. Distribution of normalized residual error.

ads. As the user's influence largely depends on the number of followers, the unfollowing action undermines the influential user's influence, thereby impairing the user's future payoff. Therefore, the influential users need to carefully strike a balance between the current payoff and future payoffs. To characterize the impact of advertising on user's followers, we take a case study on Sina Weibo.

B. Impact of Advertising in Online Social Networks

We characterize the follower increment using the user traces in Sina Weibo. We crawled and collected profiles of 22,514,394 users and their posts starting from September 2011 to September 2014. In order to collect users' profiles and posts, we started with influential users who have millions of followers, and then crawled breadth-first along the direction of the users' followers and followings. As our study requires complete post records of active users, we select 975 users with complete records over 90 days. These users continuous post new messages each day.

To avoid ad recognition errors, we manually categorize all messages into ads and non-advertising messages. We first investigate how the number of followers varies over time for

955 users who have not posted ads over the 90 days. We regard the number of followers as a function of time (one day as a time slot), and the corresponding curve exhibits strong linear properties. To validate this observation, we adopt linear regression for curve fitting. Fig. 2 draws the cumulative distribution function (CDF) of normalized root-mean-square error (NRMSE) for linear regression. The NRMSEs of over 98% users are smaller than 5×10^{-13}, which demonstrates that linear curves tightly fit the number of follower variances. For the residue errors incurred by curve fitting, we plot a histogram in Fig. 3, which consists of residual errors of all users. We see that the residual error can be well fit by a Gaussian distribution (the red curve in the figure), where the NRMSE is merely 0.0682. Therefore, for users without posting ads, the number of followers variance over time can be modeled by

$$n^t = at + b + \delta, \tag{1}$$

where n^t is the number of followers in time slot t, a, b the parameters learned by linear regression, δ the random variable following a Gaussian distribution. (1) implies that the expectation of the follower increment, i.e., $\mathbb{E}[n^{t+1} - n^t] = \mathbb{E}[a + \delta_{t+1-\delta_t}] = a$, is a constant over time. In the following, we show that this observation does not stand for users who post ads: the follower increment of a user varies when the user posts different number of ads.

Next, we study the advertising impact of 20 influential users who have posted ads in 90 days. Fig. 4 shows the scatter plot of daily follower increments with the numbers of ads posted in the corresponding days. The results reveal that the large follower increment occurs only when the number of ads is small. The follower increment diminishes as the number of ads increases. Fig. 5 depicts the CDF of Pearson correlation coefficient ρ between the follower increment and the number of ads. In the figure, we compute ρ for each user. The results report that ρ is negative for 90% users, and $\rho < -0.5$ for more than 50% users, which further validates the negative impact of advertising on follower increment.

The above case study provides the following implication: without considering the impact of advertising on follower increment, single-shot strategy based advertising systems are not sustainable for long-term operation. This observation motivates the design of TiSA, a time-dependent advertising framework that takes the impact of advertising into account, so as to maximize the long-term utilities of both the advertiser and the user.

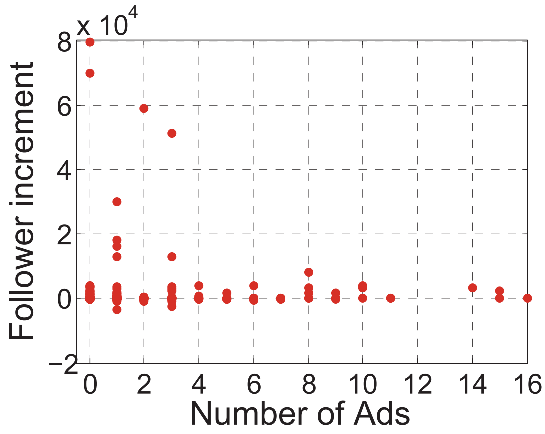

Fig. 4. Follower increments under different numbers of daily-posted ads.

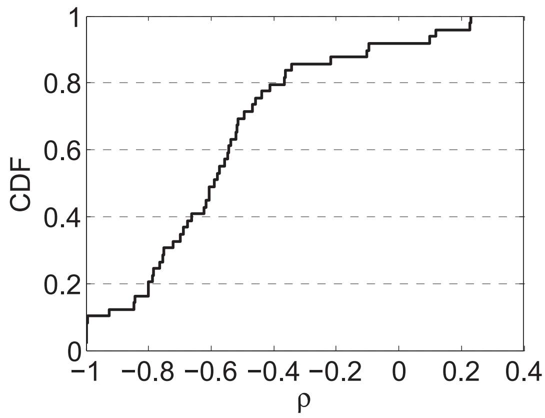

Fig. 5. CDF of Pearson correlation coefficient between follower increment and number of daily-posted ads.

III. Time-Dependent Advertising Framework

In this section, we design TiSA, a time-dependent advertising framework in OSNs. We first formulate the advertising problem as a twp-player stochastic game, and then develop a learning algorithm to derive optimal policies for TiSA.

A. Problem Formulation

Recall that the number of followers is a random variable that depends on the user's advertising behaviors, and the number of followers determines the effectiveness of advertising. Thus, the interactions in social network advertising can be modeled as stochastic game between the user and the advertiser.

System state. In each time slot, a user has a certain number of followers, which is an essential factor that determines the user's influence in the OSN. Therefore, the number of followers is considered as the system state. The number of followers at time t is denoted as $n^t \in \mathbb{N}$, and the system state at time t is defined by $S^t = n^t$.

Actions. After observing the state S^t at each stage, both the user and the advertiser decides their actions for the current stage. As we observed in Sina Weibo, users usually repeatedly post the same ad multiple times to make it easier to be seen by other users. Hence, in each time slot, the user controls the number of posts for each ad forwarded by the advertiser. Formally, the action of the user at time t is defined as $\mathbf{a}_u^t = \{m_1^t, ..., m_M^t\}$, where $m_i^t \in [0, 1]$ is the fraction of the user's posts that contain ad i within the time slot t, and M is the total number of ads. As the total number of posts that contain ads cannot exceed the number of all posts, we have the following constraints for m_i^t.

$$\sum_i m_i^t \leq 1,$$
$$0 \leq m_i^t \leq 1, \forall i. \quad (2)$$

On the other hand, the advertiser decides the unit prices it pays to the user who posts its ads. Mathematically, the adversary's actions at time t are defined as $\mathbf{a}_u^t = \{p_1^t, ..., p_M^t\}$, where $p_i^t \in (0, U_p]$ is the unit price for ad i and U_p is the upper bound for unit price.

State transitions. The system state S^t is uncertain and depends on the actions of the user and the advertiser. Thus, the state transition probability can be computed by

$$\Pr[S^{t+1}|S^t, \mathbf{a}_u^t, \mathbf{a}_a^t] = \Pr[n^{t+1}|n^t, \mathbf{a}_u^t, \mathbf{a}_a^t]. \quad (3)$$

Stage payoffs. After defining the states and actions, we give a concrete expression of stage payoffs. The payoff function of the user is defined to be reward earned by posting ads with weighted penalty on influence loss, which is expressed as

$$r_u(S^t, \mathbf{a}_u^t, \mathbf{a}_a^t) = W(\mathbf{a}_u^t, \mathbf{a}_a^t) - \omega \cdot L(S^t, \mathbf{a}_u^t), \quad (4)$$

where $W(\mathbf{a}_u^t, \mathbf{a}_a^t)$ is the reward paid by the advertiser for ad posting, ω the equivalent economic loss caused by unit influence loss, and $L(S^t, \mathbf{a}_u^t)$ the influence loss. Recall that The advertiser determines the price $\mathbf{a}_a^t$ for each ad posted by the user. Thus, we have $W(\mathbf{a}_u^t, \mathbf{a}_a^t) = M^t {\mathbf{a}_u^t}^\top \mathbf{a}_a^t$, where M^t is the total number all posts in time slot t. $L(S^t, \mathbf{a}_u^t)$ is the measure of the uninterested followers' degree of dissatisfaction with the undesired ads posted by the user they follow, which is considered as a negative impact of posting ads. Sigmoid function has been widely used to approximate the user's satisfaction/dissatisfaction with respect to service qualities [6]. Concretely, $L(\mathbf{a}_u^t, S^t)$ is measured as

$$L(\mathbf{a}_u^t, S^t) = \frac{1}{1 + e^{-\alpha(\mathbf{a}_u^t n^t - \eta_l)}}, \quad (5)$$

where α determines the steepness of the quality of service satisfactory curve, η_l the satisfaction threshold below which the followers have very limited dissatisfaction (the function curve is convex) and above which the followers' dissatisfaction rapidly approaches an asymptotic value (the function curve is concave).

In this paper, we consider users who are active in the OSN for a long time. Thus, the user cares about the current as well as the future payoffs. Normally, a recent payoff is more valuable than a payoff that will be received in the faraway future. Thereby, the user's utility is the expected sum of discounted stage payoffs

$$U_u = \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t r_u(S^t, \mathbf{a}_u^t, \mathbf{a}_a^t)\right], \quad (6)$$

where γ^t is the discount factor of the advertising payoff. The user's objective is to derive an *optimal policy* that maximizes U_u.

On the other hand, the advertiser's payoff is determined by the profit it earned minus the amount of payment $W(\mathbf{a}_u^t, \mathbf{a}_a^t)$ to the user, which is written as

$$r_a(S^t, \mathbf{a}_u^t, \mathbf{a}_a^t) = P(S^t, \mathbf{a}_u^t) - W(\mathbf{a}_u^t, \mathbf{a}_a^t), \quad (7)$$

where $P(S^t, \mathbf{a}_u^t)$ is the profit the advertiser earned from advertising, which has been widely modeled as sigmoid function [7]. Note that the advertiser earns more profit when i) the user who posts ads have more followers, and ii) the user posts more ads. Then, the advertising profit $P(S^t, \mathbf{a}_u^t)$ can be measured as

$$P(\mathbf{a}_u^t, S^t) = \frac{\theta}{1 + e^{-\beta(\mathbf{a}_u^t n^t - \eta_p)}}, \quad (8)$$

where θ is the advertiser's valuation for advertising profit with respect to the payment $W(\mathbf{a}_u^t, \mathbf{a}_a^t)$, β and η_p the steepness and threshold of the sigmoid curve. Analogous to the user's utility, the advertiser's utility can be expressed as

$$U_a = \mathbb{E}\left[\sum_{t=0}^{\infty} \gamma^t r_a(S^t, \mathbf{a}_u^t, \mathbf{a}_a^t)\right]. \quad (9)$$

B. Learning the Optimal Advertising Policy

After formulating the time-dependent advertising game, we now elaborate the algorithm to obtain the Nash Equilibrium (NE) of the advertising game, so as to derive the optimal advertising policy.

Optimal Policies in TiSA. We denote the set of system states as $\mathcal{S}$, and the user's and advertiser's policies as π_u, π_a, respectively. For a system state $s \in \mathcal{S}$, the user's utility can be written in the form of *state value* $V^\pi(s)$

$$\begin{aligned} V_u^\pi(s) &= \sum_{t=0}^{\infty} \gamma^t \mathbb{E}[r_u(S^t, \mathbf{a}_u^t, \mathbf{a}_a^t)|\pi_u, \pi_a, S^0 = s] \\ &= r_u(s, \mathbf{a}_u^\pi, \mathbf{a}_a^\pi) + \gamma \sum_{s'} \Pr[s'|s, \mathbf{a}_u^\pi, \mathbf{a}_a^\pi] V_u^\pi(s'), \quad (10) \end{aligned}$$

where $\mathbf{a}_u^\pi, \mathbf{a}_a^\pi$ are the actions determined by π_u, π_a.

Similarly, the advertiser's utility can be written as

$$V_a^\pi(s) = r_a(s, \mathbf{a}_u^\pi, \mathbf{a}_a^\pi) + \gamma \sum_{s'} \Pr[s'|s, \mathbf{a}_u^\pi, \mathbf{a}_a^\pi] V_a^\pi(s'). \quad (11)$$

The optimal policies $\{\pi_u^*, \pi_a^*\}$ are obtained at the NE point, which is defined as follows.

Definition 1 (NE in TiSA) *In the stochastic game for TiSA, an NE point is a pair of policies $\pi^* = \{\pi_u^*, \pi_a^*\}$, such that for all state $s \in \mathcal{S}$,*

$$V_u^{\pi^*}(s) \geq V_u^{\pi^a}(s) \text{ and } V_a^{\pi^*}(s) \leq V_a^{\pi^u}(s), \quad (12)$$

where $\pi^a = \{\pi_u, \pi_a^*\}, \forall \pi_u$, and $\pi^u = \{\pi_u^*, \pi_a\}, \forall \pi_a$.

The learning algorithm for TiSA. Note that the minimax algorithm [8], [9] can be used to learn the optimal policies in zero-sum stochastic game, while the minimax algorithm cannot be applied to TiSA as the user's and the advertiser's payoffs are not the opposite of each other. Therefore, we

Algorithm 1 State Value Learning Algorithm in TiSA

Input: System state set $\mathbf{S}$, state transition probability $\Pr[S^{t+1}|S^t, \mathbf{a}_u^t, \mathbf{a}_a^t]$, stage payoffs $r_u(S^t, \mathbf{a}_u^t, \mathbf{a}_a^t), r_a(S^t, \mathbf{a}_u^t, \mathbf{a}_a^t)$

Output: π_u^*, π_a^*

// *1. initialization*

$t \leftarrow 0$;

For all $s \in \mathbf{S}$, $V_u^0(s) \leftarrow 1, V_a^0(s) \leftarrow 1$;

Initialize policies π_u^0, π_a^0, which is the mixed strategy NE solution of the stage game $V_u^0(s), V_a^0(s)$;

// *2. iteration*

repeat

  Select an action pair $a_u^t, a_a^t\}$ based on π_u^t, π_a^t;

  Update S^{t+1} after both players take their actions $a_u^t, a_a^t\}$;

  Update state values according to (13);

  Update policies pair π_u^{t+1}, π_a^{t+1} according to Definition 1 such that

$$V_u^{\pi^*}(s) \geq V_u^{\pi^a}(s) \text{ and } V_a^{\pi^*}(s) \leq V_a^{\pi^u}(s); \qquad (14)$$

  $t \leftarrow t+1$;

until π_u^t, π_a^t converge

$\pi_u^*, \pi_a^* \leftarrow \pi_u^t, \pi_a^t$;

propose a new learning algorithm for TiSA to derive the optimal policies.

In order to obtain the optimal policies $\{\pi_u^*, \pi_a\}$, the user and the advertiser in TiSA update their state values according to the following rule.

$$\begin{cases} V_u^{t+1}(s) = (1-\alpha^{t+1})V_u^t(s) \\ +\alpha^{t+1}\left(r_u(s, \mathbf{a}_u^\pi, \mathbf{a}_a^\pi) + \gamma \sum_{s'} \Pr[s'|s, \mathbf{a}_u^\pi, \mathbf{a}_a^\pi] V_u^\pi(s')\right), \\ \\ V_a^{t+1}(s) = (1-\alpha^{t+1})V_a^t(s) \\ +\alpha^{t+1}\left(r_a(s, \mathbf{a}_u^\pi, \mathbf{a}_a^\pi) + \gamma \sum_{s'} \Pr[s'|s, \mathbf{a}_u^\pi, \mathbf{a}_a^\pi] V_a^\pi(s')\right), \end{cases} \qquad (13)$$

where $V_u^t(s), V_u^t(s)$ are approximates of $V_u^{\pi^*}(s), V_a^{\pi^*}(s)$, $\alpha^t \in [0, 1)$ the learning rate. In our algorithm, we set $\alpha^t = 1/t$. The update rule (13) iteratively update $V_u^t(s), V_u^t(s)$ until they converge to $V_u^{\pi^*}(s), V_a^{\pi^*}(s)$.

The detail of our learning algorithm is stated in 1. First, Algorithm 1 initializes policies and state values for $t = 0$. Then, it iteratively updates policies and state values according to (13) and Definition 1 until they converge to optimal values.

IV. Evaluation

In this section, we evaluate TiSA using trace-driven simulations. We feed Sina Weibo dataset into our simulations as the dynamic OSN model, and evaluate the effectiveness of different policies.

A. Setup

Dataset. We collected the Sina Weibo data through the API provided by Sina Weibo. The Sina company provides about 20% of the most reliable data available in China. We collected the data of 22,514,394 users from September 2011 to September 2014 (1003 days in total), from which we select a subset of 975 active users who continuous post new updates each day within a three months duration.

System parameters. We extract the follower variance from the Sina Weibo dataset, which is used to emulate the state transitions in our evaluation. Unless explicitly otherwise stated, we use the following system parameters in our evaluation. The upper bound for unit price U_p is set to 1, total number of ads set to 5, satisfaction thresholds η_l, η_p set to 0.7, steepness α, β parameters set to 0.1, discount factor γ set to 0.8.

Baselines. We compare the performance of TiSA with two typical advertising systems, i.e., *single-shot* advertising and *static* advertising. The single-short advertising strategy aims to maximize the instant utilities, and obtains policies by solving each stage game independently. The static advertising strategy use static policies for all stages. The static policies are obtained by solving the first stage game.

B. Results

We first compare the overall performance of all advertising systems in Fig. 6 and Fig. 7. Fig. 6 shows the user's sum of discounted payoffs U_u under numbers of stages. The sums of discounted payoffs under all policies converge to constant values as the number of stages increase. We see that advertising policies obtained TiSA and single-shot achieve higher sums of discounted payoffs than static policies. This is because TiSA and single-shot systems adjust actions over time based on the user's number of followers, while the static system takes actions only based on the initial state. Moreover, TiSA outperforms single-shot system in all stages. The reason is that TiSA also considers the future impact of advertising actions when learning the policies.

Fig. 7 depicts the advertiser's sum of discounted payoffs in different advertising systems. Similar to the results in Fig. 6, TiSA outperforms the other two systems in all stages. Hence, when advertisers rely on influential users in OSNs for advertising, the best strategy is to take the dynamics of the OSNs into account when making decisions.

Fig. 8 reports the advertiser's sums of discounted payoffs under different satisfaction threshold η_p. When η_p is smaller, the advertiser earns more profits for each posted ad. Thus, for all systems, the advertiser's sums of discounted payoffs diminish when the satisfaction threshold grows. We also see that the differences between the sums of discounted payoffs achieved by TiSA and the other two systems grow larger as the satisfaction threshold increases. The reason is that when the satisfaction threshold is very low, high profits can be easily earned by posting a small number of ads. As such, the profit part dominates the advertiser's utility, and stays relatively the same over different actions. Based on this observation, we conclude that the merits of TiSA is more significant in the case of high satisfaction threshold, where the advertiser's total payoff is very sensitive to different policies. Hence, the advertising system with high satisfaction threshold should be

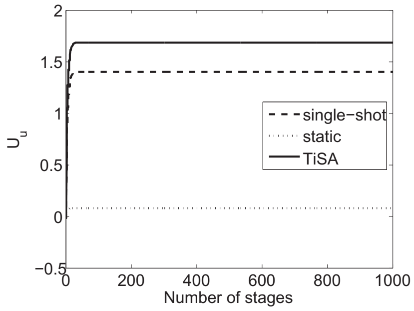

Fig. 6. User's sum of discounted payoffs under different policies.

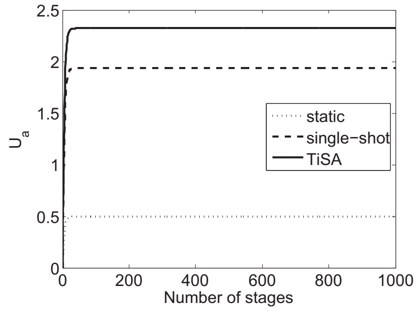

Fig. 7. Advertiser's sum of discounted payoffs under different policies.

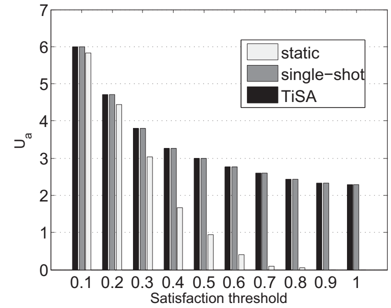

Fig. 8. Advertiser's sum of discounted payoffs under different satisfaction thresholds.

carefully designed to strike a balance between instant utility and future utility.

V. Related Work

Targeted advertising. The aim of targeted advertising is to devise effective advertising systems to find interested users or potential buyers. Chakrabarti et al. [10] propose a new class of models to display relevant ads to web pages of different contents. Kodialam et al. [4] devise a privacy-preserving protocol to deliver ads to privacy-sensitive users, who send falsified click information to an ad broker according to predetermined rules. These systems focus on single-shot scenarios and cannot be applied to long-term advertising in OSNs where a user's advertising behaviors affect its future utilities.

Social influence and advertising. Recently, online social networks are considered as a new dimension for marketing. The notion of social network advertising is first proposed as an important motivations for social influence studies. Bakshy et al. [3] conducted several field experiments to investigate how social influence affects users' responses to social network advertising, which provides profound implications for ad optimization and user interface design. Provost et al. Ning et al. [11] study the problem of ad dissemination mobile social networks, and devise an incentive scheme to motivate users to forward ads. A privacy-friendly framework is proposed in [12] to deliver different ads to users based on their browsing behaviors. The fundamental difference between these works and TiSA is that TiSA considers the time-dependency in social network when making advertising policies.

VI. Conclusion

This paper presents TiSA, the first time-dependent advertising system that considers the impact of advertising on followers' future behaviors. We first investigate the impact of advertising by taking a case study on Sina Weibo. Based on the model extracted from Sina Weibo dataset, we formulate the advertising problem as a stochastic game, which considers the dynamics of followers and future impacts of the user's and advertiser's actions. Furthermore, we devise a learning algorithm to obtain the optimal policies that maximizes long-term utilities. We hope that our investigation on the impact of advertising and the proposed time-dependent framework can provide some implications for future advertising system designs.

Acknowledgment

The research was support in part by grants from 973 project 2013CB329006, China NSFC under Grant 61173156, RGC under the contracts CERG 622613, 16212714, HKUST6/CRF/12R, and M-HKUST609/13, as well as the grant from Huawei-HKUST joint lab.

References

[1] S. Milstein, B. Lorica, R. Magoulas, G. Hochmuth, A. Chowdhury, and T. O'Reilly, *Twitter and the micro-messaging revolution: Communication, connections, and immediacy–140 characters at a time*. O'Reilly Media, Incorporated, 2008.

[2] W. Wang, L. Yang, Y. Chen, and Q. Zhang, "A privacy-aware framework for targeted advertising," *Elsevier Computer Networks*, vol. 79, pp. 17–29, 2015.

[3] E. Bakshy, D. Eckles, R. Yan, and I. Rosenn, "Social influence in social advertising: evidence from field experiments," in *ACM Proc. EC*, 2012.

[4] M. Kodialam, T. Lakshman, and S. Mukherjee, "Effective ad targeting with concealed profiles," in *Proc. IEEE INFOCOM*, 2012.

[5] T. Ning, Z. Yang, H. Wu, and Z. Han, "Self-interest-driven incentives for ad dissemination in autonomous mobile social networks," in *Proc. IEEE INFOCOM*, 2013.

[6] H. Lin, M. Chatterjee, S. Das, and K. Basu, "Arc: an integrated admission and rate control framework for competitive wireless cdma data networks using noncooperative games," *IEEE Trans. Mobile Comput.*, vol. 4, no. 3, pp. 243–258, 2005.

[7] J. K. Johansson, "Advertising and the s-curve: A new approach," *JSTOR Journal of Marketing Research*, pp. 346–354, 1979.

[8] B. Wang, Y. Wu, K. R. Liu, and T. C. Clancy, "An anti-jamming stochastic game for cognitive radio networks," *IEEE J. Sel. Areas Commun.*, vol. 29, no. 4, pp. 877–889, 2011.

[9] W. Wang and Q. Zhang, "A stochastic game for privacy preserving context sensing on mobile phone," in *Proc. IEEE INFOCOM*, 2014, pp. 2328–2336.

[10] D. Chakrabarti, D. Agarwal, and V. Josifovski, "Contextual advertising by combining relevance with click feedback," in *Proc. ACM WWW*, 2008.

[11] M. Kodialam, T. Lakshman, and S. Mukherjee, "Self-interest-drive incentives for ad dissemination in autonomous mobile social networks," in *Proc. IEEE INFOCOM*, 2013.

[12] F. Provost, B. Dalessandro, R. Hook, X. Zhang, and A. Murray, "Audience selection for on-line brand advertising: privacy-friendly social network targeting," in *Proc. ACM SIGKDD*, 2009.