Optimizing the Effectiveness of Incentivized Social Sharing

Joseph J. Pfeiffer III
Microsoft
Redmond, WA, USA
joelpf@microsoft.com

Elena Zheleva
University of Illinois at Chicago
Chicago, IL, USA
ezheleva@uic.edu

***Abstract*—Social media has become an important tool for companies interested in increasing the reach of their products and services. Some companies even offer monetary incentives to customers for recommending products to their social circles. However, the effectiveness of such incentives is often hard to optimize due to the large space of incentive parameters and the inherent tradeoff between the incentive attractiveness for the customer and the return on investment for the company. To address this problem, we propose a novel graph evolution model, *Me+N* model, which provides flexibility in exploring the effect of different incentive parameters on company's profits by capturing the probabilistic nature of customer behavior over time. We look at a specific family of incentives in which customers get a reward if they convince a certain number of friends to purchase a given product. Our analysis shows that simple monetary incentives can be surprisingly effective in social media strategies.**

I. Introduction

Online social media and social networks have transformed the ways in which people communicate. Many companies have websites where millions of users can share information with one another, including photos, music, news, products, and services. These companies encourage sharing between individuals because social recommendations can increase traffic to their websites, resulting in higher engagement and revenue, in a process known as *viral marketing*. To further utilize the benefits of viral marketing, companies can *incentivize* users by giving them monetary rewards for sharing.

There are two types of viral marketing, *direct viral marketing* and *mass-marketing sharing incentives*. *Direct viral marketing* aims to identify a relatively small set of individuals with high network centrality values, i.e., *influencers*, and then giving them a discounted or free product in the hope that they would influence many others to also adopt this product [1, 2, 3, 4, 5]. In contrast, *mass-marketing sharing incentives* reward *any* user who convinces their friend(s) to adopt a product, but withhold the reward until the friend completes the *incentive goal*. There is research which suggests that methods for predicting influencers are generally unreliable, and that targeting a wider range of users is a more cost-effective method for information diffusion [6, 7].

ASONAM '17, July 31 - August 03, 2017, Sydney, Australia

ACM ISBN 978-1-4503-4993-2/17/07?/$15.00
http://dx.doi.org/10.1145/3110025.3110147

Many examples of mass-marketing sharing incentives can be found online. Lovefilm, an online movie rental company, gives £20 to any customer whose friend signs up for their service. In addition, the friend receives the first one-month subscription for free. Fab, an online marketplace for designer items, gives a customer and her friend $25 each when the friend joins and makes his first purchase. At LivingSocial when a user buys a deal and persuades three or more friends to purchase the same deal, the user is refunded the price of the deal. With such incentives, successful recommenders self-select (or prove) themselves as influencers, and the company does not need to identify them explicitly or pay the reward to users whose sharing does not lead to increased adoption.

However, the profitability of mass-marketing sharing incentives remains largely unknown, partly due to the fact that real-world data on sharing incentives is difficult to obtain. In this paper, we propose a framework for studying their effectiveness, as well as a graph evolution model which simulates incentivized user behavior over time. In particular, we study a family of incentives, *Me+N*, in which after purchasing a product, a customer is offered a reward if she convinces N or more friends to purchase the same product.

II. Incentivized Sharing

The scenario which motivates our work is the following: a company is interested in tapping into the social network of its customers, in order to increase the adoption of its products. We assume that the company 1) is interested in implementing a *Me+N* sharing incentive only when it increases its profits, i.e., not investing in an incentive cost that is higher than the additional revenue coming from friends' adoption, and 2) has a limited budget for running controlled experiments. If the second assumption was not an obstacle, then it is best to run large-scale A/B tests, exploring the full parameter space and assessing profits directly. In most realistic scenarios, such experiments require a very large user base, they can be costly in terms of lost profits, software infrastructure to develop and maintain, as well as increased customer confusion and customer service cost.

	Margin = 10%			Margin = 50%		
	% additional shares due to incentive					
	0%	5%	50%	0%	5%	50%
Optimal Incentive	No Incentive	Maximum N		No Incentive	Maximum N	
Better Than No Incentive	No Incentive	*Me+4*, etc.	*Me+3*, etc.	No Incentive	*Me+3*, etc.	*Me+2*, etc.
Safe Incentive	Me+3, *Me+4*, etc.			Me+2, *Me+3*, etc.		

TABLE I: Incentive comparison for different incentive parameters under *Me+N model*.

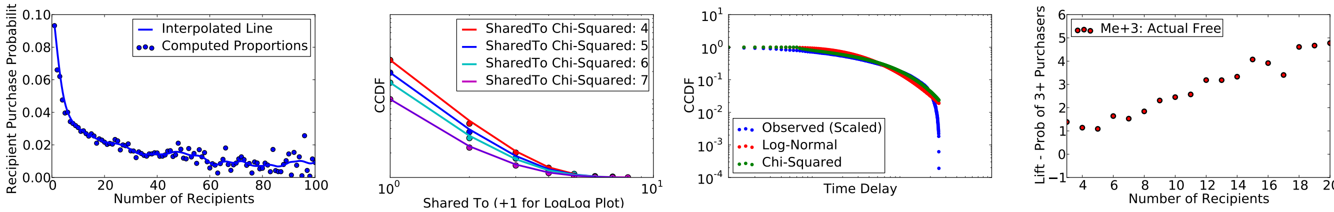

(a) (b) (c) (d)

Fig. 1: Observed properties of Me+3: (a) Recipient purchase probabilities. (b) Free Deal distributions given a share outdegree. (c) The awakening function. (d) Distribution of successful recommendations conditioned on the number of recipients

The key questions we would like to answer are:

Q1: Should a company implement a *Me+N* incentive?
Q2: For what N would the company increase its profits?
Q3: For what N would the company maximize its profits?

We break down their analysis into finding the "Optimal Incentive" (Q3), "Better Than No Incentive" (Q2) and "Safe Incentive." "Optimal Incentive" is defined as the incentive which maximizes profits. Sometimes, having no incentive is the optimal option. "Better Than No Incentive" is the incentive in which the revenue from referred purchases due to incentivized shares is higher than the reward cost. "Safe Incentive" is the incentive in which referred purchase revenue due to both organic and incentivized shares is higher than the reward cost. This analysis allows a company to bound the incentive cost by referred purchase revenue. Table I exemplifies the output from our experiments. Its entries are determined by simulating our proposed *Me+N model* under various assumptions for the profit margin and % incentivized shares.

III. Notation and Problem Statement

Let $G = \langle U, D, P, S \rangle$ refer to a graph, where U refers to the users who participate in sharing (or recommending) items, D is the set of items that can be purchased and shared, $P = U \times D$ are the item purchases by users. $S = \langle U \times U \times D \rangle$ are the item shares between users. Each 3-tuple s is ordered in terms of the share *sender* or *origin* (o) and *recipient* (r) of an item d. The set of all recipients with whom user o shared item d is $R(o, d) = \{r : \langle o, r, d\rangle \in S\}$, and the share *outdegree* is the set size, $outdegree(o, d) = |R(o, d)|$. The outdegree reflects the sharing behavior of a given sender over one given item, rather than aggregating across different shared items.

Every item d has a set of properties, such as price, $price_d$, and a profit margin, $margin_d$ (i.e., a percentage of the item price that the company keeps). Every user has a set of properties such as their arrival time into the network, purchasing and sharing preferences. Every incentive $i \in I$ is defined by two properties: a *goal* (i.e., an objective that a sharer has to reach) and a *reward* (i.e., the benefit to the sharer if he reaches the goal). The $cost_{u,d}(i)$ is the incentive cost to the company when a sharer o shares item d, achieves the incentive goal and must be given the reward. Ideally, the company picks an incentive $i^* \in I$ for a item d that maximizes profit over the set of users: $profit_d(i) = \sum_{u \in U(d)} profit_{u,d}(i)$ where,

$$profit_{u,d}(i) = \begin{cases} -cost_{u,d}(i) & \text{if } u \text{ reached } goal_i \\ margin_d \cdot price_d & \text{else.} \end{cases}$$

In practice, a company must search for an incentive that is "Better Than No Incentive."

Consider two disjoint sets of users: $A \subseteq U$ are *altruistic* users unaffected by an incentive, and $V = U \setminus A$, the set of *incentivized* users who share with the hope of achieving the incentive goal. The incentivized users are responsible for any additional profits that the company would receive from a given incentive. In contrast, the altruistic users are unaffected by the incentive but could achieve the incentive goal anyways, meaning the company incurs a cost for organic sharing. Let α_i represent the percentage of additional profits produced due to an incentive i. Our goal is to lower bound α_i for both the "Better than No Incentive" and "Safe Incentive" values.

IV. Observed behavior from Me+3

LivingSocial is an e-commerce company whose core business is about connecting customers to local merchants, such as restaurants and beauty salons, through online deals. LivingSocial *incentivizes* its users to share deals through a *Me+3* program. To participate in the program, a user must first purchase a deal and then share it with their friends through email, Facebook, or Twitter. The participant can obtain the purchased deal for free if she convinces 3 (or more) friends to purchase the same deal.

In [8], we showed that the *Me+3* program has an unusual effect on the *structure* of the resulting social network, distinguishing it from other social networks. We compared shares of users who have not purchased the deal before sharing it, *non-incentivized shares* with shares of users who have, *incentivized shares*. In the first case, the shares are inherently *altruistic* because they cannot be a part of the *Me+3* program. The outdegree distribution of non-incentivized shares fits a power

Algorithm 1 DailyActivity(days,θ)

```
1: nodes = [], purchases=[], incent=[]
2: profit = 0
3: for all day in days do
4:     arr_n = Arr(θ), awake_n = Awake(θ)
5:     for all node in arr_n do
6:         incent[node] = Incent(θ)
7:     end for
8:     daily_n = arr_n + awake_n
9:     profit += θ.purch_profit * len(daily_n)
10:    sp, sl = Share(daily_n, incent, day, θ)
11:    profit += (sp - sl)
12: end for
13: return profit
```

Algorithm 2 Share(daily_n,incent,θ)

```
1: # Initialize
2: profit=0, loss = 0,purchases = []
3: # Determine sharing for each node
4: for all node in dailynodes do
5:     r = Share(incent[node], θ)
6:     # Store purchases from shared
7:     r_p = RPrch(r, θ)
8:     profit += θ.purchase_profit * len(r_p)
9:     loss += Cost(r_p, θ)
10: end for
11: return profit, loss
```

law, mirroring the distributions found in many other social networks. In contrast, the degree distribution of incentivized shares has a noticeable 'dip' at outdegrees 1 and 2. Beginning at outdegree 3, the distribution of shares is *shifted* due to the incentive. Purchasers who share with one or two people ($outdegree \in [1, 2]$) are clearly *altruistic*, while users who share with three or more can reach the incentive even if this was not their primary reason for sharing.

Here, we examine additional properties. Figure 1.a shows that as the sender increases the number of recipients, he becomes less discriminative in choosing good recipients. Figure 1.b examines this in more detail by showing the distribution of successful recommendations conditioned on the number of recipients. A (truncated) χ^2 distribution provides the best fit for this distribution. Figure 1.c shows that the more the recipients, the higher the likelihood of a free deal. Figure 1.d shows the lag between purchases (scaled to remove proprietary information). The χ^2 and log-normal distributions empirically match the observations best.

V. Me+N: A Graph Evolution Model for Incentivized Social Sharing

While we can examine the effectiveness of an incentive in place, it is hard to gauge how other incentives would affect company's profits. To solve this problem, we propose a generalized *Me+N* graph evolution model which can simulate user behavior under a *Me+N* incentive. The model is designed to be *modular*, in order to test and compare different hypotheses with respect to incentivized sharing. Algorithm 1 lays out the model with input θ, the set of incentive parameters. The main *for loop* (Lines 3 through 12) captures the *daily* sharing behavior of customers through the following functions:

- *Arrival*: How new users *join* the network

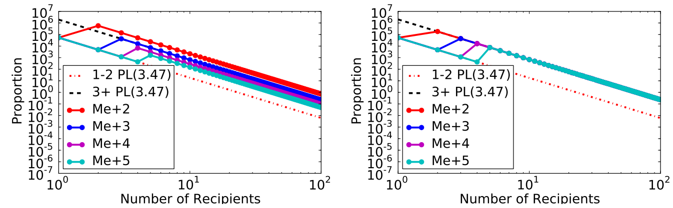

(a) (b)

Fig. 2: Two hypothetical distributions of incentivized sharing.

- *Awakening*: How long users wait between purchases
- *Sharing*: How users share
- *Recipient Purchasing*: How recipients choose to purchase
- *Profits*: How the incentive affects the overall profits

These functions can have a significant impact on the profit, and there are many nuances to consider (see Section II). To simplify this, we focus on the three main parameters: 1) $\theta.N$ (over which the company has control), 2) the percentage of additional profits $\theta.\alpha$ (which is governed by user behavior), and 3) the profit margin $\theta.margin$ (which can be supplied as input by the company).

Arrival and awakening functions: the *Arrival* and *Awakening* functions which we discussed in the previous section are explicitly incorporated in the graph model (line 4 of Algorithm 1). Arrival functions in social networks can have from exponential to sub-linear growth [9]. Two other options are the Gaussian such as the LivingSocial case, or the Poisson which are frequently used to model arrival rates [10]. Other options are to incorporate seasonal or day-of-week effects into the arrival function. As shown in Section IV, the awakening function of LivingSocial users follows a χ^2. Another possibility for the awakening function would be to awake users who are likely to purchase an offered product: $Awake(u, d, \theta) = P(u|d, \theta)$. This function would be useful if we the company has a *preference model* for each user.

Creating a *Me+N* sharing distribution: we build upon the observation that the *Me+3* incentive causes a shift in the degree distribution of user shares, resulting in nearly identical exponents for both non-incentivized and incentivized shares. Our model assumes that modifications of the incentive threshold *N* in *Me+N* will result in a similar distribution shift for the incentivized shares. While the shift can be characterized as a constant, it is unclear how large this shift (i.e., the proportion of incentivized users) becomes as the incentive parameters change. Nevertheless, it is safe to assume that as *N* increases (and the hurdle for attaining the reward gets higher), the overall percent of additional shares would decrease.

We generate a shift in an incentivized distribution using two different hypotheses. Figure 2.a shows the first one where the red dash-dot line represents the non-incentivized distribution and the distributions for various values of *N*. As *N* increases, the distribution lies closer to the original non-incentivized line, until the two converge. The volume under the dashed line corresponds to the *altruistic* sharing of individuals, while the area above the dashed red line (but under the solid line) corresponds to the additional shares resulting from the

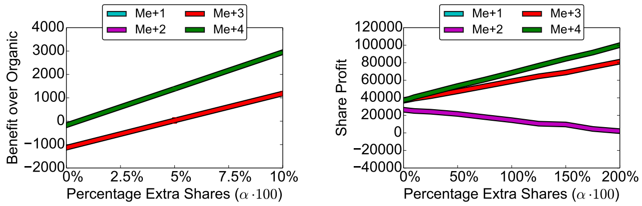

(a) (b)

Fig. 3: (a) For small margins, *Me+3* can have a negative impact. (b) However, *Me+3* is a 'safe incentive'.

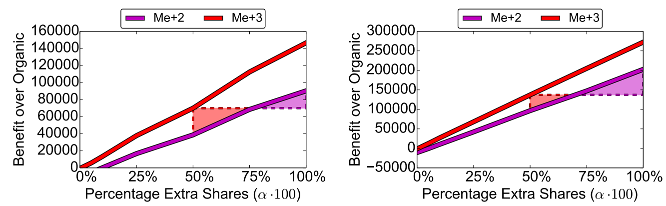

(a) (b)

Fig. 4: Lower bound of *Me+2* with *Me+3* ($\alpha_3 = .5$), for (a) Margin = .5 and (b) Margin = .9.

incentive. An alternative hypothesis is that sharing does not change the proportion of additional shares at each outdegree as we increase *N* (see Figure 2.b). Here, additional shares still decrease as *N* increases. To determine which users share because of the incentive and which ones share in an altruistic manner, we use $S \sim Bernoulli(\theta.p)$ to indicate whether a purchaser shares altruistically.

Cost of sharing incentivies: As described in Section III, the last piece that we need to define, in order to be able to calculate profits under different scenarios, is the cost to the company when a sharer achieves the incentive goal and must be given the reward. The cost depends on the reward, and it can be a percent of the item price: $cost_{u,d}(i) = \mathbb{I}[\|r_p\| \geq \theta.N] \cdot \$\theta.reward,$ where $\mathbb{I}$ is the indicator and $\theta.N$ is the incentive threshold. The cost can also be fixed (e.g., \$10).

VI. Experiments

We use the *Me+N model* to generate different networks and measure the *benefit over organic* sharing: $profit_d(i) - rev^A$. When this measure is positive, the incentive is "Better Than No Incentive." We fix the price at \$50 and simulate 50 days with 1000 new customers arriving to the site each day. The awakening function is based on $\chi^2(50)$ distribution.

Variable Profit Margins: To assess the effect of profit margins, we vary the proportion of additional shares α for different N. We find that for a *low* profit margin (i.e., a free product has a high cost), methods such as *Me+1* and *Me+2* are *worse* than not having incentives. In higher margin areas, *Me+2* generally does much better: for both margin = .5 and margin = .9, *Me+2* performs comparably to *Me+3* and *Me+4*. While *Me+1* improves, it does not become profitable, even for the high margin products.

Safe Incentives: The profitability of $Me + N$ highly depends on α. Since true value of α for any given incentive is unknown, at least we would like to check whether it is a "Safe Incentive." For these experiments, we fix the profit margin at 0.1. Figure 3.a shows that for $\alpha \geq 0.05$, *Me+3* is "Better Than No Incentive." Moreover, when we consider the organic sharing profits (Figure 3.b), we see that *Me+3* is above 0, making it a "Safe Incentive" at any level of $\alpha > 0$.

Performance of *Me+3*: So far, we assumed that different N models have the same α. However, due to the higher hurdle of higher N, it is more likely that $\alpha_N > \alpha_{N+1}$. To examine these cases, we analyze the company's *benefit over organic* for different incentives by fixing α for one of the incentives. For example, if we fix $\alpha_3 = 0.5$, we ask how large should the extra percentage of shares α_N be in order to beat *Me+3*? In Figure 4.a, we compare *Me+2* and *Me+3* at margin 0.5. The red dashed line marks the benefit of *Me+3* for $\alpha_3 = .5$. It crosses the purple *Me+2* line at approximately $\alpha_2 = .78$. This means that *Me+2* would outperform *Me+3* only if $\alpha_2 > \alpha_3 + 28\%$. At a margin of 0.9, this will hold true only if $\alpha_2 > \alpha_3 + 19\%$ (Figure 4.b). Thus, lower N is more likely to be better under higher margins. When comparing *Me+3* with *Me+4*, we find that for low margins, 4 is more likely to provide higher benefit.

VII. Conclusions

We proposed a novel problem for optimizing mass-marketing sharing incentives, and we introduced a framework for understanding when and why an incentive would be profitable to a company. We presented empirical evidence of the value of such incentives by using data from an e-commerce company. We introduced a probabilistic, graph-evolution model to simulate incentivized user behavior, which is flexible enough to accommodate multiple incentive types.

References

[1] N. Agarwal, H. Liu, L. Tang, and P. S. Yu, "Identifying the influential bloggers in a community," *WSDM*, 2008.
[2] S. Aral and D. Walker, "Identifying influential and susceptible members of social networks." *Science*, 2012.
[3] W. Chen, Y. Wang, and S. Yang, "Efficient influence maximization in social networks," *KDD*, 2009.
[4] P. Domingos and M. Richardson, "Mining the network value of customers," *KDD*, 2001.
[5] D. Kempe, J. Kleinberg, and E. Tardos, "Maximizing the spread of influence through a social network," *KDD*, 2003.
[6] E. Bakshy, J. M. Hofman, W. A. Mason, and D. J. Watts, "Everyone's an influencer: Quantifying influence on twitter," *WSDM*, 2011.
[7] M. Salganik, P. Dodds, and D. Watts, "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market," *Science*, vol. 311, no. 5762, 2006.
[8] J. J. Pfeiffer III and E. Zheleva, "Incentivized sharing in social networks," *WOSS*, 2012.
[9] J. Leskovec, L. Backstrom, R. Kumar, and A. Tomkins, "Microscopic evolution of social networks," *KDD*, 2008.
[10] R. Durrett, *Essentials of Stochastic Processes*, ser. Springer texts in statistics. Springer, 2012.