Maximizing Revenue from Strategic Recommendations under Decaying Trust[*]

Paul Dütting
EPFL Lausanne
Lausanne, Switzerland
paul.duetting@epfl.ch

Monika Henzinger
University of Vienna
Vienna, Austria
monika.henzinger@univie.ac.at

Ingmar Weber
Yahoo! Research
Barcelona, Spain
ingmar@yahoo-inc.com

ABSTRACT

Suppose your sole interest in recommending a product to me is to maximize the amount paid to you by the seller for a sequence of recommendations. How should you recommend optimally if I become more inclined to ignore you with each irrelevant recommendation you make? Finding an answer to this question is a key challenge in all forms of marketing that rely on and explore social ties; ranging from personal recommendations to viral marketing.

We prove that even if the recommendee regains her initial trust on each successful recommendation, the expected revenue the recommender can make over an infinite period due to payments by the seller is bounded. This can only be overcome when the recommendee also incrementally regains trust during periods without any recommendation. Here, we see a connection to "banner blindness", suggesting that showing fewer ads can lead to a higher long-term revenue.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics; H.1.2 [**Models and Principles**]: User/Machine Systems—*Human factors*

Keywords

recommendations, trust loss in advertising, banner blindness

1. INTRODUCTION

If you recommend a book at Amazon solely due to the monetary incentive given by Amazon's referral scheme[1] and your friends realize this, then they are likely to lose trust in your recommendations. If they regain trust whenever you make a relevant recommendation, how can you maximize your long-term revenue, and is this revenue bounded or not? The importance of "word-of-mouth" (WOM) for marketing has long been known [4, 10, 9]. According to [4], "WOM was seven times as effective as newspapers and magazines, four times as effective as personal selling, and twice as effective as radio advertising in influencing consumers to switch brands". WOM is the causal effect behind "brand congruence" where friends tend to use the same products [13, 15].

In our scenario the recommender is selfish and only makes paid recommendations to maximize her own revenue. Here, in a sense, the friend making the recommendation is no more trustworthy or altruistic than a web search engine showing sponsored search results. In these settings we believe the trust between the recommender and the recommendee to be dissipating. We see this as closely related to "banner blindness" [3, 5], where people have become so overloaded and fed up with advertisement that they stop to notice it completely. Seen from this angle, our findings show that advertisers might have to *stop* showing advertisements on a regular basis if they want to retain customers' trust without seeing click-through-rates converge to zero.

1.1 Related Work

In a recent paper truthful constant-factor approximation mechanisms for the so-called influence maximization problem were given [14]. In this problem a subset of early adopters has to be selected and incentivized to influence their peers in a given social network. Our problem is different as we consider personal recommendations and explicitly model the relationship between the recommender and the recommendee.

In our model the recommendee loses trust for each unsuccessful recommendation. This is more likely when she has the feeling that the recommendations are "dishonest". How honest recommendations can be ensured when there are several recommenders is studied in [8]. The approach suggested by the authors involves evaluating/ranking recommenders based on the rating given to their recommended items by other people. This, however, requires a public market where potential buyers can look for recommendations, which is not the case for personal recommendations.

The problem of trust decay is related to "banner blindness" [3, 6, 5] where web users become "blind" to ads due to overexposure. Cast to this setting our mathematical model suggests that, even if web users' interest is "refreshed" by a single relevant, clicked advertisement, the long term revenue of advertisers will stagnate as click-through-rates fall to zero. The only solution is to *stop* showing banner ads for a while so that users can "unlearn" to ignore all advertising. This approach is also suggested in a recent patent [12].

In typical literature on sponsored search auctions [11] it is assumed that the web search engine's expected revenue for

[*]This research was supported by a Google Research Award and the Vienna Science and Technology Fund (WWTF) grant ICT 10-002.

[1]`https://affiliate-program.amazon.com/`

CIKM'12, October 29–November 2, 2012, Maui, HI, USA.

showing a particular ad is the ad's click-through-rate (CTR) multiplied by the price the advertiser will be charged when her ad gets clicked. Usually, only a single round is considered or, when there are budget constraints [1, 7], the CTRs are assumed to be *constant* during the duration of the game. If, however, it is assumed that CTRs drop for *all* ads for each unsuccessful advertisement shown then, in the long run, this puts more emphasis on showing ads with high CTR, regardless of how much their advertisers can be charged for a single click. Although different objective functions for the search engine have been considered [1], the setting of revenue maximization with trust decay has not been studied.

1.2 Classification of Advertising Schemes

One could argue that a recommendation is, ultimately, just an advertisement and that an advertisement is just a recommendation. To highlight the differences between different kinds of advertisement in general, we present a classification scheme using the following three dimensions.

- Addressing: Personal vs. general. A recommendation is per se more personal than an advertisement and should be adapted to reflect the individual needs and interests of the potential buyer. Classic advertisement is not personalized and uses the same "message" for everyone.
- Trust: High vs. low. A recommendation should come from someone the potential buyer trusts and feels loyal or close to. This can be a personal friend or maybe a well-respected blogger. In classic advertisement the information source is viewed as less reputable, though advertisers try to use trusted icons for their purposes.
- Intention: Altruistic vs. commercial. The intention of a recommendation by a friend is generally not commercial. She might not get reimbursed at all but she still recommends something as she believes you would profit from it. In ordinary advertising the reason for the act of advertising itself is a commercial one.

To demonstrate the generality of this schema, we use it to classify a number of different advertising scenarios.

1. Billboards. A chain of pizza restaurants puts up billboards all over the country, without targeting any specific group. Addressing: general, trust: low, intention: commercial.
2. Sponsored search. A web search engine shows targeted sponsored results in addition to "organic" web search results, trying to match the searcher's intent. Addressing: personal, trust: low, intention: commercial.
3. Testimonial. You liked a book and you write a testimonial on Amazon to convince other unknown readers to read it, too. Addressing: general, trust: high, intention: altruistic.
4. Direct recommendation. A friend asks you for advice on which laptop to buy and you recommend the model which you believe is best for her. Addressing: personal, trust: high, intention: altruistic.

There are other important differences, e.g. concerning the conversion rates, but these differences are consequences of the "axiomatic" differences above and a personalized, altruistic "advertisement" from a highly trusted source will always have a higher conversion rate than a general, commercial "recommendation" from a disreputable source.

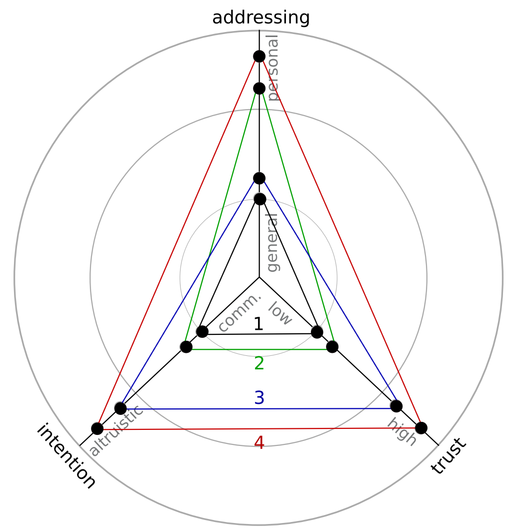

Figure 1: Visualization of the four advertising schemes discussed in the text. Direct recommendation (#4) is the most successful advertising medium as it dominates all other schemes in all dimensions.

1.3 Our Contributions and Outline

To the best of our knowledge there has been no work focusing on the strategic behavior of recommenders in a setting with decaying trust. We view the introduction of this problem as one of our contributions.

We first show that, not surprisingly, the total expected revenue of the recommender is bounded when the recommendee can only lose and does not regain trust (Section 2.2). Then we prove that the total expected revenue is still bounded over an infinite (!) sequence of recommendations, even when trust is reset to an initial level on each successful recommendation (Section: 2.3). Finally, we show that when trust is regained incrementally when no recommendations are made, the recommender's optimal total expected revenue is unbounded in the long run and that she can recommend both too aggressively and too passively (Section 2.4).

2. REVENUE MAXIMIZATION

We study the following problem: There are n products. For each product the recommender has two options: "recommend" or "*not* recommend". A recommendation is *successful* if the buyer buys the product. For a successful recommendation the recommender gets a constant reward of r and this reward is the same for all products. Initially, the probability p of success is $p_0 < 1$. With each unsuccessful recommendation this probability drops from its current value to $p = l \cdot p$, where $l < 1$ is the *loss rate*. The probability p can be seen as an estimate of the recommendee's *trust* in the recommender and a high value of l corresponds to a slow loss in trust. This basic model is analyzed in Section 2.2. We also consider extensions of this model where trust ($= p$) can increase again in two ways. First, we assume that p is reset to $p = p_0$ on each successful recommendation. This setting we refer to as "with reset" and it is analyzed in Section 2.3. Second, we introduce a factor $g \geq 1$ and each time the recommender does not recommend anything trust is regained and p is updated to $p = \min(g \cdot p, p_0)$. This setting we refer to as "with recovery" when $g > 1$ and it is analyzed in Section 2.4.

In all settings the recommender's sole goal is to maximize the overall expected reward $M_n(p_0, l, g)$ for the given parameters p_0, l and g. We are interested in the asymptotic behavior of $M_n(p_0, l, g)$, i.e. in $R(p_0, l, g) = \lim_{n\to\infty} M_n(p_0, l, g)$. Before looking at the theoretical analysis, the following section experimentally demonstrates the different behavior of the optimal total expected reward in these settings.

2.1 Experimental Results

Figure 2 gives experimental results for $n = 200$, $r = 1$, $p_0 = 0.5$, $l = 0.66$, $g = 1$ (in the setting "without recovery") and $g = 1.33$ (in the setting "with recovery"). It shows that the expected reward of the optimal strategy converges in the setting "without recovery" and diverges in the setting "with recovery". In the setting "without recovery" the expected reward converges to 2.25 if the probability of success is not reset and to 5 if it is reset to p_0 on a single successful recommendation. The figure also shows that the expected reward of the heuristic "recommend product 1, $k+1$, $2k+1$, etc." converges for $k = 2$ where $l \cdot g^k < 1$ and diverges for $k = 3$ and 4 where $l \cdot g^k > 1$. Finally, it shows that the expected reward grows faster for $k = 3$ than for $k = 4$.

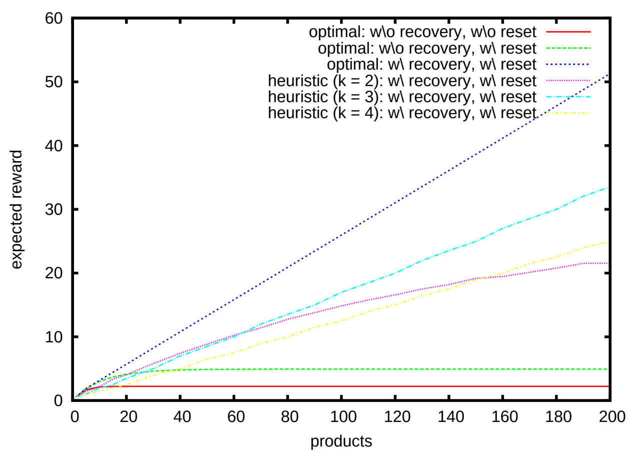

Figure 2: Without any recovery (red and green lines) the total expected reward converges. This also holds, even with recovery, for the "aggressive" heuristic (purple line) which recommends every second item. The other three settings, described in detail in Section 2.1, lead to an asymptotically unbounded reward.

2.2 Without Reset, without Recovery

Here we consider the case where $g = 1$ (= no recovery) and assume that the probability of success is *not* reset to p_0 on a successful recommendation. As the probability of success remains unchanged if no recommendation is given, the optimal strategy is to recommend *all* products. Therefore we can rewrite $R(p_0, l) := R(p_0, l, 1)$ as follows: $R(p_0, l) = p_0(r + R(p_0, l)) + (1 - p_0)R(lp_0, l)$, which we can solve analytically.

$$\begin{aligned} R(p_0, l) &= \frac{p_0}{1-p_0} r + R(p_0 \cdot l, l) \\ &= \frac{p_0}{1-p_0} r + \frac{p_0 \cdot l}{1-p_0} r + R(p_0 \cdot l^2, l) \\ &= \frac{p_0}{1-p_0} \cdot \sum_{i=0}^{\infty} l^i \cdot r = \frac{p_0}{1-p_0} \cdot \frac{1}{1-l} \cdot r < \infty. \end{aligned}$$

So, not surprisingly, if trust can only be lost and if both $p_0 < 1$ and $l < 1$, then the total expected reward the recommender can achieve is finite, even when there is an infinite sequence of items to recommend.

2.3 With Reset, without Recovery

Now let us analyze the case where still $g = 1$ (= no recovery) but each successful recommendation leads to reset of p to p_0. Again, the optimal strategy is to recommend *all* products as there is no gain from not recommending. In this setting, we can rewrite $R(p_0, l)$ as follows: $R(p_0, l) = (1 - q) \cdot (r + R(p_0, l))$, where q denotes the probability that there will be not a single successful recommendation over the infinite sequence. This recurrence can be solved (i.e. $\lim_{n\to\infty} M_n(p_0, l)$ is finite) if and only if $q > 0$.

Lemma 1. *Let* $dilog(x) = \int_x^1 \frac{\ln(t)}{1-t}\, dt$ *and* $c = \max(p_0, l)$. *Then, for all* $1 > p_0 \geq 0$, $q \geq (1 - c) exp\left(\frac{dilog(1-c)}{\ln(c)}\right) > 0$.

Proof. The probability that there will be not a single successful recommendation is:

$$q = \prod_{k=0}^{\infty} (1 - l^k \cdot p_0) \geq \prod_{k=0}^{\infty} (1 - c^{k+1}).$$

Hence it suffices to show that $\prod_{k=0}^{\infty}(1 - c^{k+1}) > 0$. Taking the $\ln(\,)$ of both sides we get

$$\ln\left(\prod_{k=0}^{\infty}(1 - c^{k+1})\right) = \sum_{k=0}^{\infty} \ln(1 - c^{k+1}) > -\infty,$$

where we need to prove this inequality. Note that the expression $\ln(1 - c^{k+1})$ is strictly increasing in k and hence $\ln(1 - c^{k+1}) \geq \int_{k-1}^{k} \ln(1 - c^{x+1})\, dx$. This gives the bound

$$\begin{aligned} \sum_{k=0}^{\infty} \ln(1 - c^{k+1}) &= \ln(1 - c) + \sum_{k=1}^{\infty} \ln(1 - c^{k+1}) \\ &\geq \ln(1 - c) + \sum_{k=1}^{\infty} \int_{x=k-1}^{k} \ln(1 - c^{x+1})\, dx \\ &= \ln(1 - c) + \sum_{k=0}^{\infty} \int_{x=k}^{k+1} \ln(1 - c^{x+1})\, dx \\ &= \ln(1 - c) + \int_{x=0}^{\infty} \ln(1 - c^{x+1})\, dx. \end{aligned}$$

Recall that $\text{dilog}(x) = \int_x^1 \frac{\ln(t)}{1-t}\, dt$. The indefinite integral of $\ln(1 - x)$ is $-\text{dilog}(1 - x)/\ln(x)$. We get

$$\int_{x=0}^{\infty} \ln(1 - c^{k+1})\, dx = \frac{\text{dilog}(1-c)}{\ln(c)} - \lim_{x\to\infty} \frac{\text{dilog}(1 - c^{x+1})}{\ln(c)}.$$

Since $\text{dilog}(x)$ is continuous and $\text{dilog}(1) = 0$ (see Lemma 2), we get

$$\int_{x=0}^{\infty} \ln(1 - c^{k+1})\, dx = \frac{\text{dilog}(1-c)}{\ln(c)} - \frac{\text{dilog}(1)}{\ln(c)} = \frac{\text{dilog}(1-c)}{\ln(c)}.$$

For $0 < x < 1$ we have $0 \leq \text{dilog}(1 - x) < 2e^{-1} + 1$ (see Lemma 2). For $0 < x < 1$ we have $\ln(x) < 0$. It follows that $\int_{x=0}^{\infty} \ln(1 - c^{k+1})\, dx > -\infty$. □

Lemma 2. *Let* $dilog(x) = \int_1^x \frac{\ln(t)}{1-t}\, dt$. *Then* $dilog(x)$ *is monotonously decreasing and* $0 \leq dilog(x) < 2e^{-1} + 1$.

In fact, the tight upper bound of $\text{dilog}(x) \leq \pi^2/6 < 2e^{-1} + 1$ is known [2], but we choose to give the following elementary proof of Lemma 2 to have a self-contained argument.

PROOF. Let $f(t) = -\ln(t)/(1-t)$. Then $f'(t) = -(\frac{1}{t}(1-t) + \ln(t))/(1-t)^2 < 0$ for $0 < t < 1$. So $\int_{t=x}^{1} f(t)dt < \int_{t=x}^{e^{-1}} f(t)dt + (1-e^{-1})f(e^{-1})$. For $0 < t \le e^{-1}$ we also have $f(t) \le -\ln(t)/(1-e^{-1})$. So, $\int_{t=x}^{e^{-1}} f(t)dt \le [t - t\cdot\ln(t)]_x^{e^{-1}}$. This is largest when $x \to 0$ where the whole expression becomes $2e^{-1}$ and so $\int_{t=x}^{1} f(t)dt < 2e^{-1}+1$ for $0 \le x < 1$. Note that $f(t)$ is continuous at $t = 1$ with $\lim_{x\to 1} f(t) = 1$ (using e.g. the l'Hopital Rule). So trivially dilog$(1) = 0$. As $f(t) > 0$ this gives the desired lower bound. □

Using Lemma 1 we can prove the following theorem.

THEOREM 1. *Let dilog(x)* $= \int_x^1 \frac{\ln(t)}{1-t}\,dt$, $c = \max(p_0, l)$, *and* $\delta(c) = (1-c)exp(dilog(1-c)/\ln(c))$. *Then, for all* $1 > p_0 \ge 0$,

$$R(p_0, l) \le \frac{1-\delta(c)}{\delta(c)} \cdot r < \infty.$$

This proves that even if the probability of success is reset to p_0 on a single successful recommendation, the total expected reward over an infinite period is bounded.

2.4 With Reset, with Recovery

Finally, we consider the setting where $g > 1$ (= with recovery). Here the probability of success is set to $\min(p_0, g\cdot p)$ if no recommendation was given. Hence it might be better not to recommend all products to avoid that that probability p converges to zero. Let $M_n(p_0, l, g)$ denote the expected reward of the optimal strategy. To obtain bounds for $M_n(p_0, l, g)$, let us consider, as a heuristic, the algorithm $A^{(k)}$ that recommends product 1, $k+1$, $2k+1$, etc. We write $A_n^{(k)}(p_0, l, g)$ to denote this algorithm's expected revenue.

THEOREM 2. *Let* ψ *be the smallest integer such that* $l \cdot g^{\psi} \ge 1$. *If* $k > \psi$, *then, for all* $1 > p_0, l > 0$ *and* $\infty > g \ge 1$,

$$A_n^{(k)}(p_0, l, g) = \lfloor \frac{n}{k} \rfloor \cdot p_0 \cdot r.$$

PROOF. The expected reward for the first recommendation is $p_0 \cdot r$. Since $k > \psi$, the expected reward for every other recommendation is also $\min(p_0, p_0 \cdot l \cdot g^{k-1}) = p_0 \cdot r$. Since there are exactly $\lfloor n/k \rfloor$ recommendations, this shows that $A_n^{(k)}(p_0, l, g) = \lfloor \frac{n}{k} \rfloor \cdot p_0 \cdot r$. □

This is instructive as it shows that (a) for $k > \psi$ the expected reward $A_n^{(k)}(p_0, l, g)$ of $A^{(k)}$ does *not* converge as n tends to infinity and (b) for $k' > k > \psi$ the expected reward $A_n^{(k')}(p_0, l, g)$ of $A^{(k')}$ grows slower (and is ultimately lower) than the expected reward $A_n^{(k)}(p_0, l, g)$ of $A^{(k)}$. Since the reward $M_n(p_0, l, g)$ of the optimal strategy is at least as high, this also shows non-convergence of $R(p_0, l, g) = \lim_{n\to\infty} M_n(p_0, l, g) = \infty$.

THEOREM 3. *Let* ψ *be the smallest integer such that* $l \cdot g^{\psi} \ge 1$. *If* $k \le \psi$, *then, for all* $1 > p_0, l > 0$ *and* $\infty > g \ge 1$, *there exist* p_0' *and* l' *such that*

$$\lim_{n\to\infty} A_n^{(k)}(p_0, l, g) \le \lim_{n\to\infty} M_n(p_0', l') < \infty.$$

PROOF. If $k < \psi$, then the revenue maximization problem with parameters p_0, l, and g on the products 1, 2, 3, etc. is equivalent to the revenue maximization problem with parameters $p_0' = p_0$, $l' = l \cdot g^{k-1} < 1$, and $g' = 1$ on the products 1, $k+1$, $2k+1$, etc. The claim follows from Theorem 1. □

Whereas Theorem 2 shows that recommending too seldomly is sub-optimal, Theorem 3 shows that recommending too often is even worse.

3. DISCUSSION AND FUTURE WORK

We believe that our work motivates a number of research questions. How exactly do web users respond to being shown irrelevant advertisements? Is it possible to revive their interest in banner ads? What are "optimal" auction mechanisms for sponsored search when the CTRs are non-constant and decay with each irrelevant advertisement being shown?

4. REFERENCES

[1] Z. Abrams, O. Mendelevitch, and J. Tomlin. Optimal delivery of sponsored search advertisements subject to budget constraints. In *EC*, pages 272–278, 2007.

[2] G. E. Andrews, R. Askey, and R. Roy. *Special functions*. Cambridge University Press, 2001.

[3] J. P. Benway and D. M. Lane. Banner blindness: Web searchers often miss "obvious" links. *ITG Newsletter*, 1(3), 1998.

[4] J. J. Brown and P. H. Reingen. Social ties and word-of-mouth referral behavior. *JCR*, 14(3):350–62, 1987.

[5] M. Burke, A. Hornof, E. Nilsen, and N. Gorman. High-cost banner blindness: Ads increase perceived workload, hinder visual search, and are forgotten. *TOCHI*, 12(4):423–445, 2005.

[6] P. Chatterjee, D. L. Hoffman, and T. P. Novak. Modeling the clickstream: Implications for web-based advertising efforts. *MktSci*, 22:520–541, 2000.

[7] J. Feldman, S. Muthukrishnan, M. Pal, and C. Stein. Budget optimization in search-based advertising auctions. In *EC*, pages 40–49, 2007.

[8] A. C. B. Garcia, M. Ekstrom, and H. Björnsson. Hyriwyg: leveraging personalization to elicit honest recommendations. In *EC*, pages 232–233, 2004.

[9] R. Grewal, T. W. Cline, and A. Davies. Early-entrant advantage, word-of-mouth communication, brand similarity, and the consumer decision-making process. *JCP*, 13(3):187–197, 2003.

[10] P. M. Herr, F. R. Kardes, and J. Kim. Effects of word-of-mouth and product-attribute information on persuasion: An accessibility-diagnosticity perspective. *JCR*, 17(4):454–462, 1991.

[11] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. Sponsored search auctions. In *AGT*, pages 699–716. Cambridge University Press, 2007.

[12] K. Patel and R. G. P. Kantamneni. Monetizing low value clickers. United States Patent Application 20080249854, 2008.

[13] P. H. Reingen, B. L. Foster, J. J. Brown, and S. B. Seidman. Brand congruence in interpersonal relations: A social network analysis. *JCR*, 11(3):771–783, 1984.

[14] Y. Singer. How to win friends and influence people, truthfully: influence maximization mechanisms for social networks. In *WSDM*, pages 733–742, 2012.

[15] A. Singla and I. Weber. Camera brand congruence in the flickr social graph. In *WSDM*, pages 252–261, 2009.