How Much Is Your Personal Recommendation Worth?*

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 you buy a new laptop and, simply because you like it so much, you recommend it to friends, encouraging them to purchase it as well. What would be an adequate price for the vendor of the laptop to pay for your recommendation?

Personal recommendations like this are of considerable commercial interest, but unlike in sponsored search auctions there can be no truthful prices. Despite this "lack of truthfulness" the vendor of the product might still decide to pay you for recommendation e.g. because she wants to (i) provide you with an additional incentive to actually recommend her or to (ii) increase your satisfaction and/or brand loyalty. This leads us to investigate a pricing scheme based on the *Shapley value* [5] that satisfies certain "axioms of fairness". We find that it is vulnerable to manipulations and show how to overcome these difficulties using the *anonymity-proof* Shapley value of [4].

Categories and Subject Descriptors

G.2.0 [**Mathematics of Computing**]: General

General Terms

Economics, Theory

Keywords

Shapley value, pricing mechanisms, recommendations

1. INTRODUCTION & RELATED WORK

Personal recommendation, or "word-of-mouth", is recognized as a highly effective marketing tool [2]. But despite its relevance the problem of determining the worth of a recommendation, and hence a "right" price for it has not been tackled yet. We present a simple model and show that (i) there can be no truthful pricing scheme unless the prices are all zero. We then (ii) use the *Shapley value* [5] to compute provably "fair" prices, but also find that (iii) these prices are vulnerable to manipulations unless we base our computations on the *anonymity-proof* Shapley value [4].

The work that is most closely related to our paper is [1]. There the authors study the sales price of an object as part of a viral marketing campaign, but they do *not* consider the problem of how the recommendation itself should be rewarded. In fact, they mention the problem of finding optimal "cashbacks" in settings where the nodes behave strategically as an open problem. It should be clear that we are *not* addressing the problem of *what* to recommend, a problem typically encountered by stores such as Amazon and usually solved using "collaborative filtering" techniques [3].

*This work was conducted as part of a EURYI scheme award. See `http://www.esf.org/euryi/` for details.

Copyright is held by the author/owner(s).
WWW 2010, April 26–30, 2010, Raleigh, North Carolina, USA.
ACM 978-1-60558-799-8/10/04.

2. A CLASSIFICATION SCHEME

To highlight the differences between different advertising schemes we present a simple classification scheme:

- Addressing: Personal vs. general.
- Trust: High vs. low.
- Intention: Altruistic vs. commercial.

To demonstrate the general applicability of this scheme, we use it to classify a number of different advertising scenarios: A (1) *billboard campaign*, e.g. for a chain of pizza restaurants, is not targeted, has low trust, and a commercial background. A web search engine showing (2) *sponsored search results* along with "organic" search results uses personal addressing, has low trust, and commercial intentions. If you write a (3) *testimonial* on Amazon to convince other customers to buy a book that you liked, then the addressing is general, the trust is high, and the intention is altruistic. If a friend asks you for advice on which laptop to buy and you (4) *recommend* the model which you believe is best for her, then this is personal, highly trusted, and altruistic.

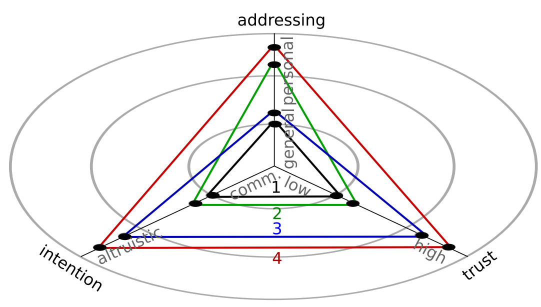

Figure 1: Recommendations are the most successful advertising medium as they dominate all other schemes in all dimensions.

3. THE COOPERATIVE MODEL

We model the pricing problem as a *coalitional game with transferable payoff* $\langle N, v\rangle$, where N is a finite set (the set of *players*) and v is a function that associates with every non-empty subset S of N (a *coalition*) a real number $v(S)$ (the *worth* of S). We use s to denote the *seller*, who is paying for recommendations, and r_i to denote the i-th recommender. There is exactly one *product* for sale, and a single *buyer*. We assume that neither the product nor the buyer belong to the set N, but we implicitly model the buyer's influence on the worth of a coalition by the purchase probability. For each coalition S the number $v(S)$ is the total payoff that is available for division among the members of S. We use $\delta \geq 0$ to denote the seller's *margin* or *gain* from selling the product, i.e. the sales price minus the production cost. We distinguish two scenarios for v:

- *Linear.* Without any recommendation the product is sold with probability $p \in [0, 1]$. The recommendation of the i-th recommender increases this probability by $q_i \in [0, 1 - \sum_{j \neq i} q_j]$. If the recommenders $R \subseteq \{r_1, .., r_n\}$ recommend the product, then the probability is $p + \sum_{i:r_i \in R} q_i$.
- *Threshold.* If less than k recommenders recommend the product, then the product is sold with probability $p \in [0, 1]$. If at least k recommenders recommend the product, then it is sold with probability $p + q$ where $q \in [0, 1 - p]$.

We refer to these scenarios as $\langle N, v\rangle$ (Linear) and $\langle N, v\rangle$ (Threshold). Our goal is to find a "fair" payoff vector $x = (x_s, x_{r_1}, \ldots x_{r_n})$, where x_s denotes the expected payoff to the seller and x_{r_i} denotes the expected payoff to the i-th recommender, i.e. the amount she is paid by the seller. Since such a payoff vector gives only expected payoffs, it can be translated into a *pricing scheme* in two ways:

1. *Pay-per-Recommendation*: The recommender gets paid by the seller for every recommendation; successful or not. That is, on every recommendation the seller s pays the i-th recommender r_i the money equivalent of x_{r_i}.
2. *Pay-per-Sale*: The recommender gets paid by the seller for successful recommendations only. That is, on every successful recommendation the seller s pays the i-th recommender r_i the money equivalent of $x_{r_i}/(p+f(\{s\}\cup R))$, where R is the set of recommenders.

In practice, the Pay-per-Sale approach might be preferable as, on a successful recommendation, one could reasonably assume $p + f(\{s\} \cup R) = 1$, sidestepping the problem of estimating $f(\{s\} \cup R)$ with very little or no data. Note that the prior probability p is easier to estimate using the seller's sales record and click-through or conversion-rates.

4. TRUTHFULNESS & FAIR PRICES

A payoff vector x is truthful w.r.t. the seller s, who holds the private information on p, f, and δ, if it gives her the incentive to reveal her information *truthfully.* Our first result is that any truthful payoff vector must have $\sum_i x_{r_i} = 0$, i.e., none of the recommenders gets paid. Intuitively, this is because the seller can always pretend to have benefited less from the recommendations then she actually did.

THEOREM 1. *There can be no truthful payoff vector* $x = (x_s, x_{r_1}, \ldots, x_{r_n})$ *that has* $\sum_i x_{r_i} \neq 0$.

To ensure high market share and long-term revenue it might still be beneficial for the seller to pay a "fair" price. The Shapley value is the unique value satisfying a certain number of "axioms of fairness" [5]. It is defined as follows: $\phi_i(v) = \sum_{S\subseteq N\setminus\{i\}}(|S|!(|N| - 1 - |S|)!)/(|N|!) \cdot (v(S \cup \{i\}) - v(S))$. Applying it to our model yields our second result.

THEOREM 2. *The Shapley value for the games (a)* $\langle N, v\rangle$ *(Linear) and (b)* $\langle N, v\rangle$ *(Threshold) is given by*

(a) $\phi_s(v) = p\delta + \frac{1}{2}\sum_i q_i\delta$ *and* $\phi_{r_i}(v) = \frac{1}{2}q_i\delta$, *and*

(b) $\phi_s(v) = p\delta + \left(1 - n\frac{k!(n-k)!}{(n+1)!}\right) q\delta$ *and* $\phi_{r_i}(v) = \frac{k!(n-k)!}{(n+1)!}q\delta$.

Thus, according to the Shapley value, each individual recommender in the game $\langle N, v\rangle$ (Linear) should receive a share of exactly one half of her contribution to the expected "extra profit" of the recommender . In the game $\langle N, v\rangle$ (Threshold) the fraction $k!(n-k)!/(n+1)!$ is exactly the fraction of times where this recommender's recommendation "makes a difference". These payoffs are fair in the sense that the payoff to each recommender is proportional to the recommender's contribution to the "extra profit" the seller can expect.

5. AN ISSUE & A WAY OUT

What would be a "fair" payoff vector if each recommendation was a *collection of arguments*? A straightforward approach would be to compute the Shapley value on the basis of arguments and to redeem recommender r_i with $\sum_{a\in S_i} \phi_a(v)$, where a is an argument from the set of arguments A and S_i is the set of arguments that recommender r_i possesses. The problem with this approach, however, is that it might be beneficial for a recommender to withhold some of her arguments:

- *Example 1.* Let $A = \{a, b, c\}$, $v(\{a, b\}) = v(\{a, c\}) = v(\{a, b, c\}) = 1$, and $v(\{a\}) = v(\{b\}) = v(\{c\}) = v(\{b, c\}) = 0$. Let $S_1 = \{a\}$ and $S_2 = \{b, c\}$. Then r_1 gets $\phi_a(v) = \frac{1}{2}$ and r_2 gets $\phi_b(v) + \phi_c(v) = \frac{1}{6} + \frac{1}{6} = \frac{1}{3}$.
- *Example 2.* Let $A' = \{a, b\}$, $v(\{a, b\}) = 1$, and $v(\{a\}) = v(\{b\}) = 0$. Let $S_1 = \{a\}$ and $S_2 = \{b\}$. Then r_1 gets $\phi_a(v) = \frac{1}{2}$, r_2 gets $\phi_b(v) = \frac{1}{2}$, and r_2 would be better off.

The *anonymity-proof Shapley value* [4] cannot be "tricked" in this way. It is defined as follows: For any set $A' \subseteq A$ of declared arguments the anonymity-proof Shapley value $\psi_a(v)$ for $a \in A'$ is: $\psi_a(v) = (\phi_a(v))/(\sum_{a'\in A'} \phi_{a'}(v))v(A')$. And so it would be better to (i) compute the value $\psi_a(v)$ for each argument $a \in A'$ and to (ii) give each recommender $\sum_{a\in S_i} \psi_a(v)$. With this approach r_1 and r_2 would get $\psi_a = 3/5$ and $\psi_b(v) + \psi_c(v) = 2/5$ in Example 1 and $\psi_a = 3/4$ and $\psi_b(v) = 1/4 < 2/5$ in Example 2.

6. FUTURE WORK

An interesting direction for future work would be to analyze the effect of pricing mechanisms - such as the ones discussed here - on the strategic behavior of recommenders.

7. REFERENCES

[1] D. Arthur, R. Motwani, A. Sharma, and Y. Xu. Pricing strategies for viral marketing on social networks. In *WINE'09*, pages 101–112, 2009.

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

[3] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*, 22(1):5–53, 2004.

[4] N. Ohta, V. Conitzer, Y. Satoh, A. Iwasaki, and M. Yokoo. Anonymity-proof shapley value. *Autonomous Agents and Multiagent Systems*, pages 927–934, 2008.

[5] L. S. Shapley. *Contributions to the Theory of Games II*, chapter A Value for n-Person Games, pages 307–317. Princeton University Press, 1953.