

# On the Pricing of Recommendations and Recommending Strategically\*

Paul Dütting<sup>†,1</sup>, Monika Henzinger<sup>‡,2,3</sup>, and Ingmar Weber<sup>§,4</sup>

<sup>1</sup>Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

<sup>2</sup>Google Switzerland, Zürich, Switzerland

<sup>3</sup>University of Vienna, Vienna, Austria

<sup>4</sup>Yahoo! Research, Barcelona, Spain

November 9, 2009

## Abstract

If you recommend a product to me and I buy it, how much should you be paid by the seller? And if your sole interest 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 these questions is a key challenge in all forms of marketing that rely on and explore social ties; ranging from personal recommendations to viral marketing.

In the first part of this paper, we show that there can be no pricing mechanism that is “truthful” with respect to the seller, and we use solution concepts from coalitional game theory, namely the Core, the Shapley Value, and the Nash Bargaining Solution, to derive provably “fair” prices for settings with one or multiple recommenders. We then investigate pricing mechanisms for the setting where recommenders have different “purchase arguments”. Here we show that it might be beneficial for the recommenders to withhold some of their arguments, unless anonymity-proof solution concepts, such as the anonymity-proof Shapley value, are used.

In the second part of this paper, we analyze the setting where the recommendee loses trust in the recommender for each irrelevant recommendation. Here we prove that even if the recommendee regains her initial trust on each successful recommendation, the expected total profit the recommender can make over an infinite period is bounded. This can only be overcome when the recommendee also incrementally regains trust during periods without any recommendation. Here, we see an interesting connection to “banner blindness”, suggesting that showing fewer ads can lead to a higher long-term profit.

**Keywords:** recommendations, pricing mechanisms, trust loss in advertising, banner blindness

---

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

<sup>†</sup>Email: [paul.duetting@epfl.ch](mailto:paul.duetting@epfl.ch)

<sup>‡</sup>Email: [monika.henzinger@univie.ac.at](mailto:monika.henzinger@univie.ac.at)

<sup>§</sup>Email: [ingmar@yahoo-inc.com](mailto:ingmar@yahoo-inc.com)

## 1 Introduction

Suppose you buy a new mobile phone and, simply because you like it so much, you recommend it to friends, encouraging them to purchase it as well. Even if you do not recommend it out of monetary reasons, what would be an adequate and *fair* price for the phone manufacturer to pay for your recommendation? If, on the other hand, you recommend a book at Amazon solely due to the monetary incentive given by Amazon’s referral scheme<sup>1</sup> and your friends realize this, then they are likely to lose trust in your recommendations. Assuming your friends regain trust whenever you make a relevant recommendation, how can you maximize your long-term profit, and is this profit bounded or not? These are the two main research questions we address in this work.

The importance of “word-of-mouth” (WOM) as a marketing channel has long been known [7, 18, 16]. According to [7], “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 both in offline [31] and online [36] social networks tend to use the same products. Recently, a platform called *Friend Vouch* [13] was founded around the idea of personal recommendations. Users of the service can become “brand ambassadors” who get paid for putting companies in touch with friends. Whether any personal touch is retained in such a system or whether the person in the middle is not simply another marketing channel is debatable and in Section 1.2 we propose a classification schema to shed light on the differences.

As far as the pricing of recommendations is concerned, one could argue that honest recommendations should always be given without any monetary recompensation and that creating financial incentives could lead to a sell-out of friends. Although this is a valid concern, we argue that it might still be worth paying recommenders, even if these are not asking to be paid. First, even though you might not be profit-maximizing in a strict sense you are probably more inclined to mention a certain product if there is some kind of recompensation: you might be honest enough not to recommend a bad product over a good one for financial reasons, but you

---

<sup>1</sup><https://affiliate-program.amazon.com/>

are still more likely to recommend a good product if you get reimbursed. Second, a *fair* compensation can lead to increased brand loyalty. If you are already satisfied with the product then the feeling that the company recompensates you in a fair and adequate manner is likely to increase your positive attitude towards the company. On the other hand, if you are only offered \$1 for recommending a particular type of sports car then this might be viewed as “offensive” and is arguably worse than not being offered any recompensation.

Especially, as the issue of trust is of utmost relevance in the realm of personal recommendations, we believe that “fair pricing” is a cornerstone [32, 24].<sup>2</sup> Google e.g. advertises its AdSense program by claiming to use a Second-Price Auction to eliminate “that feeling that you’ve paid too much”<sup>3</sup>. In a similar spirit, our work on pricing recommendations can be viewed as trying to eliminate “the feeling that you’ve been paid too little” for your recommendation.

The second problem we study, relates to scenarios where the recommender is selfish and only makes paid recommendations to maximize her own profit. 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. More concretely we assume that with every unsuccessful recommendation the recommendee becomes more and more likely to ignore any “advice” given by her friend. We see this as closely related to “banner blindness” [4, 9, 8], where people have become so overloaded and fed up with advertisement that they stop to notice it completely. Seen from this angle, our findings indicate 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

Even though recommendations can be seen as just another form of advertising, classical methods for the pricing of advertising, such as sponsored search auctions [21], are not directly applicable. This is mainly due to the fact that a true recommendation should be altruistic and so (i) the recommender is not profit maximizing and (ii) there is only a *single* seller, as an altruistic recommender will not accept “bids” from multiple sellers. The differences between various kinds of advertising are described in Section 1.2.

The work that is most closely related to our paper is [3]. There the authors study the sales price of an object as part of a viral marketing campaign. They assume that all “converted” nodes will try to convert all of their neighbors and that the conversion probability depends both on the number of neighbors converted and on the sales price. 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.

<sup>2</sup>Somewhat related is the phenomenon of pay-what-you-like pricing where people act “irrationally” and choose to pay an adequate amount [23, 12].

<sup>3</sup><http://www.google.com/adsense/afs.pdf>

The problem of optimal pricing with non-social recommender systems, where the recommendations directly come from the potential seller, was studied in [5]. Here by “non-social” we mean “computer-generated” and a typical example would be Amazon’s “Customers who bought X also bought Y”<sup>4</sup>. The somewhat surprising argument is that customers are willing to *pay* for relevant recommendations as they create “value by reducing product uncertainty for the customers”. In this paper, we consider the case where the recommendations are social and do not come from the seller directly. Though it is imaginable that the recommendee pays the recommender for a good recommendation, we do not investigate the pricing of this possible payment.

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 [35, 17]. In the first part of this paper (Section 2), we assume that the recommender recommends an item because she believes this item to be of interest to the recommendee, and the algorithm used by her to determine potential interest is irrelevant. In the second part (Section 3), the recommender is profit maximizing and now only cares about the reward offered to her by the seller and the probability  $p$  that the recommendee will buy the item. In this model the “what” is absorbed into  $p$  and the recommender simply decides on *when* to recommend.

We are also *not* addressing the topic of how rumors spread through social networks, or how to identify the best nodes to target for a viral marketing campaign [10, 19]. Our work focuses on a single atomic link in the corresponding cascades of conversions and, in the first part, we ask what a fair price should be to pay a node for activating one of her neighbors. In answering this question we limit our attention to the immediate profit of the seller due to the individual sale, and we do *not* consider the additional value due to recommendation cascades caused by the newly activated node. However, given any algorithm to compute this “higher order” profit, it can trivially be incorporated into our results. The question whether a selfish node should actually try to activate her neighbors at all is addressed in Section 3.

More generally, in the second part we look at a model where the recommendee loses trust in the recommender, i.e. for each unsuccessful recommendation she becomes less and less inclined to listen to any further suggestion. This is most likely to appear when the recommendee has the feeling that the recommendations are “dishonest”. How honest recommendations can be ensured when there are several recommenders is studied in [14]. The approach suggested by the authors involves evaluating/ranking recommenders based on the rating given to their recommended items by other people. This motivates recommenders to give good recommendations in a similar way that Ebay’s rating system gives incentives for both buyers and sellers “to behave”. This approach, however, requires a public market where potential buyers can look for recommendations. This is not the setting of personal recommendations considered here.

The problem of trust decay is related to “banner blindness” [4, 9, 8], where web users become “blind” to banner ads

<sup>4</sup><http://www.amazon.com>

due to overexposure. Cast to this setting our mathematical model suggests that, even if web users’ interest is “refreshed” by a single relevant advertisement that is clicked, the long term profit of advertisers will stagnate as click-through-rates fall to zero. The only possible way out of this dilemma 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 [30].

In typical literature on sponsored search auctions [21, 20] it is assumed that the web search engine is optimizing its expected profit and that its expected profit for 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, 11], 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 profit maximization with trust decay has not been studied and we deem this an interesting area for future work.

Finally, there is previous work that is relevant on a more technical level. In particular solution concepts such as the Core [15, 26] or the Shapley-Value [33, 26] have been studied extensively before. The exact connection to this group of work will be made clear in the sections with our technical contributions.

## 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 simple classification scheme.

- 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.

The first part of this paper (Section 2) considers the setting of personal, highly trusted and altruistic recommendations. The second part (Section 3) then investigates the case of still

personal, but commercial recommendations with a decaying amount of trust involved. To demonstrate the general applicability 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.

Of course, there are lots of other important differences, e.g. concerning the conversion rates, but we view these differences as consequences of the “axiomatic” differences above and we assume that 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.

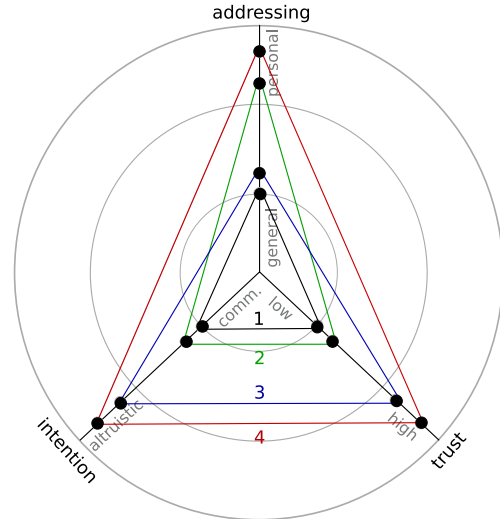


Figure 1: Visualization of the four advertising schemes discussed in the text. Direct recommendation (#4) is the most successful advertising medium as it dominates all other schemes in all dimensions.

## 1.3 Our Contributions and Outline

To the best of our knowledge there has been no work focusing on either (i) the pricing of recommendations (our Section 2) or (ii) the strategic behavior of recommenders in a setting with decaying trust (our Section 3). We view the introduction of these problems as one of our contributions.

As far as the pricing of recommendations is concerned we prove that there can be no pricing mechanism that is “truthful” with respect to the seller (Section 2.1). This shows that the seller can always pretend to profit less from the recommendations than she actually does to get a larger piece of the pie. We then apply solution concepts from coalitional game theory, namely the Core, the Shapley value, and the Nash Bargaining Solution, to determine provably “fair” prices. For the Core we find that it typically contains all “individual rational” payoff vectors, including the payoff vector where the seller gets everything and the recommenders get nothing. On the one hand, this demonstrates the weakness of the recommenders: They cannot form a coalition with non-zero value without the seller. On the other hand, it shows that the Core is essentially useless for deciding how to distribute the “extra profit” the seller can expect from being recommended among the recommenders (Sections 2.2.1 and 2.3.1). For the Shapley value we find that it not only defines unique prices, but that these prices are also “fair” in a very intuitive way: The price of a recommendation should be proportional to the “extra profit” the seller can expect from it (Section 2.2.2 and 2.3.2). For the Nash Bargaining Solution we find that it yields “fair” prices, namely those obtained by the Shapley value, only if there is a single recommender. Otherwise, especially in situations where the recommenders do *not* contribute equally to the “extra profit” of the seller, it may lead to “unfair” prices (Section 2.2.3 and 2.3.4). Finally, we also consider the case where each recommendation consists of one or more “purchase arguments”. Here the ordinary Shapley value is no longer the method of choice, as withholding arguments might be beneficial for the recommenders. We show how the anonymity-proof Shapley value from [28] can be applied to overcome this problem (Section 2.3.3).

In the second part on the strategic behavior of profit maximizing recommenders we first show that, not surprisingly, the total expected profit of the recommender is bounded when the recommendee can only lose and does not regain trust (Section 3.2). Then we prove that the total expected profit is still bounded over an infinite (!) sequence of recommendations, even when trust is reset to an initial level on each successful recommendation (Section: 3.3). Finally, we show that when trust is regained incrementally when no recommendations are made, the recommender’s optimal total expected profit is unbounded in the long run and that she can recommend both too aggressively and too passively (Section 3.4). These results are also applicable to the phenomenon of “banner blindness”.

## 2 The Pricing of Recommendations

We model the pricing of recommendations 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.<sup>5</sup>

<sup>5</sup>Note that this does *not* restrict the generality of our model. It rather says that each recommendation is for a distinct entity that we

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, and distinguish three scenarios for  $v$ :

- *General*. Without any recommendation the product is sold with probability  $p \in [0, 1]$ . If the recommenders  $R \subseteq \{r_1, \dots, r_n\}$  recommend the product, then the probability that the product is sold is  $p + f(\{s\} \cup R)$ , where  $f : 2^N \rightarrow [0, 1 - p]$  is an arbitrary function with  $f(\{s\}) = 0$ .

The following two scenarios are special cases of *General*.

- *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  $q_i \in [0, 1 - \sum_{j \neq i} q_j]$ . The joint effect of more than one recommendation is the sum of the effect of the individual recommendations. Formally, if the recommenders  $R \subseteq \{r_1, \dots, r_n\}$  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  $\langle N, v \rangle$  (General),  $\langle N, v \rangle$  (Linear), and  $\langle N, v \rangle$  (Threshold). The following table gives the worth  $v(S)$  of all  $S \subseteq N = \{s, r_1, r_2\}$  for all three scenarios.

$S$	Linear	Threshold	General
$\emptyset$	0	0	0
$\{s\}$	$p\delta$	$p\delta$	$p\delta$
$\{r_1\}$	0	0	0
$\{r_2\}$	0	0	0
$\{s, r_1\}$	$(p + q_1)\delta$	$p\delta$	$(p + f(s, r_1))\delta$
$\{s, r_2\}$	$(p + q_2)\delta$	$p\delta$	$(p + f(s, r_2))\delta$
$\{r_1, r_2\}$	0	0	0
$\{s, r_1, r_2\}$	$(p + q_1 + q_2)\delta$	$(p + q)\delta$	$(p + f(s, r_1, r_2))\delta$

Table 1: Worths  $v(S)$  of all possible coalitions  $S$  for one seller  $s$  and two recommenders  $r_1$  and  $r_2$  for our three different models.

Our goal is to find a payoff vector  $(x_s, x_{r_1}, \dots, x_{r_n})$ , where  $x_s$  denotes the expected payoff to the seller and  $x_{r_i}$  denotes the expected payoff to the  $i$ -th recommender. Suppose that the seller  $s$  is recommended by all recommenders  $r_i \in N \setminus \{s\}$ , then the worth of this coalition is  $v(N) = (p + f(N)) \cdot \delta$ . We say that the payoff vector  $(x_s, x_{r_1}, \dots, x_{r_n})$  is *feasible* if  $x_s + \sum_i x_{r_i} = v(N)$ . A feasible payoff vector, which prescribes the *expected* payoff to each player, can be translated into *prices*, i.e. payments from the seller to the recommenders, as follows:

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  $x_{r_i}$ .

refer to as a product.

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  $1/(p + f(N)) \cdot x_{r_i}$ .

In practice, the Pay-per-Sale approach might be preferable as, on a successful recommendation, one could reasonably assume  $p + f(N) = 1$ , sidestepping the problem of estimating  $f(N)$  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.

## 2.1 Impossibility Result

Ideally, the payoff vector  $(x_s, x_{r_1}, \dots, x_{r_n})$  computed by whatever mechanism should give the seller  $s$ , who holds the private information on  $p$ ,  $f$ , and  $\delta$ , the incentive to reveal her information *truthfully*. Formally, we want that for all  $p'$ ,  $f'$ , and  $\delta'$  :  $x_s(p, f, \delta) \geq x_s(p', f', \delta')$ , where  $x_s(p, f, \delta) = (p + f(N)) \cdot \delta - \sum_i x_{r_i}(p, f, \delta)$  and  $x_s(p', f', \delta') = (p' + f(N)) \cdot \delta - \sum_i x_{r_i}(p', f', \delta')$ . Unfortunately, as the following theorem shows, the only truthful payoff vector has  $\sum_i x_{r_i} = 0$ . That is, the seller gets everything and the recommenders get nothing.

**Theorem 1.** *There can be no truthful payoff vector  $(x_s, x_{r_1}, \dots, x_{r_n})$  that has  $\sum_i x_{r_i} \neq 0$  and ensures participation of the seller  $s$  and the recommenders  $r_1$  to  $r_n$ .*

*Proof.* To ensure participation for the seller, we must have  $x_s(p, f, \delta) = (p + f(N)) \cdot \delta - \sum_i x_{r_i}(p, f, \delta) \geq 0$  for all  $p$ ,  $f$ , and  $\delta$ . To ensure participation for the recommenders  $r_1$  to  $r_n$  we must have  $\sum_i x_{r_i}(p, f, \delta) \geq 0$  for all  $p$ ,  $f$ , and  $\delta$ . Now suppose  $(x_s, x_{r_1}, \dots, x_{r_n})$  with  $\sum_i x_{r_i} \neq 0$  was truthful. It follows that  $x_s(p, f, \delta) \geq x_s(p', f', \delta')$  for all  $p'$ ,  $f'$ , and  $\delta'$ , i.e.  $(p + f(N)) \cdot \delta - \sum_i x_{r_i}(p, f, \delta) \geq (p' + f(N)) \cdot \delta - \sum_i x_{r_i}(p', f', \delta')$ . And hence,  $\sum_i x_{r_i}(p', f', \delta') \geq \sum_i x_{r_i}(p, f, \delta)$ . But since  $\sum_i x_{r_i}(p, f, \delta) > 0$  there must be  $p'$ ,  $f'$ ,  $\delta'$  such that  $(p' + f(N)) \cdot \delta' < \sum_i x_{r_i}(p, f, \delta)$  with  $\sum_i x_{r_i}(p', f', \delta') \leq (p' + f(N)) \cdot \delta' < \sum_i x_{r_i}(p, f, \delta)$ . Contradiction!  $\square$

## 2.2 One Recommender

We begin by studying the problem of finding “fair” prices in the setting  $N = \{s, r\}$ , i.e., there is only one seller and one recommender. In this setting the games  $\langle N, v \rangle$  (Linear) and  $\langle N, v \rangle$  (Threshold) are equivalent. We discuss the solution concepts *Core*, *Shapley value*, and *Nash Bargaining Solution*. For a more detailed discussion of these solution concepts see [26, 29].

### 2.2.1 The Core

The *Core* [15] of a coalitional game is an outcome of cooperation among all players where no coalition of players can obtain higher payoffs for all of its members. A payoff vector in the Core is “fair” in the sense that no subset of players can justifiably argue that they are paid too little, as they are unable to achieve higher payoffs on their own.

More formally, the Core of the game  $\langle N, v \rangle$  is the set of *feasible* payoff vectors  $(x_i)_{i \in N}$  for which there is no coalition  $S \subseteq N$  and *S-feasible* payoff vector  $(y_i)_{i \in N}$  such that  $y_i > x_i$  for all  $i \in S$ . Recall that a payoff vector  $(x_i)_{i \in N}$  is *feasible* if  $\sum_{i \in N} x_i = v(N)$ . It is *S-feasible* if  $\sum_{i \in S} x_i = v(S)$ .

The Core can be shown to be non-empty by means of the Bondareva-Shapley Theorem [6, 34], which states that a game has a non-empty core if and only if it is balanced. A game  $\langle N, v \rangle$  is *balanced* if for every balanced collection of weights  $(\lambda_S)_{S \subseteq N}$ :  $\sum_S \lambda_S \cdot v(S) \leq v(N)$ . A *balanced collection of weights*  $(\lambda_S)_{S \subseteq N}$  is a collection of numbers  $\lambda_S \in [0, 1]$  (one for each coalition  $S \subseteq N$ ) such that for all  $i$ :  $\sum_{S \subseteq N: i \in S} \lambda_S = 1$ .

**Theorem 2.** *The game  $\langle \{s, r\}, v \rangle$  (General) has a non-empty core.*

*Proof.* Let  $x \in [0, 1]$ . All balanced collections of weights  $(\lambda_S)_{S \subseteq N}$  are of the form  $\lambda_S = x$  for  $S = \{s\}, \{r\}$  and  $\lambda_S = 1 - x$  for  $S = \{s, r\}$ . By the Bondareva-Shapley Theorem, the Core is non-empty if and only if for all values  $x \in [0, 1]$  :

$$x \cdot (v(\{s\}) + v(\{r\})) + (1 - x)v(\{s, r\}) \leq v(\{s, r\}).$$

For  $x = 0$  this is trivially true. Next we analyze the case  $x > 0$ . Since  $v(\{r\}) = 0$  and  $v(\{s, r\}) - v(\{s\}) = f(s, r) \cdot \delta$ ,

$$\begin{aligned} x \cdot v(\{s\}) - x \cdot v(\{s, r\}) &\leq 0 \\ \Leftrightarrow v(\{s, r\}) - v(\{s\}) &\geq 0 \\ \Leftrightarrow f(\{s, r\}) \cdot \delta &\geq 0. \end{aligned}$$

Since  $f(\{s, r\}) \geq 0$  and  $\delta \geq 0$  this is always true.  $\square$

Recall that the games  $\langle N, v \rangle$  (Linear) and  $\langle N, v \rangle$  (Threshold) are special cases of the game  $\langle N, v \rangle$  (General) and so Theorem 2 also shows non-emptiness of the Core for these games. Next we give necessary and sufficient conditions for a payoff vector  $(x_s, x_{r_1}, \dots, x_{r_n})$  to be in the Core.

**Theorem 3.** *The payoff vector  $((p + f(\{s, r\}))\delta - x, x)$  is in the Core of the game  $\langle \{s, r\}, v \rangle$  (General) if and only if:*

$$0 \leq x \leq f(\{s, r\}) \cdot \delta.$$

*Proof.* Let  $x_s = (p + f(\{s, r\}))\delta - x$  and let  $x_r = x$ . If the vector  $(x_s, x_r)$  is in the Core of the game  $\langle \{s, r\}, v \rangle$  (General), then there exists no coalition  $S \subseteq N = \{s, r\}$  and an *S-feasible* payoff vector  $y = (y_s, y_r)$  such that  $y_i > x_i$  for all  $i \in S$ . That is, for all  $S \subseteq N$  and *S-feasible* payoff vectors  $y = (y_s, y_r)$  we have that  $y_i \leq x_i$  for all  $i \in S$ . For  $S = \{s, r\}$  this means that  $y_s + y_r = v(\{s, r\}) \leq x_s + x_r = v(\{s, r\})$  (which is trivially true). For  $S = \{s\}$  this means that  $y_s = v(\{s\}) = p \cdot \delta \leq x_s = (p + f(\{s, r\}))\delta - x$ , i.e.  $x \leq f(\{s, r\}) \cdot \delta$ . For  $S = \{r\}$  this means that  $y_r = v(\{r\}) = 0 \leq x_r = x$ , i.e.  $x \geq 0$ . That is,  $0 \leq x \leq f(\{s, r\}) \cdot \delta$ .

For the reverse direction assume by contradiction that  $0 \leq x \leq f(\{s, r\}) \cdot \delta$  but that  $(x_s, x_r)$  is *not* in the Core, i.e. there exists a coalition  $S \subseteq N = \{s, r\}$  and a *S-feasible* payoff vector  $y = (y_s, y_r)$  such that  $y_i > x_i$  for all  $i \in S$ . We cannot have  $S = \{s, r\}$  as then  $y_s + y_r = v(\{s, r\}) > x_s + x_r = v(\{s, r\})$ , which gives a contradiction. But if  $S = \{s\}$ , then

$y_s = v(\{s\}) = p \cdot \delta > x_s = (p + f(\{s, r\})) \cdot \delta - x$ , i.e.  $x > f(\{s, r\}) \cdot \delta$ , which gives a contradiction. Finally, if  $S = \{r\}$ , then  $y_r = v(\{r\}) = 0 > x_r = x$ , i.e.  $x < 0$ , which also gives a contradiction.  $\square$

For the games  $\langle \{s, r\}, v \rangle$  (Linear) and  $\langle \{s, r\}, v \rangle$  (Threshold) this means that  $((p + q)\delta - x, x)$  is in the core if and only if:

$$0 \leq x \leq q \cdot \delta.$$

This implies that any “feasible” payoff vector is in the Core. The only restriction on the payoff vector is that the payoff to the recommender be non-negative and in expectation no higher than the “extra profit” the seller can expect from the recommendation. In particular, a payoff vector that gives everything to the seller and nothing to the recommender would be in the Core. This demonstrates the weakness of the recommenders: They cannot form a coalition with non-zero value without the seller.

## 2.2.2 Shapley Value

One problem with the Core is that it does *not* assign a unique payoff vector to a game. This makes it necessary to have another criterion for choosing a payoff vector. The *Shapley value* [33] is a solution concept that assigns to each game a unique, provably fair payoff vector. In general, a *value*  $\phi : v \rightarrow \mathbb{R}^{n+1}$  maps each game  $\langle N, v \rangle$  to a unique vector  $\phi(v)$ ; the  $i$ -th entry  $\phi_i(v)$  of this vector being the *expected payoff* to player  $i$ . The Shapley value is the unique value satisfying the following axioms:

1. *Symmetry*: If player  $i$  and  $j$  are interchangeable, then  $\phi_i(v) = \phi_j(v)$ . Formally, if for every  $S \subseteq N$  s.t.  $i \in S, j \notin S: v((S \setminus \{i\}) \cup \{j\}) = v(S)$ , then  $\phi_i(v) = \phi_j(v)$ .
2. *Dummy*: If player  $i$ 's contribution to any coalition  $S$  is zero, then  $\phi_i(v) = 0$ . Formally, if for every  $S \subseteq N \setminus \{i\}: v(S \cup \{i\}) = v(S)$ , then  $\phi_i(v) = 0$ .
3. *Additivity*: Player  $i$ 's entry  $\phi_i(v)$  should be additive in  $v$ . Formally, if  $\langle N, v + w \rangle$  is derived from  $\langle N, v \rangle$  and  $\langle N, w \rangle$  by defining  $(v + w)(S) = v(S) + w(S)$  for all  $S \subseteq N$ , then  $\phi_i(v + w) = \phi_i(v) + \phi_i(w)$  for all  $i \in N$ .

These axioms can be interpreted as formalizing a notion of “fairness”, that postulates that the expected payoff to player  $i$  be proportional to player  $i$ 's contribution to the outcome of the game. For an analysis along these lines see [25].

**Definition 1.** The *Shapley value*  $\phi(v) = (\phi_1(v), \dots, \phi_N(v))$  of the game  $\langle N, v \rangle$  is defined as follows:

$$\phi_i(v) = \sum_{S \subseteq N \setminus \{i\}} \frac{|S|!(|N| - 1 - |S|)!}{|N|!} \cdot (v(S \cup \{i\}) - v(S)).$$

One interpretation of this is: Suppose that all the players are arranged in some order, all orders being equally likely, then  $\phi_i(v)$  is the *expected marginal contribution* of player  $i$  to the set of players who precede her.

**Theorem 4.** Consider the game  $\langle \{s, r\}, v \rangle$  (General). The *Shapley value*  $\phi(v) = (\phi_s(v), \phi_r(v))$  is given by

$$\phi_s(v) = p\delta + \frac{1}{2}f(\{s, r\})\delta \quad \text{and} \quad \phi_r(v) = \frac{1}{2}f(\{s, r\})\delta.$$

*Proof.* The claim follows from the definition of the game  $\langle \{s, r\}, v \rangle$  (General) and Definition 1. Note that the worths  $v(S)$  of all  $S \subseteq N = \{s, r\}$  can be read from Table 1 by treating  $r_1$  as  $r$  and ignoring rows containing  $r_2$ .  $\square$

For the game  $\langle \{s, r\}, v \rangle$  (Linear, Threshold) this means that

$$\phi_s(v) = p\delta + \frac{1}{2}q\delta \quad \text{and} \quad \phi_r(v) = \frac{1}{2}q\delta.$$

This shows that the payoff to the recommender should be proportional to her contribution to the seller's expected “extra profit”. In particular, it shows that the recommender's payoff should be *linear* in her contribution to the purchase probability, i.e.  $f(\{s, r\})$ , and also in the seller's margin or gain  $\delta$ . This is consistent with “real life” pricing schemes that redeem the recommender with a certain percentage of the sales price [22], assuming that for a given product family the margin is proportional to the sales price.

## 2.2.3 The Nash Bargaining Solution

The last solution concept that we discuss in this section is the *Nash Bargaining Solution* [27].<sup>6</sup> The basic idea here is to view the game  $\langle N, v \rangle$  as a *bargaining problem* over a set  $F$  of feasible payoff vectors  $f = (f_0, \dots, f_n)$  and a dedicated payoff vector  $d = (d_0, \dots, d_n)$ ; the payoff vector in case of a *disagreement*. A *solution* is a function  $\phi : (F, d) \rightarrow F$ . The Nash Bargaining Solution is the unique solution satisfying the following axioms:

1. *Pareto Efficiency*. There is no  $f \in F$  such that  $f_i \geq \phi_i(F, d)$  for all  $i \in N$  and  $f_j > \phi_j(F, d)$  for at least one  $j \in N$ .
2. *Individual Rationality*. For all  $i \in N: \phi_i(F, d) \geq d_i$ .
3. *Scale Covariance*. If  $F' = \{(\lambda_0 \cdot f_0 + \gamma_0, \dots, \lambda_n \cdot f_n + \gamma_n) \mid (f_0, \dots, f_n) \in F\}$  and  $d' = (\lambda_0 \cdot d_0 + \gamma_0, \dots, \lambda_n \cdot d_n + \gamma_n)$ , then  $\phi(F', d') = (\lambda_0 \cdot \phi_0(F, d) + \gamma_0, \dots, \lambda_n \cdot \phi_n(F, d) + \gamma_n)$ .
4. *Independence of Irrelevant Alternatives*. If  $F' \subseteq F$  and  $\phi(F, d) \in F'$ , then  $\phi(F', d) = \phi(F, d)$ .
5. *Symmetry*. If  $(f_0, \dots, f_i, \dots, f_j, \dots, f_n) \in F$  implies  $(f_0, \dots, f_j, \dots, f_i, \dots, f_n) \in F$  and  $d_i = d_j$ , then  $\phi_i(F, d) = \phi_j(F, d)$ .

The Nash Bargaining solution is “fair” in the sense that is *Pareto efficient*, i.e. it is impossible to improve the payoff of one or more players without hurting that of others. One can show that it satisfies  $\phi(F, d) \in \operatorname{argmax}_{f \in F} \prod_i (f_i - d_i)$  [26]. We use this to prove:

<sup>6</sup>The only connection between this solution concept and the concept of a *Nash equilibrium* [26, 29] is John F. Nash.

**Theorem 5.** For  $\langle \{s, r\}, v \rangle$  (General) let  $F = \{(f_s, f_r) \mid f_s \geq 0, f_r \geq 0 \text{ and } f_s + f_r = (p + f(\{s, r\})) \cdot \delta\}$  and  $d = (d_s, d_r) = (p \cdot \delta, 0)$ . Then,

$$\phi(F, d) = \left( \left( p + \frac{1}{2} f(\{s, r\}) \right) \delta, \frac{1}{2} f(\{s, r\}) \delta \right).$$

For the game  $\langle \{s, r\}, v \rangle$  (Linear, Threshold) this means that:

$$\phi(F, d) = \left( \left( p + \frac{1}{2} q \right) \delta, \frac{1}{2} q \delta \right).$$

This demonstrates that the Nash Bargaining Solution and the Shapley value coincide in our model. On the one hand, this is surprising as the axioms used to define the Shapley value and the Nash Bargaining Solution are quite different. On the other hand, this is intuitive as for two players there is only one non-trivial coalition to be considered for the Shapley value. So if this two-player coalition leads to a bigger payoff  $P$  than the sum of its two non-cooperative ‘‘atoms’’, then the different symmetry axioms present in both solution concepts imply that this surplus should be divided 50-50. This holds as long as the feasible payoff vectors are  $F = \{(f_0, f_1) \mid f_0 + f_1 \leq P\}$ , and so  $(f_0, f_1) \in F \Leftrightarrow (f_1, f_0) \in F$ , but would stop to hold if, e.g. there were additional constraints on  $F$  such as  $f_0 \geq 2 \cdot f_1$ . In such cases, the symmetry axiom of the Nash Bargaining Solution no longer applies.

## 2.3 Many Recommenders

Next we study the problem of finding ‘‘fair’’ prices in the more general setting  $N = \{s, r_1, \dots, r_n\}$ , i.e. there is one seller and  $n \geq 1$  recommenders. Note that in this setting the games  $\langle N, v \rangle$  (Linear) and  $\langle N, v \rangle$  (Threshold) are no longer equivalent. As in the setting where  $N = \{s, r\}$  we study the solution concepts Core, Shapley value, and Nash Bargaining Solution.

### 2.3.1 The Core

Recall that the Core of the game  $\langle N, v \rangle$  comprises all feasible payoff vectors with which no coalition  $S \subseteq N$  is ‘‘unhappy’’ meaning that the players in  $S$  cannot break away to obtain a higher payoff on their own. For a formal definition of the Core (and related definitions) see Section 2.2.1.

**Theorem 6.** The game  $\langle N, v \rangle$  (General) has a non-empty core iff for every balanced collections of weights  $(\lambda_S)_{S \subseteq N}$ :

$$\sum_{S \subseteq N: s \in S} [\lambda_S (f(N) - f(S))] \geq 0.$$

*Proof.* Let  $(\lambda_S)_{S \subseteq N}$  be a balanced collection of weights. Since  $v(S) = 0$  whenever  $s \notin S$ , applying the Bondareva-Shapley Theorem [6, 34] to the game  $\langle N, v \rangle$  (General) gives:

$$\sum_{S \subseteq N: s \in S} [\lambda_S (p + f(S))] \leq p + f(N).$$

Since  $(\lambda_S)_{S \subseteq N}$  is a balanced collections of weights, we have  $\sum_{S \subseteq N: s \in S} \lambda_S = 1$  and  $\lambda_N = 1 - \sum_{S \subseteq N: s \in S} \lambda_S$ . It follows

that:

$$\begin{aligned} \sum_{S \subseteq N: s \in S} \lambda_S \cdot f(S) &\leq f(N). \\ \Leftrightarrow \sum_{S \subseteq N: s \in S} [\lambda_S (f(S) - f(N))] + f(N) &\leq f(N) \\ \Leftrightarrow \sum_{S \subseteq N: s \in S} [\lambda_S (f(N) - f(S))] &\geq 0. \quad \square \end{aligned}$$

The condition given by Theorem 6 holds trivially for  $\langle N, v \rangle$  (General) if  $f(N) \geq f(S)$  for all  $S \subseteq N$  since all the  $\lambda_S$  values are non-negative. For the game  $\langle N, v \rangle$  (Linear) and  $\langle N, v \rangle$  (Threshold) this means that the core is *always* non-empty since  $q_i \geq 0$  for all  $i$  respectively  $q > 0$ .

**Theorem 7.** Consider the game  $\langle N, v \rangle$  (General). The payoff vector  $(x_s = (p + f(N)) \cdot \delta - \sum_i x_{r_i}, x_{r_1}, \dots, x_{r_n})$  is in the Core if and only if for all  $T \subseteq N$  s.t.  $s \notin T$ :

$$0 \leq \sum_{r_i \in T} x_{r_i} \leq (f(N) - f(N \setminus T)) \cdot \delta.$$

*Proof.* Assume that the payoff vector  $(x_s = (p + f(N)) \cdot \delta - \sum_i x_{r_i}, x_{r_1}, \dots, x_{r_n})$  is in the Core. Since  $v(S) = 0$  for all coalitions  $S \subseteq N$  such that  $s \notin S$ , it follows that:

1. For all  $S \subseteq N$  such that  $s \in S$ :

$$\begin{aligned} (p + f(N))\delta - \sum_{r_i \in N} x_i + \sum_{r_i \in S} x_i &\geq v(S) \\ \Leftrightarrow (p + f(N))\delta - \sum_{r_i \in N \setminus S} x_i &\geq (p + f(S))\delta \\ \Leftrightarrow (f(N) - f(S))\delta &\geq \sum_{r_i \in N \setminus S} x_i \end{aligned}$$

2. For all  $S \subseteq N$  such that  $s \notin S$ :

$$\sum_{r_i \in S} x_i \geq v(S) \Leftrightarrow \sum_{r_i \in S} x_i \geq 0.$$

With  $T = N \setminus S$  in 1. and  $T = S$  in 2. it follows that  $0 \leq \sum_{r_i \in T} x_{r_i} \leq (f(N) - f(N \setminus T)) \cdot \delta$  for all  $T \subseteq N$  s.t.  $s \notin T$ .

For the reverse direction assume by contradiction that  $0 \leq \sum_{r_i \in T} x_{r_i} \leq (f(N) - f(N \setminus T)) \cdot \delta$  for all  $T \subseteq N$  s.t.  $s \notin T$  but that  $(x_s = (p + f(N)) \cdot \delta - \sum_i x_{r_i}, x_{r_1}, \dots, x_{r_n})$  is not in the Core, i.e. there exists a coalition  $S \subseteq N$  and an  $S$ -feasible payoff vector  $y = (y_s, y_{r_1}, \dots, y_{r_n})$  in which  $y_k > x_k$  for all players  $k$  in  $S$ . Since  $y$  is  $S$ -feasible, the total payoff  $\sum_{k \in S} y_k$  to the players in  $S$  must equal  $v(S)$ . Since  $y_p > x_p$  for all players  $p$  in  $S$ , we must have  $\sum_{k \in S} y_k > \sum_{k \in S} x_k$ . Thus,  $v(S) > \sum_{k \in S} x_k$ .

*Case 1:* If  $s \in S$ , since  $v(S) = (p + f(S)) \cdot \delta$  and  $\sum_{k \in S} x_k = (p + f(N))\delta - \sum_{r_i \in N \setminus S} x_{r_i}$ , this means that:

$$\begin{aligned} (p + f(S)) \cdot \delta &> (p + f(N))\delta - \sum_{r_i \in N \setminus S} x_{r_i} \\ \Leftrightarrow \sum_{r_i \in N \setminus S} x_{r_i} &> (f(N) - f(S)) \cdot \delta. \end{aligned}$$

With  $T = N \setminus S$  this gives a contradiction to the fact that for all  $T \subseteq N$  s.t.  $s \notin T$ :  $\sum_{r_i \in T} x_{r_i} \leq (f(N) - f(N \setminus T)) \cdot \delta$ .

Case 2: If  $s \notin S$ , since  $v(S) = 0$  and  $\sum_{k \in S} x_k = \sum_{r_i \in S} x_{r_i}$ , this means that:

$$v(S) = 0 > \sum_{r_i \in S} x_{r_i}.$$

With  $T = S$  this gives a contradiction to the fact that for all  $T \subseteq N$  s.t.  $s \notin T$ :  $\sum_{r_i \in T} x_{r_i} \geq 0$ .  $\square$

For the game  $\langle N, v \rangle$  (Linear) this means that the payoff vector  $(x_s = (p + \sum_i q_i) \cdot \delta - \sum_i x_{r_i}, x_{r_1}, \dots, x_{r_n})$  is in the Core if and only if for all  $T \subseteq N$  s.t.  $s \notin T$ :

$$0 \leq \sum_{r_i \in T} x_{r_i} \leq \left( \sum_{r_i \in T} q_i \right) \cdot \delta.$$

For the game  $\langle N, v \rangle$  (Threshold) this means that the payoff vector  $(x_s = (p + q)\delta - \sum_i x_{r_i}, x_{r_1}, \dots, x_{r_n})$  is in the Core if and only if for all  $T \subseteq N$  s.t.  $s \notin T$ :

$$0 \leq \sum_{r_i \in T} x_i \leq \begin{cases} 0 & , \text{ if } |T| \leq n - k \\ q \cdot \delta & , \text{ if } |T| > n - k \end{cases}.$$

This means that for  $\langle N, v \rangle$  (Linear) and  $\langle N, v \rangle$  (Threshold) with  $k = n$  a certain payoff vector is in the core precisely if no coalition of recommenders receives more than their joint contribution to the seller's expected "extra profit". For the game  $\langle N, v \rangle$  (Threshold) with  $k < n$  this means that  $x_{r_i} = 0$  for all  $i$  (with  $T = \{r_i\}$ ) and, thus, the only payoff vector in the Core is  $((p + q)\delta, 0, \dots, 0)$ , i.e. the seller gets everything and the recommenders get nothing.

### 2.3.2 Shapley Value

Recall that the Shapley value assigns to each game a unique payoff vector that is "fair" as it satisfies the *Symmetry*, *Dummy*, and *Additivity* axioms. For a formal definition of the Shapley value (and related definitions) see Section 2.2.2.

**Theorem 8.** Consider the game  $\langle N, v \rangle$  (General). The Shapley value  $\phi(v) = (\phi_s(v), \phi_{r_1}(v), \dots, \phi_{r_n}(v))$  is given by

$$\begin{aligned} \phi_s(v) &= p\delta + \sum_{S \subseteq N \setminus \{s\}} \frac{|S|!(|N| - 1 - |S|)!}{|N|!} f(S \cup \{s\})\delta \\ \phi_{r_i}(v) &= \sum_{\substack{S \subseteq N \setminus \{r_i\}: \\ s \in S}} \frac{|S|!(|N| - 1 - |S|)!}{|N|!} (f(S \cup \{r_i\}) - f(S))\delta. \end{aligned}$$

*Proof.* The claim follows from the definition of the game  $\langle N, v \rangle$  (General) and Definition 1. The worths  $v(S)$  of all  $S \subseteq N = \{s, r_1, r_2\}$  are given explicitly in Table 1. For  $|N| > 3$  the worths  $v(S)$  are obtained similarly.  $\square$

For the game  $\langle N, v \rangle$  (Linear) this means that

$$\begin{aligned} \phi_s(v) &= p\delta + \frac{1}{2} \sum_i q_i \delta & \text{and} \\ \phi_{r_i}(v) &= \frac{1}{2} q_i \delta & \text{for all } i. \end{aligned}$$

For the game  $\langle N, v \rangle$  (Threshold) this means that

$$\begin{aligned} \phi_s(v) &= p\delta + \left(1 - n \frac{k!(n-k)!}{(n+1)!}\right) q\delta & \text{and} \\ \phi_{r_i}(v) &= \frac{k!(n-k)!}{(n+1)!} q\delta & \text{for all } i. \end{aligned}$$

This suggests that in the game  $\langle N, v \rangle$  (Linear) each individual recommender should receive a share of exactly one half of her contribution to the expected "extra profit" of the recommender. For 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". So all in all the Shapley value does not only give a unique payoff vector, but it also yields "fair" payoffs in the sense that the payoff to each recommender is proportional to the recommender's contribution to the "extra profit" the seller can expect.

### 2.3.3 Anonymity-Proof Shapley Value

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 sets  $S_i$  being disjoint. 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  $\phi_a(v) = \frac{1}{2}$  and  $r_2$  gets  $\phi_b(v) = \frac{1}{2}$ . I.e.  $r_2$  would be better off.

The *anonymity-proof Shapley value* [28] cannot be "tricked" in this way. It is defined as follows:

**Definition 2.** For any set  $A' \subseteq A$  of declared arguments the anonymity-proof Shapley value  $\psi_a(v)$  for  $a \in A'$  is:

$$\psi_a(v) = \frac{\phi_a(v)}{\sum_{a' \in A'} \phi_{a'}(v)} v(A').$$

So a better way to redeem the recommenders would be to compute the anonymity-proof Shapley value  $\psi_a(v)$  for each argument  $a \in A'$  and to 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.

### 2.3.4 The Nash Bargaining Solution

Recall that the Nash Bargaining Solution is defined as the unique bargaining solution that satisfies the axioms listed in Section 2.2.3.

**Theorem 9.** For  $\langle N, v \rangle$  (General) let  $F = \{(f_s, f_{r_1}, \dots, f_{r_n}) \mid f_s \geq 0, f_{r_i} \geq 0 (\forall i) \text{ and } f_s + \sum_i f_{r_i} = (p + f(N)) \cdot \delta\}$  and  $d = (d_s, d_{r_1}, \dots, d_{r_n}) = (p \cdot \delta, 0, \dots, 0)$ . Then,

$$\begin{aligned} \phi_s(F, d) &= \left(p + \frac{1}{n+1} f(N)\right) \cdot \delta & \text{and} \\ \phi_{r_i}(F, d) &= \frac{1}{n+1} f(N) \cdot \delta & \text{for all } i. \end{aligned}$$

*Proof.* The claim follows from the fact that  $\phi(F, d) \in \operatorname{argmax}_{f \in F} \prod_i (f_i - d_i)$  [26].  $\square$

One problem with the Nash Bargaining Solution is that it completely ignores the possibility of cooperation among subsets of players. To see that this may lead to “unfair” prices, consider the game  $\langle N, v \rangle$  (Linear) with one seller  $s$  and two recommenders  $r_1$  and  $r_2$ . Suppose that  $q_1 = 1 - p - \epsilon$  and that  $q_2 = \epsilon$  for some small  $\epsilon > 0$ . It follows that  $\phi_{r_1}(F, d) = \phi_{r_2}(F, d) = \frac{1-p}{2}$ , i.e. the expected payoff to both recommenders is the same. But since  $r_1$ ’s contribution to the expected worth of the grand coalition is significantly higher than that of  $r_2$  (especially as  $\epsilon \rightarrow 0$ ), this cannot be regarded as “fair”. We conclude that for  $|N| > 2$  it is not advisable to use the Nash Bargaining Solution to guide the pricing of recommendations.

### 3 Recommending Strategically

We study the following problem: There are  $n$  products. For each product the recommender has two options: “recommend it” or “not recommend it”. 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 3.2. 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 3.3. 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 3.4.

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  $R(p_0, l, g) = \lim_{n \rightarrow \infty} M_n(p_0, l, g)$ . Before looking at the theoretical analysis, the following section experimentally demonstrates the different behavior of the optimal total expected reward in these settings.

#### 3.1 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 3.1, lead to an asymptotically unbounded reward.

#### 3.2 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 the initial trust  $p_0$  and the loss rate  $l$  are strictly smaller than 1, then the total expected reward the recommender can achieve is finite, even when there is an infinite sequence of items to recommend.

### 3.3 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.  $\lim_{n \rightarrow \infty} M_n(p_0, l)$  is finite) if and only if  $q > 0$ .

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

*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  $\ln(\cdot)$  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

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

Since  $\text{dilog}(x)$  is continuous<sup>7</sup> and  $\text{dilog}(1) = 0$  (see Lemma 2), we get

$$\begin{aligned} \int_{x=0}^{\infty} \ln(1 - c^{x+1}) dx &= \frac{\text{dilog}(1-c)}{\ln(c)} - \frac{\text{dilog}(1)}{\ln(c)} \\ &= \frac{\text{dilog}(1-c)}{\ln(c)}. \end{aligned}$$

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  $\int_{x=0}^{\infty} \ln(1 - c^{x+1}) dx > -\infty$ .  $\square$

<sup>7</sup>It is even differentiable as it is defined as an indefinite integral.

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

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

Using Lemma 1 we can prove the following theorem.

**Theorem 10.** *Let  $\text{dilog}(x) = \int_x^1 \frac{\ln(t)}{1-t} dt$ ,  $c = \max(p_0, l)$ , and  $\delta(c) = (1-c) \exp(\text{dilog}(1-c)/\ln(c))$ . Then, for all  $1 > p_0 \geq 0$ ,*

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

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.

### 3.4 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 profit.

**Theorem 11.** *Let  $\psi$  be the smallest integer such that  $l \cdot g^\psi \geq 1$ . If  $k > \psi$ , then, for all  $1 > p_0, l > 0$  and  $\infty > g \geq 1$ ,*

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

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

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  $R(p_0, l, g) = \lim_{n \rightarrow \infty} M_n(p_0, l, g) = \infty$ .

**Theorem 12.** Let  $\psi$  be the smallest integer such that  $l \cdot g^\psi \geq 1$ . If  $k \leq \psi$ , then, for all  $1 > p_0, l > 0$  and  $\infty > g \geq 1$ , there exist  $p'_0$  and  $l'$  such that

$$\lim_{n \rightarrow \infty} A_n^{(k)}(p_0, l, g) \leq \lim_{n \rightarrow \infty} M_n(p'_0, l') < \infty.$$

*Proof.* If  $k < \psi$ , then the profit maximization problem with parameters  $p_0, l$ , and  $g$  on the products 1, 2, 3, etc. is equivalent to the profit 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 10.  $\square$

Whereas Theorem 11 shows that recommending too seldomly is sub-optimal, Theorem 12 shows that recommending too often is even worse.

## 4 Discussion and Future Work

Suppose you recommend a product to a friend and the seller of the product offers to pay you for your recommendation. What would be a “good” price? Our first finding was that the only “truthful” price would be zero. The problem with this, however, is that if you do not get paid, then you might as well decide *not* to recommend the product. And so the seller might be willing to pay you a “fair” price. We approached the problem of finding “fair” prices by studying solution concepts from coalitional game theory such as the Core, the Shapley value, and the Nash Bargaining Solution. Since each of these solution concepts formalizes some notion of “fairness”, these prices can be regarded as *provably* “fair”. We view such an “axiomatic” foundation of “fairness” to be the only viable basis for truly “fair” prices in practice.

Now suppose that you get paid for each successful recommendation you make, and that you want to maximize the amount of money paid to you. At first sight, it might appear that the best strategy for you is to send out as many recommendations to as many friends as possible. But, then, just as you get “blind” when being shown too many ads, your friends will probably start to ignore your “recommendations”. We addressed this problem by modeling the loss in “trust” by a drop in “purchase probability” on each unsuccessful recommendation. Our main finding here was that, even if the “trust” in you is reset to the initial level on a single successful recommendation, the total expected profit you can make over an infinite period of time is bounded. This can only be overcome if the recommendee also incrementally regains “trust” over periods without any recommendation.

We believe that our work motivates a number of interesting research questions. E.g.: What are “good” pricing mechanisms in settings where the seller has objectives such as maintaining the buyer’s “trust”? 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 click-through-rates are non-constant and decay with each irrelevant advertisement being shown?

## References

- [1] Z. Abrams, O. Mendeleevitch, and J. Tomlin. Optimal delivery of sponsored search advertisements subject to budget constraints. In *Conference on Electronic commerce (EC'07)*, pages 272–278, 2007.
- [2] G. E. Andrews, R. Askey, and R. Roy. *Special functions*. Cambridge University Press, 2001.
- [3] D. Arthur, M. Motwani, A. Sharma, and Y. Xu. Pricing strategies for viral marketing on social networks. In *Workshop on Internet and Network Economics (WINE'09)*, page to appear, 2009.
- [4] J. P. Benway and D. M. Lane. Banner blindness: Web searchers often miss “obvious” links. *ITG Newsletter*, 1(3), 1998. [http://www.internettg.org/newsletter/dec98/banner\\_blindness.html](http://www.internettg.org/newsletter/dec98/banner_blindness.html).
- [5] D. Bergemann and D. Ozmen. Optimal pricing with recommender systems. In *Conference on Electronic commerce (EC'06)*, pages 43–51, 2006.
- [6] O. N. Bondareva. Some applications of linear programming methods to the theory of cooperative games. *Problemy Kybernetiki*, 10:119–139, 1963.
- [7] J. J. Brown and P. H. Reingen. Social ties and word-of-mouth referral behavior. *Journal of Consumer Research: An Interdisciplinary Quarterly*, 14(3):350–62, 1987.
- [8] M. Burke, A. Hornof, E. Nilsen, and N. Gorman. High-cost banner blindness: Ads increase perceived workload, hinder visual search, and are forgotten. *ACM Transactions on Computer-Human Interaction*, 12(4):423–445, 2005.
- [9] P. Chatterjee, D. L. Hoffman, and T. P. Novak. Modeling the clickstream: Implications for web-based advertising efforts. *Marketing Science*, 22:520–541, 2000.
- [10] P. Domingos and M. Richardson. Mining the network value of customers. In *SIGKDD international conference on Knowledge discovery and data mining (KDD'01)*, pages 57–66, 2001.
- [11] J. Feldman, S. Muthukrishnan, M. Pal, and C. Stein. Budget optimization in search-based advertising auctions. In *Conference on Electronic commerce (EC'07)*, pages 40–49, 2007.
- [12] J. Fernandez and B. Nahata. Pay what you like. Technical Report 16265, Munich Personal RePEc Archive, 2009.
- [13] Friend vouch, 2008. <http://www.friendvouch.com>.
- [14] A. C. B. Garcia, M. Ekstrom, and H. Björnsson. Hyriwyg: leveraging personalization to elicit honest recommendations. In *Conference on Electronic commerce (EC'04)*, pages 232–233, 2004.

- [15] D. Gillies. *Contributions to the Theory of Games IV*, chapter Solutions to general non-zero-sum games, pages 47–85. Princeton University Press, 1959.
- [16] R. Grewal, T. W. Cline, and A. Davies. Early-entrant advantage, word-of-mouth communication, brand similarity, and the consumer decision-making process. *Journal of Consumer Psychology*, 13(3):187–197, 2003.
- [17] 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.
- [18] P. M. Herr, F. R. Kardes, and J. Kim. Effects of word-of-mouth and product-attribute information on persuasion: An accessibility-diagnostics perspective. *Journal of Consumer Research*, 17(4):454–462, 1991.
- [19] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence through a social network. *KDD*, pages 137–146, 2003.
- [20] S. Lahaie. An analysis of alternative slot auction designs for sponsored search. In *Conference on Electronic commerce (EC’06)*, pages 218–227, 2006.
- [21] S. Lahaie, D. Pennock, A. Saberi, and R. Vohra. *Algorithmic Game Theory*, chapter Sponsored Search Auctions, pages 699–716. Cambridge University Press, 2007.
- [22] J. Leskovec, L. A. Adamic, and B. A. Huberman. The dynamics of viral marketing. In *Conference on Electronic commerce (EC’06)*, pages 228–237, 2006.
- [23] A. Mantzaris. Pay-what-you-like restaurants, 2008. <http://www.budgettravel.com/bt-dyn/content/article/2008/02/29/AR2008022902761.html>.
- [24] S. Maxwell. *The Price is Wrong: Understanding What Makes a Price Seem Fair and the True Cost of Unfair Pricing*. John Wiley and Sons, 2008.
- [25] H. Moulin. *Axioms of Cooperative Decision Making (Econometric Society Monographs)*. Cambridge University Press, July 1991.
- [26] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1997.
- [27] J. Nash, J. F. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.
- [28] N. Ohta, V. Conitzer, Y. Satoh, A. Iwasaki, and M. Yokoo. Anonymity-proof shapley value: Extending shapley value for coalitional games in open environments. *Autonomous Agents and Multiagent Systems*, pages 927–934, 2008.
- [29] M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.
- [30] K. Patel and R. G. P. Kantamneni. Monetizing low value clickers. United States Patent Application 20080249854, 2008.
- [31] P. H. Reingen, B. L. Foster, J. J. Brown, and S. B. Seidman. Brand congruence in interpersonal relations: A social network analysis. *Journal of Consumer Research*, 11(3):771–783, 1984.
- [32] J. J. Rotemberg. Fair pricing. Technical Report 10915, National Bureau of Economic Research, 2004.
- [33] L. S. Shapley. *Contributions to the Theory of Games II*, chapter A Value for n-person Games, pages 307–317. Princeton University Press, 1953.
- [34] L. S. Shapley. On balanced sets and cores. *Naval Research Logistics Quarterly*, 14:453–460, 1967.
- [35] U. Shardanand and P. Maes. Social information filtering: algorithms for automating “word of mouth”. In *SIGCHI conference on Human factors in computing systems (CHI’95)*, pages 210–217, 1995.
- [36] A. Singla and I. Weber. Camera brand congruence in the flickr social graph. In *Conference on Web Search and Data Mining (WSDM’09)*, pages 252–261, 2009.