

The Dynamics of Viral Marketing

Jurij Leskovec^{1,3}

Lada A. Adamic²

Bernardo A. Huberman²

¹ Carnegie Mellon University, Pittsburgh, PA 15213

jure@cs.cmu.edu

² HP Labs, Palo Alto, CA 94304

{lada.adamic,bernardo.huberman}@hp.com

³The research was done while at HP Labs.

The Dynamics of Viral Marketing

Abstract

We present an analysis of a person-to-person recommendation network, consisting of 4 million people who made 16 million recommendations on half a million products. We observed the propagation of recommendations and the cascade sizes, which can be explained by a stochastic model. We then established how the recommendation network grows over time and how effective it is from the viewpoint of the sender and receiver of the recommendations. While on average recommendations are not very effective at inducing purchases and do not spread very far, there are product and pricing categories for which viral marketing seems to be very effective.

1 Introduction

With consumers showing increasing resistance to traditional forms of advertising such as TV or newspaper ads, marketers have turned to alternate strategies, including viral marketing. Viral marketing exploits existing social networks by encouraging customers to share product information with their friends. Until recently, it has been difficult to measure how influential person-to-person recommendations actually are over a wide range of products. We were able to directly measure the effectiveness of recommendations by studying one online retailer's incentivised viral marketing program. The website gave discounts to customers recommending any of its products to others, and then tracked the resulting purchases and additional recommendations.

Previously, a few in depth studies have shown that social networks affect the adoption of individual innovations and products. Ryan and Gross [20] developed a model of innovation diffusion based on the adoption of a new type of hybrid seed corn among Iowa farmers in the first half of the 20th century. Subsequent diffusion studies focused on everything from doctors prescribing the drug tetracycline [6] to the adoption of defense mechanisms against hostile takeovers by the boards of directors of major companies [7] (for a review see [19] or [21]).

Although word of mouth can be a powerful factor influencing purchasing decisions, it can be tricky for advertisers to tap into. Some services used by individuals to communicate are natural candidates for viral marketing, because the product can be observed or advertised as part of the communication. Free email services such as Hotmail and Yahoo had very fast adoption curves because every email sent through them contained an advertisement for the service. Hotmail spent a mere \$50,000 on traditional marketing and still grew from zero to 12 million users in 18 months [10]. Google's Gmail captured a significant part of market share in spite of the fact that the *only* way to sign up for the service is through a referral.

Most products cannot be advertised in such a direct way however. At the same time the choice of products available to consumers has increased manyfold thanks to online retailers who can offer a much wider variety than traditional brick-and-mortar stores. Not only is the variety of products larger, but one observes a 'fat tail' phenomenon, where a large fraction of purchases are of relatively obscure items. On Amazon.com, somewhere between 20 to 40 percent of unit sales fall outside of its top 100,000 ranked products [3]. Rhapsody, a streaming-music service, streams more tracks outside than inside its top 10,000 tunes [2]. Effectively advertising these niche products using traditional advertising approaches is impractical. Therefore using more targeted marketing approaches is advantageous to both to the merchant and the consumer, who would benefit from learning about new products.

The problem is partly addressed by the advent of online product and merchant reviews, both at retail sites such as EBay and Amazon, and specialized product comparison sites such as Epinions and CNET. These

websites have given new power to word-of-mouth advertising, because a product review by a single consumer can now be accessed by all other users searching for the product on the site. For the case of consumer electronics, over half of individuals research products via a search engine before making a purchase [4]. The rating of products and merchants has been shown to effect the likelihood of an item being bought [17, 5].

Further refinements to online review sites include allowing users to rate the usefulness of reviews and to specify whose reviews they trust. Richardson and Domingos [18] used Epinions’ trusted reviewer network to construct models of network influence, as well as an algorithm to maximize viral marketing efficiency assuming that individuals’ probability of purchasing a product depends on the opinions on the trusted peers in their network. Kempe, Kleinberg and Tardos [11] evaluate the efficiency of several algorithms for maximizing cascades given various models of adoption. While these models address the question of maximizing the spread of influence in a network, they were based on assumed rather than measured influence effects.

Another very effective way to recommend products to customers has been to mine purchase patterns to either present user-specific recommendations based on prior purchases or product-centric recommendations such as Amazon’s “people who bought x also bought y ” feature [15]. These refinements help consumers discover new products and receive more accurate evaluations, but they cannot completely substitute personalized recommendations that one receives from a friend or relative. It is human nature to be more interested in what a friend buys than what an anonymous person buys, to be more likely to trust their opinion, and to be more influenced by their actions. Our friends are also acquainted with our needs and tastes, and can make appropriate recommendations. A Lucid Marketing survey found that 68% of individuals consulted friends and relatives before purchasing home electronics – more than the half who used search engines to find product information [4].

Given the natural tendency of individuals to trust their friends’ opinions, several services have tried to facilitate the spread of recommendations between friends. Netflix, a DVD mail-rental company, encourages its customers to use their website to recommend movies they viewed to their friends. Yahoo Music gives its users the option of displaying the music tracks they are listening to on their instant messenger client, which can then be seen by all their Yahoo messenger contacts. It is this kind of direct recommendation via an individual’s true social network that we study in this paper.

Our analysis focuses on the recommendation referral program run by a large retailer. The program rules were as follows. Each time a customer purchases a book, music, or a movie he or she is given the option of sending emails recommending the item to friends. The first person to purchase the same item through a referral link in the email gets a 10% discount. When this happens the sender of the recommendation receives a 10% credit on their purchase.

This study is novel in that we are able, for the first time, to directly observe the effectiveness of person to

person word of mouth advertising for hundreds of thousands of products. We can see what kind of product is more likely to be purchased as a result of this type of recommendation, as well as describe the size of the cascade that results from an initial purchase and subsequent recommendations.

Although the data gives us a detailed and accurate view of recommendation cascades, it does have its limitations. The only indication of the success of a recommendation is the observation of the recipient purchasing the product through the same vendor. We have no way of knowing if the person had decided instead to purchase elsewhere, borrow, or otherwise obtain the product. The delivery of the recommendation is also somewhat different from one person simply telling another about a product they enjoy, possibly in the context of a broader discussion of similar products. The recommendation is received as a form email including information about the discount program. Someone reading the email might consider it spam, or at least deem it less important than a recommendation given in the context of a conversation. The recipient may also doubt whether the friend is recommending the product because they think the recipient might enjoy it, or are simply trying to get a discount for themselves. Finally, because the recommendation takes place before the recommender receives the product, it might not be based on a direct observation of the product. Nevertheless, we believe that these recommendation networks are reflective of the nature of word of mouth advertising, and give us key insights into the influence of social networks on purchasing decisions.

1.1 Paper roadmap

In Section 2, we describe the dataset, giving overall statistics for different product groups. Section 3 describes the observed sizes of information cascades and ties them to a simple propagation model that accounts for both the predominance of short chains and the occasional large cascades. Section 4 describes the growth of the recommendation network over time and for different kinds of products. In Section 7 we measure the effect of receiving recommendations from multiple individuals, as well as how the effectiveness of recommendations drops as more and more of them are exchanged between the same two people. We also report the dependence between the number of recommendations sent and number of resulting purchases. Section 8 shows an interesting time lag between when recommendations are made and when they are acted upon, corresponding to different times of day when individuals are likely to be shopping online or reading email. In Section 9 we show that some categories of products (corresponding to different communities of interest) have varying response to recommendations. We also show that the preference of individuals to make personal recommendations compared to posting a review on the website varies by category. Section 10 shows the influence of various product attributes on the recommendation success rate. Finally, Section 11 discusses the results and concludes.

2 The recommendation network

First we describe the recommendation network dataset.

The following information is recorded for each recommendation

1. Sender Customer ID (shadowed)
2. Receiver Customer ID (shadowed)
3. Date of Sending
4. Purchase flag (*buy-bit*)
5. Purchase Date (error-prone due to asynchrony in the servers)
6. Product identifier
7. Price

We represent this data set as a directed graph. The nodes represent customers, and a directed edge contains all the information about the recommendation. The edge (i, j, p, t) indicates that customer i recommended product p to customer j at time t .

The typical process generating edges in the recommendation network is as follows: a node i first buys a product p at time t and then it recommends it to nodes j_1, \dots, j_n . The j nodes can they buy the product and further recommend it. The only way for a node to recommend a product is to first buy it. Note that even if all nodes j buy a product, only the edge to the node j_k that first made the purchase (within a week after the recommendation) will be marked by a purchase flag. We refer to the purchase flag also as a *buy-bit*. Because the purchase flag (buy-bit) is set only for the first person who acts on a recommendation, we identify additional purchases by the presence of outgoing recommendations for a person, since all recommendations must be preceded by a purchase. We call this type of evidence of purchase a *buy-edge*.

Next we define *successful recommendations*. These are the recommendations for which we can say that they influenced the purchase of a product. All recommendations that reached a node before the *first* purchase are *successful recommendations*. To determine the time of the earliest purchase we take the minimum of the time of the earliest incoming edge marked with a buy-bit and the time of earliest outgoing edge.

In order to identify cascades, i.e. the “causal” propagation of recommendations, we focus only on the first purchase of an item. There are many cases when a person made multiple purchases of the same product, and in between those purchases she may have received new recommendations. In this case one cannot say that recommendations following the first purchase really influenced the later purchases.

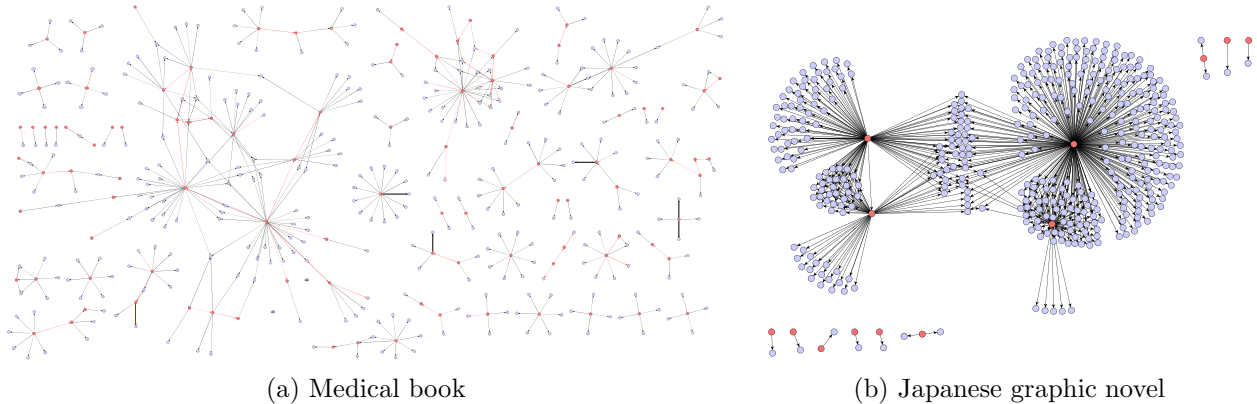


Figure 1: Examples of two product recommendation networks: (a) First aid study guide *First Aid for the USMLE Step*, (b) Japanese graphic novel (manga) *Oh My Goddess!: Mara Strikes Back*.

To avoid these issues we identify cascades by deleting *late recommendations*. Given a product recommendation network for each node, we delete all incoming recommendations that happened after the earliest purchase of the product. This way we make the network *time increasing* or *causal* — for each node all incoming edges happened before all outgoing edges. Now each connected component represents a time obeying propagation of recommendations.

Figure 1 shows two typical product recommendation networks: (a) a medical study guide and (b) a Japanese graphic novel. Throughout the dataset we observe very similar patterns. Most product recommendation networks consist of a large number of small disconnected components where we do not observe cascades. Then there is usually a small number of relatively small components where we observe recommendations propagating. We also notice bursts of recommendations and collisions (figure 1(b)). Some nodes recommend to many friends, forming a star like pattern. Sometimes nodes also receive recommendations from two or more sources. A detailed enumeration and analysis of all observed cascades for this dataset is made in [14].

2.1 Dataset characteristics

The recommendation dataset consists of 15,646,121 recommendations made among 3,943,084 distinct users. The data was collected from June 5 2001 to May 16 2003. In total, 548,523 products were recommended. 98.9% (or 542,719) of all products belong to 4 main product groups: Books, DVDs, Music and Videos.

In addition to recommendation data, we also crawled the retailer’s website to extract detailed product information. For each product we obtained the categorization into the product hierarchy, a list of similar products, sales ranks and all reviews and ratings together with time of posting. The lower sales rank, the

Group	p	n	e	e_u	n_c
Book	103,161	2,863,977	5,741,611	2,097,809	8,333,508
DVD	19,829	805,285	8,180,393	962,341	4,313,848
Music	393,598	794,148	1,443,847	585,738	2,110,959
Video	26,131	239,583	280,270	160,683	433,468
Full	542,719	3,943,084	15,646,121	3,153,676	15,191,783

Table 1: Product group recommendation network statistics. p : number of products, n : number of nodes, e : number of edges (recommendations), e_u : number of unique edges, n_c : cumulative number of nodes.

higher the sales for that particular item. There were a total of 7,556,174 reviews on these half million products, with an average of 14 per product. The average product rating was 4.3 (on a 1 to 5 scale). There were a total of 519,138 different product categories and a product belonged to 4.8 categories on the average. We also note that after two years, out of 548,523 products 5813 (1%) were discontinued (the retailer no longer provided any information about them).

Table 1 shows the sizes of various product group recommendation networks. For each product group we took recommendations on all products from the group and created a graph. Column p shows the total number of products in the product group, n is total number of nodes spanned by the group recommendation network and e is the number of edges (recommendations). The column e_u shows the number of unique edges – disregarding multiple recommendations between the same source and recipient. The last column n_c shows the cumulative number of nodes – we created individual product recommendation networks and summed-up their sizes.

Note that the number of edges e sums to the number of edges of the full network, while the number of nodes n does not, since a node can make recommendations on multiple products.

In terms of the number of different items, there are by far the most music CDs, followed by books and videos. There is a surprisingly small number of DVD titles. On the other hand DVDs account for almost half of all recommendations in the dataset. The DVD graph is also the most dense, having about 10 recommendations per node, while books and music have about 2 recommendations per node and videos have only a bit more than 1 recommendation per node. On the whole there are about 3.69 recommendations per node on 3.85 different products.

Music recommendations reached about the same number of people as DVDs but used more than 5 times fewer recommendations to achieve the same coverage of the nodes. Book recommendations reached by far the most people – 2.8 million. Notice that all networks have a very small number of unique edges. For books, videos and music the number of unique edges is smaller than the number of nodes – this suggests that the networks are highly disconnected [8]. Even if we compose a network using all the recommendations in the dataset, the largest weakly connected component contains less than 2.5% (100,420) of the nodes, and the




Figure 2: Distribution of the number of recommendations and number of purchases made by a node in the network.

second largest component has only 600 nodes.

Dividing by the cumulative number of nodes we see that the average person in a product group network sent or received recommendations on 3 book titles, 5.3 DVDs, 2.6 music titles and 1.8 videos.

Figure 2 shows the distribution of the recommendations and purchases made by a single node in the recommendation network. Notice the power-law distributions and long flat tails. The most active person made 83,729 recommendations and purchased 4,416 different items.

Even with these simple counts we can make the first few observations. It seems that some people got quite heavily involved in the recommendation program, that they tended to recommend a large number of products to the same set of friends (since the number of unique edges is so small). Books tend to have very sparsely linked communities, meaning that a small number of edges connected a lot of nodes. The same is also true for music and video. DVDs on the other hand are densely linked both in terms of the total number of recommendations and the number of unique edges. This shows that people tend to buy more DVDs and also like to recommend them to their friends, while they seem to be more conservative with books. One possible reason is that a book is bigger time investment than a DVD: one usually needs several days to read a book, while a DVD can be viewed in a single evening.

One external factor which may be affecting the recommendation patterns for DVDs is the existence of referral websites (www.dvdtalk.com). On these websites people who want to buy a DVD and get a discount, would ask for recommendations. This way there would be recommendations made between people who don't really know each other but rather have an economic incentive to cooperate. We were not able to find similar referral sharing sites for books or CDs.

3 The recommendation propagation model

A simple model that tries to capture the propagation of recommendations throughout the network assumes that each recipient of a recommendation will forward it to others if its value exceeds an arbitrary threshold that the individual sets for herself. Since exceeding this value is a probabilistic event, let's call p_t the probability that at time step t the recommendation exceeds the threshold. In that case the number of recommendations N_{t+1} at time $(t + 1)$ is given in terms of the number of recommendations at a time earlier by

$$N_{t+1} = p_t N_t. \quad (1)$$

where the probability p_t is defined over the unit interval.

Notice that, because of the probabilistic nature of the threshold being exceeded, one can only compute the final distribution of recommendation chain lengths, which we now proceed to do.

Subtracting from both sides of this equation the term N_t and dividing by it we obtain

$$\frac{N_{(t+1)} - N_t}{N_t} = p_t - 1 \quad (2)$$

If we sum both sides from the initial time to some very large time T and assume that for long times the numerator is smaller than the denominator (a reasonable assumption) we get

$$\frac{dN}{N} = \sum p_t \quad (3)$$

The left hand integral is just $\ln(N)$, and the right hand side is a sum of random variables, which in the limit of a very large uncorrelated number of recommendations is normally distributed (central limit theorem).

This means that the logarithm of the number of messages is normally distributed, or equivalently, that the number of messages passed is log-normally distributed. In other words the probability density for N is given by

$$P(N) = \frac{1}{N\sqrt{2\pi\sigma^2}} \exp \frac{-(\ln(N) - \mu)^2}{2\sigma^2} \quad (4)$$

which, for large variances describes a behavior whereby the typical number of recommendations is small (the mode of the distribution) but there are unlikely events of large chains of recommendations which are also observable.

Furthermore, for large variances, the lognormal distribution can behave like a power law for a range of values. In order to see this, take the logarithms on both sides of the equation (equivalent to a log-log plot)




Figure 3: Size distribution of cascades (size of cascade vs. count). Bold line presents a power-fit.

and one obtains

$$\ln(P(N)) = -\ln(N) - \ln(\sqrt{2\pi\sigma^2}) - \frac{(\ln(N) - \mu)^2}{2\sigma^2} \quad (5)$$

So for large σ the last term of the right hand side goes to zero, and since the the second term is a constant one obtains a power law behavior with exponent value of minus one. There are other models which produce power-law distributions of cascade sizes, but we present ours for its simplicity, since it does not depend on network topology [9] or critical thresholds in the probability of a recommendation being accepted [23].

3.1 Distribution of cascade sizes

The model predicts that sizes of cascades obey a lognormal distribution which for large variance leads to a power-law distribution with exponent 1. To identify the cascades we deleted late recommendations (section 2) and then measured the sizes of cascades (size of connected component) without the nodes that did not buy.

From figure 3 we see that the distribution of cascade sizes is heavy-tailed, but it is only for DVDs that we observe a power-law exponent close to 1. The reason is the large variance associated with the DVDs, which is not observed in the other product groups. We further investigate this in next sections.

3.2 Propagation characteristics

The stochastic model makes the general assumption that the recommendation behavior does not change with the depth in the cascade. We did find, however, that people who had received no prior recommendations for the product through the recommendation referral program behaved differently from those who had. We further examine the behavior of individuals who received a recommendation after 2, 3 and 4 propagations.

Figure 4 shows the cumulative out-degree distribution, that is the number of people who sent out at least k_p recommendations, for a product. It shows that the deeper an individual is in the cascade, if they choose to make recommendations, they tend to recommend to a greater number of people on average (the




Figure 4: The distribution of the number of recommendations sent by a user for a book with each curve representing a different depth of the user in the recommendation chain. A power law exponent γ is fitted to all but the tail of each distribution.

level	prob. buy & forward	average out-degree
0	N/A	1.99
1	0.0069	5.34
2	0.0149	24.43
3	0.0115	72.79
4	0.0082	111.75

Table 2: Statistics about individuals at different levels of the cascade.

distribution has a higher variance). This effect is probably due to only very heavily recommended products producing large enough cascades to reach a certain depth. We also observe, as is shown in Table 2, that the probability of an individual making a recommendation at all (which can only occur if they make a purchase), declines after an initial increase as one gets deeper into the cascade.

4 Recommendation network over time

We traced the evolution of the recommendation network from the start to the end of the program. The growth of the customer base was linear, adding on average 165,000 new users each month, which is an indication that the service itself was not spreading epidemically. Further evidence of non-viral spread is provided by the relatively high percentage (94%) of users who made their first recommendation without having previously received one.

We find that, over time, customers are recommending to more of their friends. This results in a densi-

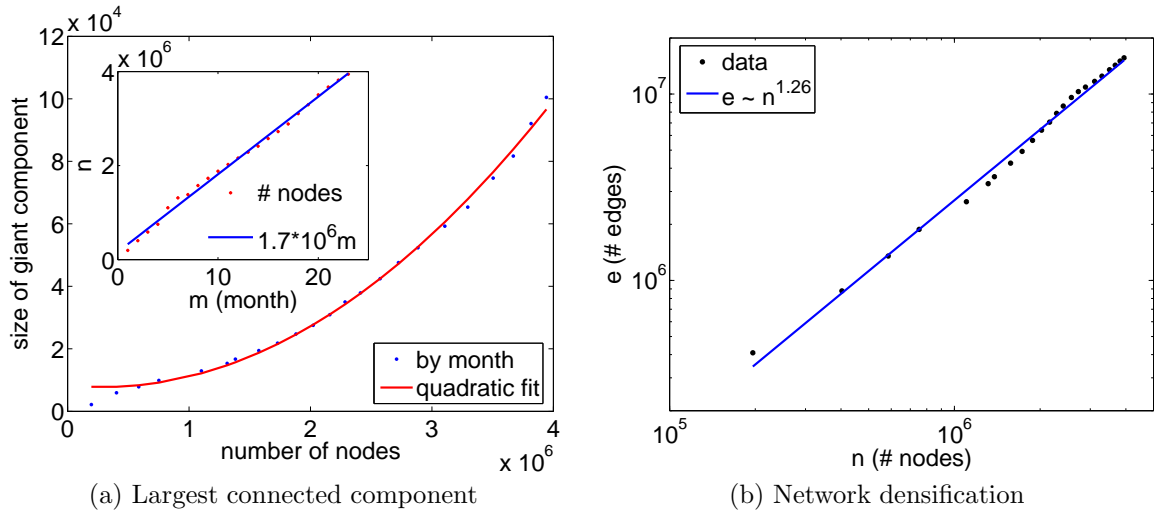


Figure 5: The size of the largest connected component of customers as a function of the cumulative number of customers using the service. The inset shows the linear growth in the number of customers over time

fication of the network [13], with the average degree of a node increasing over time proportionally to $n^{0.26}$, where n is the total number of nodes. In spite of this, the network overall stays largely disconnected.

As we saw in section 3, individual products produced large recommendation cascades only rarely, and those reached on the order of a thousand customers. Even if we aggregate recommendations over all products, the social network of customers stays largely disconnected. The largest connected component, although growing quadratically with the number of nodes, still only increases from 1% to 2.5% of the network. The diameter of this component increases from 7 hops for the 200,000 customers participating in the first month, to 16 hops for the 4,000,000 customers in aggregate over the full two years. Since the largest connected component is very small, one could attribute the increasing diameter to the formation of the giant connected component.

With the total network being largely disconnected, we sought out smaller communities, which we identified by the category of products that they bought. We found large variation in the size of the largest connected component within each category. Of the 180 thousand customers who bought books in the home and garden category, fewer than 3% belonged in the giant connected component. On the other hand, some of the top level DVD genre categories had a giant component that occupied about a fifth of the graph: 24% out of 18,000 users for westerns, 26% of 25,000 for classics, and 19% of 47,000 for anime (Japanese animated film). Other DVD categories were less well connected. For example, the giant components for the fitness and kids categories contained only 2 and 7 percent of their customers. The anime community stands out further, because this relatively small number of people was responsible for 2.5 of the 16 million recommendations in the system. We will discuss the recommendation patterns in this category and others in more detail in

Group	b_t	b_s	b_b	b_e
Book	2,859,096	83,113	65,344	17,769
DVD	837,300	75,421	17,232	58,189
Music	712,673	10,576	7,837	2,739
Video	165,109	1,376	909	467
Total	4,574,178	170,486	91,322	79,164

Table 3: Product group recommendation network purchase statistics.

section 9.

5 Purchase statistics

We now describe the purchase statistics over the product groups. We create per product recommendation networks, discard late recommendations and then count the number of nodes that purchased a product.

Table 3 shows the counts of the total number of purchases (b_t) and the total number of purchases that were made as a result of a recommendation (b_s). Note that b_t only includes purchases relating to the program: either a purchase resulting for a recommendation or a purchase preceding a recommendation. A purchase can be determined to be the result of a recommendation either via a buy-bit b_b (a person got the discount) or buy-edge (b_e), i.e. an incoming edge which happened after the outgoing edge.

We observe books having by far the highest number of purchases. But the number of purchases through recommendations is much smaller. DVDs and books have both about the same number of purchases through recommendations, and books have more than 3 times more total purchases. For books only 3% of all purchases were can be attributed to a recommendation, for DVDs this increases to 9%, for music and video it drops again to 1.4% and 0.8%, respectively.

Looking at the number of purchases that got discount, we observe that for books, music and video 78.6, 74.1 and 66.0 percent of purchases through recommendations resulted in a discount. For DVDs there existed special websites where people listed email addresses to receive DVD recommendations (and eventually got discounts). Based on this we expected that a large majority of DVD purchases would result in a discount. Comparing the number of purchases through recommendations and purchases with discounts, we observe that only 22.8% of the DVD purchases made through the recommendation system resulted in a discount.

Given the total number of recommendations and purchases influenced by recommendations we can estimate how many recommendations need to be independently sent over the network to induce a new purchase. Using this metric books have the most influential recommendations followed by DVDs and music. For books one out of 69 recommendations resulted in a purchase. For DVDs it increases to 108 recommendations per purchase and further increases to 136 for music and 203 for video. Note that here we are only comparing

Group	Number of nodes		
	Purchases	Forward	Percent
Book	65,391	15,769	24.2
DVD	16,459	7,336	44.6
Music	7,843	1,824	23.3
Video	909	250	27.6
Total	90,602	25,179	27.8

Table 4: Fraction of people that purchase and also recommend forward. *Purchases*: number of nodes that purchased. *Forward*: nodes that purchased and then also recommended the product.

the total number of purchases that can be attributed to recommendations with the total number of recommendations made. In these first observations we are not yet considering the network structure, so the lower effectiveness of DVD recommendations can be attributed to the high density of the recommendation network.

6 Forward recommendations

Not all people who make a purchase also decide to give recommendations. So we estimate what fraction of people that purchase also decide to recommend forward. To obtain this information we can only use the nodes with purchases that resulted in a discount.

Table 4 shows that only about a third of people that purchase also recommend the product forward. The ratio of forward recommendations is much higher for DVDs than for other kinds of products. Videos also have a higher ratio of forward recommendations, while books have the lowest. This shows that people are most keen on recommending movies, while more conservative when recommending books and music.

7 Effectiveness of recommendations

So far we only looked into the aggregate statistics of the recommendation network. Next we ask questions about the effectiveness of recommendations in the recommendation network itself. First we analyze the probability of purchasing as one gets more and more recommendations. Next we measure the effectiveness of recommendations as two persons exchange more and more recommendations. One might expect that the recommendation channel between persons who exchange more recommendations gets more efficient. Lastly we check the recommendation network from the perspective of the sender of the recommendation. Does a node that makes more recommendations also influence more purchases? Do people that are more involved in the recommendation network also tend to make more purchases?

In the following sections we answer these and some additional questions. For all experiments we first create

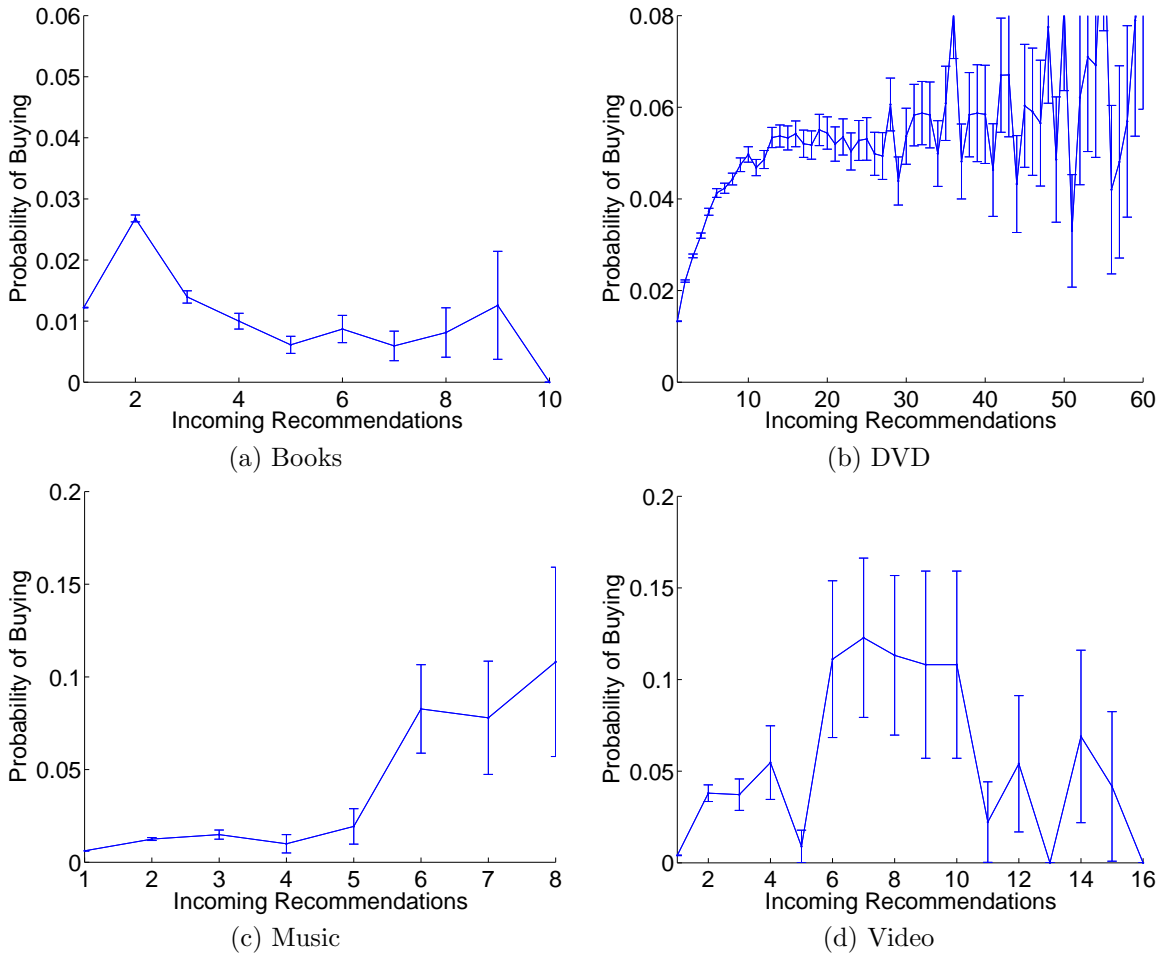


Figure 6: Probability of buying given a number of incoming recommendations.

individual product recommendation networks and delete late recommendations. With this preprocessing we ensure that all incoming recommendations happened before the first purchase of a product.

7.1 Probability of buying in the number of incoming recommendations

First we examine how the probability of purchasing changes as one gets more and more recommendations. We would expect that a person is more likely to buy a product the more recommendations she gets for that particular product. On the other hand one would also think that there is a saturation point – if a person hasn't bought a product after a given number of recommendations, they are not likely to change their minds after receiving even more of them. One can ask how many recommendations are too many.

Figure 6 shows the probability of purchasing a product as a function of the number of incoming recommendations on the product. As we move to higher numbers of incoming recommendations, the number of observations drops rapidly. For example, there were 5 million cases with 1 incoming recommendation on a




Figure 7: Probability of buying given a number of incoming recommendations. Zoom-in for DVD recommendations (Figure 6(b)).

book, and only 58 cases where a node got 20 incoming recommendations on a single book. The maximum was 30 incoming recommendations. For these reasons we cut off the the plot when the number of observations becomes too small and the error bars become too large.

Figure 6(a) shows that overall, book recommendations are rarely followed, and even more surprisingly as more and more recommendations are received their success decreases. We observe a peak in probability of buying at 2 incoming recommendations and then a slow drop.

For DVDs (Figure 6(b) and Figure 7) we observe a saturation around 10 incoming recommendations. This means that after a person gets a 10 recommendations they become immune to them – their probability of buying does not increase anymore. The number of observations is 2.5 million at 1 incoming recommendation and 100 at 60 incoming recommendations. The maximal number of incoming recommendations for a person is 172 (and that person did not buy), and someone purchased a DVD after 169 recommendations.

In the case of music and videos the number of observations was very small (10 cases with 10 incoming recommendations). Music seems to be immune to recommendations – we observe a very low effectiveness of recommendations at the beginning (up to 5 recommendations), then the probability of buying increases up to 0.1. Since the number of observations drops bellow 10, it is hard to draw any conclusions. Videos show different behavior. The probability of buying reaches a plateau between 6 to 10 recommendations (50 observations for each point). As the number of recommendations increases the probability drops. We also observe video recommendations to be quite influential as one receives more than 5 recommendations.




Figure 8: The effectiveness of recommendations as a function of the total number of exchanged recommendations.

7.2 Effectiveness of subsequent recommendations

We showed that getting more recommendations generally increases the chance of purchases, we also showed that it saturates at around 10 incoming recommendations. Next we analyze how the effectiveness of recommendations changes as two persons exchange more and more recommendations. A large number of exchanged recommendations between two persons can be a sign of trust and influence, but a sender of recommendations can be also perceived as a spammer, who recommends everything she buys. The premise is that if a person recommends a few products to a friend, then she may consider buying them, but if a person floods the friend with all sorts of recommendations, then she probably won't react to them anymore.

We measured the effectiveness of recommendations as a function of the total number of previously exchanged recommendations between the two nodes. We conducted the experiment in the following way. For every recommendation r on some product p between nodes u and v we first determine how many recommendations were exchanged between u and v before r , and then we check whether v , the recipient of recommendation, purchased p after recommendation r arrived. For the experiment we consider only node pairs (u, v) , where there were at least 10 recommendations sent from u to v . We perform the experiment using only recommendations from the same product group.

Figure 8 shows the probability of buying as a function of total number of exchanged recommendations between two persons. For books we observe that the effectiveness of recommendation remains about constant up to 3 exchanged recommendations. As the number of exchanged recommendations increases, the probability of buying starts to decrease to about half of the original value and then levels off. For DVDs we observe an immediate and consistent drop. At one exchanged recommendation, the probability of buying is

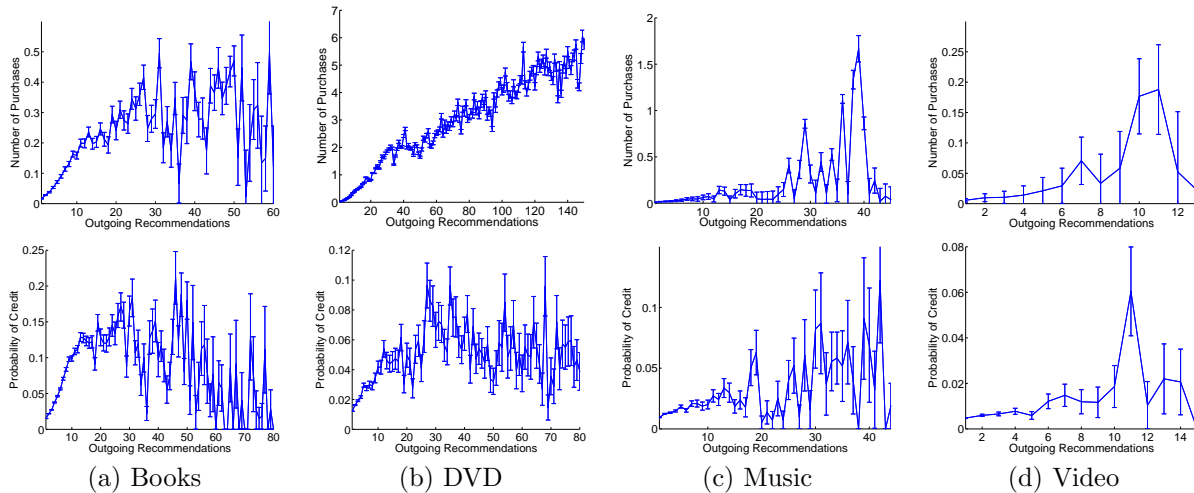


Figure 9: Top row: Number of resulting purchases given a number of outgoing recommendations. Bottom row: Probability of getting a credit given a number of outgoing recommendations.

6.3%, at five exchanged recommendations it drops to 3.7% and later even further down to around 3%. This experiment shows that recommendations start to lose effect after more than two or three are passed between the persons. We performed the experiment also for video and music, but the number of observations was too low and the measurements were noisy.

7.3 Effectiveness of outgoing recommendations

So far we examined the data from the viewpoint of the receiver of the recommendation. Now we look from the viewpoint of the sender. The two interesting questions are: how does the probability of getting a 10% credit change with the number of outgoing recommendations; and given a number of outgoing recommendations, how many purchases will they influence?

One would expect that recommendations would be the most effective when recommended to the right subset of friends. If one is very selective and recommends to too few friends then the chances of success are slim. On the other hand recommending to everyone we know and spamming them with recommendations may have limited returns as well.

Figure 9 shows how the average number of people that purchased changes with number of outgoing recommendations (top row). The bottom row shows the probability of getting a 10% credit given a number of outgoing recommendations.

For books, music, and videos the number of purchases soon saturates and starts to drop (figure 9 top row). It grows fast up to 10 outgoing recommendations and then the trend either slows or starts to drop. DVDs exhibit different behavior. The expected number of purchases increases throughout.

A possible explanation for this phenomenon is the following: the DVD recommendation network is densely connected: it has about 10 recommendations per node and around 400 per product, which is an order of magnitude more than other product groups. The recommendations very soon start to collide, meaning that a purchase was preceded by multiple recommendations. Every sender who sent a recommendation to the person then records a purchase through the recommendation. So we find that the more outgoing recommendations there are, the greater the number of collisions and hence the higher the number of purchases.

The bottom row of figure 9 plots the probability of getting a 10% credit as a function of the number of outgoing recommendations. Except for DVDs we observe the same qualitative behavior as in the top row. For books, DVDs and music the the probability reaches a maximum at around 30 outgoing recommendations. The probability of getting credit is highest for books where it climbs up to 15%. Notice that the DVD curve now saturates, which is further evidence that the increasing trend of number of purchases can be attributed to high density and collisions of DVD recommendations. This means that many different individuals are recommending to the same person, and after that person makes a purchase, even though all of them made a 'successful recommendation' by our definition, only one of them receives a credit.

For books and DVDs we actually observe super-linear growth in the number of purchases. Figure A-1 shows the power fits of the probability of buying and number of purchases curves. The super-linear trend stops at around 10 for books and 40 for DVDs. For the number of purchases the coefficient of power fit is 1.09 for books and 1.34 for DVDs. For the probability of getting a credit, the power coefficient for books remains about the same, while for DVDs the situation changes dramatically – the probability of getting a credit (figure A-1(b)) increases sub-linearly.

7.4 Probability of buying given the total number of incoming recommendations

The collisions of recommendations are the dominant feature of the DVD recommendation network. Book recommendations have the highest chance of getting a credit, but DVD recommendations cause the most purchases. So far it seems people are very keen on recommending various DVDs, while very conservative on recommending books. But how does the behavior of customers change as they get more involved into the recommendation network? We would expect that most of the people are not heavily involved, so their probability of buying is not high. In the extreme case we would expect to find people who buy almost everything they get recommendations on.

There are two ways to measure the involvedness of a person in the network: by the total number of incoming recommendations (on all products) or the total number of different products they were recommended. For every purchase of a book at time t , we count the number of different books (DVDs, ...) the person




Figure 10: The probability of buying a product given a number of different products a node got recommendations on.

received recommendations for before time t . As in all previous experiments we delete late recommendations.

We show the probability of buying as a function of the number of different products recommended on figure 10. Figure A-2 plots same data but with the total number of incoming recommendations on the X axis.

We observe two trends. For books and videos (figure 10(a) and (c)) the probability of buying is the highest when a person got recommendations on just 1 item, as the number of incoming recommended products increases to 2 or more the probability of buying quickly decreases and then stabilizes.

Movies exhibit different behavior (figure 10(b) and (d)). A person is more likely to buy the more recommendations she gets. For DVDs the peak is at around 15 incoming products, while for videos there is no such peak – the probability slowly and steadily increases. Interestingly for DVDs the distribution reaches its low at 2 and 3 items, while for videos it lies somewhere between 3 and 8 items.

The results suggest that books and music buyers tend to be conservative and focused. On the other hand there are people who like to buy movies in general. Buying a book seems to be a larger investment of time than buying a movie. One can see a movie in an evening, while reading a book requires more effort.

The other difference between the book and music recommendations in comparison to movies are the recommendation referral websites where people could go to get recommendations. One could see these websites as recommendation subscription services – posting one’s email on a list results in a higher number of incoming recommendations. For movies, people with a high number of incoming recommendations “subscribed” to them and thus expected/wanted the recommendations. On the other hand people with high numbers of incoming book or music recommendations did not “sign up” for them, so they may perceive recommendations as spam and thus the influence of recommendations drops.

Another evidence of the existence of recommendations referral websites includes the DVD recommendation network degree distribution. The DVDs follow a power law degree distribution with an exception of a peak at out-degree 50. Other plots of DVD recommendation behavior also exhibited abnormalities at around 50 recommendations. We believe these can be attributed to the recommendation referral websites.

8 The time lag between recommendation and purchase

The recommendation referral program encourages people to purchase as soon as possible after they get a recommendation, which maximizes the probability of getting a discount. We study the time lag between the recommendation and the purchase of different product groups. For every purchase through a recommendation we find the time of the earliest and latest incoming recommendation before the purchase.

We present the histograms of the time lag between the purchase and the last recommendation with a bin size of 1 day (figure 11). Around 35% of book and music purchases occurred within a day after the last recommendation was received. For movies the percentage of purchases within a day of the recommendations increases to about 40%. Between day 2 and 7 only 42% of the DVD purchases occur (55% for other groups). For DVDs 16% purchases occur more than a week after last recommendation, while this drops to 10% for other product groups.

By taking the last recommendation a person receives on the product, we measured the behavior of the “time to think”. Observing the time lag between the first recommendation and purchase gives the time needed to gather the sufficient number of recommendations to trigger a purchase. Figure A-3 shows these plots.

We observe about 35% purchases to happen within a day after the first recommendation was received. For DVDs this percentage is only 23%, and more strikingly, about 40% of the purchases occurred more than




Figure 11: The time between the recommendation and the actual purchase. The bucket size is 1 day. We use all purchases.

a week after the first recommendation was received. For all products the amount of sales from day 2 to day 7 was around between 50% and 55% (55% for books, 51% for music and 50% for video). For DVDs it was much smaller, only 32%.

This shows that for DVDs people need to accumulate more incoming recommendations to react, but when they decide to purchase, they do this sooner than for other products. One could also hypothesize that DVD recommendations were not that successful, since the lag between recommendation and purchase tends to be large. This suggests that people tend to purchase DVDs independently of recommendations. Maybe the density of recommendations in the network was so high that we observed many successful recommendations just by chance – a person would buy the DVD regardless of the recommendation, which is in agreement with results from table 3 (columns b_b and b_e) where DVDs in comparison to other groups have a small number of purchases with a discount and a large number of purchases through the network.




Figure 12: Time of day for purchases and recommendations for the whole recommendation network. (a) shows the distribution of recommendations over the day, (b) shows all purchases and (c) shows only purchases that resulted in getting discount.

We also examine the distribution of time lags for purchases that resulted in a discount (figure A-4). When a person recommends the product, among all the receivers only the first one to buy a product would get a discount. Notice the nice daily periodicity and fast decay. The logarithmic fit (dashed line) shows how fast the buying rate is decaying. For books the coefficient in front of the logarithm is -118 , for DVD -25.4 , Music -14.5 and Video -1.45 . This shows that people who decided to buy books made their decision soon. On the other hand people could buy videos quite late and still get the discount.

8.1 Time of purchase

We measure how the number of recommendations varies over the day. For this purpose we created 3 experiments: first, we examined how the total number of recommendations varies by hour of day (figure 12(a)). Next we examined how the purchases vary over the day (figure 12(b)). And lastly we plotted the number of purchases which resulted in a discount as a function of an hour of the day (figure 12(c)).

The recommendations and purchases follow the same pattern. The only little difference is that purchases reach a sharper peak in the afternoon (after 3pm Pacific Time, 6pm Eastern time). The purchases that resulted in a discount look like a negative image of the first two figures. This means that most of discounted purchases happened in the morning when the traffic (number of purchases/recommendations) on the retailer’s website was low. This makes a lot of sense since most of the recommendations happened during the day, and if the person wanted to get the discount (had to be the first one to purchase), she had the highest chances when the traffic on the website was the lowest.

9 Recommendation effectiveness by category

Social networks are a product of circumstances that bring people together. We tend to choose our friends and acquaintances from the pool of people we had a chance to spend time with, whether it be in the neighborhood, in school, at work, at religious activities, or in pursuit of a hobby. But not all circumstances are equally influential in the formation of a friendship. A study [1] of student homepages found that individuals participating in group activities such as sports were more likely to know one another than individuals who share an activity which is usually pursued individually. The same study found that the smaller the number of people sharing a particular interest, the more likely they are know one another (whether it is the friendship that arose from the shared interest, or the interest that was shared between friends was not differentiated).

In certain contexts affiliation relating to a professional activity is more effective at conducting an action. For example, in small world experiments, where participants attempt to reach a target individual through their chain of acquaintances, profession trumped geography, which in turn was more useful in locating a target than attributes such as religion or hobbies [12, 22]. In the context of product recommendations, we can ask whether a recommendation for a work of fiction, which may be made by any friend or neighbor, is more or less influential than a recommendation for a technical book, which may be made by a colleague at work or school.

By looking at recommendation behavior for various product categories, we can discern a few contexts where recommendations are more influential. The categories are assigned by the merchant's site and each product may belong to several categories which are organized hierarchically. We find that books of a professional or technical nature are recommended and the recommendation is accepted more often than for books tied to leisure activities such as fictional novels or hobbies such as gardening which may be pursued individually. There are sometimes of course exceptions, for example, Japanese anime DVDs have a strong following in the US, and this is reflected in their frequency and success in recommendations. Another example is that of gardening. In general, recommendations for books relating to gardening have only a modest chance of being accepted, which agrees with the individual prerogative that accompanies this hobby. At the same time, orchid cultivation can be a highly organized and social activity, with frequent 'shows' and online communities devoted entirely to orchids. Perhaps because of this, the rate of acceptance of orchid book recommendations is twice as high as those for books on vegetable or tomato growing.

9.1 Results by book category

Table 5 shows recommendation trends for all top level book categories by subject. For clarity, we grouped the results by 4 different category types: fiction, personal/leisure, professional/technical, and nonfiction/other.

category	n_p	n	cc	r_{p1}	$r_{p1}/$ r_{p2}	v_{av}	$c_{av}/$ r_{p1}	p_m	$b * 100$
Books general	370230	2,860,714	1.87	5.28	10.50	4.32	1.41	14.95	3.12
Fiction									
Children's Books	46,451	390,283	2.82	6.44	7.38	4.52	1.12	8.76	2.06**
Literature & Fiction	41,682	502,179	3.06	13.09	10.68	4.30	0.57	11.87	2.82*
Mystery and Thrillers	10,734	123,392	6.03	20.14	14.71	4.08	0.36	9.60	2.40**
Science Fiction & Fantasy	10,008	175,168	6.17	19.90	11.22	4.15	0.64	10.39	2.34**
Romance	6,317	60,902	5.65	12.81	15.52	4.17	0.52	6.99	1.78**
Teens	5,857	81,260	5.72	20.52	7.08	4.36	0.41	9.56	1.94**
Comics & Graphic Novels	3,565	46,564	11.70	4.76	11.94	4.36	2.03	10.47	2.30*
Horror	2,773	48,321	9.35	21.26	12.23	4.16	0.44	9.60	1.81**
Personal/Leisure									
Religion and Spirituality	43,423	441,263	1.89	3.87	8.93	4.45	1.73	9.99	3.13
Health Mind and Body	33,751	572,704	1.54	4.34	9.66	4.41	2.39	13.96	3.04
History	28,458	28,3406	2.74	4.34	11.81	4.30	1.27	18.00	2.84
Home and Garden	19,024	180,009	2.91	1.78	11.69	4.31	3.48	15.37	2.26**
Entertainment	18,724	258,142	3.65	3.48	12.03	4.29	2.26	13.97	2.66*
Arts and Photography	17,153	179,074	3.49	1.56	12.27	4.42	3.85	20.95	2.87
Travel	12,670	113,939	3.91	2.74	12.15	4.26	1.87	13.27	2.39**
Sports	10,183	120,103	1.74	3.36	13.39	4.34	1.99	13.97	2.26**
Parenting and Families	8,324	182,792	0.73	4.71	8.39	4.42	2.57	11.87	2.81
Cooking Food and Wine	7,655	146,522	3.02	3.14	10.79	4.45	3.49	13.97	2.38*
Outdoors & Nature	6,413	59,764	2.23	1.93	11.69	4.42	2.50	15.00	3.05
Professional/Technical									
Professional & Technical	41,794	459,889	1.72	1.91	9.87	4.30	3.22	32.50	4.54**
Business and Investing	29,002	476,542	1.55	3.61	11.51	4.22	2.94	20.99	3.62**
Science	25,697	271,391	2.64	2.41	9.62	4.30	2.42	28.00	3.90**
Computers and Internet	18,941	375,712	2.22	4.51	21.06	3.98	3.10	34.95	3.61**
Medicine	16,047	175,520	1.08	1.41	9.99	4.40	4.19	39.95	5.68**
Engineering	10,312	107,255	1.30	1.43	11.76	4.14	3.85	59.95	4.10**
Law	5,176	53,182	2.64	1.89	13.28	4.25	2.67	24.95	3.66*
Nonfiction-other									
Nonfiction	55,868	560,552	2.03	3.13	11.35	4.29	1.89	18.95	3.28**
Reference	26,834	371,959	1.94	2.49	9.96	4.19	3.04	17.47	3.21
Biographies and Memoirs	18,233	277,356	2.80	7.65	11.49	4.34	0.90	14.00	2.96

Table 5: Statistics by book category: n_p :number of products in category, n number of customers, cc percentage of customers in the largest connected component, r_{p1} av. # reviews in 2001 – 2003, r_{p2} av. # reviews 1st 6 months 2005, r_{p1}/r_{p2} attrition rate, v_{av} average star rating, c_{av} average number of people recommending product, c_{av}/r_{p1} ratio of recommenders to reviewers, p_m median price, b ratio of the number of purchases resulting from a recommendation to the number of recommenders. The symbol ** denotes statistical significance at the 0.01 level, * at the 0.05 level.

Fiction encompasses categories such as Sci-Fi and Romance, as well as children's and young adult books. Personal/Leisure encompasses everything from gardening, photography and cooking to health and religion.

We calculated several different statistics for each category. Unfortunately we did not have the sales data for the products during the time period that the recommendation program was run. Accurate sales statistics would have allowed us to approximate the percentage of purchases that resulted in or were a result

of a recommendation. We did have a sample of sales ratings for 10,000 books at the very beginning of the recommendation program, and the full list of sales ratings in June 2005 for products that were still available. Because the sales rating of each product naturally varies over time, the sales were only weakly correlated with recommendation volume, so we instead focused our analysis on reviews which were written during the three calendar years, 2001-2003, that the program was active.

As an indication of the longevity of interest in a product, we took the ratio of the number of reviews written for a product in the category for the period from January 2001 to December 2003 and the number of reviews written for that same product in the first six months of 2005. If interest in books did not fade over time, we would expect to find six times as many reviews in the three year period than in the six month one. Some book categories aged much more gracefully than others. The most robust categories are Children's books (7.38), Teen (7.08), Parenting and Families (8.39) and Religion and Spirituality (8.93). Technical books, especially those on the subject of Computers and the Internet (21.06), seemed to date most quickly. Also, categories consisting mostly of page-turners, such as mystery and thrillers (14.71) and Romance (15.52) faded more rapidly than general literature and fiction (10.68).

A second feature we examined was the relative number of recommendations to reviews posted on the site. Surprisingly, we found that the number of people making personal recommendations was only a few times greater than the number of people posting a public review on the website. For some product categories, such as some VHS genres, the ratio was even below one. For books, we observe that fiction books have relatively few recommendations compared to the number of reviews, while professional and technical books have more recommendations than reviews. This could reflect several factors. One is that people feel more confident reviewing fiction than technical books. Another is that they hesitate to recommend a work of fiction before reading it themselves, since the recommendation must be made at the point of purchase. Yet another explanation is that the median price of a work of fiction is lower than that of a technical book. This means that the discount received for successfully recommending a mystery novel or thriller is lower and hence people have less incentive to send recommendations.

Finally, we measured the efficacy of recommendations by observing the ratio of the number of purchases occurring following a recommendation to the number of recommenders for each book subject category. On average, only 2% of the recommenders of a book received a discount because their recommendation was accepted, and another 1% made a recommendation that resulted in a purchase, but not a discount. We did not consider purchases that occurred more than a week after a recommendation, because those did not qualify for a discount, and we can be less certain that they were a response to a recommendation.

We observed marked differences in the response to recommendation for different categories of books. Fiction in general was not very effectively recommended, with only around 2% of recommenders succeeding.

The efficacy was a bit higher (around 3%) for non-fiction books dealing with personal and leisure pursuits, but was significantly higher in the professional and technical category. Medical books had nearly double the average rate of recommendation acceptance. This could be in part attributed to the higher median price of medical books and technical books in general. As we will see in Section 10, a higher product price increases the chance that a product recommendation will be accepted.

Recommendations were also more likely to be accepted for certain religious categories: 4.3% for Christian living and theology and 4.8% for Bibles. In contrast, books not tied to organized religions, such as ones on the subject of new age (2.5%) and occult (2.2%) spirituality, had lower recommendation effectiveness. These results raise the interesting possibility that individuals have greater influence over one another in an organized context, for example through a professional contact or a religious one.

Finally, we can control for the possibility that individuals purchase professional books because they are more likely to accept recommendations on a topic that they are passionate about. We examined the personal and leisure category type, where books dealing with leisure pursuits, such as arts, photography, gardening, cooking, and travel were often recommended, but the recommendation was not accepted as often as it was for technical books.

9.2 Results by DVD, VHS and Music category

Unlike books, DVDs, VHS and Music CDs (with the exception of possibly the educational category for videos), fell predominantly in the leisure and entertainment category. We therefore did not expect the same distinction between professional and personal recommendations, but rather were seeking out product categories which may be enjoyed by smaller, niche communities of viewers and listeners. Categories such as movie classics and international and art house movies were more successfully recommended than comedy or westerns. One of the most striking features of DVD purchase patterns is the dominance and success of anime recommendations. While only 2% of DVDs sold in America are anime¹, fully 22% of the DVD recommenders in our dataset recommended anime titles. There was not only a disproportionate number of anime recommenders, but these recommenders were unusually successful with 30 accepters per 100 anime recommenders. For DVDs in other categories the success rate was about 8%. This observation suggests that the anime community is either very closely knit or was able to use online forums effectively to exchange email addresses in order to obtain discounts. One anime fan commented on the existence of a specialized list specifically for anime DVDs: “If you buy a lot of anime dvds, there is a large list over at the AnimeonDvd forums. I get the discounts for anything new that comes out.”²

¹<http://www.channel4.com/film/reviews/feature.jsp?id=113768&page=2>

²<http://www.cheapassgamer.com/forums/archive/index.php/t-19982.html>

DVD and Video recommendations were accepted at substantially different rates, possibly because of the rising popularity of the DVD format during the past few years. Recommenders were 10 times less likely to successfully recommend a video cassette than a DVD, which may reflect several factors. The first is that customers may be unwilling to invest in a purchase in the declining VHS format. An exception were educational videos, with a 1.5% rate of success, which may reflect teachers purchasing videos for older VCR players at school. Another possibility is that video content in general has a low conversion rate for purchases, but that the observed purchases are a reflection of the well-organized DVD websites which allowed strangers to grant one another discounts. To our knowledge, there were no such websites for titles in VHS format. Anime was the only VHS category that had a significantly higher success rate at 2%. Music similarly did not show notable difference in recommendation success by top-level category, which is to say that customers were about equally successful in recommending classical CDs and rap. Tables A-1, A-2 and A-3 in the Appendix show the statistics for DVD, VHS, and Music titles broken down by category.

10 Regressing the recommendation success

So far we have examined how various aspects of a product recommendation network individually influence purchases. Now, we ask how the features of the product interact with the success of recommendations. We describe each product with a set of attributes and fit a regression line to model the success of recommendations.

Each product is described with the following attributes:

- r : number of recommendations
- n_s : number of senders of recommendations
- n_r : number of recipients of recommendations
- p : price of the product
- v : number of reviews of the product
- t : average product rating

The dependent variable s is the success rate of individual product recommendations. We obtain s the same way as in section 9, by dividing the total number purchases made through recommendations with the number of senders of the recommendations. We decided to use this kind of normalization, rather than normalizing by the total number of recommendations sent, in order not to penalize communities with lots of big stars, i.e. a few individuals sending out many recommendations (figure 1(b)).

Variable	Coefficient
const	-0.940 (0.025)***
$\ln(r)$	0.426 (0.013)***
$\ln(n_s)$	-0.782 (0.004)***
$\ln(n_r)$	-1.307 (0.015)***
$\ln(p)$	0.128 (0.004)***
$\ln(v)$	-0.011 (0.002)***
$\ln(t)$	-0.027 (0.014)*
R^2	0.74

Table 6: Regressing per product recommendation success rate. A dependent variable is the log of the recommendation success rate, $\ln(s)$. For each coefficient we provide the standard error and the symbol *** denotes statistical significance at the 0.01 level, ** at the 0.05 level and * at the 0.1 level.

The pairwise correlation matrix is given in table A-4. From the original set of half a million products we remove all products that have no purchases made through recommendations or for which the price was not given. We end up with 48,218 data points. Since most of the variables follow a heavy tailed distribution, we take their logarithms.

Table 6 shows the regression coefficients. With exception of the average rating, they are all significant. The only two attributes with a positive coefficient are the number of recommendations and price. This shows that more expensive and more recommended products have higher success rate. The number of senders and receivers have large negative coefficients.

All this shows that successful products are the more likely to be not so widely popular products, which have fewer reviews. While on the other hand they have lots of recommendations with a small number of senders and receivers, which suggests a very dense recommendation network where lots of recommendations were exchanged between a small community of people.

11 Discussion and Conclusion

Our analysis paints a fairly complete and varied picture of incentivised recommendation patterns for the recommendation program. Although the retailer may have hoped to boost its revenues through viral marketing, the additional purchases that resulted from recommendations, are just a drop in the bucket of sales that occur through the website. Nevertheless, we were able to obtain a number of interesting insights into how viral marketing works that challenge common assumptions made in epidemic and rumor propagation modeling.

Firstly, it is frequently assumed in epidemic models that every time individuals interact they have equal probability of being infected. Contrary to this we observe that the probability of infection decreases with repeated interaction. Marketers should take heed that even if viral marketing works initially, providing

excessive incentives for customers to recommend products could backfire by weakening the credibility of the very same links they are trying to take advantage of.

Traditional epidemic and innovation diffusion models also often assume that individuals either have a constant probability of 'converting' every time they interact with an infected individual or that they convert once the fraction of their contacts who are infected exceeds a threshold. In both cases, an increasing number of infected contacts results in an increased likelihood of infection. Instead, we find that the probability of purchasing a product increases with the number of recommendations received, but quickly saturates to a constant and relatively low probability. This means that individuals are often impervious to the recommendations of their friends, and will resist buying items that they do not want.

In network-based epidemic models, extremely highly connected individuals play a very important role. For example, in needle sharing and sexual contact networks these nodes become the "super-spreaders" by infecting a large number of people. For power-law networks it has been shown that the presence of high degree hubs results in an absence of an epidemic threshold: no matter how low the probability of transmission, the disease will persist in the network [16]. But these models assume that a high degree node has an equal probability of infecting each of its neighbors as a low degree node does. In contrast, we find that there are limits to how influential high degree nodes are in the recommendation network. As a person sends out more and more recommendations past a certain number on a particular product, the success per recommendation declines. This would seem to indicate that individuals have influence over a few of their friends, but not everybody they know.

We also presented a simple multiplicative stochastic model that allows for the presence of relatively large cascades for a few products, but reflects well the general tendency of recommendation chains to terminate after just a short number of steps. Finally, we saw that the effectiveness of recommendations varies by category, with more successful recommendations being made on technical or religious books, which presumably are placed in the social context of a school, workplace or place of worship. So despite the relative ineffectiveness of the viral marketing program in general, we were able to find a number of new insights which we hope will have general applicability to future models of viral information spread.

References

- [1] L. A. Adamic and E. Adar. Friends and neighbors on the web. *Social Networks*, 25(3):211–230, 2003.
- [2] Anonymous. Profiting from obscurity: What the long tail means for the economics of e-commerce. *Economist*, 2005.

- [3] Erik Brynjolfsson, Yu Hu, and Michael D. Smith. Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science*, 49(11):1580–1596, 2003.
- [4] Kevin Burke. As consumer attitudes shift, so must marketing strategies. study suggests rethinking the customer experience. <http://www.audioholics.com/contact/ConsumerAttitudesMarketingStrategiesStudy.html>, 2003.
- [5] Judith Chevalier and Dina Mayzlin. The effect of word of mouth on sales: Online book reviews. <http://gsbwww.uchicago.edu/research/workshops/wae/judy.pdf>, 2004.
- [6] J.S. Coleman, H. Menzel, and E. Katz. The diffusion of an innovation among physicians. *Sociometry*, 20:253–270, 1957.
- [7] Gerald F. Davis and Henrich R. Greve. Corporate elite networks and governance changes in the 1980s. *American Journal of Sociology*, 103:1–37, 1997.
- [8] P. Erdős and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 5:17–61, 1960.
- [9] D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins. Information diffusion through blogspace. In *WWW '04*, 2004.
- [10] S. Jurevics. What exactly is viral marketing? *Red Herring*, 78:110–112, 2000.
- [11] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of influence in a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 137–146, Washinton DC, 2003. ACM Press.
- [12] P. Killworth and H. Bernard. Reverse small world experiment. *Social Networks*, 1:159–192, 1978.
- [13] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. Graphs over time: Densification laws, shrinking diameters and possible explanations. In *Proceedings of the Eleventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Chicago, IL, 2005. ACM Press.
- [14] Jurij Leskovec, Ajit Singh, and Jon Kleinberg. Patterns of influence in a recommendation network. 2005.
- [15] G. Linden, B. Smith, and J. York. Amazon.com recommendations: item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1):76–80, 2003.

- [16] R. Pastor-Satarros and A. Vespignani. Epidemic spreading in scale free networks. *Physical Review Letters*, 86:3200, 2001.
- [17] Paul Resnick and Richard Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebays reputation system. www.si.umich.edu/~presnick/papers/ebayNBER/RZNBERBodegaBay.pdf, 2001.
- [18] Matt Richardson and Pedro Domingos. Mining knowledge-sharing sites for viral marketing. In *Proceedings of the Eight International Conference on Knowledge Discovery and Data Mining*, pages 61–70, Edmonton, Canada, 2002. ACM Press.
- [19] Everett M. Rogers. *Diffusion of Innovations*. Free Press, New York, fourth edition, 1995.
- [20] B. Ryan and N.C. Gross. The diffusion of hybrid corn in two iowa communities. *Rural Sociology*, 8:15–24, 1943.
- [21] David Strang and Sarah A. Soule. Diffusion in organizations and social movements: From hybrid corn to poison pills. *Annual Review of Sociology*, 24:265–290, 1998.
- [22] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 32:425–443, 1969.
- [23] D.J. Watts. A simple model of global cascades on random networks. *PNAS*, 99(9):4766–5771, Apr 2002.

Appendix

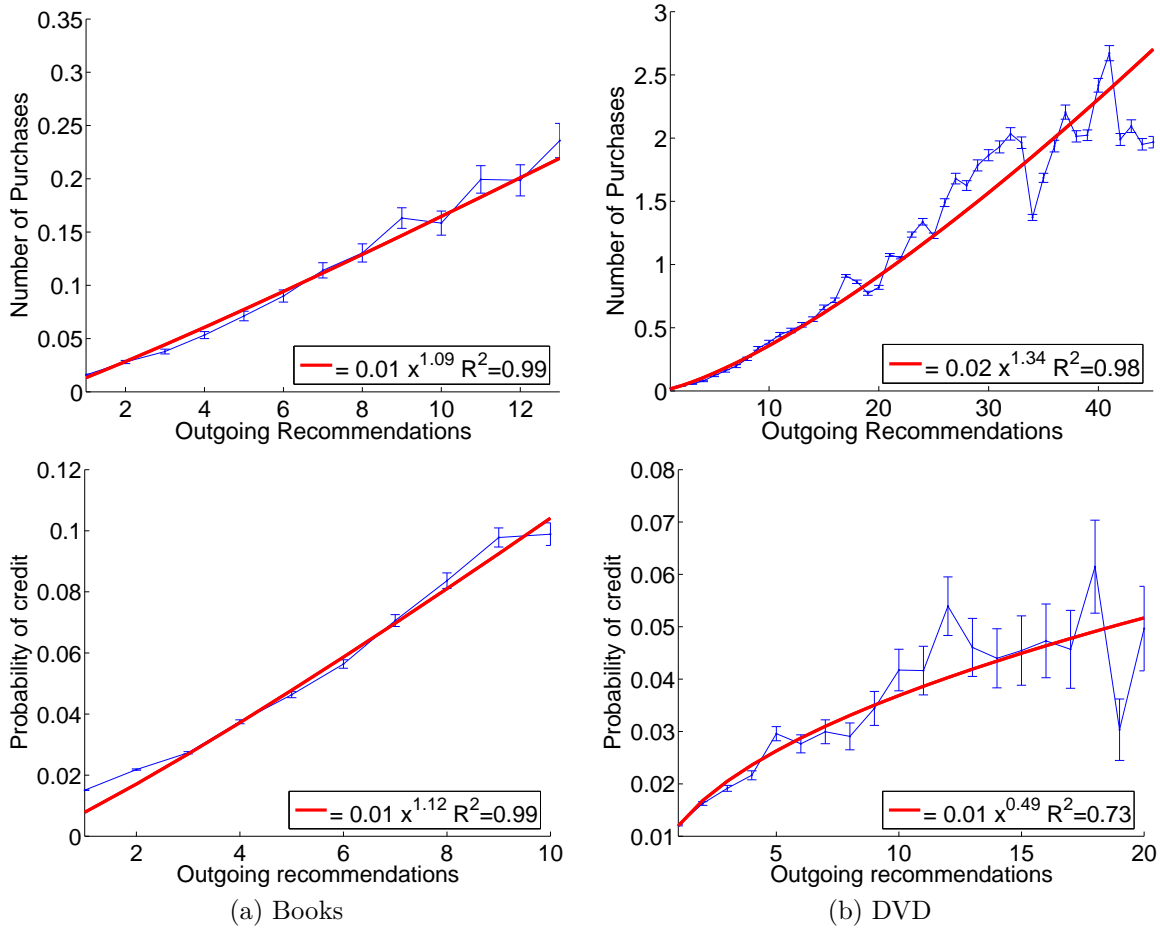

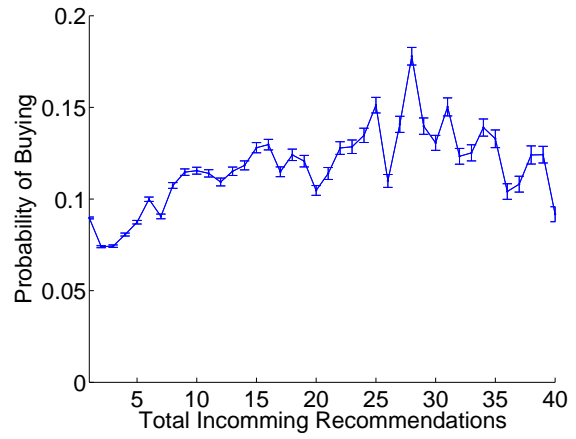



Figure A-1: Top row: Power fit to the non-linear part of the number of resulting purchases given a number of outgoing recommendations. Bottom row: Power fit to the probability of getting a credit given a number of outgoing recommendations.




(a) Books



(b) DVD



(c) Music



(d) Video

Figure A-2: Probability of buying a product given a total number of incoming recommendations on all products.




Figure A-3: The time between the first recommendation and the purchase. The histograms show how long does it take to accumulate sufficient number of recommendations to trigger a purchase. The bin size is 1 day. We use all purchases through recommendations.




Figure A-4: The time between the recommendation and the purchase taking only the recommendations that resulted in a 10% discount. The bin size is 3 hours. The dashed line presents a logarithmic fit.

category	n_p	n	cc	r_{p1}	r_{p1}/r_{p2}	v_{av}	c_{av}/r_{p1}	p_m	b
Anime and Manga	1301	46941	18.92	14.40	17.17	4.19	2.96	26.96	28.44**
Classics	266	24922	25.59	9.68	6.66	4.18	4.16	22.49	11.22**
Animation	237	80092	11.99	41.90	19.17	4.03	3.88	22.49	10.43
Science Fiction & Fantasy	1410	317420	6.61	59.18	16.66	3.85	2.51	17.99	9.62
Art House & International	3185	276142	7.37	24.35	12.97	3.95	2.22	22.46	9.43*
Television	1133	195948	8.17	18.95	11.68	4.22	5.32	17.99	8.90
Horror	1125	79744	13.15	30.00	9.10	3.59	1.37	17.98	8.72
Action and Adventure	2058	248674	7.00	39.80	15.11	3.80	1.96	17.96	8.42**
Mystery and Suspense	1683	151101	9.28	26.73	10.45	3.82	2.20	17.98	7.57
Military and War	379	69180	12.53	39.31	11.14	4.12	2.26	17.96	7.41
Cult Movies	324	94049	11.24	37.93	8.45	3.89	3.34	17.98	7.28
Kids and Family	1357	230300	6.70	30.96	12.81	4.12	3.35	17.98	6.75
Drama	3376	255544	7.12	25.14	11.02	3.98	2.10	17.98	6.72*
Comedy	2455	312033	6.08	26.25	11.14	4.02	3.30	17.98	6.01**
Musicals & Performing Arts	1091	88665	10.24	17.07	11.11	4.09	2.34	22.48	4.93
Westerns	234	17612	24.40	11.76	7.30	3.94	2.72	13.48	4.71*
Sports	484	23191	16.92	8.64	7.89	3.97	2.49	17.98	4.55*
Documentary	1058	53538	15.24	6.12	9.08	3.95	3.70	17.99	4.24
Educational	89	5532	19.60	3.39	2.63	3.97	5.48	19.95	3.99
Music Video and Concerts	2222	91657	8.44	8.06	11.16	4.09	2.88	17.99	3.85
Special Interests	963	43225	10.42	5.83	7.45	3.99	3.43	18.74	2.62
Fitness and Yoga	223	17160	2.23	14.65	6.66	3.88	2.93	17.96	1.98
African American Cinema	81	10609	17.92	16.00	9.06	4.15	3.41	17.98	1.56

Table A-1: Statistics by DVD genre. * denotes significance at the 0.05 level, ** at the 0.01 level

category	n_p	n	cc	r_{p1}	r_{p1}/r_{p2}	v_{av}	c_{av}/r_{p1}	p_m	$b * 100$
Anime and Manga	962	5081	9.64	13.98	18.76	4.39	0.26	17.99	1.99*
Educational	607	6569	1.64	1.97	10.75	4.17	3.01	19.95	1.59
Fitness	920	24627	0.43	8.41	12.09	4.09	1.92	14.95	1.48
Animation	171	9500	4.04	61.83	19.58	4.29	0.36	17.99	1.36
Kids and Family	4736	84608	1.13	14.26	12.11	4.29	0.85	12.98	1.16
Special Interests	3769	36862	1.45	3.19	12.73	4.14	1.65	19.95	1.09
Mystery and Suspense	1514	13459	9.90	30.09	9.83	4.01	0.14	14.95	1.01
Art House & International	2459	24713	3.52	17.54	10.09	4.18	0.28	17.99	0.84
Science Fiction and Fantasy	1583	29565	2.54	51.92	13.76	4.01	0.18	13.99	0.83
Documentary	2936	18884	1.15	3.33	9.83	4.21	0.95	19.95	0.82
Television	3632	31475	0.95	5.13	12.11	4.33	1.01	14.95	0.71
Music Video & Concerts	1595	14360	4.46	8.75	11.26	4.40	0.49	16.99	0.70
Musicals & Performing Arts	1621	22539	3.13	13.22	9.39	4.20	0.51	19.95	0.69
Sports	1251	7987	0.49	4.07	9.83	4.15	0.91	16.99	0.69
Comedy	3645	55868	2.13	22.26	10.60	4.13	0.36	13.99	0.59
Drama	4837	52691	1.87	21.72	9.25	4.15	0.26	14.95	0.56
Military and War	829	10859	1.13	28.54	9.39	4.22	0.21	14.95	0.56
Westerns	487	3743	1.58	9.42	6.01	4.12	0.43	9.99	0.56
Classics	326	3029	0.56	8.73	8.15	4.12	0.51	14.94	0.49
African American Cinema	87	1861	0.64	15.53	7.59	4.10	0.61	9.99	0.49
Horror	935	6728	1.07	36.38	9.02	3.81	0.10	12.99	0.40
Action and Adventure	2390	25921	1.84	33.13	11.90	3.96	0.17	13.99	0.31
Cult Movies	401	5260	0.65	32.06	7.63	3.90	0.18	9.99	0.30

Table A-2: Statistics for videos in VHS format by genre

category	n_p	n	cc	r_{p1}	r_{p1}/r_{p2}	v_{av}	c_{av}/r_{p1}	p_m	b
Broadway and Vocalists	5423	104396	4.25	6.03	13.86	4.49	1.68	14.49	2.01
Country	5876	98069	4.67	5.50	18.45	4.56	1.76	13.99	1.87
Rock	10717	196852	4.10	11.00	10.18	4.40	0.99	14.99	1.87
Alternative Rock	13405	216324	5.12	13.20	11.24	4.41	0.81	13.99	1.87
Soundtracks	4491	133507	4.81	7.92	13.82	4.38	1.77	14.99	1.87
Classical	14223	116937	5.34	2.65	11.60	4.52	1.82	15.49	1.83
Folk	5244	87580	5.33	4.40	13.54	4.60	2.05	14.99	1.81
Pop	16764	322431	3.30	9.55	13.19	4.43	1.22	13.99	1.78
Opera and Vocal	5402	61643	6.08	3.32	12.90	4.48	1.69	15.99	1.73
Miscellaneous	5823	80243	5.71	3.54	12.31	4.35	1.90	13.98	1.62
Blues	2987	31199	6.62	2.76	11.53	4.59	1.89	14.99	1.54
Hard Rock and Metal	4787	63893	4.96	18.23	7.92	4.33	0.42	14.99	1.52
Christian and Gospel	2977	37554	2.02	5.41	16.75	4.67	1.20	14.99	1.51
Jazz	11868	113078	4.49	2.91	11.40	4.59	1.99	14.99	1.50
Classic Rock	5711	117255	4.74	13.62	6.78	4.29	0.87	13.99	1.50
Children s Music	1755	37015	4.89	3.96	12.52	4.53	2.94	12.32	1.47
Dance and DJ	11332	139787	5.16	7.14	14.64	4.38	1.05	14.99	1.42
New Age	4219	60951	5.90	3.92	13.79	4.54	1.95	14.99	1.42
International	13139	130499	5.02	3.54	9.52	4.57	1.51	14.99	1.32
Latin Music	4634	38725	5.06	2.57	16.76	4.60	1.75	13.99	1.30
Rap and Hip Hop	3996	60135	3.67	12.23	9.64	4.38	0.67	14.99	1.14
R&B	5965	85380	2.78	8.49	12.90	4.48	0.89	13.98	1.13

Table A-3: Statistics by Music Style

	$\ln(s)$	$\ln(r)$	$\ln(n_s)$	$\ln(n_r)$	$\ln(p)$	$\ln(v)$	$\ln(t)$
$\ln(s)$	1.0000						
$\ln(r)$	-0.5816	1.0000					
$\ln(n_s)$	-0.7943	0.8812	1.0000				
$\ln(n_r)$	-0.6383	0.9923	0.9066	1.0000			
$\ln(p)$	0.1001	0.0127	-0.0017	0.0027	1.0000		
$\ln(v)$	-0.5377	0.5804	0.6456	0.6024	-0.1146	1.0000	
$\ln(t)$	0.0199	-0.0294	-0.0319	-0.0311	-0.0468	-0.0852	1.0000

Table A-4: Pairwise correlation matrix of the product attributes. $\ln(s)$: log recommendation success rate, $\ln(r)$: log number of recommendations, $\ln(n_s)$: log number of renders of recommendations, $\ln(n_r)$: log number of receivers, $\ln(p)$: log price, $\ln(v)$: log number of reviews, $\ln(t)$: log average rating.