The Dynamics of Viral Marketing *

Jure Leskovec†
Carnegie Mellon University
Pittsburgh, PA 15213
jure@cs.cmu.edu

Lada A. Adamic‡
University of Michigan
Ann Arbor, MI 48109
ladamic@umich.edu

Bernardo A. Huberman
HP Labs
Palo Alto, CA 94304
bernardo.huberman@hp.com

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 observe the propagation of recommendations and the cascade sizes, which we explain by a simple stochastic model. We then establish 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, we present a model that successfully identifies product and pricing categories for which viral marketing seems to be very effective.

Categories and Subject Descriptors

J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Economics

Keywords

E-commerce, Recommender systems, Viral marketing

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. Previously, a few in depth studies have shown that social networks affect the adoption of individual innovations and products (for a review see [15] or [16]). But 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 and model 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.

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. Email services such as Hotmail and Yahoo had very fast adoption curves because every email sent through them contained an advertisement for the service and because they were free. Hotmail spent a mere $50,000 on traditional marketing and still grew from zero to 12 million users in 18 months [7]. 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 was through a referral.

Most products cannot be advertised in such a direct way. At the same time the choice of products available to consumers has increased manyfold thanks to online retailers who can supply a much wider variety of products 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 [2]. Rhapsody, a streaming-music service, streams more tracks outside than inside its top 10,000 tunes [1]. Effectively advertising these niche products using traditional advertising approaches is impractical. Therefore using more targeted marketing approaches is advantageous 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. Quantitative marketing techniques have been proposed [12], and the rating of products and merchants has been shown to effect the likelihood of an item being bought [13, 4]. Of further help to the consumer are collaborative filtering recommendations of the form "people who bought x also bought y" feature [11]. These refinements help consumers discover new products

*A longer version of this paper can be found at http://arxiv.org/abs/physics/0509039

†This research was done while at HP Labs.

‡Research likewise done while at HP Labs.

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 [3].

Several studies have attempted to model just this kind of network influence. Richardson and Domingos [14] used Epinions' trusted reviewer network to construct 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 [8] evaluate the efficiency of several algorithms for maximizing the size of influence set given various models of adoption. While these models address the question of maximizing the spread of influence in a network, they are based on *assumed* rather than *measured* influence effects.

In contrast, in our study we are able to directly observe the effectiveness of person to person word of mouth advertising for hundreds of thousands of products for the first time. We find that most recommendation chains do not grow very large, often terminating with the initial purchase of a product. However, occasionally a product will propagate through a very active recommendation network. We propose a simple stochastic model that seems to explain the propagation of recommendations. Moreover, the characteristics of recommendation networks influence the purchase patterns of their members. For example, individuals' likelihood of purchasing a product initially increases as they receive additional recommendations for it, but a saturation point is quickly reached. Interestingly, as more recommendations are sent between the same two individuals, the likelihood that they will be heeded decreases. We also propose models to identify products for which viral marketing is effective: We find that the category and price of product plays a role, with recommendations of expensive products of interest to small, well connected communities resulting in a purchase more often. We also observe patterns in the timing of recommendations and purchases corresponding to times of day when people are likely to be shopping online or reading email. We report on these and other findings in the following sections.

2. THE RECOMMENDATION NETWORK

2.1 Dataset description

Our analysis focuses on the recommendation referral program run by a large retailer. The program rules were as follows. Each time a person 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.

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, 99% of them belonging to 4 main product groups: Books, DVDs, Music and Videos. In addition to recommendation data, we also crawled the retailer's website to obtain product categories, reviews and ratings for all products. Of the products in our data set, 5813 (1%) were discontinued (the retailer no longer provided any information about them).

Although the data gives us a detailed and accurate view of recommendation dynamics, 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.

2.2 Recommendation network statistics

For each recommendation, the dataset included the product and product price, sender ID, receiver ID, the sent date, and a *buy-bit*, indicating whether the recommendation resulted in a purchase and discount. The sender and receiver ID's were shadowed. We represent this data set as a directed multi graph. The nodes represent customers, and a directed edge contains all the information about the recommendation. The edge (i, j, p, t) indicates that 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 $j_1, \ldots, j_n$. 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 *buy-bit*. Because the 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*. Note that buy-edges provide only a lower bound on the total number of purchases without discounts. It is possible for a customer to not be the first to act on a recommendation and also to not recommend the product to others. Unfortunately, this was not recorded in the data set. We consider, however, the buy-bits and buy-edges as proxies for the total number of purchases through recommendations.

For each product group we took recommendations on all products from the group and created a network. Table 1

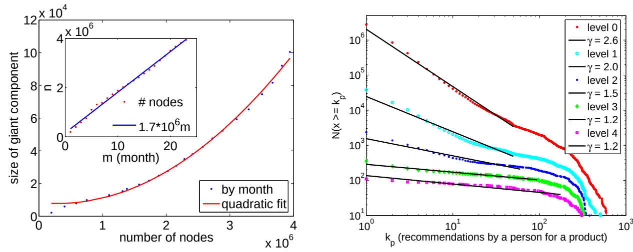

(a) Network growth (b) Recommending by level

Figure 1: (a) The size of the largest connected component of customers over time. The inset shows the linear growth in the number of customers n over time. (b) The number of recommendations sent by a user 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.

(first 7 columns) shows the sizes of various product group recommendation networks with p being the total number of products in the product group, n the total number of nodes spanned by the group recommendation network and e the number of edges (recommendations). The column e_u shows the number of unique edges – disregarding multiple recommendations between the same source and recipient.

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 more half of all recommendations in the dataset. The DVD network 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.

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 [5].

Figure 1(a) shows the fraction of nodes in largest weakly connected component over time. Notice the component is very small. Even if we compose a network using all the recommendations in the dataset, the largest connected component contains less than 2.5% (100,420) of the nodes, and the second largest component has only 600 nodes. Still, some smaller communities, numbering in the tens of thousands of purchasers of DVDs in categories such as westerns, classics and Japanese animated films (anime), had connected components spanning about 20% of their members.

The insert in figure 1(a) shows the growth of the customer base over time. Surprisingly it 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.

Back to table 1: given the total number of recommendations e and purchases $(b_b + b_e)$ 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.

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, and 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). 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.

2.3 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.

The last 3 columns of table 1 show that only about a third of the 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.

Figure 1(b) 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 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 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.

2.4 Identifying cascades

As customers continue forwarding recommendations, they contribute to the formation of cascades. In order to identify cascades, i.e. the "causal" propagation of recommendations, we track *successful recommendations* as they influence purchases and further recommendations. We define a recommendation to be successful if it reached a node before its *first* purchase. We consider only the first purchase of an item, because there are many cases when a person made multiple

Group	p	n	e	e_u	b_b	b_e	Purchases	Forward	Percent
Book	103,161	2,863,977	5,741,611	2,097,809	65,344	17,769	65,391	15,769	24.2
DVD	19,829	805,285	8,180,393	962,341	17,232	58,189	16,459	7,336	44.6
Music	393,598	794,148	1,443,847	585,738	7,837	2,739	7,843	1,824	23.3
Video	26,131	239,583	280,270	160,683	909	467	909	250	27.6
Total	542,719	3,943,084	15,646,121	3,153,676	91,322	79,164	90,602	25,179	27.8

Table 1: Product group recommendation statistics. p: number of products, n: number of nodes, e: number of edges (recommendations), e_u: number of unique edges, b_b: number of buy bits, b_e: number of buy edges. Last 3 columns of the table: 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.

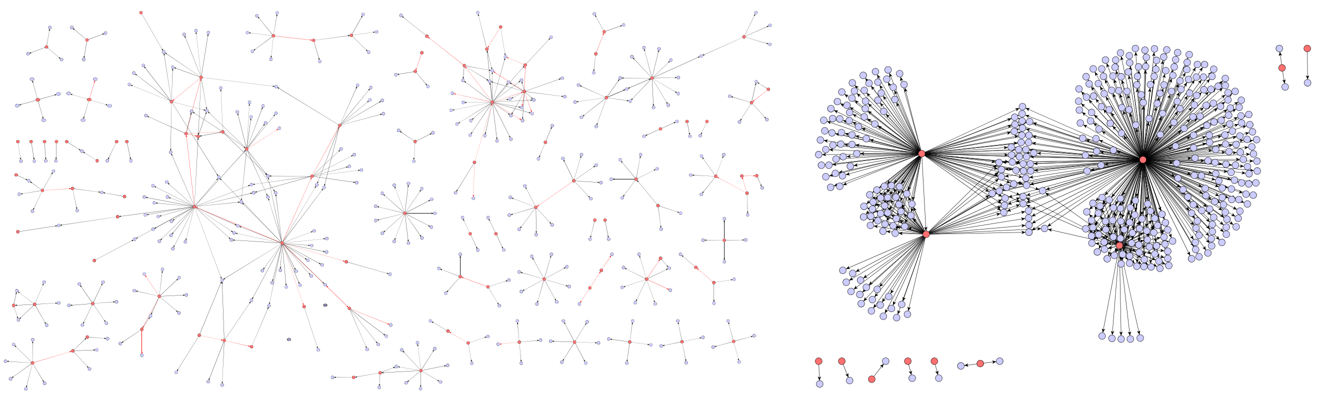

(a) Medical book

(b) Japanese graphic novel

Figure 2: 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*.

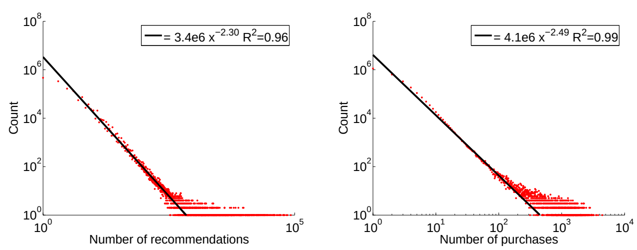

(a) Recommendations

(b) Purchases

Figure 3: Distribution of the number of recommendations and number of purchases made by a node.

purchases of the same product, and in between those purchases she may have received new recommendations. In this case one cannot conclude that recommendations following the first purchase influenced the later purchases.

Each cascade is a network consisting of customers (nodes) who purchased the same product as a result of each other's recommendations (edges). We delete *late recommendations* — all incoming recommendations that happened after the first purchase of the product. This way we make the network *time increasing* or *causal* — for each node all incoming edges (recommendations) occurred before all outgoing edges. Now each connected component represents a time obeying propagation of recommendations.

Figure 2 shows two typical product recommendation networks: (a) a medical study guide and (b) a Japanese graphic novel. Throughout the dataset we observe very similar patters. 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 with recommendations successfully propagating.

This observation is reflected in the heavy tailed distribution of cascade sizes (see figure 4), having a power-law exponent close to 1 for DVDs in particular.

We also notice bursts of recommendations (figure 2(b)). Some nodes recommend to many friends, forming a star like pattern. Figure 3 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. Finally, we also sometimes observe 'collisions', where nodes receive recommendations from two or more sources. A detailed enumeration and analysis of observed topological cascade patterns for this dataset is made in [10].

2.5 The recommendation propagation model

A simple model can help explain how the wide variance we observe in the number of recommendations made by individuals can lead to power-laws in cascade sizes (figure 4). The model 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 thresh-

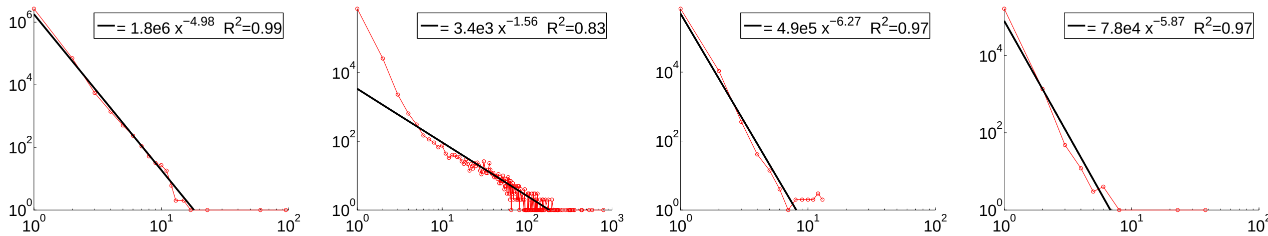

(a) Book (b) DVD (c) Music (d) Video

Figure 4: Size distribution of cascades (size of cascade vs. count). Bold line presents a power-fit.

old. In that case the number of recommendations N_{t+1} at time $(t+1)$ is given in terms of the number of recommendations at an earlier time by

$$N_{t+1} = p_t N_t \tag{1}$$

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 diving by it we obtain

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

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

$$\frac{dN}{N} = \sum p_t \tag{3}$$

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

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

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) and one obtains

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

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 [6] or critical thresholds in the probability of a recommendation being accepted [18].

3. SUCCESS 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 recommendation effectiveness as two people exchange more and more recommendations. Lastly, we observe the recommendation network from the perspective of the sender of the recommendation. Does a node that makes more recommendations also influence more purchases?

3.1 Probability of buying versus number of incoming recommendations

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

Figure 5 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 book, and only 58 cases where a person got 20 incoming recommendations on a particular book. The maximum was 30 incoming recommendations. For these reasons we cut-off the plot when the number of observations becomes too small and the error bars too large.

Figure 5(a) shows that, overall, book recommendations are rarely followed. 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 5(b)) we observe a saturation around 10 incoming recommendations. This means that after a person gets 10 recommendations on a particular DVD, 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 received recommendations is 172 (and that person did not buy)

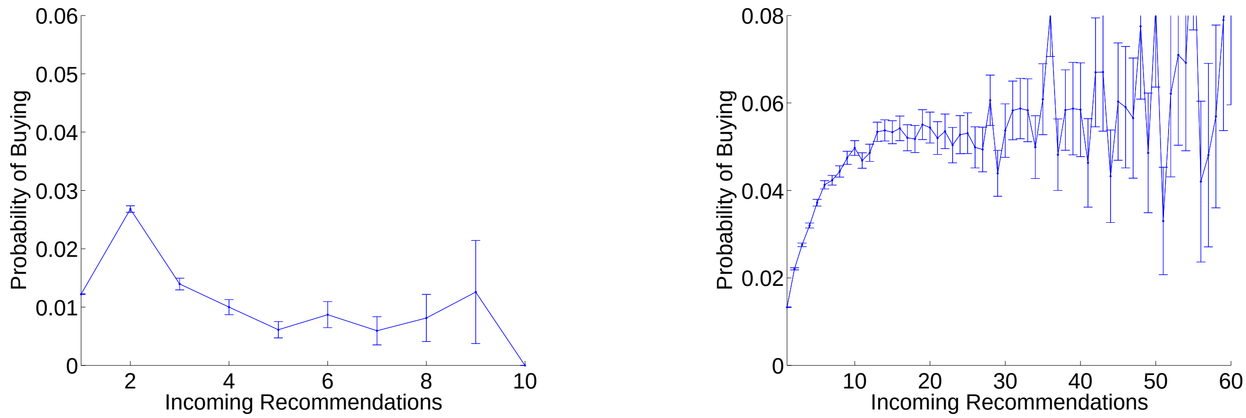

(a) Books (b) DVD

Figure 5: Probability of buying a book (DVD) given a number of incoming recommendations.

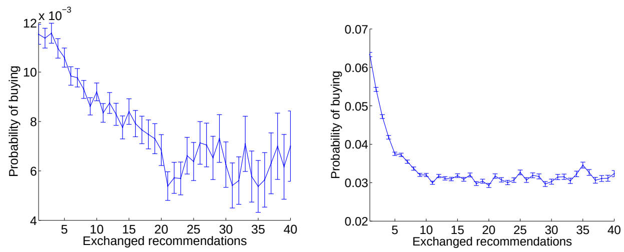

(a) Books (b) DVD

Figure 6: The effectiveness of recommendations with the total number of exchanged recommendations.

3.2 Success of subsequent recommendations

Next, we analyze how the effectiveness of recommendations changes as two persons exchange more and more recommendations. A large number of exchanged recommendations can be a sign of trust and influence, but a sender of too many recommendations can be perceived as a spammer. A person who recommends only a few products will have her friends' attention, but one who floods her friends with all sorts of recommendations will start to loose her influence.

We measure the effectiveness of recommendations as a function of the total number of previously exchanged recommendations between the two nodes. We construct 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 recommendation r. 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 a total of 10 recommendations sent from u to v. We perform the experiment using only recommendations from the same product group.

Figure 6 shows the probability of buying as a function of the total number of exchanged recommendations between two persons up to that point. 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. This experiment shows that recommendations start to lose effect after more than two or three are passed between two people. We performed the experiment also for video and music, but the number of observations was too low and the measurements were noisy.

3.3 Success of outgoing recommendations

In previous sections 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. One the other hand, recommending to everyone and spamming them with recommendations may have limited returns as well.

The top row of figure 7 shows how the average number of purchases changes with the number of outgoing recommendations. For books, music, and videos the number of purchases soon saturates: it grows fast up to around 10 outgoing recommendations and then the trend either slows or starts to drop. DVDs exhibit different behavior, with the expected number of purchases increasing throughout. But if we plot the probability of getting a 10% credit as a function of the number of outgoing recommendations, as in the bottom row of figure 7, we see that the success of DVD recommendations saturates as well, while books, videos and music have qualitatively similar trends. The difference in the curves for DVD recommendations points to the presence of collisions in the dense DVD network, which has 10 recommendations per node and around 400 per product — an order of magnitude more than other product groups. 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'

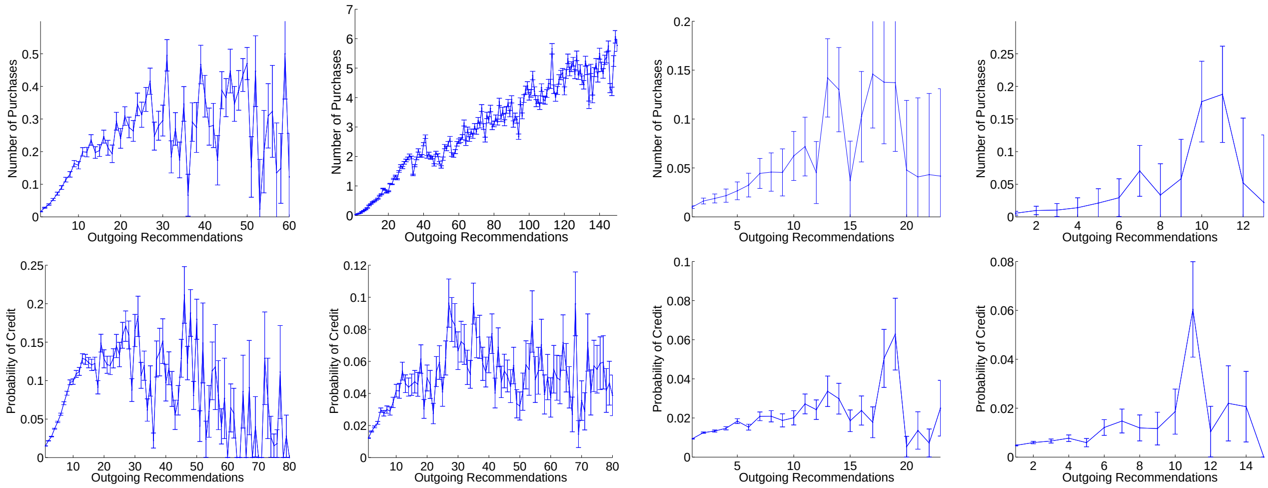

(a) Books (b) DVD (c) Music (d) Video

Figure 7: 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.

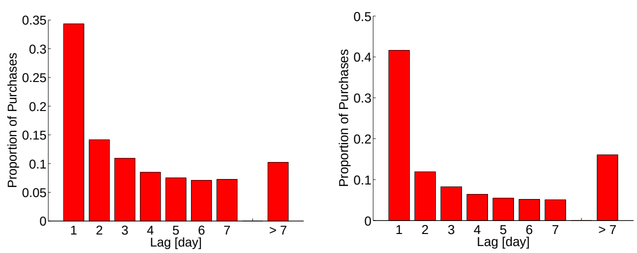

(a) Books (b) DVD

Figure 8: The time between the recommendation and the actual purchase. We use all purchases.

by our definition, only one of them receives a credit.

4. TIMING OF RECOMMENDATIONS AND PURCHASES

The recommendation referral program encourages people to purchase as soon as possible after they get a recommendation, since this maximizes the probability of getting a discount. We study the time lag between the recommendation and the purchase of different product groups, effectively how long it takes a person to both receive a recommendation, consider it, and act on it.

We present the histograms of the "thinking time", i.e. the difference between the time of purchase and the time the last recommendation was received for the product prior to the purchase (figure 8). We use a bin size of 1 day. Around 35%-40% of book and DVD purchases occurred within a day after the last recommendation was received. For DVDs 16% purchases occur more than a week after last recommendation, while this drops to 10% for books. In contrast, if we consider the lag between the purchase and the *first* recommendation, only 23% of DVD purchases are made within a day, while the proportion stays the same for books. This reflects a greater likelihood for a person to receive multiple recommendations for a DVD than for a book. At the same time, DVD recommenders tend to send out many more recommendations, only one of which can result in a discount. Individuals then often miss their chance of a discount, which is reflected in the high ratio (78%) of recommended DVD purchases that did not a get discount (see table 1, columns b_b and b_e). In contrast, for books, only 21% of purchases through recommendations did not receive a discount.

We also measure the variation in intensity by time of day for three different activities in the recommendation system: recommendations (figure 9(a)), all purchases (figure 9(b)), and finally just the purchases which resulted in a discount (figure 9(c)). Each is given as a total count by hour of day.

The recommendations and purchases follow the same pattern. The only small 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 by being the first one to purchase, she had the highest chances when the traffic on the website was the lowest.

5. RECOMMENDATION EFFECTIVENESS BY BOOK CATEGORY

Social networks are a product of the contexts that bring people together. Some contexts result in social ties that are 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 [9, 17]. 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

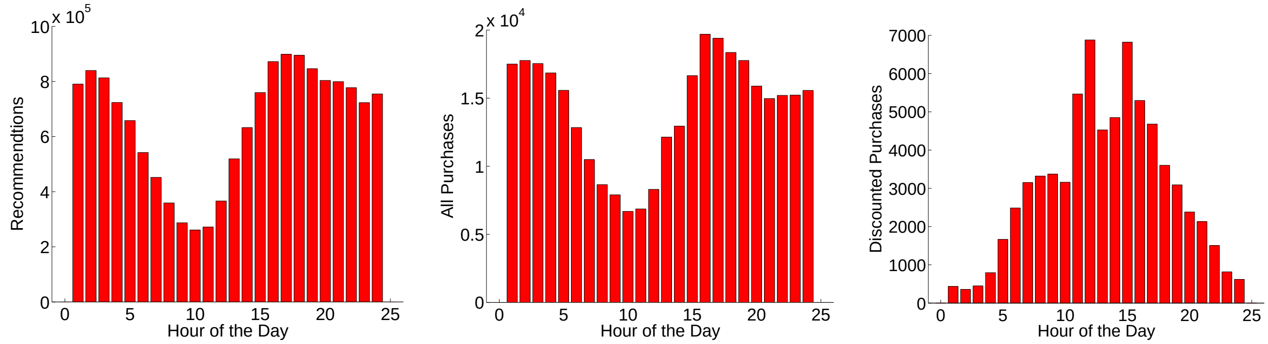

(a) Recommendations (b) Purchases (c) Purchases with Discount

Figure 9: Time of day for purchases and recommendations. (a) shows the distribution of recommendations over the day, (b) shows all purchases and (c) shows only purchases that resulted in getting discount.

more or less influential than a recommendation for a technical book, which may be made by a colleague at work or school.

Table 2 shows recommendation trends for all top level book categories by subject. An analysis of other product types can be found in the extended version of the paper. For clarity, we group the results by 4 different category types: fiction, personal/leisure, professional/technical, and nonfiction/other. 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.

First, we compare the relative number of recommendations to reviews posted on the site (column c_{av}/r_{p1} of table 2). Surprisingly, we find 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. 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.

Next, we measure the per category efficacy of recommendations by observing the ratio of the number of purchases occurring within a week following a recommendation to the number of recommenders for each book subject category (column b of table 2). 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 observe marked differences in the response to recommendation for different categories of books. Fiction in general is 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 is significantly higher in the professional and technical category. Medical books have 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 6, a higher product price increases the chance that a recommendation will be accepted.

Recommendations are 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, have 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. There are exceptions of course. 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.

6. MODELING THE RECOMMENDATION SUCCESS

We have examined the properties of recommendation network in relation to viral marketing, but one question still remains: what determines the product's viral marketing success? We present a model which characterizes product categories for which recommendations are more likely to be accepted. We use a regression of the following product attributes to correlate them with recommendation success:

- 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

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

Table 2: 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, 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.

From the original set of half a million products, we compute a success rate s for the 48,218 products that had at least one purchase made through a recommendation and for which a price was given. In section 5 we defined recommendation success rate s as the ratio of the total number purchases made through recommendations and 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 where a few individuals send out many recommendations (figure 2(b)). Since the variables follow a heavy tailed distribution, we use the following model:

$$s = \exp(\sum_i \beta_i \log(x_i) + \epsilon_i)$$

where x_i are the product attributes (as described on previous page), and ϵ_i is random error.

We fit the model using least squares and obtain the coefficients β_i shown on table 3. With the 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 a higher success rate. The number of senders and receivers have large negative coefficients, showing that successfully recommended products are more likely to be not so widely popular. They have relatively many 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.

These insights could be to marketers — personal recommendations are most effective in small, densely connected communities enjoying expensive products.

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

Table 3: Regression using the log of the recommendation success rate, $\ln(s)$, as the dependent variable. For each coefficient we provide the standard error and the statistical significance level (:0.01, *:0.1).**

7. DISCUSSION AND CONCLUSION

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 individuals have equal probability of being infected every time they interact. Contrary to this we observe that the probability of infection decreases with repeated interaction. Marketers should take heed that 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 individuals are often impervious to the recommendations of their friends, and 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. But these models assume that a high degree node has as much of a 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 for a 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 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.

We saw that the characteristics of product reviews and effectiveness of recommendations vary by category and price, 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.

Finally, we presented a model which shows that smaller and more tightly knit groups tend to be more conducive to viral marketing. So despite the relative ineffectiveness of the viral marketing program in general, we found a number of new insights which we hope will have general applicability to marketing strategies and to future models of viral information spread.

8. REFERENCES

[1] Anonymous. Profiting from obscurity: What the long tail means for the economics of e-commerce. *Economist*, 2005.

[2] E. Brynjolfsson, Y. Hu, and M. D. Smith. Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science*, 49(11), 2003.

[3] K. Burke. As consumer attitudes shift, so must marketing strategies. 2003.

[4] J. Chevalier and D. Mayzlin. The effect of word of mouth on sales: Online book reviews. 2004.

[5] P. Erdös and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci.*, 1960.

[6] D. Gruhl, R. Guha, D. Liben-Nowell, and A. Tomkins. Information diffusion through blogspace. In *WWW '04*, 2004.

[7] S. Jurvetson. What exactly is viral marketing? *Red Herring*, 78:110–112, 2000.

[8] D. Kempe, J. Kleinberg, and E. Tardos. Maximizing the spread of infuence in a social network. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2003.

[9] P. Killworth and H. Bernard. Reverse small world experiment. *Social Networks*, 1:159–192, 1978.

[10] J. Leskovec, A. Singh, and J. Kleinberg. Patterns of influence in a recommendation network. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD)*, 2006.

[11] G. Linden, B. Smith, and J. York. Amazon.com recommendations: item-to-item collaborative filtering. *IEEE Internet Computing*, 7(1):76–80, 2003.

[12] A. L. Montgomery. Applying quantitative marketing techniques to the internet. *Interfaces*, 30:90–108, 2001.

[13] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebays reputation system. In *The Economics of the Internet and E-Commerce*. Elsevier Science, 2002.

[14] M. Richardson and P. Domingos. Mining knowledge-sharing sites for viral marketing. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, 2002.

[15] E. M. Rogers. *Diffusion of Innovations*. Free Press, New York, fourth edition, 1995.

[16] D. Strang and S. A. Soule. Diffusion in organizations and social movements: From hybrid corn to poison pills. *Annual Review of Sociology*, 24:265–290, 1998.

[17] J. Travers and S. Milgram. An experimental study of the small world problem. *Sociometry*, 1969.

[18] D. Watts. A simple model of global cascades on random networks. *PNAS*, 99(9):4766–5771, Apr 2002.