

December 2006

Network Structure and the Long Tail of ECommerce Demand

Gal Oestreicher-Singer
New York University

Arun Sundararajan
New York University

Follow this and additional works at: <http://aisel.aisnet.org/icis2006>

Recommended Citation

Oestreicher-Singer, Gal and Sundararajan, Arun, "Network Structure and the Long Tail of ECommerce Demand" (2006). *ICIS 2006 Proceedings*. 26.
<http://aisel.aisnet.org/icis2006/26>

This material is brought to you by the International Conference on Information Systems (ICIS) at AIS Electronic Library (AISEL). It has been accepted for inclusion in ICIS 2006 Proceedings by an authorized administrator of AIS Electronic Library (AISEL). For more information, please contact elibrary@aisnet.org.

NETWORK STRUCTURE AND THE LONG TAIL OF ECOMMERCE DEMAND

Economics and Information Systems

Gal Oestreicher-Singer
New York University
goestrei@stern.nyu.edu

Arun Sundararajan
New York University
arun@stern.nyu.edu

Abstract

We report on a research project that studies how network structures affect demand in electronic commerce, using daily data about the graph structure of Amazon.com's co-purchase network for over 250,000 products, gathered over this year. We describe how the presence of such network structures alters demand patterns by changing the distribution of traffic between ecommerce Web pages. When this traffic distribution generated by the presence of the network is less skewed than the intrinsic or "real world" traffic distribution, such network structures will even out demand across products, leading to a demand distribution with a longer tail. We estimate an econometric model to validate this theory, and report on preliminary confirmation by contrasting the demand distributions of products in over 200 distinct categories on Amazon.com. We measure the overall extent to which a product influences the network by adapting Google's PageRank algorithm, applying it to a weighted composite of graphs over four distinct seven-day periods, and we characterize the demand distribution of each category using its Gini coefficient. Our results establish that categories whose products are influenced more by the network structure have significantly flatter demand distributions, which provides an additional explanation for the widely documented phenomenon of the long tail of ecommerce demand. Our research in progress aims to extend these results by using a model of consumer behavior more consistent with the economics of electronic commerce, one we term the "strategic surfer" model.

Keywords: electronic commerce, Internet, networks, network structure, econometrics, demand estimation, long tail, PageRank, Gini coefficient, SalesRank.

Introduction

There are numerous networks associated with electronic business. Some of these can describe the relationships between consumers who communicate product information and influence each others' purchasing, others can describe how the demand for different products are related based on shared purchasing patterns, and yet others may describe the patterns of trade between firms. Many recent studies in information systems have recognized the role that such network structure play in IS research as diverse as measuring the productivity of information workers (Aral et al., 2006), the effectiveness of information systems for knowledge (Kane and Alavi, 2006), and modeling the adoption of products that display local network effects (Sundararajan, 2006).

In this research summary, we report on a project that aims to model and measure the effect that such network structures have on outcomes in electronic commerce. A good example of such a structure is the network of product pages on an ecommerce site. Each product on an ecommerce site has a *network position*, which is determined by the products it links to, and those that link to it. If one imagines the process of browsing an ecommerce site as being analogous to walking the aisles of a physical store, then the *ecommerce aisle structure* is the this graph of interconnected products, and the network position of a product in this graph is analogous to its *virtual shelf placement*. A product that is linked to by an intrinsically popular one is likely to enjoy an increase in sales on account of this aspect of its network position. A product linked to by hundreds of others is likely to get more

“network traffic” more than one linked to by just a few. Thus, both the structure of the networks and the nodes that comprise them seem to matter.

We measure the extent to which the position of a product in such a network structure will affect its demand, based on the idea that the network structure redirects the flow of consumer attention, which results in a redistribution of traffic and demand. The manner in which attention is redistributed can be measured using certain properties of the network structure, and these properties can be associated with observed variations in both individual and aggregate product demand. One specific prediction of our theory is that network structures with common degree distributions will even out traffic between products (nodes), thereby reducing demand inequity between products.

Our preliminary evidence is based on econometric estimates of how the intensity of the network structure affects the demand distributions of over 250,000 products within over 200 distinct categories on Amazon.com. Briefly, we compute an adapted version of the PageRank coefficient for each node of a composite of seven consecutive daily instances of Amazon.com’s co-purchase network. We contrast four such composites, one corresponding to each week of February 2006. We characterize the demand distribution of each category by constructing its Lorenz curve and measuring its Gini coefficient. We show that when network structure has a greater influence (when the average Weighted PageRank is higher) on a category, its demand distribution displays significantly lower inequity (its Gini coefficient is significantly lower). In other words, the presence of the network structure flattens ecommerce demand distribution. This provides a new explanation for the widely documented long tail of ecommerce demand.

Many recent studies in the information systems field have recognized the role that such network structures play in IS. Those studies have begun to introduce concepts from social network and complex network research to different streams of IS research. For example, researchers have studied the evolution and development of network structures in online communities over time (Schoberth et al., 2003) and the internal dynamics of communities (Butler, 2001). Models of network structure have been recently applied to studying knowledge management by Alavi and Kane (2006), whose multimodal knowledge networks include both people and information systems as nodes within a single network, and facilitate exploring more complex interactions (like IS centrality) and their effects on knowledge sharing. Aral et al. (2006) study the influence of network structures on IT worker productivity, relating technology use and social network characteristics to economic measures, and providing evidence that the structure and size of workers' communication networks, including such social network metrics as betweenness and structural holes, are highly correlated with performance. Features based on network structure have been shown to improve the predictions of data mining models used for targeted marketing (Hill et al., 2006).

There has been some prior research in marketing that aimed to assess the structure of preferences for products based on purchase data. These studies have used scanner panel data, and are based on a similar notion: that such data contain important information based on revealed preference about the structure of brand preferences both within and across product categories. For example, the time series of purchasing data has been used to compute segment-product distances (Ramaswamy and DeSarbo, 1991), segment consumers with respect to brand preferences (Russell and Kamakura, 1997, Matthias et al., 2002) and build probabilistic models that provide spatial representations of product structure (Erdem et al., 1999). These are based on co-purchase bundles, and none of these papers has used a co-purchase network, or exploited any structural properties of these networks in making inferences about the nature of demand.

Beyond the scope of traditional IS research, network structures have received a significant amount of attention from sociologists studying relationships between people, from physicists and computer scientists studying the Web. Our work is also related to a growing literature on using network structures to create sophisticated ranking algorithms, one well-known contribution being the PageRank algorithm of Brin and Page, 1998 (for further information, see Langville and Meyer, 2005).

Data

We collect daily product, pricing, demand, and “network” information for over 250,000 books sold on Amazon.com. Each product on Amazon.com has an associated Web page. Those pages each has a set of “co-purchase links”, which are hyperlinks to the set of products that were co-purchased most frequently with this product on Amazon.com. This set is listed under the title "Customers who bought this also bought" and is limited to five items (see Figure 1).

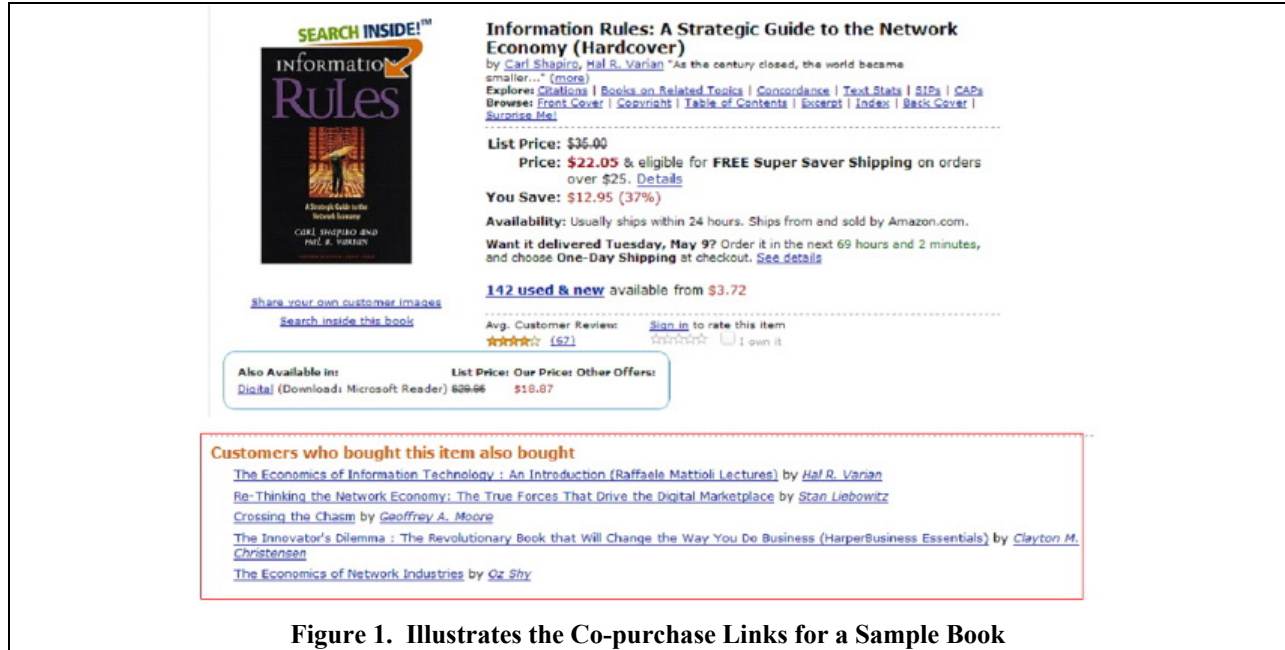


Figure 1. Illustrates the Co-purchase Links for a Sample Book

Conceptually the co-purchase network is a directed graph in which nodes correspond to products, and edges to directed co-purchase links. We collect data about this graph using a Java-based crawler, which starts from a popular book and follows the co-purchase links using a depth-first algorithm. At each page, the crawler gathers and records information for the book whose Web page it is on, as well as the co-purchase links on that page, and terminates when the entire connected component of the graph is collected. This process is repeated daily.

We have chosen to focus on books because they are in the product category with by far the largest number of individual titles, whose product set is relatively stable (compared to electronics, for instance), and it seems to be a class of products for which the network we study would actually matter.

Apart from the co-purchases, each book’s ISBN, list price, sale price, category affiliation, secondary market activity, author, publisher, publication date, and consumer ratings are gathered. A sample part of the graph is illustrated in Figure 2. We use data from February 1st through 28th, divided into four week-long samples: 2/1-2/7, 2/8-2/14, 2/15-2/21 and 2/22-2/28.

An additional script collects the demand information for all books on the graph every three hours for the 24-hour period *following* the collection of the graph.

The following data are gathered:

ASIN: a unique serial number given to each book by Amazon.com. Different editions and different versions have different ASIN numbers.

List Price: The publisher’s suggested price.

Sale Price: The price on the Amazon.com Web site that day.

Co-purchases: ASINs of the books that appear as its co-purchases.

SalesRank: The sales rank is a number associated with each product on Amazon.com that measures its demand of relative to other products. The lower the number is, the higher the sales of that particular product. The sales rank of a book is updated each hour to reflect recent and historical sales of every item sold on Amazon.com.

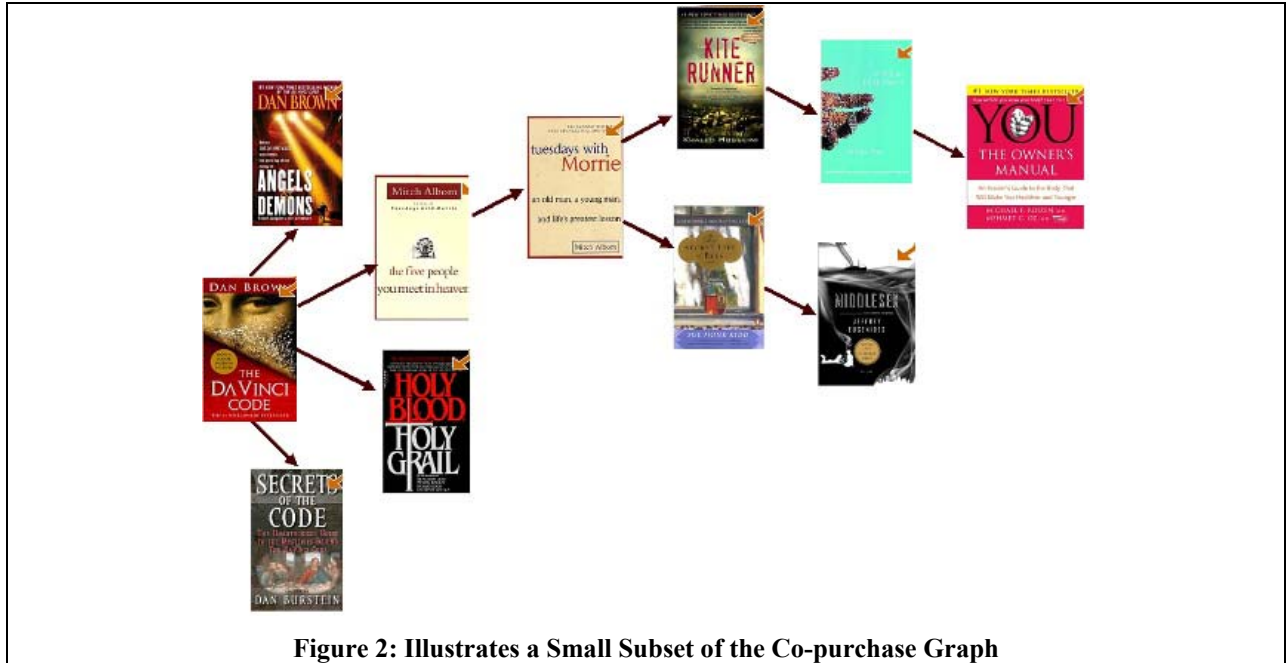


Figure 2: Illustrates a Small Subset of the Co-purchase Graph

Category Affiliation: Amazon.com uses a hierarchy of categories to classify its books. Thus, each book is associated with one or more hierarchical lists of categories, starting with the most general category affiliation and ending with the most specific one. For example:

Subjects > Business & Investing > Biographies & Primers > Company Profiles

(for “The Search” by John Batelle).

We use the second-most general partition defined by this hierarchy to define category affiliation. There are 1472 such categories in all, spanning the millions of books sold on Amazon.com. Of these, between 212 and 233 categories have over 100 products on our co-purchase graph. If a book has multiple category lists associated with it, we use the first one listed on the book’s Web page.

Author: The name of the author or authors of the book.

Publisher: The name of the publisher of the book.

Publication date: The date of publication of the book (by that publisher).

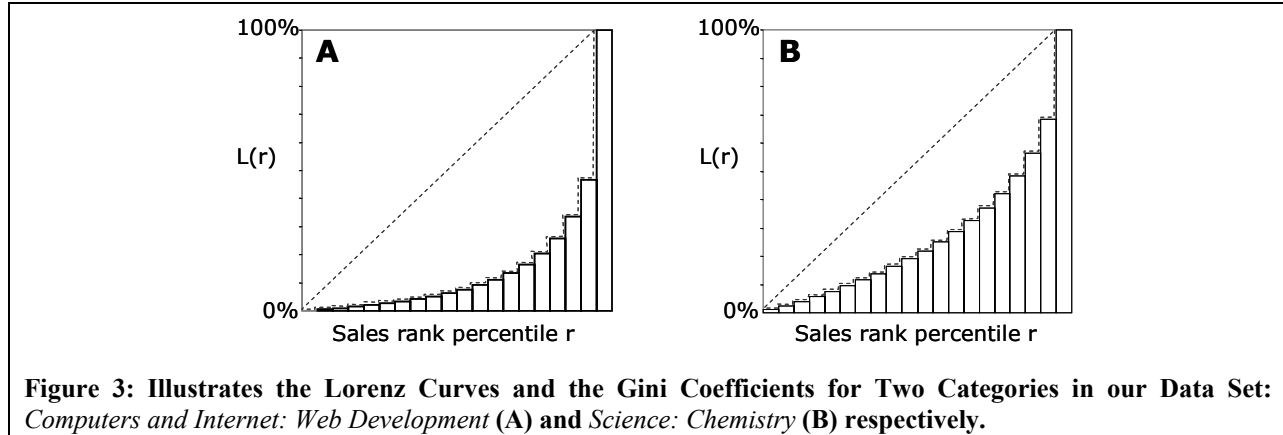
Characterizing Ecommerce Demand and Its Distribution

In order to relate the network position of a product to variation in its demand, we do the following:

- (a) Infer demand levels from the SalesRank data reported by Amazon.com, thereby associating a periodic demand level with each product.
- (b) Characterize the extent to which the network structure influences a product based on its network position.
- (c) Associate variation in (b) with variation in (a) at both a product-specific level of analysis and at a group-specific level of analysis.

Estimating Demand from Amazon.com SalesRank Data

To estimate the actual level of demand ($Demand(j)$) of a book from its sales rank ($SalesRank(j)$), we use a conversion model suggested by Goolsbee and Chevalier (2003) and by Brynjolfsson et al. , (2003).



$$\text{Log}[\text{Demand}(j)] = a + b\text{Log}[\text{SalesRank}(j)]$$

This formula to convert SalesRank information into demand information was first introduced by Chevalier and Goolsbee (2003). Their goal was to estimate demand elasticity. Their approach was based on making an assumption about the probability distribution of book sales, and then fitting some demand data to this distribution. They choose the standard distributional assumption for this type of rank data, which is the Pareto distribution (i.e., a power law).

In a later study, Brynjolfsson et al., (2003) use data provided by a publisher selling on Amazon.com to conduct a more robust estimation of the parameters of the formula. They estimate the parameters:

$$a = 10.526, b = -0.871.$$

We have used the latter estimate in our study. In future work, we propose to conduct an independent purchasing and demand estimation experiment in order to update these estimates. Note, however, that since our results are all based on comparisons between categories, the fact that these parameters are dated are unlikely to affect our results directionally.

Quantifying the Distribution of Demand: The Gini Coefficient

Next, we compute the Gini coefficient of each of the 200+ categories of books, for each of our four weeks. The Gini coefficient is a measure of distributional inequality, a number between 0 and 1, where 0 corresponds with perfect equality (in our case, where all the books in that category have the same demand) and 1 corresponds with perfect inequality (where one book has all the demand, and all other books have zero demand).

The Gini coefficient, based on the Lorenz curve, is a graphical representation of the cumulative distribution function of a probability distribution; it is a graph showing the proportion of the distribution assumed by the bottom $y\%$ of the values. It is often used to represent income distribution, where it shows for the bottom $x\%$ of households, what percentage $y\%$ of the total income they have. The percentage of households is plotted on the x-axis, the percentage of income on the y-axis. It can also be used to show distribution of assets.

To compute the Gini coefficient of each category one first ranks the books in the category in increasing order of their demand, thereby constructing the Lorenz curve. The Gini coefficient is calculated as a ratio of the area under the Lorenz curve to the area under the “perfect equality” line. Figure 3 illustrates this computation for two “extreme” categories in our data set. The dotted area is one half of the Gini coefficient. The category's Gini coefficients are 0.751 (A) and 0.502 (B) respectively. (The data has been binned for illustrative purposes in the figure).

The Gini coefficient is especially suitable for this study for a variety of reasons. Most importantly, it measures inequality in the demand distribution, regardless of the category's average demand (or popularity), which facilitates comparing different categories despite their intrinsic differences and independent of their scale.

Two Measures of Network Influence: Immediate Influence and Weighted PageRank

We have developed two different measures of network influence:

ImmediateInfluence: This is a measure of the traffic that flows into a product's Web page from its neighbors in the network. It is based on the assumption that the influence exerted by each product is proportionate to its total incoming traffic, is divided equally, and flows to those products it has direct co-purchase links to (note that this model does not allow back clicking). It therefore captures the influence of a product's immediate neighbors. Therefore, the influence that the co-purchase network has on demand depends on two factors: the local structure of the network and the amount of traffic associated with each link in the network. To combine the demand information with the structure of the network, we construct the *ImmediateInfluence* variable in the following way:

$$ImmediateInfluence(i) = \sum_{j \in G(i)} \frac{Demand(j)}{OutDegree(j)}$$

where $G(i)$ is the set of books that link to book i .

WeightedPageRank: This is based on Google's PageRank algorithm, and iteratively computes the influence of the entire network on each product over time, although ignoring variations in intrinsic traffic across pages. It operates on an "average graph", constructed as a weighted composite of a time series of co-purchase networks. The original PageRank algorithm provides a ranking of the "importance" of Web pages based on the link structure of the "Web" created by the hyperlinks between the pages. This ranking forms the basis for Google's search engine. The PageRank algorithm is based on a simple model of behavior – a random surfer. This surfer follows any one of the links on a page with equal probability or jumps to a random page with probability $[1 - \alpha]$ (this probability is also referred to as the "dumping factor"). The algorithm divides a page's PageRank evenly among its successors in the network. The ranking of a page ends up being the long-run steady-stage probability that a random surfer who starts at a random page will visit the specific page. Thus, a page can gain a high ranking by either having many pages pointing to it or having few highly ranked pages pointing to it. The PageRank of all pages in the network is computed iteratively, until some convergence estimator is met.

We adapt the PageRank algorithm to account for the fact that we wish to measure the average influence the network has on a product over a two-week period. In our adapted model:

$$WeightedPageRank(i) = \frac{1-\alpha}{n} + \alpha \sum_{j \in G(i)} Weight(j,i) \left(\frac{WeightedPageRank(j)}{\sum_{k \in F(j)} Weight(j,k)} \right)$$

It is important to note that while this measure is widely used in ranking algorithms (such as Google's), we use the fact that fundamentally Weighted PageRank measures the probability that a "random surfer" will arrive at a hyperlinked page if he were to traverse *just the hyperlinks* of the network. In other words, a product with a higher Weighted PageRank is more likely to get traffic from the network than one with a lower Weighted PageRank, and this therefore measures the *extent* to which the network structure we are interested in – the co-purchase network – *influences* the product in question.

The Weighted PageRank and the Immediate Influence measures described above are two different measures, which differ in the following key ways:

1. Weighted PageRank does not take the demand or intrinsic traffic variation across books into account. It is based only on the structure of the network. In contrast, Immediate Influence is based on both the structure of the network and the demand associated with each page.
2. The two measures use information about the structure of the network differently. Immediate Influence only includes the information about the immediate neighbors of the page, while Weighted PageRank measures the influence of the entire network.

How Network Structure Influences Demand and Its Distribution

The results we discuss in what follows were obtained using data for four distinct one-week periods between February 1st and February 28th. There is a seasonal demand pattern associated with sales around Valentine's Day in




Figure 4: Distribution of Log(ImmediateInfluence) and Log(Demand)

our data set, and we observed substantial changes in the edges of the co-purchase graph during this period (close to 20% of the edges changed). As will be seen, despite such changes in the identities of the books linking to each other, our results remained relatively stable.

We first study the variation between demand for each individual product, and the corresponding ImmediateInfluence of that product. Our results (available on request) indicate that the immediate influence of a product explains a significant amount of the variation in the demand for the product. Our final section reports on refinements we are working on that account for potential endogeneity in these estimates.

What is more pertinent to our main results is the contrast between the distribution of demand and Immediate Influence, which is illustrated in Figure 4. A casual examination of the distribution of demand and influence suggests the effect that the network structure might have on the demand distribution. As required by the definition of ImmediateInfluence, both have the same mean (influence is simply demand being redistributed), but the range and standard deviation of influence are larger than that of demand. This leads one to suspect that suspect that the network redistributes demand in a more “equal” manner, a hypothesis we report on in the next section.

Evidence of How Network Structure Intensifies the Long Tail” of ECommerce Demand

We now “zoom out” to an aggregated level (the category affiliation) and study how the network structure affects the market and the demand distribution. To study the effect of the network structure on the demand distribution, we group the books according to category affiliation. We test the hypothesis that a higher average Weighted PageRank for a category will be associated with a lower Gini coefficient for that category. Following our interpretation of PageRank, a category whose products collectively have a higher average Weighted PageRank, is, all else being equal, one whose products are influenced more by the network. The summary statistics for average Weighted PageRank (and other variables we use as controls) are presented in Table 1.

Variable	Range	Mean	StdDev
GINI	0.39 to 0.96	0.55	0.12
AVGDEMAND	0.92 to 22.02	3.04	3.07
AVGPAGERANK	1.93 $\times 10^{-8}$ to 6.03 $\times 10^{-8}$	3.36 $\times 10^{-8}$	6.32 $\times 10^{-7}$
PAGERANKVAR	3.97 $\times 10^{-11}$ to 4.23 $\times 10^{-10}$	5.74 $\times 10^{-11}$	6.73 $\times 10^{-11}$
SIZE	100 to 11,179	1,087	1,722
MIXING	0.01 to 0.80	0.32	0.18

Table 1: Representative Summary Statistics, for 2/1-2/7. The statistics of the other weeks are similar.

SIZE is the number of books in the category, and MIXING is a control for the extent of assortative mixing in the category’s co-purchase links. We tested the average price and the variance in price independently as controls, but neither was significant in any of our estimations.

Using this data, we estimate the following reduced-form econometric model:

$$\begin{aligned}
 \text{Log}[\text{GINI}] = & a + b_1 \text{Log}[\text{AVGDEMAND}] + b_2 \text{Log}[\text{AVGPAGERANK}] \\
 & + b_3 \text{Log}[\text{PAGERANKVAR}] + b_4 \text{Log}[\text{SIZE}] + b_5 \text{Log}[\text{MIXING}]
 \end{aligned}$$

Variable	Coefficient	Estimated Values (Standard Error)			
		2/1-2/7	2/8-2/14	2/15-2/21	2/22-2/28
constant	<i>a</i>	-1.97 (0.39) ***	-1.93 (0.37) ***	-2.19 (0.45) ***	-2.05 (0.43) ***
log[AVGDEMAND]	<i>b1</i>	0.26 (0.00) ***	0.25 (0.00) ***	0.23 (0.00) ***	0.24 (0.00) ***
log[AVGPAGERANK]	<i>b2</i>	-0.15 (0.04) ***	-0.15 (0.04) ***	-0.18 (0.04) ***	-0.17 (0.04) ***
log[PAGERANKVAR]	<i>b3</i>	0.03 (0.00) ***	0.03 (0.00) ***	0.03 (0.00) ***	0.03 (0.00) ***
log[SIZE]	<i>b4</i>	0.03 (0.00) ***	0.03 (0.00) ***	0.03 (0.00) ***	0.04 (0.00) ***
log[MIXING]	<i>b5</i>	-0.01 (0.00)	-0.01 (0.00)	-0.02 (0.00) *	-0.02 (0.01) *
	Adj. R-squared	86.85%	82.72%	81.04%	82.36%

indicates significance at the * 5%, ** 1%, and ***0.1% levels

Table 2: How Network Structure Affects the Distribution of ECommerce Demand

We chose a logarithmic specification because it facilitates ease of interpretation of the coefficients (in percentage terms), and because the empirical distributions of the transformed variables were more suitable for OLS regression.

The results of this estimation, presented in Table 2 are striking. Based on a comparative analysis across over 200 categories of products, they establish that an increase in the extent to which the network structure is influential leads to a flattening of demand, or an increase in the relative demand for niche (rather than blockbuster) products. That is, we find that the average *Weighted PageRank* of the books in the category (the AVGPAGERANK variable) is negatively associated with the Gini coefficient of the category. This confirms that an increase in the extent to which network structure influences demand flattens the distribution of demand, or leads to a longer tail for demand, a phenomenon widely observed in electronic commerce (Anderson, 2004). Notice that this coefficient is not just statistically significant but is economically significant as well. The highest average PageRank is generally about three times the lowest average PageRank. A doubling of average PageRank decreases the Gini coefficient by between 15% and 18% (since this is a log-log regression), which is pretty close to one standard deviation of Gini relative to its mean. Surprisingly, these results persist across four different weeks, one of which had substantial seasonal variation.

Moreover, the variance of the *Weighted PageRank* of different books within a category is positively correlated with the category's Gini coefficient. That is, after controlling for differences in average *Weighted PageRank*, a higher variance in the ranking (measured by PAGERANKVAR) is associated with increased inequality. To understand this result, consider two categories, both with the same average *Weighted PageRank*. Category A, where all books has the same *Weighted PageRank* and Category B, where there are a few books with a much higher than average *Weighted PageRank*, and correspondingly a number of books with a lower than average *Weighted PageRank*. It seems reasonable to expect that the flattening effect will be stronger for category A than for category B. After all, most of the traffic that goes into category B goes to the same few books and is likely to enhance the inequality in demand, thus increasing the Gini coefficient. In contrast, all books in category A get the same additional traffic from the network, so the relative differences in demand decrease, thus flattening the demand distribution.

The number of books in a category has a positive effect on the Gini coefficient. The categories in our data had between 100 and over 10,000 books in them. It is natural to assume that when all else is equal, a category with over 10,000 books is more likely to have higher variance in the demand for its books than a category with about 100 books.

Further, the average demand of the category has a positive effect on the Gini coefficient of the category. A straightforward interpretation of these results is that as the intrinsic demand increases, the added demand due to network traffic has a lower *relative* effect on the distribution of demand. To understand this result, consider two categories, both with the same average *Weighted PageRank*. Category A, with low average demand and Category B, with high average demand. Since both categories have the same average *Weighted PageRank*, they receive the same traffic from the co-purchase network (same number of consumers "flowing in"). This means they sell the same number of books to consumers who arrived at the books' pages via the co-purchase network. The network traffic has a flattening effect in both cases. In other words, the fraction of demand, which can be attributed to the best-selling books, is lower. However, the impact that same number of additional copies sold will have on the fraction of demand that come from the best-selling books will be lower for category A. Thus, since the traffic from the network accounts for a smaller fraction of category A's sales, the flattening effect will be smaller in magnitude.

Summary and Ongoing Work

We have briefly outlined a new economic theory of how network structures in electronic commerce might affect demand and cause ecommerce demand to be different from what is observed in traditional bricks-and-mortar commerce. We have gathered a new and unique data set comprising daily observations of a giant component of the co-purchase network of Amazon.com, along with the relevant economic variables for each of its constituent products. We have provided the first evidence that the presence of these network structures can cause changes in demand patterns that are consistent with the observed “long tail” of ecommerce demand. We do so by adapting the PageRank algorithm to measure the influence that the network structure has on each product, and then contrasting variations in the average such measure across categories that have different demand distributions. To the best of our knowledge, ours is the first study of its kind.

Our research in progress will extend these results in the following salient ways:

- Rather than being the “random” surfers used by the PageRank model, ecommerce consumers are strategic. They tend to visit more popular products more often, and their purchasing is affected by other economic variables like price, customer reviews and product age. Our first extension aims to develop a model of a “strategic surfer” that is grounded in more familiar economic theory but with sufficient structure to allow the iterative estimation of the “importance” of the network. We have made substantial progress on this front, solving a first model, and its estimation is in process.
- The methods we have used so far do not explicitly separate demand that is caused by the presence of the hyperlinks associated with the network structure with the demand variation that complementary products might naturally experience together. Identifying these distinctly requires appropriate instrumental variables. We have begun experimenting with constructing suitable variables of this kind (using lagged demand, and contrasting demand for identical products over successive days over which the link appeared). Since we have a time series of over 300 days of data, with well over 200,000 products per day, we are optimistic that we can identify this suitably for a subset of our data. This will facilitate a clearer understanding of the value of such network structures as strategic IT design variables that are unique to ecommerce.

References

- Anderson, C. The Long Tail. http://www.wired.com/wired/archive/12.10/tail_pr.html
- Aral, S. Brynjolfsson, E., and Van Alstyne, M., 2006. Information, Technology and Information Worker Productivity: Task Level Evidence. Mimeo, MIT.
- Brin, S., and Page, L., 1998. The anatomy of a large-scale hypertextual Web search engine. *Computer Networks and ISDN Systems*, 33:107–117.
- Brin, S., Page L., Motwami, R., and Winograd, T., 1999. The PageRank citation ranking: bringing order to the Web. Technical Report 1999-0120, CS Department, Stanford University.
- Brynjolfsson, E., Smith, M. D., and Hu, Y. J., 2006. From Niches to Riches: Anatomy of the Long Tail. *Sloan Management Review* 47(4), 67-71.
- Brynjolfsson, E., Smith, M. D., and Hu, Y. J., 2003. Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers. *Management Science*, Vol. 49, No. 11.
- Butler, B., 2001. Membership Size, Communication Activity, and Sustainability: The Internal Dynamics of Networked Social Structures. *Information Systems Research* 12, 346--362.
- Erdem, T., Imai, S., Kean, M. P., 1999. Econometric Modeling of Competition: A Multi-Category Choice-Based Mapping Approach. *Journal of Econometrics* 89.
- Goolsbee, A., and Chevalier, J. A., 2003. Measuring Prices and Price Competition Online: Amazon and Barnes and Noble. *Quantitative Marketing and Economics* (1) P. 203-222.
- Hill, S. F. Provost and C. Volinsky, 2006. Network-based Marketing: Identifying Likely Adopters via Consumer Networks. *Statistical Science* 21, 256-276.
- Kane, J. and Alavi., M., 2006. Casting the Net: Towards a Theory of Multimodal Knowledge Networks. Mimeo, Emory University.
- Langville, A., and Meyer, C., 2005. Deeper Inside PageRank. *Internet Mathematics*, 1(3):335--380.

- Matthias, S., Bauer, H. H., and Hammerschmidt, M., 2002. Structuring Product Markets: An Approach Based on Customer Value. *Journal of the American Marketing Association*.
- Ramaswamy, V., DeSarbo, W. S., 1990, SCULPTRE: A New Methodology for Deriving and Analyzing Hierarchical Product-Market Structures from Panel Data. *Journal of Marketing Research* 27(4)
- Russell, G. J., Kamakura, W. A. 1997, Modeling Multiple Category Brand with Household Basket Data. *Journal of Retailing* 73(4)
- Schoberth, T. Preece, J. and Heinzl, A., 2003. Online communities: a longitudinal analysis of communication activities. *Proceedings of the 36th Hawaii International Conference on Systems Sciences*.
- Sundararajan, A., 2006. Local Network Effects and Complex Network Structure. *Contributions to Theoretical Economics* 6 (1), in press.
- Watts, D. J. and Strogatz, S. H., 1998, Collective dynamics of 'small-world' networks. *Nature* 393