MIT Sloan School of Management

MIT Sloan School Working Paper 4655-07
7/21/2007

Long Tail or Steep Tail? A Field Investigation into How Online Popularity Information Affects the Distribution of Customer Choices

Catherine Tucker, Juanjuan Zhang

Long Tail or Steep Tail?
A Field Investigation into How Online Popularity Information Affects the Distribution of Customer Choices

July 21, 2007

Catherine Tucker

Assistant Professor of Marketing
MIT Sloan School of Management
1 Amherst Street, E40-167
Cambridge, MA 02142
Phone: (617) 252-1499
Fax: (617) 258-7597
Email: cetucker@mit.edu

Juanjuan Zhang

Assistant Professor of Marketing
MIT Sloan School of Management
1 Amherst Street, E40-171
Cambridge, MA 02142
Phone: (617) 452-2790
Fax: (617) 258-7597
Email: jjzhang@mit.edu

Long Tail or Steep Tail?
A Field Investigation into How Online Popularity Information Affects the Distribution of Customer Choices

Abstract

The internet has made it easier for customers to find and buy a wide variety of products. This may lead to a "long tail" effect as more customers buy low-volume products. However, the internet has also made it easier for customers to find out which products are most popular. This could lead to a "steep tail" effect as customers flock towards the most popular products. Using data from a field experiment with a website that lists wedding service vendors, we find empirical evidence that a steep tail exists. The most popular vendors become more popular when customers can easily observe previous customers' click-through behavior. Then, we ask whether this steep tail effect "complements" the long tail, by attracting customers who would otherwise have chosen nothing, or "competes with" the long tail, by shifting customers from less popular vendors to popular ones. We find evidence of a complementary effect, where the steep tail indicates new interest in the most popular vendors from outside, with negligible cannibalization of interest for less popular vendors. The findings suggest that popularity information can serve as a powerful marketing tool that facilitates product category growth. They also explain the prevalence of firm practices to highlight bestsellers.

Keywords: Long Tail, Steep Tail, Customer Learning, Decisions Under Uncertainty, Internet Marketing, Category Management

1 Introduction

The internet has changed how customers shop. In particular, websites' ability to automate the process of matching customers with products has made buying previously niche products easier. This has led some researchers to speculate that there is a "long tail" for e-commerce, where more customers buy previously unpopular products. However, websites are often designed to attract customer attention to their most popular products and listings. Among the top 100 U.S. websites, 60 percent display popularity information about what products or web links previous customers chose. This information could lead to bandwagon behavior or a "steep tail" effect, if customers are attracted by the more popular products.

In this paper, we investigate whether the display of information on previous customer choices leads to a steep tail. We then investigate whether the steep tail competes against or complements the long tail. The steep tail would compete with the long tail if customers who would have bought less popular products second-guessed their choice and switched to a more popular product. The steep tail would complement the long tail if marginal customers, who without a strong quality signal were content not to buy anything, were convinced to buy a popular product because they inferred high quality from high popularity.

We use data from a website portal that lists different wedding vendors. This website experimented with shifting from a traditional, yellow pages style of alphabetical listing, where no popularity information was displayed, to a more contemporary style where both popularity information was displayed and listings were ranked by their popularity. The website measured vendor popularity by how many times customers clicked on its URL. We compare the shift in the distribution of clicks in this new format with the original format. We find that while the more popular vendors experienced an increase in popularity, the less popular vendors experienced an asymmetric and nearly negligible decrease in popularity. One concern is that instead of measuring bandwagons, we may merely be capturing "website real estate effects", where a more salient page location leads to a high volume of clicks before and after the experiment. To control for this, we use data from a parallel test where the vendors were reordered by popularity but popularity information was not displayed. This result confirms that, while there is a website real estate effect for those vendors hidden at the bottom, the effect of displaying popularity information dominates.

These findings could be encouraging to firms, because we find that the steep tail effect is a result of generating new interest from outside, rather than of cannibalizing interest from less popular products. In particular, our results are encouraging for websites that bring together multiple firms and customers and that are considering whether to embrace "Web 2.0" tools that give prominence to online social learning. This message is also relevant for off-line firms who sell multiple similar products within a product line and are wondering whether to highlight their best-sellers. It also helps us understand more broadly the strategies of websites such as Amazon, that offer an extensive product range but also publicize popularity information heavily on their webpages.

This paper contributes to several lines of research. The first is empirical research that studies the long tail of e-commerce (as popularized by Anderson (2006)). This research overwhelmingly finds that online retailers sell more products that are less popular than do traditional retailers. Brynjolfsson, Hu, and Smith (2003) examines how customers benefit from the increased variety of books online. Brynjolfsson, Hu, and Simester (2007) study the long tail for a multi-channel retailer and find evidence that customers with more internet experience are more likely to buy more obscure products. Recently, Oberholzer-Gee and Elberse (2007) study the long tail for movies and argue that customers with the aid of the internet have easier access to their favorite choice. They call the resulting demand boost for blockbuster movies the "superstar effect", based on theoretical work by Rosen (1981) that predicts a "winner-takes-all" outcome. Unlike these studies, we do not study how the internet has made it easier to find products. Instead, we study how the internet molds customers decisions where there is uncertainty and where customers can use online popularity information as a shopping tool.

The idea of the long tail is that there is an under-exploited spectrum of customer tastes to which pre-internet retailers could not cost-effectively cater. Internet retailers' low cost of finding and serving these previously hard-to-find customers means that they can serve them profitably. This expands the market for less popular products. The idea of the steep tail, on the other hand, is that internet retailers' low cost of automating their display of product popularity information facilitates observational learning by customers. Customers use popularity information to reduce uncertainty about product quality. This expands the market for more popular products. If the long tail and the steep tail effects complement each other, the overall market will expand. We provide empirical evidence that this is possible.

Our study also contributes to the literature on observation learning. Observational learning was first studied theoretically by Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992), who demonstrate the striking result that individuals may rationally repeat their predecessors' actions regardless of their own information. Evidence of observational learning has been documented both in the lab (e.g., Anderson and Holt (1997), Çelen and Kariv (2004)) and in the field (Zhang 2007). Most of these empirical studies isolate an environment where observational learning is likely to occur, and then identify observational learning through appropriate convergence measures. We introduce to this line of research the first online field experiment that exogenously enables or disables observational learning and then quantifies the impact on customer choices. To do so, we first generalize the traditional observational learning model to allow for multiple-unit choices, and then develop a differences-in-differences empirical framework within which to quantify the effects of observational learning.

Popularity information can affect other people's choices in two ways. First, when quality is hard to gauge, people may take quality cues from other people's prior decisions. This observational learning process can lead to repetitive social choices. Second, customers could also receive utility from conformity due to, for example, preference for social identification, sanction of deviants, and network externalities. This second mechanism seems less relevant in our empirical setting, because we are studying an industry where quality inference is important and where individual choices are rarely subject to pressure to conform. Therefore, we interpret our results as suggesting that websites can make observational learning easier by publishing popularity information.

Other disciplines outside of marketing and economics have also looked at how customers respond to information about others' choices. In IT, Fleder and Hosanagar (2007)'s analytical model of recommender systems suggests that such systems not only reinforce popular products, but also help customers discover new products and thereby increase sales diversity. In sociology, a series of experiments by Matthew J. Salganik (2006) studies how customers react to information about previous music downloads. Their results suggest that "stars" are made rather than born. We extend this line of research by focusing on the underlying marketing strategy of popularity information and by providing a formal framework of observational learning to analyze the results.

The paper is organized as follows. Section 2 analytically illustrates how observational learning can lead to a steep tail, and how a steep tail can complement or compete with the long tail. Section 3 discusses the field experiment design and implementation. Section 4 describes the data and presents the estimation results. Section 6 concludes the paper and discusses future research.

2 Analytical Illustration

This section develops an analytical model to illustrate how publicizing product popularity information can generate a steep tail effect, which may either complement or compete with the long tail through observational learning about product quality.

Assume there are J competing vendors, indexed by j, within the same category. Each vendor offers one product and privately knows its intrinsic value v. Customers are uncertain about product values. However, customers do know that v takes the value of either 0 or 1 with equal prior probability. In addition, each customer receives a private signal about how valuable each product is. These private signals are identically and (conditional on the true value of a product) independently distributed. They take either of two values: good (G) or bad (B). Suppose the conditional signal probabilities are $p(G|v_j = 1) = p(B|v_j = 0) = q$, where $1/2 < q < 1$, so that the signals are informative yet imperfect. Each customer incurs cost c in viewing that vendor's website. In this simplified illustration, we treat c as exogenously given. Customers are allowed to visit multiple vendors, and will visit vendor j if and only if $E(v_j) > c$.

We first allow the market to evolve naturally without publicizing any information about previous choices. In other words, each customer makes her purchase decisions based on her private signals. By Bayes' rule, the posterior belief about v_j after observing a good signal on product j is:

$$p(v_j = 1|G) = \frac{p(G|v_j = 1)p(v_j = 1)}{p(G|v_j = 1)p(v_j = 1) + p(G|v_j = 0)p(v_j = 0)} = \frac{q/2}{q/2 + (1-q)/2} = q$$

Therefore, $E(v_j|G) = q$ and $E(v_j|B) = 1 - q$. If the cost of visiting each vendor is low enough (i.e., $c < 1 - q$), a customer will visit every vendor regardless of the private signal she receives, since $E(v_j) \geq 1-q > c$ for any j. Similarly, if c is high enough (i.e., $c > q$), no customer will visit any vendor regardless of the private signal. For the rest of the analysis we focus on the nontrivial case where $1-q \leq c \leq q$, so that a customer will visit a vendor if and only if she receives a good private signal about its product.

We take a snapshot of the market outcome after N customers have individually decided which vendors to visit. Let n_j denote the number of visits, and equivalently the number of good signals, that vendor j has received. Note that when $N \to \infty$, $\frac{n_j}{N} \to q$ for vendors whose product is truly valued at 1, and $\frac{n_j}{N} \to 1-q$ for vendors whose product is truly valued at 0. In other words, before any previous choice information is publicized, a vendor's popularity is ultimately determined by the underlying value of its product and is independent of the actual visits it has received. However, subsequent customers' choices can change dramatically if we introduce online tools that publicize vendors' popularity information. Specifically, suppose N and n_j $(j = 1, \cdots, J)$ is now public information. If the $(N+1)^{th}$ customer receives a good (G) signal about vendor j, she knows that vendor j is associated with $n_j + 1$ good signals and $N - n_j$ bad signals. Therefore, her posterior expectation of v_j is:

$$\begin{aligned} E(v_j|n_j, N, G) &= \frac{p(n_j,N,G|v_j=1)p(v_j=1)}{p(n_j,N,G|v_j=1)p(v_j=1) \ + \ p(n_j,N,G|v_j=0)p(v_j=0)} \\ &= \frac{q^{n_j+1}(1-q)^{N-n_j}/2}{q^{n_j+1}(1-q)^{N-n_j}/2 \ + \ (1-q)^{n_j+1}q^{N-n_j}/2} \\ &= \frac{1}{1+(\frac{q}{1-q})^{N-2n_j-1}} \end{aligned}$$

Similarly, if the $(N+1)^{th}$ customer receives a bad (B) signal about vendor j, her posterior expectation of v_j is:

$$\begin{aligned} E(v_j|n_j, N, B) &= \frac{p(n_j,N,B|v_j=1)p(v_j=1)}{p(n_j,N,B|v_j=1)p(v_j=1) \ + \ p(n_j,N,B|v_j=0)p(v_j=0)} \\ &= \frac{q^{n_j}(1-q)^{N-n_j+1}/2}{q^{n_j}(1-q)^{N-n_j+1}/2 \ + \ (1-q)^{n_j}q^{N-n_j+1}/2} \\ &= \frac{1}{1+(\frac{q}{1-q})^{N-2n_j+1}} \end{aligned}$$

The $(N+1)^{th}$ customer will visit vendor j if her posterior expectation of v_j is greater than the cost c, where $1-q \leq c \leq q$. An interesting scenario is when $E(v_j|n_j, N, B) > q$, so that the $(N+1)^{th}$ customer will always visit vendor j regardless of her private signal. The visit of this customer to vendor j thus contains no information to subsequent customers, who will all visit vendor j regardless of their private signals. Similarly, when $E(v_j|n_j, N, G) < 1-q$, no subsequent customer will visit vendor j regardless of her private signal. We state the above results formally:

Proposition 1. *All subsequent customers will visit vendor j if $n_j > \frac{N}{2}+1$; no subsequent customer will visit vendor j if $n_j < \frac{N}{2}-1$.*[1]

In the above example, information cascades arise when customers disregard their own private signals and follow their predecessors' choices. A "downward cascade" arises when interest in a product dwindles due to its poor track record (i.e., $\frac{n_j}{N} \to 0$ after the information $n_j < \frac{N}{2}-1$ is publicized), as future customers hold less confidence in their positive private opinions, if any, about the product. On the other hand, an "upward cascade" arises when a popular item benefits from bandwagon effects and receives a further popularity boost (i.e., $\frac{n_j}{N} \to 1$ after the information $n_j > \frac{N}{2}+1$ is publicized). When either cascade happens, the resulting visit distribution across vendors exhibits a steep tail effect, compared to the natural market where $\frac{n_j}{N} \to q$ or $1-q$. When both the downward and the upward cascades arise, the steep tail effect competes with the long tail effect, because the most popular sellers get all the attention. However, when only an upward cascade arises, the steep tail effect can complement the long tail effect and generate extra interest in the category. Before popularity information becomes available, a fraction of customers would choose not to buy because no vendor's quality is perceived as being above the cost of buying. When popularity

[1] Proof: $E(v_j|n_j, N, B) > q \Leftrightarrow \frac{1}{1+(\frac{q}{1-q})^{N-2n_j+1}} > q \Leftrightarrow (1-q)^{N-2n_j+2} > q^{N-2n_j+2} \Leftrightarrow N-2n_j+2 < 0$ since $p > 1/2$. Similarly, $E(v_j|n_j, N, G) < 1-q \Leftrightarrow n_j < \frac{N}{2}-1$.

information is publicized, and when there exists a clearly well-received product, such strong social proof may outweigh the negative personal opinions and buying costs of those previously uninterested customers, converting them to the idea of making a purchase.

Whether a steep tail occurs and whether it competes with or complements the long tail effect, however, depends on a variety of factors, including vendor visiting costs (c), customers' knowledge about potential demand (N), taste heterogeneity, etc. These factors make it an open empirical question. In the rest of the paper, we examine the effects of publicizing popularity information on choice distribution across vendors through a set of controlled field experiments.

3 Field Experiment

3.1 Experimental Setting

We want to know how the availability of popularity information affects the distribution of customer choices within a category. Using historical data would present problems, because websites' decision whether to publicize popularity information can be linked to product characteristics in unobservable ways. For example, while Tiffany.com does not display any popularity information, it is not credible to assume that any differences in the distribution of customer choices for Tiffany's products compared to Amazon.com's products is due to Amazon's easily accessible sales ranking information. It is also hard to precisely measure the steep tail effect from longitudinal historical data by comparing the distribution of customer choices for a website before and after it displayed popularity information. For example, suppose Apple.com began publicizing bestsellers within its iPod line. This may be accompanied by Apple shelving less successful iPod models and launching other promotions that are unobservable to the researcher.

To get around this endogeneity problem, we use data from an internet portal website. This website was experimenting with ways to update their alphabetical yellow page listing style. They tested different types of popularity information format for different categories for two months. The website provided wedding vendor listings to brides in a New England state. [2] This website is representative of many online information portals that bring together vendors and customers.

Theoretically, the wedding industry is attractive to study, because there is little prior consumption experience. [3] Consequently, customers are likely to be equipped with imperfect information about vendors, and observational learning may become a more important part of their decision making. This is also an industry where vendor selection is a high-stakes decision that customers take seriously. On average, 2.3 million weddings take place in the US each year, accounting for $72 billion in annual wedding expenditures. Most brides invest considerable efforts in selecting vendors. During an average 13-month engagement, eight hours a week is spent planning.[4] A bride spends on average $967 on flowers, $1,276 on bridal attire, and $7,330 on the reception.[5]

Another reason for the wedding industry to be attractive to study is that different categories of wedding services can be used as natural controls for each other in a field experiment. This cross-category control would be problematic if we were studying an apparel retailer and we were trying to use interest in, say, sweaters to predict the interest in bathing suits. How-

[2] The geographic area that the website covered was representative of the US. The number of marriages in the state they operate in is proportional to the national average. The only observable deviation from national trends is that wedding costs are around $10,000 greater in this state than the national average of $27,000.

[3] Even if an individual organizes successive weddings, there is anecdotal evidence to suggest that they prefer to select different vendors in order to differentiate the current wedding from their previous wedding.

[4] Source: Association of Bridal Consultants from Bride's Magazine reader survey

[5] Source: American Wedding Study, 2002

ever, in the wedding industry different categories of services, such as catering and florists, are complementary components of the ultimate end wedding product. Therefore, it seems plausible to assume that these categories share similar unobserved shocks over time, and provide a “level playing field” upon which to measure treatment effects.

It is important to be clear that we are not expecting popularity information to increase the demand for weddings. Instead, we are interested in how popularity information affects customers’ decisions to click on the URL of the listed vendors. In particular, we first speculate that popularity information will influence whether a bride chooses to visit a particular vendor within a category. In fact, confidential information provided to us by the website suggests that a sizable proportion of browsers visit the listing page without ever following a link to a vendor, suggesting that brides are making a nontrivial to-visit-or-not-to-visit decision. Second, we speculate that popularity information may attract customers who would otherwise have chosen to seek wedding services from other channels, such as a national chain or a department store, rather than visiting one of the stand-alone vendors listed on the website. In other words, even if the total number of weddings stayed constant, popularity information might shift the number of vendors brides are interested in within a category, and across categories.

There were other advantages of this particular website setting. The website does not provide substantial information about the vendors, such as detailed vendor descriptions or customer reviews, which could decrease quality uncertainty and reduce the need for popularity information as a quality signal. Instead, it lists only the vendors’ name, location and telephone number. Crucially, vendor prices are not displayed. In other words, price is not a concern when customers make their click decision, and so cannot act as a alternative quality signal. Vendors, correspondingly, would have no incentive to strategically change the price to manipulate customer observational learning. This feature rules out the price endogeneity problem which would have been a key concern if the experiment had been run on a price-grabber style website.

3.2 Experiment Design

The website explored a variety of ways that it could present information about the popularity of each vendor. The site owners measured the popularity of a vendor by the number of previous clicks its link had received. As discussed by Baye, Morgan, Gatti, and Kattuman (2006), the number of clicks puts an upper limit on the distribution of demand. The website experimented with three different ways of presenting this information: Displaying click information alongside the previous alphabetical ranking; ranking the vendors by popularity but not revealing the number of clicks; and finally, ranking the vendors by popularity *and* displaying the number of clicks. Our empirical analysis focuses on comparing the two latter conditions.

The website conducted this experiment using four out of a total of 19 wedding service categories that received the most traffic and had the largest number of listings: Caterers, Reception Halls, Bridal Attire and Florists. These four categories were randomly allocated into the four experimental conditions. Table 1 summarizes the assignment and the pre-experiment traffic level of the four categories. We witnessed and verified the randomness of this allocation. Florists were the control category and retained their alphabetic ordering with no display of clicks. Reception halls retained their alphabetical ordering but displayed information about previous clicks. Caterers displayed no information about previous clicks but were listed by the number of previous clicks, with the vendor receiving the most clicks being listed first. Bridal shops not only had the number of prior clicks displayed, but were also ranked in descending order of popularity.

The field experiment ran for two months, from August to September 2006. The number of previous clicks was calculated using a base date of six months prior to the field experiment. The website did not disclose to visitors any information about the start date for this stock of clicks. This lack of disclosure resembles the practice of other firms, and it also means that customers are not confused by additional cues of seasonality, etc. The number of clicks is put in an extra cell of the html table for each vendor, in a column entitled "clicks". In the control condition, this column was unlabeled and empty. There is no difference in the webpage format across conditions, except for the display of click information. Every three days we ran a screen-scraping program to ensure that there were no glitches in the experiments and to verify the data provided to us by the website.

Popularity-based positions are based on the current stock of clicks, including the clicks collected while the experiment ran. Therefore, theoretically, the ordering of vendors could have changed as the experiment proceeded. On average each vendor moved a maximum of 3.36 places during the experiment. The largest shift was 8 places. Therefore, we use daily data to reflect the fact that the ranking of a vendor could change regularly.

4 Data

The firm collected data on browsing behavior based on their Apache Web Server logs. To protect the privacy of their users, they removed IP address information and created a dataset they could share with us. In this dataset, each observation is a time-stamp for when a link received a click, alongside the vendor and category that received this click. The data span the two months prior to the field experiment (June and July 2006) and the two months of the field experiment (August and September 2006). During these four months, there were 860,675 clicks across all categories. The four categories in our experiment accounted for 515,121 of these clicks. There is a total of 346 vendors listed within the four selected categories: 52 in florists, 155 in reception halls, 66 in caterers, and 73 in bridal shops. Figure 1 illustrates how the dependent variable of daily clicks varied across vendors. While the average vendor received 4.9 clicks each day, there were a few "popular" vendors who received over 15 clicks a day, together with a long tail of less popular vendors receiving only 1 click a day.

There are several challenges in processing the data. The first challenge comes from unintentional clicks, due to, for example, slow website response time. Since privacy rules prevented our accessing the IP addresses, we could not identify repeat clicks by the same user. As an alternative strategy, we dropped 60,925 observations where there were multiple requests for the same link within the same minute. To check the sensitivity of our results to this procedure, we also tried dropping observations where there were more than five requests for the same link within the same minute. There was no substantial change in our findings.

The second processing challenge was that there was a small amount of vendor entry into and exit from this information portal during the period we study. In the reception hall category, there was one change where a reception hall with a name beginning with "O" was removed during the second month of the experiment and a vendor which began with "T" was put in its place. This shifted the position of all reception halls with first letters "P" to "S" up one place for the second month of the experiment. The reception halls affected ranked initially between 95 and 114 and subsequently ranked between 93 and 113. In the florists category, there was one change where a florist beginning with "V" was removed in the second month of the experiment. One florist beginning with "W" was affected by this change and moved position from 58 to 57. We have experimented both with incorporating these slight changes in position in our estimation and excluding them. The results are almost identical, so for simplicity of presentation, we ignore changes in position due to vendor entry and exit in the tables below.

Each vendor moved a different amount up or down the page depending on the alphabetical ordering of their name relative to the number of clicks they had previously generated. Figure 2 makes it clear that the shift from alphabetical to popularity-based ranking led to a large shift in the way vendors were presented on the page. It is not the case, for example, that the most popular vendors began with the letter A, and that therefore the experiment merely reinforced a pre-existing relationship between position in the alphabet and popularity. We exploit this variation in our empirical analysis to tease apart the confound from vendor page position.

5 Empirical Analyses

5.1 Initial Graphical Analysis

We want to learn whether there is a steep tail effect, and whether it competes with the long tail effect by shifting clicks from less-popular vendors to the popular ones, or complements the long tail effect by boosting the popular vendors without sacrificing the less popular ones.

As an initial exploration of the data, we performed graphical analysis of the clicks in the different conditions. Figure 3 compares the effect on the distribution of clicks for the control condition (corresponding to the tradition yellow-page style of listing) with that for the condition, corresponding to the most contemporary listing practice, where both click information was available and the vendors were arranged by popularity on the page. For ease of comparison we detrend the category data, order the vendors by number of cumulative clicks, and compare pre- and post-clicks on the same graph. While there was little difference for the florists, there was an increase in interest for the most popular bridal shop vendors. This popularity boost appears to dwindle by the twentieth most popular vendor.

One possible criticism of this simple comparison is that it ignores the possible confound of page position. This page location, or "Website Real Estate" effect, could occur for two reasons. First, customers might visit only the highest listing because they incur search costs from scrolling down. Second, there may be subconscious psychological forces at work drawing customers' to the listings at the top. Eyetracker studies conducted on yellow pages by (Lohse 1997), for example, have confirmed increased customer eyeball attention to initial listings. The website real estate effect can confound observational learning when customers repeatedly click on the most popular vendors who are listed at the top of the webpage.

In order to control for this page location confound, we use as our control the experimental condition where vendors were ranked by popularity but visitors received no information about why they were ranked in that way. This allows us to separate out the effect on interest that was due merely to the vendor moving up the page, and the effect that is due to the explicit display of popularity information.

Figure 4 compares the condition where vendors are listed by popularity but the rationale behind the rearrangement (the number of clicks previously received) is not revealed, with the condition where vendors are listed by popularity and where information on number of previous clicks is available. It appears that we are not picking up a positional effect, since the main boost in clicks comes for the top vendors when popularity information is revealed. In the next section, we use difference-in-difference regression analysis to quantify more precisely the effect of different display and ranking policies on click-through behavior. We compare the two conditions just mentioned to control for the confound of page position.

5.2 Exploring the Steep Tail

In all our regressions, the dependent variable is the daily amount of clicks that a vendor receives. We want to find out how this is affected by popularity information. However, the random assignment of treatment conditions occurred across categories rather than within categories, so in our regressions we need to account for seasonal effects or other systematic differences across categories. We use a differences-in-differences approach to control for these systematic differences in customer click behavior across categories and across time.[6]

We cannot simply compare differences in the amount of interest expressed for the most popular caterers and the most popular bridal attire shops. If brides systematically browse bridal attire shops more than caterers, we would wrongly interpret this category effect as a steep tail effect. Therefore we include vendor-specific fixed effects α_j for each vendor j to control for static differences in base demand.

We also need to control for changes in how many brides are shopping for vendors in August and September compared to June and July. If we only study demand before and after these periods for bridal attire shops, we could wrongly confound a change in seasonal demand with a steep tail effect. Therefore, we use the baseline condition for caterers as a control for the general time trend for the category. We capture this time trend by a vector X_t of weekly dummies and day-of-week dummies.

With these controls in place, we can go on to estimate the interactions of interest to us, that explore whether and how popularity information affects relative customer choices across categories. We estimate a model specification that captures the effect of being popular (i.e. a vendor moving up the page when we switch from alphabetical ordering to popularity-based

[6]Chevalier and Mayzlin (2006) use a similar methodology to study the effect of online word of mouth on book purchases. Anderson, Fong, Simester, and Tucker (2007) use this approach to study the effect of sales taxes on apparel purchases.

ordering) on how interested customers are in that vendor. We estimate the coefficients for interactions between vendor movement ($MovesUp_j/MovesDown_j$ equals 1 if the vendor moves up/down), an indicator for whether clicks are displayed $display_j$, *and* an indicator variable $test_t$ for whether an observation takes place during the experiment. Equation 1 summarizes our specification, where α_j and the βs are parameters to be estimated.

$$
\begin{aligned}
clicks_{jt} &= \alpha_j + \beta_1\ X_t + \beta_2\ MovesUp_j * test_j + \beta_3\ MovesDown_j * test_j \qquad (1)\\
s &+ \beta_4\ display_j * MovesUp_j * test_j + \beta_5\ display_j * MovesDown_j * test_j
\end{aligned}
$$

Table 2 shows the results for an OLS estimation of equation 1, together with the regression results with robust standard errors and clustering of standard errors, which we shall discuss in the robustness section. Overall, the already-popular vendors receive a significant further boost in clicks ($p < 0.01$) when the prior number of clicks is displayed. The relatively less popular vendors (those who moved down the page), on the other hand, suffer an insignificant reduction in clicks ($p > 0.10$) when prior click information is displayed. This asymmetry suggests that first, there does exist a steep tail effect that widens the gap in clicks between the popular and less popular vendors, and second, the steep tail complements the long tail by drawing extra interest in the most popular vendors, rather than cannibalizing interest from vendors who are less popular.

There is a small negative effect for less popular vendors that appear further down the page independent of the display of clicks. There is no statistically significant effect, however, from a vendor moving upward on a page, if the rationale behind that movement (increased number of clicks) is not displayed. These results imply that customers only react positively to an improvement in page location if such an improvement is justified by higher popularity.

In other words, the importance of webpage location might have been overstated in practice, as page location in isolation might not be sufficient to attract visitors who might well be driven by other quality cues such as popularity.

A slightly modified specification reinforces this finding. Instead of using the MoveUp and MoveDown dummy variables, we include on the right-hand side a variable "*ChangePosUpwards*" and "*ChangePosDownwards*" that reflects the absolute increase or decrease in a vendor's page position. Therefore, the interactions now measure the linear effect of the number of page positions a vendor leapfrogs due to popularity. Table 3 shows the results. While there is a slight boost from moving up in page position for the popular vendors, again it is the combination of moving up the page *and* popularity information being displayed that leads to a boost in interest. Again, the effect is asymmetric, and there is a far smaller negative effect from moving down the page due to a lack of popularity. Unsurprisingly, given that this specification assumes a completely linear relationship between page location movement and clicks, the significance of our estimates is slightly lower relative to the previous specification 1.

5.3 Robustness

Previous researchers have highlighted how carefully the significance of differences-in-differences estimates should be interpreted, if the specification covers multiple time periods (e.g., Bertrand, Duflo, and Mullainathan (2004)). The reason is that the repeated use of the same exogenous change in variables can lead researchers to overstate the significance of the estimates. To address this concern, we used three broadly accepted techniques. First, we clustered the standard errors by whether the observation was in the pre- or post-test time period. Second, we repeated our estimation using only single blocks of the pre- and post-test data. Third, we repeated our estimation using a Poisson QML regression with fixed effects and errors clustered at the category level, as suggested by (Hausman, Hall, and Griliches 1984). Our results were reassuringly similar in size and significance.

One concern with studying the wedding industry is that any experiments could be confounded by seasonal changes in level of interest in wedding vendors. This is why we have such a rich set of controls to capture changes in cross-category interest over time. The monthly controls did not show any large variation. Table 4 provides some explanation. This table shows that entry and exit for this industry is more evenly spread across the year than the conventional belief in the prevalence of summer weddings would suggest. The largest monthly shock is in December, when 19 percent of engagements happen. By contrast, there is less variation in how many weddings take place each month. June and July, commonly supposed to be the most popular months for weddings, only account for 10.5 percent of the interest in wedding vendors on average.

In any differences-in-differences specification it is important that the "control" condition shares a similar time-trend to the treatment condition. In other words, we need the caterers category to have similar time-variant demand shocks to the bridal attire category. According to industry experts we asked, the demand for wedding services does vary somewhat with the season, but it is reasonable to believe that various categories are subject to similar levels of seasonal shocks due to the high degree of complementarity across categories for the end product.

Additional confounds could arise if a rival website started providing listings of, for example, wedding caterers during our field experiments, which would plausibly reduce the visits to caterers on the website running our experiment. That would lead us wrongly to label the decrease in popularity compared with the bridal shop category as a steep tail effect. Fortunately, during the time period we study, this website had no significant local competitors in the state it operates in. National competitors, such as "TheKnot.com" and "WeddingChannel.com", did not change their listing policies.

5.4 Analysis of Aggregate Effect

To investigate the steep tail, we use information on the distribution of clicks across vendors for two of the experimental conditions. However, websites who want to know which format they should use are interested in comparing the aggregate effects of all the experimental conditions. To evaluate this, we employed a similar identification strategy, that contained interactions for the aggregate effects of all the different conditions.

$$\begin{aligned} clicks_{jt} = \alpha_j + \beta_1\ X_t + \ test_j + \beta_2\ display_j * test_j + \beta_3\ ranked_j * test_j \\ + \beta_4\ ranked_j * display_j * test_j \end{aligned} \tag{2}$$

$display_j$ equals 1 if the number of prior clicks that vendor j has received is displayed, and 0 otherwise. Similarly, $ranked_j$ equals 1 if vendor j belongs in the condition where listings are ranked in descending order of popularity. While $display_j$ and $ranked_j$ capture the effects of click display and popularity-based ranking, $display_j * ranked_j$ captures the interaction effect.

The results of the regressions are reported in Table 5. The listing strategy which boosts clicks most is the one that both displays the number of prior clicks and ranks vendors by popularity. It seems that publicizing popularity information and making such information salient through popularity-based ranking is the best way to increase overall interest in the category. This is consistent with our steep-tail hypothesis.

6 Conclusion

The Internet can generate a long tail effect, where a large number of small-volume vendors coexists with a few high-volume ones. In this paper, we ask whether the internet also generates a steep tail effect through the growing practice of publicizing product popularity information. A steep tail happens when customers infer high quality from high popularity which leads to an increase in interest in popular products. We explore this steep tail effect empirically, using data from a controlled field experiment conducted by an online portal for wedding vendors. The website experimented with whether the number of prior clicks was displayed, and with whether vendors were positioned on the page alphabetically or in decreasing order of popularity.

We find strong evidence for a steep tail effect, where customers are more likely to click on the most popular vendors when the popularity information is publicized and made salient through ranking the vendors on the page by popularity. We also find that this steep tail effect complements rather than competes with the long tail effect, because the boost for popular vendors does not cannibalize the clicks on less popular vendors, for whom the change in display format changes little. We infer that this boost in interest comes from customers who were initially uncertain about pursuing any of the vendors.

Our findings can partly explain the widespread practice of web-based popularity information display. If a steep tail effect exists, and if it complements the long tail, websites such as Google.com and Digg.com can increase overall number of clicks at little cost to the less popular listings. Though we have focused on implications for e-commerce, our results also support the common strategy of multi-product firms that reveal information about which products are their best-sellers. Our findings help us to understand the apparent contradiction between the many websites that seem to focus on pushing their most popular products, and a host of internet strategists who emphasize how important it is for a website's success to sell niche products.

There are several potential ways of building on this research. One possibility is to explore how the publication of popularity information affects the dynamics of a two-sided market, where sellers can strategically manipulate and react to popularity information. In our setting, no price information is published, but in many consumer retail settings, price can be used strategically by firms to win back customers and increase their popularity. Another possible avenue would be to analyze the effectiveness of various popularity information tools websites use, such as separating off and publicizing a list of the top ten most popular products sold on the website.

References

Anderson, C. (2006, July). *The Long Tail: Why the Future of Business Is Selling Less of More*. Hyperion.

Anderson, E., N. Fong, D. Simester, and C. Tucker (2007). Do internet tax policies place local retailers at a competitive disadvantage. Working paper, MIT Sloan School of Management.

Anderson, L. R. and C. A. Holt (1997). Information cascades in the laboratory. *American Economic Review 87*(5), 847–862.

Banerjee, A. V. (1992, August). A simple model of herd behavior. *The Quarterly Journal of Economics 107*(3), 797–817.

Baye, M., J. Morgan, J. R. J. Gatti, and P. Kattuman (2006, November). Clicks, discontinuities, and firm demand online. Mimeo, Berkeley.

Bertrand, M., E. Duflo, and S. Mullainathan (2004, February). How much should we trust differences-in-differences estimates? *The Quarterly Journal of Economics 119*(1), 249–275.

Bikhchandani, S., D. Hirshleifer, and I. Welch (1992, October). A theory of fads, fashion, custom, and cultural change in informational cascades. *Journal of Political Economy 100*(5), 992–1026.

Brynjolfsson, E., Y. Hu, and M. D. Smith (2003, November). Consumer surplus in the digital economy: Estimating the value of increased product variety at online booksellers. *Management Science 49*(11), 1580–1596.

Brynjolfsson, E., Y. J. Hu, and D. Simester (2007). Goodbye Pareto Principle, Hello Long Tail: the Effect of Search Costs on the Concentration of Product Sales. Mimeo, MIT.

Çelen, B. and S. Kariv (2004). Observational learning under imperfect information. *Games and Economic Behavior 47*(1), 72–86.

Chevalier, J. A. and D. Mayzlin (2006, August). The effect of word of mouth on sales: Online book reviews. *Journal of Marketing Research*.

Fleder, D. M. and K. Hosanagar (2007). Blockbuster Culture's Next Rise or Fall: The Impact of Recommender Systems on Sales Diversity. *SSRN eLibrary*.

Hausman, J. A., B. H. Hall, and Z. Griliches (1984). Econometric models for count data with an application to the patents-r&d relationship. *Econometrica 52*, 909–938.

Lohse, G. (1997). Consumer eye movement patterns on yellow pages advertising.

Matthew J. Salganik, Peter Sheridan Dodds, D. J. W. (2006). Experimental study of inequality and unpredictability in an artificial cultural market. *Science 311*(5762), 854–856.

Oberholzer-Gee, F. and A. Elberse (2007). Superstars and underdogs: An examination of the long tail phenomenon in video sales. Harvard Business School Working Paper Series, No. 07-015.

Rosen, S. (1981). The economics of superstars. *American Economic Review 71*(5), 845–858.

Zhang, J. (2007). The Sound of Silence - Evidence of Observational Learning from the U.S. Kidney Market. Working paper, MIT Sloan School of Management.

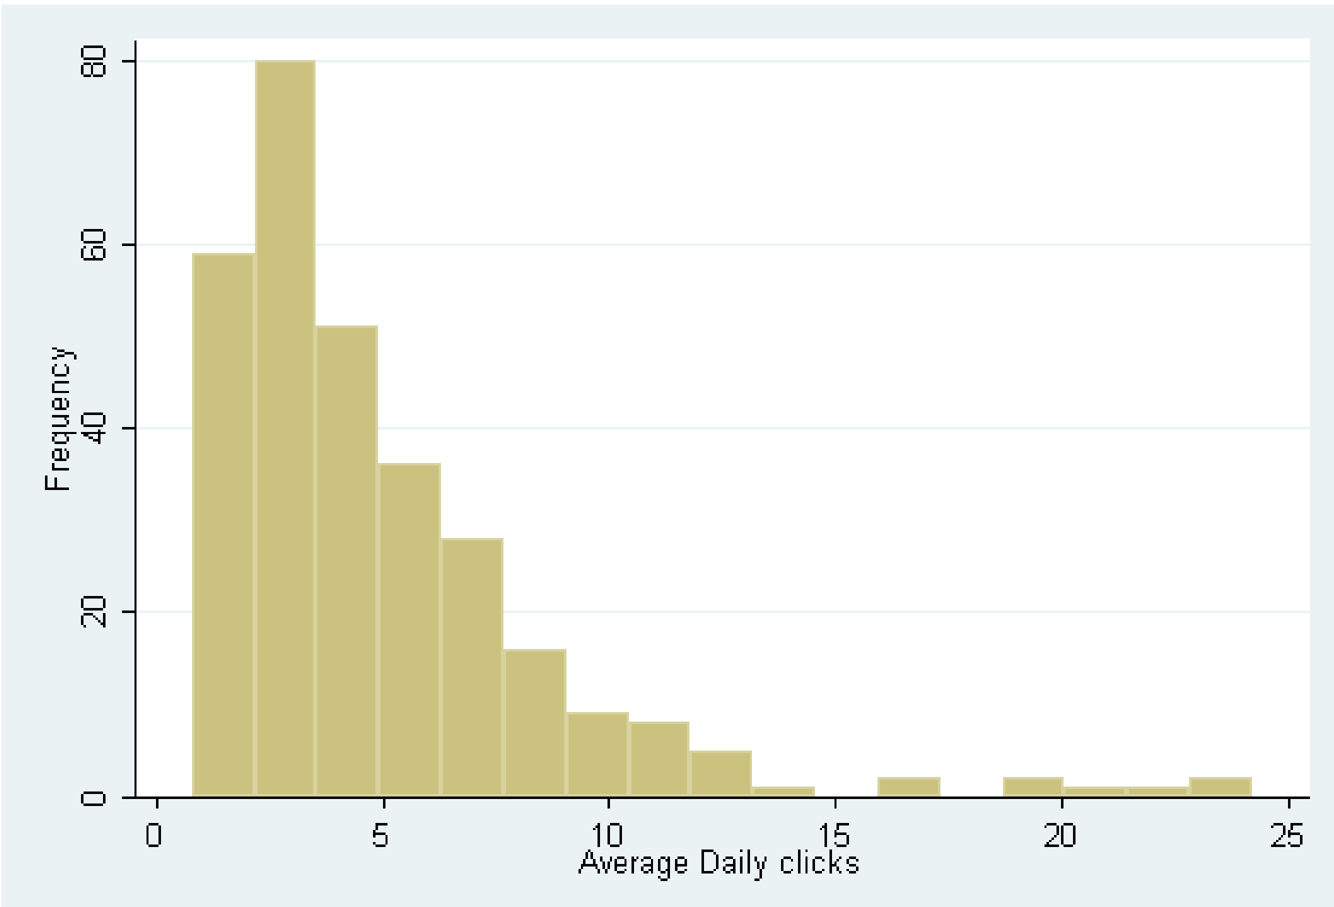

Figure 1: Distribution of Average Daily Clicks Across Vendors

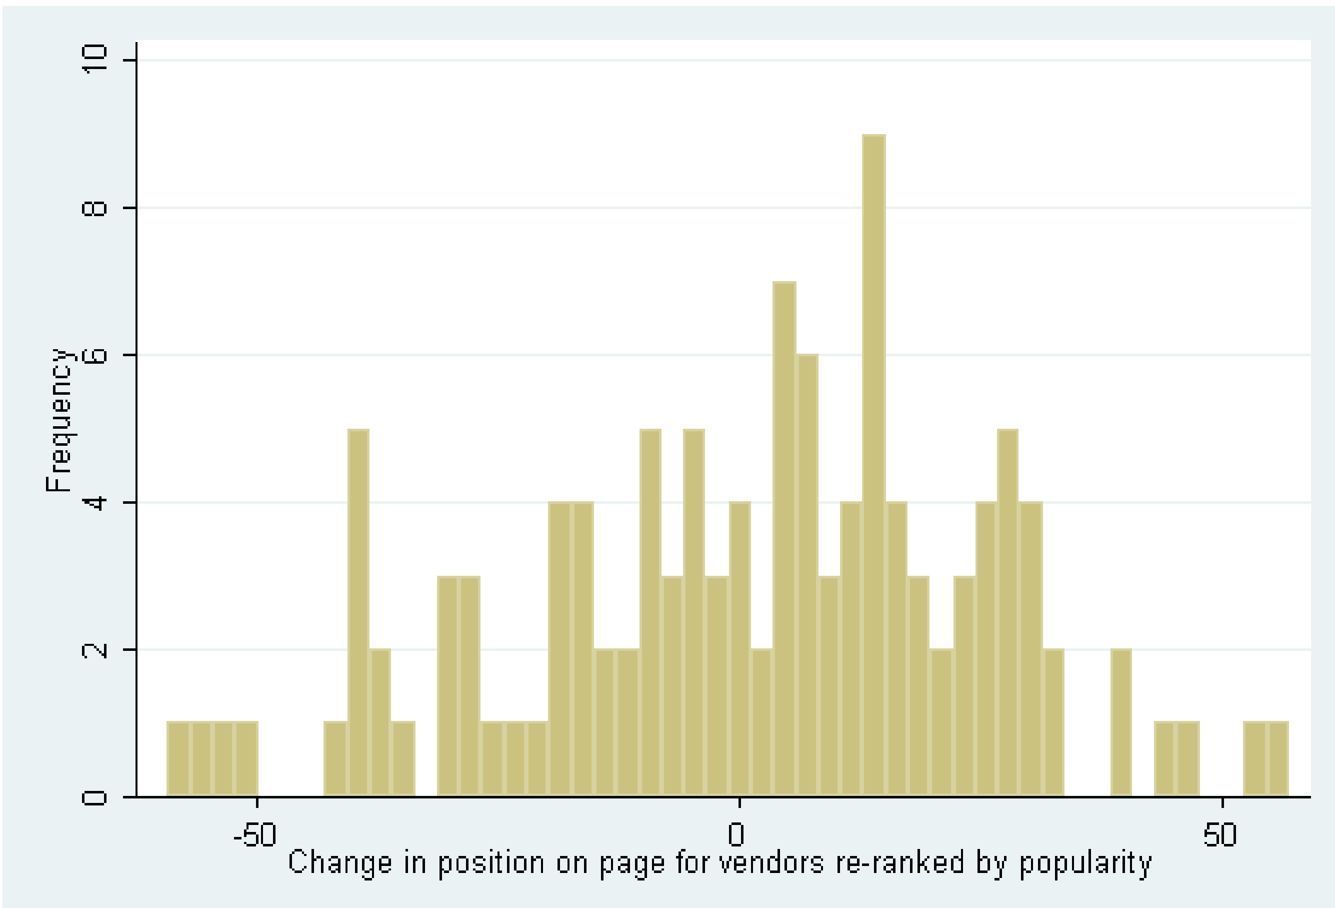

Figure 2: Average shift in vendors position in page after moving from alphabetical ordering to popularity ordering

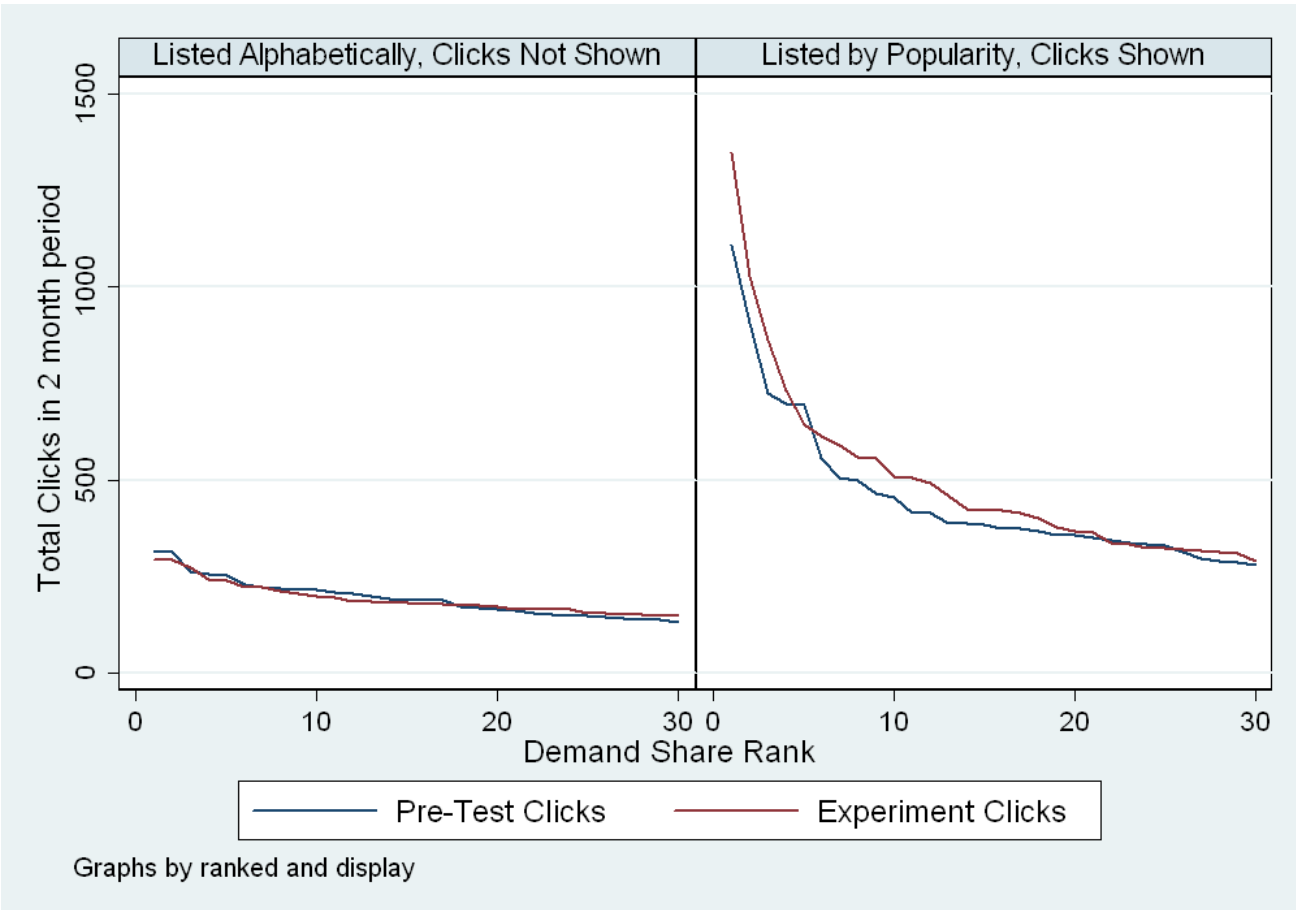

Figure 3: Overall effect of popularity information

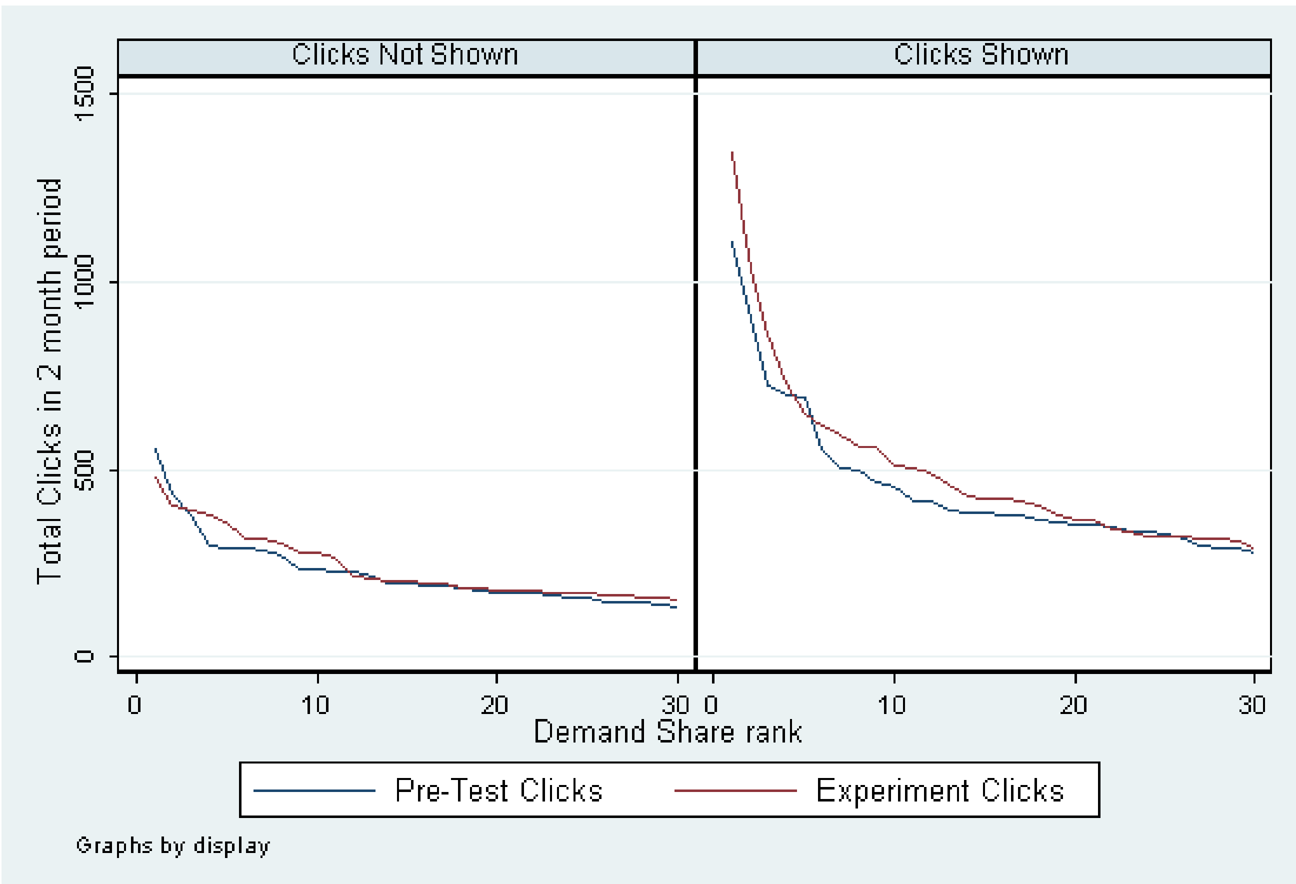

Figure 4: The effect of popularity information after controlling for page location effects

Table 1: Experiment Design

	Popularity Ranking	Clicks Displayed	Mean PreTest Clicks
Bridal Shops	Yes	Yes	310.1
Reception Halls	No	Yes	394.5
Caterers	Yes	No	175.1
Florists	No	No	153.4

Table 2: Differences in Differences: Exploring the Steep Tail

	OLS	**OLS, Robust**	**OLS, Cluster**
ClicksDisplayed*MovesUp	0.714***	0.714***	0.714**
	(0.141)	(0.160)	(0.318)
ClicksDisplayed*MovesDown	-0.086	-0.086	-0.086
	(0.117)	(0.089)	(0.125)
MovesUp	-0.182	-0.182	-0.182
	(0.230)	(0.232)	(0.277)
MovesDown	-0.425*	-0.425**	-0.425**
	(0.217)	(0.207)	(0.197)
Constant	0.870***	0.870***	0.870***
	(0.260)	(0.171)	(0.125)
Observations	13221	13221	13221
Vendor Fixed Effects	Yes	Yes	Yes
Weekly Dummies	Yes	Yes	Yes
Day of Week Dummies	Yes	Yes	Yes
Standard Errors	Standard	Robust	Clustered by Category

Dependent Variable: Number of Daily clicks
Sample: Vendors whose alphabetical rankings were replaced with popularity rankings
* $p<0.10$, ** $p<0.05$, *** $p<0.01$

Table 3: Differences in Differences: Exploring the Steep Tail (Alt.)

	OLS	**OLS, Robust**	**OLS, Cluster**
ClicksDisplayed*UpwardsChangePos	0.020***	0.020***	0.020
	(0.006)	(0.007)	(0.014)
ClicksDisplayed*DownwardsChangePos	-0.005	-0.005	-0.005
	(0.005)	(0.004)	(0.005)
UpwardsChangePos	0.004	0.004	0.004
	(0.006)	(0.006)	(0.011)
DownwardsChangePos	-0.010*	-0.010**	-0.010
	(0.005)	(0.004)	(0.006)
Constant	0.928***	0.928***	0.928***
	(0.273)	(0.178)	(0.139)
Observations	13221	13221	13221
Vendor Fixed Effects	Yes	Yes	Yes
Weekly Dummies	Yes	Yes	Yes
Day of Week Dummies	Yes	Yes	Yes
Standard Errors	Standard	Robust	Clustered by Category

Dependent Variable: Number of Daily clicks
Sample: Vendors whose alphabetical rankings were replaced with popularity rankings
* p<0.10, ** p<0.05, *** p<0.01

Table 4: Monthly Entry and Exit of Wedding Customers

Month	Percentage of Engagements	Percentage of Marriages
January	5 %	6 %
February	8 %	7 %
March	4 %	7 %
April	6 %	8 %
May	6 %	8 %
June	8 %	11 %
July	9 %	10 %
August	9 %	10 %
September	7 %	10 %
October	9 %	9 %
November	9 %	7 %
December	19 %	7 %

Source: Fairchild Bridal Infobank, American Wedding Study, 2002; National Center for Health Statistics, 2004

Table 5: Differences in Differences: Listing Strategies and Total Clicks

	OLS	**OLS,Robust**	**OLS,Cluster**
ListedByPopularityandClicksDisplayed	0.456***	0.456***	0.456***
	(0.091)	(0.080)	(0.078)
ClicksDisplayedOnly	-0.345***	-0.345***	-0.345**
	(0.080)	(0.070)	(0.078)
ListedByPopularityOnly	0.131	0.131*	0.131
	(0.102)	(0.077)	(0.078)
Constant	8.970***	8.970***	8.970***
	(0.266)	(0.337)	(0.362)
Observations	35217	35217	35217
Vendor Fixed Effects	Yes	Yes	Yes
Weekly Dummies	Yes	Yes	Yes
Day of Week Dummies	Yes	Yes	Yes
Standard Errors	Standard	Robust	Cluster Cat

Dependent Variable: Daily Clicks for that vendor
Sample: All four Categories * p<0.10, ** p<0.05, *** p<0.01