How Does Popularity Information Affect Choices? A Field Experiment

May 28, 2009

Catherine Tucker

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

How Does Popularity Information Affect Choices?
A Field Experiment

Abstract

Popularity information is usually thought to reinforce existing sales trends by encouraging customers to flock to mainstream products. We propose an opposite hypothesis: popularity information may benefit niche products disproportionately, because the same level of popularity implies higher quality for a niche product than for a mainstream product. We examine this hypothesis empirically using field experiment data from a web site that lists wedding service vendors. Consistent with our hypothesis, we find that popular niche vendors receive more visits than popular mainstream vendors, across several definitions of niche.

Keywords: Observational Learning, Niche Marketing, Long Tail, Internet Marketing

1 Introduction

Imagine an MBA student who wants to choose which class to attend. She sees that 90 students are enrolled in "Branding," and 90 are enrolled in "Applied Stochastic Discrete Choice Models." How might this information influence her decision?

Previous research on popularity predicts that the information about equal past enrollment will affect future enrollment across classes equally as well (see for example Salganik et al. (2006); Cai et al. (2009); Chen et al. (2009)). We will argue that this is not always the case. If the student perceives that the stochastic model course is more of a "niche" topic, she may interpret an enrollment of 90 in this course as a stronger signal of course quality than an enrollment of 90 in Branding.

We formalize this notion by distinguishing between two sources of popularity: quality and natural market size. An item may be popular either because quality is perceived to be high or because it caters to a larger market. We use "niche" as the label for products that appeal to smaller target markets and therefore have a lower chance of being chosen when all products offer the same quality. We use a simple analytical model to illustrate that if both a mainstream and a niche product appear equally popular, then popularity information will increase consumers' attraction to the niche product more.

We evaluate this hypothesis using field experiment data from a web site that lists wedding service vendors. This web site experimented with shifting from a traditional "yellow pages" style of alphabetical listing where no popularity information is provided, to a more contemporary "best-seller list" style where a vendor's previous number of clicks is displayed prominently and listings are ranked by the number of clicks that vendor has received. We classify vendors as "mainstream" or "niche" along three different dimensions: (1) scope of services as implied by the name, where "mainstream" corresponds to names that suggest generic services (for example, "Elegant Bouquets") and "niche" corresponds to names that suggest specific services ("Wild Flower Bouquets");

(2) familiarity of the words used in the name, where "mainstream" corresponds to common words ("Fine Food Catering") and "niche" corresponds to unusual words ("Etruscan Catering"), as defined by data on word usage frequency in English (Pastizzo and Carbone (2007)); (3) geographic location, where "mainstream" corresponds to more densely populated areas (Boston, MA) and "niche" corresponds to more thinly populated areas (Plymouth, MA). Across all three taxonomies, we find that if customers can easily access popularity information, then popular niche vendors receive more visits than popular mainstream vendors. We verify the robustness of these results using alternative specifications.

These results are important because it is becoming common for businesses to publicize popularity information online, in part due to the lower costs of information display produced by Internet automation (Shapiro and Varian, 1998).[1] Our findings suggest that vendors of popular niche products benefit from being listed on web sites that make popularity information highly salient. The findings also suggest ways for Internet portals, category managers, and multi-product firms to redirect sales. Highlighting the popularity of well-received niche products can boost their sales disproportionately.

This paper draws on the literature of observational learning, but yields a new set of predictions. Classic analytical models of observational learning focus on how customers infer product quality from peer choices (Banerjee (1992); Bikhchandani et al. (1992)). Empirical studies in this direction have also emphasized evidence of quality inference, either in the lab (see Anderson and Holt (1997); Boğaçhan Çelen and Shachar Kariv (2004)) or in the field (see Cai et al. (2009); Chen et al. (2009); Zhang (2009)). All these studies make strong winner-takes-all conclusions, where popularity information benefits high-volume items. By introducing underlying market size into the inference process, we find that higher-volume products do not necessarily fare better. Indeed, popularity does not benefit a product if its high volume is driven by its naturally wide appeal to the mainstream market; on the other hand, even moderate sales can signal high quality if the product only targets a narrow segment of consumers.

This study also contributes to the understanding of the "long tail" concept in e-commerce, which refers to the success of a long array of low-volume niche products online (e.g, Brynjolfsson et al. (2003); Anderson (2006); Brynjolfsson et al. (2007); Oberholzer-Gee and Elberse (2007)). One leading explanation of the long tail is awareness and access, where the Internet lowers customers' search costs and helps them find otherwise obscure items. Our results suggest that the increasing availability of popularity information on the Internet might have further promoted high-quality niche products and therefore increased the profitability of selling products composing the long tail.

The paper is organized as follows. Section 2 develops an analytical model to illustrate why popularity information may affect the choices of mainstream and niche products differently. We derive our central hypothesis using this illustration. Section 3 discusses the design and implementation of a field experiment that aims to evaluate the hopothesis in a real-world setting. Section 4 presents the field experiment data and estimation results. Section 5 concludes the paper and discusses directions for future research.

2 The Hypothesis and A Theoretical Illustration

In this section we use a simple model to illustrate our central hypothesis that niche products benefit more than mainstream products from the same level of received popularity. The model is based on an observational learning mechanism, whereby consumers infer product quality by observing other consumers' product choices.

Products are both horizontally and vertically differentiated, where horizontal product attributes, such as taste-related features, are observed by all customers but vertical quality is unobservable. Taking MBA classes as an example, one horizontal attribute is the topic (Branding vs. Stochastic Models), and one vertical attribute is the quality of teaching. We label a product that targets a small market as "niche", and label a product that targets a large market as "mainstream". Popularity information is information on the relative frequency with which the product is chosen by a set of customers. Popularity can be driven by both quality and match, and a niche product can be popular if its quality is believed to be high. Each customer possesses private information about quality, and her product choice reflects that information.[2] Therefore, product popularity information can be used by subsequent customers to update their knowledge of quality. Crucially, however, each product's popularity is interpreted relative to customers' expectations about the product's natural market size. Therefore, niche products may benefit more from popularity information than mainstream products do, conditional on both achieving the same level of popularity.

2.1 The Setup

Suppose there are two vendors within the same category, each carrying one product. Customers are heterogeneous in their product tastes and are divided into two types with share θ and $1-\theta$ respectively. Assume $1/2 < \theta < 1$ such that one vendor carries a mainstream (denoted as m) product and the other vendor supplies a niche (denoted as n) product. A customer derives match utility $t \geq 0$ by choosing the vendor that matches her taste and 0 otherwise, where t measures the degree of taste heterogeneity. Each customer knows her own taste but does not observe other customers' tastes. The values of θ and t are common knowledge.

The quality of the two products, denoted as v_m and v_n respectively, can be either 0 or 1 with equal prior probability. Customers are uncertain about quality. However, each customer receives a private quality signal which can be either high (H) or low (L). We assume these private signals are identically and, *conditional on the true quality*, independently distributed. Suppose the conditional signal probabilities are $p(H|v_j = 1) = p(L|v_j = 0) = q$, $j \in \{m, n\}$, where $1/2 < q < 1$ so that private signals are informative yet imperfect.

Each customer incurs an exogenous "search cost" of c when visiting a vendor. In the field experiment context of this study, c can be a web viewer's costs of clicking on each vendor. Let $I(\cdot)$ be an indicator variable which equals 1 if the statement inside holds true and 0 otherwise. Let U_{ij}

denote the net utility enjoyed by a customer of taste $i \in \{m, n\}$ when visiting vendor j:

$$U_{ij} = v_j + t \cdot I(i = j) - c. \quad (1)$$

Customers are allowed to visit multiple vendors. This assumption is consistent with the settings in our experiment. Nevertheless, the intuition underlying our hypothesis remains valid when customers are restricted to visiting a single vendor. A customer of type i will visit vendor j if and only if $E(U_{ij}) \geq 0$, where $E(U_{ij}) = 1 \cdot p(v_j = 1) + t \cdot I(i = j) - c$.

2.2 Choices without Popularity Information

In the absence of popularity information, each customer infers quality using her private signal. By Bayes' rule, the posterior belief about v_j after observing an H signal on product j is:

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

Therefore, the expected quality of product j upon receiving an H signal is $E(v_j|H) = q$. Similarly, the expected quality upon receiving an L signal is $E(v_j|L) = 1 - q$. It follows from Equation (1) that the expected utility a type i customer receives from visiting vendor j is $E(U_{ij}|H) = q+t\cdot I(i = j) - c$ upon an H signal, and $E(U_{ij}|L) = 1 - q + t \cdot I(i = j) - c$ upon an L signal.

The Appendix contains a full presentation of the resulting vendor choices without popularity information. In summary, such choices are jointly determined by private quality signals and taste match when $c \in [\underline{c}, \min(c_S, c_M)]$ or $c \in [\max(c_S, c_M), \bar{c}]$, where $\underline{c} = 1-q$, $c_S = q$, $c_M = 1-q+t$, and $\bar{c} = q+t$. For other values of c, choices are determined by private signals alone, or taste match alone, or neither. The rest of the illustration will focus on the more interesting case where choices are jointly shaped by quality signals and tastes.[3]

2.3 Choices with Popularity Information

To illustrate the impact of popularity information, we consider a two-period model. In the first period, one customer makes her choice independently as modeled in the previous section. In the second period, the first customer's choice is observed by another customer. We explore how the choice of the second customer is influenced by the information of her predecessor's decisions.

When search costs are low ($c \in [\underline{c}, \min(c_S, c_M)]$), the first customer will always visit a vendor upon receiving an H signal, but will only visit a matching vendor upon receiving an L signal (see the Appendix). Match is less likely if the vendor is of the niche type. Therefore, from subsequent customers' perspective, the first customer's visit to a niche vendor is more indicative of an H signal, and therefore implies higher quality.

Formally, if v_m equals 1, the probability that the first customer visits vendor m is $p(\textit{visit}|v_m = 1) = \theta \cdot p(\textit{visit}|v_m = 1, \textit{match}) + (1-\theta) \cdot p(\textit{visit}|v_m = 1, \textit{mismatch}) = \theta \cdot 1 + (1-\theta) \cdot p(H|v_m = 1)$ $= \theta + (1 - \theta)q$. Similarly, $p(\textit{visit}|v_m = 0) = \theta \cdot p(\textit{visit}|v_m = 0, \textit{match}) + (1 - \theta) \cdot p(\textit{visit}|v_m = 0, \textit{mismatch}) = \theta \cdot 1 + (1-\theta) \cdot p(H|v_m = 0) = \theta + (1-\theta)(1-q)$. By Bayes' rule, after observing the first customer's visit to vendor m and before receiving her own signal, the second customer's updated belief that v_m equals 1 is given by $p(v_m = 1|\textit{visit}) = \frac{p(\textit{visit}|v_m=1)p(v_m=1)}{p(\textit{visit}|v_m=1)p(v_m=1)+p(\textit{visit}|v_m=0)p(v_m=0)}$. We can similarly derive the second customer's expected quality of either vendor for either previous choice:

$$E(v_m|\textit{visit}) = \frac{\theta + (1-\theta)q}{1+\theta}, \quad E(v_n|\textit{visit}) = \frac{(1-\theta) + \theta q}{2-\theta}, \quad E(v_m|\textit{no visit}) = E(v_n|\textit{no visit}) = 1 - q,$$

It can be verified that $E(v_m|\textit{visit}) < E(v_n|\textit{visit})$. In other words, a visit implies higher quality for the niche than for the mainstream due to the mainstream's higher chance of match.

When search costs are high ($c \in [\max(c_S, c_M), \bar{c}]$), the first customer will visit a vendor unless it is a mismatch and the signal is L. It can be similarly shown that $E(v_m|\textit{no visit}) < E(v_n|\textit{no visit})$. That is, a decision not to visit a mainstream vendor carries more negative quality implication than a decision not to visit a niche vendor, due to the lower chance of match for niche products.

In summary, quality inferences from observations of choices are asymmetric between the mainstream and the niche. The apparent disadvantage of the niche in matching customer tastes becomes its advantage in quality inferences: people partially attribute the popularity of the mainstream to its wide appeal, but "forgive" the unpopular niche given its naturally narrow reach. We state this intuition with the following hypothesis.

Hypothesis. *When customer choices are jointly determined by quality and tastes, the same level of popularity benefits niche products more than mainstream products.*

Note that the hypothesis is a "conditional" statement. Conditional on achieving the same level of popularity, niche products benefit more from popularity information than mainstream products. Niche products are less likely to be popular. Therefore, whether popularity information benefits niche products *ex ante* is not clear. However, our focus is on empirically understanding whether customers do actively use a product's mainstream/niche status to moderate the amount of quality inference they draw from the product's popularity. Similarly, while the release of popularity information can signal product quality, we do not investigate this question in this paper.

Next we empirically evaluate our hypothesis using data from a controlled field experiment. The field experiment approach allows us to observe customer choices conditional on a given level of popularity. It also ensures that the provision of popularity information is an exogenous experimental manipulation rather than an endogenous firm decision. See Anderson and Simester (2004) and Lim et al. (2009) for more discussions of advantages of field experiments, Charness et al. (2007) for a discussion of Internet experiments, and Greenstein (2007) for a discussion of how such experiments have been crucial for firms online.

3 Field Experiment

3.1 Experimental Setting

We use data from an Internet-based field experiment to evaluate our main hypothesis. The web site that conducted the field experiment tried out ways to update their alphabetical "yellow pages" listing style to a contemporary "bestseller list" format which presents popularity information saliently. The web site provided wedding service vendor listings for a New England state. The number of marriages in the geographic area that the web site covered is in line with the national average.[4]

Theoretically, the wedding industry is attractive to study because customers in this industry generally have little prior consumption experience. Even if an individual organizes successive weddings, they prefer to select different vendors in order to differentiate the current wedding from its predecessor. Consequently, customers are likely to have imperfect information about vendors. At the same time, brides may have private quality signals from other weddings they previously attended, from various referral sites ((Chen et al., 2002)), or from third-party reviews ((Chen and Xie, 2005)). As a result of quality uncertainty and the existence of private signals, observational learning is likely to influence brides' decisions. This is also an industry in which customers take vendor selection seriously. On average, 2.3 million weddings take place in the U.S. 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 are spent planning.[5]

We are interested in how popularity information affects customers' decisions to click on the URL of a listed vendor on this web site.[6] Popularity information may attract clicks from 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 web site. A sizable proportion of visitors go to the list-of-vendors page without eventually clicking on any vendor's link. This suggests that for brides the vendor visit decision is not trivial or automatic.

The web site provides minimal information about vendors on the list-of-vendors page. It lists only the vendors' name, location and telephone number. (See the Appendix for a mockup of the Webpage.) We exploit the information on name and location to develop three complementary definitions of which vendors brides may view as niche.

The first definition of niche is based on the vendor's service scope as suggested by its name. Some names may suggest that the vendor provides general services, while other names may suggest the provision of specific services. We define as niche a vendor whose name suggests specific scope of services. For example, the name "Elegant Bouquets" suggests a wide selection of bouquets, while the name "Wild Flower Bouquets" suggests a limited selection of bouquets and may be considered as niche by brides. We conducted a pre-test where we asked independent evaluators to classify the vendors as mainstream or niche based on vendor name. Each vendor's classification follows the majority vote.

The second definition of niche is based on the rarity of the words in the vendor name. The idea is that a vendor with an unusual word in their name (such as "Etruscan Catering") might appear to brides to serve a smaller market than a vendor with a more generic name (such as "Fine Food Catering"). Kucera and Francis (1967) demonstrate that word usage frequency is highly predictive of word familiarity. Therefore, to measure word rarity we integrated our data with Pastizzo and Carbone (2007)'s data on usage frequency of 1.6 million words in the English language. We define a niche vendor as one where on average each of the words in the name is used fewer than 50 times (excluding prepositions and definite articles).[7] The cross-correlation table in the Appendix shows that our niche name definition and niche word definition are highly correlated (0.93). The major differences between the two definitions come from vendors whose name contains words like "sumptuous" which, while uncommon in English, are widely used in the wedding industry and imply no particular limitation in scope of vendor services.

The third type of niche categorization is based on the remoteness of the vendor's location. Using location to define niche status resembles spatial models of horizontal differentiation, where customers incur "transportation costs" by choosing products away from their location on the hotelling line. We define mainstream vendors as those located in the only metropolitan area in the state. We define niche vendors as those located outside the metropolitan area. This definition of niche is uncorrelated to both niche name (0.03) and niche word (0.01). This means that the location definition of niche gives us an orthogonal way of defining niche, increasing our confidence in the findings.

3.2 Experimental Design and Procedures

The web site measured the popularity of a vendor by the number of clicks that vendor's link had received.[8] The site explored three different ways of presenting such popularity information: displaying the number of clicks, ranking vendors in descending order of click volume, and both. This, together with the baseline condition where clicks were not displayed and vendors not ranked by popularity, generated a two-by-two experimental design. The web site conducted this experiment using four out of a total of 19 wedding service categories. These four categories were randomly allocated into the four conditions. We witnessed and verified the randomness of this allocation.

The categories with the highest traffic were selected for the experiment: Florists, Reception Halls, Caterers, and Bridal Shops. Table 1 summarizes the assignment and the pre-experiment traffic level of the four categories. 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.

[Table 1 about here.]

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 web site did not disclose to visitors any information about the start date for this stock of clicks. This lack of disclosure is consistent with industry norms, and prevents customers from being confused by additional cues such as seasonality. The number of clicks was displayed as 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. Except for the display of click information and ordering of vendors, there was no difference in the webpage format across conditions. Every three days we ran a screen-scraping program to verify our data and to ensure that there were no glitches in the experiments.

The experiment could be contaminated if subjects visited categories sequentially. For example, brides could first visit the bridal attire listings and then visit listings for caterers but at that stage guess that these listings were ordered by popularity. Such behavior would lead us to underestimate the effect of popularity information. Aggregate-level web site statistics suggest, however, that most visitors to a list-of-vendors page arrived there directly from search engines rather than navigating there from within the web site. This keeps the field experiment close to a between-subjects design.

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 from the data. In this dataset, each observation is a time-stamp for when a link received a click, alongside the vendor details 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 total clicks across all 19 categories. The four categories in our experiment accounted for 515,121 of these clicks. There was 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. 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. Niche vendors on average received 0.3 fewer clicks per day than mainstream vendors (significant at the 1 percent level). Table 2 provides summary statistics. The Appendix contains the cross-correlation table, and the figure which presents the distribution of average daily clicks across vendors.

[Table 2 about here.]

3.3 Data Processing

There were several challenges in processing the data. The first challenge came from unintentional clicks due to, for example, slow web site response time. Since privacy rules prevented us from 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 on when 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 the existence of a small amount of vendor entry and exit from this web site during the period we study. In the reception hall category, there was one change: during the second month of the study, a reception hall with a name beginning with "O" exited while another reception hall beginning with the letter "T" entered. This shifted the position of all reception halls with first letters "P" to "S" up one place for the second month of the experiment. Another similar change happened in the florist category. For the empirical analysis we tried both incorporating and excluding these vendors. The results were almost identical. Therefore, we will present results for a balanced panel where we exclude the few vendors that exited and entered the web site.

4 Empirical Analysis

4.1 Main Results

We want to find out how clicks are affected by popularity information, and how the effect of popularity information is moderated by niche status.

Equation (2) summarizes our specification, where α and β are vectors of parameters to be estimated:

$$
\begin{aligned}
& clicks_{jt} = \alpha_j + \beta_0\, X_{jt} + \beta_1\, PagePos_{jt} \\
& +\beta_2\, RankedDisplayed_{jt} + \beta_3\, Ranked_{jt} + \beta_4\, Displayed_{jt} \\
& +\beta_5\, Niche_j * RankedDisplayed_{jt} + \beta_6\, Niche_j * Ranked_{jt} + \beta_7\, Niche_j * Displayed_{jt} \\
& +\beta_8\, RankedDisplayed_{jt} * Top20pc_j + \beta_9\, Ranked_{jt} * Top20pc_j + \beta_{10}\, Displayed_{jt} * Top20pc_j \\
& +\beta_{11}\, Niche_j * RankedDisplayed_{jt} * Top20pc_j + \beta_{12}\, Niche_j * Ranked_{jt} * Top20pc_j \\
& +\beta_{13}\, Niche_j * Displayed_{jt} * Top20pc_j + \epsilon
\end{aligned}
\tag{2}
$$

The dependent variable ($clicks_{jt}$) is the number of clicks that vendor j receives during day t. We present results from a linear specification, but the results are similar if we use a poisson regression.

On the right-hand side, we include vendor-specific fixed effects α_j for each vendor j to control for static differences in base demand across vendors. A bride's propensity to make vendor selections may change over time. (See the Appendix for a review of seasonality in the wedding industry.) We capture this time trend in interest for vendors by a vector X_{jt} of weekly dummies and day-of-week dummies. X_{jt} also contains a dummy for whether this is the treatment period or not, and an interaction between this dummy and whether vendor j is a niche vendor.

To capture the level effect of popularity information for both niche and mainstream vendors, we include a dummy $Displayed_{jt}$ that is set to one for vendors whose click information was displayed without re-ordering. We also include the dummy $Ranked_{jt}$ that is set to one for vendors who were re-ordered by their popularity but with no click information displayed. Last, we include a dummy $RankedDisplayed_{jt}$ that is set to one for vendors who were subject to a bestseller format where click information was displayed *and* vendors' page position was re-ordered according to their popularity. This "bestseller" format makes popularity information available and easy to process, and is thus expected to have the largest impact on subsequently choices.

Our major variables of interest are three-way interaction terms between a vendor's niche status, its treatment condition, and whether it is in the top 20 percent of vendors (as measured by clicks) in the pre-test period. The interaction $Niche_j * RankedDisplayed_{jt} * Top20pc_j$ is set to one if vendor j is niche by the pertinent definition of niche, if the vendor is subject to the bestseller format at time t, and if vendor j is in the top 20 percent of vendors in its category prior to the experiment. This three-way interaction measures whether the niche status moderates the impact of popularity information display formats on popular vendors. This allows us to examine the central hypothesis developed in section 2. We also include all potential lower order interactions between the indicator variable for niche status ($Niche_j$), the popularity threshold $Top20pc_j$, and the three experimental conditions $Displayed_{jt}$, $Ranked_{jt}$ and $RankedDisplayed_{jt}$. Because no vendor's niche status changed over time, we are not able to estimate $Niche_j$ separately from vendor-level fixed effects. Similarly, we cannot estimate $Top20pc_j$ separately from vendor fixed effects. We can only estimate time-varying interactions which can be identified from vendor fixed effects.

The results reported in Table 3 are based on a threshold of 20 percent. We checked the robustness of our results to other choices of threshold as reported in section 4.2.[9]

Last, we include the variable $PagePos_{jt}$ to pick up the "web site real estate" effect for vendor j's average page position on day t. This mere page location effect could occur either because customers incur high search costs from scrolling, or because customers' eyes are drawn to the top listings, as suggested by eye-tracking studies. By including the reception halls and florists categories which were not re-ordered by popularity, we are able to separately identify the effect of page position and the effect of popularity.

Our identifying assumption for the time trend is that all categories would have had similar time trends in clicks had it not been for the experimental intervention. Our control condition (Florists) establishes a weekly and daily time trend for visits to the site. This approach could be problematic if we were studying an apparel retailer and we were trying to use interest in, say, sweaters as a control for the interest in bathing suits. However, in the wedding industry different categories of services, such as catering and florists, are complementary components of the same ultimate wedding, so interest in one category is likely to be similar in timing to another category. We tested this by looking at time trends in aggregate clicks in the four categories during the previous year. There was no statistically significant evidence of different time trends. Furthermore, because our main coefficients of interest rest on interactions between niche status, a popularity threshold and the vendor's experimental condition, even if there were category-wide differences in the time trend, so long as these differences were not restricted exclusively to either niche or mainstream products (across a variety of definitions of niche), our relative results hold.

[Table 3 about here.]

Table 3 shows the results for estimating equation (2). $Niche * RankedDisplayed * Top20pc$ is positive and significant for each of our different definitions of niche. When the bestseller format is used, a popular vendor that may be viewed by brides as niche because of the scope of services implied by its name gains 1.5 clicks per day relative to a similarly popular mainstream vendor; a popular vendor that may be viewed as niche by brides because it has unfamiliar words in its name gains an incremental 1.6 clicks per day; a popular vendor that may be viewed by brides as niche because of its location gains an incremental 1.9 clicks per day. $Niche * Displayed * Top20pc$, which measures the effect for popular niche vendors to have click information displayed (with no popularity based re-ordering) is positive, though insignificant for niche locations. $Niche * Ranked * Top20pc$ is insignificant; this interactive term measures the effect of popularity-based ranking on popular niche vendors without telling customers where the rankings coming from. In summary, other display formats, such as mere ranking by popularity or mere display of clicks, have a less significant effect on customer click behavior than the bestseller format. This could be because the bestseller format reduces the cost of processing popularity information, thus encouraging customers to use such information. It could also be that the bestseller format makes popularity information particularly salient, thus increasing both the opportunity and the motivation for information processing, according to the elaboration-likelihood theory.

The insignificant coefficients of $RankedDisplayed * Top20pc$ and $Ranked * Top20pc$ suggest that mainstream vendors do not benefit from popularity information even when they are among the top 20 percent. The negative and significant coefficient on $Displayed * Top20pc$ suggests that there is actually a negative effect for mainstream vendors when the number of clicks are displayed but vendors are not re-ordered by popularity. This may be because a web site which displays clicks but does not rank clicks appears unprofessional or outdated, signaling an unattractive lack of technical know-how which is particularly off-putting to brides seeking mainstream vendors.

These conditional results are consistent with our hypothesis. Customers expect mainstream vendors to be busier than niche vendors. Therefore, when customers see a vendor with an unusual name or located in a rural area receive a large number of clicks, they are more likely to infer high quality than when they see a vendor with a common name or a convenient location receive a similar number of clicks.

Among the remaining variables, $Ranked_{jt}$, $Displayed_{jt}$, and $RankedDisplayed_{jt}$ capture the average aggregate effect for mainstream vendors from receiving different display formats. These main effects are generally insignificant, with the exception of a positive $RankDisplayed_{jt}$ when we define niche by location. This positive effect represents an increase in overall clicks when the bestseller form is used, which may result from the more professional appearance of this contemporary display format. The interactions of the niche dummy with the condition dummies are largely insignificant.

The variable $PagePos_{jt}$ is negative. This suggests that vendors who are listed first on the page (thus having the lowest value of $PagePos$) receive more clicks than vendors displayed lower down the page, independently of popularity. Last, the (unreported) coefficients for the time trend are much as expected. They indicate a decrease in activity over the Labor Day weekend and a high level of web-surfing on Mondays. The time trend for niche vendors is negative and marginally significant compared to mainstream vendors.

4.2 Robustness Checks

To check the robustness of the results to using 20 percent as a popularity threshold, we re-ran our analysis with 10 alternative thresholds. Figure 1 displays the coefficient estimates for $Niche * RankedDisplayed * TopXpc$ across these choices of thresholds, where X stands for the popularity threshold percentile. Figure 1 reveals a constant pattern: the higher the popularity threshold, the larger the comparative advantage popular niche vendors receive. Intuitively, customers perceive it as more improbable that a niche vendor could be in the top 5 percent without being extremely high quality than if it were in the top 50 percent.[10]

[Figure 1 about here.]

When using a panel dataset where there is only one policy experiment, such as in our experiment, the level of significance of the estimates should be interpreted with care (see Bertrand et al. (2004)). 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 two broadly accepted techniques. First, as suggested by Hausman et al. (1984), we used a Poisson quasi-maximum likelihood specification with conditional fixed effects and clustering at the vendor level. The results show a similar magnitude and significance. Second, we examine the effect of using different time frames. In Table 4, we present results from regressions where we combined our daily data into a pre-period and a post-period. The results, in particular the positive and significant coefficient on $NicheRankedDisplayedTop20pc$, support our previous finding that the display of popularity information has the largest benefit for popular niche vendors.

[Table 4 about here.]

One other potential concern is that the results in Table 3 could be subject to serial correlation. For example, if a rival web site started providing listings of, for example, urban bridal shops during our field experiments, which would plausibly reduce the visits to urban bridal vendors on the web site running our experiment and confound our interpretation of $Niche_j Displayed_{jt} Top20pc_j$. Fortunately, during the time period we study, this web site 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.

However, there could be alternative and unobserved sources of a time-varying shock or serial correlation which affects niche bridal shops and no other vendors in the experimental period. For example, there could have been growing word-of-mouth about the bargains to be had at non-urban bridal stores. This would increase both the stock of clicks and the propensity to click through for rural bridal stores. To address this concern, we use a regression discontinuity approach (Black (1999), Hahn et al. (2001), and Busse et al. (September 2006)). The identification logic is that by taking a very short time window we reduce the likelihood that such a time-varying shock (other than the experimental treatment) could explain the results. In the estimates presented in Table 5, we reduce the time window of estimation and evaluate changes in click behavior for the week before the field experiment and the week after the experiment. These estimates are similar to our previous results in Table 3, alleviating concerns about serial correlation.

[Table 5 about here.]

One last concern is that vendors could have reacted strategically to the field experiment. We examined the data for evidence of suspicious clicks (or “click-fraud”) but could find no patterns of successive clicks that suggested vendors were artificially boosting their own popularity rankings. Vendor prices are not displayed and therefore cannot act as an alternative quality signal. Vendors, correspondingly, would have no incentive to strategically change their 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 web site.

5 Conclusion

Previously, researchers have perceived popularity information as a marketing tool that reinforces the *status quo*, and consequently reinforces the dominance of mainstream products. This perception is based on the belief that mainstream products are high volume, and consequently benefit from the bandwagon of sales. We propose an opposing view: that popularity information may actually be of greater benefit to niche products. The fact that niche products are less likely to attract customers, as they appeal to a smaller segment of customers, means that when they are chosen this conveys a greater quality signal to future customers.

We explore this insight using data from a field experiment conducted with a web site that lists wedding service vendors. We find that releasing easy-to-digest popularity information in a bestseller format brings the greatest benefits to popular vendors who appear to serve a minority market, either because of their name or their location. Brides are more likely to infer that a niche vendor is high quality compared to a mainstream vendor of similar popularity, because the niche vendor's natural market is smaller. We check the robustness of our results in a number of ways and find that this main result holds.

These findings contribute to the understanding of how the common practice of displaying popularity information affects customer choices. Our results suggest that the bestseller format benefits popular niche products disproportionately. The findings also help to understand the long tail phenomenon of e-commerce. Contrary to the belief that automated "web 2.0" type tools which highlight previous customer choices promote mainstream products, our results suggest that these display tools can actually strengthen the long tail if it is composed of a sufficient of popular niche products.

There are several potential ways of building on this research. We explore whether *ex post* conditional on achieving the same popularity, a niche product benefits more from the same level of popularity than a mainstream product. Based on this understanding, future research can investigate the *ex ante* effect of releasing popularity information. Also, it would be interesting to model the endogenous release of popularity information as a quality signal. Another possibility is to explore whether popularity information can be similarly moderated by other marketing mix variables. For example, will popular products with higher prices benefit more from the release of popularity information than popular products with lower prices? If indeed customers infer superior quality to justify the high price tag, what would be the firm's optimal pricing strategy? Last, if customers are uncertain about their preferences (see Wernerfelt (1995)), they may infer mainstream status from popularity. It would be interesting to explore the dynamics by which popularity redefines, and is redefined by the perceived boundary between "niche" and "mainstream."

References

Anderson, Chris, *The Long Tail: Why the Future of Business is Selling Less of More*, Hyperion, NY, 2006.

Anderson, Eric and Duncan Simester, "Long Run Effects of Promotion Depth on New Versus Established Customers: Three Field Studies," *Marketing Science*, 2004, *23* (1), 4–20.

Anderson, Lisa R. and Charles A. Holt, "Information Cascades in the Laboratory," *American Economic Review*, 1997, *87* (5), 847–862.

Banerjee, Abhijit V, "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, August 1992, *107* (3), 797–817.

Baye, Michael, John Morgan, J. Rupert J. Gatti, and Paul Kattuman, "Clicks, Discontinuities, and Firm Demand Online," November 2006. Mimeo, Berkeley.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan, "How Much Should We Trust Differences-in-Differences Estimates?," *Quarterly Journal of Economics*, February 2004, *119* (1), 249–275.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "A Theory of Fads, Fashion, Custom, and Cultural Change in Informational Cascades," *Journal of Political Economy*, October 1992, *100* (5), 992–1026.

Black, Sandra E., "Do Better Schools Matter? Parental Valuation of Elementary Education," *Quarterly Journal of Economics*, 1999, *114*, 577799.

Boğaçhan Çelen and Shachar Kariv, "Observational Learning Under Imperfect Information," *Games and Economic Behavior*, 2004, *47* (1), 72–86.

Brynjolfsson, Erik, Yu Hu, and Duncan Simester, "Goodbye Pareto Principle, Hello Long Tail: the Effect of Search Costs on the Concentration of Product Sales," 2007. Mimeo, MIT.

Brynjolfsson, Erik., Yu. Hu, and Michael D. Smith, "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science*, November 2003, *49* (11), 1580–1596.

Busse, Meghan, Jorge Silva-Risso, and Florian Zettelmeyer, "$1,000 Cash Back: The Pass-Through of Auto Manufacturer Promotions," *American Economic Review*, September 2006, *96*, 1253–1270(18).

Cai, Hongbin, Yuyu Chen, and Hanming Fang, "Observational Learning: Evidence from a Randomized Natural Field Experiment," *American Economic Review*, 2009. forthcoming.

Charness, Gary, Ernan Haruvy, and Doron Sonsino, "Social distance and reciprocity: An Internet experiment," *Journal of Economic Behavior and Organizations*, 2007, *63*, 88–103.

Chen, Yubo and Jinhong Xie, "Third-Party Product Review and Firm Marketing Strategy," *Marketing Science*, 2005, *24* (2), 218–240.

_ , Qi Wang, and Jinhong Xie, "Online Social Interactions: A Natural Experiment on Word of Mouth versus Observational Learning," 2009. Working paper, University of Arizona, Binghamton University, University of Florida.

Chen, Yuxin, Ganesh Iyer, and V. Padmanabhan, "Referral Infomediaries," *Marketing Science*, 2002, *21* (4), 412–434.

Greenstein, Shane, "Economic Experiments and Neutrality in Internet Access," NBER Working Papers 13158, National Bureau of Economic Research, Inc June 2007.

Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw, "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 2001, *69*, 201209.

Hausman, Jerry A., Bronwyn H. Hall, and Zvi Griliches, "Econometric Models for Count Data with an Application to the Patents-R&D Relationship," *Econometrica*, 1984, *52*, 909–938.

Kucera, Henry and W. Nelson Francis, *Computational Analysis of Present-day American English*, Brown University Press, 1967.

Lim, Noah, Michael Ahearne, and Sung H. Ham, "Designing Sales Contests: Does the Prize Structure Matter?," *Journal of Marketing Research*, 2009, *46* (3), 356–371.

Lo, Alison KC, John G. Lynch Jr., and Richard Staelin, "How to Attract Customers by Giving Them the Short End of the Stick," *Journal of Marketing Research*, 2007, *44*, 128–141.

Oberholzer-Gee, Felix and Anita Elberse, "Superstars and Underdogs: An Examination of the Long Tail Phenomenon in Video Sales," 2007. HBS Working Paper Series, No. 07-015.

Pastizzo, Matthew J. and Robert F. Carbone, "Spoken word frequency counts based on 1.6 million words in American English," *Behavior Research Methods*, November 2007, *39*, 1025–1028(4).

Salganik, Matthew J., Peter Sheridan Dodds, and Duncan J. Watts, "Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market," *Science*, 2006, *311* (5762), 854–856.

Shapiro, Carl and Hal R. Varian, *Information Rules: A Strategic Guide to the Network Economy*, Harvard Business School Press, November 1998.

Wernerfelt, Birger, "A Rational Reconstruction of the Compromise Effect: Using Market Data to Infer Utilities," *Journal of Consumer Research: An Interdisciplinary Quarterly*, March 1995, *21* (4), 627–33.

Zhang, Juanjuan, "The Sound of Silence: Observational Learning from the U.S. Kidney Market," *Marketing Science*, 2009. forthcoming.

Notes

[1]We conducted a survey that confirmed that 60 of the top 100 U.S. web sites display information about what products past customers have chosen.

[2]In this model, customers draw quality inferences from others' actual product choices. In comparison, Lo et al. (2007) explore quality inferences from what products are offered to other customers. They find in the lab that a customer will infer high quality if a product is associated with a promotion that is a poor fit to herself but a good fit to another group of expert customers.

[3]If first-period choices are solely determined by private signals, subsequent release of popularity information generates the classic bandwagon effect, benefiting the popular products and hurting the unpopular products (e.g., Banerjee (1992); Bikhchandani et al. (1992)). However, a visit (i.e., incremental popularity) gives a mainstream vendor and a niche vendor the same boost in perceived quality, and the lack of visit hurts them to the same extent. If first-period choices are solely determined by taste match, they contain no information on private signals and thus do not affect subsequent customers' choices. Similarly, choices do not affect subsequent customers if they are driven by neither private signals or taste match (for example, when search costs are zero). In this sense, the field experiment can be seen as a high power test of our central hypothesis.

[4]The only observable deviation from national statistics is that weddings in that state cost \$10,000 more than the national average of \$27,000.

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

[6]We do not study how popularity information affects the number of weddings.

[7]We have verified the robustness of the results to different thresholds.

[8] Baye et al. (2006) discuss how the number of clicks on a shopping web site can be thought of as an upper bound on demand.

[9]We also moderated the interaction term using a continuous measure of previous clicks. The results are similar.

[10]In fact, based upon the model in Section 2 it can be shown that $E(v_n|k, K) - E(v_m|k, K)$ increases with k, where $E(v_j|k, K)$ refers to the updated prior quality perception about a vendor of type j (where j refers to either mainstream or niche) when k out of K customers have visited this vendor. The proof is available upon request.

Figure 1: Estimates for the coefficient $Niche_j RankedDisplayed_{jt} TopX0pc_j$ using different thresholds X (in percentage values)

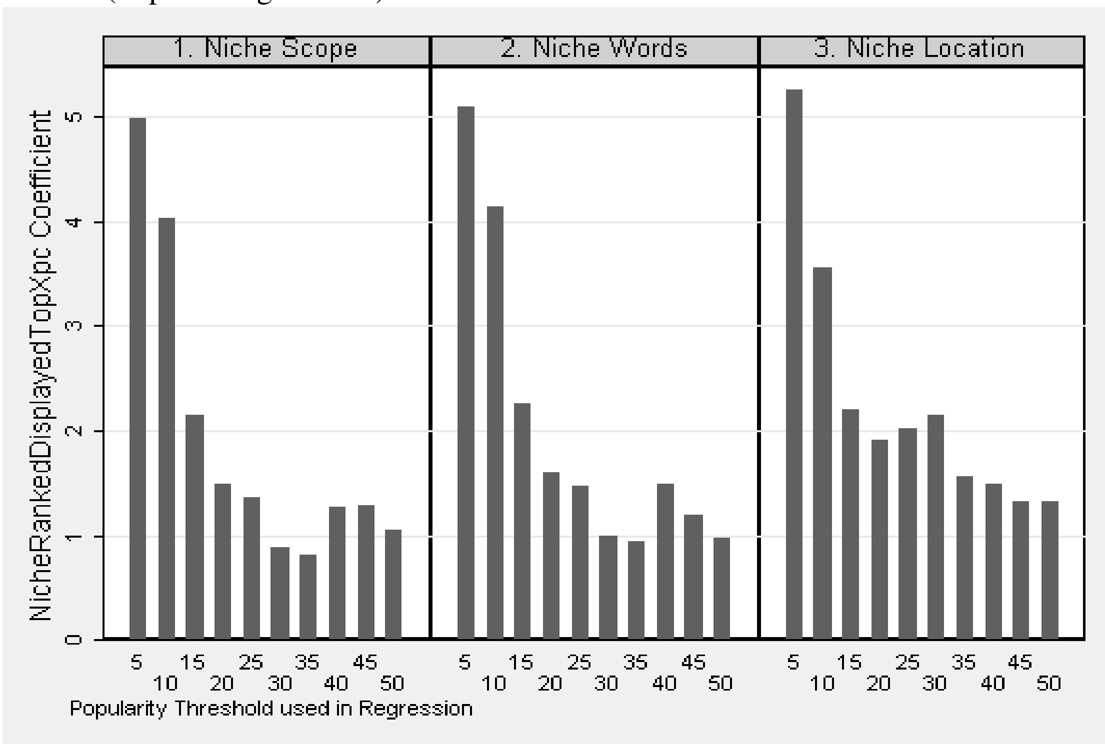

Note: All Coefficients are significant at the $p < 0.05$ level.

Table 1: Experimental Design

	Popularity Ranking	Clicks Displayed	Mean Daily Pretest Clicks	Mean Cumulative Clicks
Florists	No	No	2.6	201.8
Reception Halls	No	Yes	6.6	540.5
Caterers	Yes	No	3.0	243.7
Bridal Shops	Yes	Yes	5.2	441.7

Table 2: Summary Statistics

Variable	Mean	Std. Dev.	Min.	Max.
DailyClicks	4.76	4.664	0	51
Previous Clicks	278.3	293.1	0.00	2763.0
Page Position	52.226	41.262	1	159
NicheScope (1 if service scope is niche)	0.741	0.438	0	1
NicheWords (1 if word is niche)	0.766	0.423	0	1
NicheLocation (1 if location is niche)	0.525	0.499	0	1
Ranked (1 if vendors are ranked & clicks not displayed)	0.073	0.26	0	1
Displayed (1 if clicks are displayed & vendors not ranked)	0.232	0.422	0	1
RankedDisplayed (1 if vendors are ranked & clicks displayed)	0.107	0.309	0	1
Observations		36656		

Table 3: Popular Niche Vendors Benefit More than Popular Mainstream Vendors from Popularity Information

Model	**Niche Scope**	**Niche Word**	**Niche Location**
Niche*RankedDisplayed*Top20pc	1.497**	1.607**	1.903**
	(0.395)	(0.405)	(0.316)
Niche*Ranked*Top20pc	-0.251	-0.181	-0.623
	(0.566)	(0.571)	(0.448)
Niche*Displayed*Top20pc	1.155**	1.167**	0.300
	(0.224)	(0.225)	(0.210)
RankedDisplayed*Top20pc	-0.162	-0.256	-0.204
	(0.355)	(0.366)	(0.243)
Ranked*Top20pc	0.200	0.156	0.184
	(0.548)	(0.552)	(0.238)
Displayed*Top20pc	-1.448**	-1.460**	-0.815**
	(0.185)	(0.187)	(0.151)
Niche*RankedDisplayed	0.336	0.226	-0.141
	(0.217)	(0.235)	(0.134)
Niche*Ranked	0.343***	0.277	0.160
	(0.185)	(0.200)	(0.192)
Niche*Displayed	-0.076	-0.088	-0.104
	(0.097)	(0.101)	(0.096)
RankedDisplayed	-0.010	0.084	0.359**
	(0.216)	(0.235)	(0.117)
Ranked	-0.125	-0.091	0.088
	(0.174)	(0.191)	(0.117)
Displayed	-0.122	-0.110	-0.105
	(0.106)	(0.110)	(0.105)
PagePos	-0.007***	-0.007**	-0.006***
	(0.003)	(0.003)	(0.003)
Observations	36656	36656	36656
R-Squared	0.668	0.668	0.668
Vendor Fixed Effects	Yes	Yes	Yes
Weekly Dummies	Yes	Yes	Yes
Day of Week Dummies	Yes	Yes	Yes

Dependent Variable: Number of clicks
OLS. Robust Standard Errors: * $p<0.1$, ** $p<0.05$, *** $p<0.01$

Table 4: Popular Niche Vendors Benefit More from Popularity Information (Pre-Experiment Period and Post-Experiment Period)

Model	**Niche Scope**	**Niche Word**	**Niche Location**
Niche*RankedDisplayed*Top20pc	83.050**	89.098**	104.355**
	(19.090)	(20.376)	(6.501)
Niche*Ranked*Top20pc	-8.222***	-5.123	-26.760**
	(4.358)	(4.912)	(5.747)
Niche*Displayed*Top20pc	69.688**	70.075**	14.149
	(25.058)	(26.853)	(37.631)
RankedDisplayed*Top20pc	-19.642	-24.814	-21.309**
	(21.817)	(23.291)	(8.148)
Ranked*Top20pc	-14.565**	-16.565**	-11.926
	(4.469)	(4.351)	(7.520)
Displayed*Top20pc	-102.183**	-102.690**	-61.649**
	(19.783)	(18.592)	(18.965)
Niche*RankedDisplayed	19.434	13.367	-7.923
	(19.134)	(20.414)	(6.503)
Niche*Ranked	19.667**	16.871**	7.947**
	(2.812)	(3.380)	(3.740)
Niche*Displayed	-5.532	-5.915	-3.428
	(7.206)	(5.959)	(6.861)
RankedDisplayed	-3.891	1.281	17.292**
	(19.562)	(21.142)	(6.222)
Ranked	-5.668	-4.452	6.871**
	(3.898)	(4.349)	(2.447)
Displayed	-10.756	-10.250	-12.224
	(7.599)	(6.290)	(7.580)
PagePos	-0.864**	-0.883**	-0.833**
	(0.105)	(0.105)	(0.117)
Observations	632	632	632
R-Squared	0.988	0.988	0.988
Vendor Fixed Effects	Yes	Yes	Yes
Weekly Dummies	Yes	Yes	Yes
Day of Week Dummies	Yes	Yes	Yes

Dependent Variable: Number of clicks
OLS. Robust Standard Errors: * p<0.1, ** p<0.05, *** p<0.01
Sample data collapsed into two periods: one before and one after experiment.

Table 5: Popular Niche Vendors Benefit More than Popular Mainstream Vendors from Popularity Information (for 7 Days Before and After the Experiment)

Model	**Niche Scope**	**Niche Word**	**Niche Location**
Niche*RankedDisplayed*Top20pc	3.805**	3.553***	2.907***
	(1.913)	(1.908)	(1.525)
Niche*Ranked*Top20pc	1.604**	1.683**	-0.195
	(0.578)	(0.628)	(1.080)
Niche*Displayed*Top20pc	2.859**	2.948**	-0.030
	(1.371)	(1.373)	(1.264)
RankedDisplayed8Top20pc	-2.166	-1.936	-1.022
	(1.798)	(1.792)	(1.300)
Ranked*Top20pc	-1.230**	-1.293**	0.292
	(0.572)	(0.613)	(0.705)
Displayed*Top20pc	-2.913**	-2.996**	-0.879
	(1.197)	(1.200)	(0.889)
Niche*RankedDisplayed	0.024	0.277	-0.077
	(0.389)	(0.362)	(0.365)
Niche*Ranked	-0.040	-0.120	0.500
	(0.395)	(0.467)	(0.366)
Niche*Displayed	-0.489***	-0.579**	0.351
	(0.266)	(0.276)	(0.272)
RankedDisplayed	0.120	-0.110	0.179
	(0.375)	(0.347)	(0.299)
Ranked	0.999**	1.064**	0.846**
	(0.411)	(0.480)	(0.232)
Displayed	0.592**	0.676**	0.038
	(0.283)	(0.294)	(0.291)
PagePos	-0.025**	-0.025**	-0.024**
	(0.009)	(0.009)	(0.010)
Observations	4424	4424	4424
R-Squared	0.706	0.706	0.704
Vendor Fixed Effects	Yes	Yes	Yes
Weekly Dummies	Yes	Yes	Yes
Day of Week Dummies	Yes	Yes	Yes

Dependent Variable: Number of clicks
OLS. Robust Standard Errors: * p<0.1, ** p<0.05, *** p<0.01
Sample limited to 1 week before and 1 week after experiment.

A Appendix

A.1 Choices without Popularity Information

If search costs are low enough (i.e., $c < \underline{c} = 1 - q$), a customer will visit both vendors regardless of her tastes and her private signals. Meanwhile, if c is high enough (i.e., $c > \bar{c} = q + t$), a customer will visit neither vendor. In either case, a customer's decision reveals no information about her private signal to subsequent customers. Releasing popularity information therefore would not affect subsequent choices. For the rest of the analysis, we focus on the non-degenerate case where $c \in [\underline{c},\ \bar{c}]$.

A customer will visit a matching vendor despite an L signal if $c \leq c_M = 1 - q + t$, where c_M represents the cost threshold below which match alone guarantees a visit. Similarly, a customer will visit a vendor upon an H signal despite mismatch if $c \leq c_S = q$, where c_S denotes the cost threshold below which an H signal alone guarantees a visit. Figure A-1 summarizes customer choices in the absence of popularity information. When c is sufficiently low (i.e., $c < \min(c_S, c_M)$), a customer visits a vendor if quality signal is high or if the tastes match. On the other hand, when c is sufficiently high (i.e., $c > \max(c_S, c_M)$), a customer visits a vendor if and only if it is the matching type and the signal is H. The sufficient and necessary condition of visit is match when $c_S < c < c_M$, and is an H signal when $c_M < c < c_S$. Note that $c_S < c_M$ if and only if $1 + t > 2q$. The intuition is that match is more likely to determine choices when customer tastes are heterogeneous and private signals are noisy. And quality is more likely to determine choices when customer tastes are homogeneous and private signals are accurate.

In sum, choices are solely determined by match if $c_S < c < c_M$. On the other hand, choices are solely determined by private signals if $c_M < c < c_S$. Finally, choices are jointly determined by private signals and taste match when $c \in [\underline{c}, \min(c_S, c_M)]$ or $c \in [\max(c_S, c_M), \bar{c}]$.

Figure A-1: Choices without Popularity Information

Heterogeneous tastes, noisy quality signals $(1 + t > 2q)$

	High signal	Low signal
Match	1	1
Mismatch	1	0

	High signal	Low signal
Match	1	1
Mismatch	0	0

	High signal	Low signal
Match	1	0
Mismatch	0	0

$\underline{c}$ $\qquad$ c_S $\qquad$ c_M $\qquad$ $\overline{c}$

Homogeneous tastes, accurate quality signals $(1 + t < 2q)$

	High signal	Low signal
Match	1	1
Mismatch	1	0

	High signal	Low signal
Match	1	0
Mismatch	1	0

	High signal	Low signal
Match	1	0
Mismatch	0	0

$\underline{c}$ $\qquad$ c_M $\qquad$ c_S $\qquad$ $\overline{c}$

Note: The figure summarizes a customer's decisions of whether to visit a vendor given her quality signal, taste match, and search costs. 1 represents a visit, and 0 represents no visit.

A.2 Mockups of the Webpage

Due to confidentiality agreements with the web site, we are not permitted to reprint the actual webpages concerned. However, to give a basic idea of what they looked like before and during the experiment, we constructed the two mockup webpages shown in Figure A-2.

A.3 Cross-Correlation Table

Table A-1 presents the cross-correlations among the variables used in the empirical analysis.

Table A-1: Cross-Correlation Table

Variables	Daily Clicks	Niche Location	Niche Words	Niche Name	Displayed	Ranked	Ranked& Displayed	Top20pc	PagePos
Daily Clicks	1.000								
NicheLocation	-0.003	1.000							
NicheWords	-0.016	0.013	1.000						
NicheName	0.004	0.030	0.934	1.000					
Displayed	0.102	0.101	-0.079	-0.087	1.000				
Ranked	-0.139	-0.164	0.028	0.016	-0.154	1.000			
RankedDisplayed	0.016	-0.013	0.097	0.102	-0.190	-0.097	1.000		
Top20pc	0.534	-0.003	-0.020	-0.004	0.481	-0.036	0.187	1.000	
PagePos	0.017	0.152	-0.066	-0.084	0.342	-0.192	-0.145	0.019	1.000

Figure A-2: Mock-Up Webpage: Before and During the Experiment

(Before the experiment) (During the experiment)

A.4 Distribution of Daily Clicks across Vendors

Figure A-3 presents the distribution of average daily clicks across vendors.

A.5 Industry-Level Robustness Checks

One concern with studying the wedding industry is that any experiment could be confounded by seasonal changes in the level of interest in weddings. This is why we use a rich set of controls to capture the time trend. Meanwhile, Table A-2 provides additional assurance that the interest in the wedding industry is more evenly spread across the year than the conventional belief in "summer weddings" would suggest. The largest monthly shock is in December, when 19 percent of engagements take place. By contrast, there is less variation in how many weddings take place each month. June and July, commonly assumed to be the most popular months for weddings, only account on average for 10.5 percent of the interest in wedding vendors.

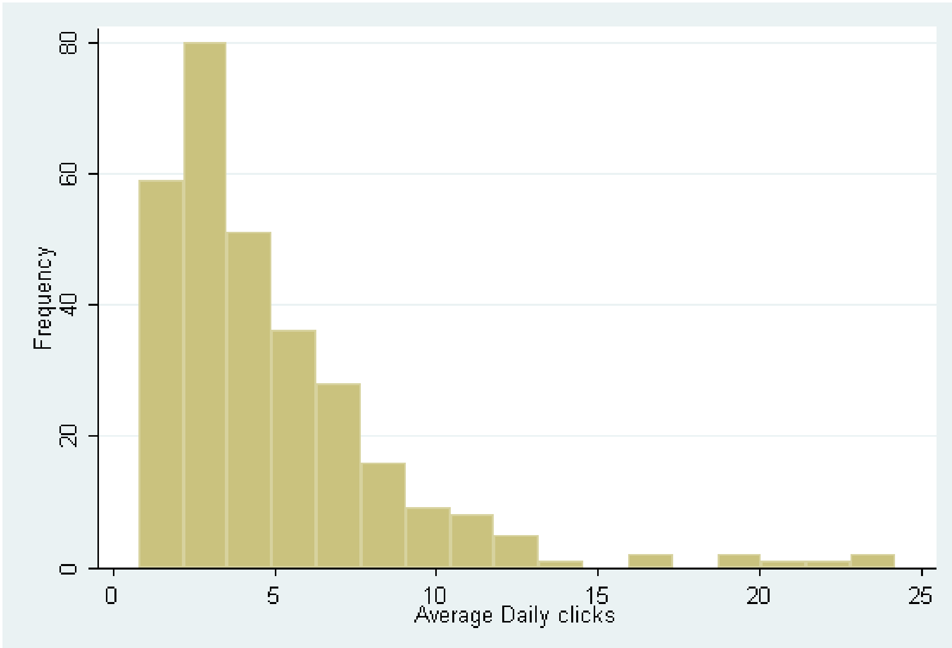

Figure A-3: Distribution of Average Daily Clicks across Vendors

Table A-2: Seasonality in the Wedding Industry

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