

How Does Popularity Information Affect Choices? A Field Experiment

October 23, 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 with broad appeal. We propose an opposite hypothesis: popularity information may benefit niche products with narrow appeal disproportionately, because the same level of popularity implies higher quality for narrow-appeal products than for broad-appeal products. 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 narrow-appeal vendors receive more visits than popular broad-appeal vendors after the introduction of popularity information.

Keywords: Popularity Information, Observational Learning, Internet Marketing, Long Tail.

1 Introduction

Imagine an MBA student who wants to choose which class to attend. She sees that 90 students are enrolled in “Strategy,” and 89 are enrolled in “Applied Stochastic Discrete Choice Models.” How might this information influence her decision?

Previous research predicts that this class enrollment information makes the Strategy class more attractive, as popularity tends to be self-reinforcing (see for example Salganik, Dodds and Watts, 2006; Cai, Chen and Fang, 2009; Chen, Wang and Xie, 2009). We will argue that this is not always the case. If the student perceives that the Stochastic Modeling course covers a topic with naturally narrower appeal, she may interpret an enrollment of 89 in this course as a stronger signal of course quality than an enrollment of 90 in a class with inherently broad appeal such as Strategy.

We formalize this notion by distinguishing between two drivers of popularity: quality and natural appeal. An item may be popular either because quality is perceived to be high or because it caters to a broader range of tastes. We use “narrow-appeal” as a label for products that serve only a small niche of the market and consequently have a lower chance of being chosen when all products offer the same quality. Similarly, we use “broad-appeal” to refer to products that suit the mainstream tastes and therefore enjoy a high chance of being choice among products of the same quality. We use a simple analytical model to illustrate that if both a broad-appeal and a narrow-appeal product appear equally popular, then popularity information will increase consumers’ attraction to the narrow-appeal product more.

We evaluate this hypothesis using a field experiment 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 “bestseller 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 either broad-appeal or narrow-appeal by whether they are located in a town with a large population. This classification is similar to Hotelling models of horizontal differentiation where transportation costs affect market coverage (Hotelling, 1929). We find that if customers can easily access popularity information, then popular narrow-appeal vendors receive more visits than popular broad-appeal vendors. We verify the robustness of these results using a wide range of alternative specifications. In addition, we check for robustness with respect to the definition of appeal by classifying vendors based on whether their names sound unique or familiar (Pastizzo and Carbone, 2007).

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 narrow-appeal, or 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 popular niche products can boost their sales disproportionately, compared to popular mainstream products.

This paper draws on the literature of observational learning and 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, Hirshleifer and Welch, 1992). Empirical studies in this direction have also emphasized evidence of quality inference, either in the lab (see Anderson and Holt, 1997; Celen and Kariv, 2004) or in the field (see Cai, Chen and Fang, 2009; Chen, Wang and Xie, 2009; Zhang, 2009). All these studies

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

make strong winner-takes-all conclusions, where popularity information benefits high-volume items. By introducing natural appeal 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” paradigm in e-commerce, which refers to the success of a long array of low-volume niche products online (e.g, Brynjolfsson, Hu and Smith, 2003; Anderson, 2006; Brynjolfsson, Hu and Simester, 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 high-quality niche titles composing the long tail.

The rest of the paper is organized as follows. Section 2 develops an analytical model to illustrate why popularity information may affect the choices of broad-appeal and narrow-appeal 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 hypothesis in a real-world setting. Section 4 presents the analysis of the field experiment data. Section 5 concludes the paper and discusses directions for future research.

2 Hypothesis and A Theoretical Illustration

In this section we use a simple model to illustrate our central hypothesis that narrow-appeal products benefit more than broad-appeal 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 (Strategy vs. Stochastic Models), and one vertical attribute is the quality of teaching. We label a product that matches the tastes of few consumers as “narrow-appeal”, and label a product that caters to the tastes of the majority consumers as “broad-appeal”. 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 narrow-appeal 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 appeal. Therefore, narrow-appeal products may benefit more from popularity information than broad-appeal 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 broad-appeal (denoted as b) product and the other vendor carries a narrow-appeal (denoted as n) product.

²In this model, customers draw quality inferences from others' actual product choices. In comparison, Lo, Lynch and Staelin (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.

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_b 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 $p(H|v_j = 1) = p(L|v_j = 0) = q$, $j \in \{b, n\}$, 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 $i \in \{b, n\}$ when visiting vendor j :

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

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

³If 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, Hirshleifer and Welch, 1992). However, a visit (i.e., incremental popularity) gives a broad-appeal vendor and a narrow-appeal 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.

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 ($c \in [c, \min(c_S, c_M)]$), the first customer will always visit a vendor when 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 has narrow-appeal. Therefore, from subsequent customers' perspective, the first customer's visit to a narrow-appeal vendor is more indicative of an H signal, and therefore implies higher quality.

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

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

It can be verified that $E(v_b|\text{visit}) < E(v_n|\text{visit})$. In other words, a visit implies higher quality for a product with narrow appeal than for a product with broad appeal due to the latter'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 $E(v_b|no\ visit) < E(v_n|no\ visit)$. That is, a decision not to visit a broad-appeal vendor has a larger negative quality implication than a decision not to visit a narrow-appeal vendor, due to the lower chance of match for narrow-appeal products.

In summary, quality inferences from observations of choices are asymmetric between broad-appeal and narrow-appeal products. The apparent disadvantage of products with only narrow-appeal in matching customer tastes becomes an advantage in quality inferences: people partially attribute the popularity of a product with broad-appeal due to the large number of customers it may attract. 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 narrow-appeal products more than broad-appeal products.*

Note that the hypothesis is a “conditional” statement. Conditional on achieving the same level of popularity, narrow-appeal products benefit more from popularity information than broad-appeal products. Narrow-appeal products are less likely to be popular. Therefore, whether popularity information benefits narrow-appeal products *ex ante* is not clear. However, our focus is on empirically understanding whether customers do actively use product appeal to moderate how much quality they infer 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.

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 ex-

ogenous experimental manipulation rather than an endogenous firm decision. See Lucking-Reiley (1999), Anderson and Simester (2004) and Lim, Ahearne and Ham (2009) for more discussions of advantages of field experiments, Charness, Haruvy and Sonsino (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 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, Iyer and Padmanabhan, 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.

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

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 displays only the vendors' name, location and telephone number. (See the Appendix for a mockup of the Webpage.) However, we will subsequently exploit the information on name and location when defining which vendors have narrow appeal.

Our primary definition of appeal is based upon the size of the population in a vendor's town using 2000 census data. Using location to define appeal resembles spatial models of horizontal differentiation, where customers incur "transportation costs" by choosing products away from their location on the Hotelling line. We define narrow-appearance vendors as those located in towns with a population of less than 50,000.⁷ Narrow-appearance vendors on average received 0.3 fewer clicks per day than broad-appearance vendors (significant at the 1 percent level).

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

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

⁷The results are robust if we use 40,000 or 60,000 as the cut-off.

3.2 Experimental Design and Data

The web site measures the popularity of a vendor by the number of clicks that vendor’s link has received previously.⁸ There are a number of vendors in each category, and the web site consists of 19 categories. The website selected a few frequently visited categories for the field study. Random assignment of experimental conditions occurred at the category level. First, the “Bridal Shops” category received the treatment of interest, where the number of previous clicks was displayed for each vendor, and where the vendors in this category were ranked in descending order of popularity. Second, the “Florists” category served as the control group which maintained the original yellow-page style of display—previous clicks information was not displayed, and vendors were ranked alphabetically. Third, the “Caterers” category served as an additional control, where clicks information was not displayed but vendors were ranked in descending order of popularity. As we shall discuss, the Caterers category helps to disentangle whether the changes in clicks are caused by vendors’ page location or their popularity information.⁹

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”, and updated

⁸Baye, Morgan, Gatti and Kattuman (2006) discuss how the number of clicks on a shopping web site represents an upper bound on demand.

⁹There was initially another category where the web site management attempted to present popularity information but keep the alphabetical ranking of vendors. The implementation experienced unexpected category-specific difficulties. We exclude this category from all analysis to achieve precise interpretation of the estimates. Previous versions of this paper included this category in a more complex 5-way interaction specification and the main results for the effect of popularity information on narrow- versus broad-appeal vendors are qualitatively similar.

instantaneously. In the control conditions, this column was unlabeled and empty (see the Appendix for the mockup page design). 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 the data and to ensure that there were no glitches in the experiments.

Given that there were different formats used in different categories during the experiment, our results could be contaminated if subjects visited categories sequentially. For example, brides could first visit the Bride Shop listings and then visit Caterers listings 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 the list-of-vendors page arrived directly from search engines rather than navigating from within the web site. This feature keeps the field experiment close to a between-subjects design.

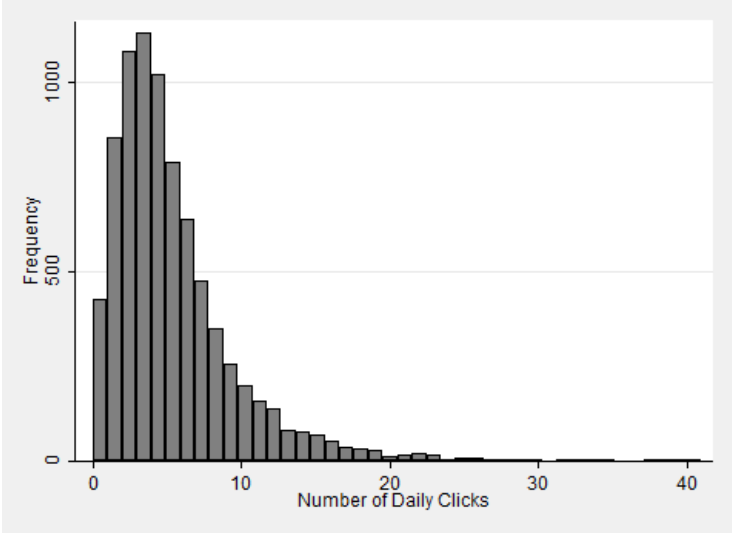
The firm collected data on browsing behavior based on their Apache Web Server logs. To protect the privacy of the users, IP address information was removed from the data. In this dataset, each observation is a time-stamp for when a link received a click, alongside the vendor information and category affiliation. 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).¹⁰

During these four months, there were 860,675 total clicks across all 19 categories. The focal category of interest, Bridal Shops, accounted for 121,682 of these clicks which were

¹⁰One challenge of processing this data 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 the 14,757 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.

spread among 73 vendors over the four months period. 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. Figure 1 displays the frequency distribution of daily clicks in the Bridal Shops category.

Figure 1: Distribution of Daily Clicks in the Bridal Shops Category



Note: The vertical axis measures the number of vendors who receive a certain number of clicks each day, aggregated over all four months of the data.

In addition, the Florists category (which we use as a control) contained 51 vendors and received 44,813 clicks. Similarly, the Caterers category (which we use as a second control) had 66 vendors and received 44,147 clicks. Table 1 provides summary statistics for the Bridal Shops, Caterers, and Florists categories. The average daily clicks is 3.55, while the average number of cumulative previous clicks is 205. In subsequent regressions, we measure previous clicks in thousands for ease of interpretation.

Table 1: Summary Statistics for the Bridal Shops, Caterers, and Florists Categories

	Mean	Std Dev	Min	Max	Observations
Daily Clicks	3.55	3.58	0	41	19372
Previous Clicks (1,000)	0.21	0.21	-0.022	2.35	19372
Initial Letter Position in Alphabet	11.50	5.99	3	23	19372
Narrow Appeal Location	0.35	0.48	0	1	19372
Narrow Appeal Name	0.26	0.44	0	1	19372

Note: Each observation is taken at the vendor/day level.

4 Empirical Analysis

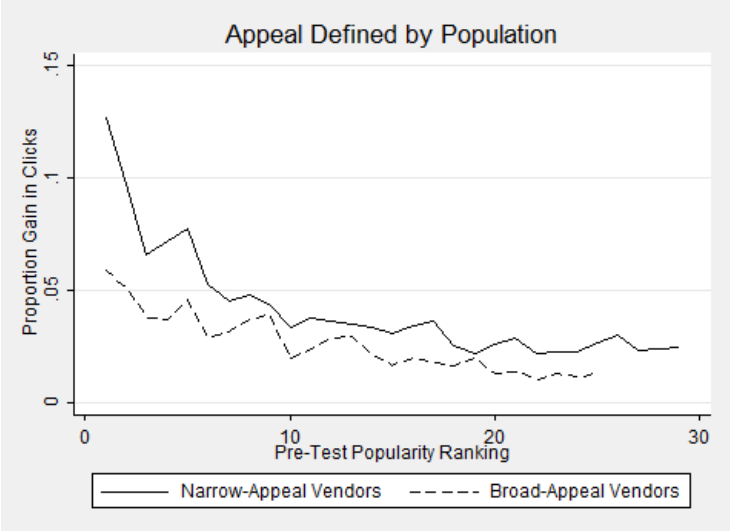
In order to evaluate the effect of popularity information and its comparative impact on broad- and narrow-appeal vendors, we proceed in three stages. We start by presenting graphical evidence that describes the raw impact of popularity information. We then use a differences-in-differences specification, comparing the treatment category (Bridal Shops) with the control categories (Caterers and Florists), and comparing the two months in the experiment with the two months prior to the experiment. Last, we focus on the treatment category using instrumental variables analysis to tackle any endogeneity problem with vendor popularity.

4.1 Initial Graphical Evidence

We first graphically assess the raw impact of popularity information. Figure 2 displays the proportional gain in the number of clicks in the treatment category (Bridal Shops) after popularity information is released. The horizontal axis is the popularity ranking of vendors in the treatment category prior to the experiment. In addition, we compare the relative change in clicks for broad-appeal vendors (vendors located in large towns) and narrow-appeal vendors (vendors located in small towns) separately. There are two patterns to note.

First, releasing popularity information on average increases clicks in the treatment category, especially for vendors who were already popularity prior to the experiment. Second, the proportional gain in clicks is larger for vendors with narrow appeal than vendors with broad appeal, and especially greater for narrow-appeal vendors who were already popular.

Figure 2: The Effect of Popularity Information on the Treatment Category (Bridal Shops) When Appeal is Defined by Vendor Location.



It would be premature to draw causal interpretations from the raw correlations between previous popularity and current popularity (Manski, 1993). For example, there may be unobserved time trends that increase the demand for wedding services during the months of the experiment. (See the Appendix for a review of seasonality in the wedding industry.) To better understand causality, we use two established policy evaluation techniques: the differences-in-differences method, and instrumental variables analysis.

4.2 Differences-in-Differences Specifications

We have data from both the treatment category and the control categories, which allows us to isolate the incremental effect of the treatment. We also have data both before and after

the experiment, which allows us to control for unobservable time trends. In combination, we employ the differences-in-differences method to evaluate the impact of releasing popularity information. Furthermore, in each category and each period there is variation in vendors' locations, from which we can identify the relative effect of popularity information on broad-appeal vendors and narrow-appeal vendors. Since the randomization of the field experiment took place at the category level rather than at the customer level, we employ more complex empirical strategies than usually found in the analysis of field experiments to ensure that we control for unobserved heterogeneity across vendors and categories.

As Figure 1 suggests, in this count data setting, a Poisson specification seems a reasonable approximation for the distribution of clicks.¹¹ In particular, we assume that the number of clicks on day t for vendor j , $Clicks_{jt}$, is drawn from a Poisson distribution where the mean of the distribution is represented by the parameter λ_{jt} :

$$Prob(Clicks_{jt} = c) = \frac{\lambda_{jt}^c e^{-\lambda_{jt}}}{c!} \quad (2)$$

where $c = 0, 1, 2, \dots$ and where the mean parameter λ_{jt} reflects the differences-in-differences specification:

$$\begin{aligned} \ln(\lambda_{jt}) = & \alpha_j + \beta_0 X_{jt} + \beta_1 PagePos_{jt} \\ & + \beta_2 Bridal_j * Test_t * PrevClicks_{jt} * NarrowAppeal_j \\ & + \beta_3 Bridal_j * Test_t * PrevClicks_{jt} + \beta_4 Bridal_j * Test_t * NarrowAppeal_j \\ & + \beta_5 Bridal_j * PrevClicks_{jt} * NarrowAppeal_j + \beta_6 Test_t * PrevClicks_{jt} * NarrowAppeal_j \\ & + \beta_7 Bridal_j * Test_t + \beta_8 Bridal_j * PrevClicks_{jt} + \beta_9 Test_t * PrevClicks_{jt} \\ & + \beta_{10} Test_t * NarrowAppeal_{jt} + \beta_{11} PrevClicks_{jt} * NarrowAppeal_j + \beta_{12} PrevClicks_{jt} + \epsilon_{jt} \end{aligned} \quad (3)$$

¹¹Although care should be taken in interpreting interactions in non-linear models (Ai and Norton, 2003), more recent research has suggested that for measuring a treatment effect it is still appropriate to focus on the size and direction of the interaction (Puhani, 2008). Nevertheless, we have estimated all our models using a linear OLS specification, with very similar results.

On the right-hand side of the above specification, we include vendor-specific fixed effects α_j for each vendor j to control for static differences in base demand across vendors. Meanwhile, a bride’s propensity to make vendor selections may change over time. We capture the time trend effect by a vector X_{jt} that consists of weekly dummies and day-of-week dummies. We include the variable $PagePos_{jt}$ for vendor j ’s average page position on day t . This variable helps to pick up any “web site real estate effect” which 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 (Lohse, 1997).

Our key variables of interest are the two interactive terms $Bridal_j * Test_t * PrevClicks_{jt}$ and $Bridal_j * Test_t * PrevClicks_{jt} * NarrowAppeal_j$. The dummy variable $Bridal_j$ indicates whether the vendor is a bridal shop and consequently belongs in the treatment category; $Test_t$ is an indicator for whether day t occurred during the experiment; $PrevClicks_{jt}$ is a continuous variable of the number of previous cumulative clicks vendor j has received until day t ; and the indicator $NarrowAppeal_j$ equals 1 if the vendor has narrow appeal for being located in a town with a small population. In combination, $Bridal_j * Test_t * PrevClicks_{jt}$ captures the effect of a treated vendor’s past popularity on its current popularity in the test period, and $Bridal_j * Test_t * PrevClicks_{jt} * NarrowAppeal_j$ measures the incremental effect for narrow-appeal vendors. For completeness we also include all lower-order interactive terms except those collinear with the vendor fixed effects and time dummies.

Our identifying assumption is that all categories would have had similar time trends in clicks had it not been for the experimental intervention. The differences-in-differences approach would be problematic if we were studying an apparel retailer and we were trying to use interest in sweaters as a control for the interest in bathing suits. However, in the wedding industry different categories of services, such as bridal shops and florists, are complementary components of the same ultimate wedding product, so interest in one category is likely to

be similar in timing to another category. Indeed, we examine time trends in aggregate clicks in the three categories prior to the experiment. There is no statistically significant evidence of different time trends. Furthermore, because our main coefficients of interest rest on interactions between vendor’s category, time period, previous popularity, and the vendor’s appeal, even if there were category-wide differences in the time trend, so long as these differences were not restricted exclusively to either narrow-appeal or broad-appeal vendors, our relative results hold.

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, Duffo and Mullainathan, 2004). Repeated use of the same exogenous change in variables can lead researchers to overstate the significance of the estimates. To address this concern, as suggested by Hausman, Hall and Griliches (1984), we use a Poisson quasi-maximum likelihood specification with conditional fixed effects and clustering at the vendor level.

Column (1) of Table 2 reports the estimation results when we compare the treatment category with the first control category Florists, where no changes to the web site display were made. The coefficient of $Bridal_j * Test_t * PrevClicks_{jt}$ is positive and significant, suggesting that treated vendors’ popularity does increase with their previous clicks. This result is in line with the main finding of the herding literature that popularity tends to be self-reinforcing. The coefficient $Bridal_j * Test_t * PrevClicks_{jt} * NarrowAppeal_j$ captures the incremental effect for narrow-appeal vendors. It is positive and significant. This result is consistent with our hypothesis. Web site visitors expect broad-appeal vendors to be busier than narrow-appeal vendors. Therefore, when customers see a vendor located in a low-population area receive a large number of clicks, they are more likely to infer high quality than when they see a large-city vendor receive a similar volume of clicks.

Table 2: The Effect of Popularity Information and the Moderating Effect of Appeal: Appeal Defined by Vendor Location

	All Four Months of Data		Short Window	
	(1)	(2)	(3)	(4)
	Florists as Control	Caterers as Control	Florists as Control	Caterers as Control
Bridal * Test * Prev Clicks * Narrow Appeal	0.348*** (0.108)	0.466*** (0.116)	1.617*** (0.353)	1.539*** (0.362)
Bridal * Test * Prev Clicks	0.162*** (0.0263)	0.180*** (0.0250)	0.534*** (0.0672)	0.166** (0.0685)
Bridal * Test * Narrow Appeal	-0.0407 (0.0417)	-0.100** (0.0416)	-0.0628 (0.120)	-0.317*** (0.120)
Bridal * Prev Clicks * Narrow Appeal	0.00276 (0.0150)	-0.00200 (0.0148)	-0.00432 (0.0316)	-0.0408 (0.0317)
Test * Prev Clicks * Narrow Appeal	-0.0225 (0.0433)	-0.00465 (0.0447)	-0.157 (0.115)	-0.145 (0.120)
Bridal * Test	-0.0128 (0.00989)	-0.00524 (0.00978)	-0.0190 (0.0212)	-0.00597 (0.0214)
Bridal * Prev Clicks	-0.0147 (0.0151)	-0.0159 (0.0148)	0.0222 (0.0314)	0.0443 (0.0312)
Test * Prev Clicks	0.00795 (0.0169)	0.00699 (0.0166)	0.0343 (0.0462)	0.0463 (0.0452)
Test * Narrow Appeal	0.00569 (0.0147)	0.0229 (0.0159)	0.00341 (0.0306)	0.0241 (0.0338)
Prev Clicks * Narrow Appeal	-0.156 (0.0974)	-0.204* (0.109)	-4.965*** (0.532)	-3.257*** (0.638)
Previous Clicks	0.0329 (0.0252)	-0.00323 (0.0242)	-0.186** (0.0742)	-0.199*** (0.0722)
Page Position	-0.00371*** (0.00111)	-0.00394*** (0.000617)	0.0116 (0.00826)	-0.00727*** (0.00170)
Vendor Fixed Effects	Yes	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes	Yes
Day of Week Fixed Effects	Yes	Yes	Yes	Yes
Observations	13920	13456	3360	3248
Log-Likelihood	-28457.4	-28194.1	-6473.5	-6337.6

Note: Poisson differences-in-differences specification. Dependent variable: the number of daily clicks a vendor receives. Previous clicks are measured in thousands. In the Bridal Shops treatment category, previous clicks information is displayed, and vendors are ranked in descending order of popularity. In the Florists control category no previous clicks information is displayed, and vendors are ranked alphabetically. In the Caterers control category, no clicks information is displayed, and vendors are ranked in descending order of popularity. Robust standard errors clustered at the vendor level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

The lower-order interactions are insignificant, providing reassurance that there was no unobserved heterogeneity in time trends across categories. The (unreported) coefficients for the weekly fixed effects and day of week fixed effects 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 variable $PagePos_{jt}$ is negative and significant. This suggests that vendors who are listed first on the page (thus having the lowest value of $PagePos_{jt}$) receive more clicks than vendors displayed lower down the page, independently of popularity.

We extend our analysis by using the Caterers category as a second control. In this category, vendors were ranked by popularity but no information about the number of clicks was given. This manipulation allows us to further separate the effect of popularity information from the mere page location effect. Column (2) of Table 2 reports the results. The interactive term $Bridal_j * Test_t * PrevClicks_{jt}$ is positive and significant, suggesting that it is the popularity information rather than reranking which is driving our results. The other key interactive term $Bridal_j * Test_t * PrevClicks_{jt} * NarrowAppeal_j$ is again positive and significant, confirming that our results are not driven by differences in the page position effect for narrow-appeal versus broad-appeal vendors.

One potential concern is that the results in columns (1) and (2) of Table 2 could be subject to serial correlation. For example, a rival web site could have started providing listings of urban bridal shops during the experiment period, which would plausibly reduce the visits to urban bridal vendors in our sample and confound our interpretation of $Bridal_j * Test_t * PrevClicks_{jt} * NarrowAppeal_j$. Fortunately, during the time period we study, the web site that ran the experiment 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 unobserved sources of time-varying shocks which

affect narrow-appeal bridal shops and no other vendors in the experimental period. For example, there could be growing awareness about the bargains to be had at non-urban bridal shops. This would increase both the stock of clicks and the current propensity to click for small-town bridal shops. To address this concern, we use a regression discontinuity approach (Black, 1999; Hahn, Todd and der Klaauw, 2001; Busse, Silva-Risso and Zettelmeyer, 2006). The identification logic is that by taking a very short time window we reduce the likelihood that time-varying shocks (other than the experimental treatment) could explain the results. Columns (3) and (4) of Table 2 report the estimation results when we reduce the time window of evaluation to only include the week before the field experiment and the week into the experiment. The sign and significance of the two key interactive terms are similar to those in columns (1) and (2) of Table 2. The effect size is larger. One interpretation is that the stock of cumulative previous clicks are naturally smaller within the short time window than at the end of the experiment. Another interpretation is that the impact of popularity is diminishing over time, as customers near the end of the experiment rationally know that people are herding into a vendor for its past popularity and the herd does not reflect much new information (Banerjee, 1992; Bikhchandani, Hirshleifer and Welch, 1992).

Yet another concern is that vendors could have reacted strategically to the field experiment. We examine the data for evidence of suspicious clicks (or “click-fraud”) but find no patterns of successive clicks which could suggest vendor manipulation of its clicks volume. Vendor prices are not displayed. Therefore, prices cannot act as an alternative quality signal, and the vendors have no incentive either to strategically change their price to manipulate 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.

4.3 Robustness to Alternative Definition of Appeal

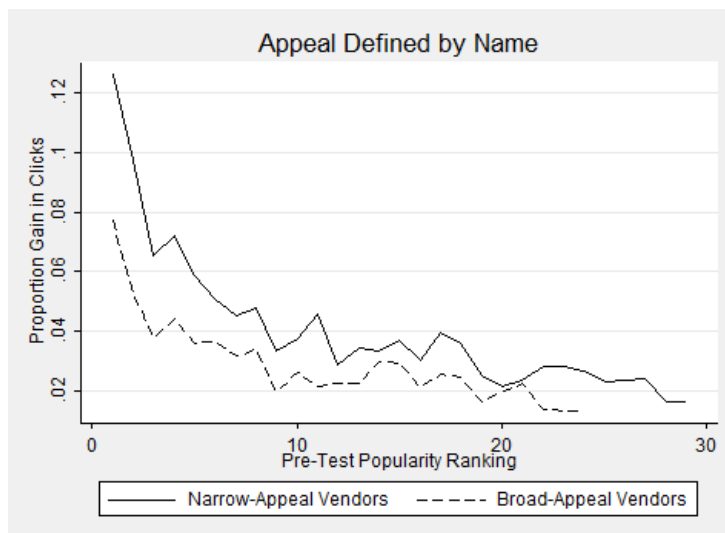
We want to ensure that the results are robust to different definitions of narrow and broad appeal. When brides look at this web site, there are two major cues about the nature of the vendor: vendor location and vendor name. Having established that location moderates the effect of popularity information, we turn to examine whether vendor name also serves as a moderator. The idea is that a vendor with an unfamiliar word in their name (such as “Medieval Brides”) might appear to brides to serve a narrower set of tastes than a vendor with a more generic name (such as “Beautiful Brides”). Kucera and Francis (1967) demonstrate that word usage frequency is highly predictive of word familiarity. Therefore, we augment our data with Pastizzo and Carbone (2007)’s dataset on usage frequency of 1.6 million words in the English language. We define a narrow-appeal vendor as one where each of the words in its name (excluding prepositions and definite articles) is on average used fewer than 50 times.¹²

Figure 3 shows the raw impact of popularity information on the treatment category as measured by proportional gains in clicks after popularity information is released. The patterns are similar to those in Figure 2: past popularity generates further popularity, and the effect is more pronounced for narrow-appeal vendors with unfamiliar names.

We also estimate all specifications with the new definition of appeal. Table 3 reports the results, which are similar to those of Table 2 for all variables of interest except that $Bridal_j * Test_t * PrevClicks_{jt} * NarrowAppeal_j$ in column (1) becomes less significant.

¹²We have verified the robustness of the results with respect to different thresholds.

Figure 3: The Effect of Popularity Information on the Treatment Category (Bridal Shops) When Appeal is Defined by the Familiarity of Vendor Name.



4.4 Instrumental Variables

To further confirm the causal relationship between previous popularity and current popularity, in this section we focus on the treatment category of Bridal Shops. Instead of using other categories as controls, we look for an exogenous shifter of bridal shop popularity to serve as an instrumental variable for the potentially endogenous variables $PrevClicks_{jt}$ and $PrevClicks_{jt} * Narrow_j$.

One such shifter is a bridal shop’s alphabetical position prior to the experiment. With the original yellow-page style of web site display, vendors’ alphabetical position should affect their click volume before the experiment due to the web site real estate effect. In our sample, vendors whose first letter is at the beginning of the alphabet indeed receive more clicks in the pre-test period. However, once vendors are reranked by popularity, alphabetical position is unlikely to affect the popularity of a vendor directly except through the effect of past popularity information that we initially study. Therefore, vendors’ alphabetical position

Table 3: The Effect of Popularity Information and the Moderating Effect of Appeal: Appeal Defined by Vendor Name

	All Four Months of Data		Short Window	
	(1)	(2)	(3)	(4)
	Florists as Control	Caterers as Control	Florists as Control	Caterers as Control
Bridal * Test * Prev Clicks * Narrow Appeal	0.245* (0.140)	0.361*** (0.130)	1.787*** (0.500)	1.051** (0.494)
Bridal * Test * Prev Clicks	0.180*** (0.0256)	0.194*** (0.0248)	0.573*** (0.0657)	0.199*** (0.0676)
Bridal * Test * Narrow Appeal	-0.0423 (0.0532)	-0.110** (0.0529)	-0.171 (0.166)	-0.352** (0.165)
Bridal * Prev Clicks * Narrow Appeal	-0.00249 (0.0150)	0.00101 (0.0148)	0.00652 (0.0317)	0.0163 (0.0317)
Test * Prev Clicks * Narrow Appeal	-0.0464 (0.0531)	-0.0539 (0.0501)	-0.241* (0.142)	-0.307** (0.134)
Bridal * Test	-0.0133 (0.00989)	-0.00602 (0.00977)	-0.0116 (0.0212)	-0.00257 (0.0214)
Bridal * Prev Clicks	-0.0150 (0.0151)	-0.0160 (0.0148)	0.0111 (0.0314)	0.0359 (0.0312)
Test * Prev Clicks	0.00680 (0.0169)	0.00640 (0.0166)	0.0175 (0.0462)	0.0345 (0.0451)
Test * Narrow Appeal	0.0197 (0.0172)	-0.00291 (0.0176)	0.0750** (0.0359)	0.0364 (0.0376)
Prev Clicks * Narrow Appeal	-0.0893 (0.118)	-0.0932 (0.0984)	-5.567*** (0.635)	-1.449** (0.658)
Previous Clicks	0.0243 (0.0250)	-0.00924 (0.0242)	-0.208*** (0.0739)	-0.222*** (0.0720)
Page Position	-0.00388*** (0.00111)	-0.00414*** (0.000622)	0.0177** (0.00823)	-0.00783*** (0.00172)
Vendor Fixed Effects	Yes	Yes	Yes	Yes
Week Fixed Effects	Yes	Yes	Yes	Yes
Day of Week Fixed Effects	Yes	Yes	Yes	Yes
Observations	13920	13456	3360	3248
Log-Likelihood	-28465.3	-28203.2	-6478.9	-6346.3

Note: same as Table 2.

satisfies the exclusion restriction for instrumental variables. We use this alphabetical position as an instrument for $PrevClicks_{jt}$ and $PrevClicks_{jt} * Narrow_j$. We also use its squared term as an instrument to allow for non-linear effects.

We follow Mullahy (1997) in using a GMM instrumental variables procedure for Poisson count data models with endogenous explanatory variables. In our Poisson specification, we assume a multiplicative error to reflect the potential endogeneity of previous clicks. We cluster the standard errors at the vendor level using a bootstrap replication method to adjust for the possibility of vendor level serial correlation.¹³

Note that with this instrumental variables approach we will not be able to separately identify vendor fixed effects. We cannot incorporate instruments into a panel setting with vendor fixed effects because the instrument is not time-varying, rendering any fixed effect perfectly collinear with the instrument. Meanwhile, dynamic panel techniques which circumvent such problems by using lags of the dependent measure as instruments (Arellano and Bond, 1991) are not applicable in our setting, because our lagged dependent measures are precisely reflected in our endogenous variable of interest (previous clicks) and therefore do not satisfy the exclusion restriction. However, measuring vendor fixed effects is not the main focus of our analysis. Furthermore, we will be able to include $Narrow_j$ in the specification, which is otherwise collinear with vendor fixed effects.

Column (1) of Table 4 reports the results of this Poisson specification. The dependent variable is the number of daily clicks that a vendor in the Bridal Shops category receives during the test period. The endogenous variable $PrevClicks_{jt}$, which is now instrumented for, is positive and significant, a result that herding theories would predict. This result also parallels the positive and significant coefficient of $Bridal_j * Test_t * PrevClicks_{jt}$ in Table 2.

¹³See Mullahy (1997) and Windmeijer (2006) for more details about the construction of the two-step weight matrix.

$Narrow_j$ has a negative coefficient, suggesting that *ceteris paribus* vendors that are situated in these low-population towns do indeed receive fewer clicks. This result supports our initial labelling of these vendors as narrow appeal.

The interactive term $PrevClicks_{jt} * Narrow_j$ captures how vendor location moderates the effect of previous clicks. It is positive and significant. The magnitude of the coefficient suggests that vendors at less populous locations receive an 19 percent increase in clicks for the same popularity level, relative to vendors located in populous areas. This finding echoes the positive and significant coefficient of $Bridal_j * Test_t * PrevClicks_{jt} * NarrowAppeal$ in Table 2, providing further empirical support for the hypothesis discussed in Section 2.

In addition, we also include vendor page position as a control. However, the estimate should be interpreted with caution as it is endogenous to the ranking mechanism.¹⁴ The weekly dummies and day of week dummies are much as expected.

We must ensure that our instrumental variables were sufficiently correlated with the endogenous variable—previous clicks. However, there are no well-developed test statistics for the GMM framework. Therefore, we repeat our estimation within a linear regression framework and perform a traditional Anderson canonical correlation test for under-identification on the first-stage regression. As shown in column (2) of Table 4, the first-stage regression is strongly identified. Since there are potentially two endogenous variables of interest (previous clicks and its interaction with whether the vendor has narrow-appeal) and two instruments (alphabet and alphabet squared), the equation is precisely identified and we do not test for over-identification. The linear probability model also allows us to obtain intuitive estimates for the first-stage regressors for the endogenous variables $PrevClicks_{jt}$ and $PrevClicks_{jt} * Narrow_j$. For $PrevClicks_{jt}$ the first-stage estimate for the alphabetical position is -0.052 (p -value<0.001). This result is intuitive; a vendor that begins with Z and

¹⁴We have estimated this specification with and without page location and receive similar results.

was consequently in 26th position is expected to receive a lower number of clicks than a vendor that begins with an A. The coefficient on the squared term of alphabetical position is 0.0002 (p -value <0.001), which suggests that this effect declines down the alphabet.

We again, for robustness, repeat the estimation using a different definition of appeal, which is based on the familiarity with the vendor's name. Columns (3) and (4) of Table 4 report the results, which are similar to those of columns (1) and (2), though the point estimates suggest that vendors with unfamiliar names receive 25 percent more clicks than vendors with familiar names, conditional on the same degree of popularity.

Table 4: Instrumental Variables Estimation for the Treatment Category: Bridal Shops

	Appeal Defined by Location		Appeal Defined by Name	
	(1) Poisson	(2) OLS	(3) Poisson	(4) OLS
Prev Clicks	8.556*** (3.067)	20.82*** (2.147)	5.193*** (1.753)	16.32*** (2.301)
Narrow Appeal Location	-0.569* (0.317)	-1.438** (0.668)		
Prev Clicks * Narrow Appeal Location	1.675** (0.847)	4.839*** (1.667)		
Narrow Appeal Name			-0.406** (0.201)	-1.643** (0.698)
Prev Clicks * Narrow Appeal Name			1.444*** (0.535)	4.743** (1.888)
Page Position	-0.0574** (0.0259)	-0.110*** (0.0256)	-0.0268* (0.0153)	-0.0511* (0.0286)
Week Fixed Effects	Yes	Yes	Yes	Yes
Day of Week Fixed Effects	Yes	Yes	Yes	Yes
Observations	3933	3933	3933	3933
Log-Likelihood		-10219.1		-9904.8
Anderson Under-Identification Test Stat		83.65		64.53
Anderson Test p-value		5.05e-18		6.32e-14

Note: Poisson GMM instrumental variables specification. Dependent variable: the number of daily clicks a vendor receives. Previous clicks are measured in thousands. Sample: the Bridal Shops treatment category, where previous clicks information is displayed and vendors are ranked in descending order of popularity. Endogenous variables: Prev Clicks and Prev Clicks*Narrow Appeal. Instrumental variables: vendor's alphabetical position and its square term. Robust standard errors clustered at the vendor level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

5 Conclusion

Previously, researchers have perceived popularity information as a marketing tool that reinforces the *status quo*, and consequently reinforces the dominance of products that naturally have broader appeal. This perception is based on the belief that broad-appeal 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 narrow-appeal products. It is precisely because narrow-appeal products are less likely to attract customers that when they are actually chosen this choice 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 popularity information in a bestseller format brings the greatest benefits to popular vendors who appear to serve a narrow market, either because of their less populous location or their unfamiliar name. Brides are more likely to infer that a narrow-appeal vendor is high quality compared to a broad-appeal vendor of similar popularity, because the narrow-appeal 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 popularity information benefits popular narrow-appeal 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 broad-appeal products, our results suggest that these display tools can actually strengthen the long tail if it is composed of a sufficient number of popular narrow-appeal products.

There are several potential ways of building on this research. We explore whether a

narrow-appeal product benefits more than a broad-appeal product from popularity information *conditional* on achieving the same popularity. 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 (Wernerfelt, 1995), they may infer broad appeal from popularity. It would be interesting to explore the dynamics by which popularity redefines and is redefined by the perceived appeal of products.

References

- Ai, Chunrong and Edward C. Norton**, “Interaction terms in logit and probit models,” *Economics Letters*, July 2003, *80* (1), 123–129.
- 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.
- Arellano, Manuel and Stephen Bond**, “Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations,” *Review of Economic Studies*, April 1991, *58* (2), 277–97.
- 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*, 577–799.
- 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*, 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, *99* (3), 864–882.
- Celen, Boğaçhan and Shachar Kariv**, “Observational Learning Under Imperfect Information,” *Games and Economic Behavior*, 2004, *47* (1), 72–86.
- 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.
- Hotelling, Harold**, “Stability in Competition,” *The Economic Journal*, 1929, *39*, 41–57.
- 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, and Richard Staelin**, “How to Attract Customers by Giving Them the Short End of the Stick,” *Journal of Marketing Research*, 2007, *44*, 128–141.
- Lohse, Gerald**, “Consumer eye movement patterns on Yellow Pages advertising,” 1997.
- Lucking-Reiley, David**, “Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet,” *American Economic Review*, 1999, *89* (5), 1063–1080.

- Manski, Charles F.**, “Identification of Endogenous Social Effects: The Reflection Problem,” *The Review of Economic Studies*, 1993, 60 (3), 531–542.
- Mullahy, John**, “Instrumental-Variable Estimation of Count Data Models: Applications to Models of Cigarette Smoking Behavior,” *The Review of Economics and Statistics*, 1997, 79 (4), 586–593.
- 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).
- Puhani, Patrick A.**, “The Treatment Effect, the Cross Difference, and the Interaction Term in Nonlinear ‘Difference-in-Differences’ Models,” *SSRN eLibrary*, 2008.
- 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.
- Windmeijer, Frank**, “GMM for Panel Count Data,” October 2006. Discussion Paper No. 06/591, University of Bristol.
- Zhang, Juanjuan**, “The Sound of Silence: Observational Learning from the U.S. Kidney Market,” *Marketing Science*, 2009. forthcoming.

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 deter-

Figure A-1: Choices without Popularity Information

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



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



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.

mined by private signals and taste match when $c \in [\underline{c}, \min(c_S, c_M)]$ or $c \in [\max(c_S, c_M), \bar{c}]$.

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 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-1 provides additional assurance that the interest in the wedding industry is more evenly spread across the year than the

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

(Before the experiment)		(During the experiment)	
Boutique La Reine Boston, MA	(781) 899-0348	Jenn & Jills Bridal Boutique Boston, MA	Clicks 1587
Bridals by Rochelle Uxbridge, MA	(508) 278-9166	Bridals by Valerie Reading, MA	1563
Bridals by Valerie Reading, MA	(781) 942-2525	Boutique La Reine Boston, MA	800
Flair Boston, MA	(617) 247-2828	Yolanda's Waltham, MA	796
Gowns by Jane Nowell, MA	(781) 878-2050	Gowns by Jane Nowell, MA	648
Grandasia Bridal & Fashion Worcester, MA	(508) 328-6380	Flair Boston, MA	573
Jenn & Jills Bridal Boutique Boston, MA	(617) 998-0779	Grandasia Bridal & Fashion Worcester, MA	489
Yolanda's Waltham, MA	(781) 398-1027	Bridals by Rochelle Uxbridge, MA	134

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.

Table A-1: 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