How Does Popularity Information Affect Choices?
A Field Experiment

Catherine Tucker, Juanjuan Zhang

MIT Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142
{cetucker@mit.edu, jjzhang@mit.edu}

Popularity information is usually thought to reinforce existing sales trends by encouraging customers to flock to mainstream products with broad appeal. We suggest a countervailing market force: 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 website that lists wedding service vendors. Our findings are consistent with this hypothesis: narrow-appeal vendors receive more visits than equally popular broad-appeal vendors after the introduction of popularity information.

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 et al. 2006, Cai et al. 2009, Zhang 2010, Chen et al. 2011). 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 breadth of 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 a small niche of the market and consequently have a lower likelihood of being chosen when all products offer the same quality. Similarly, we use "broad appeal" to refer to products that suit mainstream tastes and therefore enjoy a high chance of being chosen among products of the same quality. We show with a simple analytical model that consumers will infer greater quality from a narrow-appeal product than from an equally popular broad-appeal product.

We evaluate this hypothesis using a field experiment from a website that lists wedding service vendors. This website 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 broad or narrow appeal by whether they are located in a town with a large or small population. We find that if customers can easily access popularity information, narrow-appeal vendors receive more visits than equally popular broad-appeal vendors. We verify the robustness of these results by conducting a regression discontinuity analysis, restricting attention to nonadjacent towns, incorporating vendors' price image, and examining the correlation between town size and vendor names. In addition, we check for robustness with respect to the definition of appeal by classifying vendors based on whether their names sound unusual or common (Pastizzo and Carbone 2007). We also present evidence that the effect is economically significant relative to a typical wedding vendor's online advertising budget.

These results are important because it is becoming common for businesses to publicize popularity information online, in part because of the lower costs of information display produced by Internet automation (Shapiro and Varian 1998). Our findings suggest that vendors of popular narrow-appeal, or niche, products

benefit from being listed on websites that make popularity information highly salient. The findings also suggest ways for Internet portals, category managers, and multiproduct firms to redirect sales. Highlighting the popularity of niche products can boost their sales disproportionately, compared to equally popular mainstream products.

This paper draws on the literature of observational learning. Classic analytical studies in this literature find that decision makers tend to follow peer choices as they infer product quality from what their peers have chosen (Banerjee 1992, Bikhchandani et al. 1992). Empirical studies in this direction have also emphasized evidence of quality inference, either in the lab (Anderson and Holt 1997, Celen and Kariv 2004) or in the field (Cai et al. 2009, Zhang 2010, Chen et al. 2011). These studies make winner-takes-all conclusions, that 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, whereas even moderate sales can signify high quality if the product targets only a narrow segment of consumers.

2. Hypothesis and a Theoretical Illustration

We start with a simple model to illustrate our central hypothesis that consumers perceive narrow-appeal products to be of greater quality than broad-appeal products that are equally popular. Products are both horizontally and vertically differentiated, where horizontal 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 versus stochastic models), and one vertical attribute is the quality of teaching. We define "appeal" based on horizontal attributes; a narrow-appeal product is likely to match the tastes of fewer consumers than a broad-appeal product, holding quality constant. 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. Customers use popularity information to update their knowledge of quality. Crucially, however, each product's popularity is interpreted relative to its breadth of appeal.[1]

Suppose there are J vendors, each carrying one product, and a continuum of customers. Let U_{ij} denote the utility customer i derives when visiting vendor j:

$$U_{ij} = v_j + t \cdot I_{ij} + \epsilon_{ij}.$$

The term v_j denotes the quality of vendor j. Customers are uncertain about quality. They hold the prior belief that v_j equals 1 with probability μ and 0 with probability $1 - \mu$. Customer i derives utility $t > 0$ by visiting vendor j that matches her taste (the match indicator variable I_{ij} equalling 1). The error term ϵ_{ij} captures customers' idiosyncratic utility shocks and is independent and identically distributed across customers and vendors. Customer i will visit vendor j if and only if her expected utility $E(U_{ij}) > 0$.[2]

A customer privately knows whether a match occurs. Vendor j matches a customer's taste with ex ante probability m_j. Intuitively, a vendor with broader appeal is likely to match more customers. However, the match probability that a vendor eventually achieves with a certain group of customers may depend on random factors besides its breadth of appeal. For example, the same MBA course may exhibit different taste match probabilities across years, depending on the student cohort and the market environment. To capture this idea, we let a_j denote vendor j's *publicly observable* breadth of appeal, and let $f(m_j \mid a_j)$ denote the conditional probability density function of achieving *unobservable* match probability m_j given a_j. We assume that $f(m_j \mid a_j)$ satisfies the monotonic likelihood ratio property (MLRP) in a_j; for any two match probabilities, the relative chance of achieving the higher match probability increases with the breadth of appeal.[3]

To see how popularity information affects choices, we consider a two-period model. In the first period, a continuum of customers make vendor visit decisions based on their private signals about vendor quality. Allowing a continuum of customers ensures that our results do not rely on vendors "getting lucky" with a few first-period customers. Specifically, we assume that each of these first-period customer receives a private quality signal that can be either high (H) or low (L). Suppose the conditional signal probabilities are $\text{prob}(H \mid v_j = 1) = \text{prob}(L \mid v_j = 0) = q$, where $1/2 < q < 1$, so that signals contain noisy information

---

[1] 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.

[2] This assumption is consistent with the experimental setting. Nevertheless, the intuition underlying our hypothesis remains valid when customers are restricted to visiting a single vendor. The same intuition also holds if customers enjoy a outside utility different from zero when not visiting any vendor.

[3] Analogously assuming that buyer interest satisfies the MLRP in marketing efforts, Miklós-Thal and Zhang (2010) investigate whether "demarketing" will be a profitable seller strategy.

on quality. By Bayes' rule, it is straightforward to see that a first-period customer's posterior quality belief is greater after observing a high signal: $E(v_j \mid H) > E(v_j \mid L)$. The customer's visit decision is further influenced by whether there is a match. Let y_j denote the measure of first-period customers (out of a normalized mass of 1) that choose to visit vendor j.

To illustrate the key effect of interest, we look at second-period customers' quality beliefs after observing vendor j's popularity among first-period customers y_j and its breadth of appeal a_j: $E(v_j \mid y_j, a_j)$. We establish the following result (see the appendix for proof):

$$E(v_j \mid y_j, a_j) \leq E(v_j \mid y_j, a'_j), \quad \forall a_j > a'_j, \qquad (1)$$

where the inequality holds strictly unless vendor popularity is too high or too low.

The intuition is as follows. Match is less likely if the vendor has a narrower appeal. Therefore, from a second-period customers' perspective, a narrow-appeal vendor must have generated more good signals to have achieved the same level of popularity as a broad-appeal vendor. More good signals in turn imply higher quality. In this way, the apparent disadvantage of narrow appeal in matching customer tastes can become an advantage in quality inferences.[4] We state this result with the following hypothesis.

HYPOTHESIS. *For a given level of vendor popularity, customers infer higher quality if the vendor has a narrower appeal.*

The hypothesis is a "conditional" statement. Conditional on achieving the same level of popularity, narrow-appeal products are perceived to offer greater quality than broad-appeal products. Admittedly, narrow-appeal products are less likely to become popular. However, our focus is on empirically understanding whether customers do use product appeal to moderate the quality inference they draw

from popularity.[5] Fortunately, the field experiment approach does allow us to observe customer choices conditional on a given level of popularity. Similarly, while releasing popularity information might signal quality, the field experiment approach ensures that the provision of popularity information is an exogenous experimental manipulation rather than an endogenous firm decision.[6]

3. Field Experiment

3.1. Experimental Setting
We use data from an Internet-based field experiment to evaluate our hypothesis. The website that conducted the field experiment tried out ways to update their alphabetical "yellow pages" listing style to a contemporary "bestseller list" format that presents popularity information saliently. The website provided wedding service vendor listings for a New England state. The number of marriages in the geographic area that the website covers is in line with the national average, suggesting that the activeness of the wedding service market we study is representative of the national wedding service market.

The wedding industry is attractive to study because customers in this industry generally have little prior consumption experience. Even if a bride organizes a second wedding, it is likely that she will select different vendors in order to psychologically differentiate her experience from the previous one. Consequently, customers tend to have imperfect information about vendor quality. At the same time, brides may have private quality signals from other weddings they attended, from various referral sites (Chen et al. 2002), or from third-party reviews (Chen and Xie 2005). Given quality uncertainty and the existence of private signals, observational learning is likely to influence brides' decisions. Finally, this is an industry in which customers take vendor selection seriously. On average, 2.3 million weddings take place in the United States 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.[7] Because of the importance of vendor choices, brides are likely motivated to take cognitive

---

[4] For parsimony, we do not explicitly model the effect of vendor prices. In general, vendor prices can affect customer visits in several ways. For example, posted price may be a quality signal (Simester 1995). However, as we will discuss in §3, the website does not display vendor price information, thus ruling out its signalling effect. Nevertheless, if customers arrive at the website with a preconception of vendors' price image, it may affect utilities and moderate the impact of popularity information—a vendor's popularity despite its high-price image likely implies good quality. We examine this possibility in §5.2. Last, even if customers have no preconception of vendor prices, they may still form equilibrium expectations of prices based on perceived vendor quality. To the extent that high-quality vendors are associated with higher prices, our experiment may be a conservative test of the impact of popularity on quality inference (see Tucker and Zhang 2010 for a study of an analogous trade-off sellers face when entering a popular yet competitive market).

[5] See Miklós-Thal and Zhang (2010) for a theoretical analysis of the ex ante effect of demarketing, a strategy that includes choosing less favorable market environments.

[6] See Lucking-Reiley (1999), 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 (2008) for a discussion of how such experiments have been crucial for firms online.

[7] *Source.* Association of Bridal Consultants from *Brides Magazine* reader survey.

efforts and engage in active quality inferences (Petty and Cacioppo 1981).

We are interested in how popularity information affects customers' decisions to click on the URL of a listed vendor on this website.[8] Popularity information may attract clicks from customers who would otherwise have chosen to seek wedding services from alternative channels, such as a national chain or a department store, rather than visiting one of the stand-alone vendors listed on the website. Approximately 40% of visitors go to the list-of-vendors page without eventually clicking on any vendor's link. This number suggests that for brides the decision to click on the link to a vendor is not trivial or automatic.

The website provides minimal information about vendors on the listing page. It displays only the vendors' name, location, and telephone number. (For a mockup of the webpage, see the electronic companion of this paper on http://www.SSRN.com/.) We will exploit the information on vendor location and name when defining which vendors have narrow appeal.

One concern is that vendors could have reacted strategically to the field experiment. In particular, they could click on their own URL to inflate their popularity. We examine the data for disproportionate successions of clicks during the experiment in the treatment categories relative to the control category. We find no evidence that suggests vendor manipulation.

3.2. Experimental Design and Data
The website lists vendors from 19 categories. Management selected three frequently visited categories for the field study.[9] Random assignment of experimental conditions occurred at the category level. The fact that allocation was random is helpful. For example, it would be troubling if the website had selected a bestseller list display format for a category where it thought brides were uncertain about quality and wanted to give them incremental help, and a yellow pages display format for a category where it thought brides were well informed and just wanted straightforward access to contact information.

The website measures the popularity of a vendor by the number of clicks that vendor's link has received previously. Based on this measure, the following experimental conditions were established. 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 that maintained the original yellow pages 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 (Lohse 1997) or their popularity information.[10]

Given that different formats were applied to different categories, our results could be contaminated if subjects visited categories sequentially. For example, brides could first visit the bridal shops listings and then visit caterers listings but at that stage guessed that these listings were ordered by popularity. Such behavior would lead us to underestimate the treatment effect. Aggregate-level website statistics suggest, however, that over three-quarters of visitors to the list-of-vendors page arrived directly from search engines rather than navigating from within the website. This fact keeps the field experiment close to a between-subjects design.

The field experiment ran for two months, from August to September 2006. The number of previous clicks displayed was calculated using a start date in April 2006. The website 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 was updated instantaneously. In the control conditions, 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 the data and to ensure that there were no glitches in the experiments.

The firm collected data on click behavior based on their Apache Web Server logs. To protect the privacy

---

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

[9] The categories excluded from the experiment were bands, bridal consultants, cakes, DJs, invitations, limos, officiants, photographers, videographers, and "unique ideas" that provide services such as dove release. Website management excluded these categories because of the small number of vendors in each of them.

[10] There was initially another category (reception halls) where management attempted to present popularity information but keep the alphabetical ranking of vendors. The implementation experienced unexpected difficulties because of problems concerning where to place vendors that had a number rather than an alphabetical character at the beginning of their name. We exclude this category from all analysis to achieve precise interpretation of the estimates. Previous versions of this paper included this category; the main results for the effect of popularity information on narrow- versus broad-appeal vendors were qualitatively similar.

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

	Mean	Std. dev.	Min	Max	Observations
Pretest clicks (bridal shops)	337.78	202.25	111	1,230	69
Pretest clicks (florists)	170.67	71.17	66	370	51
Pretest clicks (caterers)	197.10	108.50	44	618	47
Narrow appeal defined by location	0.35	0.49	0	1	167

of the users, IP address information was removed. In this data set, each observation is a time stamp for when a link received a click, alongside the vendor information (including its previous click count) and category affiliation. Our empirical analysis compares a two-month *pretest period* (June and July 2006) with a two-month *test period* (August and September 2006).

There were 366,350 clicks across all the 19 categories in the four-month pretest and test periods.[11] Of these clicks, 48,401 went to 69 bridal shop vendors; 19,613 clicks went to 51 florist vendors; and 20,775 clicks went to 47 caterer vendors. The first three rows of Table 1 report summary statistics on pretest clicks for each of the three categories. Across all three categories, an average vendor received 247 clicks, but there were a few "popular" vendors who received over 1,200 clicks during the pretest period. Because of the disparity in average clicks across different categories, in subsequent regression analysis we mean-center and standardize previous clicks to ensure comparability.

Our primary definition of "appeal" is based on the population in a vendor's town according to 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 home location on the Hotelling line. We define narrow-appeal vendors as those located in towns with a population of less than 50,000.[12] The fourth row of Table 1 reports the summary statistics on the narrow-appeal dummy variable as defined by location. Narrow-appeal vendors on average received 0.3 fewer clicks per day than broad-appeal vendors (significant at the 1% level). Figure 1 displays the frequency distribution of pretest clicks

with all three categories pooled together. Note that most narrow-appeal vendors do have equally popular broad-appeal counterparts. This is a helpful observation because our hypothesis is based on narrow- and broad-appeal vendors receiving the same level of popularity.

4.   Main Analysis

4.1.   Initial Graphical Evidence

We first graphically describe the raw treatment effect. Figure 2 displays the before–after change in the number of clicks in the three categories. For each condition, vendors are grouped based on whether their total pretest clicks were above or below the mean for that category. In addition, we compare the relative change in clicks for broad-appeal vendors and narrow-appeal vendors separately. In the bridal shops treatment category, popular narrow-appeal vendors experience a notable gain in clicks after the treatment, whereas popular broad-appeal vendors do not show an obvious difference. Meanwhile, unpopular broad-appeal vendors suffer a greater loss in clicks than unpopular narrow-appeal vendors. In the florists category, there is a general decrease in clicks within the category, which might suggest a baseline decline in aggregate demand for wedding vendors, absent the experimental intervention, in the late summer. In addition, popular florists, especially those with narrow appeal, experience a notable decline in clicks compared with popular bridal shops. In the caterers category, where listings were reordered according to popularity but no popularity information was revealed, there were no obvious changes observed for unpopular vendors. Interestingly, however, broad-appeal popular vendors witness an increase in clicks whereas narrow-appeal popular vendors experience a decrease, as contrasted with bridal shops. These raw data patterns are consistent with our hypothesis. To control for factors such as time trends and to better assess the economic significance of the effects, we next turn to a regression analysis.

4.2.   Regression Analysis

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 a difference-in-differences method with panel fixed effects to evaluate the treatment effect. Though we have daily data, we follow Bertrand et al. (2004) and divide our data into two periods—a pretest period and a test period—to avoid downward

---

[11] One challenge of processing this data came from unintentional clicks due to, for example, slow website response time. Because 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. We also tried dropping observations when there were more than five requests for the same link within the same minute. There was no substantial change in our findings.

[12] The results are still robust if we set cutoffs at 40,000 or 60,000.

Figure 1   Distribution of Pretest Clicks by Appeal (Location)
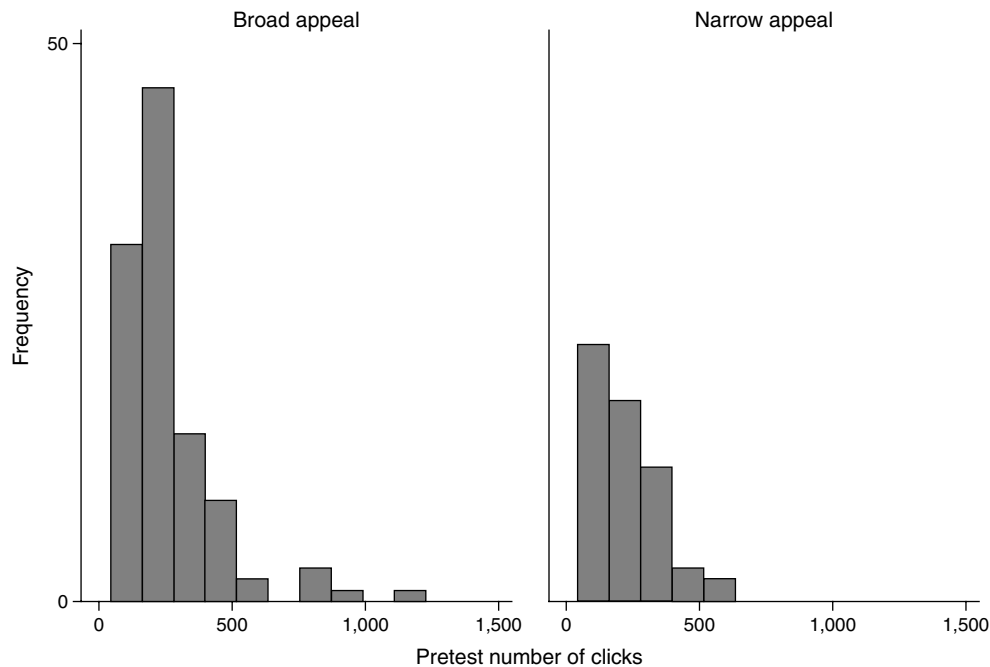
Notes. The sample includes vendors from all three categories. The horizontal axis measures the number of clicks received during the two-month pretest period. The vertical axis measures the number of vendors who receive the corresponding number of pretest clicks.

bias for the standard errors when using multiple data periods.[13]

This section is organized as follows. First, we ask whether there is an effect of popularity information on future clicks. We expect heterogeneous treatment effects because more popular vendors are more likely to attract future clicks when popularity information is released. We identify these effects from the variation in vendors' previous clicks at the start of the experiment. Second, we evaluate our central hypothesis and examine how vendors' breadth of appeal moderates the impact of popularity information. We identify this moderating effect from the variation in vendors' breadth of appeal (as defined by location) for each condition and each popularity level. Finally, we investigate to which extent the effects are driven by the display of popularity information, and to which extent they are caused by the mere reordering of vendors.

4.2.1.   The Effect of Popularity Information. To examine the effect of popularity information, we analyze an ordinary-least-squares specification.[14] We assume that the number of clicks in period t ($t = 1$ or 2)

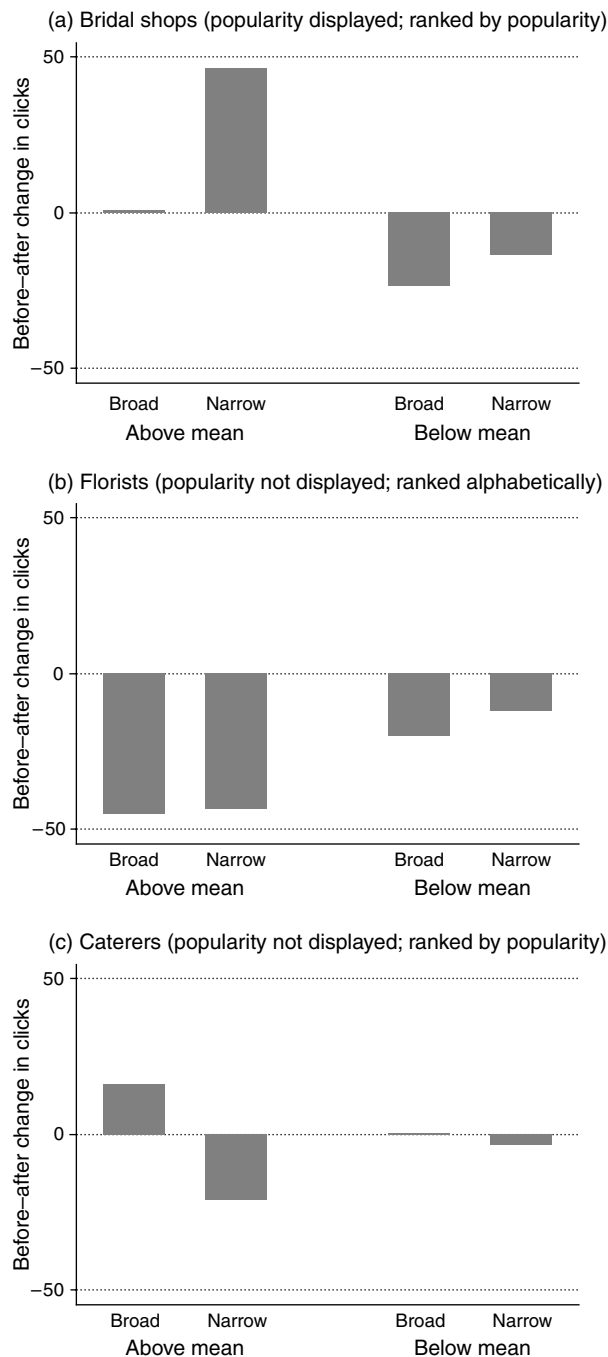for vendor j, $Clicks_{jt}$ can be modeled as

$$
\begin{aligned}
Clicks_{jt} = {} & \alpha_j + \beta_1 Test_t + \beta_2 PrevClicks_{jt} \\
& + \beta_3 Test_t \times Bridal_j \times PrevClicks_{jt} \\
& + \beta_4 Test_t \times Bridal_j + \beta_5 Test_t \times PrevClicks_{jt} \\
& + \beta_6 Bridal_j \times PrevClicks_{jt} + e_{jt}.
\end{aligned} \tag{2}
$$

On the right-hand side, we include a fixed effect α_j for each vendor j to control for static differences in clicks across vendors. Meanwhile, a bride's propensity to make vendor selections may change over time. (See the electronic companion for a review of seasonality in the wedding industry.) We capture the time effect with $Test_t$, an indicator variable denoting whether t is the test period. $PrevClicks_{jt}$ is a continuous measure of cumulative clicks received by vendor j (since April 2006) until the start of period t, which is mean centered and standardized across vendors within the same category over the same pretreatment period.[15] Using a centered and standardized measure both ensures comparability across the different conditions and facilitates interpretation of the effects, because an increase from 0 to 1 in $PrevClicks_{jt}$ simply reflects a one standard-deviation gain in previous clicks. The dummy variable $Bridal_j$ indicates whether vendor j is a bridal shop and consequently belongs in the treatment category. Our key variable of interest is the interactive term $Test_t \times Bridal_j \times PrevClicks_{jt}$, which

---

[13] We have also run our regressions using daily data and have obtained qualitatively similar results.

[14] Our results are also robust to a Poisson distribution that explicitly reflects the fact we use count data. Results are available upon request.

[15] Note that $PrevClicks_{jt}$ is different from pretest clicks that are measured over the two-month pretest period.

Figure 2 **Before–After Changes in Clicks by Popularity and by Appeal (Location)**

(a) Bridal shops (popularity displayed; ranked by popularity)

(b) Florists (popularity not displayed; ranked alphabetically)

(c) Caterers (popularity not displayed; ranked by popularity)

Notes. The vertical axis is the total number of clicks in the test period minus that in the pretest period. In each category, vendors are grouped by whether their pretest clicks are "above mean" or "below mean," and by their breadth of appeal as defined by vendor location.

captures the effect of a treated vendor's past popularity on its clicks in the test period. We also include all lower-order interactive terms except $Bridal_j$, which is collinear with the vendor fixed effects.[16]

[16] We have also run similar specifications that include the variable $PagePos_{jt}$ for vendor j's average page position. This variable

Our identifying assumption is that all categories would have had similar time trends in clicks had it not been for the experimental intervention. The difference-in-differences approach would be problematic if we were studying an apparel retailer and 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 interest in another. Indeed, we examined time trends in aggregate clicks in the three categories prior to the experiment using nonparametric controls for each month interacted with category indicators. A joint F-test on the interactions reveals no statistically significant evidence of different time trends.

Column (1) in Table 2 reports the estimation results when we compare all vendors in the treatment category (bridal shops) and the first control category where no changes to the website display were made (florists). The coefficient of $Test \times Bridal \times PrevClicks$ is positive and significant, suggesting that treated vendors' test-period clicks do increase with their previous popularity. Bridal shops whose popularity is one standard deviation above category mean enjoy 30.53 more clicks in the experimental period. This result is in line with the main finding of the observational learning literature that popularity tends to be self-reinforcing. *Test* is negative and significant, meaning that there is an overall decline in clicks in the bridal shops and florists categories during the test period, consistent with the patterns observed in Figure 2. *PrevClicks* and its interaction with *Bridal* are positive and significant, possibly because past clicks affect choices through other channels of social influence as well. The experiment treatment, however, captures the incremental effect of saliently displaying popularity information. Finally, $Test \times Bridal$ is positive and significant, meaning that treated vendors in the bridal shops category overall attract more clicks in the test period compared with those in the florists control category.

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 (4) in Table 2 reports the results. The key interaction term $Test \times Bridal \times PrevClicks$ is smaller in size

helps to pick up any "website real-estate effect" that 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). Its inclusion makes little difference to the results, as might be expected given its collinearity with $PrevClicks_{jt}$.

Table 2   The Effect of Popularity Information (Appeal Defined by Location)

	Florists as control			Caterers as control		
	(1) Pooled	(2) Narrow appeal	(3) Broad appeal	(4) Pooled	(5) Narrow appeal	(6) Broad appeal
Test × Bridal × PrevClicks	30.53***	67.02***	23.09**	14.80*	60.92**	9.128
	(8.528)	(16.64)	(10.63)	(8.896)	(25.87)	(9.328)
Test	−30.63***	−30.41***	−29.47***	4.383	−5.087	5.297
	(6.302)	(10.72)	(7.804)	(6.765)	(16.09)	(7.675)
PrevClicks	55.02***	51.96***	64.47***	120.7***	79.73	138.8***
	(14.55)	(18.82)	(24.32)	(35.01)	(61.15)	(40.36)
Test × Bridal	24.15***	48.87***	13.25	−10.86	23.55	−21.52**
	(8.311)	(14.37)	(10.01)	(8.771)	(18.61)	(10.17)
Test × PrevClicks	−9.611	−10.44	−7.456	6.122	−4.345	6.501
	(6.532)	(9.829)	(9.016)	(6.871)	(22.28)	(7.148)
Bridal × PrevClicks	186.1***	175.4***	181.9***	120.3***	147.6*	107.6**
	(32.14)	(56.53)	(39.85)	(45.80)	(80.40)	(52.54)
Vendor fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	240	90	150	232	76	156
Log-likelihood	−1,081.7	−403.6	−661.1	−1,052.4	−338.1	−697.6

Notes. Linear panel specification with vendor fixed effects. Dependent variable: the total number of clicks a vendor receives during the pretest versus the test period. Previous clicks are standardized and mean centered. 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.

   $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

and estimated less precisely. This result suggests that some, but not all, of the positive effect of previous popularity comes through premium page positions that popular vendors tend to occupy.

4.2.2. The Moderating Effect of Appeal. To evaluate our main hypothesis, we want to know how vendor appeal moderates the effect of popularity information. As a first exploration, we divide the sample into broad versus narrow-appeal vendors and re-estimate Equation (2). Columns (2), (3), (5), and (6) in Table 2 report the results. The variable of interest *Test × Bridal × PrevClicks* is positive and significant for narrow-appeal vendors, and smaller and less significant for broad-appeal vendors. This result is consistent with our hypothesis: Website 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 certain number of clicks, they are likely to infer higher quality than when they see a large-city vendor receive a same volume of clicks.

To specifically test the statistical significance of the moderating effect of appeal, we use a four-way interaction. We interact the indicator variable *NarrowAppeal$_j$*, which equals 1 if vendor j has narrow appeal, with the right-hand side variables of Equation (2). Column (1) in Table 3 reports the results when we compare the treatment category with Florists. The key interactive term of interest *Test × Bridal × PrevClicks × NarrowAppeal* captures the incremental effect of popularity information for narrow-appeal vendors. It is positive and significant, supporting our main hypothesis. Quantitatively, it suggests that, among bridal shops whose popularity is one standard deviation above category mean, those located in a remote location enjoyed 43.93 more clicks in the experimental period than those in a highly populated location.

Among the lower-order interaction terms, *Test × Bridal × NarrowAppeal* is positive and significant, which suggests that remotely located bridal shops receive an overall boost in clicks in the test period. It could be that releasing popularity information increases clicks on narrow-appeal vendors regardless of their actual level of established popularity, which we examine in §4.3. It could also indicate unobserved heterogeneity in time trends across categories and across appeal, a possibility we shall explore using the regression discontinuity approach.

Column (2) in Table 3 reports the results using the caterers category as a second control. The key interactive term *Test × Bridal × PrevClicks × NarrowAppeal* is positive and significant, suggesting that our results are not simply driven by differences in the page position effect for narrow-appeal versus broad-appeal vendors.

One potential concern is that the results could be subject to serial correlation. For example, a rival web-

Table 3   The Moderating Effect of Appeal (Appeal Defined by Location)

	Full panel		Short window	
	(1) Florists as control	(2) Caterers as control	(3) Florists as control	(4) Caterers as control
Test × Bridal × PrevClicks × NarrowAppeal	43.93**	51.79*	16.01*	20.13*
	(19.11)	(27.14)	(8.627)	(12.12)
Test × Bridal × PrevClicks	23.09**	9.128	13.98***	1.576
	(11.08)	(9.396)	(4.876)	(3.879)
Test	−29.47***	5.297	−13.68***	14.36***
	(8.134)	(7.731)	(3.604)	(3.217)
Test × Bridal	13.25	−21.52**	15.78***	−12.27***
	(10.43)	(10.25)	(4.639)	(4.262)
Test × NarrowAppeal	−0.937	−10.38	−1.272	−3.805
	(12.92)	(17.62)	(5.792)	(8.236)
Test × Bridal × NarrowAppeal	35.63**	45.07**	6.726	9.259
	(17.02)	(20.99)	(7.649)	(9.525)
PrevClicks	64.47**	138.8***	−35.31	−1.666
	(25.35)	(40.66)	(34.15)	(34.17)
Test × PrevClicks	−7.456	6.501	−2.760	9.644***
	(9.397)	(7.200)	(4.131)	(2.983)
Bridal × PrevClicks	181.9***	107.6**	196.6***	163.0***
	(41.54)	(52.92)	(40.18)	(39.73)
NarrowAppeal × PrevClicks	−12.51	−59.04	45.68	26.55
	(30.87)	(72.63)	(45.90)	(95.14)
Test × NarrowAppeal × PrevClicks	−2.986	−10.85	−0.532	−4.651
	(13.15)	(23.08)	(5.654)	(10.39)
Bridal × NarrowAppeal × PrevClicks	−6.458	40.07	−132.5*	−113.3
	(67.29)	(95.19)	(72.80)	(109.4)
Vendor fixed effects	Yes	Yes	Yes	Yes
Observations	240	232	240	232
Log-likelihood	−1,065.0	−1,035.8	−871.4	−831.8

Notes. Linear panel specification with vendor fixed effects. Dependent variable: the total number of clicks a vendor receives during the pretest versus the test period. Previous clicks are standardized and mean centered. 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.

     $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

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 *Test × Bridal × PrevClicks × NarrowAppeal*. Fortunately, during the time period of our study, the website 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 that 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 nonurban 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 et al. 2001, Busse et al. 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 influence the results. Columns (3) and (4) in Table 3 report the estimation results when we reduce the time window of evaluation to only include the week before and the week into the experiment. The key interactive term *Test × Bridal × PrevClicks × NarrowAppeal* remains positive and significant.

Finally, to test directly the extent to which popularity-based reordering affects clicks, we rerun the estimation of Table 3 with data from all three categories pooled together. We examine changes in the bridal shops versus caterers categories, using the florists category as a common control. We report the full results in the electronic companion. In particular,

$Test \times Caterers \times PrevClicks$ is significant and positive, whereas $Test \times Caterers \times PrevClicks \times NarrowAppeal$ is insignificant. This result suggests that popularity does reinforce itself *through* occupying premium page locations, but that popularity-based ranking per se has little differential effect on broad versus narrow-appeal vendors. Moreover, the difference between $Test \times Caterers \times PrevClicks \times NarrowAppeal$ and $Test \times Bridal \times PrevClicks \times NarrowAppeal$ is positive and significant ($p = 0.0159$), suggesting that displaying popularity information has a positive incremental effect on popular narrow-appeal vendors.

4.3. Unconditional Effect of Popularity Information

Our hypothesis and the experiment focuses on the conditional effect of popularity information. We show that narrow-appeal vendors attract more subsequent visits than broad-appeal vendors after the release of popularity information, conditional on vendors receiving the same level of popularity. It is nevertheless interesting to examine the unconditional effect of popularity information across different realizations of popularity. Figure 3 reports the before–after change in clicks across experimental conditions and across breadths of appeal. This echoes Figure 2 but has popular and unpopular vendors in each category pooled together. In the bridal shops treatment category, clicks increase for narrow-appeal vendors and decrease for broad-appeal vendors. The result is remarkable given that the opposite pattern is observed in the caterers control category where vendors are merely reordered by popularity. Clicks decrease in the florists category as discussed before.

Figure 3 **Unconditional Effect of Popularity Information (Appeal Defined by Location)**
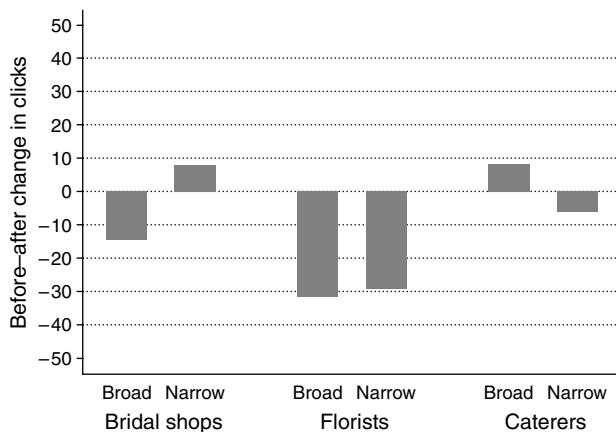
Notes. The vertical axis is the total number of clicks in the test period minus that in the pretest period. In the bridal shops treatment category, previous clicks information is displayed, and vendors are ranked in descending order of popularity. In the florists control category, no previous clicks 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.

Note that two factors affect the unconditional effect of popularity information on narrow versus broad-appeal vendors: the conditional effect we focus on that favors narrow-appeal vendors, and the countervailing fact that narrow-appeal vendors are less likely to become popular. Therefore, the unconditional effect is a high-power test of the conditional effect—the unconditional result that narrow-appeal vendors fare better after the release of popularity information lends strong support to our hypothesized conditional effect. Whether popularity information always benefits narrow-appeal vendors, however, depends on market-specific factors such as the match probability function discussed in §2.

4.4. Economic Significance of the Effect

It will be helpful to put the economic significance of the effects of popularity information into context. The fact that previous clicks are mean centered and standardized in the regressions makes such assessment straightforward. For example, the coefficient of *PrevClicks* indicates the number of clicks a vendor should have gained had it improved its previous clicks by one standard deviation among vendors in the same category. Similarly, for such a one standard deviation improvement in previous clicks, the coefficient $Test \times Bridal \times PrevClicks \times NarrowAppeal$ measures a treated narrow-appeal vendor's gain in clicks during the test period relative to a treated broad-appeal vendor, whereas the sum of $Test \times Bridal \times PrevClicks \times NarrowAppeal$ and $Test \times Bridal \times PrevClicks$ measures a treated narrow-appeal vendor's absolute gain in clicks during the test period.

Using the estimates from column (2) in Table 3 that controls for page location effects, we can infer that, for a one standard deviation increase in previous clicks, a narrow-appeal bridal shop would gain 61 clicks during the test period. In other words, a narrow-appeal bridal shop whose previous clicks are one standard deviation above category mean will gain 61 clicks during the test period after the release of popularity information. To set this effect in context, such a narrow-appeal bridal shop receives 587 clicks in the two-month pretest period, meaning that displaying popularity information increases its clicks by over 10%. This percentage gain is only bigger for less popular narrow-appeal bridal shops. For example, a bridal shop with previous clicks one standard deviation below mean receives 131 clicks in the pretest period. If this vendor increases its previous clicks to the category average, displaying popularity information will give it approximately 50% incremental gain of clicks.

A 10% increase in clicks, with just a redesign of the website, is a remarkable improvement for any website or vendor trying to grow user interest. In this particular setting, the revenue implications could

Table 4 The Effect of Popularity Information and the Moderating Effect of Appeal (Appeal Defined by Location; Nonadjacent Towns)

	Full panel		Short window	
	(1) Florists as control	(2) Caterers as control	(3) Florists as control	(4) Caterers as control
Test × Bridal × PrevClicks × NarrowAppeal	81.46*** (29.00)	57.93** (25.77)	35.00*** (12.62)	32.85*** (11.04)
Test × Bridal × PrevClicks	31.19 (22.00)	23.55 (20.74)	16.46* (9.710)	0.356 (7.685)
Test	−34.05*** (6.416)	23.65** (8.950)	−12.77** (5.717)	20.35*** (2.612)
Test × Bridal	20.26 (12.21)	−37.44*** (13.90)	16.24** (7.447)	−16.88*** (5.538)
Test × NarrowAppeal	−6.362 (14.02)	−22.40* (11.28)	−4.873 (7.505)	−8.844** (3.708)
Test × Bridal × NarrowAppeal	56.38** (21.80)	72.42*** (20.47)	21.34** (10.43)	25.31*** (8.289)
PrevClicks	32.16 (22.70)	162.6*** (31.96)	−44.63 (28.60)	−127.3*** (26.45)
Test × PrevClicks	−12.22 (9.847)	−4.578 (4.979)	1.910 (6.619)	18.01*** (2.487)
Bridal × PrevClicks	306.8*** (86.93)	176.3* (91.64)	245.4*** (43.76)	328.0*** (42.99)
NarrowAppeal × PrevClicks	−12.45 (32.41)	−107.5 (87.11)	−3.715 (46.48)	221.7*** (44.94)
Test × NarrowAppeal × PrevClicks	−17.35 (15.86)	6.180 (6.812)	−9.675 (8.117)	−7.528 (4.910)
Bridal × NarrowAppeal × PrevClicks	−118.1 (113.0)	−23.05 (141.0)	−99.10 (107.9)	−324.6*** (109.3)
Vendor fixed effects	Yes	Yes	Yes	Yes
Observations	126	92	126	92
Log-likelihood	−557.7	−411.5	−461.4	−328.1

Notes. Linear panel specification with vendor fixed effects. Dependent variable: the total number of clicks a vendor receives during the pretest versus the test period. Previous clicks are standardized and mean centered. 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.

$^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

be sizable given that wedding vendors are bidding approximately \$1.67 per click for ads displayed next to Google search results.[17] That is, with a one standard deviation increase in previous clicks, the display of popularity information would save narrow-appeal bridal shops \$611 per year, compared with what they would have paid to grow clicks via Google search ads. This translates into cost savings of 8% of the \$7,500 online spending a typical small business annually incurs (Hopkins 2009).

A natural question of interest is whether there are any effects for broad-appeal vendors. The effect of

Test × Bridal × PrevClicks is smaller and often imprecisely estimated. This suggests that popularity information will be more effective as a tool if used in conjunction with vendors for whom their popularity represents more of a surprise given their natural market appeal.

5. Robustness Checks

5.1. Nonadjacent Towns
An issue with our use of location in defining appeal is the potential spillover in demand across towns (see also Gentzkow et al. 2011). In particular, if some towns with small populations are located disproportionately closely, their actual appeal may be greater than what their population indicates. If this is true, it could contaminate our effect of interest. On the

---

[17] *Data source*. "Google Traffic Estimator" data reflecting midpoint cost per click for top placement in sponsored search results for search terms associated with the vendor category in the specific geographical region we study.

Table 5   The Effect of Popularity Information (Appeal Defined by Name)

	Florists as control			Caterers as control		
	(1) Pooled	(2) Narrow appeal	(3) Broad appeal	(4) Pooled	(5) Narrow appeal	(6) Broad appeal
Test × *Bridal* × *PrevClicks*	30.53***	75.14***	21.72**	14.80*	55.69**	7.341
	(8.528)	(19.72)	(10.42)	(8.896)	(22.36)	(10.14)
Test	−30.63***	−32.67***	−32.44***	4.383	10.41	1.667
	(6.302)	(11.48)	(7.912)	(6.765)	(13.36)	(7.686)
PrevClicks	55.02***	73.25***	42.96**	120.7***	144.8**	110.5***
	(14.55)	(23.91)	(20.12)	(35.01)	(70.29)	(40.52)
Test × *Bridal*	24.15***	36.87**	22.64**	−10.86	−6.217	−11.47
	(8.311)	(14.97)	(10.37)	(8.771)	(17.17)	(10.16)
Test × *PrevClicks*	−9.611	−16.22	−7.726	6.122	3.226	6.656
	(6.532)	(11.67)	(8.358)	(6.871)	(13.48)	(8.035)
Bridal × *PrevClicks*	186.1***	85.86	170.9***	120.3***	14.33	103.4*
	(32.14)	(61.90)	(45.22)	(45.80)	(95.11)	(57.09)
Vendor fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
Observations	240	84	156	232	76	156
Log-likelihood	−1,081.7	−373.3	−698.9	−1,052.4	−345.8	−697.8

Notes. Linear panel specification with vendor fixed effects. Dependent variable: the total number of clicks a vendor receives during the pretest versus the test period. Previous clicks are standardized and mean centered. 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.

*$p < 0.1$; **$p < 0.05$; ***$p < 0.01$.

other hand, it could also weaken the quality signal of popularity enjoyed by vendors from these small but easy-to-get-to towns. To rule out these concerns, we repeat our estimation on a subset of vendors located in towns that are the only town in their county and are not part of a larger metropolitan statistical area. Table 4 reports the results, which are qualitatively similar to those in Table 3. The point estimates for the main effects of interest are higher, which would be expected if indeed proximity to other towns is dampening the signal strength of popularity.

5.2.   Price Image

Vendor prices are not displayed on the website. Therefore, prices cannot act as a direct signal of vendor quality (Simester 1995). This feature rules out the price endogeneity problem, which would have been a key concern if the experiment had been run on a price-grabber style website (Baye and Morgan 2009). However, if customers arrive at the website with a preconceived impression of vendor prices, it could also moderate the impact of popularity information in addition to appeal—high popularity despite steep prices is likely a sign of good vendor quality.

To investigate this issue, we collect data from research done after our experiment by the same website that divided vendors into "bargain" and "non-bargain" categories. Approximately 12% of vendors in our sample are rated as bargain vendors. We rerun the regressions in Table 3 but with an additional

layer of interaction with the *Bargain* dummy. The full results are presented in the electronic companion to this paper. The key effect of interest *Test* × *Bridal* × *PrevClicks* × *NarrowAppeal* remains significant and positive. Meanwhile, *Test* × *Bridal* × *PrevClicks* × *Bargain* and *Test* × *Bridal* × *PrevClicks* × *NarrowAppeal* × *Bargain* are both insignificant. One possible interpretation is that customers have limited knowledge about vendor prices on this website because of the lack of prior experience with wedding services.

5.3.   Town Size and Vendor Name

There is a subtle effect to consider when using the yellow pages listing style as a benchmark. If highly populated towns have a large mass of vendors, some might have to adopt names such as "AA Weddings." In this case, there may be a disproportionate concentration of broad-appeal vendors, as defined by location, near the top of a yellow page listing. Consequently, the experiment treatment that reorders vendors by popularity may move a disproportionate number of broad-appeal vendors down the list, thus confounding our effect of interest.

To investigate this possibility, we check all vendors in our data and find no adoption of "AA" or similar strings at the beginning of vendor names. We further compare the proportions of vendors with initial "A" in high versus low-population towns. These proportions are 0.065 and 0.085, respectively, and the difference between the two is insignificant ($p = 0.640$).

Table 6   The Moderating Effect of Appeal (Appeal Defined by Name)

	Full panel		Short window	
	(1) Florists as control	(2) Caterers as control	(3) Florists as control	(4) Caterers as control
Test × *Bridal* × *PrevClicks* × *NarrowAppeal*	53.42**	48.35**	21.62**	18.74**
	(22.30)	(23.15)	(9.424)	(9.303)
Test × *Bridal* × *PrevClicks*	21.72**	7.341	10.64**	0.0797
	(10.42)	(10.56)	(4.579)	(4.218)
Test	−32.44***	1.667	−13.94***	12.17***
	(7.914)	(8.009)	(3.335)	(3.093)
Test × *Bridal*	22.64**	−11.47	13.73***	−12.37***
	(10.37)	(10.59)	(4.395)	(4.096)
Test × *UnusualName*	−0.235	8.741	0.877	5.592
	(13.94)	(14.69)	(5.755)	(5.668)
Test × *Bridal* × *UnusualName*	14.23	5.252	6.933	2.218
	(18.21)	(19.04)	(7.572)	(7.311)
PrevClicks	42.96**	110.5**	−16.52	−29.46
	(20.13)	(42.22)	(24.64)	(37.65)
Test × *PrevClicks*	−7.726	6.656	−2.907	7.656**
	(8.359)	(8.372)	(3.525)	(3.205)
Bridal × *PrevClicks*	170.9***	103.4*	64.83	77.77
	(45.23)	(59.49)	(47.85)	(53.84)
UnusualName × *PrevClicks*	30.30	34.26	32.26	93.05
	(31.25)	(77.31)	(58.25)	(67.06)
Test × *UnusualName* × *PrevClicks*	−8.494	−3.429	−0.149	2.732
	(14.35)	(14.98)	(5.595)	(5.994)
Bridal × *UnusualName* × *PrevClicks*	−85.08	−89.04	33.57	−27.22
	(76.64)	(105.9)	(78.33)	(83.13)
Vendor fixed effects	Yes	Yes	Yes	Yes
Observations	240	232	240	232
Log-likelihood	−1,072.2	−1,043.9	−867.4	−823.2

Notes. Linear panel specification with vendor fixed effects. Dependent variable: the total number of clicks a vendor receives during the pretest versus the test period. Previous clicks are standardized and mean centered. 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.

$^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Last, we summarize the frequency of the 26 letters appearing as the first letter in the names of vendors from high versus low-population towns in our sample. Reassuringly, these two measures are positively and significantly correlated, with the correlation coefficient being 0.718 ($p = 0.000$).

5.4. Alternative Definition of Appeal

Finally, we want to ensure that the results are robust to different definitions of appeal. When brides look at this website, there are two major cues about the nature of the vendor: location and 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 serve a nar-

rower set of tastes than a vendor with a generic name (such as "Beautiful Brides").

We follow Kucera and Francis (1967), who demonstrate that word usage frequency is highly predictive of word familiarity. We augment our data with the Pastizzo and Carbone (2007) data set on usage frequency of 1.6 million words in the English language, which in turn is based on the Simpson et al. (2002) recording and transcription of 190 hours of speech at the University of Michigan between 1997 and 2001. 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 during the 190 hours of recording by Simpson et al. (2002).[18] We also check that the word

---

[18] We have verified the robustness of the results with respect to different thresholds.

frequency measure we use is consistent with a newer corpus of the English language recently collected by Brysbaert and New (2009) based on film and television subtitles.

Under the new definition of appeal, 33% of vendors are categorized as narrow-appeal vendors, similar to that under the location-based definition of appeal. Nevertheless, there is only a -0.0208 correlation between whether a vendor has an unusual name and remote location, suggesting that our alternative definition of appeal does capture different vendors. We repeat the panel regressions with the new definition of appeal. Tables 5 and 6 report the results, which parallel those where appeal is defined by location and indicate a similar effect size. Finally, the electronic companion presents figures of pretest clicks distribution, the before–after changes in clicks across popularity buckets and across appeal, and the unconditional effect of popularity information by this new definition of appeal. These figures reveal patterns similar to their counterpart that defines appeal based on location.

6. Conclusion

Popularity information has been perceived as a marketing tool that reinforces the status quo and strengthens 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 choices. We propose a countervailing effect that popularity information may actually be of greater benefit to narrow-appeal products. Just because narrow-appeal products are less likely to attract customers, when they are actually chosen this choice conveys a greater quality signal to other customers.

We evaluate this hypothesis using data from a field experiment conducted with a website that lists wedding service vendors. We find that a bestseller format, which displayed popularity information and ranks vendors by popularity, brings more clicks to vendors who appear to serve a narrow market—either because of their less populous location or their unfamiliar name—than equally popular vendors with a broad appeal. This happens if brides think that a narrow-appeal vendor must offer high enough quality to overcome its naturally smaller market and become as popular as a vendor with broad appeal. We verify the robustness of our results in a number of ways.

There are several limitations to this research. First, the data contain a single experiment from a specific product market in a single geographic area. It will be useful to study how popularity information affects choices in other contexts. For instance, we do not explore whether popularity information can be similarly moderated by other marketing variables.

Second, the experimental treatments occur at the category level. In the absence of within-category variation induced by the experimental design, our ability to make inferences relies on these categories being good controls for each other. The method should be generalized with caution if the categories are subject to inherently different dynamics unobservable to the researcher. Third, the outcome variable in this data is clicks. To fully quantify the economic consequences of Internet-based information campaigns, it would be helpful to investigate actual purchase decisions. Last, we assume that customers know which vendor matches their tastes. It would be interesting to explore the effect of popularity information when customers infer their preferences from product offerings (Wernerfelt 1995), or if they have to incur an evaluation cost to identify a matching product (Kuksov and Villas-Boas 2010). Notwithstanding these limitations, this paper suggests that popularity information has asymmetric effects on products of different breadth of appeal, a finding that is relevant to firms who are interested in displaying popularity information as a marketing tool.

Appendix. Proof of Result (1)
From a second-period customer's perspective, a first-period customer matches vendor j with probability m_j, and receives a good signal with probability q if vendor j is high quality and $1 - q$ otherwise. Given match and a high signal, first-period customer i visits vendor j if and only if $E(v_j \mid H) + t + \epsilon_{ij} > 0$. For notational simplicity, let ϵ_{ij} be symmetrically distributed around zero and have positive support everywhere, following a cumulative distribution function G. (The intuition underlying the proof can apply to other specifications of ϵ_{ij}.) The chance of a visit given match and a high signal can then be written as $1 - G[-E(v_j \mid H) - t] = G[E(v_j \mid H) + t]$. Similarly, this visit probability becomes $G[E(v_j \mid L) + t]$ given match and a low signal; $G[E(v_j \mid H)]$ given mismatch and a high signal, and $G[E(v_j \mid L)]$ given mismatch and a low signal. It follows that

if vendor j is high quality, its popularity is a function of its realized match probability m_j in the following way:

$$y_j = \phi_H(m_j) = qm_jG[E(v_j \mid H) + t] + (1-q)m_jG[E(v_j \mid L) + t]$$
$$+ q(1-m_j)G[E(v_j \mid H)] + (1-q)(1-m_j)G[E(v_j \mid L)].$$

Similarly, if vendor j is low quality, its popularity should be a function of m_j specified below:

$$y_j = \phi_L(m_j) = (1-q)m_jG[E(v_j \mid H) + t] + qm_jG[E(v_j \mid L) + t]$$
$$+ (1-q)(1-m_j)G[E(v_j \mid H)] + q(1-m_j)G[E(v_j \mid L)].$$

Functions ϕ_H and ϕ_L are continuous and linearly increasing in m_j. Therefore, a second-period customer can unambiguously infer that vendor j is low quality for any $y_j < \phi_H(0)$, and high quality for any $y_j > \phi_L(1)$. If the set $[\phi_H(0), \phi_L(1)]$ is nonempty, a second-period customer remains uncertain about vendor j's quality for any y_j belonging in this set. Specifically, her expected quality of vendor j is given by Bayes' rule as

$$E(v_j \mid y_j, a_j) = \frac{f(\phi_H^{-1}(y_j) \mid a_j)\mu}{f(\phi_H^{-1}(y_j) \mid a_j)\mu + f(\phi_L^{-1}(y_j) \mid a_j)(1-\mu)}$$
$$= \frac{1}{1 + (f(\phi_L^{-1}(y_j)\mid a_j)/f(\phi_H^{-1}(y_j)\mid a_j))((1-\mu)/\mu)},$$

where ϕ^{-1} denotes the inverse function of ϕ.

Note that $\phi_H(0) = qG[E(v_j \mid H)] + (1-q)G[E(v_j \mid L)] > \phi_L(0) = (1-q)G[E(v_j \mid H)] + qG[E(v_j \mid L)]$ because $q > 1/2$ and $E(v_j \mid H) > E(v_j \mid L)$. Similarly, $\phi_H(1) > \phi_L(1)$. Therefore, $\phi_H(m_j) > \phi_L(m_j)$ for all $m_j \in [0, 1]$, and hence $\phi_L^{-1}(y_j) > \phi_H^{-1}(y_j)$ for any $y_j \in [\phi_H(0), \phi_L(1)]$.

Finally, the MLRP assumption requires that

$$\frac{f(m_j \mid a_j)}{f(m_j' \mid a_j)} > \frac{f(m_j \mid a_j')}{f(m_j'\mid a_j')}, \quad \forall m_j > m_j', \ a_j > a_j',$$

which implies that

$$E(v_j \mid y_j, a_j) < E(v_j \mid y_j, a_j'), \quad \forall a_j > a_j', \ y_j \in [\phi_H(0), \phi_L(1)]$$

if the set $[\phi_H(0), \phi_L(1)]$ is nonempty. Q.E.D.

References

Anderson, L. R., C. A. Holt. 1997. Information cascades in the laboratory. *Amer. Econom. Rev.* **87**(5) 847–862.

Anderson, E. T., D. I. Simester. 2004. Long-run effects of promotion depth on new versus established customers: Three field studies. *Marketing Sci.* **23**(1) 4–20.

Banerjee, A. V. 1992. A simple model of herd behavior. *Quart. J. Econom.* **107**(3) 797–817.

Baye, M. R., J. Morgan. 2009. Brand and price advertising in online markets. *Management Sci.* **55**(7) 1139–1151.

Bertrand, M., E. Duflo, S. Mullainathan. 2004. How much should we trust differences-in-differences estimates? *Quart. J. Econom.* **119**(1) 249–275.

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

Black, S. E. 1999. Do better schools matter? Parental valuation of elementary education. *Quart. J. Econom.* **114**(2) 577–799.

Brysbaert, M., B. New. 2009. Moving beyond Kucera and Francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for American English. *Behav. Res. Methods* **41**(4) 977–990.

Busse, M., J. Silva-Risso, F. Zettelmeyer. 2006. $1,000 cash back: The pass-through of auto manufacturer promotions. *Amer. Econom. Rev.* **96**(4) 1253–1270.

Cai, H., Y. Chen, H. Fang. 2009. Observational learning: Evidence from a randomized natural field experiment. *Amer. Econom. Rev.* **99**(3) 864–882.

Celen, B., S. Kariv. 2004. Observational learning under imperfect information. *Games Econom. Behav.* **47**(1) 72–86.

Charness, G., E. Haruvy, D. Sonsino. 2007. Social distance and reciprocity: An Internet experiment. *J. Econom. Behav. Organ.* **63**(1) 88–103.

Chen, Y., J. Xie. 2005. Third-party product review and firm marketing strategy. *Marketing Sci.* **24**(2) 218–240.

Chen, Y., G. Iyer, V. Padmanabhan. 2002. Referral infomediaries. *Marketing Sci.* **21**(4) 412–434.

Chen, Y., Q. Wang, J. Xie. 2011. Online social interactions: A natural experiment on word of mouth versus observational learning. *J. Marketing Res.* Forthcoming.

Gentzkow, M., J. M. Shapiro, M. Sinkinson. 2011. The effect of newspaper entry and exit on electoral politics. *Amer. Econom. Rev.* Forthcoming.

Greenstein, S. 2008. Economic experiments and neutrality in Internet access markets. A. Jaffe, J. Lerner, S. Stern, eds. *Innovation Policy and the Economy*, Vol. 8. University of Chicago Press, Chicago, 59–110.

Hahn, J., P. Todd, W. Van der Klaauw. 2001. Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica* **69**(1) 201–209.

Hopkins, H. 2009. How much should your company spend on Internet marketing? Technical report, July 11. Avalanche Internet Marketing. http://www.avalancheinternetmarketing.com/determine-internet-marketing-budget/.

Kucera, H., W. N. Francis. 1967. *Computational Analysis of Present-Day American English*. Brown University Press, Providence, RI.

Kuksov, D., J. M. Villas-Boas. 2010. When more alternatives lead to less choice. *Marketing Sci.* **29**(3) 507–524.

Lim, N., M. Ahearne, S. H. Ham. 2009. Designing sales contests: Does the prize structure matter? *J. Marketing Res.* **46**(3) 356–371.

Lo, A. K. C., J. G. Lynch, R. Staelin. 2007. How to attract customers by giving them the short end of the stick. *J. Marketing Res.* **44**(1) 128–141.

Lohse, G. 1997. Consumer eye movement patterns on yellow pages advertising. *J. Advertising* **26**(1) 61–73.

Lucking-Reiley, D. 1999. Using field experiments to test equivalence between auction formats: Magic on the Internet. *Amer. Econom. Rev.* **89**(5) 1063–1080.

Miklós-Thal, J., J. Zhang. 2010. Demarketing. Working paper, University of Rochester, Rochester, New York.

Pastizzo, M. J., R. F. Carbone. 2007. Spoken word frequency counts based on 1.6 million words in American English. *Behav. Res. Methods* **39**(4) 1025–1028.

Petty, R. E., J. T. Cacioppo. 1981. *Attitudes and Persuasion: Classic and Contemporary Approaches*. Westview Press, Boulder, CO.

Salganik, M. J., P. S. Dodds, D. J. Watts. 2006. Experimental study of inequality and unpredictability in an artificial cultural market. *Science* **311**(5762) 854–856.

Shapiro, C., H. R. Varian. 1998. *Information Rules: A Strategic Guide to the Network Economy*. Harvard Business School Press, Boston.

Simester, D. 1995. Signaling price image using advertised prices. *Marketing Sci.* **14**(2) 166–188.

Simpson, R. C., S. L. Briggs, J. Ovens, J. M. Swales. 2002. *The Michigan Corpus of Academic Spoken English*. The Regents of the University of Michigan, Ann Arbor.

Tucker, C., J. Zhang. 2010. Growing two-sided networks by advertising the user base: A field experiment. *Marketing Sci.* **29**(5) 805–814.

Wernerfelt, B. 1995. A rational reconstruction of the compromise effect: Using market data to infer utilities. *J. Consumer Res.: An Interdisciplinary Quart.* **21**(4) 627–633.

Zhang, J. 2010. The sound of silence: Observational learning from the U.S. kidney market. *Marketing Sci.* **29**(2) 315–335.