



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Ad Revenue and Content Commercialization: Evidence from Blogs

Monic Sun, Feng Zhu

To cite this article:

Monic Sun, Feng Zhu (2013) Ad Revenue and Content Commercialization: Evidence from Blogs. Management Science 59(10):2314-2331. <http://dx.doi.org/10.1287/mnsc.1120.1704>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2013, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Ad Revenue and Content Commercialization: Evidence from Blogs

Monic Sun

Graduate School of Business, Stanford University, Stanford, California 94305; and
University of Southern California, Los Angeles, California 90089, monic@stanford.edu

Feng Zhu

Marshall School of Business, University of Southern California, Los Angeles, California 90089,
fzhu@marshall.usc.edu

Many scholars argue that when incentivized by ad revenue, content providers are more likely to tailor their content to attract “eyeballs,” and as a result, popular content may be excessively supplied. We empirically test this prediction by taking advantage of the launch of an ad-revenue-sharing program initiated by a major Chinese portal site in September 2007. Participating bloggers allow the site to run ads on their blogs and receive 50% of the revenue generated by these ads. After analyzing 4.4 million blog posts, we find that, relative to nonparticipants, popular content increases by about 13 percentage points on participants’ blogs after the program takes effect. About 50% of this increase can be attributed to topics shifting toward three domains: the stock market, salacious content, and celebrities. Meanwhile, relative to nonparticipants, participants’ content quality increases after the program takes effect. We also find that the program effects are more pronounced for participants with moderately popular blogs, and seem to persist after participants enroll in the program.

Key words: ad-sponsored business model; media content; blog; revenue sharing; user-generated content; platform-based markets

History: Received May 13, 2011; accepted July 26, 2012, by Pradeep Chintagunta, marketing. Published online in *Articles in Advance* April 4, 2013.

1. Introduction

Media consumption today is characterized by three patterns. First, information acquisition is moving online. For example, in 2010, for the first time, more people obtained news online than from print newspapers (Pew Research Center 2012). Second, a significant portion of the content consumed online is generated by the consumers themselves (Zhang and Sarvary 2011, Sun 2012).¹ All three media sites among the top 10 most-visited websites in the world—YouTube, Blogger, and Twitter—are based on user contributions.² Third, consumers are increasingly expecting their media consumption to be free (Wray 2010). As a result, content providers are under increasing pressure to monetize their content through ad revenue.

Indeed, ad-sponsored business models appear to be increasingly prevalent among online content sites (e.g., Casadesus-Masanell and Zhu 2010, 2013; Goldfarb 2004). Leading content sites, such as YouTube and Hulu, rely entirely on ad revenue to finance their operations. Small content sites can take advantage of programs from advertising aggregators, such as Google’s AdSense, to generate ad revenue without finding advertisers themselves.³ At the same time, ad-sponsored business models are no longer limited to website owners. Individual content providers today can also earn ad revenue from the content they provide. For example, the most popular video-sharing site, YouTube, started sharing ad revenue with its top contributors in 2007, and it recently extended the ad-revenue-sharing program to all contributors.⁴ Many other content sites based on user-generated content (e.g., blog sites such

¹ For example, in 2008, 12% of Internet users (9% of all adults) blogged, and 33% of Internet users (24% of all adults) read blogs (Mayzlin and Yoganasimhan 2012). A blog (a blend of the term Web log) is a type of website or part of a website. Blogs are usually maintained by an individual with regular entries of commentaries, descriptions of events, or other material, such as pictures or video clips. Entries are commonly displayed in reverse chronological order.

² Source: <http://www.alexa.com> (accessed March 2011).

³ Website owners can enroll in such programs to enable advertisements on their websites. These advertisements are delivered by advertising aggregators such as Google and generate revenue for the website owners on either a per-click or per-impression basis.

⁴ See <http://www.youtube.com/partners> (accessed September 2010). See also Yoganasimhan (2012) for an interesting discussion on how content propagation depends on the social network structure in a context such as YouTube.

as Blogger and WordPress) have adopted similar practices.

Many scholars criticize the use of ad-sponsored business models in media industries (e.g., Baker 1994, Cross 1994, Herman and McChesney 1997, Turow 1998, Hamilton 2004, McChesney 2004, Anderson and Gabszewicz 2006). They argue that when supported by advertising revenue, media firms, both online and offline, have incentives to cater content production to popular tastes so that they can attract the maximal number of eyeballs. As a result, popular content will be duplicated and excessively supplied, leaving viewers with niche preferences underserved (e.g., Anderson and Gabszewicz 2006). Anecdotal evidence seems to support these criticisms. Broadcast television networks in the United States, for example, are frequently blamed for abolishing advertising-unfriendly programs and sticking with redundant ones (e.g., Brown and Cavazos 2003, McChesney 2004, Wilbur 2008). Similarly, most newspapers and magazines are charged with being designed for advertising rather than fundamental editorial content (Bagdikian 2004, pp. 241–246).

In addition, many communications scholars (e.g., Steinem 1990; Herman and McChesney 1997, p. 137) argue that popular content is often not the most consequential to readers and may promote unintended social norms. For example, because violence and sex generally sell well, many content providers routinely employ them. Although regulations such as the Fairness Doctrine require commercial broadcasters to present an ample number of issues of public importance, McChesney (2004, p. 44) points out that these regulations have never been enforced. Because media are an important driver of culture, some critics go so far as to argue that the advertising–media relationship is effectively destroying the culture and that we in society are “amusing ourselves to death” (Postman 2005).

The theoretical literature in economics repeatedly offers support for the claim that ad revenue induces content providers to cater to the majority taste by producing popular, duplicated content. Early theoretical studies (e.g., Steiner 1952, Beebe 1977) show that when TV broadcasters are sponsored by advertisers and thus have the sole objective of maximizing viewership, they are likely to choose the same program type and split the market in equilibrium, a result often referred to as the “Principle of Duplication” (Anderson and Gabszewicz 2006). As an example, if 70% of the population watches sports and the rest watches history, in a duopolistic market, the two competing TV broadcasters will both offer sports programs and split the market. In recent studies, scholars (e.g., Gabszewicz et al. 2001, Gal-Or and Dukes 2003, Anderson and Gabszewicz 2006, Gabszewicz et al.

2006, Peitz and Valletti 2008) extend the analysis by explicitly modeling media markets as two-sided markets and find similar results.

Although the relationship between ad-sponsored business models and content providers’ incentives has received great attention in the theoretical literature, it has received surprisingly little empirical evaluation. The lack of empirical evidence is perhaps due to the difficulty of establishing a causal relationship: Providers of popular content are more likely to seek ad revenue, and as a result, the causal relationship could be in the opposite direction.⁵ In this study, we empirically evaluate the impact of ad-sponsored business models on content providers’ incentives by taking advantage of the introduction of an ad-revenue-sharing program by a major Chinese portal site in September 2007. Participating bloggers allow the site to run ads on their blogs and, in return, they receive 50% of the revenue generated by these ads. Our empirical setting goes beyond the duopolistic or oligopolistic setups in the theoretical literature (e.g., Steiner 1952, Beebe 1977, Gal-Or and Dukes 2003), as there are millions of bloggers. It is therefore interesting to examine whether ad revenue still motivates participating bloggers to shift toward popular content at the expense of an elevated level of competition with each other, and if so, how strong this effect is.

We use a difference-in-differences approach to compare the content shift of 4,200 participants before and after the program takes effect to that of 26,974 non-participants. We also employ fixed effects and instrumental variables approaches to account for bloggers’ endogenous decisions to participate in the program. After analyzing 4.4 million blog posts, we find that, relative to nonparticipants, popular content increases by about 13 percentage points on participants’ blogs after the program takes effect. About 50% of this increase can be attributed to topics shifting toward three domains: the stock market, salacious content, and celebrities.

We also examine the shift in content quality, an aspect not discussed in the theoretical literature, because it is unclear whether participants will devote more or less effort into producing content in which they are not intrinsically interested. We find a significant quality improvement for participants’ blog posts.

In addition, we find that the program effect differs across different bloggers. In particular, participants with moderately popular blogs shift their content popularity, topics, and quality more than both non-participants and those participants with very popular blogs. This result may reflect that nonparticipants

⁵ In a similar vein, Kind et al. (2009) show that the degree of content differentiation between media firms’ products may affect their dependency on ad revenue.

derive a high level of disutility when deviating from their natural tastes, whereas very popular bloggers have always covered popular topics and maintained a high level of quality for their content, leaving little room for improvement on these dimensions.

Our panel data also allow us to examine the persistence of the program effect. We find that the program effect does not diminish after bloggers begin participating in the program. In fact, during the period under study, the program effect seems to become stronger over time for some of the outcome variables on content popularity, topics, and quality.

In addition to the instrumental variables approach, we explore the robustness of our results against potential selection bias with three alternative approaches: propensity-score matching, Rosenbaum bounds, and an approach first developed by Altonji et al. (2002, 2005) and recently extended by Sen et al. (2011). Each of these approaches operates on different assumptions, and thus they provide complementary evidence that the revenue-sharing program indeed has a significant impact on content popularity, topic domains, and quality.

Our study offers important implications for both consumers and advertisers. From the consumers' perspective, our results suggest that popular content will be more dominant once ad-revenue-sharing programs take place, which makes it harder for consumers with niche tastes to find free content that matches their interests. Meanwhile, we also find that the quality of posts increases because of ad revenue. The trade-off hence suggests that the welfare implication of adopting ad-sponsored business model is ambiguous.

For advertisers, it is important to understand how monetary incentives affect content providers' behavior, which in turn influences the types of consumers that their ads will reach. In particular, as content providers have incentives to change content topics, advertisers need to monitor the types of content their ads are associated with, because such associations may affect their brand images.

It is tempting for intermediaries or platform providers in media industries such as YouTube and Blogger to adopt ad-revenue-sharing programs. These intermediaries are operating in multisided markets (e.g., Anderson and Gabszewicz 2006, Chen et al. 2012, Seamans and Zhu 2011), which are prone to winner-take-all dynamics because of significant indirect network effects (e.g., Eisenmann et al. 2006, Zhu and Iansiti 2012). As a result, they are under tremendous pressure to grow their traffic by encouraging contributions from content providers and attracting eyeballs. Our results show that sharing ad revenue indeed incentivizes content providers to contribute more frequently. In addition, although consumers may react negatively to ads, sharing ad revenue leads

to more popular and higher quality content that many consumers enjoy. At the same time, however, intermediaries are likely to see a decrease in the share of content on less popular or niche topics after the launch of such ad-revenue-sharing programs. Consumers with niche tastes may switch to content sites that are subscription based.

Broadly speaking, this paper contributes to the growing literature that examines factors influencing media content. Scholars have examined how the positioning of media content is affected by the entry of competitors (e.g., George and Waldfogel 2006, Seamans and Zhu 2012) and the mix of consumer types (e.g., George and Waldfogel 2003), how content quality changes with the emergence of the Internet (e.g., Frijters and Velamuri 2010), and how content variety changes as media firms consolidate (e.g., Berry and Waldfogel 2001; George 2002, 2007). They have also identified sources of media bias, such as pressure from advertisers or the government (e.g., Price 2003, Reuter and Zitzewitz 2006, Rinallo and Basuroy 2009), readers' desire to reinforce their own prior beliefs (e.g., Mullainathan and Shleifer 2005, Gentzkow and Shapiro 2006, Xiang and Sarvary 2007, Gentzkow and Shapiro 2010, Gal-Or et al. 2012), and aggregation of voluntary contributions (e.g., Greenstein and Zhu 2012a, b). Our paper complements these studies by providing empirical evidence on the impact of ad revenue on the popularity and quality of media content.

This paper proceeds as follows. Section 2 provides details on the empirical setting. Section 3 describes the data. Section 4 presents empirical results. Section 5 explores how the program effects vary across participants and over time. Section 6 presents robustness checks. Section 7 concludes.

2. Background

Our empirical setting is a Chinese portal site, Sina.com, which was founded in 1998. Sina.com is the 16th most popular website in the world and receives more than 1.4 billion daily page views.⁶ It offers many services, including news, emailing, blogging, photo and video sharing, microblogging, and instant messaging. Our analysis focuses on its blogging service. Sina started hosting blogs for free in September 2005. It is a late mover in the blogging business, as the first Chinese blog-hosting site (<http://bokee.com>) appeared in 2002, with many other websites providing blogging services since then. For the first two years, the portal site did not place any ads on individual bloggers' content pages. Then, on September 11, 2007, Sina announced an ad-revenue-sharing program. The general public, including bloggers on Sina,

⁶ See <http://www.alexa.com/siteinfo/sina.com.cn> (accessed February 2011).

were not aware of this program before the announcement, as the company had kept the program's development strictly confidential to avoid competitive responses by its rivals.

From September 2007 to March 2008, the company conducted a test run of the ad-revenue-sharing program and invited about 3,000 bloggers to participate. About 1,000 bloggers joined the program during this period. In April 2008, the test period ended, and Sina started accepting applications from all bloggers. As indicated in the application guidelines, for an application to be successful, a blog must have a minimum of 700 page views per week for four consecutive weeks prior to the application date. Once approved, the site places ads on the blog, and the blogger receives 50% of the ad revenue generated by the traffic to her blog pages. To participate in the program, the blogger also needs to provide the site with basic personal information, such as her real name, home address, and bank information. Payments are deposited to participants' bank accounts on a monthly basis whenever the balance exceeds RMB 100 (equivalent to about USD 15).

On the advertiser side, Sina uses a pay-per-impression mechanism: At the beginning of each quarter, it announces a fixed price per thousand impressions, and advertisers decide on the number of impressions to purchase. The site started selling impressions in October 2007, one month after the program was announced. In November 2007, program participants started noticing ads on their blog pages. Bloggers cannot choose the specific ads to be displayed on their blogs, and they receive the same amount of money for each impression at a given time. At the beginning of the program, a blogger would make RMB 4.5 (equivalent to about USD 0.69) per 1,000 impressions generated by her blog. To avoid annoying viewers, ads are displayed as a small pop-up window in the lower right corner of the screen, and the pop-up window automatically disappears within two to three seconds after the Web page finishes loading.

The blog-hosting site offers an ideal setting for our study for multiple reasons. First, the site is the largest media website in China. When the ad-revenue-sharing program was introduced, blogs on this site generated about 0.3 billion page views per day, and on a single day, a popular blog post could generate more than 100,000 page views. Given the amount of attention the blogs receive, any systematic change in the content is economically important. Second, unlike many video-sharing sites, our target site offers unlimited storage space to content providers. As a result, bloggers have little incentive to delete their old posts, which allowed us to collect data on the complete history of blog posts from each blogger in our sample. Third, perhaps the most important advantage of

our empirical setting is the change in the site's business model: It initially did not compensate content providers, but suddenly introduced the program. The setting hence enables us to observe the change in content production for each participant and estimate the *influence* of ad-sponsored models on the content providers' incentives. Because not every blogger participated in the program, we can use those nonparticipants as a control group in our analysis. Finally, because the site uses a pay-per-impression mechanism on the advertiser side, we do not have to worry about differences among advertisements and the possibility that bloggers tweak their content to target different audiences to get a higher click-through rate.


3. Data

The company provided us with a data set that contains a complete list of all bloggers enrolled in the ad-revenue-sharing program as of January 31, 2009, and the dates each blogger joined the program. Each blogger is associated with a unique 10-digit ID. In total, our data include 5,140 participants, of which 4,200 joined the program after April 2008. We focus our analysis on the bloggers who joined after April 2008, because the motivation of invited participants during the test period could be different. Figure 1 shows the number of bloggers enrolled in the program in each month since April 2008. More than 1,700 bloggers enrolled in the program right after it became open to the general public, and a few hundred bloggers enrolled in the program every month thereafter. In our data, participants are enrolled in the program for 6.25 months on average.

To control for general trends in the content of all blogs, we create a control group by randomly generating another 50 million 10-digit ID numbers. Many of these IDs are mapped to users without blogs: They are users of the portal site's other services.⁷ For the bloggers, we first drop those who started blogging after January 2009 and then select those who write more than one blog post per month on average. We apply this last criterion to focus our analysis on active bloggers: Many bloggers create only one or two, often very short, posts right after setting up their blogs and never blog again. It seems that these bloggers want to experience what blogging is like but are not serious about producing any content. In the end, we obtain a list of 26,974 nonparticipants.

We download every blog post that each of the 4,200 participants and 26,974 nonparticipants had written on the site by January 31, 2009. For each blog post, we collect information on the date it was posted; the title;

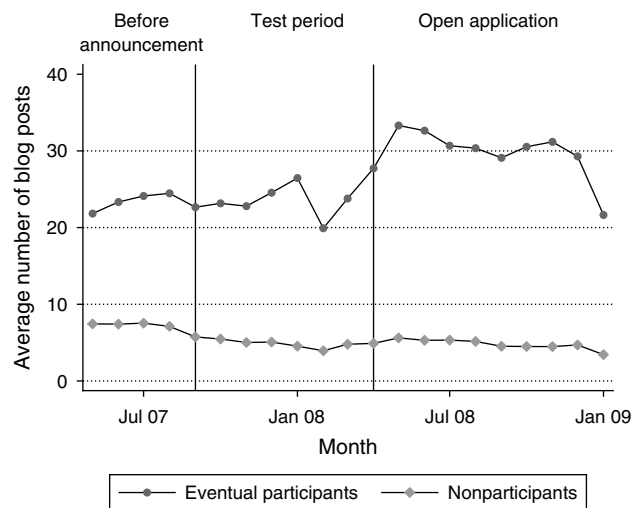
⁷ Sina does not disclose the total number of blogs that it hosts.

Figure 1 Number of New Participants in the Program in Each Month

the number of characters,⁸ pictures, and videos in the post; as well as the number of times the post had been read and bookmarked by its viewers. We also collect the tags supplied by the bloggers for each post. For each post, a blogger could supply multiple tags. For example, for a post on a basketball game, the blogger could use such tags as “Lakers,” “Rockets,” and “basketball.” Similarly, such tags as “travel,” “museum,” “spa,” and “diving” could be used on a post about travel.

We focus our analysis between May 2007 and January 2009. Because we rely on tags to identify popular topics for each month in China, it is critical that we aggregate tags from a sufficient number of posts in each month. After its launch in September 2005, the site’s blogging service experienced accelerated growth in 2006. In the first quarter of 2007, it became the largest blog-hosting site (as measured by the number of visitors) in China. In addition, the site did not introduce the tagging feature until April 2007. In May 2007, around 50% of the blog posts in our sample had tags, and this percentage increased to more than 90% in January 2009. For blog posts with no tags, we use post titles to generate tags.⁹

From May 2007 to January 2009, the bloggers in our data set composed 4,359,197 blog posts. Eventual participants in the program contributed 1,904,609 (43.7%) posts, and nonparticipants contributed 2,454,588 (56.3%) posts. Figure 2 shows the average number of blog posts in each month by participants and nonparticipants. We find that participants blogged much more frequently, on average, than nonparticipants. The number of blog posts

Figure 2 Average Number of Blog Posts in Each Month

Note. The two vertical lines indicate September 2007 and April 2008, respectively.

per month increased for participants over time, and the increase was most pronounced when the program became open to all bloggers. In contrast, the average number of blog posts for nonparticipants declined slightly over time. The pattern suggests that the program motivated participants to be more active. We also find that the average number of blog posts dropped significantly in February 2008 and January 2009. These drops coincide with the Chinese New Year holidays, which typically last 7 to 10 days.¹⁰

4. Empirical Analysis

4.1. Shift in Content Popularity

We first consider the popularity of the blog posts. A natural way to consider post popularity is to check whether or not the post is associated with a popular tag. To gauge interest in the tags, we define a tag’s popularity in a certain month by the total number of page views of blog posts containing the tag in that month.¹¹ The distribution of tag popularity is highly skewed. Take the tags in August 2007 as an example. In total, we have 47,273 tags in that month and, on average, each tag receives 17,938 page views. The most popular tag, “stock market,” receives more than 15 million page views, or 1.8% of the total page views of all tags. Other popular tags in that month include “gossip,” “stock index,” “history,” “Fan Bingbing,”¹²

¹⁰ The dates for the Chinese New Year in these two years are February 7, 2008, and January 26, 2009. The 7–10 holidays after the Chinese New Year’s eve are typically marked by family gatherings and visits to relatives and friends.


¹¹ It is important to analyze data on a monthly basis, because a tag’s popularity may change over time. The tag “Chinese New Year,” for example, is popular only at the beginning of a year.

¹² Fan Bingbing is a popular Chinese actress and singer.

⁸ In Chinese, characters form the basic unit of meaning. Most Chinese words are formed by two or three characters.

⁹ We use *Pan Gu Segment*, an open source software that divides Chinese sentences into a set of keywords, to generate these tags. The software is based on a library of more than 170,000 Chinese keywords and has been used by many commercial firms to build Chinese search engines.

Figure 3 Log-Log Plot of Tags' Page Views Against Their Ranks



“investment,” and “beauty.” The top 150 tags receive 39% of the total page views. In Figure 3, we plot the logarithm of the tags’ page views against the logarithm of their ranks. The log-log plot shows that the relationship between tags’ page views and tags’ ranks follows a power law when the logarithm of rank is less than 8 (i.e., when tag rank is less than 2,980). For tags with ranks above 2,980, the relationship becomes even more skewed, as a large number of tags receive very few page views: 3,561 tags receive fewer than ten page views. We rank all the tags based on their popularity and consider the top 150 tags in each month as popular tags. We choose this threshold to have a reasonable set of popular tags. Other thresholds, such as 100, 300, or the top 0.5% of all tags, provide similar results.

We then identify blog posts associated with the popular tags in each month as popular posts. On average, 23% of all blog posts are classified as popular posts, and these popular posts generate 63% of the page views.

Next, we compute the percentage of popular blog posts for each blogger i in each month t and denote this variable by $\%Popular_{it}$.¹³ Panel A of Table 1 shows the summary statistics of this variable, and Figure 4 shows how the value of this variable evolves over time for both participants and nonparticipants in our data. We find that, on average, participants are more likely to post popular content. The percentages of popular content for participants and nonparticipants diverged even more upon the launch of the program:

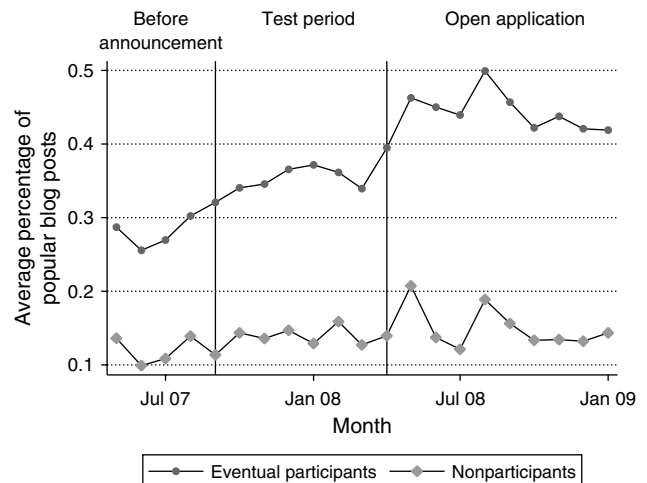
¹³ When a blogger does not contribute any post in a given month, there may be two interpretations. In one interpretation, the blogger does not contribute any popular content, and thus we can set $\%Popular_{it}$ to 0. In another interpretation, it means that a blogger’s content popularity cannot be measured, and we can set $\%Popular_{it}$ to missing. We report results based on the first interpretation in the paper. Results are qualitatively unchanged when we use the second interpretation.

Table 1 Summary Statistics of Outcome Variables and Instrumental Variables

Variable	Mean	Std. dev.	Min	Max
Panel A: Popularity and topic domains				
$\%Popular$	0.169	0.288	0	1
$\%Stock$	0.013	0.091	0	1
$\%Salacious$	0.031	0.117	0	1
$\%Celebrity$	0.016	0.085	0	1
Panel B: Quality				
$\%Bookmark$	0.000	0.002	0	0.518
Num_Chars	553.190	1,177.973	0	215,927.700
Num_Pics	0.884	2.312	0	106.667
Num_Videos	0.001	0.023	0	1.100
Panel C: Instrumental variables				
$Blogging_Age$	13.850	8.813	0	42
$Blogging_Freq$	8.010	35.936	0	8,150

whereas the percentage of popular content for non-participants stayed around 13%, the percentage for participants increased slightly upon the program’s announcement and then increased significantly when the program became open to all bloggers. We also notice month-specific effects on the percentages of popular content for both participants and nonparticipants. In May 2008, for example, the percentage of popular content for all bloggers had a sudden increase. This increase resulted from the Wenchuan earthquake, which occurred on May 12, 2008, in China’s Sichuan province and killed more than 69,000 people. The earthquake was the most discussed topic in all media in that month, and the tag “earthquake” was the most popular one in that month in our data. The percentage dropped back to its average level for the nonparticipating group right after May 2008 but remained at a high level for the participating group.

Figure 4 Average Percentage of Popular Blog Posts in Each Month



Note. The two vertical lines indicate September 2007 and April 2008, respectively.

Similarly, the increase in the percentages of popular content for both groups in August 2008 resulted from the opening of the Summer Olympic Games in Beijing.

We observe similar patterns in panel A of Table 2, where we compute the means and the gaps between the means of the outcome variables for these two groups during three time periods: before the announcement, during test period, and during the time of open application. We find that the average content popularity is higher for eventual participants than for nonparticipants in all periods. Moreover, for eventual participants, the average content popularity increases over time, whereas for nonparticipants, the average content popularity stays roughly the same. As a result, the gap between the two means widens over time. These model-free results suggest that eventual participants indeed increase the share of their posts on popular content, relative to nonparticipants, as the program takes effect.¹⁴

We now turn to the regression framework to detect the shift in the content popularity of program participants relative to that of nonparticipants. Regression frameworks allow us to take advantage of our panel data to control for time-specific and blogger-specific effects. We employ a difference-in-differences approach with the specification below:¹⁵

$$\%Popular_{it} = \beta_0 + \beta_1 EP_i + \beta_2 EP_i \times After_{it} + \sum_{j=2}^{21} \gamma_j MonthDummy_j + \epsilon_{it}, \quad (1)$$

where EP_i is a dummy that takes the value of 1 if blogger i is an eventual participant in the program, and 0 otherwise. The dummy captures the systematic difference between program participants and nonparticipants. The variable $After_{it}$ is 1 if blogger i is an eventual participant and has already enrolled in the program in month t , and 0 otherwise. The variable β_2 is our difference-in-differences estimator that captures the program's effect on content popularity for participants. We also include dummies for each month from May 2007 to January 2009 to control for changes in all bloggers' propensity to produce popular content. We

¹⁴ The significant differences between participants and nonparticipants before the program's announcement, as illustrated in Figure 2, Figure 4, and Table 2, suggest a potential selection problem, which will be addressed in our regression analysis and robustness checks.

¹⁵ Note that our specification uses a percentage as the dependent variable. Although this makes the coefficients more interpretable, the approach requires the assumption that the marginal effects are linear over the relevant part of the distribution. As a robustness check, we transform all dependent variables that are percentages in our analysis as logarithms of the odds (e.g., we use $\log(\%Popular_{it}/(1 - \%Popular_{it}))$ instead of $\%Popular_{it}$). Our results are qualitatively unchanged.

Table 2 Comparison of Means Between Eventual Participants and Nonparticipants

Variable	Period	Eventual participants	Non-participants	Gap
Panel A: Popularity and topic domains				
%Popular	Before announcement	0.279	0.121	0.158
	Test period	0.331	0.124	0.208
	Open application	0.412	0.125	0.287
%Stock	Before announcement	0.030	0.005	0.025
	Test period	0.035	0.004	0.031
	Open application	0.050	0.005	0.046
%Salacious	Before announcement	0.057	0.023	0.033
	Test period	0.072	0.022	0.050
	Open application	0.083	0.019	0.064
%Celebrity	Before announcement	0.039	0.012	0.027
	Test period	0.047	0.010	0.037
	Open application	0.053	0.009	0.044
Panel B: Quality				
%Bookmark	Before announcement	0.019	0.000	0.019
	Test period	0.024	0.000	0.024
	Open application	0.112	0.000	0.111
Num_Chars	Before announcement	1,002.874	534.374	468.500
	Test period	1,013.984	449.295	564.689
	Open application	1,088.085	419.333	668.752
Num_Pics	Before announcement	0.386	0.949	-0.563
	Test period	0.600	0.751	-0.151
	Open application	2.437	0.727	1.709
Num_Videos	Before announcement	0.000	0.000	0.000
	Test period	0.000	0.000	0.000
	Open application	0.003	0.002	0.001

cluster the error terms at the blogger level to account for autocorrelation in the data across bloggers and over time (Bertrand et al. 2004).

We need to address two problems in our specification. First, we need to account for a potential endogeneity problem, because those who apply to join the program are not randomly selected. In other words, some unobserved heterogeneity among bloggers in the error term may be correlated with their decisions to participate in the program, leading to biased estimates. These unobserved factors could lead to overestimation or underestimation of the program effect. For example, participants, in general, could like to blog about popular content more than nonparticipants. As a result, their blog posts are more popular, and it is easier for them to qualify for the program. If, at the same time, these participants are more (less) likely to respond to financial incentives, the program effect could be overestimated (underestimated).

We take two approaches to address this problem. First, we introduce blogger-level fixed effects to control for time-invariant, unobserved blogger characteristics. Fixed effects allow us to focus on changes in content over time for any given blogger, rather than the absolute levels. Fixed effects, however, do not control for time-variant unobservables that may

be correlated with the decision to participate in the program. These time-variant unobservables could lead, for example, to different trends over time for participants and nonparticipants. Given this concern, we also construct two instrumental variables by taking advantage of the minimum number of page views required to participate in the program. Valid instruments need to correlate with the decision to participate and affect the dependent variable only through the participation decision. Our first instrument is the number of months since a blogger's first post (*Blogging_Age_{it}*). Our second instrument is the average number of posts per month for the blogger in the past (*Blogging_Freq_{it}*). The idea is that the longer and more frequently a blogger has been blogging, the more likely it is she would have cultivated an audience base generating more than 700 page views per week, which would make her eligible for the program. At the same time, the two variables are unlikely to be directly correlated with the content of the posts. The two instruments, therefore, can help us control for unobservable blogger characteristics that may affect program enrollment and content popularity at the same time. Panel C of Table 1 provides summary statistics, which suggest great variation for each instrument.

Second, bloggers may tailor their content to improve their chances of getting approved for the program. Hence the program's impact may take effect before they join it. More generally, because it takes time to increase the popularity of one's blog, some bloggers may choose to shift toward popular content right after the program's announcement and wait until their page views meet the requirement before applying to the program. As a result, our regressions could underestimate the program's impact. Indeed, the increase in content popularity for the eventual participants right after the program's announcement, as shown in Figure 4, suggests that such effects

may exist. Furthermore, some nonbloggers might be incentivized by ad revenue to start blogging and focus on popular content. As a result, bloggers with start dates after September 2007 could be systematically different from those who joined before the program's announcement. Finally, some nonparticipants may also be incentivized by ad revenue. Such bloggers may have increased their content popularity but still failed to meet the program requirement.

To minimize these effects, we take September 2007 as the breakpoint for all participants and include only those bloggers who started blogging on the site before that month. To ensure that the announcement of the program is truly exogenous, we search baidu.com, the top search engine in China, for news related to the ad-revenue-sharing program. All the news is dated on or after the day of the program's announcement. We also search the text of all blog posts in our data set, because bloggers on the site are likely to discuss this program once they become aware of it. All posts that mention this program are also dated after the program's announcement.

Table 3 reports our regression results. In the first three models, we use bloggers' enrollment dates as breakpoints. Model (1) reports the results based on ordinary least squares (OLS) regression. On average, a participating blogger's percentage of popular posts before she joins the program is higher than that of a nonparticipant by 22.0 percentage points. This percentage increases by an additional 7.1 percentage points after she joins the program. Model (2) reports the results with fixed effects (FE), which are similar to those in model (1). The variable, EP_i , drops from the regression, because its value does not vary over time. In both models (1) and (2), the coefficients of the interaction variable, $EP_i \times After_{it}$, reflect the average effect of the program on the treated group. Model (3) reports the results with both fixed effects and instrumental variables using two-stage least-squares (2SLS).

Table 3 The Impact of Revenue Sharing on Content Popularity

Model:	Enrollment dates as breakpoints			September 2007 as the breakpoint		
	(1)	(2)	(3)	(4)	(5)	(6)
Dep. var.:	%Popular	%Popular	%Popular	%Popular	%Popular	%Popular
<i>EP</i>	0.220*** (0.004)			0.158*** (0.005)		
<i>EP</i> × <i>After</i>	0.071*** (0.004)	0.070*** (0.004)	0.078*** (0.004)	0.096*** (0.004)	0.093*** (0.004)	0.130*** (0.017)
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	544,209	544,209	544,209	448,973	448,973	448,973
Adjusted R^2	0.097	0.016	0.016	0.093	0.017	0.016
Number of IDs		31,174	31,174		21,792	21,792
Specification	OLS	FE	FE/2SLS	OLS	FE	FE/2SLS

Note. Heteroskedasticity-adjusted standard errors are in parentheses.

***Significant at 1%.

The results in model (3) show that the program's effect becomes stronger after we correct for endogeneity with the two instrumental variables.

In the next three models, we repeat the analysis in models (1)–(3) using September 2007 as the breakpoint for all participants. We redefine the dummy variable $After_{it}$ to be 1 if month t is on or after September 2007, and 0 otherwise. We find that the systematic difference in the percentages of popular posts between participants and nonparticipants becomes smaller (15.8%). As expected, the program's impact is more pronounced: An eventual participant's percentage of popular posts increases by as much as 13.0 percentage points after the program's announcement, which is equivalent to an increase of 65%.¹⁶ The p -value of the overidentification test statistic is 0.26. The results are consistent with our conjecture that many participants started providing popular content in preparation for enrolling in the program after the program's announcement. In the rest of the analysis, we use September 2007 as the breakpoint for all participants.

Table 4 provides first-stage estimation results for model (6) in Table 3 to illustrate the instrumental variables' relevance. In models (1) and (2), we include the two instruments, $Blogging_Age_{it}$ and $Blogging_Freq_{it}$, separately. In model (3), we include both of them together. We find that both instruments are highly correlated with becoming a program participant, and these results are statistically significant at the 1% level in all three models. The overall Wald chi-squared test or F -test for the instruments in each model is also highly significant.

4.2. Shift in Content Topics

We next examine the shift in the topics of participants' blog posts after the program takes effect. After speaking with several frequent bloggers in China, we decided to focus on the three most-mentioned topics: the stock market, salacious content, and celebrities. China's stock market started in early 1990 and has been notorious for its fluctuations (see, e.g., Einhorn 2010). To maximize their returns, many people regularly read blog posts related to the stock market for free opinions and recommendations. Hence, blogging about the stock market is likely to be an effective strategy in attracting traffic. The other two topics, salacious content and celebrities, are universally considered as hot topics.¹⁷

Two research assistants independently examined the top 150 tags in each month and classified each tag

¹⁶ Before September 2007, 20% of all posts are classified as popular posts.

¹⁷ Although it is illegal in China to post images or videos that contain nudity or text containing explicit descriptions of sexual acts, bloggers can include images or text that are sexually suggestive.

Table 4 First-Stage Regressions and Instrument Relevance

Model:	(1)	(2)	(3)
Dep. var.:	$EP \times After$	$EP \times After$	$EP \times After$
$Blogging_Age \times After$	0.080*** (0.003)		0.052*** (0.005)
$Blogging_Freq \times After$		0.009*** (0.001)	0.008*** (0.001)
Month dummies	Yes	Yes	Yes
Observations	448,973	448,973	448,973
Adjusted R^2	0.144	0.204	0.209
Number of IDs	21,792	21,792	21,792
F -statistics	154.4	173.1	167.1
Specification	FE	FE	FE

Note. Heteroskedasticity-adjusted standard errors are in parentheses.

***Significant at 1%.

into one of four domains: the stock market, salacious content, celebrities, and others. The classifications were highly consistent, and the few discrepancies were resolved by a meeting of the research assistants. On average, in each month, 12% of the popular tags (e.g., "stock market index" and "stock recommendation") are classified as being related to the stock market, 13% (e.g., "nude photo scandal" and names of Japanese adult-video idols) are classified as being salacious content, 9% (e.g., "celebrity gossip" and names of the celebrities) are classified as being related to celebrities, and 66% (e.g., "earthquake" and "Chinese New Year") are classified as others.

We identify all blog posts with tags that fall into one or more of the first three domains. A blog post may be classified into multiple domains. For example, a post on a nude photo scandal involving celebrities is classified as both salacious and related to celebrities. In the end, 6.2%, 6.5%, and 3.0% of all the posts are classified as being related to the stock market, salacious content, and celebrities, respectively. We then compute the percentage of posts in each of these three domains for each blogger in each month and denote these variables by $\%Stock_{it}$, $\%Salacious_{it}$, and $\%Celebrity_{it}$. Panel A of Table 1 provides summary statistics of these three variables. Among the three topic domains, we have a greater proportion of salacious content than content related to stock markets and celebrities.¹⁸

Similar to the previous analysis on popularity, we provide a comparison of the mean levels of the outcome variables across the three different periods between participants and nonparticipants in panel A of Table 2. As before, we find that for all the three topic domains and in all three periods, participants

¹⁸ Note that when we average the percentages of posts in each content domain across bloggers, the means are lower than those on the post level, which is because many bloggers do not write posts in any of these three domains in a given month.

Table 5 The Impact of Revenue Sharing on Content Topics

Model:	(1)	(2)	(3)	(4)	(5)	(6)
Dep. var.:	%Stock	%Stock	%Salacious	%Salacious	%Celebrity	%Celebrity
<i>EP</i> × <i>After</i>	0.012*** (0.002)	0.024** (0.012)	0.024*** (0.002)	0.026*** (0.007)	0.014*** (0.002)	0.016* (0.009)
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations	448,973	448,973	448,973	448,973	448,973	448,973
Adjusted <i>R</i> ²	0.002	0.001	0.003	0.003	0.001	0.001
Number of IDs	21,792	21,792	21,792	21,792	21,792	21,792
Specification	FE	FE/2SLS	FE	FE/2SLS	FE	FE/2SLS

Note. Heteroskedasticity-adjusted standard errors are in parentheses.
 *Significant at 10%; **significant at 5%; ***significant at 1%.

allocate a higher fraction of their posts to each domain than nonparticipants. Furthermore, as the program moves into the test period and then the period of open application, the gap between the two groups of bloggers becomes larger, suggesting that the program indeed motivates participants to increase the fraction of their posts in these three domains, relative to nonparticipants.

To detect the shift in content topics in a regression framework, we replace the dependent variable, %*Popular*_{*it*}, with each of these three variables and repeat the difference-in-differences analysis. Table 5 reports the regression results. For each of the three percentages above, we report the results with fixed effects, as well as the results with both fixed effects and instrumental variables. The results demonstrate a significant shift of content toward all three domains. In total, the blog posts of participants in these three domains increase by 6.6% percentage points (based on the specification with both fixed effects and instrumental variables) relative to nonparticipants.

4.3. Shift in Content Quality

Finally, we consider the impact of ad revenue on content quality. Although high-quality content attracts eyeballs, it is unclear whether participating bloggers will expend more effort into blogging content in which they may not be intrinsically interested. To identify the extent of such effects, we develop several measures on post quality. For each post, we first compute the percentage of viewers who bookmark the post as one of their favorites.¹⁹ Bloggers cannot tell who bookmarked their posts and hence cannot reciprocate by visiting the blogs of their patrons. Therefore, the only benefit of bookmarking a post is the convenience of reaccessing it in the future. For each blogger in each month, we compute the average of the percentages for all her posts and denote this measure as

%*Bookmark*_{*it*}, which reflects the satisfaction that readers derive from reading the posts.

We also measure each blogger’s effort by the average number of characters, pictures, and video clips in her posts. For any given blogger, the more effort she devotes to writing, the more likely it is that the blog post is of higher quality. In general, a higher number of characters requires more effort from the blogger and allows the post to go deeper into the focal topic, thus increasing the quality of the post. Pictures and videos often make a post more attractive, although they may also require more effort from the blogger. We denote these measures as *Num_Chars*_{*it*}, *Num_Pics*_{*it*}, and *Num_Videos*_{*it*}, respectively. As the summary statistics in panel C of Table 1 suggest, on average, very few visitors bookmark blog posts, although some bloggers’ posts are bookmarked as frequently as 51.8% of the time. On average, a blog post has 0.884 pictures and 0.001 video clips. The low value for *Num_Videos*_{*it*} is perhaps due to the fact that preparing and uploading videos takes effort and time, and many Internet users in China did not have fast connections during the study period.

In panel B of Table 2, we compare the mean levels of all quality measures across the same three periods as before. We find that for all measures except the number of pictures, the average quality of participants’ posts always exceeds that of nonparticipants’ posts. For the number of pictures, while participants initially have fewer pictures in their posts relative to nonparticipants, once the program moves into the period of open application, the relationship reverses, and participants have more pictures per post. Consistent with our hypothesis that the program motivates participants to increase the quality of their posts, the gap between the average quality of participants’ posts and that of nonparticipants’ posts, for all the quality measures, keeps increasing as the program gets tested and becomes open to all bloggers.

To provide further evidence on the program’s effect on quality, we apply the same difference-in-differences approach using each of the four quality measures mentioned above as the dependent variable.

¹⁹ For each blog post, the site provides a button that any reader with an account on the site can use to bookmark the blog post, which puts the post in her personal collection.

Table 6 The Impact of Revenue Sharing on Content Quality

Model: Dep. var.:	(1) %Bookmark	(2) %Bookmark	(3) Num_Chars	(4) Num_Chars	(5) Num_Pics	(6) Num_Pics	(7) Num_Videos	(8) Num_Videos
<i>EP</i> × <i>After</i>	0.054*** (0.003)	0.094*** (0.010)	0.854*** (0.032)	1.181*** (0.146)	1.392*** (0.039)	1.973*** (0.138)	0.001*** (0.000)	0.001* (0.001)
% <i>Popular</i>	0.009*** (0.001)	0.008*** (0.001)	3.138*** (0.021)	3.130*** (0.021)	0.907*** (0.024)	0.893*** (0.024)	0.002*** (0.000)	0.002*** (0.000)
% <i>Stock</i>	0.028 (0.021)	0.026 (0.021)	1.638*** (0.088)	1.619*** (0.088)	−0.706*** (0.062)	−0.740*** (0.064)	−0.002*** (0.000)	−0.002*** (0.000)
% <i>Salacious</i>	0.003 (0.003)	0.002 (0.003)	1.462*** (0.044)	1.454*** (0.044)	0.330*** (0.058)	0.316*** (0.058)	−0.001 (0.001)	−0.001 (0.001)
% <i>Celebrity</i>	0.000 (0.003)	−0.001 (0.003)	1.336*** (0.063)	1.326*** (0.064)	1.018*** (0.090)	1.001*** (0.090)	0.004*** (0.001)	0.004*** (0.001)
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	448,973	448,973	448,973	448,973	448,973	448,973	448,973	448,973
<i>R</i> ²	0.006	0.005	0.161	0.161	0.031	0.030	0.003	0.003
Number of IDs	21,792	21,792	21,792	21,792	21,792	21,792	21,792	21,792
Specification	FE	FE/2SLS	FE	FE/2SLS	FE	FE/2SLS	FE	FE/2SLS

Notes. We multiply %*Bookmark* by 100 for the ease of displaying coefficients. We take the logarithm of *Num_Chars* to minimize the impact of outliers. Heteroskedasticity-adjusted standard errors are in parentheses.

*Significant at 10%; ***significant at 1%.

Note that the shift in content quality might also reflect a shift in content topics. For example, readers may be more likely to bookmark popular posts, and hence %*Bookmark*_{*it*} may increase as the content becomes more popular. Similarly, as a blogger increases her content on celebrities, it may be easier for her to find relevant pictures and videos, and she might include more pictures and video clips in her blog posts. Therefore, in our regressions, we include content popularity and topics as controls. We also take the logarithm of the average number of characters to minimize the effect of outliers. Table 6 reports the results. We find significant and positive program effects on all the quality measures.

In addition, in a supplementary analysis, we first measure the quality of blog posts that fall into each of the three topic domains for each blogger in each month and then examine the quality shift within each topic domain. Except for the case of *Num-Videos*_{*it*}, where sometimes the coefficients are positive but insignificant, we find that the program increases content quality in all three domains.

5. Program Effects Across Participants and Over Time

Our difference-in-differences approach so far has examined only the average program effects on bloggers' behavior. One may expect, however, that the program effects to vary across different bloggers. In particular, as bloggers need to achieve 700 page views for four consecutive weeks to be eligible for the program, we expect bloggers with page views below the threshold to shift their blogging behavior more than those who have already met the requirement.

We therefore partition the eventual participants into two groups, those with more than 3,000 page views in the month before the program's announcement (i.e., in August 2007) and those who do not have this number of page views.²⁰

We create two dummy variables, *Below_Threshold*_{*i*}, which is 1 if a blogger is an eventual participant and has fewer than 3,000 page views in August 2007, and *Above_Threshold*_{*i*}, which is 1 if a blogger is an eventual participant and has 3,000 page views or more in August 2007. Of the 2,722 eventual participants who joined the blogging site before the announcement of the program, 1,984 (73%) have more than 3,000 page views in August 2007. We alter our difference-in-differences analysis by replacing the interaction term, *EP*_{*i*} × *After*_{*it*}, with two new interactions, *Below_Threshold*_{*i*} × *After*_{*it*} and *Above_Threshold*_{*i*} × *After*_{*it*}. Table 7 reports the regression results. The main effects of *Below_Threshold*_{*i*} and *Above_Threshold*_{*i*} are absorbed by the blogger fixed effects.

Consistent with our hypothesis, those eventual participants who have not met the program requirement before the announcement shift their topics and increase their content quality more than other participants, except for *Num_Pics*. Nonetheless, compared to

²⁰ Seven hundred page views per week is roughly equivalent to 3,000 page views per month. Note that we observe only the number of page views each blog post has received at the time we collect the data and do not observe the actual number of page views in each month. Because the popularity of blog posts drops quickly after they are posted, as an approximation, we count the total number of page views that blog posts created in August 2007 have received and use this number as a proxy for the total number of page views for each blogger in that month.

Table 7 Program Effects on Different Types of Participants

Panel A: Popularity and topic domains								
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. var.:	%Popular	%Popular	%Stock	%Stock	%Salacious	%Salacious	%Celebrity	%Celebrity
<i>Below_Threshold</i> × <i>After</i>	0.148*** (0.009)	0.153*** (0.009)	0.019*** (0.004)	0.024*** (0.004)	0.041*** (0.004)	0.042*** (0.004)	0.025*** (0.003)	0.025*** (0.003)
<i>Above_Threshold</i> × <i>After</i>	0.073*** (0.005)	0.129*** (0.018)	0.009*** (0.002)	0.021* (0.011)	0.018*** (0.002)	0.027*** (0.008)	0.010*** (0.002)	0.016 (0.010)
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	448,973	448,973	448,973	448,973	448,973	448,973	448,973	448,973
Adjusted <i>R</i> ²	0.017	0.016	0.002	0.001	0.003	0.003	0.002	0.001
Number of IDs	21,792	21,792	21,792	21,792	21,792	21,792	21,792	21,792
Specification	FE	FE/2SLS	FE	FE/2SLS	FE	FE/2SLS	FE	FE/2SLS
Panel B: Quality								
Model:	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Dep. var.:	%Bookmark	%Bookmark	Num_Chars	Num_Chars	Num_Pics	Num_Pics	Num_Videos	Num_Videos
<i>Below_Threshold</i> × <i>After</i>	0.056*** (0.003)	0.090*** (0.011)	1.496*** (0.078)	1.545*** (0.080)	0.902*** (0.056)	0.955*** (0.059)	0.001* (0.001)	0.001* (0.000)
<i>Above_Threshold</i> × <i>After</i>	0.048*** (0.008)	0.051*** (0.009)	0.617*** (0.030)	1.086*** (0.154)	1.572*** (0.049)	2.087*** (0.149)	0.000 (0.000)	0.000 (0.001)
%Popular	0.009*** (0.001)	0.008*** (0.001)	3.134*** (0.021)	3.127*** (0.021)	0.910*** (0.024)	0.902*** (0.024)	0.002*** (0.000)	0.002*** (0.000)
%Stock	0.028 (0.021)	0.027 (0.022)	1.628*** (0.087)	1.612*** (0.087)	-0.699*** (0.062)	-0.716*** (0.062)	-0.002*** (0.000)	-0.002*** (0.000)
%Salacious	0.003 (0.003)	0.002 (0.003)	1.457*** (0.044)	1.451*** (0.044)	0.334*** (0.058)	0.327*** (0.058)	-0.001 (0.001)	-0.001 (0.001)
%Celebrity	0.000 (0.003)	-0.001 (0.003)	1.329*** (0.063)	1.322*** (0.063)	1.023*** (0.090)	1.015*** (0.090)	0.004*** (0.001)	0.004*** (0.001)
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	448,973	448,973	448,973	448,973	448,973	448,973	448,973	448,973
Adjusted <i>R</i> ²	0.006	0.005	0.161	0.161	0.032	0.031	0.003	0.003
Number of IDs	21,792	21,792	21,792	21,792	21,792	21,792	21,792	21,792
Specification	FE	FE/2SLS	FE	FE/2SLS	FE	FE/2SLS	FE	FE/2SLS

Note. Heteroskedasticity-adjusted standard errors are in parentheses.
 *Significant at 10%; ***significant at 1%.

nonparticipants, the program effects on eventual participants above the threshold are significant, except for %Celebrity and Num_Videos. These results are consistent with our intuition that the program effects on content topics and quality would be most pronounced for moderately popular bloggers, and provide further support for the causal effects of the program.

In addition to varying across different types of bloggers, the program effects may vary over time. In particular, we are interested in the program's long-term impact on bloggers' behavior. The impact could diminish over time for multiple reasons. For example, participants might change their behavior mainly for the purpose of becoming eligible for the program, and upon enrolling in the program, they might lose the incentive to maintain the same blogging patterns. It is also possible that as more participants compete for a share of eyeballs for popular content, participants feel a need to differentiate their content to target different audience groups. If these intuitions are true, the

program effects would be short-lived. On the other hand, it may take time for participants to learn how to make their blog posts more attractive. As they gain more experience, they may shift their behavior even further. In this case, the shift may increase over time, and thus the program effects may become more pronounced over time.

Our current specification only detects shifts in the levels. To examine the slope of the shifts after a participant enrolls in the program, we create a new variable, *Num_Months_Enrolled_{it}*, which measures the number of months since a participant has enrolled in the program. We then add an interaction term *EP_i × Num_Months_Enrolled_{it}* to Equation (1) to detect the slope of the shifts after enrollment. Table 8 provides the regression results. We find no evidence that the program's impact on participants diminishes after their enrollment. To the contrary, for example, the shift toward the stock market actually increases with the number of months in the program. We also find

Table 8 Persistence of Program Effects After Enrollment

Panel A: Popularity and topic domains								
Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dep. var.:	%Popular	%Popular	%Stock	%Stock	%Salacious	%Salacious	%Celebrity	%Celebrity
<i>EP</i> × <i>After</i>	0.090*** (0.004)	0.111*** (0.013)	0.009*** (0.001)	0.016** (0.008)	0.025*** (0.002)	0.024*** (0.005)	0.014*** (0.002)	0.011*** (0.004)
<i>EP</i> × <i>Num_Months_Enrolled</i>	0.002 (0.001)	0.001 (0.001)	0.002*** (0.000)	0.001*** (0.000)	−0.000 (0.000)	−0.000 (0.000)	0.000 (0.000)	0.000 (0.000)
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	448,973	448,973	448,973	448,973	448,973	448,973	448,973	448,973
Adjusted <i>R</i> ²	0.017	0.016	0.003	0.003	0.003	0.003	0.001	0.001
Number of IDs	21,792	21,792	21,792	21,792	21,792	21,792	21,792	21,792
Specification	FE	FE/2SLS	FE	FE/2SLS	FE	FE/2SLS	FE	FE/2SLS
Panel B: Quality								
Model:	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)
Dep. var.:	%Bookmark	%Bookmark	Num_Chars	Num_Chars	Num_Pics	Num_Pics	Num_Videos	Num_Videos
<i>EP</i> × <i>After</i>	0.035*** (0.003)	0.048*** (0.008)	0.807*** (0.032)	1.192*** (0.130)	1.001*** (0.033)	1.213*** (0.101)	0.000 (0.000)	0.000 (0.001)
<i>EP</i> × <i>Num_Months_Enrolled</i>	0.012*** (0.001)	0.011*** (0.001)	0.029*** (0.006)	0.013* (0.008)	0.241*** (0.011)	0.230*** (0.012)	0.000*** (0.000)	0.000** (0.000)
%Popular	0.009*** (0.001)	0.009*** (0.001)	3.137*** (0.021)	3.129*** (0.021)	0.904*** (0.024)	0.900*** (0.024)	0.002*** (0.000)	0.002*** (0.000)
%Stock	0.024 (0.021)	0.023 (0.021)	1.627*** (0.088)	1.612*** (0.088)	−0.799*** (0.064)	−0.806*** (0.064)	−0.002*** (0.000)	−0.002*** (0.000)
%Salacious	0.003 (0.003)	0.003 (0.003)	1.463*** (0.044)	1.454*** (0.044)	0.338*** (0.058)	0.333*** (0.058)	−0.001 (0.001)	−0.001 (0.001)
%Celebrity	−0.000 (0.003)	−0.001 (0.003)	1.335*** (0.064)	1.325*** (0.064)	1.012*** (0.090)	1.007*** (0.090)	0.004*** (0.001)	0.004*** (0.001)
Month dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	448,973	448,973	448,973	448,973	448,973	448,973	448,973	448,973
Adjusted <i>R</i> ²	0.011	0.011	0.161	0.161	0.041	0.041	0.003	0.003
Number of IDs	21,792	21,792	21,792	21,792	21,792	21,792	21,792	21,792
Specification	FE	FE/2SLS	FE	FE/2SLS	FE	FE/2SLS	FE	FE/2SLS

Note. Heteroskedasticity-adjusted standard errors are in parentheses.

*Significant at 10%; **significant at 5%; ***significant at 1%.

strong evidence that the shift toward high-quality content increases with the number of months in the program. Although our data could cover only a limited time frame, these results suggest that the program effects persist after participants' enrollment.

6. Exploring Robustness

6.1. Selection Problem

The instrumental variables approach that we have taken to address the selection problem relies on validity of our instruments. Because we are unable to empirically evaluate the exogeneity of the instruments, we also undertake three other approaches: propensity-score matching (DiPrete and Gangl 2004, Leuven and Sianesi 2003), Rosenbaum bounds (Rosenbaum 2002), and the AET-SSS approach (an approach first developed by Altonji et al. (2002, 2005)

and recently extended by Sen et al. (2011)), to evaluate the impact of potential selection effects.²¹

The instrumental variables approach and these three alternative approaches rely on different assumptions. The instrumental variables approach relies on exogenous variables to purge the effects of unobservables on the decision to participate. Propensity-score

²¹ In additional robustness checks, we repeat the difference-in-differences analysis after dropping nonparticipants whose monthly page views never exceed 3,000 after the program's announcement. Of 19,070 nonparticipants who start blogging before the program's announcement, 12,733 (66.8%) never meet the program requirement and thus are dropped from our control group. Alternatively, we repeat the difference-in-differences analysis after restricting both participants and nonparticipants to those bloggers who qualify for the program before its announcement. Among all bloggers, 2,348 participants and 6,142 nonparticipants meet the program requirement before its announcement and are used in the analysis. The program effects remain significant when we use these two different samples.

Table 9 Propensity-Matched Results

Model: Matching var.:	(1) $\Delta\% \text{Popular}$	(2) $\Delta\% \text{Stock}$	(3) $\Delta\% \text{Salacious}$	(4) $\Delta\% \text{Celebrity}$	(5) $\Delta\% \text{Bookmark}$	(6) $\Delta \text{Num_Chars}$	(7) $\Delta \text{Num_Pics}$	(8) $\Delta \text{Num_Videos}$
No matching	0.095*** (0.003)	0.012*** (0.001)	0.025*** (0.001)	0.014*** (0.001)	0.038*** (0.002)	0.557*** (0.035)	1.496*** (0.028)	0.001*** (0.000)
$\text{Page_Views}_i^{\text{pre}}$	0.099*** (0.005)	0.014*** (0.002)	0.028*** (0.002)	0.015*** (0.002)	0.038*** (0.005)	0.629*** (0.043)	1.604*** (0.050)	0.001** (0.000)
All variables	0.158*** (0.006)	0.018*** (0.002)	0.032*** (0.003)	0.022*** (0.002)	0.038*** (0.005)	0.292*** (0.045)	1.126*** (0.043)	0.001*** (0.000)

Notes. Matching is implemented using the single nearest-neighbor method. Standard errors are in parentheses.

Significant at 5%, *significant at 1%.

matching corrects for selection bias by matching participants with nonparticipants based on observables. Rosenbaum bounds characterize the magnitude of selection on unobservables, which is required to nullify the program effects identified by propensity-score matching. Finally, the AET-SSS approach evaluates the effects of selection on unobservables by assuming that the selection on unobservables is the same as that on observables. As a result, these approaches provide complementary information about underlying causal relationships.

As all three alternative approaches use one pre-event and one postevent observation for each blogger, we first collapse each of our outcome variables, Y_{it} , where $Y \in \{\% \text{Popular}, \% \text{Stock}, \% \text{Salacious}, \% \text{Celebrity}, \% \text{Bookmark}, \text{Num_Chars}, \text{Num_Pics}, \text{Num_Videos}\}$, into simple averages before and after the program's announcement for each blogger i , and denote these averages as Y_i^{pre} and Y_i^{post} . We then take the difference between the averages and generate new variables, ΔY_{it} , where $\Delta Y_{it} = Y_{it}^{\text{post}} - Y_{it}^{\text{pre}}$. We also create additional control variables such as $\text{Blogging_Freq}_i^{\text{pre}}$ and $\text{Page_Views}_i^{\text{pre}}$, which are averages of monthly blogging frequency and page views for blogger i before the program's announcement, and $\text{Blogging_Age}_i^{\text{pre}}$, which is the number of months blogger i has been blogging when the program is announced.

We first control for potential selection on observables using propensity-score matching (DiPrete and Gangl 2004, Leuven and Sianesi 2003). The propensity score is operationalized as the predicted probability of becoming a participant estimated from a logistic regression of EP_i on the characteristics of blogger i before the announcement of the program, including the averages of the outcome variables (Y_i^{pre} , $\text{Blogging_Freq}_i^{\text{pre}}$, $\text{Page_Views}_i^{\text{pre}}$ and $\text{Blogging_Age}_i^{\text{pre}}$). Each program participant is then compared to nonparticipants with similar propensity scores. Table 9 reports the results from propensity-score matching.

In the first row, we report the estimated changes without any propensity matching. The estimated program effects are similar to the fixed-effects results

reported in Tables 3, 5, and 6. In the second row, we report the estimates when we use $\text{Page_Views}_i^{\text{pre}}$ as the only matching variable, because it is clear from the program requirement that the number of page views is a critical criterion. The estimated program effects are similar to those without propensity matching. In the third row, we report the results using all observed characteristics of the bloggers above as matching variables. We find that the estimated program effects are greater than those in the first row for content popularity and content topics, but remain the same or become smaller for content quality. In all cases, the program effects remain significantly positive.

Propensity-score matching operates on a strong assumption that observable characteristics fully account for the selection of bloggers into the treatment and control conditions. We next conduct a sensitivity analysis by estimating Rosenbaum bounds (Rosenbaum 2002, Leuven and Sianesi 2003), which measure how strongly an unobservable must influence the selection process to completely nullify the causal effects identified in the propensity-matching analysis. We find that an unobservable variable would have to change the odds of selection into the program by at least 50% for the program effect to disappear for any outcome variable.²² Hence, our overall conclusion that ad revenue leads to more popular and higher-quality content is unlikely to be overturned by unobserved selection effects.

As a third test, we employ the AET-SSS approach. This approach relies on the notion that *selection on observables is the same as selection on unobservables*, and examines the treatment effect under the assumption that unobservables have the same effects on the outcome as observables. In our setting, because blogging behavior before the announcement is a good predictor of future blogging behavior, the explanatory power of observed variables is likely to exceed the explanatory power of unobservables. The AET-SSS approach is hence likely to produce conservative estimates of the program effects.

²² Note that the threshold we find is on the same order of magnitude as the Rosenbaum bound results reported by DiPrete and Gangl (2004) and Sen et al. (2011).

Following the AET-SSS approach, we specify the outcome and selection equations as follows (we omit the subscript i for simplicity):

$$\begin{aligned}\Delta Y &= \alpha \cdot EP + X' \gamma + \varepsilon, \\ EP &= 1(X' \beta + u > 0), \\ \text{where } \begin{pmatrix} \varepsilon \\ u \end{pmatrix} &\sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2 & \rho\sigma \\ \rho\sigma & 1 \end{pmatrix} \right).\end{aligned}$$

Similar to the propensity-score matching test, we include the bloggers' characteristics prior to the program's announcement— Y_i^{pre} , $\text{Blogging_Freq}_i^{\text{pre}}$, $\text{Page_Views}_i^{\text{pre}}$, and $\text{Blogging_Age}_i^{\text{pre}}$ —in vector X . The indicator function $1(\cdot)$ is 1 when $X' \beta + u > 0$ and 0 otherwise.

In unreported estimations, we find that except in the cases of $\Delta \text{Num_Pics}$ and $\Delta \text{Num_Videos}$, we underestimate the program effects in our OLS regression. In the case of $\Delta \text{Num_Videos}$, the estimated bias is insignificant. In the case of $\Delta \text{Num_Pics}$, we overestimate the program effect, but even after taking the selection effect into account, the program effect is still positive.

In sum, with all three alternative approaches, propensity-score matching, Rosenbaum bounds, and the AET-SSS approach, we find positive and significant program effects. Similar to our findings with the instrumental variables approach, for most outcome variables, we find that the OLS regressions are likely to underestimate the program effects.

6.2. Pretrend Analysis

The blog-hosting site's sudden announcement of the program gives us an opportunity to employ a useful falsification test. If our assumption on the orthogonality between blogger-specific unobservables and their decisions to participate in the program is violated after employing blogger-level fixed effects and instrumental variables, our data should produce diverging patterns for participants and nonparticipants in periods even before the program's announcement. We thus regress each outcome variable in Tables 3, 5, and 6 on dummies for each month between May 2007 and August 2007 and their interactions with EP_i . In unreported regressions, we find little evidence that the trends for participants and nonparticipants differ significantly prior to the program's announcement in fixed-effects specifications with or without instrumental variables. The absence of such false positives increases our confidence that the observed shifts for participants are caused by the program.

6.3. Definition of Popular Tags

We are concerned about our approach to defining popular tags, because tag popularity can be affected by quality and topic choices. In particular, even

though a greater fraction of our data comes from non-participants, the shift in participants' blogging behavior may still affect tag popularity and consequently affect our analysis for content popularity and topics. In unreported regressions, we conduct two robustness checks. First, we use blog posts contributed by only nonparticipants and their page views to determine tag popularity. Second, we use blog posts contributed by both participants and nonparticipants in the month right before the program's announcement (i.e., August 2007) to define tag popularity and then use this definition for all time periods. In both sets of regressions, our results are qualitatively the same.

We are also concerned about potential strategic manipulation of tags by program participants. For example, participants may attach more tags to each of their posts to attract readers after the program's announcement. They may also attach popular tags even when these tags do not accurately describe their posts. Such strategic manipulation could have contributed to our finding that program participants are more likely to produce popular content or content in one of the three topic domains. To address this concern, we generate tags for all blog posts based on the text in each post.²³ For each blog post, we tokenize the text into individual words and use the four most frequently mentioned nouns as the tags.²⁴ We follow the same procedure to identify popular posts and the domains to which each popular post belongs, and then repeat the analysis in Tables 3 and 5. We obtain similar results.

6.4. Removal of Outliers

Finally, we are concerned that some bloggers may write much more frequently than others, and as a result, their blog posts may have a disproportionately large influence in determining whether certain content is popular or not: If a small number of prolific bloggers mention the same tag in every blog post they write, this tag is likely to be classified as a popular tag even if none of the other bloggers use the tag. We repeat our analysis after excluding the 133 bloggers whose average monthly number of posts is more than four standard deviations above the mean, and we obtain similar results.

7. Discussion and Conclusion

In this paper, we empirically evaluate the impact of ad-sponsored business models on the incentives of online content providers. The theoretical literature on media content provision (e.g., Steiner 1952;

²³ One disadvantage of this approach is that we may not analyze blog posts correctly if they contain mostly pictures or videos.

²⁴ If a post contains fewer than four nouns, we use all the nouns as tags.

Beebe 1977; Gabszewicz et al. 2001, 2006; Gal-Or and Dukes 2003; Anderson and Gabszewicz 2006; Peitz and Valletti 2008) repeatedly find that content providers are likely to cater their content production to popular tastes, although little empirical evidence exists for that prediction.

We find that, consistent with the theoretical literature, content providers sponsored by ad revenue are more likely to generate popular content. Meanwhile, we find that the ad-revenue-sharing program leads to an increased effort toward generating content and making it more likeable. Moreover, the program's impact tends to vary across content providers, being the most pronounced on participants with moderately popular content prior to the program announcement. We also demonstrate that the program effect persists after participants enroll in the program.

Although we demonstrate the potential impact that ad revenue may have on social welfare, we are unable to evaluate the net effect. To measure the net effect, we need to consider the impact of the ad-revenue-sharing program on both content providers and blog readers. A first step would be to quantify changes in the bloggers' utility. Studies in other contexts have shown that content contributors in online social media are generally motivated by joy, social images, and social ties (Toubia and Stephen 2011, Shriver et al. 2013). We therefore need to understand how the introduction of ads affects these motivations. Furthermore, we need to quantify the bloggers' gain from ad revenue and their cost in contributing different types of content.

Similarly, to measure the welfare change for blog readers, we need to evaluate the overall shift in content diversity on the site, and how readers value content diversity and quality. With free entry of bloggers, it is possible that new bloggers or existing nonparticipants choose to cover niche topics vacated by participants after the program's announcement. If this is the case, the ad-revenue-sharing program may have little impact on the site's overall content diversity. Although we do not have data on all bloggers on the site, in unreported regressions, we examine the change in content popularity over time for nonparticipants in our sample. We find that, on average, nonparticipants' content popularity increases after the program's announcement. Among all nonparticipants, those who start blogging after the program's announcement are more likely to contribute popular content than the others, suggesting that many bloggers are incentivized by the prospect of joining the program in the future. In light of these results, there seems to be an overall shift on the site toward popular content.

To fully account for changes in readers' welfare, we also need to consider how the content shift may influence their behavior (e.g., Gopinath et al. 2010,

Dewan and Ramaprasad 2012). For example, as more bloggers are discussing the stock market, are readers receiving better stock recommendations, and as a result, making wiser investment decisions?²⁵ Similarly, are celebrities becoming more popular among readers?²⁶ Another interesting question along these lines has to do with readers' perceptions of the websites dominated by popular content and ads. They can either be discouraged by the fact that the sites are created for ad revenue and hence become less motivated to read the content (Porter and Donthu 2008), or encouraged by the fact that other readers are also interested in the same content (Tucker and Zhang 2010).

Several other limitations of this study are also noteworthy. First, in our empirical setting, content providers do not face a capacity constraint, but in offline media, such as television and newspapers, the capacity constraint can be critical. In such cases, the substitution of unpopular content with popular content may be more pronounced when content providers rely more on ad revenue.

Second, in China, pornography and politically controversial issues are generally not allowed. Therefore, similar studies conducted in other cultural settings may yield different results.

Third, content providers in our setting are compensated by the number of ad impressions they serve. It would be interesting to compare our results to those settings where content providers are paid by the number of ad clicks (e.g., Chen et al. 2009, Ghose and Yang 2009, Goldfarb and Tucker 2011, Liu et al. 2010, Zhang and Feng 2011). In such cases, while content providers will still have incentives to maximize the number of eyeballs, they may also try to match their content with advertisements to increase readers' incentives to click on ads.²⁷

Finally, our work provides implications for future theoretical studies of ad-sponsored business models. In prior studies examining location choices under ad-sponsored business models, quality is often assumed to be exogenous and identical for competing content providers (e.g., Gabszewicz et al. 2006, Godes et al.

²⁵ To explore the consequences of content shifts, one may need to consider content providers' incentives to directly copy and paste others' popular content to attract traffic. See Desai et al. (2010) for a discussion on the effect of digital rights management in the context of digital music.

²⁶ It may also be interesting to consider the changes in content that is related to products. See discussions on product information provision and its consequences on consumer choices in Branco et al. (2012a, b), Sun (2011), Sun and Tyagi (2012), and Xu and Zhang (2013), for example.

²⁷ See Zettelmeyer (2000) for a detailed discussion on firms' pricing and communications strategies when they compete in both online and offline channels.

2009). As a result, these models do not predict the effects of ad-sponsored business models on content quality. Our results suggest that both content quality and location choices need to be endogenized to fully understand the impact of ad-sponsored business models.

Acknowledgments

The authors thank Pradeep Chintagunta, the associate editor, and three referees for their generous input, which greatly improved this paper. For helpful discussion and insightful comments, the authors thank Ron Adner, Victor Bennett, Erik Brynjolfsson, Shantanu Dutta, Kathy Eisenhardt, Chuck Eesley, Lisa George, Avi Goldfarb, Wesley Hartmann, Yinghua He, Donna Hoffman, Jim Lattin, Robin Lee, Josh Lerner, Albert Ma, Michael Manove, Harikesh Nair, John Matsusaka, Marc Rysman, Catherine Tucker, and Dennis Yao; participants of the Organization and Strategy workshop and the Applied Economics workshop at the University of Southern California; and seminar participants at Stanford University, the University of British Columbia, UC Davis, UC Riverside, UCSD, UCLA, Harvard University, the 33rd INFORMS Marketing Science Conference, the 6th Biannual Conference of the Software and Internet Industries at the Toulouse School of Economics, the 9th ZEW Conference on the Economics of Information and Communication Technologies, the 2011 Yale Customer Insights Conference, and the NBER Conference on Economics of Digitization. The authors are also grateful to Sina.com for providing data; to *eaglet*, the author of the open source software *Pan Gu Segment*, for answering many questions related to the use of the software; to Boudhayan Sen, K. Sudhir, and Todd Elder for sharing code and answering questions regarding the Altonji, Elder, and Taber (AET) approach; to Yinfeng Qin and Rita Ren for research assistance; and to the NET Institute (<http://www.NETinst.org>) for financial support.

References

- Altonji JG, Elder TE, Taber CR (2002) Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools. Working paper, Yale University, New Haven, CT.
- Altonji JG, Elder TE, Taber CR (2005) Selection on observed and unobserved variables: Assessing the effectiveness of Catholic schools. *J. Political Econom.* 113(1):151–184.
- Anderson SP, Gabszewicz JJ (2006) The media and advertising: A tale of two-sided markets. Ginsburgh VA, Throsby D, eds. *Handbook of the Economics of Art and Culture*, Vol. 1, Chap. 18 (Elsevier B.V., Amsterdam), 568–614.
- Bagdikian BH (2004) *The New Media Monopoly* (Beacon Press, Boston).
- Baker CE (1994) *Advertising and a Democratic Press* (Princeton University Press, Princeton, NJ).
- Beebe JH (1977) Institutional structure and program choices in television markets. *Quart. J. Econom.* 91(1):15–37.
- Berry ST, Waldfogel J (2001) Do mergers increase product variety? Evidence from radio broadcasting. *Quart. J. Econom.* 116(3):1009–1025.
- Bertrand M, Duflo E, Mullainathan S (2004) How much should we trust differences-in-differences estimates? *Quart. J. Econom.* 119(1):249–275.
- Branco F, Sun M, Villas-Boas JM (2012a) Optimal search for product information. *Management Sci.* 58(11):2037–2056.
- Branco F, Sun M, Villas-Boas JM (2012b) Too much information? Information gathering and search costs. Working paper, University of California, Berkeley, Berkeley.
- Brown KS, Cavazos RJ (2003) Empirical aspects of advertiser preferences and program content of network television. Media Bureau Staff Research Paper 2003-1, Federal Communications Commission, Washington, DC.
- Casadesus-Masanell R, Zhu F (2010) Strategies to fight ad-sponsored rivals. *Management Sci.* 56(9):1484–1499.
- Casadesus-Masanell R, Zhu F (2013) Business model innovation and competitive imitation: The case of sponsor-based business models. *Strategic Management J.* 34(4):464–482.
- Chen J, Fan M, Li M (2012) Advertising versus brokerage model for online trading platform. *INFORMS Conf. Inform. Systems Tech. (CIST)*, Phoenix.
- Chen J, Liu D, Whinston AB (2009) Auctioning keywords in online search. *J. Marketing* 73(4):125–141.
- Cross DW (1994) *Media Speak: How Television Makes Up Your Mind* (Penguin Group, New York).
- Desai P, Purohit D, Vernik D (2010) Digital music set free: The flip side of DRM protection. Working paper, Duke University, Durham, NC.
- Dewan S, Ramaprasad J (2012) Music blogging, online sampling, and the long tail. *Inform. Systems Res.* 23(3, part 2 of 2): 1056–1067.
- DiPrete TA, Gangl M (2004) Assessing bias in the estimation of causal effects: Rosenbaum bounds on matching estimators and instrumental variables estimation with imperfect instruments. *Sociol. Method.* 24(1):271–310.
- Einhorn B (2010) China tries to tame its swinging stock market. *Bloomberg Businessweek* (January 6), http://www.businessweek.com/globalbiz/content/jan2010/gb2010016_835230.htm.
- Eisenmann TR, Parker G, Van Alstyne MW (2006) Strategies for two-sided markets. *Harvard Bus. Rev.* 84(10):92–101.
- Frijters P, Velamuri M (2010) Is the Internet bad news? The online news era and the market for high-quality news. *Rev. Network Econom.* 9(2), doi: 2202/1446-9022.1187.
- Gabszewicz JJ, Laussel D, Sonnac N (2001) Press advertising and the ascent of the “pensée unique.” *Eur. Econom. Rev.* 45(4–6):641–651.
- Gabszewicz JJ, Laussel D, Sonnac N (2006) Competition in the media and advertising markets. *Manchester School* 74(1):1–22.
- Gal-Or E, Dukes AJ (2003) Minimum differentiation in commercial media markets. *J. Econom. Management Strategy* 12(3):291–325.
- Gal-Or E, Geylani T, Yildirim TP (2012) The impact of advertising on media bias. *J. Marketing Res.* 49(1):92–99.
- Gentzkow M, Shapiro JM (2006) Media bias and reputation. *J. Political Econom.* 114(2):280–316.
- Gentzkow M, Shapiro JM (2010) What drives media slant? Evidence from U.S. daily newspapers. *Econometrica* 78(1):35–71.
- George LM (2002) Ownership concentration and product variety in daily newspaper markets. Cranor LF, Greenstein S, eds. *Communications Policy and Information Technology: Promises, Problems, Prospects*, Chap. 9 (MIT Press, Cambridge, MA), 235–253.
- George LM (2007) What’s fit to print: The effect of ownership concentration on product variety in daily newspaper markets. *Inform. Econom. Policy* 19(3–4):285–303.
- George LM, Waldfogel J (2003) Who affects whom in daily newspaper markets? *J. Political Econom.* 111(4):765–785.
- George LM, Waldfogel J (2006) The *New York Times* and the market for local newspapers. *Amer. Econom. Rev.* 96(1):435–447.
- Ghose A, Yang S (2009) An empirical analysis of search engine advertising: Sponsored search in electronic markets. *Management Sci.* 55(10):1605–1622.
- Godes D, Ofek E, Sarvary M (2009) Content vs. advertising: The impact of competition on media firm strategy. *Marketing Sci.* 28(1):20–35.
- Goldfarb A (2004) Concentration in advertising-supported online markets: An empirical approach. *Econom. Innovation New Tech.* 13(6):581–594.

- Goldfarb A, Tucker C (2011) Search engine advertising: Channel substitution when pricing ads to context. *Management Sci.* 57(3):458–470.
- Gopinath S, Chintagunta PK, Venkataraman S (2010) Blogs and local-market movie box-office performance. Working paper, University of Utah, Salt Lake City.
- Greenstein S, Zhu F (2012a) Collective intelligence and neutral point of view: The case of Wikipedia. Working paper, Northwestern University, Evanston, IL.
- Greenstein S, Zhu F (2012b) Is Wikipedia biased? *Amer. Econom. Rev.* 102(3):343–348.
- Hamilton JT (2004) *All the News That's Fit to Sell* (Princeton University Press, Princeton, NJ).
- Herman ES, McChesney RW (1997) *The Global Media: The Missionaries of Global Capitalism* (Cassell, New York).
- Kind HJ, Nilssen T, Sørsgard L (2009) Business models for media firms: Does competition matter for how they raise revenue? *Marketing Sci.* 28(6):1112–1128.
- Leuven E, Sianesi B (2003) PSMATCH2: Stata module to perform full Mahalanobis and propensity score matching, common support graphing, and covariate imbalance testing. Statistical Software Components S432001, Boston College, Chestnut Hill, MA. Last revised July 19, 2012, <http://ideas.repec.org/c/boc/bocode/s432001.html>.
- Liu D, Chen J, Whinston AB (2010) Ex ante information and the design of keyword auctions. *Inform. Systems Res.* 21(1):133–153.
- Mayzlin D, Yoganasimhan H (2012) Link to success: How blogs build an audience by promoting rivals. *Management Sci.* 58(9):1651–1668.
- McChesney RW (2004) *The Problem of the Media: U.S. Communication Politics in the Twenty-First Century* (Monthly Review Press, New York).
- Mullainathan S, Shleifer A (2005) The market for news. *Amer. Econom. Rev.* 95(4):1031–1053.
- Peitz M, Valletti TM (2008) Content and advertising in the media: Pay-tv versus free-to-air. *Internat. J. Indust. Organ.* 26(4):949–965.
- Pew Research Center (2012) In changing news landscape, even television is vulnerable—Trends in news consumption: 1991–2012 (September 27), <http://www.people-press.org/2012/09/27/in-changing-news-landscape-even-television-is-vulnerable/>.
- Porter CE, Donthu N (2008) Cultivating trust and harvesting value in virtual communities. *Management Sci.* 54(1):113–128.
- Postman N (2005) *Amusing Ourselves to Death: Public Discourse in the Age of Show Business* (Penguin, New York).
- Price CJ (2003) Interfering owners or meddling advertisers: How network television news correspondents feel about ownership and advertiser influence on news stories. *J. Media Econom.* 16(3):175–188.
- Reuter J, Zitzewitz E (2006) Do ads influence editors? Advertising and bias in the financial media. *Quart. J. Econom.* 121(1):197–227.
- Rinallo D, Basuroy S (2009) Does advertising spending influence media coverage of the advertiser? *J. Marketing* 73(6):33–46.
- Rosenbaum PR (2002) *Observational Studies*, 2nd ed. (Springer, New York).
- Seamans R, Zhu F (2011) Responses to entry in multi-sided markets: The impact of Craigslist on local newspapers. Working Paper 10-11, NET Institute, New York University, New York.
- Seamans R, Zhu F (2012) Competition and platform positioning. Working paper, New York University, New York.
- Sen B, Shin J, Sudhir K (2011) Demand externalities from co-location: Evidence from a natural experiment. Working paper, Yale University, New Haven, CT.
- Shriver SK, Nair HS, Hofstetter R (2013) Social ties and user-generated content: Evidence from an online social network. *Management Sci.* 59(6):1425–1443.
- Steinem G (1990) Sex, lies and advertising. *Ms. Magazine* (July/August):18–28.
- Steiner PO (1952) Program patterns and preferences, and the workability of competition in radio broadcasting. *Quart. J. Econom.* 66(2):194–223.
- Sun M (2011) Disclosing multiple product attributes. *J. Econom. Management Strategy* 20(1):195–224.
- Sun M (2012) How does the variance of product ratings matter? *Management Sci.* 58(4):696–707.
- Sun M, Tyagi R (2012) When does a manufacturer disclose product match information? Working paper, University of Southern California, Los Angeles.
- Toubia O, Stephen AT (2011) Intrinsic versus image-related motivations in social media: Why do people contribute content to Twitter? Working paper.
- Tucker C, Zhang J (2010) Growing two-sided networks by advertising the user base: A field experiment. *Marketing Sci.* 29(5):805–814.
- Turow J (1998) *Breaking Up America: Advertisers and the New Media World* (University of Chicago Press, Chicago).
- Wilbur KC (2008) A two-sided, empirical model of television advertising and viewing markets. *Marketing Sci.* 27(3):356–378.
- Wray R (2010) Media consumption on the increase. *Guardian* (April 18), <http://www.guardian.co.uk/business/2010/apr/19/media-consumption-survey>.
- Xiang Y, Sarvary M (2007) News consumption and media bias. *Marketing Sci.* 26(5):611–628.
- Xu S, Zhang XM (2013) Impact of Wikipedia on market information environment: Evidence on management disclosure and investor reaction. *MIS Quart.* Forthcoming.
- Yoganasimhan H (2012) Impact of social network structure on content propagation—A study using YouTube data. *Quant. Marketing Econom.* 10(1):111–150.
- Zettelmeyer F (2000) Expanding to the internet: Pricing and communications strategies when firms compete on multiple channels. *J. Marketing Res.* 37(3):292–308.
- Zhang K, Sarvary M (2011) A competitive analysis of Web 2.0 communities: Differentiation with user-generated content. Working paper, INSEAD, Fontainebleau, France.
- Zhang XM, Feng J (2011) Cyclical bid adjustments in search-engine advertising. *Management Sci.* 57(9):1703–1719.
- Zhu F, Iansiti M (2012) Entry into platform-based markets. *Strategic Management J.* 33(1):88–106.