

Peitz, Martin; Reisinger, Markus

Working Paper

The Economics of Internet Media

Working Paper Series, No. 14-23

Provided in Cooperation with:

University of Mannheim, Department of Economics

Suggested Citation: Peitz, Martin; Reisinger, Markus (2014) : The Economics of Internet Media, Working Paper Series, No. 14-23, University of Mannheim, Department of Economics, Mannheim, <https://nbn-resolving.de/urn:nbn:de:bsz:180-madoc-371163>

This Version is available at:

<https://hdl.handle.net/10419/129576>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

University of Mannheim / Department of Economics

Working Paper Series

The Economics of Internet Media

Martin Peitz Markus Reisinger

Working Paper 14-23

October 2014

The Economics of Internet Media¹

Martin Peitz²

University of Mannheim

Markus Reisinger³

WHU – Otto Beisheim School of Management

Chapter prepared for the *Handbook of Media Economics*

Edited by Simon Anderson, David Stromberg, and Joel Waldfoegel

To be published by Elsevier

Abstract: We survey the economics literature on media as it applies to the Internet. The Internet is an important driver behind media convergence and connects information and communication technologies. While new Internet media share some properties with traditional media, several novel features have appeared: On the content side, aggregation by third parties that have no editorial policy and user-generated content have become increasingly important. On the advertiser side, fine-tuned tailoring and targeting of ads based on individual user characteristics are common features on many Internet media and social networks. On the user side, we observe increased possibilities of time-shifting, multi-homing, and active search. These changes have gone hand-in-hand with new players entering media markets, including search engines and Internet service providers. Some of these players face novel strategic considerations, such as how to present search results. In response to these changes, an emerging economics literature focuses on the allocative and welfare implications of this new media landscape. This paper is an attempt to organize these contributions and provide a selective account of novel economic mechanisms that shape market outcomes of Internet media. A large body of work has focused on the advertising part of the industry, while some studies also look at content provision and the interaction between the two.

Keywords: Internet, media economics, digital media, targeting, news aggregation, search advertising, display advertising, two-sided markets

JEL-Classification: L82, L86, M37, L13, D21, D22

¹ We thank Alessandro Bonatti, Emilio Calvano, Alexandre de Cornière, Greg Taylor, Tim Thomes, and Ken Wilbur for helpful comments.

² Department of Economics, University of Mannheim, D-68131 Mannheim; email: martin.peitz@gmail.com

³ Department of Economics, WHU – Otto Beisheim School of Management, D-56179 Vallendar;
e-mail: markus.reisinger@whu.edu

Table of Contents

1. Introduction.....	3
2. Media and advertising on the Internet: Some Facts.....	6
2.1 Facts on Internet media use.....	6
2.2 Facts on Internet advertising.....	11
3. Providing media content	15
3.1 Internet media and the choice of news	15
3.2 Media platforms as gatekeepers.....	16
3.3 News aggregators and the choice of news.....	18
3.4 Search engines and media content	22
3.5 ISPs and media content.....	22
4. Users choosing media content.....	24
4.1 Consumer choice with multi-homing	26
4.2 Search engines and search bias.....	34
4.3 Information spreading on the Internet	43
5. Media platforms matching advertising to content.....	46
5.1 How to formalize targeting	47
5.2 Keyword advertising.....	56
6. Media platforms matching advertising to users.....	61
6.1 Tracking and personalized offers	61
6.2 Advertising congestion and limited attention.....	70
7. Conclusion.....	72

1. Introduction

The Internet has changed the lives of most people, in both business and leisure activities. It allows for a previously unknown immediacy of news coverage from a large number of sources. In a sense, Internet media cover unexpected events in real time over the whole world (see, e.g., Salaverría, 2005).⁴ Searching on Google, social networking on Facebook, and video streaming on YouTube are ways to spread and receive such news. The providers of these services have become increasingly important, affecting media markets in general and Internet media in particular.

Due to technological progress, firms have new and unparalleled opportunities to advertise their products to Internet users, through methods such as fine-tuned targeting. Therefore, Internet companies use business models that are different from the ones used in more traditional media markets and are often purely ad-financed. This raises questions on the efficiency of media markets on the Internet and (potentially) their optimal regulation.⁵

These are only two of several reasons why media on the Internet is an important and highly-debated topic. Other reasons include Internet media's influence on political diversity and their ability to create (or destroy) civic movements more easily than traditional media. This paper provides a guide to the recent literature on the economics of Internet media. After reporting some facts in Section 2, we provide four themes along which Internet media markets can be analyzed.

First, as covered in Section 3, Internet affects content provision of media players. We note a tendency of different media to converge, which makes it increasingly difficult to classify a media outlet on the Internet as belonging purely to one of the traditional media formats. For instance, the web presence of several newspapers and television channels offer text and video. This is more than a labeling issue to the extent that different types of media are subject to different legal rules and that different traditions for media design may become closer substitutes on the Internet since they are only one click away.

While traditional media businesses have built a presence on the Internet (e.g., BBC and CNN) and some have adjusted their business models, new actors are appearing. Big Internet players such as network operators and software platforms have become infomediaries by aggregating information or signing specific contracts with content providers. These aggregators do not have editorial policies but make content suggestions based on algorithms and user feedback (e.g., Google and Facebook). Neither Google nor Facebook started out as news aggregators, but they have evolved into them. For some users, Facebook has become a personalized magazine, as they read the posts on Facebook pages they have subscribed to. Given the reach of Facebook, this can have a fundamental impact on media consumption.⁶ Similarly, YouTube started out providing short amateur videos, which could be considered pure entertainment and outside the media world. However, nowadays, YouTube can be seen as a source of information that functions like media. In addition to revamped old media and big Internet companies moving into media, new media models and players, such as personal blogs and threads, have appeared.

⁴ This holds as long as governments and commercial providers (ISPs, search engines) do not restrict the flow of information. In some countries, certain topics do not appear as search results when using a search engine. To allow for a wide circulation of this article, we do not give concrete examples.

⁵ For a survey on antitrust issues in Internet platforms industries, see Calvano and Jullien (2012).

⁶ See Section 2, on some facts of Internet media consumption.

These include new formats (e.g., messaging networks such as Twitter and blogs) and pure Internet formats of newspapers (e.g., Huffington Post).

To qualify as media, a website has to update information frequently and to replace old content with new content (or to give more visibility to new than to old content).⁷ This distinguishes media from, for example, an encyclopedia. Hence, we would not call Wikipedia media, at least in its current form. However, if Wikipedia adjusted to include news items or propose updated Wikipedia entries according to recent developments, Wikipedia could become Internet media.

Second, as covered in Section 4, the Internet affects user consumption of media content. The Internet has attracted many users, and their behavior online may be different compared to behavior documented in traditional media markets. We observe that for many people in OECD countries, the Internet has become the main source of media consumption. According to GlobalWebIndex, in a large number of countries, the average user spends more time consuming media online than offline (see Section 2).

While it is true that paper has simply been replaced by a screen on a laptop, electronic reader or smart phone, this change in technology, though noteworthy, may not change the underlying economics of media markets. If this were the only change, the economics of Internet media would probably not deserve a survey. Similarly, if people watch a live sports event on their laptop via some streaming service instead of using a television screen, this also constitutes a change of device but would not be a reason by itself to reconsider television media. Finally, if somebody listens only to Internet radio instead of relying on radio frequencies, this constitutes a change in habit, but for our purposes is otherwise irrelevant.

However, users may act differently on the Internet compared to how they consume linear television programming, radio broadcasting and newspapers. While multi-homing is also often a feature of consumer behavior for television programming, the ad-financed media on the Internet are particularly prone to encountering multi-homing users who click themselves through different websites. Thus, the impact of multi-homing on media competition is particularly relevant on the Internet.

Users search for media content primarily via search engines. Thus, the functioning of these search engines is likely to affect media consumption. In particular, an important question is whether search engines bias their results to search queries and whether this affects market outcomes.

In general, consumers play a more active part in media consumption. Media content on the Internet can be subject to reader feedback, which affects the further diffusion of content. This can take the form of comments and recommendations. These also play out on social networks, where user decisions determine the spread of content (e.g., users can share an article or video with their friends). Here, users become, in a sense, curators of the media environment. In addition, they become creators by uploading their own images or videos, which may have news content for small online communities and, thus, may constitute “local” news. This relates to the observation that, in general, a key characteristic of media, in contrast to communication, is that information moves in only one direction, from one sender to many receivers; and this distinction becomes less clear-cut on the Internet, as social networks have resulted in (interactive) user-generated content and limits to exposure. The former implies that some Internet

⁷ This is, among other things, driven by an inherent ephemerality in content value.

media have incorporated elements of interactions. The latter implies that only a select group (e.g., the Facebook friends of a particular person or cause) obtains access to the available information.

Third, as covered in Section 5, media on the Internet match advertising to content. Most Internet media are primarily and exclusively financed through advertising, which can be display advertising or search advertising, among other forms. The former is similar to advertising in traditional media. However, in traditional media, content and advertising tend to be unrelated (with the caveat that some media have narrow target groups – in specialized magazines, for example – that allow for targeted advertising). On the Internet, however, it is feasible to have context-specific ads that depend on the requested content. Such tailoring of ads affects media competition. For example, targeting allows firms with a potentially small customer group to actively participate in the advertising market – a phenomenon known as the “long tail” hypothesis of Internet advertising (Anderson, 2006). This affects advertising prices and, thus, the advertising intensity of big companies with a broad customer base. In addition, better targeting has implications for the interaction between offline and online media.

Of particular relevance is keyword advertising, which consists of advertisements linked to a specific word or phrase. Keyword advertising increases the precision with which advertisers convey their “message” to the right audience, but the general welfare consequences of this type of advertising are not obvious. Also, the incentive of a search engine to offer the best match to consumers is not evident. The matching precision may have consequences for competition between advertisers, thereby affecting the revenue stream of the search engine. Another important question is the effectiveness of such keyword advertising, measured by conversion and click-through rates. Again, it is not clear whether search engines are interested in displaying the results in the most efficient rank order.

We note a great reduction in classified ads in print media. In many cases, their substitutes have become disconnected from media (as in the case of Craigslist) and, thus, are outside the scope of this survey.⁸

Fourth, as covered in Section 6, media on the Internet match users to advertising directly. With the increasing amount of data available about Internet users (including information obtained from cookies), the Internet allows for tracking and individual ad targeting, disconnecting advertising from media content. Again, this has implications for the functioning of media markets. For example, holding ad and product prices constant, better tracking should be beneficial to advertisers because their messages are wasted with a lower probability. Thus, websites should be able reap higher advertising revenues. At the same time, better tracking may increase competition between advertisers since consumers become better informed. This tension affects the outcome in media markets and, in turn, has consequences for Internet users. An obvious one is that users are less likely to block ads if ads, per se, provide a better match. At the same time, there can be more subtle effects. For example, in order to reap more advertising revenues, platforms have an incentive to prevent users from switching to other platforms. A way to achieve this could be to offer higher-quality content, so that, for example, a user consuming news becomes fully informed on a single website and, therefore, has no incentive to consult an alternative website. Hence, content is also affected by the tracking technology.

⁸ Classified ads on the Internet are similar to ads in print yellow pages, but they allow for better tailoring and search.

The improved tracking technology also creates an additional revenue source for websites, given that websites provide data to advertisers. Websites not only can charge users for content and producers for placing ads, but also can offer data to producers. Since this helps producers to make advertising more effective, producers are willing to pay for these data. In addition to providing a better match, better tracking has further advantages to producers. For example, it allows them to retarget users, thereby increasing users' attention span for a particular ad.

With the increasing amount of time users spend on the Internet and websites' strong reliance on ad revenues, users appear to be more likely not to recall all advertising; thus, congestion issues – already a problem in traditional media – appear to be even more relevant on the Internet since users often visit a large number of ad-financed websites.

The literature on media economics, both in general and in how it applies to the Internet, is evolving rapidly. This paper does not aim to provide an exhaustive overview of the economics of Internet media, but presents a selection of recent works on the topics addressed above.

2. Media and advertising on the Internet: Some Facts

2.1 Facts on Internet media use

Internet media include traditional media going online (such as the *New York Times*); pure online media, which in its editorial policies resemble traditional media (such as *Huffington Post*); and pure online platforms (such as Yahoo) that may lack a clear editorial policy and, instead, rely on sources such as Reuters news. Internet media are becoming increasingly important as a source of news. The Pew Research Center runs surveys asking people whether they got news “yesterday” from a particular type of media (data are made available by Pew Research Center). While, in 1991, more than two thirds of respondents said that they got news from television (68%), in 2012, only 55% said so. A similar pattern holds for radio: in 1991, more than half of the respondents said that they got news from radio (54%), while only 29% said so in 2012. Even stronger is the change with respect to newspapers. Here, 56% of respondents said that they got news from newspapers in 1991, and only 29% said so in 2012. So-called digital news has only recently been included in the survey: in 2012, 50% of the respondents said that they got news as digital news.

To the extent that users replace a subscription to a print edition with a subscription to an electronic submission, this merely reflects a change in technology. However, the switch to digital media certainly may affect the type of news and the way that it is consumed. The move to the Internet also affects the production technology of the Internet; in particular, it is less costly to provide more frequent updates, and the variable costs are reduced, compared to print (no printing and low delivery costs).

For the economics of Internet media, it is more important that there are changes in the way that media on the Internet operate. Many newspapers have not made strong inroads into Internet distribution, and those that do must seriously alter their business model. In addition, emerging players in Internet media have attracted a lot of attention, as we document next and treat more systematically in Section 3.

Several online measurement firms (in particular, comScore, Nielsen, and Experian Hitwise) provide information on digital traffic. Since they compute traffic in different ways, their numbers look somewhat different. Here, we report only the numbers released by Experian Hitwise (taken from Pew Research Center, 2013).

Table 1 provides information on the most frequented news websites based on hits in 2012. We find that new and traditional media are both popular. Three general observations can be made. First, the 2012 numbers suggest that different types of traditional media (newspapers such as the *New York Times* and television such as CNN) co-exist on the Internet. Second, new media players have entered – in particular, news portals such as Yahoo and Microsoft’s MSN. Third, Google, as a news aggregator and search engine, has entered the game. To the extent that users consider both types of media as sources of news or information, we observe that the online market of news platforms is more diverse than traditional media markets. In addition, we observe media convergence in the sense that some media that, in the offline world, belong to separate markets now arguably belong to the same market. In Table 1, we observe that the newspapers *New York Times* and *USA Today* have a strong online presence, as do television broadcasters CNN and Fox News.

Table 1: Internet news sites by visits in the U.S. in 2012

	Share	Total Visits
news.yahoo.com	9.96%	5,061,468,962
www.huffingtonpost.com	5.45%	2,772,623,500
www.weather.com	3.8%	1,937,665,856
www.msnbc.msn.com	2.13%	1,096,295,903
www.cnn.com	2.02%	1,006,425,806
gma.yahoo.com	1.70%	871,463,749
www.foxnews.com	1.70%	838,014,419
news.google.com	1.38%	694,071,244
usnews.msnbc.msn.com	1.21%	643,412,705
weather.yahoo.com	1.06%	541,483,467
cityguides.msn.com	1.05%	536,032,331
www.nytimes.com	1.04%	522,225,600
www.drudgereport.com	0.98%	481,470,727
home.now.msn.com	0.95%	497,616,015
www.usatoday.com	0.82%	420,809,971
www.accuweather.com	0.82%	417,186,293
www.weatherunderground.com	0.81%	412,512,889
abcnews.go.com	0.80%	404,226,736
usnews.nbcnews.com	0.72%	356,936,414
local.yahoo.com	0.72%	367,295,141
www.people.com	0.67%	341,541,090
www.newser.com	0.63%	318,767,914
www.washingtonpost.com	0.57%	286,614,807
www.foxnews.com/us	0.57%	291,019,495

www.dailyfinance.com

0.56%

283,381,243

Source: PEW (2013a) using Experian Hitwise

There are multiple ways in which a user may access the media platform. An obvious, necessary condition is that the user has access to the Internet. Here, the user has a contractual relationship with an Internet service provider (ISP). The standard business model of the (user's) ISP is to charge the user a tariff. This tariff can be a flat rate that depends on quality characteristics (such as upload and download speed). More generally, the tariff may be non-linear and depend on the volume of traffic. While not a common feature, an ISP may receive payments from the content/media side for delivering traffic. We return to this issue in Section 3 in the context of the net neutrality debate.

News aggregators provide a means to find news. For example, Google News can be the first stop for a user looking for news on a particular event. While there is (currently) no financial transaction between media platform and media aggregator, this issue is looming in the public debate. For instance, there is strong lobbying by newspapers in Germany to make Google pay for indexing content of German newspapers. In the course of this debate, Google offered to withdraw newspapers from their listing if they did not agree with its policy. Newspapers see themselves as providing content to Google without receiving any monetary payment, allowing Google to make money on advertising. Google, by contrast, claims to deliver additional traffic to newspapers and does not charge directly for the service. Following this logic, newspapers can derive benefit from this traffic (e.g., via advertising, pay-per-view or offering subscriptions). Therefore, it is not clear which economic mechanism would warrant public intervention to satisfy the newspapers' demands for payments by Google.

Users may not turn directly to a news aggregator and may, instead, initiate their search using a general search engine. The user then finds search results by news aggregators and media platforms. Thus, the set of actors include content providers, advertisers, media platforms, news aggregators, search engines, Internet service providers and users. Figure 1 illustrates how content and advertising can reach users. On the advertiser side, our illustration is a bit simplistic. Large media platforms do, indeed, have direct contact with advertisers. However, many websites sell ad space to advertising networks that act as intermediaries between media platforms and advertisers. Additional players, not included in our analysis, are firms offering advertising software tools; these firms offer separate tools to media platforms and advertisers.⁹

⁹ For more details, we refer to Evans (2008). This part of the industry is also characterized by vertical integration. For instance, Google, which is also a major advertising network for display advertising, acquired DoubleClick in 2008. DoubleClick offers software that allows advertisers and media platforms to track users (with the help of cookies) and to organize advertising campaigns. DoubleClick also offers software for search advertising.

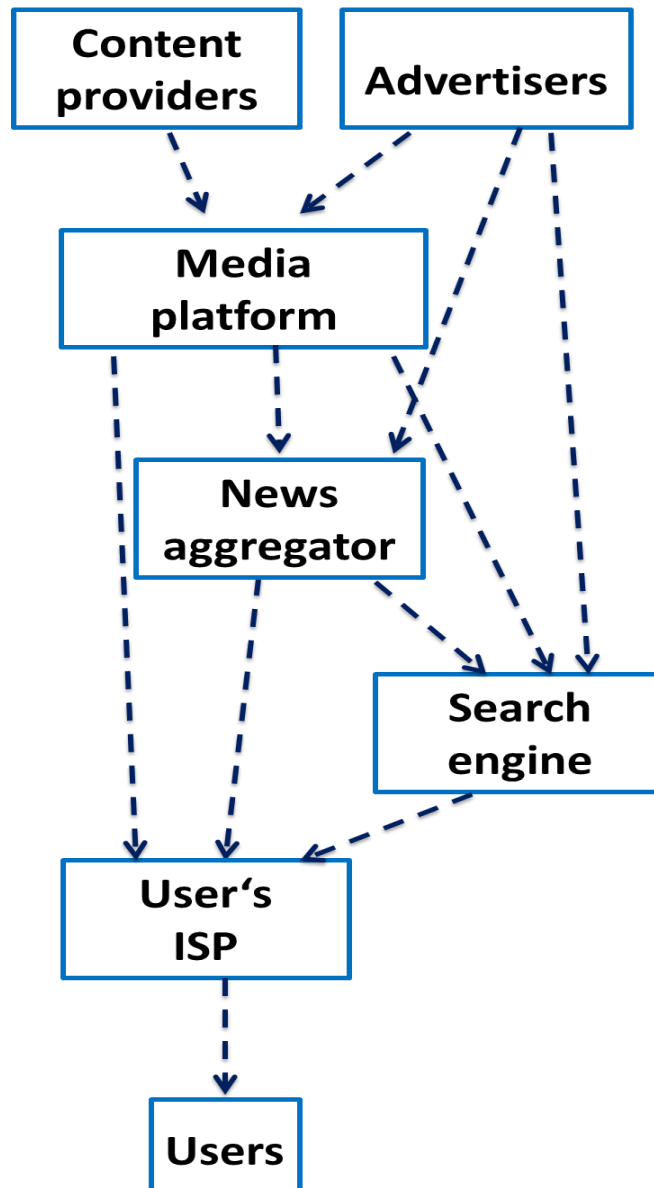


Figure 1: A stylized Internet media market

Media operate in a richer environment on the Internet than in the offline world. As in many traditional media, a media platform on the Internet typically combines content and ads. When the platform does not vertically integrate with the content provider side, it has contracts with content providers and advertisers. The most common business model, then, is to offer such bundles of content and advertising to Internet users. The Internet user derives a benefit from the content offer, while advertising may also affect her utility.

Advertisers pay media platforms for placing their ads and delivering them to Internet users. Content providers receive payments from the media platform for delivering content. In purely ad-financed media,

Internet users do not make monetary payments to the media platform. However, viewers of these ads are valuable, as they pay with their eyeballs, similar to free-to-air television.¹⁰

In the presence of multiple media platforms, the standard assumption is that users make a discrete choice among the platforms. Although this assumption appears to be appropriate in the traditional newspaper market, it is questionable how well it fits in the case of television, and even more so in the context of Internet media, where the alternative media platform is just one click away. We discuss alternative models of user behavior (in particular, multi-homing) and their implications in Section 4.

News consumption via social network plays an increasingly important role for a large number of users. News is defined here as “information about events and issues that involve more than just your friends or family” (see, e.g. PEW, 2013c). To document the user base, PEW (2013b) provides survey results for the U.S. adult population. Among adult Internet users, 71% are active on Facebook, while the corresponding numbers are 46% for those over age 65 and 90% for those aged between 18 and 29. A more recent phenomenon is social networking on mobile phones; in September 2013, 40% of all Internet users were active with social networking on mobile phones. The engagement on several social networking sites, especially on Facebook, is strong: 63% of Facebook users visit the site at least once a day, and 40% do so multiple times throughout the day.

In particular, Facebook has become an important platform for users to get news. According to PEW (2013c), with survey data from September 2013, 64% of U.S. adults use Facebook, and 30% of U.S. adults use Facebook to get news on the site; thus, for around half of its users, Facebook is a news site. Another popular site is YouTube. While 51% of U.S. adults use YouTube, only 10% of U.S. adults get news from it, which is around one in five. Only Twitter exhibits a ratio similar to that on Facebook: 16% of U.S. adults are on Twitter, and 8% get news from it.

With respect to the role of social networks in the creation and spread of information, we note that some users may become creators by posting their own images and videos, while others become curators by reposting and sharing existing material. PEW (2013d), based on a survey data from 2013, classified 54% of social network users as creators and 47% as curators. According to a July-August 2012 survey, two thirds of social network users have shown political engagement on the network – e.g., by encouraging others to vote, expressing their political opinion or sharing others’ political expressions (PEW, 2013e). This shows that social networks do not only carry information relevant for a small set of people, but also contribute to the general debate on civic issues. In our survey, we address the information-spreading aspect of Internet media in Section 4.

¹⁰ If users consider advertising to be a nuisance, a media model on the Internet can be developed along the lines of traditional media, as is formalized, for instance, in the Anderson and Coate (2005) model (for a recent, more-general treatment, see Anderson and Peitz, 2014b). This model formalizes the interaction among advertisers, media platforms, and users. This media model connects to the literature on two-sided markets, with seminal contributions by Rochet and Tirole (2003) and Armstrong (2006), among others. For a textbook treatment, see Belleflamme and Peitz (2010).

2.2 Facts on Internet advertising

The previous subsection provided some facts about Internet media use and the content side. Next, we present some facts on Internet advertising, which started only in 1994 with the sale of the first banner ad (see Kaye and Medoff, 2000). As pointed out in the introduction, the Internet has changed the landscape for advertisers. To the extent that advertisers have replaced some of their advertising on traditional media with Internet advertising, this has a direct impact on those media. A striking example has been the move of classified ads from print media to electronic platforms. While classified ads used to bring in substantial advertising revenues for newspapers (which possibly cross-subsidized other parts of the newspaper), Internet platforms have drastically cut those revenues for newspapers unless they have been able to dominate the respective market segment on the Internet. This loss in revenues on the advertising side has implications for the pricing of newspapers. In particular, newspapers increase subscription prices, sacrifice circulation and set lower ad prices (see Seamans and Zhu, 2014).

After television advertising, Internet advertising has become the most important medium in terms of ad revenues, at least in the U.S. As Figure 2 illustrates, in 2012, advertising revenues for Internet media were close to 37 billion US\$. As in other media, advertising can play different roles for advertisers. It can inform about product availability, price and product characteristics; it may allow consumers to draw inferences about product characteristics (advertising as a signal: Nelson, 1974 and Milgrom and Roberts, 1986); it may change consumer preferences to the benefit of the advertiser and, thus, be persuasive; or it may serve as a complement to the product (Becker and Murphy, 1988).¹¹ We return to these views on advertising when discussing different Internet advertising formats.

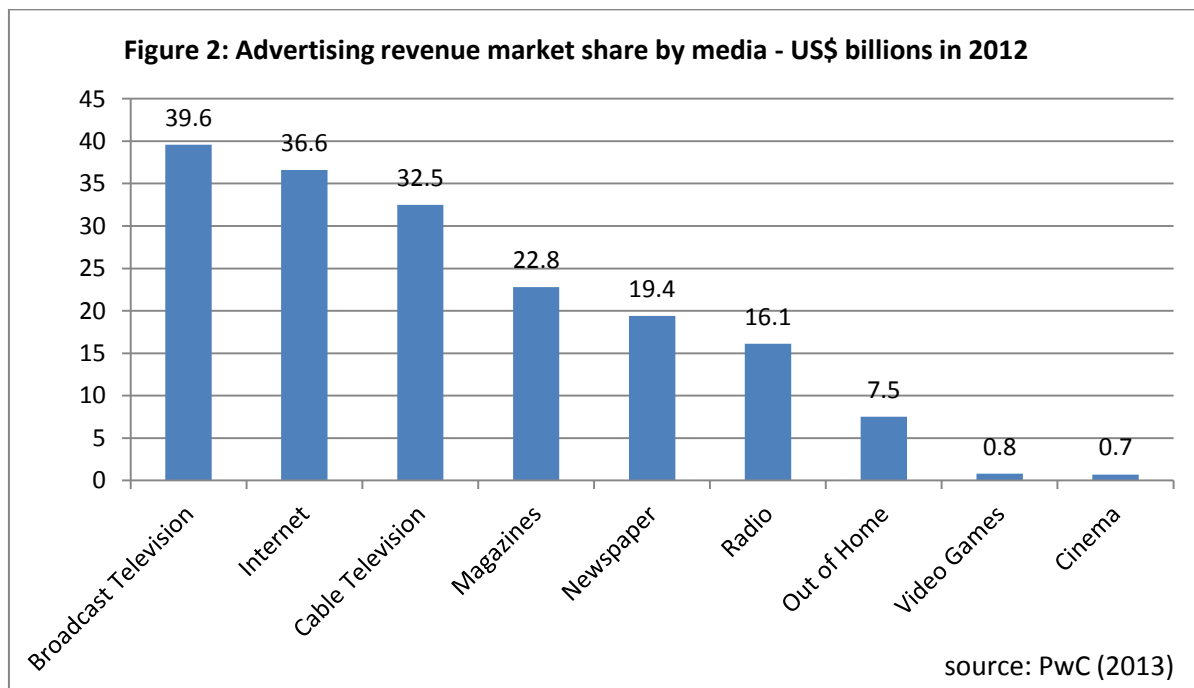
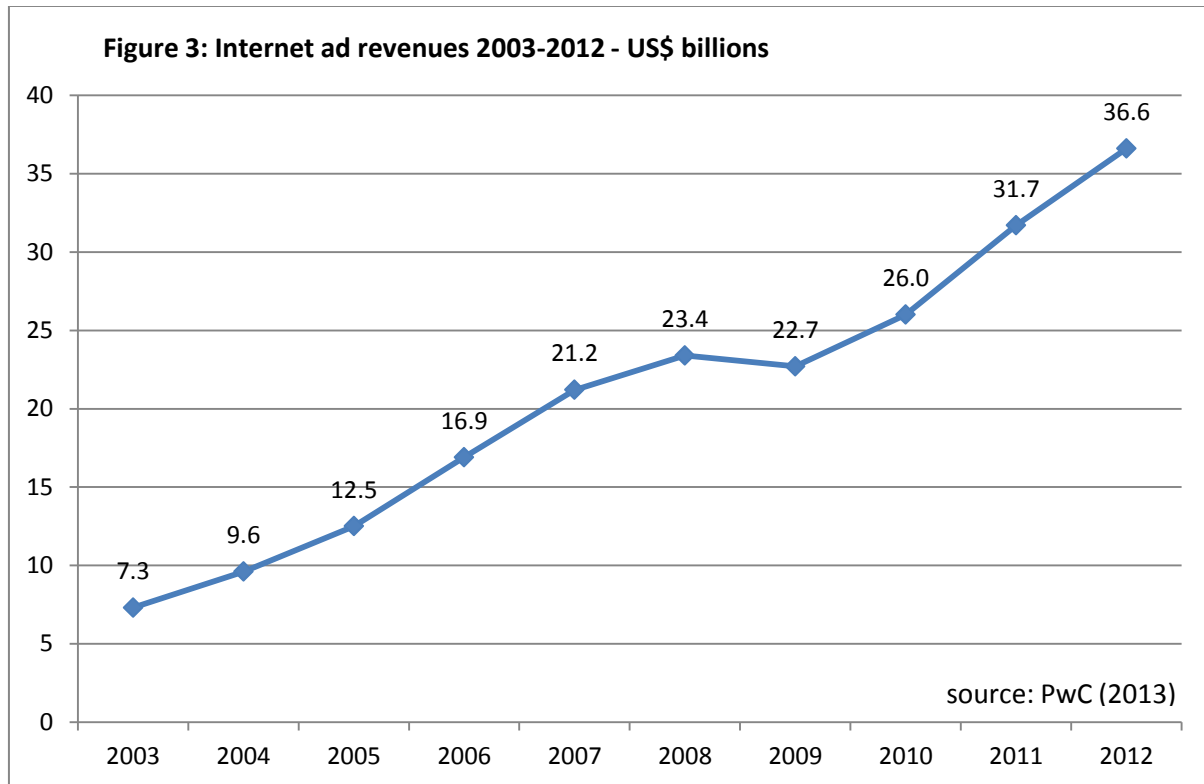


Figure 3, in turn, shows the so far unstoppable increase in advertising spending on Internet media, with an annual growth rate of about 20% over the last ten years. Only at the height of the financial crisis in

¹¹ Bagwell (2007) provides an excellent survey of the economics of advertising.

the U.S. from 2008 to 2009 can one observe a small dip in ad revenues. This contrasts with the development of newspaper advertising revenues. In 2000, ad revenues stood at around 49 Billion US\$; in 2012, they totaled less than 20 Billion US\$.¹²



Internet advertising can take different forms.¹³ Here, we mention the most common: *search advertising* accounts for a large fraction of total Internet ad revenues. To the extent that it makes consumers aware of certain offerings, this type of advertising can be seen as directly informative. Consumers may also consider it to be indirectly informative about the quality of a match through a signaling role of the ad – e.g., if an offering appears at the top of the page of sponsored search results.

The dominant platform for search advertising is Google; in the U.S. it first took the top position in 2003, relegating Yahoo to second place. In 2013, around two thirds of all search requests on general search engines were made on Google, while 18% were made on Microsoft sites (using Bing) and 11% on Yahoo (according to comScore qSearch). In Europe, Google is even more dominant, accounting for more than

¹² There is compelling evidence that Internet advertising can be effective in stimulating purchases. This does not hold only for online purchases, but also for offline ones. For example, Lewis and Reiley (2014) find in a controlled field experiment with a large retail store that more than 90% of purchase increases occur in brick-and-mortar stores. Similarly, Johnson, Lewis and Reiley (2014) demonstrate that repeated exposure to online ads was highly profitable for the retail store whose data were analyzed.

¹³ We refer to Evans (2008) for a detailed explanation of search and display advertising. See Goldfarb (2014) for a discussion of the different forms of online advertising,

90% of searches in 2013. Google is stronger yet in terms of revenues. An explanation for this finding is that the larger platform is more attractive to advertisers.

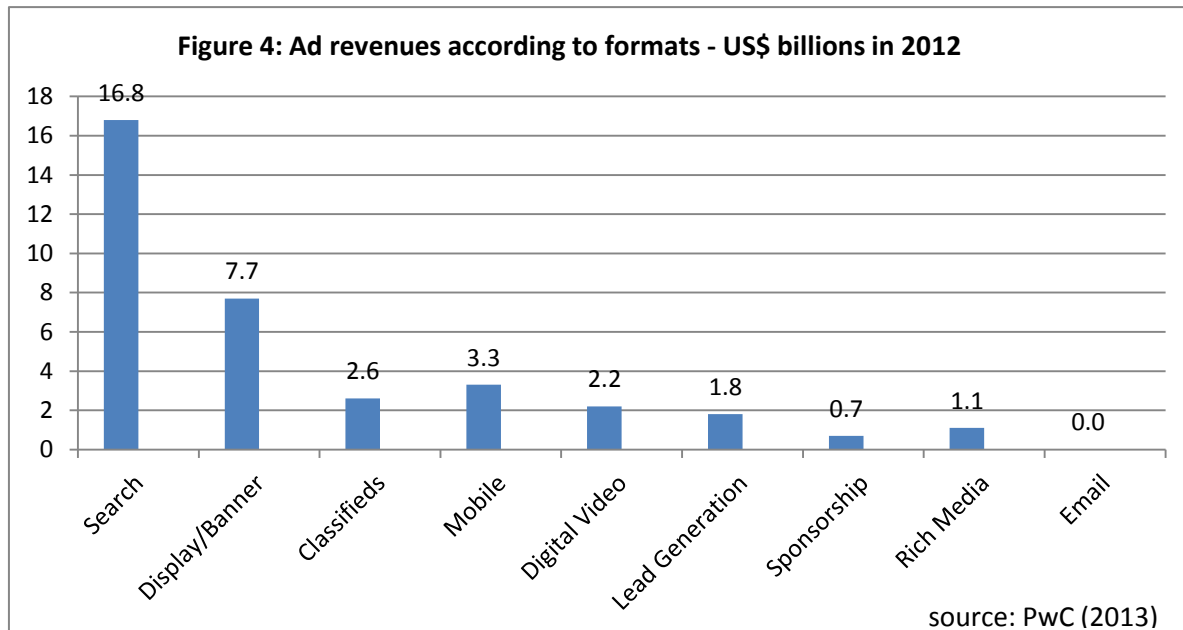
The second big category of ad revenues is *display advertising*. Display advertising can be contextual, as it can be linked to particular keywords or phrases. We can also refer to such a strategy as tailoring. In addition, display advertising on the Internet can be personalized if the advertising platform has knowledge about consumer characteristics. This then leads to targeted advertising on the Internet,¹⁴ which improves the match between the advertised product and the consumer. In general, such advertising can be informative, persuasive or complementary.

The third major category of advertising revenues used to be *classified ads*; however, revenues have been decreasing over the last years in the U.S. In the early 2000s, classified ads moved quickly from newspapers to Internet platforms (e.g., Craigslist in the U.S.) – one reason for the fast decline of advertising revenues at many newspapers. Classified ads on the Internet allow users to apply individual searches rather than using a predetermined classification scheme. Otherwise, the economics of classified ads did not change because of the move to the Internet. One would often consider this advertising format informative, as it makes consumers aware of an offering.

Additional formats, which are listed in Figure 4, include mobile advertising and digital video advertising. Mobile advertising, which saw a quick increase between 2010 and 2012, has the potential to add another tailoring dimension: the displayed ad may depend on a consumer being physically close to a particular location at which a product or service is available. Advertisers' hope is that mobile advertising is a means to generate immediate purchases. Here, advertising may play mainly an informative role, as it makes consumers aware of a product or service that they may be interested in at a particular location. Digital video advertising is akin to television advertising, with the important difference that it allows for tailoring and targeting (we note that television advertising also allows for some tailoring). This format may be more attractive for advertising that is persuasive or serves as a complement. Figure 4 reports advertising

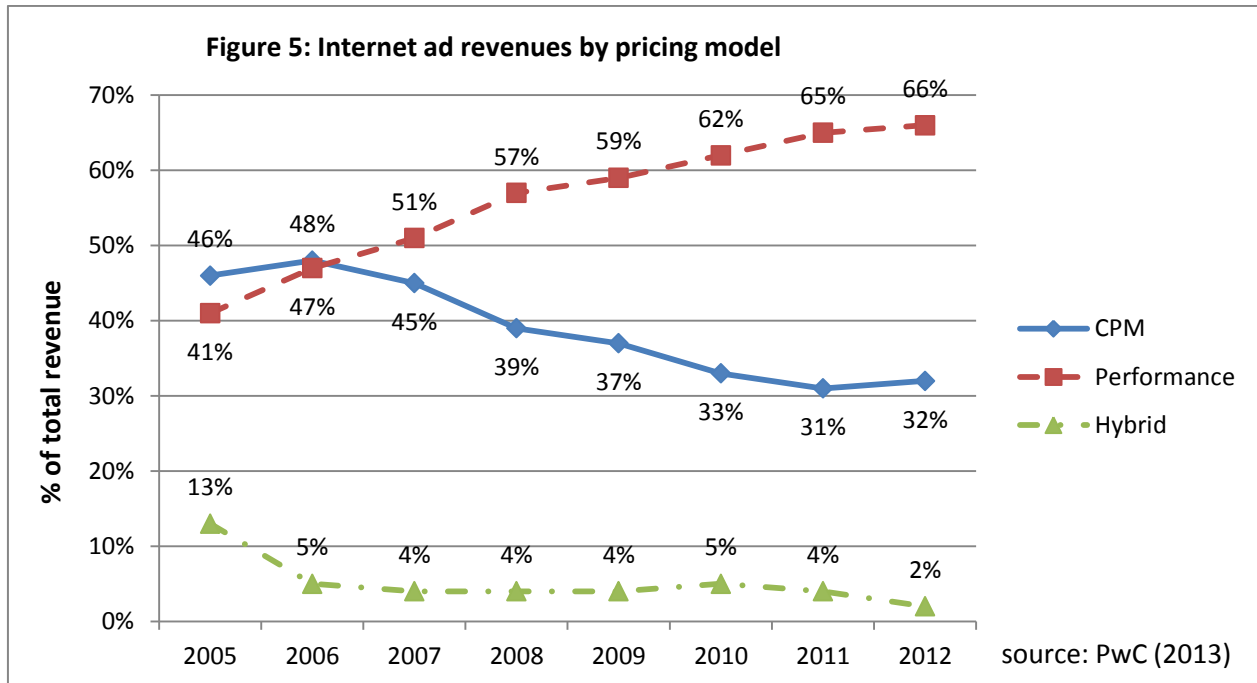
¹⁴ Athey and Gans (2010) make the distinction between tailored and targeted advertising. See Sections 5 and 6 for more details.

revenues according to the ad format for the U.S. in 2012.



Media platforms on the Internet offer the possibility of measuring the impact of advertising – by counting the number of clicks an ad generates, for example. This has opened up the possibility of using a pricing model that is different from simply counting the number of impressions, as is done traditionally in media. The measure for the latter is CPM (cost per mille – i.e., the cost per one thousand impressions). If the ad price depends, instead, on the number of clicks (or possibly even the number of purchases), the traditional pricing model is replaced by one that is performance-based. As Figure 5 illustrates, performance-based ad pricing has been gaining the upper hand on the Internet, accounting for close to two thirds of all Internet ad revenues in 2012. Ad revenues accruing to Google, through sponsored search results, make up the majority of these revenues.

Internet advertising may affect not only own-product sales. Several researchers have conducted natural experiments on Internet advertising to analyze spillovers of advertising campaign on competitors. For example, Lewis and Nguyen (2014) used display advertising on Yahoo!’s front page. On some days, Yahoo! sells ads on its front page as ad-splits, which is, two display ads are alternately shown throughout the day, one every even second and the other every odd second. This allows naturally for exogenous variation because there is no systematic difference between even-second and odd-second visitors, thereby providing a test group (visitors exposed to the target ad) and a control group (visitors exposed to the other ad-split ad). Lewis and Nguyen (2014) find that display ads caused an increase in searches for the advertised brand (35-40%) but also for competitors’ brands (1-6%). The effect on searches by brand can be up to 23% for competing brands. Sanhi (2013) finds similar results using data from restaurant search websites. He also shows that the extent of the spillover depends on the advertising intensity. If the intensity is low, spillovers are particularly large, perhaps because advertising then has mainly the effect of reminding consumers about similar options. By contrast, if the intensity is high, spillovers disappear and the advertiser receives more searches.



The business model of most news websites is ad-financed. However, there are a few exceptions, primarily newspapers. For example, in the UK, almost all major newspapers require readers to subscribe to their online content. *The Times* has a pay-wall for all of its content. Other newspapers, such as *The New York Times* or *The Wall Street Journal*, offer limited free access with some number of free articles per month. Therefore, these newspapers have a mixed business model—that is, partly ad-financed and partly subscription-financed.

3. Providing media content

Media content on the Internet reaches users through multiple channels. In this section, we discuss some of the issues of media content provision and how users access this content. We address the various layers of the value chain, as illustrated in Figure 1, with a focus on the provision of content. While a large part of this section will be informal, the section contains a formal analysis of media platforms as gatekeepers and of the role of news aggregators. We take a broad view and also shortly discuss the role of search engines (formally investigated in Section 4) in the context of user choice, and the role of ISPs with a guide to the net neutrality debate.

3.1 Internet media and the choice of news

The availability of media content on the Internet dilutes or even removes the boundaries between newspaper and television channels. For instance, in the case of news programming, a consumer can visit the website of a news channel or one of a newspaper. Each of these media typically has combined offerings of electronic articles (containing text and often photos) and videos. This describes the convergence of different media. We may view this convergence as market integration, implying that offerings that used to be independent (newspaper versus television) become substitutes. How this convergence affects market structure is an issue that deserves investigation.

One issue is the coverage of different topics. Media platforms have to choose topics that differ in their success probability, with success meaning that the topic attracts users' attention and generates a reward (e.g., through ad revenues). In such markets, the quality of the media platform and its ability to predict the success of a topic affects its choice of topic.¹⁵

Another issue (for empirical research) is whether the Internet alters the way news is consumed. In particular, news on the Internet facilitates the combined consumption of news and background information. For instance, a particular news item may lead a user to consult Wikipedia for additional information or background. While we typically would not classify Wikipedia as media, the overall portfolio of information on a topic consumed over the Internet may look very different from the product consumed in traditional media.

The convergence of different types of traditional media also leads to important policy questions. As newspapers and television channels on the Internet become closer substitutes, policy makers have to tackle the fact that newspapers and television channels are, in many places, subject to different regulations. Furthermore, in many countries, public service broadcasters play an important role in traditional radio and television markets. As television channels develop an Internet presence, they start to compete with newspapers on the Internet. While we do not intend to address these policy issues directly, this section helps to understand the functioning of Internet media markets and may, therefore, also be helpful from a policy perspective.

3.2 Media platforms as gatekeepers

The gatekeeper role of media appears to be of particular relevance on the Internet – the issue also arises in traditional media, as discussed below. Media platforms on the Internet (and traditional media) can be seen as managing the amount and type of information a user can digest. In particular, if a user has a limited attention span for news and is unable or unwilling to push herself to read more news, a media platform, by recommending a selection of news, can emphasize the most relevant news items. This is especially the case if the platform has information on the user's tastes and tailors the news selection to them, as is the case with Internet media. To address the role of Internet media platforms as gatekeeper, we reinterpret the model by Anderson and de Palma (2009), in that we include individual news providers that compete for users' attention.

Consider a media platform that selects among alternative news providers and offers a selection of news items to a user. Suppose that there are n potential news providers, each contributing up to one piece of news. A news provider of type θ obtains advertising revenue $\pi(\theta)$. Advertising revenue is increasing in the value of content θ because we assume that a news provider with more-valuable content is more likely to deliver advertising to the user.

The user has an attention span φ for news items, which will be derived below; i.e., she randomly selects φ news items if she is presented with more than φ . In particular, if there is no selection among news, each available news item is seen with probability $\min\{\varphi/n, 1\}$. Suppose that delivering a news item

¹⁵ For a theory contribution on this point, see Katona, Knee, and Sarvary (2013).

costs κ . Under free entry and appropriate boundary conditions, there is a marginal type $\theta^* = n$ such that $\frac{\pi(\theta^*)\varphi}{n} - \kappa = 0$. This determines the number of news items $n(\varphi; \kappa)$ under free entry.

The user's cost of sampling messages is $C(\varphi)$. Given the total number of news items n , each randomly sampled news item gives an expected surplus of

$$S^e(n) = \frac{1}{n} \int_0^n s(\theta) d\theta.$$

The utility-maximizing attention span is $\arg \max_{\varphi} \varphi S^e(n) - C(\varphi)$. The first-order condition $S^e(n) = C'(\varphi)$ determines the chosen attention span when $\varphi \leq n$. Hence, $\varphi(n) = \min\{n, C'^{-1}(S^e(n))\}$. In equilibrium, for low values of κ , there may be information overload – i.e., $\varphi^* \leq n^*$.

The media platform can manage the amount of news being offered to users. In other words, the media platform can become a gatekeeper. Traditional media partly filled this role by selecting articles provided by news agencies and other sources. In the case of newspapers, they traditionally charge readers for the active selection of content. While they may also provide unique content, an important role of both traditional and new media is this selection role. Whether this selection is based on an editorial policy or is software-based is irrelevant in this context. Limiting the amount of content on the platform can reduce or even eliminate information overload. In particular, a monopoly gatekeeper would price out information overload. It may do so by charging news providers for access.

To the extent that users frequent multiple media platforms, these platforms do not fully internalize that additional content on their platform reduces the probability that content on other platforms will receive the users' attention. Hence, information overload remains an issue for competing media platforms. A characterizing feature of the Internet is that users often visit many media platforms.¹⁶ Therefore, the Internet is particularly prone to the problem of information overload.

The media platform's role as gatekeeper is not restricted only to the amount of content, but also to content selection and quality. These two dimensions are also highly influenced by the ad-financed business model. For example, Sun and Zhu (2013) conduct a study on how the content and quality of blogs are affected by an ad-revenue-sharing program of a Chinese portal site. In particular, this portal site launched the program and invited around 3000 bloggers to participate. Bloggers who allowed the portal site to run ads on their blogs received 50 percent of the ad revenues generated on the site. Around 1000 bloggers decided to participate.

Sun and Zhu (2013) find that the decision to participate in this program has led to a shift to more-popular content by around 13%. Around 50% of this increase comes from shifts to topics from three domains: stock market blogs, salacious content, and blogs about celebrities. The blog posts of participants in these domains increased by 6.6% relative to non-participants.

In addition, participating bloggers also increased their content quality. This is measured, for example, by the number of users who bookmark a post as one of their favorites and by the average number of

¹⁶ See Section 4 for a more detailed discussion of this issue.

characters, pictures, and video clips. (For example, more characters means that the blogger invests more effort in writing and goes deeper into the focal topic, whereas more pictures often make the blog more attractive.) This suggests that bloggers exert more effort on blogging content that they are not necessarily intrinsically interested in, and that they do so to obtain higher revenues from advertising.

Sun and Zhu (2013) also show that these effects are strongest for moderately popular bloggers and, in particular, stronger than for very popular bloggers (and non-participating bloggers). A possible explanation is that very popular bloggers have always covered popular topics and/or maintained a high level of quality. Hence, there is not much space for improvement. Non-participating bloggers, by contrast, may have a large disutility from blogging about content that does not reflect their tastes, and, thus, they choose not to participate.

3.3 News aggregators and the choice of news

News aggregators such as Google News have added another layer in the market for Internet media. Users may use an aggregator as the main access point and select news items based on this aggregator's listings. A few academic studies have tried to shed light on the link between quality choice of media platforms and the presence of news aggregators.

Here, the role of a news aggregator such as Google News is to help users to easily find high-quality content. As Dellarocas, Rand, and Katona (2013) point out, absent news aggregator, users may find their way to different news because media platforms may provide links to a rival's content. They present a model in which users are interested in a particular event, and different media platforms cover this event with different quality. Users are not informed *ex ante* about the quality and, therefore, visit media platforms randomly. While users appreciate the provision of external links (towards higher-quality content), and this increases the overall attractiveness of a media platform, users reduce their time spent on a particular platform, which decreases ad revenues for the media platform. Despite such links, there is a role for the news aggregator, as it actively selects among content, allowing users to avoid the hassle of moving from one media platform to another that provides higher quality. The news aggregator improves the overall performance of the market, which is in the overall interest of media platforms, as this increases user participation in the market; however, the news aggregator also absorbs some of the rents generated in the market, thereby reducing what is on the table of media platforms. Dellarocas, Rand, and Katona (2013) show that if content providers offer links to each other's content, entry of an aggregator may lead to less competition among content providers to provide high quality. If, by contrast, content providers cannot link to rivals' content, entry of an aggregator tends to lead to more competition among content providers.

Rutt (2011) analyzes quality choice by media platforms when some users are loyal to a particular media platform, and others search for free high-quality content. Loyal users are willing to pay for content, but searchers are also valuable, as additional traffic generates advertising revenue. The presence of an aggregator is assumed to be essential for searchers to identify the quality of content. A platform may charge for content but then loses searchers. Platforms simultaneously set price on the user side and the quality of content. Under some conditions, a mixed-strategy symmetric equilibrium exists, with the feature that media platforms randomize over price and quality by choosing from a probability

distribution among qualities at a price of zero or by setting a particular quality at a positive price. If a platform ends up doing the latter, it serves only loyal users (and extracts all the surplus from them), while in the former equilibrium, it competes with other platforms for searchers. The main finding is that as the fraction of searchers increases, platforms with free content increase their content quality (in the sense of first-order stochastic dominance), while platforms with positive prices decrease their content quality. In addition, as the fraction of searchers increases, content is provided more often for free. Thus, existing searchers benefit from more searchers. Also, loyal searchers benefit, as it becomes more likely that they can enjoy content for free.

Jeon and Nasr (2013) study media platform competition, where users can access content either by going directly to a media platform or by accessing news via an aggregator, such as Google News. The presence of such a news aggregator affects competition between media platforms, especially in the long run, when they react to changes by adjusting the quality of their news items. To analyze this issue, Jeon and Nasr (2013) propose a stylized model of competing media platforms that offer news items from a full set of news categories. Each news item is either of high or of low quality.

Media platforms may operate in an environment without a news aggregator or in an environment in which the news aggregator selects high-quality news items for all news categories for which high quality is available. Low quality is selected only for those news categories for which media platforms fail to make high quality available.

The formal setup is as follows. There is a continuum of categories $[0,1]$ that each of two media platforms covers. Each platform chooses a subset $I_i \subseteq [0,1]$ of categories for which it offers high-quality news. We can then associate the quality of the media platform with the measure of high-quality news issues $q_i = \mu(I_i)$, where it is assumed that $\mu([0,1]) = 1$. The cost of quality-provision is assumed to satisfy that a media platform always chooses a policy with $\mu([0,1]) \leq 1/2$. Users are identical with respect to their valuation of high-quality coverage.

In addition to differences in high-quality coverage, platforms are horizontally differentiated, where differentiation may reflect different political views or different styles (such as using British or American English). This horizontal difference applies to each news category. Users are heterogeneous with respect to these horizontal platform characteristics. This is formalized following the standard Hotelling representation with platforms being located at the extreme ends of the $[0,1]$ -interval and users being uniformly distributed over this interval. A user located at x incurs a disutility of τx if she consumes all news from platform 1 and $\tau(1 - x)$ if she consumes all news from platform 2. Then, the utility of a user x choosing platform 1 is

$$v_1(x) = u_0 + q_1 - \tau x,$$

where the marginal utility from increasing high-quality coverage is normalized to 1. Correspondingly, the utility of x choosing platform 2 is

$$v_2(x) = u_0 + q_2 - \tau(1 - x).$$

In the absence of a media aggregator, users face a discrete choice problem between two media platforms.

Platforms obtain revenues from advertising. Jeon and Nasr (2013) assume that platform revenues are proportional to the amount of time users spend on a platform, which is implied by a constant advertising price per exposure and exposure being proportional to the amount of time spent on the platform. It is assumed that users spend a total of one unit of time on low-quality media. High-quality coverage makes them investigate a category longer, increasing the time spent on news consumption by δ per category. Hence, a user who consumes news only on platform 1 spends time $1 + \delta q_1$ on platform 1. Each unit of time spent on the platform generates ad revenues A . Denoting the number of users of platform 1 by D_1 in the absence of a news aggregator, the profit of platform 1 is

$$\pi_1(I_1) = AD_1(1 + \delta\mu(I_1)) - C(\mu(I_1)),$$

where the cost of increasing high-quality coverage is convex (in particular, $C(s) = cq^2$ for $q \leq 1/2$ and infinity for larger q). Here, the parameter c is assumed to be sufficiently large such that in all environments, platforms choose high-quality coverage with q_i strictly lower than $1/2$.

In the absence of a media aggregator, platforms simultaneously choose their high-quality coverage I_i across news categories. Then, users make a discrete choice between the two media platforms. Because users cannot combine news from different platforms, for any given q_i , each media platform i is indifferent to which particular category it offers high quality. Straightforward calculations show that qualities are strategic substitutes; i.e., if the competing platform increases its amount of quality coverage, the best response by the media platform is to decrease its own quality coverage. One can then show that in a symmetric equilibrium, quality coverage is decreasing in the differentiation parameter τ , while profits are increasing. In other words, if users consider platforms weak substitutes, media platforms invest less in quality coverage. This finding is in line with the basic intuition that more differentiation makes competition less intense.

How does the presence of a news aggregator affect competition between media platforms? Its presence changes the picture considerably because the aggregator proposes a mix of news from different platforms. In this respect, users see the news aggregator as a device to multi-home – that is, it allows them to, indeed, mix across platforms.¹⁷ In the presence of a news aggregator, media platforms are no longer indifferent about which particular category contains high-quality news, as it is relevant whether there is duplication of high-quality coverage for the different categories. In particular, if platforms fully specialize – i.e., $I_1 \cap I_2 = \emptyset$ – we have that $\mu(I_1 \cap I_2) = 0$. By contrast, if platforms choose maximal overlap, we have that $\mu(I_1 \cap I_2) = \min\{\mu(I_1), \mu(I_2)\}$.

To illustrate the functioning of a news aggregator, suppose that there are six (instead of a continuum of) news categories. Furthermore, suppose that platform 1 offers the vector $(1,0,1,1,0,0)$, where 1 stands for high quality and 0 for low quality, while platform 2 offers $(0,1,0,1,1,0)$. By choosing maximal quality, the news aggregator then offers $(1,1,1,1,1,0)$, where, in the case of the same quality, we postulate that each media platform is listed with probability $1/2$. Thus, the news aggregator provides higher quality than

¹⁷ For an analysis of the effect of multi-homing users on advertising revenues, see Section 4.

each individual media platform. From the user's perspective, this comes at the cost of a worse fit (for $x \neq 1/2$).

Returning to the model with a continuum of categories, a user located at x who obtains news through the news aggregator receives utility

$$v_{12}(x) = u_0 + \mu(I_1 \cup I_2) - \eta_1 \tau x - \eta_2 \tau (1 - x),$$

where η_i is the fraction of news items that are linked to media platform i . This fraction is the sum of the fraction of categories with exclusive high-quality news items on media platform i , $\mu(I_i) - \mu(I_1 \cap I_2)$; one half of the fraction of news categories with two high-quality news items, $(1/2)\mu(I_1 \cap I_2)$; and one half of the fraction of news categories that do not contain any high-quality news items, $(1/2)(1 - \mu(I_1 \cup I_2))$.

Thus, we can write

$$v_{12}(x) = u_0 + \mu(I_1 \cup I_2) - \frac{\tau}{2} + \tau \left(x - \frac{1}{2} \right) (\mu(I_2) - \mu(I_1)).$$

A user $x < 1/2$ prefers the news aggregator to media platform 1 if $v_{12}(x) > v_1(x)$. This is equivalent to

$$\mu(I_1 \cup I_2) - \mu(I_1) > \tau \left(\frac{1}{2} - x \right) (1 + \mu(I_2) - \mu(I_1)).$$

The left-hand side contains the gain due to higher quality from the news aggregator and the right-hand side the loss due to the larger preference mismatch with respect to horizontal characteristics. Whenever there are some categories for which only platform 1 offers high quality and some others where the reverse holds, a user at $x = 1/2$ strictly prefers the mix provided by the news aggregator over the offers by the two media platforms. Hence, users fall into up to three sets: users around $1/2$ rely on the news aggregator, while users at the extreme points tend to rely on the respective media platform. By contrast, if $I_1 = I_2$, there is no room for a news aggregator.

It can be shown that media platforms either choose full specialization such that $\mu(I_1 \cap I_2) = 0$, or that they provide maximal overlap such that $\mu(I_1 \cap I_2) = \min\{\mu(I_1), \mu(I_2)\}$. In any symmetric equilibrium with $\mu(I_1) = \mu(I_2) = \mu$, media platforms choose full separation if exposure due to high-quality news is sufficiently large, $\delta\mu > 1$. It can then be shown, that for large δ , there is a unique symmetric equilibrium in which platforms choose full separation and the news aggregator is active. In this environment, qualities are strategic complements, and the market-expansion effect due to higher quality dominates the business-stealing effect. In the reverse case, the business-stealing effect dominates the market-expansion effect, and there are equilibria in which media platforms choose the same categories for high-quality news items. This suggests that the viability of the news aggregators depends on the demand expansion of high-quality news items.

When the media platforms fully separate their high-quality coverage, users benefit from the presence of the news aggregator. While the effect on media platforms' profits is ambiguous, total surplus is also higher.

The overall message that emerges from the analysis of news aggregators is that they affect the media platforms' incentives to invest in the quality of content. The above discussed works have identified

situations in which the presence of news aggregators is beneficial for society; however, the opposite result may hold true, in particular since the news aggregator is an additional player extracting – also at the margin – rents from the market.

3.4 Search engines and media content

Search engines are an important entry point for readers. Readers may be interested in a certain topic or event and simply use Google or some other search engine (e.g., Bing, Baidu in China, Yandex in Russia, or Naver in South Korea) to click on a particular news item. This traffic generates profits for the search engines, as it allows them to place ads together with the organic search results. As news items are linked, a click by the reader on a particular news item moves the reader to a particular news site. When searching on Google, for example, the reader receives information on the news provider (e.g., the Internet site of a newspaper or television channel) and a snippet from the news item, which provides some context in which the search item appears. The distinction between a search engine and a news aggregator is, in some cases, a bit blurred, as a reader may use Google or Google News to access news, where we would label the former a search engine and the latter a news aggregator. An issue in both cases is, first, whether search engines have the right to provide links or need an explicit agreement from the website owner to provide a link, and, second, whether search engines have an obligation to treat all content in a transparent and “non-biased” way.

Concerning the former, some interested parties have asked to be compensated for the extraction of snippets. For instance, the industry association of German newspapers has asked to receive payments. We note that, before this request, Google had already offered newspapers the option to delist their content. In this case, neither snippets nor links are provided in the organic search results on Google. Under a new law,¹⁸ a group of German media companies hopes to extract license fees from Google for making snippets available (it is apparently unclear what length of the snippet would justify such a license fee). Essentially, this group of media companies aims to sustain a positive price vis-à-vis Google by coordinating their actions; this would not be possible if they acted independently.

The second issue has been analyzed in the context of search neutrality. While we are not aware of academic work on the first issue, several contributions have considered search neutrality, which applies not only to searches for news items, but also to broader searches, including those for products. We discuss search neutrality in Section 4.2.

3.5 ISPs and media content

To enjoy media content, users need an Internet connection. Thus, a user obtains her consumption utility from jointly consuming both the content and the connectivity service. If the user is not using public wifi, she typically will have a contract with an Internet service provider (ISP). This ISP offers her download and upload services at a contractually agreed-upon speed.

When content travels from the content provider to the consumer, the provider accesses the Internet via its ISP. Content is then sent through the Internet to the consumer’s ISP. Traditionally, the content

¹⁸ This is the ancillary copyright for press publishers (“Leistungsschutzrecht für Presseverleger”), which came into force on August 1, 2013. In its initial draft, it was intended to introduce a fee even for short snippets, but this has been removed from the final version.

provider makes payments to its ISP. The ISP then ensures that content is delivered to the consumer's ISP. The consumer pays her ISP for the access product. There are no payments from the content provider to the consumer's ISP. In addition, all material is treated equally according to the best-effort principle.

Due to the explosion in data volume, a new issue is congestion, which leads to delays at certain times or to the breakdown of some services. Internet media are part of the congestion issue; according to Sandvine (2014), real-time entertainment, which includes media, constitutes a large fraction of the traffic. For instance, on mobile networks in Europe, YouTube contributes 20.62% and Facebook 11.04% to downstream traffic; as reported in Section 2.1, a large fraction of this traffic stems from news accessed by users. The OECD predicts that video streaming and IP-based television will increase traffic volumes (OECD, 2014).

Congestion issues are particularly relevant with mobile access where capacities are lower, but may also take place on landlines (DSL, cable). Some content providers have opted for the possibility of bypassing the public Internet and the risk of delay at interconnection points by operating content-delivery networks. Also, some ISPs offer media products (e.g., tv) that are treated differently from other content. Furthermore, as part of the net neutrality debate, there is discussion about whether a consumer's ISP can also charge on the content-provider side, thus introducing two-sided pricing. In addition, ISPs may inspect the data that they are handling and decide – based on the characteristics of the data in question – which type should receive priority treatment (deep packet inspection). Furthermore, as a number of countries are currently considering, content providers might self-select into different service classes, as ISPs offer both a slow and a fast lane. Such tiering would be legal according to the European Commission's proposal. Content providers could pay for prioritized access (while the "slow" lane is typically considered to be free). It has become mostly a political question whether these more flexible approaches should be allowed.

Proponents of strict net neutrality want to rule out such approaches, forcing ISPs to obtain all their revenues on the consumer side and not allowing them to deviate from the best-effort principle, which treats all traffic symmetrically. Critics of strict net neutrality point out that a one-sided price structure in a two-sided market tends to lead to rents on one side while reducing rents on the other. In particular, ruling out payments by content providers to consumer ISPs may lead to low overall revenues for user ISPs and may reduce their incentives to invest in a more powerful access network. While the exact competitive effects of imposing net neutrality rules are complex, when investment incentives by ISPs are considered (see, e.g., Choi and Kim, 2010; Economides and Hermalin, 2012; Krämer and Wiewiorra, 2012; Bourreau, Kourandi, and Valletti; 2014), a general issue is that unrestricted transmission may lead to congestion problems. Here, the capacity at a particular point (e.g., a switch close to the user) constitutes a common property resource.

An important observation is that some types of traffic are time-sensitive, meaning that consumer's utility is strongly negatively affected by delay (e.g, video calls or online gaming), while other traffic is not (e.g., video-on-demand or emails). For media, live sports events tend to be time-sensitive, whereas most other types of content can be delivered with short delays without costs for consumers, provided that they have equipment at home that allows for buffering.

A first basic result can be obtained in a model with a monopoly ISP. If the ISP is not allowed to discriminate between time-sensitive and time-insensitive traffic (by treating these types of traffic differently or by charging for priority access), both types of traffic are treated in the same way based on best effort. For any given composition of time-sensitive and time-insensitive traffic, this implies that in times of congestion, the allocation could be improved if time-sensitive traffic received priority. Since suppliers of time-sensitive traffic have an incentive to deliver on time, they also have an incentive to obtain prioritized access even if it carries a positive price. Taking into account that some traffic is more time-sensitive than other traffic makes the introduction of a priority lane potentially welfare-improving. Prioritized delivery may be secured through a price charged on the priority lane. Then, the priority lane serves as a screening device, and ISPs do not need to know the type of data they are delivering. Alternatively, ISPs may inspect the data packages and decide whom to give priority. This typically requires ISPs to look into the packet, which, however, raises privacy and data protection issues.

As Choi, Jeon, and Kim (2013) and Peitz and Schuett (2014) point out, inframarginal content providers can also adjust their traffic volume. For instance, to deal with congestion, they may invest in compression technologies to reduce the volume of traffic.¹⁹ While deep packet inspection addresses the inefficiency caused by treating time-sensitive and time-insensitive traffic equally, it is not a useful tool to tackle the incentive issues faced by content providers (see Peitz and Schuett, 2014). This holds since deep packet inspection does not put a price on congestion, in contrast to charging a price for prioritized delivery under tiering. By putting a price on time-sensitive traffic, which might be passed on to final users, the ISP creates an incentive for content providers to avoid such payments. In particular, content providers have an incentive to reduce the volume of time-sensitive traffic that they send.

Summarizing the state of the net neutrality debate is beyond the scope of this article (for a recent survey on the academic analysis of net neutrality, see Krämer, Wiewiorra and Weinhardt, 2013). The important lesson emerging from the debate is that regulatory decisions affect the rent distribution between content providers and ISPs. This, in turn, may affect the strategy of media platforms. In particular, if congestion is not priced, content providers may add a lot of traffic stemming from advertising (e.g., advertising preceding videos). If congestion is priced, media platforms may obtain a larger fraction of revenues from charging consumers directly.²⁰

4. Users choosing media content

In traditional audiovisual media, such as television, consumers usually need to choose which content to consume at any given point in time and which to dismiss. Consider a television viewer who is interested in two different movies and one sports game broadcast by different stations on the same evening. This viewer has to choose one program to watch but, by making this choice, misses the other programs. This problem of linearly progressing content of TV or radio is absent in online media offers.

¹⁹ Alternatively, this can be done by making traffic less time-sensitive, which can be achieved through buffering and, thus, removing the need to obtain prioritized delivery.

²⁰ Moreover, the pricing of congestion may affect the choice of format of media content. For instance, video consumes more bandwidth than text, and with video, a higher resolution requires more transmission capacity.

Content provided by Internet platforms can be quite durable. For example, media libraries allow users to access content at any time. Thus, the Internet is a non-linear medium in which each consumer can choose her preferred time and order of content consumption.

In this respect, Internet media offers content at the individually preferred time, in contrast to the predefined time slots of traditional audiovisual media. So it shares important features of (digital) VCRs and on-demand content with respect to time shifting, but unlike these other forms of on-demand content, the cost to obtain the same content at different points in time on the Internet is almost negligible. In fact, another website is “just one click away,” and accessing it does not require costly hardware. For traditional media, time shifting requires special hardware devices such as VCR or PVR set-up boxes (e.g., offered by TiVo or DirecTV). In addition, accessing multiple websites is often an almost mindless activity, while deciding which content to videotape is a more conscious decision.

Consequently, online media consumers are usually multi-homers, whereas in traditional media markets, consumers are more likely to choose one outlet and stick to it. Consider, for example, the market for newspapers: Most consumers read only one daily newspaper (if any) due to time constraints and often stick to this choice for a long time.²¹ Also, during the course of an evening, TV viewers who want to watch a movie usually choose one and single-home on the channel showing the movie. For the other side of the market, this implies that an advertiser can reach a particular consumer only by placing ads in the particular newspaper that this consumer is reading, or by placing commercials during the movie that the viewer is watching. To inform a large number of consumers, advertisers need to buy ads on multiple outlets due to consumers’ single-homing behavior. This problem is captured by the seminal competitive bottleneck model of Anderson and Coate (2005) and follow-ups.

If, instead, consumers choose multiple outlets, an advertiser can reach a consumer not only on a single platform but on multiple ones. In this respect, platforms lose their monopoly power of delivering consumers’ attention to advertisers. In the competitive bottleneck model, in order to attain such a monopoly position, platforms fight for the exclusive turf of consumers, thereby capturing rents on the advertiser side but dissipating parts of these rents to consumers. In a market with multi-homing on both sides, this is no longer necessarily true. This implies that the well-known force that competition intensity is determined by the strength of business stealing on the consumer side is less relevant in online media markets. Since consumers are active on multiple platforms, new forces come into play and old ones are probably disabled. This can affect platforms’ content choice.

In Section 4.1, we provide different formalizations of multi-homing on the consumer side. We focus particularly on implications for competition and distinguish them from models with single-homing consumers. Before doing so, we note that if advertisers could perfectly coordinate their messages, then multi-homing would be equivalent to single-homing. Consider the situation in which advertising has decreasing returns-to-scale – that is, the first impression is very valuable, but further impressions are less valuable because there is a probability that the consumer has already noticed the ad on another

²¹ However, Gentzkow, Shapiro, and Sinkinson (2014), using historical data on newspaper readership find that even in the newspaper market, multiple readership is quantitatively important. In particular, in their data, 15% of households that read a daily newspaper read two or more newspapers.

platform. If advertisers can perfectly coordinate their messages, they can prevent a consumer from being exposed to the ad multiple times. For example, on TV, this requires that an advertiser choose the same time slot for its ads on each station. Therefore, competition in this model is equivalent to competition in a model with single-homing consumers. Anderson and Peitz (2014a) use the formulation to study advertising congestion. We discuss this paper in more detail in Section 6.2.

4.1 Consumer choice with multi-homing

One of the first attempts to allow consumers to combine consumption of multiple products was done in the Hotelling (1929) framework in a one-sided market. Suppose that there are two platforms, 1 and 2. The content provided by each platform is interpreted as its location on the Hotelling line. Platform 1 offers content α and platform 2 offers content $1 - \beta$. Most of the literature works under the assumption that a consumer subscribes to only one platform. This implies that the disutility of a consumer located at x from not consuming the preferred content is $g(|x - \alpha|)$ or $g(|1 - \beta - x|)$, depending on which platform the consumer is active, where g is an increasing function. In most papers, this function is assumed to be linear or quadratic. Anderson and Neven (1989) extend this formulation by allowing a consumer to consume any mix of the two contents of

$$\omega\alpha + (1 - \omega)(1 - \beta),$$

with $0 \leq \omega \leq 1$. The disutility incurred by a consumer at x under the assumption of quadratic disutility is $(\omega\alpha + (1 - \omega)(1 - \beta) - x)^2$. Therefore, the consumer can obtain her optimal content by combining the existing content in the right way.²² The Hotelling model with content mixing can be straightforwardly interpreted in the media market context. Suppose that each consumer has some amount of time that she can allocate between the two platforms. Then, the consumer spends a share $\omega(x)$ of this amount on platform 1 and $1 - \omega(x)$ on platform 2.

Suppose that consumers are uniformly distributed on the interval between 0 and 1. Although the utility formulation is very different from the standard single-homing Hotelling model, the resulting aggregate demand function is exactly the same if consumers pay for the amount of time they spend on a platform. To see this, recall, first, that in a traditional Hotelling model, the aggregate demand of firm 1 is

$$D_1 = \frac{1 + \alpha - \beta}{2} + \frac{p_2 - p_1}{2(1 - \alpha - \beta)}. \quad (1)$$

Let us now briefly derive the aggregate demand in the mixing model. The utility function of a consumer located at x is

$$v(p_1, p_2, \alpha, \beta, x) = u_0 - (\omega\alpha + (1 - \omega)(1 - \beta) - x)^2 - \omega p_1 - (1 - \omega)p_2,$$

where u_0 is the gross utility from using the platform. Maximizing with respect to ω , we obtain

²² Anderson and Neven (1989) consider linear pricing and find that the model gives similar results as obtained in the standard Hotelling model. Hoernig and Valletti (2007) allow for two-part tariffs and find that consumers, then, do not necessarily choose their preferred product because they need to pay fixed fees to both platforms.

$$\omega(x) = \frac{2(1 - \alpha - \beta)(1 - \beta - x) - p_1 + p_2}{2(1 - \alpha - \beta)^2}$$

for

$$\alpha + \frac{p_2 - p_1}{2(1 - \alpha - \beta)} < x < 1 - \beta + \frac{p_2 - p_1}{2(1 - \alpha - \beta)},$$

while $\omega = 1$ for $x \leq \alpha + (p_2 - p_1)/(2(1 - \alpha - \beta))$ and $\omega = 0$ for $x \geq 1 - \beta + (p_2 - p_1)/(2(1 - \alpha - \beta))$. Therefore, consumers whose preference is close to the content of one of the platforms do not mix, while those located at less extreme positions choose to mix the content. Determining the aggregate demand of firm 1, we obtain

$$D_1 = \alpha + \frac{p_2 - p_1}{2(1 - \alpha - \beta)} + \int_{\alpha + \frac{p_2 - p_1}{2(1 - \alpha - \beta)}}^{1 - \beta + \frac{p_2 - p_1}{2(1 - \alpha - \beta)}} \omega(x) dx = \alpha + \frac{p_2 - p_1}{2(1 - \alpha - \beta)} + \frac{1}{2(1 - \alpha - \beta)}.$$

Simplifying this expression, it is easy to see that

$$D_1 = \frac{1 + \alpha - \beta}{2} + \frac{p_2 - p_1}{2(1 - \alpha - \beta)},$$

which equals (1), implying that the two formulations are equivalent. In other words, single-homing in the Hotelling model can also be interpreted as multi-homing of consumers who mix content. As a consequence, competition plays out in exactly the same way in the two models.

We now extend the framework of mixing content to a two-sided ad-financed media model – i.e., platforms obtain their revenues from advertisers instead of consumers. That is, prices p_1 and p_2 are equal to zero, but consumers view advertising levels a_1 and a_2 on the platforms as a nuisance. Therefore, the utility function of a consumer located at x is

$$v(a_1, a_2, \alpha, \beta, x) = u_0 - (\omega\alpha + (1 - \omega)(1 - \beta) - x)^2 - \gamma\omega a_1 - \gamma(1 - \omega)a_2,$$

where γ represents the nuisance parameter of advertising.

Gal-Or and Dukes (2003) use this framework to analyze content choice in media markets. They find that platforms choose the same location on the Hotelling line, a result in stark contrast to the one obtained in the traditional framework with quadratic transport costs and a uniform distribution of consumers on the unit interval, in which firms locate at the extreme points of the interval. In the model of Gal-Or and Dukes (2003), advertisers compete in the product market and inform consumers about their products via advertising. A lower advertising intensity leads to less-intense product market competition, implying that advertisers' prices and profits are higher. By choosing minimal differentiation, platforms reduce the amount of advertising in equilibrium because advertising is a nuisance to consumers. Hence, intense competition for consumers in the media market results in low advertising levels. With minimal

differentiation, platforms commit to a low advertising intensity, allowing them to demand higher prices from advertisers.²³

Gabszewicz, Laussel and Sonnac (2004) do not consider advertiser competition in the product market but assume that the disutility of consumers is convex in the advertising level – that is, the disutility from advertising is a_i^θ , with $\theta \geq 1$. As they show, in equilibrium, platforms may choose a location in the interior range of the Hotelling line; that is, the content is relatively similar.²⁴

These papers are based on the idea that consumers mix the time that they spend on different platforms, keeping the total amount of time fixed. However, in most markets, the availability of content increases consumption. These features have been incorporated into several models; we begin with those that keep the assumption of consumers being distributed on the Hotelling line and incurring a disutility in distance.

In Ambrus and Reisinger (2006), platforms are located at the endpoints of the Hotelling line, with platform 1 located at point 0 and platform 2 located at point 1. The utility of a consumer located at x is then $u_0 - \gamma a_1 - tx$ when consuming from platform 1 and $u_0 - \gamma a_2 - t(1 - x)$ when consuming from platform 2. A consumer can be active on both platforms and optimally does so if the utility she obtains on each platform is positive.

Deriving the demand functions, consumers located close to platform 1 will be exclusive consumers of platform 1, while consumers who are located close to platform 2 will be exclusive consumers of platform 2. However, consumers in the middle segment of the Hotelling line enjoy a positive utility on each platform (given that advertising levels are not too high). These consumers are, therefore, overlapping consumers active on both platforms. Denoting the demand of exclusive and overlapping consumers of platform $i = 1, 2$ by D_i we obtain $D_1 = (u_0 - \gamma a_1)/\tau$ and $D_2 = 1 - (u_0 - \gamma a_2)/\tau$. Overall, there are $D_{12} = [2u_0 - \gamma(a_1 + a_2)]/\tau - 1$ overlapping consumers, $D_1 - D_{12} = 1 - (u_0 - \gamma a_2)/\tau$ exclusive consumers on platform 1 and $D_2 - D_{12} = 1 - (u_0 - \gamma a_1)/\tau$ exclusive consumers on platform 2. The demand configuration is displayed in Figure 6.

²³ Exclusive advertising contracts are a different means to mitigate competition between advertisers (and allowing platforms to charge higher advertising prices). These contracts are standard practice, e.g., in the U.S. television industry. By offering single-category advertising rights, a platform guarantees not to sell another slot in the same advertising break to any close competitor. Therefore, consumers are less informed about competing products, yielding higher profits for advertising firms. For a detailed analysis, see Dukes and Gal-Or (2003).

²⁴ Peitz and Valletti (2008) also find that platforms do not choose “maximal” differentiation to obtain higher surplus from advertisers. Their model is cast in a framework in which all consumers single-home and advertisers multi-home. However, as they note, their results carry over to a setting in which consumers mix content.

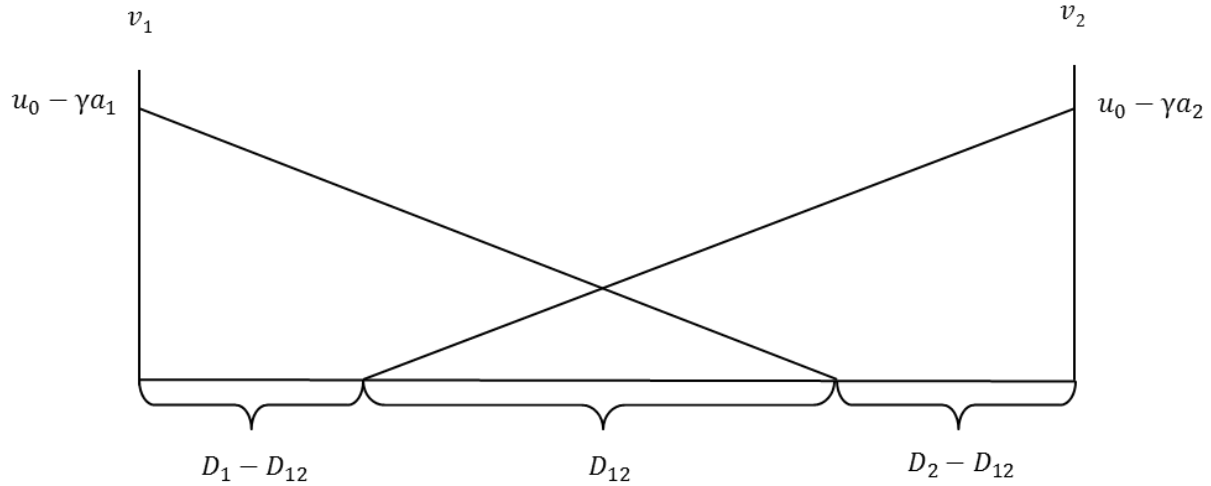


Figure 6: Consumer demand structure with multi-homing in Ambrus and Reisinger (2006)

An advertiser's profit from reaching a consumer can be different for overlapping and exclusive consumers. For example, an overlapping consumer may spend less time on a platform than a consumer who is exclusively active on a single platform. Similarly, an advertiser who buys advertising space on both platforms may receive a lower benefit from an overlapping consumer than an advertiser who displays ads on only one platform because the overlapping consumer may become aware that the advertisement is already on the other platform. This implies that an (additional) advertisement on platform i is less valuable for an advertiser who is also active on platform j than for an advertiser who is active only on platform i .

Ambrus and Reisinger (2006) characterize the equilibrium advertising levels in the cases of monopoly and duopoly. They find that advertising levels are socially excessive if advertisers are homogeneous due to the fact that multi-homing reduces the competitive pressure in advertising levels more than single-homing does.

Anderson, Foros and Kind (2010) also analyze multi-homing consumers in the Hotelling setting. In their model, media platforms do not carry advertising but, rather, charge viewers directly. (Thus, their model is not a two-sided market model.) They distinguish between two types of overlapping consumers, those who prefer platform 1 over platform 2 but subscribe to both, and those with the opposite preference. More formally, suppose that a consumer located at x obtains a utility of

$$v_1 = (r - \tau x)q_1 - p_1$$

when subscribing to platform 1 only. Here, q_1 represents the coverage of platform 1 and p_1 the price of platform 1. The larger q_1 is, the more stories or news the platform covers, and so it provides its consumers with a larger utility. In this respect, high coverage can be interpreted as high quality. The parameter r represents the reservation value of a consumer while τ is, as before, the transportation parameter. Therefore, a consumer benefits more from high coverage if her preferences are better aligned with the content of the platform – that is, if her location is closer to that of the platform. Correspondingly, the utility from being active on platform 2 is $v_2 = (r - \tau(1 - x))q_2 - p_2$.

When active on both platforms, the consumer obtains an additional utility. In particular, Anderson, Foros and Kind (2010) explicitly allow multi-homing consumers to first consume the content of the platform they prefer most and then consume the content of the other platform. What is the additional value a consumer obtains from subscribing to the second platform? This value depends on the extent of overlap in the content of the two platforms. This overlap is given by q_1q_2 , as platforms cannot coordinate their content coverage. Therefore, platform 2 covers $(1 - q_1)q_2$ stories that are not covered by platform 1. The consumer also obtains an additional value when reading the same story twice because each platform presents it in a different way. This additional value can be captured by $1 - b$, with $b \in [0,1]$; that is, the lower b , the higher the additional value of reading the same story a second time. This implies that the additional benefit of consuming an amount q_2 of news or stories when first consuming an amount q_1 is

$$(1 - q_1)q_2 + (1 - b)q_1q_2 = (1 - bq_1)q_2.$$

Therefore, the utility of a consumer who visits platform 1 and then platform 2 is given by

$$v_{12} = v_1 + (r - \tau(1 - x))(1 - bq_1)q_2 - p_2.$$

In the same way, we can determine the utility of a consumer who first visits platform 2 and then platform 1. It is given by

$$v_{21} = v_2 + (r - \tau x)(1 - bq_2)q_1 - p_1.$$

The locations of consumers indifferent between different pairs of options are illustrated in Figure 7, in which the marginal consumers are denoted by $x_{1,12}$, $x_{12,21}$, and $x_{21,2}$.

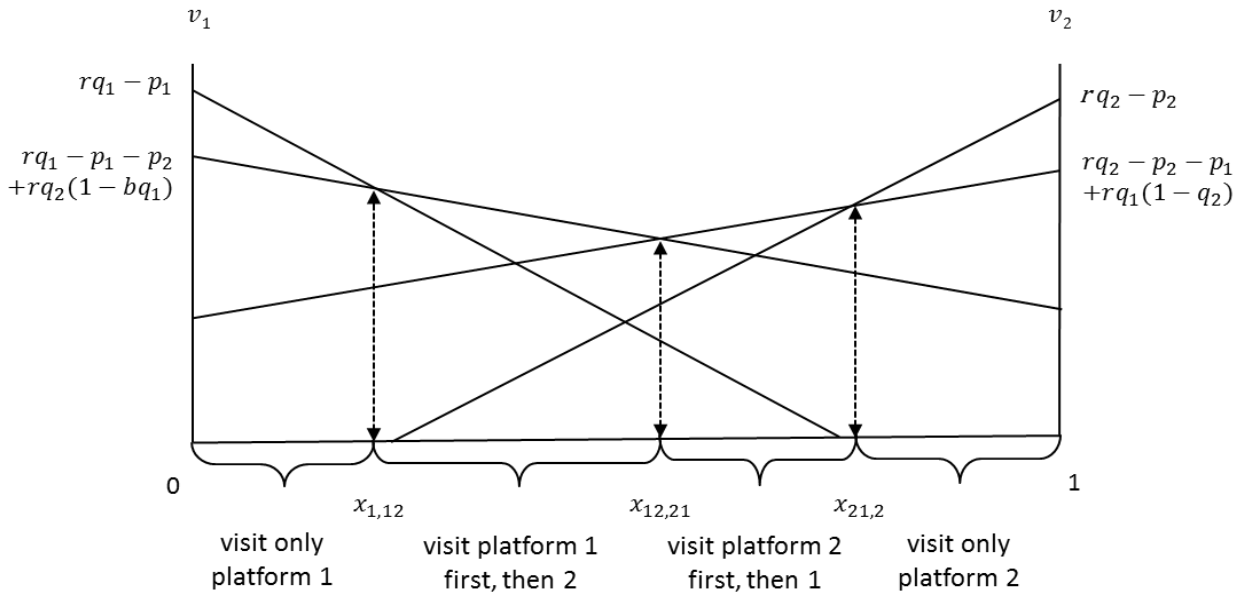


Figure 7: Consumer demand structure with multi-homing in Anderson, Foros and Kind (2010)

In this framework, Anderson, Foros and Kind (2010) solve for the equilibrium prices and profits. They find that profits are decreasing in the transportation cost parameter τ . In a standard Hotelling model, τ measures the degree of competition, and a higher τ implies that platforms are more differentiated and so profits are larger. By contrast, with multi-homing, the result is reversed because the total demand of platform 1 is independent of the price charged by platform 2. To see this, note that the total demand of platform 1 includes all consumers between 0 and $x_{21,2}$. When determining the consumer who is indifferent between, on the one hand, visiting platform 2 first and platform 1 next and, on the other hand, visiting only platform 2, we have $v_1 = v_{12}$. In this expression, the price of platform 2 is canceled out. Therefore, each platform's total demand is independent of the rival's price, and platforms act as monopolists in their pricing decisions. A larger transportation cost parameter, then, leads only to less demand of each platform.

While the assumption that consumers' preferences are distributed along the Hotelling line has the advantage of tractability and straightforward interpretation, it imposes a straitjacket on preferences and, thus, on demand. In particular, the sum of demand of both firms is perfectly price-inelastic (up to some upper bound of prices). Therefore, a price cut leads only to business stealing but does not increase overall demand in the market. Second, the preference structure is, by definition, perfectly negatively correlated; that is, the consumer who has the strongest preference for the content of platform 1 has the weakest preference for the content of platform 2, and vice versa. Therefore, it is impossible to analyze effects of content or ideology correlation on competition, which is often at the heart of media markets.

Two recent papers consider different preferences and demand. Ambrus, Calvano and Reisinger (2014) propose a model in which a consumer's choice to visit one platform is independent of the choice to visit another. This implies that a consumer chooses to visit the platform if her utility is positive – that is, $v = u_i - \gamma a_i \geq 0$, where u_i is the gross utility from being active on platform i , whereas γa_i represents the disutility from advertising, with γ being the nuisance parameter and a_i being the total advertising level of platform i . Advertisers are homogeneous, and the value of informing a consumer is normalized to 1. Platforms offer contracts to advertisers consisting of an advertising level m_i in exchange for a transfer payment t_i . An exclusive consumer of platform i becomes informed about an advertiser's product with probability $\phi_i(m_i)$, with ϕ_i strictly increasing and concave. This captures the idea that the consumer might be exposed to the same ad twice, implying decreasing marginal returns from advertising. Similarly, a consumer active on both platforms becomes informed with probability $\phi_{12}(m_1, m_2)$, with ϕ_{12} strictly increasing and concave in each argument and a negative cross-partial derivative.

In this framework, Ambrus, Calvano and Reisinger (2014) show that due to the concavity of the advertising technology, all advertisers accept the contracts of both platforms, so it can never be optimal for a platform to exclude some advertisers. With a mass 1 of advertisers, we have $m_i = a_i$. In equilibrium, a platform cannot extract the full surplus from advertisers even though advertisers are homogeneous. Instead, a platform can extract only the incremental surplus, which is the surplus that it delivers in addition to the surplus that an advertiser can obtain when rejecting the contract and being active only on the other platform. More formally, denote the demand of exclusive consumers by

$D_i(a, a_2)$, $i = 1, 2$, and the demand of multi-homing users by $D_{12}(a_1, a_2)$. An advertiser who accepts both contracts gets a benefit of

$$D_1(a_1, a_2)\phi_1(a_1) + D_2(a_1, a_2)\phi_2(a_2) + D_{12}(a_1, a_2)\phi_{12}(a_1, a_2) - t_1 - t_2.$$

Instead, when accepting the contract of platform j only, he gets

$$(D_j(a_1, a_2) + D_{12}(a_1, a_2))\phi_j(a_j) - t_j$$

because the advertiser reaches consumers only on platform j . As a consequence, platform i can demand only

$$t_i = D_i(a_1, a_2)\phi_i(a_i) + D_{12}(a_1, a_2)(\phi_{12}(a_1, a_2) - \phi_j(a_j)).$$

This leaves a profit of $D_{12}(a_1, a_2)(\phi_1(a_1) + \phi_2(a_2) - \phi_{12}(a_1, a_2))$ to each advertiser, which is positive due to the concavity of the advertising technology.

To understand how competition works in this model, compare the market with duopolistic platforms with a monopoly market. The monopoly platform can extract the full surplus of advertisers. There are two effects from entry. First, each multi-homing consumer can now get informed about an advertiser's product on both platforms. Therefore, the single ad is worth less, a *duplication effect*. Due to this effect, the advertising intensity in duopoly falls. However, there is also a more subtle countervailing effect. For the single platform, all consumers are exclusive consumers. Relative to overlapping consumers, these consumers are more valuable, and losing them is relatively costly for the platform. This curbs its incentive to increase the advertising level. By contrast, in duopoly, some of the lost business due to increased advertising levels comes from overlapping consumers. The duopolist shares business with its rival, and losing this shared business is less detrimental than losing exclusive consumers. As a consequence, due to this *business-sharing effect*, the duopolist has a greater incentive to increase the advertising level. This shows that competition is driven by the *composition* of consumer demand and not just by the mere *size* of the demand. The paper demonstrates that if the business-sharing effect dominates the duplication effect, advertising levels in a duopoly are larger than in a monopoly, and vice versa.

It is also possible to relate the strength of the business-sharing and the duplication effects to consumers' preferences and the content provided by the platforms. Suppose, for example, that platforms' contents are relatively similar in the sense that consumers who like the content of platform 1 have a high probability of also liking the content of platform 2. In this case of a positive consumer preference correlation, platforms have many overlapping consumers. If platforms are symmetric and have a similar advertising intensity, a platform, by reducing its advertising intensity, attracts many exclusive consumers relative to its current demand composition. Therefore, the gains from reducing advertising levels are relatively high, implying that platforms with positively correlated content are likely to have lower advertising levels. By contrast, suppose that platforms are news platforms and that one caters to a more right-wing audience and the other to a more left-wing audience. Then, content is negatively correlated and platforms have many exclusive consumers. By reducing the advertising intensity, a platform attracts mainly consumers with moderate preferences, and some of them will visit both platforms. These

consumers are, on average, less valuable than the existing ones, who are mainly exclusives. Therefore, each platform has only little incentive to lower the advertising intensity, which leads to a high equilibrium level of ads. Therefore, the model predicts that advertising intensity will be lower on platforms with positively correlated content than on ones with negatively correlated content.

Anderson, Foros and Kind (2014) analyze a different model with multi-homing consumers. Let us again focus on the case of two platforms and suppose that $D_i(a_1, a_2), i = 1, 2$, denotes the exclusive consumers of platform i , whereas $D_{12}(a_1, a_2)$ denotes the overlapping consumers. An advertiser can choose whether or not to advertise on a platform – that is, n_i equals either 0 or 1. Suppose that the value of informing an exclusive consumer equals ϑ , while the value of informing a multi-homing consumer equals $\vartheta(1+\lambda)$, with $\lambda \in [0, 1]$. A possible interpretation is that the probability with which a consumer becomes aware of the ad is x on each platform. An ad to the same consumer on another platform raises the chance that he becomes aware of the ad by $x(1-x)$. Therefore, normalizing the value of informing a consumer to 1, we have that $\vartheta = x$ and $\lambda = 1 - x$.

In Anderson, Foros and Kind's (2014) model, first, consumers form an expectation about the advertising intensity on each platform. Then, platforms set a price per ad, and advertisers rationally anticipate the number of consumers and choose to buy ads. Finally, consumers decide which platform to join. Due to the game structure, consumers do not react to the actual number of ads, which simplifies the solution of the game.²⁵

In equilibrium, each platform charges a price of (omitting arguments)

$$D_i\vartheta + D_{12}\vartheta\lambda,$$

and all advertisers are active on both platforms. Hence, platforms can extract the full surplus an advertiser obtains when informing an exclusive consumer, but only part of the surplus that an advertiser obtains when informing an overlapping consumer. An advertiser obtains a positive profit of $D_{12}\vartheta(1 - \lambda)$. Therefore, the principle of incremental pricing reemerges.

Anderson, Foros and Kind (2014) then focus on several important aspects of media markets that are influenced by the presence of multi-homing consumers. Consider the well-known problem of content duplication and suppose that each media platform has the choice of providing content A or B. If consumers only single-home and more than two thirds of them are interested in content A instead of content B, then both platforms will specialize in content A. This is because if one half of the consumers choose platform 1 and the other half platform 2 when both have the same content, then each platform

²⁵ The game involves passive expectations as in the classic model by Katz and Shapiro (1985). This form of expectation building implies that consumers form their expectations before observing the prices set by platforms. Hagiu and Halaburda (2014) and Belleflamme and Peitz (2014) use related assumptions in two-sided market models, in which all or a fraction of the agents of one side cannot observe the prices charged on the other side. The opposite assumption involves agents forming expectations after observing all prices and is (usually) denoted responsive expectations. For a discussion of the different implications of the two assumptions on termination charges of communication networks, see Hurkens and Lopež (2013).

gets more than one third of the consumers. By contrast, when choosing content B, a platform gets less than one third of the consumers. In general, duplication of content occurs in equilibrium if and only if

$$\frac{D_A}{2} > D_B,$$

where $D_j, j = A, B$, denotes the consumership for content j .

Now, suppose that consumers can multi-home. If both platforms have the same content, a fraction d_{12} of consumers multi-home. If both platforms choose content A, then each of them obtains a profit of

$$\frac{D_A}{2} \vartheta((1 - d_{12}) + d_{12}\lambda).$$

By contrast, if one platform chooses content B, it obtains a profit of $D_B\vartheta$. Comparing the two profits shows that choosing content A is preferred if and only if

$$\frac{D_A}{2}(1 - d_{12}(1 - \lambda)) > D_B.$$

Compared to the situation without multi-homing consumers, it is evident that content duplication occurs under a strictly smaller parameter range. Hence, if multi-homing consumers are present, the problem of content duplication is less severe.

To sum up, the work on multi-homing provides two important lessons. First, although platforms do not compete for advertisers when consumers are multi-homing, they cannot extract the full surplus from advertisers. This is because the consumer's first impression is usually more valuable than the second, and consumers can now be reached on multiple platforms, which leads to incremental pricing. Second, platforms do not care only about the size of demand, but also about how it is composed of single- and multi-homing consumers; thus, the composition of demand affects market outcomes.

4.2 Search engines and search bias

So far, we have simplified our presentation by assuming that users know which websites' content they are interested in and can access it directly. This implies that no intermediary is needed to help users select their preferred websites. However, the Internet offers a multitude of information, and finding the most relevant bits is often impossible without the help of a search engine. In fact, worldwide, there are billions of queries each day on different search engines. As described in Section 2, the most prominent one is clearly Google, with a market share of more than 90% in European countries and a global average of more than 80%. However, other search engines dominate in some countries; for example, in China, the search engine Baidu has a market share of more than 75%, whereas Google China has only slightly more than 15%.²⁶

If the only role of search engines were to efficiently allocate users to their preferred websites, then search engines would not be particularly relevant to this paper. In that limited role, the information

²⁶ See <http://www.chinainternetwatch.com/category/search-engine/> (accessed on April 25, 2014).

gatekeeper directs users to the appropriate media content, and its presence has no economic implications. However, the business model of a search engine, like that of most media platforms, centers around attracting users and obtaining revenues from advertisers. It is, therefore, not obvious that the incentives of users and search engines are perfectly aligned. In particular, search engines may bias their search results to obtain high revenues from advertisers. Suppose, for example, that a user wants to watch the video of a particular song and searches for it via Google. The video is available through multiple video portals, such as YouTube, MyVideo, or Clipfish, and Google can choose the order in which to display the search results. Since Google owns YouTube but not the other video portals, Google may have an incentive to bias its search results in favor of its own video portal and away from others.

Establishing such a search engine bias empirically is not always straightforward. Tarantino (2011) reports that, in response to a query with the keyword “finance,” Google lists Google Finance first, whereas Yahoo! lists Yahoo! Finance first. This suggests that at least one of the two search engines is biased if consumers on one search engine are comparable to those on the other.

Edelman and Lai (2013) consider the following quasi-experiment: In December 2011, Google introduced a tool called Google Flight Search, which helps users to search for a flight from A to B. When Google Flight Search appeared, it always appeared in a box at the top position. However, the appearance of Google Flight Search was very unsystematic, and minor changes in the entry could lead to the appearance or disappearance of the box.²⁷ Edelman and Lai (2013) estimate the change in the click-through rate when the Google Flight Search box appeared. They find that with the box, the click-through rate for paid advertising increased by 85%, whereas the click-through rates for non-paid search of competing travel agencies decreased by 65%. Therefore, the study provides evidence that search engines are able to influence user behavior.

Search engines usually have two different kinds of links, organic (or non-paid-for) links and sponsored (or paid-for) links. The organic links reflect the relative importance or relevance of listings according to some algorithm. The sponsored links are paid for by advertisers. As outlined in Section 2, selling those advertising slots represents the largest revenue source of Internet advertising.

The major commercial search engines sell the sponsored links via second-price auctions with a reserve price for each auction. Since a search engine observes whether or not a user clicks on a sponsored link, advertisers pay per click. Thus, the price is called the per-click-price (PCP). If an advertiser bids a higher PCP, this secures a rank closer to the top. However, the advertiser with the highest bid does not necessarily receive the first slot on top of page one of the search results. The search engine’s goal is to maximize revenue from selling slots and, therefore, it also takes into account the number of times users click on an ad. As a consequence, the search engine needs to estimate the click-through rate (CTR) and may put ads with a lower PCP in a higher position if their CTR is high. Google uses a quality score that reflects the estimated CTR to determine the slots for the respective advertisers.²⁸ There are several studies analyzing the auction mechanism in detail, including the seminal papers by Edelman, Ostrovsky

²⁷ For example, the box was shown when typing in “flights to orlando,” but it did not appear when searching for “flights to orlando fl.”

²⁸ For a more detailed discussion, see Evans (2008).

and Schwarz (2007) and Varian (2007). More-recent papers are Katona and Sarvary (2010), who analyze the interaction between sponsored and organic links, and Börgers, Cox, Pesendorfer and Petricek (2013), who explore the bidding behavior for sponsored links on Yahoo's search pages.

Do search engines list search results in the best interest of consumers? The economics literature has uncovered several reasons why search engines may have an incentive to bias their search results. We start with reasons that are to be considered even if a search engine is not integrated with a media platform. First, distinguishing between organic and sponsored links can provide one answer to why search engines bias their search results. As Xu, Chen and Whinston (2012), Taylor (2013a), and White (2013) point out, organic links give producers a free substitute to sponsored links on the search engine. Therefore, if the search engine provides high quality in its organic links, it cannibalizes its revenue from sponsored links. At the same time, providing better (i.e., more reliable) organic search results makes the search engine more attractive. If consumers have search costs, a more attractive search engine obtains a larger demand. However, if the latter effect is (partially) dominated by self-cannibalization, a search engine optimally distorts its organic search results.

Chen and He (2011) and Eliaz and Spiegel (2011) provide a further reason why search engines may bias their search results. Since the search engine obtains profits from advertisers, it is in its best interest that advertisers' valuation of sponsored links is high. This valuation increases if product market competition between advertisers is relatively mild. Therefore, the search engine may distort search results to relax product market competition between advertisers. In Chen and He (2011) and Eliaz and Spiegel (2011), the search engine has an incentive to decrease the relevance of its search results, thereby discouraging users from searching extensively. This quality degradation leads to lower competition between producers and, therefore, to higher prices.²⁹

We now turn to the case in which the search engine is integrated with a media platform (as is the case with YouTube and Google). Does this lead to additional worries about search engine bias, or can integration possibly reduce search engine bias? In what follows, we present the models of de Cornière and Taylor (2014) and Burguet, Caminal, and Ellman (2014) to systematically analyze the costs and benefits of search engine integration.

In de Cornière and Taylor's (2014) model, there are a monopoly search engine $i = 0$ and two media platforms $i = 1, 2$. The media platforms are located at the ends of a Hotelling line, with platform 1 located at point 0 and platform 2 at point 1. Users are distributed on the unit interval, but before deciding to search, they are not aware of their location. This implies that without searching, a user cannot identify which media platform has the content she is interested in most. A user incurs search costs s when using the search engine, where s is distributed according to a cumulative distribution function denoted by F .

Both the media platforms and the search engine obtain revenues exclusively from advertising. The quantity of advertising on website i is denoted by a_i . Users dislike advertising, implying that the disutility of a user who will be directed by the search engine to website i is $\gamma_i(a_i)$, which is strictly increasing. A

²⁹ See Xu, Chen and Whinston (2010, 2011) for related models.

user's utility is also decreasing in the distance between her location and the location of website i . The utility a user receives from website i is

$$v(d, a_i) = u(d) - \gamma_i(a_i) - s,$$

where d denotes the distance between the location of website i and the location of the user and $u'(d) < 0$.

The search engine works as follows: If a user decides to use the search engine, she enters a query. The search engine then maps the user's query into a latent location on the Hotelling segment and directs the user to one of the platforms. The search engine's decision rule is a threshold rule such that all users with $x < \bar{x}$ are directed to platform 1 and those with $x \geq \bar{x}$ are directed to platform 2.

Advertising is informative, and there is a representative advertiser. The expected per-user revenue of the advertiser is

$$R(a_0, a_1, a_2, \bar{x}) = r_0(a_0) + \bar{x} r_1(a_0, a_1) + (1 - \bar{x}) r_2(a_0, a_2),$$

where $r_i, i = 0, 1, 2$ represents the revenue from contacting users on the search engines or the respective media platform. A key assumption is that ads on the search engine and on the media platforms are imperfect substitutes. That is, the marginal value of an ad on one outlet decreases as the number of advertisements on the other outlet increases. Formally,

$$\frac{\partial^2 r_i(a_0, a_i)}{\partial a_0 \partial a_i} \leq 0.$$

This implies that the advertising revenue generated by a media platform falls if a_0 rises. The advertiser pays platforms on a per-impression basis, and the respective prices are denoted by p_i . Therefore, the expected per-user profit of an advertiser is

$$\pi_a = R(a_0, a_1, a_2, \bar{x}) - a_0 p_0 - \bar{x} a_1 p_1 - (1 - \bar{x}) a_2 p_2.$$

Given that a fraction of users D use the search engine, the profit of the search engine is $\pi_0 = D a_0 p_0$, while the profits of the media platforms are $\pi_1 = \bar{x} D a_1 p_1$ and $\pi_2 = (1 - \bar{x}) D a_2 p_2$, respectively. To simplify the exposition, de Cornière and Taylor (2014) keep a_0 fixed and focus on the choice of a_1, a_2 , and \bar{x} .

The timing of the game is as follows: First, media platforms choose their advertising levels a_1 and a_2 and the search engine chooses \bar{x} . In the second stage, the advertising market clears – that is, p_0, p_1 , and p_2 equalize demand and supply for each outlet. In the third stage, consumers decide whether or not to use the search engine. Finally, those consumers who use the search engine type in a query and visit the media platform suggested by the search engine.

When deciding whether or not to use the search engine, a consumer knows \bar{x} and has an expectation about the advertising levels on the media platforms, denoted by a_i^e . The expected utility of a consumer from using the search engine is, therefore, given by

$$E[v] = \int_0^{\bar{x}} u(z)dz + \int_{\bar{x}}^1 u(1-z)dz - \bar{x}\gamma_1(a_1^e) - (1-\bar{x})\gamma_2(a_2^e) - s.$$

Since the outside option of a consumer is normalized to 0, she will use the search engine as long as $E[v] \geq 0$, which determines the user demand D . We denote the value of \bar{x} that maximizes the expected consumer utility and, thus, the participation rate by x^D .

A *search bias* can then be defined as follows: The search engine is biased as long as $\bar{x} \neq x^D$. In particular, the search engine is biased in favor of media outlet 1 (media outlet 2) when $\bar{x} > x^D$ ($\bar{x} < x^D$).

When deciding about its optimal cutoff value \bar{x} , the search engine faces the following problem. First, it wants to have high user participation. Other things equal, a larger number of search engine users leads to higher profits because the willingness-to-pay of advertisers increases. Therefore, the search engine cares about relevance to users, but this is not the only important characteristic of concern to users. Users also dislike advertising, implying that they prefer to be directed to a site that shows a low number of ads. As a consequence, the fewer the advertisements on outlet 1 relative to outlet 2, the higher the search engine sets \bar{x} , and vice versa. These considerations align the incentives of the search engine with those of users. However, the search engine obtains profits from advertisers and, therefore, aims to maintain a high price p_0 for its own links. With that, a strategic consideration comes into play. If ads on media platform i are particularly high substitutes for ads on the search engine, the search engine prefers to bias results against this platform. This allows the search engine to keep p_0 high and to obtain higher profits. Formally, clearing of the advertising market implies that the price of ads is equal to the advertiser's marginal willingness-to-pay; that is (omitting arguments),

$$p_0 = \frac{\partial r_0}{\partial a_0} + \bar{x} \frac{\partial r_1}{\partial a_0} + (1-\bar{x}) \frac{\partial r_2}{\partial a_0}.$$

Therefore, p_0 is non-increasing in a_1 and a_2 . It follows that the search engine biases its results against the media outlet for which $\partial r_i / \partial a_0$ is more negative. As a consequence, too many consumers are directed to media outlet 1 if $\partial r_1 / \partial a_0 > \partial r_2 / \partial a_0$, and too many consumers are directed to media outlet 2 if the reverse holds true. If the two media platforms are symmetric, or if the advertising demands are independent, then no bias occurs.

De Cornière and Taylor (2014) then analyze the effects of integration of the search engine with one of the media platforms. Suppose that there is partial integration without control of ad levels. That is, media platform i shares a fraction ρ_i of its profit with the search engine but obtains full control with respect to the ad level a_i . In reality, this case is relevant for two reasons: First, Google owns subsidiaries, such as DoubleClick, which sell advertising technologies to media outlets. Therefore, the search engine already obtains some revenues from media outlets. Second, even when fully integrated, a media outlet might still be independently managed.

The profit of the media outlet is, then, $\pi_i = (1 - \rho_i)\bar{x}D a_i p_i$, which implies that the optimal advertising level is the same as in the non-integrated case since the profit function is just multiplied by a constant. The search engine's profit is now $\pi_0 = D (a_0 p_0 + \rho_i \bar{x} a_i p_i)$. There are two immediate consequences: First, the search engine has an incentive to bias its result in favor of media platform i because it benefits

directly from this platform's revenues. Second, it also benefits more from higher consumer participation, implying that the search engine wants to implement higher quality (i.e., less-biased results).

Because of these two potentially opposing forces, partial integration can increase or decrease the level of bias. In particular, if the search engine were biased to the detriment of platform i without integration, partial integration might mitigate this bias. However, even if the search engine were biased in favor of media outlet i without integration, partial integration could lead to a reduction in the bias because the search engine would care more about quality. De Cornière and Taylor (2014) show that this occurs if the elasticity of user participation with respect to quality is large.

When considering full instead of partial integration, similar effects are at work. Here, the ad level of the integrated media platform will fall after integration due to the internalization of the price effect – i.e., advertising levels are substitutes. With respect to the bias, the search engine again has the incentive to increase consumer participation. Therefore, it may benefit from a reduction in the bias. This result is most likely if the two media platforms are very different with respect to the substitutability of their ads to the search engine's ads; in other words, one is a close substitute, while the other is a mild substitute for the search engine's ads. In general, if the media platforms are symmetric, partial or full integration always leads to an increase in bias but can still benefit consumers through lower ad levels.

Burguet, Caminal, and Ellman (2014) propose a different setup to analyze the problem of search engine bias and integration. They do not allow for ad nuisance but explicitly model consumer search for advertisers' products. In what follows, we present a simplified version of their model, which nevertheless captures the main trade-off.

There is a mass one of users, indexed by i . Users are interested in the content of a website. There are N websites, indexed by n , and each user i has a specific content that matches her interests and generates net utility $u > 0$. This content is denoted by $n(i)$. Any other content generates a net utility of zero. Users do not know which website matches their interests. Therefore, they need the help of a search engine, denoted by 0 . The search engine can perfectly identify the relevant website $n(i)$ after a user has typed in the query. Websites are symmetric in the sense that each website's content interests the same fraction of users, $1/N$. When using the search engine, a user incurs search cost s , where s is distributed according to a c.d.f. denoted by F .

The search engine displays a link to a website after a user has typed in the query. The search engine can choose the probability that the link leads to the content matching the user's interest. We denote this probability by ϕ^O . The superscript O stands for organic link, representing the fact that the links to websites are non-paid and, therefore, organic.

The search engine and websites obtain profits from producers who pay to advertise their products. There are J different product categories, indexed by $j \in \{1, 2, \dots, J\}$. User i values only one category $j(i)$ and buys one unit. There are two producers of products $k \in \{1, 2\}$ in each category. Producer 1 provides the best match to a user, leading to a net utility of v_1 . Producer 2 provides only the second-best match, leading to a net utility of v_2 , with $0 < v_2 < v_1$. The margins earned by the producers are given by m_1 and m_2 , respectively. To simplify the exposition, Burguet, Caminal, and Ellman (2014) assume that users'

and producers' interests are misaligned, in the sense that for each category, $m_2 > m_1$.³⁰ Moreover, to simplify the welfare analysis in all categories, $m_1 + v_1 > m_2 + v_2$, implying that the social optimum involves only transactions of the best-match product. As above, categories are symmetric in the sense that each category's products interest the same fraction of users, $1/J$.

The search engine provides a single link after a user has typed in a query for product search in a particular category.³¹ It also has full market power and sets a pay-per-click price. The search engine chooses to display the link of producer 1 with probability ϕ^S and the link of producer 2 with probability $1 - \phi^S$. The superscript S stands for sponsored link, representing the fact that producers pay for the links to websites.³²

Users can also buy the products on websites. In particular, each website offers a "display-ad" slot for a link to a producer. If a user i visits website $n(i)$ (i.e., the website with the content she is interested in), she notices the ad with probability α , with $0 < \alpha \leq 1$. By contrast, if she is directed to a website different from the one with the content that interests her, the probability is $\alpha\beta$, with $0 < \beta < 1$. The targeting of websites is less accurate than that of the search engine. Formally, the website gets user i 's product category right with only probability $\sigma < 1$. That is, the website obtains a signal of each visiting user's product category interest, and this signal equals $j(i)$ with probability σ . The websites also have full bargaining power vis-à-vis producers.

The timing of the game is as follows: In the first stage, the search engine and the website announce prices to merchants. In addition, the search engine announces its design variables ϕ^O and ϕ^S . In the second stage, users decide whether or not to use the search engine. If they participate, they first search for content and can visit the website displayed by the search engine. When consuming the content, they may click on the ad for the producer displayed by the website and may buy the producer's product. They then either leave the market or type in a product query and can visit the website of the producer displayed by the search engine and buy the product.³³ Therefore, the search order is first content and then product.

Figure 8 depicts all actors in the model and the interaction between users' search for content and products. The solid arrows represent products and the dashed arrow content. Downward arrows

³⁰ This is an extreme assumption. However, for the main result to hold, it is sufficient that there are some categories for which this misalignment holds true.

³¹ In Burguet, Caminal and Ellman (2014) and in de Cornière and Taylor (2014), users visit only a single website after typing in a query. However, users sometimes click on multiple search results (in sequential order) broadly following the respective ranking of the results. This implies that advertisers exert externalities on each other, e.g., through bidding for more prominent placement. Athey and Ellison (2011) and Kempe and Mahdian (2008) provide models that explore the effects of these externalities on the optimal selling mechanism of the search engine.

³² The description of the game departs from the original model considered by Burguet, Caminal, and Ellman (2014). Instead of having only two producers in each category, Burguet, Caminal, and Ellman (2014) consider four, two for the best-match product and two for the second-best product. These firms are in Bertrand competition for links. The search engine runs a second-price auction with the twist that it partially discounts second-best products, to allow best-match products to win the auction. Burguet, Caminal, and Ellman (2014) show that the discount is set in such a way that all four producers choose the same bid after discounting. Then, ϕ^S represents the probability that the search engine chooses a best-match product as the winner.

³³ Note that users pay the search costs only once – that is, the search participation decision is a single one.

indicate links of advertisers (producers) and media platforms (websites), and upward arrows depict users visiting the respective websites.

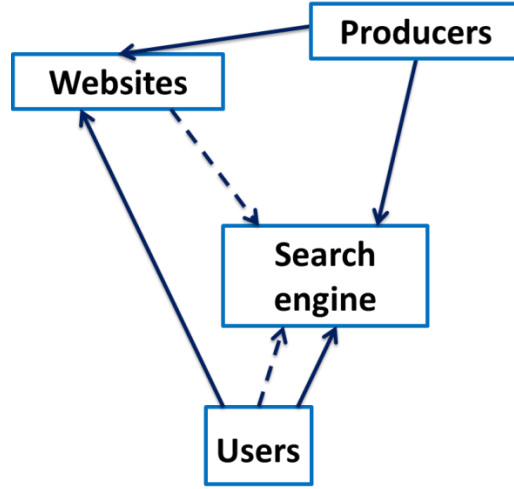


Figure 8: The media market in Burguet, Caminal, and Ellman (2014)

In this model, the *search bias* is defined as follows: The search engine is biased as long as $\phi^O < 1$ and/or $\phi^S < 1$. In particular, since $u > 0$, (i.e., directing a user to her preferred content provides a higher benefit than directing the user to any other content), a social planner optimally chooses $r^O = 1$. Moreover, since $m_1 + v_1 > m_2 + v_2$, a social planner optimally chooses $r^S = 1$.

The game can be solved by backward induction: Users' choices are immediate once stated, except for the participation decision. In particular, users click on the links provided by the search engine. In the product search stage, they buy the advertised product, whether it is the best match or the second-best, provided that it is in their preferred category. When a link to a preferred product is displayed on the website, a user buys the product only if it is the best match because she anticipates that she can buy at least the second-match product and, with probability ϕ^S , the best-match product in the next stage.

To determine the mass of participating users, we first determine the margins of the search engine and the websites. Since the search engine and the website make take-it-or-leave-it offers to producers, they will charge a price of m_1 to type-1 producers and m_2 to type-2 producers. This implies that the average margin of the search engine equals $\phi^S m_1 + (1 - \phi^S) m_2$. Turning to the website, it will display ads only for best-match products since users will never buy second-best match products when clicking on the link from websites. Since a website offers the product category a user is interested in with probability σ , and the user becomes aware of the link with probability α (given that the user is interested in the content of the website), the expected margin of the website is $\sigma \alpha m_1$. By contrast, if the user is not interested in the content of the website, she realizes the ad only with probability $\alpha \beta < \alpha$, implying that the corresponding margin is $\sigma \alpha \beta m_1$.

The share of users who buy from organic display advertising on websites is, then,

$$\mu = \sigma \alpha (\phi^O + (1 - \phi^O) \beta).$$

The utility of a user from using the search engine is given by

$$\phi^O u + (\mu + (1 - \mu)\phi^S)v_1 + (1 - \mu)(1 - \phi^S)v_2 - s.$$

Let us denote the critical value of s , at which the last expression equals zero, by \bar{s} . The mass of participating users is, then, $F(\bar{s})$. This yields the profit of a search engine

$$\pi_0 = F(\bar{s})(1 - \mu) (\phi^S m_1 + (1 - \phi^S)m_2). \quad (2)$$

What are the levels ϕ^O and ϕ^S that the search engine wants to set? Let us start with ϕ^O . Increasing ϕ^O raises participation because the search engine becomes more reliable on content search (the first term of (2), $F(\bar{s})$, increases). At the same time, increasing ϕ^O makes advertising on platforms more effective. Since the search engine and the websites compete for advertisers (producers), this reduces π_0 (the second term, $(1 - \mu)$, falls). Turning to ϕ^S , an increase in ϕ^S raises the reliability of the search engine with respect to product search and, therefore, increases user participation (the first term increases). However, because $m_2 > m_1$, the average margin of the search engine falls (the third term, $(\phi^S m_1 + (1 - \phi^S)m_2)$, decreases).

Burguet, Caminal, and Ellman (2014) show that the incentive of the search engine to distort content search and product search, starting from $\phi^O = \phi^S = 1$, depends on the ratio of the following terms:

$$\frac{u}{\sigma\alpha(1-\beta)m_1} \quad \text{versus} \quad \frac{v_1 - v_2}{m_2 - m_1}.$$

The first expression refers to the costs and benefits of distorting content search, while the second refers to the costs and benefits of distorting product search. When distorting content search (with no distortion of product search), consumer surplus falls by a rate u , but the advertising revenues of the search engine rise at a rate $\sigma\alpha(1 - \beta)m_1$.³⁴ Instead, distorting product search reduces consumer surplus at a rate $v_1 > v_2$ but increases the value for the search engine by $m_2 - m_1$. Burguet, Caminal, and Ellman (2014) show that, generically, the search engine will distort at most one type of search, setting the other at the optimal value. Specifically, if the expression on the left is larger than the one on the right, only product search is distorted, whereas only content search is distorted if the reverse holds true. Only if both expressions are the same might both searches be distorted.

Overall, this shows that even without integration of a website with the search engine, the search engine might have an incentive to distort search due to competition with websites for advertising. The question is, again, whether vertical integration with a website exacerbates this distortion or reduces it.

To see the incentives of the search engine under integration, suppose, first, that the search engine is integrated with all websites. Then, the profit of the search engine becomes

$$\pi_0 = F(\bar{s})[(1 - \mu) (\phi^S m_1 + (1 - \phi^S)m_2) + \mu m_1].$$

³⁴ This is because a reduction of r^S reduces the probability that a user buys a product through a display link on a website from $\sigma\alpha$ to $\sigma\alpha\beta$.

The search engine internalizes the externality exerted on websites by distorting ϕ^O or ϕ^S because it fully participates in the profits of the websites. This induces the search engine to improve its reliability, for both content and product search. Thus, the effect of integration is positive.

To understand the negative effects, consider the more realistic case in which the search engine is integrated with only a fraction of the websites. Then, it has an incentive to divert search from non-affiliated websites to affiliated ones. This leads to a different level of ϕ^O for affiliated websites than for non-affiliated ones. For example, if $\phi^O = \phi^S = 1$ without integration, then integration lowers consumer surplus because it may induce the search engine to reduce ϕ^O for non-affiliated websites.

To sum up, the literature on search engines shows that even without integration of a search engine with content providers, the search engine may have an incentive to bias search results. This bias occurs due to competition for advertisers between the search engine and content providers. Integration between the search engine and a content provider affects the way that competition for advertisers plays out; integration leads to higher or lower social welfare, depending on the circumstances.

4.3 Information spreading on the Internet

In the previous subsection, we restricted our attention to search engines as the only intermediaries between users and content-providing websites. However, there exist several other online channels allowing users to find out which website they are potentially interested in. In what follows, we discuss some of these channels and mechanisms, with a particular focus on their implications for the spread of information across the Internet. Specifically, we are interested in whether different users receive the same or differing information, according to the channel they use. Because few papers in the literature analyze these issues, we will confine our discussion to a description of the phenomenon and the tentative implications for competition and plurality, without presenting a rigorous analysis.

A popular way that users access content apart from using a search engine is to visit a news website and search for “most-read news” or “most-popular stories.” This device is offered by most news websites, such as BBC or Bloomberg, the websites of most newspapers, and also by video-sharing websites, such as YouTube. The standard way in which websites decide to classify content as most popular or as must-read news is by counting the absolute number of clicks on this content in the past (correcting for up-to-dateness and other factors). In this respect, the popularity of stories is similar to a classic network effect; that is, the more people read a story, the more attractive it will be to others.³⁵ The effect of most-popular stories is that users are more likely to obtain the same information. Even if users are heterogeneous and are interested in different content *ex ante*, the pre-selected content of websites is the same, and users access only content within this pre-selected sample. Therefore, users obtain the same information, which implies that they become more homogeneous regarding their information. This exerts a negative effect on plurality.³⁶ This issue is not (or to a much smaller extent) present in

³⁵ See Katz and Shapiro (1985, 1994) for seminal papers on network effects.

³⁶ Another effect is that the selection of the content usually depends on the absolute number of clicks but not on the time users spend on the website. Therefore, it is not clear if websites accurately measure how interesting the respective content is to users.

traditional media, in which the tool of counting the number of clicks and, therefore, a direct measure of popularity is not feasible.

Additionally, most-popular stories often have a tendency to be self-reinforcing as most-popular. If a story is recommended as highly popular, then more users will read it, implying that the number of clicks increases, thereby making the story even more popular. This effect is known as observational learning and is documented by, among others, Cai, Chen, and Fang (2009), Zhang (2010) and Chen, Wang, and Xie (2011). As a consequence, it is not obvious whether users read the same stories because they are actually interesting for a majority of users or if users read them merely because they are recommended.

Contrasting this hypothesis, Tucker and Zhang (2011) present a mechanism for why listing “most-read” stories can benefit stories with niche content or narrow appeal. Users usually have an ex-ante expectation if particular content is of broad versus narrow appeal. In this respect, a story with content that appeals to a majority of users is more likely to make it onto the most-read list. Now suppose that a story of broad-appeal content is ranked fourth on the most-read list, whereas a story of narrow-appeal content is ranked fifth. Since the broad-appeal story has a higher probability of being part of the most-read list, users will infer from this ranking that the narrow-appeal story is probably of higher quality or has particularly interesting insights. Therefore, if both stories are ranked almost equally, users will be more attracted by the story with narrow-appeal content. Tucker and Zhang (2011) test this hypothesis in a field experiment. A website that lists wedding service vendors switched from an alphabetical listing to one in which listings are ranked by the number of clicks the vendor received. They measure vendors as broad-appeal ones when located in towns with a large population and as narrow-appeal ones when located in small towns. Tucker and Zhang (2011) find strong evidence that narrow-appeal vendors, indeed, receive more clicks than broad-appeal vendors when ranked equally.

Oestreicher-Singer and Sundararajan (2012) also conduct an analysis to determine if popular or niche items benefit most from recommendations. In particular, they analyze the demand effects in recommendation networks by using data about the co-purchase network of more than 250,000 products sold on Amazon.com. They use the feature of Amazon.com to provide hyperlinks to connected products. To identify the effect that the visible presence of hyperlinks brings about, the authors control for unobserved sources of complementarity by constructing alternative sets of complementary products. For example, they construct a complementary set using data from the co-purchase network of Barnes & Noble (B&N). The B&N website provides a recommendation network similar to Amazon.com’s, but the product links might be different, and those on the B&N website are invisible to Amazon.com customers. Therefore, the products linked on the former website but not on the latter provide an alternative complementary set.³⁷ Oestreicher-Singer and Sundararajan (2012) find that visibility of the product network has very large demand effects – i.e., the influence that complementary products can have on the demand for each other can be up to a threefold average increase. Newer and more popular products use the attention induced by their network position more efficiently.

³⁷ Similarly, products that are linked in the future on the Amazon.com website but not today can be assumed to be complementary to the focal product today and can be used to construct an alternative complementary set.

The results of Oestreicher-Singer and Sundararajan (2012) differ from those of Tucker and Zhang (2011). In particular, the former paper finds that popular products benefit more from recommendation, while the latter find that niche products receive larger benefits. A potential explanation is that consumers may be more inclined to browse niche websites when looking for products for a special occasion (such as wedding dresses) than when looking for more-standard products. The contrasting findings could also reflect different reasons underlying the demand effect – i.e., attention in Oestreicher-Singer and Sundararajan (2012) and observational learning in Tucker and Zhang (2011).

Another way for users to access content is to read what other users recommend. For example, via the “share” command on twitter or other social media, users recommend content to their friends or followers (for some facts on users as curators, see Section 2.1). These friends are highly likely to read what the recommenders “like,” which is not necessarily what the majority of users are interested in or what friends of other users like. Therefore, in contrast to the “most-popular” stories, sharing content leads to different users obtaining different information and, therefore, does not necessarily lead to a reduction in plurality. However, users may access only content of a particular type because they largely ignore or are not aware of recommendations by users who are not their friends or whom they do not follow. In this respect, sharing content can lead to narrow or exaggerated views and is, therefore, prone to media bias.

It is evident that the flow and diversity of the information depends on the architecture of the (social) network. For example, the architecture of Twitter is similar to the star network, in which the user in the middle spreads information to all its followers. However, two followers may not necessarily exchange information directly with each other but only through the user they follow. By contrast, on Facebook, mostly groups of users interact, implying that there are more-direct links and direct information sharing among these users.³⁸

Most-read news and individual users sharing news are two extreme forms of spreading information on the Internet. Whereas the former depends only on the absolute number of clicks, the latter depends on a user’s subjective evaluation.³⁹ In between these two forms are recommendations provided by websites. These recommendations are based partly on content (as in the case of most-popular stories) and partly on the specific user (as in case of sharing information by users).

Regarding content, a website has many different forms of selecting recommendations to users. An extreme one is based purely on an algorithm, such as the absolute number of clicks in the past, and does not involve any editorial selection.⁴⁰ The other extreme is a purely curated selection, based on editorial policy. While the latter is evidently more subjective, it usually involves real journalism – that is, journalists becoming well informed about particular topics. An interesting question regards the benefits of these two forms of news selection for different classes of content. In particular, it is interesting to

³⁸ For detailed analyses of network formation, see, e.g., Jackson and Wolinsky (1996) or Jackson (2010). Banerjee, Chandrasekhar, Duflo, and Jackson (2014) provide a recent analysis of how gossip spreads within a network.

³⁹ Thus, to formalize the former, standard models with aggregate network effects can be used, whereas for the latter, the link structure of the social network has to be taken into account.

⁴⁰ For example, this is the case with Google News. For a more detailed discussion on news aggregators, see Section 3.3.

explore whether both types can survive for a particular content category since users value differentiation and/or multi-home, or whether one type tends to become dominant. Regarding recommendations based on the specific user, a website is informed about the history of the user's browsing behavior. Therefore, it can tailor its recommendation to this behavior and recommend stories or news according to the user's past preferences. In contrast to the case when other users recommend stories or news, here, a user's own past behavior determines the stories that she becomes aware of. Again, this may lead to a loss in plurality.

These devices to obtain information compete with search engines. As an example, consider Amazon versus Google. Many users searching for books now directly use Amazon's website and no longer search on Google. In the case of media, a similar pattern can be observed, with users who are looking for news bypassing the search engines and immediately visiting the website of their preferred news provider. An interesting question is how such behavior affects the bias of search engines and (potentially) of news websites.

5. Media platforms matching advertising to content

The success of a firm's advertising campaign is driven mainly by the effectiveness of its ads. Foremost, the recipients of the ads (i.e., the potential consumers) should be primarily individuals or companies with an inherent interest in the firm's products. Otherwise, informing potential consumers about characteristics of the product is unlikely to lead to actual purchases. To increase advertising effectiveness and reduce wasted impressions, firms match their advertisements to content on media outlets in such a way that consumers who are interested in the content are also likely to be interested in the advertised product. This practice is called *content matching* or *tailoring* and is a particular instance of targeting.

Consider, for example, a local bookstore. The store has higher returns from placing ads in a local newspaper than in a global one. The local newspaper is read by local audience, which consists of the potential consumer group for the bookstore. By contrast, a large portion of ads in the global newspaper are wasted since many readers do not live in the vicinity of the bookstore. Similar examples apply to content instead of geography. The advertisements of a cosmetics company are usually more effective in a women's magazine than in a computing magazine, and a sports apparel manufacturer's ads are likely to be more effective during televised sports than during a comedy show. However, in traditional media, the degree and effectiveness of such tailoring is limited. As argued by Goldfarb (2014), for example, the distinguishing feature of Internet advertising is its reduction in targeting costs compared to traditional media.

On the Internet, targeting is not limited to linking advertising to specific content. Advertisements can be targeted to the intentions of the consumer (reader/viewer/listener) as inferred from past behavior or based on specific circumstances, such as the weather conditions at the consumer's location. For example, media platforms can expose different users to different advertisements, even when those users browse the same website at the same time. The particular advertisement can be conditioned on many different parameters. For instance, the website may engage in geo-targeting and display advertisements relevant to the user's geography (inferred from IP addresses). Similarly, the website may keep track of ad exposure to users, thereby reducing repetitious exposure of ads, or search engines may

display ads conditional on queries conducted, a practice called keyword advertising. Clearly, both practices lower the number of wasted impressions, allowing the website or search engine to charge higher prices to advertisers, everything else given. A highly debated form of targeting is called behavioral targeting. Here, a website customizes the display advertisement to information collected in the past about a user. The website uses cookies based on pages that the user has visited and displays ads that could be of particular interest to the user;⁴¹ cookies are small pieces of data sent from a website, which track the user's activities. These cookies give precise information about the user's past web-browsing behavior and, therefore, about her preferences. We analyze implications of behavioral targeting in Section 6.

In Subsection 5.1, we discuss different formalizations of targeting (in the wider sense), focusing on tailoring on the Internet and how it differs from general tailoring. In Section 5.2, we then discuss the practice of "keyword targeting" in more detail.

5.1 How to formalize targeting

In the economics literature, targeted advertising has been shown to be able to segment the market. Esteban, Gil and Hernandez (2001) consider targeted advertising by a monopolist and show that the monopolist will target primarily consumers with high reservation values, thereby extracting a higher surplus.

Targeting also affects market outcomes under imperfect competition between advertisers. In particular, segmentation due to targeting may relax product market competition and, thus, allow firms to charge higher equilibrium prices. Iyer, Soberman and Villas-Boas (2005) consider a model with two competing firms that need to advertise to inform consumers about the existence of their products. There are three different consumer segments. Consumers belonging to the first have a high preference for the first firm in that its members consider buying only from that firm; those belonging to the second have a high preference for the second firm; and those belonging to the third are indifferent between the firms, and buy the lower-priced product.⁴² Advertising is costly to firms. Iyer, Soberman and Villas-Boas (2005) show that without targeting, equilibrium profits are zero because firms spend their entire product market profit to inform consumers. By contrast, with the possibility of targeting consumers, firms advertise with a higher probability to the market segment that prefers the firm's product than to the indifferent consumers, enabling the firms to reap strictly positive profits. Roy (2000) and Galeotti and Moraga-Gonzalez (2008), analyzing different models, also show that targeting can lead to full or partial market segmentation, allowing firms to obtain positive profits.⁴³

In what follows, we provide a more detailed discussion of the models by Athey and Gans (2010) and Bergemann and Bonatti (2011); both works explicitly consider targeting strategies on the Internet. The

⁴¹ For an in-depth discussion and analysis of behavioral targeting, see Chen and Stallaert (2014).

⁴² Chen, Narasimhan, and Zhang (2001) consider a similar demand structure to analyze the implications of targeting. In contrast to Iyer, Soberman and Villas-Boas (2005), they assume that firms can charge different prices to the consumers in different segments and show that imperfect targeting softens competition.

⁴³ Other models of targeted advertising include van Zandt (2004), who analyzes information overload; Gal-Or and Gal-Or (2005), who analyze targeting by a common marketing agency; and Johnson (2013), who studies ad-avoidance behavior by consumers when targeting is possible. The latter will be analyzed in Section 6.1 below.

former focuses on the supply side and keeps consumer demand simple, whereas the latter explicitly models the demand side and keeps the supply side simple.⁴⁴ We then briefly discuss the model by Rutt (2012).

Athey and Gans (2010) present a model that is cast in terms of geo-targeting. However, it can easily be adjusted (within limits) to other forms of targeting. Specifically, consider a set of localities $x \in \{1, \dots, X\}$, where each locality consists of N consumers. In each locality, there is one local media outlet. There is also one general outlet denoted by g , which is active in all locations. Consumers single-home – that is, they visit only one outlet. The market shares for local and global outlets in each locality are the same and given by n_g for the global outlet and $N - n_g$ for the local outlet.

Each advertiser i is only local and, therefore, values only impressions to consumers in the respective locality. The value to advertiser i of informing a consumer is v_i . Outlets track advertisers, which implies that they offer each advertiser a single impression per consumer. There is a continuum of advertisers with values $v_i \in [0,1]$ with cumulative distribution function $F(v_i)$. Each outlet chooses the number of ads, $a_j, j \in \{1, \dots, X, g\}$, that can be impressed on a consumer. We denote by p_j the impression price of outlet j . Finally, the probability of informing a consumer with an impression on outlet j is given by ϕ_j . There is no nuisance of advertising.

For each local outlet l , the probability of informing a consumer equals 1. Instead, for the global outlet, this probability depends on targeting being possible or not. If targeting is not possible, this probability is $\phi_g = 1/X$ because market shares are the same in all localities. By contrast, if targeting is possible, the probability is $\phi_g = 1$.

Solving for the equilibrium number of ads when targeting is not possible, the first observation is that an advertiser will buy impressions on outlet j if $\phi_j v_i \geq p_j$. Since $\phi_l = 1$ for local outlets, the total demand for impressions to a given consumer is $1 - F(p_l)$ for a local outlet l . In equilibrium, demand equals supply, implying that $1 - F(p_l) = a_l$ or $p_l = F^{-1}(1 - a_l)$ for a local outlet in a given locality. The profit function of outlet l is, therefore, $a_l F^{-1}(1 - a_l)$, which is to be maximized over a_l . We now turn to the global outlet. An advertiser will buy impressions on outlet g if $v_g \geq X p_g$. Since a_g are the impressions *per consumer* and there are X localities, the overall demand is $X(1 - F(X p_g))$, leading to an equilibrium that is characterized by $X(1 - F(X p_g)) = a_g$ or $p_g = \left(\frac{1}{X}\right) F^{-1}\left(1 - \left(\frac{1}{X}\right) a_g\right)$. The profit function is $a_g \left(\frac{1}{X}\right) F^{-1}\left(1 - \left(\frac{1}{X}\right) a_g\right)$, which is to be maximized over a_g .

It is easy to see that the problem of the global outlet is the same as that of the local outlets, adjusted by a scaling factor. Hence, $a_g^* = X a_l^*$ and $p_l^* = X p_g^*$, implying that per-consumer profits are the same. If targeting is possible, the problem of the global outlet becomes exactly the same as that of the local outlet. In this case, the number of ads and the ad price are the same for both types of outlets.⁴⁵ Athey and Gans (2010) obtain that without targeting, the global outlet expands its number of advertisements

⁴⁴ For an in-depth discussion of the different parameters influencing supply and demand of Internet advertising, see Evans (2009).

⁴⁵ Note that the problems of local and global outlets are separated, implying that a change in the number of ads of one outlet does not affect the number of ads on the other.

to X times what a local outlet would provide due to wasteful impressions. Therefore, the price it charges is only $1/X$ times that of a local outlet. However, per-consumer profits are not affected, and the global outlet replicates the outcome of the local outlet.

To easily grasp the trade-off in the model of Athey and Gans (2010), it is instructive to look at the model for the case in which advertising space is fixed for the global outlet. Is the advertising price of the global outlet with targeting higher or lower than without? The obvious effect is that advertising on the global outlet is less effective since advertisements are mismatched with probability $(X - 1)/X$. By contrast, advertising on a local outlet is effective with probability 1. In addition to this *efficiency effect*, there is also a *scarcity effect*. Without targeting, an advertiser from a locality competes with advertisers from other localities for scarce advertising space. This increases the price on the global outlet. Formally, comparing the advertising price on the global outlet with and without targeting gives

$$F^{-1}(1 - a_g) > \frac{1}{X} F^{-1}\left(1 - \frac{1}{X} a_g\right).$$

The advertising price with targeting is higher than that without targeting only if the last inequality is satisfied. As shown by Athey and Gans (2010), this holds true as long as a_g is not particularly high.

Athey and Gans's (2010) model demonstrates that targeting primarily allows an outlet to reduce wasteful impressions. As long as there are no costs of these impressions, targeting does not help an outlet to achieve higher profits. However, under many circumstances, there are such costs. For example, in most of the models discussed above, consumers dislike advertising. If there are nuisance costs to advertising, consumer demand is lower, the larger the number of ads. Targeting then reduces this problem and allows the global outlet to realize higher demand. Athey and Gans (2010) provide other reasons for such costs of impressions. Suppose, for example, that there is a constraint on advertising space that prevents the global outlet from just raising its number of impressions. Targeting then makes the use of the scarce advertising space more effective and allows the global outlet to reap higher profits. (A similar reasoning holds for the case in which providing advertising space is costly.) Alternatively, in the model presented above, demand across localities was assumed to be homogeneous. However, a more realistic model would consider heterogeneous demand, so that the global outlet has higher demand in some localities than in others. This implies that advertisers in these localities have a higher willingness-to-pay for advertising space. Thus, targeting allows the global outlet to price discriminate between advertisers of different localities and obtain higher profits.

It is worth mentioning that targeting does not necessarily increase profits for the global outlet. Consider an extension of the basic model in which outlets compete for advertisers. This could be due to the fact that advertisers value, at most, one consumer impression. As Athey and Gans (2010) show, targeting can spur competition between local and global outlets because the two types of outlets are vertically differentiated without targeting. When implementing targeting, both outlets provide a similar service to advertisers, leading to reduced prices. As a consequence, profits may fall with targeting. Anecdotal evidence of excessively fine targeting reported by Levin and Milgrom (2010) supports the relevance of this result.

Athey and Gans's (2010) model focuses on the supply side and reveals that increasing the supply of advertising can be a substitute for targeting. Therefore, targeting is particularly effective if an outlet can increase its advertising space only by incurring a cost.

Bergemann and Bonatti (2011) pursue a different route by modeling the demand side in a detailed way and keeping the supply side as simple as possible. In particular, they explicitly introduce the idea that targeting on the Internet allows for unbundling of content, thereby splitting a single advertising market into multiple ones. For example, readers of a traditional newspaper have to buy the whole newspaper to access the content they are interested in. Therefore, advertisers with niche products will probably find it too expensive to place an ad. By contrast, online consumers may access (and pay for) only selected articles. This implies that a producer of a niche product may find it profitable to pay for an ad that targets only the consumer group interested in the particular article.⁴⁶ A similar effect holds for Internet-TV. Major broadcasting networks usually focus on the taste of the masses to increase their advertising revenues. For example, sports channels do not devote much air time to niche sports, such as shot put or weight lifting. However, followers of these sports can access reports online and watch them at any time. This allows small businesses whose target groups are people interested in such niche sports to access these advertising markets.

Bergemann and Bonatti (2011) examine these effects and their implications for offline versus online media. Suppose that there is a continuum of products and a continuum of advertising markets. A product is denoted by y and is produced by a single firm y with $y \in [0, \infty)$. Similarly, advertising markets are denoted by $z \in [0, \infty)$. There is a continuum of buyers with a mass of one. Each buyer has a preference for a particular product and is located in one advertising market. The joint distribution of consumers across advertising markets and product markets is $F(z, y)$ with density $f(z, y)$. The fraction of consumers interested in product y can be written as $f(y) = \int_0^\infty f(z', y) dz'$, when integrating over all advertising markets. Similarly, the size of advertising market z can be written as $f(z) = \int_0^\infty f(z, y') dy'$, integrating over all products. The conditional distribution of advertising markets for a given product y is $f(z|y) = f(z, y)/f(y)$. Product differences can be expressed by differences in the size $f(y)$ to distinguish mass from niche products.

Each firm y can inform consumers about its product by sending a number of advertising messages $a_{z,y}$ in advertising market z . Each message reaches a random consumer with a uniform probability, as in the model of Butters (1977): With probability

$$\text{pr}(a_{z,y}, f(z)) = 1 - e^{-a_{z,y}/f(z)},$$

a given consumer in advertising market z of size $f(z)$ becomes aware of product y .

In each advertising market z , the supply of messages M_z is fixed. This supply is proportional to the size $f(z)$ of the advertising market – that is, $M_z = f(z)M$. Here, M can be interpreted as the average time a consumer spends on advertising messages. In each advertising market, there are a large number of

⁴⁶ This phenomenon has been called the “long tail of advertising”; see Anderson (2006). It also applies to keyword advertising and behavioral targeting.

media outlets. Outlets act as price takers, implying that a firm y can purchase messages at a price p_z in each market. The profit of firm y can then be written as⁴⁷

$$\pi_y = \int_0^\infty [f(z, y)\text{pr}(a_{z,y}, f(z)) - p_z a_{z,y}] dz.$$

To easily distinguish between mass and niche products, Bergemann and Bonatti (2011) impose that $f(y) = \alpha e^{-\alpha y}$. Here, a larger parameter value α represents a more concentrated product market. With this formulation, firms can be ranked in decreasing order of market size – i.e., a firm with a higher index y is smaller in the sense that fewer consumers are interested in its product.

The advertising markets are also ranked according to the mass of consumers interested in the market. Advertising market 0 is a large market in which all advertisers are interested, and advertising markets become smaller and more specialized with an increasing index. To formalize this, suppose that firm y is interested only in consumers in markets with $z \leq y$. For each firm y , the advertising market $z = y$ is the one with the highest density of consumers, conditional on market size. The conditional distribution of consumers with interest in product y over advertising markets z is given by the following truncated exponential distribution:

$$\frac{f(z, y)}{f(y)} = \begin{cases} \beta e^{-\beta(y-z)} & \text{if } 0 < z \leq y, \\ 0 & \text{if } y < z < \infty, \end{cases}$$

for all advertising markets $z > 0$. There is a mass point at $z = 0$ and the conditional distribution is $f(z, y)/f(y) = e^{-\beta z}$ if $z = 0$. The parameter β measures the concentration of consumers in advertising markets. We will explain below how the possibility of targeting consumers can be measured by β . Combining the definition of market size with the conditional distribution gives the unconditional distribution

$$f(z, y) = \begin{cases} \alpha \beta e^{-(\alpha+\beta)y} e^{\beta z} & \text{if } 0 < z \leq y, \\ 0 & \text{if } y < z < \infty, \end{cases}$$

with a mass point at $z = 0$, where the unconditional distribution is $f(z, y) = \alpha e^{-(\alpha+\beta)y}$. The market size can then be calculated by integrating over the population shares. Since consumers who are potential buyers of product y are present in all advertising markets $z \leq y$ but not in advertising markets $z > y$, we have

$$f(z > 0) = \int_z^\infty \alpha \beta e^{-(\alpha+\beta)y} e^{\beta z} dy = \frac{\alpha \beta}{\alpha + \beta} e^{-\alpha z}$$

and

$$f(z = 0) = \int_0^\infty \alpha e^{-(\alpha+\beta)y} dy = \frac{\alpha}{\alpha + \beta}.$$

The distribution of consumers across product and advertising space has a natural interpretation in terms of specialization of preferences and audiences. First, a product with a larger index is a more specialized product in the sense that there are fewer potential buyers. Similarly, an advertising market with a higher

⁴⁷ The value of informing a consumer is normalized to 1.

index z is a market with a more narrow audience. Second, potential consumers of larger firms are distributed over a smaller number of advertising markets. This can be seen by the assumption above that $f(z, y)$ is positive only for $z \leq y$. For example, potential buyers of product $y = 0$ are concentrated in the advertising market $z = 0$. Interpreting advertising markets as media outlets, this implies that a consumer interested in a mass product does not visit a website with advertisements for niche products.

Third, the variable β , ranging from 0 to ∞ , captures in a simple way the ability of firms to target consumers. For example, as $\beta \rightarrow 0$, all consumers are concentrated in advertising market 0, implying that there is single advertising market. By contrast, as $\beta \rightarrow \infty$, then all potential buyers of product y are in advertising market y , and so there is perfect targeting. In general, an increase in β implies that consumers are spread over more advertising markets and can be better targeted by firms. Overall, the highest conditional density of potential consumers of firm y is in advertising market $z = y$. As β gets larger, more consumers move away from the large advertising markets (near $z = 0$) to the smaller advertising markets (near $z = y$).

To illustrate how the model works, let us look at the benchmark case in which all consumers are present in a single advertising market $z = 0$. We solve for the equilibrium amount of advertising and the equilibrium price. Since there is a single advertising market, we drop the subscript z in the notation. The profit function of firm y is then $\pi_y = f(y)\text{pr}(a_y) - pa_y$. Determining the first-order conditions and using the definition of $f(y)$ yields

$$a_y = \begin{cases} \ln(f(y)/p) & \text{if } f(y) \geq p, \\ 0 & \text{if } f(y) < p. \end{cases} \quad (3)$$

As is evident, firms with a larger market size optimally choose a larger amount of advertising. Therefore, in equilibrium, only firms with the largest market size find it optimal to advertise. Let M be the total number of advertising messages and denote by Y the marginal advertiser. The market-clearing condition is given by $\int_0^Y a_y dy = M$. Using demand for ads given by (3) and $f(y) = \alpha e^{-\alpha y}$ yields

$$\int_0^Y \left(\ln\left(\frac{\alpha}{p}\right) - \alpha y \right) dy = M.$$

Using $a_Y = 0$ together with the last equation, we can solve for the equilibrium price and the marginal advertiser. This gives $p^* = \alpha e^{-\sqrt{2\alpha M}}$ and $Y^* = \alpha\sqrt{2M/\alpha}$. Inserting back into the demand function of advertiser y yields

$$a_y^* = \begin{cases} \alpha\sqrt{2M/\alpha} - \alpha y & \text{if } y \leq Y^*, \\ 0 & \text{if } y > Y^*. \end{cases} \quad (4)$$

Therefore, only the largest firms advertise, and the equilibrium number of advertising messages is linearly decreasing in the rank y of the firm. As the concentration in the product market measured by α increases, fewer advertising messages are wasted, leading to an increase in social welfare. In particular, the allocation adjusts to firms with a larger market size, implying that fewer firms advertise as α increases.

To analyze the effect of targeting, Bergemann and Bonatti (2011) examine the situation with a continuum of advertising markets and a positive targeting parameter $\beta \in (0, \infty)$. The allocation of advertising messages is then given by a generalization of (4):

$$a_{z,y}^* = \begin{cases} \alpha\beta e^{-\alpha z} (\sqrt{2M/(\alpha + \beta)} - (y - z)) & \text{if } z > 0, \\ \alpha\sqrt{2M/(\alpha + \beta)} - \alpha y & \text{if } z = 0. \end{cases}$$

Does targeting improve social welfare and do firms benefit from targeting? Bergemann and Bonatti (2011) show that the social value of advertising is increasing in the targeting ability β . The intuition is that targeting increases the value of advertising for a firm y in its “natural” advertising markets $z \approx y$. This leads to an increased volume of matches between firms and potential consumers, which improves social welfare.

However, looking at the cross-sectional implications of targeting, not all firms benefit from improved targeting. In particular, only the small firms that are not active in the mass advertising market $z = 0$ and the largest firms, which are primarily active in that market, benefit. By contrast, medium-sized firms, which are active in the mass market and also in several others markets $0 < z \leq y$, are hurt. To grasp the intuition behind this result, observe, first, that for small firms (those not active in market $z = 0$), the mass of potential buyers in their natural advertising markets $z \approx y$ increases, allowing them to reach a larger fraction of consumers. A similar effect is present for large firms. Their customers are concentrated in a small number of markets, and an increase in the targeting ability increases the chances of achieving a match. By contrast, medium-sized firms are hurt by the decrease in consumers participating in market $z = 0$, and this decline cannot be compensated by the rise in participation in their natural markets $z \approx y$.

The model can be used to analyze the implications of targeting for “online” versus “offline” media. In the offline medium, there is only a single advertising market, whereas there is a continuum of advertising markets in the online medium. For simplicity, suppose that the online medium allows for perfect targeting of advertising messages to consumers. Consumers are dual-homing and spend a total amount of M_1 on the offline medium and M_2 on a single online market z . More specifically, $f(z)M_2$ is the supply of advertising messages in each targeted market z . So, the online medium consists of a continuum of specialized websites that display firms’ advertisements. There is competition between the two media because each firm views the advertising messages sent online and offline as substitutes due to the risk of duplication. Bergemann and Bonatti (2011) show that the price for offline advertising decreases in M_2 , reflecting the decreased willingness-to-pay for regular ads if a better-targeted market is present. The price for online advertising decreases in M_1 only on those websites that carry advertisements of firms that are also active offline. However, advertising markets with a high index z carry only the advertisements of niche firms, which are not affected by the allocation in the offline medium because they do not advertise there. This implies that online advertising reaches new consumer segments that are distinct from the audience reached by offline advertising.

Suppose, in addition, that each consumer is endowed with an amount of time equal to M and allocates a fraction σ of this time to the online medium. This implies that $M_1 = (1 - \sigma)M$ and $M_2 = \sigma M$. It is now possible to analyze what happens when consumers spend more time online – that is, when σ increases.

The effect on the offline advertising price is then non-monotonic. If σ is low (i.e., online exposure is low), the marginal willingness-to-pay for offline advertising falls because online advertising is more efficient. This induces a decrease in the offline advertising price, although the supply of offline advertising messages decreases. However, as σ increases further, the composition of firms active in offline advertising changes. In particular, only the largest firms display advertising messages offline, implying that the marginal advertiser has a high willingness-to-pay. This leads to an increase in the offline advertising price with σ . With regard to firm revenue, this implies that if consumers spend more time online, firms that are active solely in the online market unambiguously benefit. These are rather small firms. The effect on large firms is ambiguous: Since they are active on the offline medium, they may pay a larger advertising price, which reduces their profits.

Rutt (2012) proposes a different formalization of targeting. He considers a model with n platforms which are distributed equidistantly on a circle, single-homing users, and multi-homing advertisers. A user's valuation for an advertiser's product is binary, namely either of high or of low valuation (with the low valuation being set equal to 0).⁴⁸ Advertisers are uncertain about the true valuation. In particular, advertiser j does not know consumer i 's valuation for her product with certainty but only has an expectation. Each advertiser receives an informative signal about a consumer's true valuation. The realization of the signal induces an advertiser to update the expectation. The targeting technology can now be modelled as a change in the informativeness of the signal.⁴⁹ In the extreme case that targeting is impossible and signals are pure noise, the updated expectation equals the prior expectation. By contrast, when the signal is perfect, the advertiser knows the consumer's valuation with certainty. As a result, an increase in the informativeness leads some advertisers to revise their beliefs upwards, while others revise them downwards. Advertisers receiving a positive signal become more optimistic, whereas advertisers with a negative signal become more pessimistic.⁵⁰

With regard to the timing, platforms first select their advertising levels and advertisers then submit their conditional demands specifying which types of users they wish to be matched with for a given advertising price. For example, advertisers announce that they want to have ads displayed only to users who show an interest with a probability higher than some cut-off level. After users have decided which media platform to consume from, advertisers observe the signals and update their valuations. The advertising price then adjusts to clear the market. Afterwards, consumers enjoy the media content, observe the advertisement, and make purchases.

In this setting, Rutt (2012) shows that targeting increases the advertising price and, thereby, allows platforms to receive higher profits. This effect is the stronger, the more competitive is the market and the more users are averse to ads. This is because in markets satisfying these conditions, the equilibrium

⁴⁸ Pan and Yang (2014) consider a similar market structure with two platforms to analyze the effects of targeting; they specify user demand and, thus, improved targeting differently.

⁴⁹ A simple example of a signal structure which fits this description is the truth-or-noise information structure: suppose that the prior expectation about consumer i having a high valuation for the product is X , the signal is Y , and the signal reveals the true consumer valuation with probability Z . Then, the posterior expectation is $ZY + (1 - Z)X$.

⁵⁰ In this respect, the information structure is similar to the demand rotation considered in Johnson and Myatt (2006).

features only few advertisements, implying that the marginal advertiser received a particularly high signal. As a consequence, improved targeting increases this advertiser's willingness-to-pay, resulting in a strong increase in the advertising price. In case of free entry of platforms (at some set-up costs), improved targeting may exacerbate excessive entry and leads to insufficient advertising due to high prices. Overall, the model therefore predicts a heterogeneous effect of targeting in different media markets.

The models of Athey and Gans (2010) and Bergemann and Bonatti (2011) share the assumption that offline and online advertising are substitutes.^{51, 52} If a potential consumer can be contacted via an advertisement through one channel, a conditional contact through the other channel is worth less. Goldfarb and Tucker (2011a) provide empirical support. They use data on estimated advertising prices paid by lawyers to contact potential clients with recent personal injuries. Goldfarb and Tucker (2011a) exploit state-level variation in the ability of lawyers to solicit those customers. In particular, this "ambulance-chasing" behavior is regulated in some states by the state bar associations, which forbid written communication (including direct electronic communication via e-mail) with potential clients for 30-45 days after the accident. Goldfarb and Tucker (2011a) use data on estimated auction prices of 139 Google search terms on personal injury in 195 regional city markets to analyze the effects of these regulations.⁵³ They find that, compared to the prices for personal injury keywords in non-regulated states, in states with solicitation restrictions, such keyword prices are between 5% and 7% higher. Therefore, when offline marketing communication is not possible, firms appear to switch to online advertising. This suggests that there is substitution between the two forms of advertising. In addition, Goldfarb and Tucker (2011a) demonstrate that this effect is larger in locations with a small number of potential clients. One interpretation is that mass-media advertising may not be cost-effective when consumers are hard to reach. Therefore, the possibility of direct offline advertising is particularly valuable. If this advertising channel is closed, online advertising prices rise. This implies that in markets with fewer customers, online advertising allows firms to reach the hardest-to-find customers, thereby lending support to the "long tail" hypothesis in Internet advertising (Anderson, 2006).

We also note that targeting on the Internet can take many different forms beyond geographical or contextual targeting. For example, social targeting has become increasingly popular. Socially-targeted ads, when displaying the ad to a particular user, refer to another user (e.g., a friend on Facebook). More precisely, a social ad is an online ad that "incorporates user interactions that the consumer has agreed to display and be shared. The resulting ads display these interactions along with the user's persona (picture and/or name) within the ad's content" (IAB, 2009). Tucker (2012) compares the effectiveness of socially-targeted ads to that of conventionally (demographically) targeted ads and non-targeted ads. She

⁵¹ Rutt (2012) considers only online platforms. When interpreting some platforms as online and others as offline media outlets in his model, online and offline advertising would be independent. This is because users only visit a single platform and so advertisers can reach each a particular user exclusively via online or via offline advertising.

⁵² More generally, there is a link between offline and online in the sense that advertising offline affects consumer purchase online. In particular, Liukonyte, Teixeira, and Wilbur (2014) study the effect of television commercials on actual purchase behavior on the Internet. They find a positive effect, even immediately after a viewer is exposed to the commercial.

⁵³ Search terms on personal injury can be identified objectively because there is a precise legal definition by the bar association.

conducted a field experiment on Facebook involving a non-profit charity organization that provides educational scholarships for girls to attend high school in East Africa. The non-profit organization launched a standard advertising campaign and a social variant of it. In this social variant, the recipients of the ads were only Facebook users who are friends of existing fans of the charity. Tucker (2012) finds that these socially-targeted ads were more effective than regular display advertising. This holds both for randomly selected users and users who previously expressed their interest in either charity or education. For example, the aggregate click-through rate of the socially-targeted ads was around twice as large as that of non-socially-targeted ads.

5.2 Keyword advertising

A very effective form of targeted advertising is keyword advertising. Keyword targeting refers to any form of advertising that is linked to specific words or phrases and is displayed when a user is looking for information. Therefore, such advertising is not necessarily a nuisance to users and is wasted with a lower probability than, say, TV advertising or banner ads, as the ad is relevant to the user's query and, therefore, valuable. Nowadays, almost all search engines offer keyword advertising. The most well-known form is probably Google AdWords.⁵⁴ It also engages in contextual advertising; that is, Google's system scans the text on the websites that are most relevant to the search query and displays ads based on the keywords found in the respective texts. A main question is whether this form of targeting is welfare-enhancing.

In this subsection, we focus on keyword advertising on search engines. The question has been raised whether search engines have an incentive to present the most relevant ads according to the keyword entered by the user. To address this question, de Cornière (2013) proposes a model with a single search engine that matches potential consumers and producers. Consumers are uniformly distributed along a circle with circumference 1. Each consumer is described by a two-dimensional vector: first by the consumer's location on the circle, which describes her favorite product denoted by $\omega \in [0; 1]$; and second, by her willingness-to-pay, denoted by $\theta \in [0, \bar{\theta}]$. In particular, in each location, there is a continuum of mass 1 of consumers whose reservation price θ is distributed with c.d.f. F . Both variables ω and θ are private information. Each consumer buys, at most, one unit and obtains a utility of

$$v(\theta, d, P) = \theta - c(d) - P$$

where P is the price of the product and $d = d(\chi, \omega)$ measures the distance between the product's location χ and the consumer's location ω . The function $c(d)$, therefore, measures the mismatch costs and is assumed to be increasing and weakly convex.

Products are continuously distributed on the circle. There is a continuum of entrants for each product. Each product can be described by a keyword, which is denoted by $\chi \in [0; 1]$; that is, the keyword is identical to the location of the product. The parameter χ is private information to the producer, implying that consumers know neither the position of a firm on the circle nor the price; hence, they need to use the search engine. When a firm wants to advertise on the search engine, it incurs a fixed cost of C to launch an advertising campaign. (This cost is not a payment to the search engine.)

⁵⁴ The study by Goldfarb and Tucker (2011a) discussed above uses keyword advertising by Google.

The search process works as follows: Firms select a set of keywords that they want to target. The set is assumed to be symmetric around χ – that is, $\Sigma(\chi) = [\chi - D_\chi; \chi + D_\chi]$. Consumers enter the keyword of their preferred product ω . After entry of the keyword ω , the search engine randomly selects a firm χ , such that $\omega \in \Sigma(\chi)$. Once a consumer has decided to use the search engine, she incurs search costs of s and learns the price and location of the firm selected by the search engine. The firm then pays an amount p to the search engine; therefore, p represents the price-per-click. The consumer can then buy the product or not buy it and stop searching, or she can hold the offer and continue searching. That is, recall is costless. For each additional search, the consumer again incurs costs of s .

The timing of the game is as follows: In the first stage, the search engine chooses the per-click price p ,⁵⁵ which is public information to producers and consumers. In the second stage, producers make their decisions. They first decide whether to be active on the search engine; if so, they incur the fixed cost C . Each active firm located at χ then chooses a price P_χ for its product and an advertising strategy D_χ . The mass of active firms is denoted by h . Finally, consumers decide whether to use the search engine. If they do so, they incur search costs s , type in the keyword of the most preferred product ω , and start a sequential search among firms $d(\chi, \omega) \leq D_\chi$. The search engine draws firms in the respective range with equal probability.

What is the perfect Bayesian equilibrium with free entry of firms for this game? Once a consumer has decided to use the search engine and has entered the keyword ω , he obtains a search result showing the link to a firm in the support $[\omega - D^*; \omega + D^*]$. Suppose that all firms charge an equilibrium price of P^* . Then, the expected utility that a consumer gets from this search is

$$\int_{\omega-D^*}^{\omega+D^*} \frac{v(\theta, d(\chi, \omega), P^*)}{2D^*} d\chi = \int_0^{D^*} \frac{v(\theta, x, P^*)}{D^*} dx.$$

The consumer's optimal search behavior is a cutoff rule. That is, the consumer will buy the product of a firm χ if the distance $d(\omega, \chi)$ is lower than or equal to a reservation distance, denoted by R . If a consumer decides not to buy the product, she can improve her utility only by finding a firm that is located closer to her most preferred product because all firms charge the equilibrium price P^* . Therefore, for R^* to be an equilibrium reservation distance, the consumer must be indifferent between buying the product and continue searching; that is,

$$\int_0^{R^*} \frac{v(u, x, P^*) - v(u, R^*, P^*)}{D^*} dx = \int_0^{R^*} \frac{c(R^*) - c(x)}{D^*} dx = s.$$

The left-hand side is the expected gain from continuing to search, and the right-hand side represents the search costs. By totally differentiating this expression, one can show that R^* increases with s and D^* .

To determine the firms' optimal targeting strategy, note that a consumer will never come back to a firm if she does not buy from this firm immediately. This is because the consumer's stopping rule is

⁵⁵ For a model of click fraud, in which publishers affiliated with the search engine's advertising network or competing advertisers artificially drive up clicks (e.g., by impersonating consumers) without increasing sales, see Wilbur and Zhu (2009).

stationary, and she will keep on searching as long as her match is at (weakly) lower distance than her reservation distance. This implies that the conditional purchase probability of a consumer after clicking on a firm's link is either 0 or 1. Because the firm has to pay only for consumers who click on the link, the optimal targeting strategy is simple and equal to $R^* = D^*$. Therefore, in equilibrium, consumers will not search more than once. This allows us to deduce the consumers' participation decisions regarding whether or not to use the search engine. The cutoff reservation value, such that a consumer is indifferent between using the search engine or not, is given by $\theta - P^* - s - E[c(d)|d \leq R^*] = 0$.

We now turn to a firm's optimal pricing decision. If a firm charges price P different from the candidate equilibrium price P^* , it will optimally also change its targeting strategy. In particular, it will target consumers located at a distance smaller than the new reservation distance $R(P, P^*, D^*)$, taking into account that all other firms follow the candidate equilibrium strategy P^* and D^* . Given this new strategy, the firm faces a mass of $2R(P, P^*, D^*)h^*$ competitors in equilibrium. This is because every consumer within distance $R(P, P^*, D^*)$ is targeted by exactly this mass of firms. Since all consumers buy without searching a second time, the firm's profit function per consumer equals

$$\pi(P, P^*, p) = (P - p) \frac{R(P, P^*, D^*)}{R(P^*, P^*, D^*)h^*}.$$

The equilibrium price P^* is given by the first-order condition of this expression with respect to P . In equilibrium, the mass of participating consumers is given by $1 - F(\theta^*)$. Therefore, a firm's expected profit is $(1 - F(\theta^*))\pi(P, P^*, p)$. Since all firms charge the same price in equilibrium, the free-entry condition determines the mass of entering firms. Explicitly accounting for the dependence of P^* , θ^* , and h^* on p , we can write the free-entry condition as

$$(P^*(p) - p) \frac{1 - F(\theta^*(p))}{h^*(p)} = C.$$

Finally, we turn to the profit of the search engine. Since every consumer searches only once, the profit of the search engine is given by $\pi^{SE}(p) = p(1 - F(\theta^*(p)))$. This profit function shows the trade-off faced by the search engine. Everything else equal, the search engine obtains a higher revenue when increasing p . However, such an increase in the per-click fee leads to a higher price on the product market. Since consumers anticipate this, fewer of them will use the search engine. Maximizing $\pi^{SE}(p)$, we obtain that the optimal per-click fee is implicitly given by $p^* = \frac{1 - F(\theta^*(p^*))}{f(\theta^*(p^*))}$. It is evident that the search engine sets a positive p^* . Instead, the socially optimal fee equals zero. Hence, the equilibrium implies a distortion, as consumer participation is too low and product prices are too high.

Within de Cornière's (2013) model, we can now analyze the effects of targeting. Suppose that targeting is not possible. In the model, this corresponds to the case in which $D = 1/2$ for all products. The optimal reservation distance for consumers, R^* , is then implicitly given by $\int_0^{R^*} \frac{c(R^*) - c(x)}{1/2} dx = s$. Therefore, consumers may search more than once, and targeting reduces the expected number of clicks. Since the reservation distance is increasing in D , the expected mismatch costs are also lower with targeting. As a consequence, targeting reduces the search frictions.

Targeting has more-subtle effects on the price of the final good. First, targeting changes the pool of firms from which consumers sample. In particular, with targeting, this pool is composed of firms that are expected to be a better match for consumers. This implies that the continuation value of searching is higher for consumers, inducing a downward pressure on prices. By contrast, without targeting, a firm cannot adjust its advertising strategy D along with its equilibrium price P . The per-click price is, therefore, considered a fixed cost and is not passed through into the final good price.⁵⁶ As a consequence, the overall effect is ambiguous, and de Cornière (2013) shows that targeting can lead to a welfare loss.

Another important question regards the incentive of the search engine to choose the most relevant ads after a consumer has entered a keyword. To this end, suppose that the search engine can choose the value of accuracy of its search results by choosing D itself. So, the search engine has two choices, D and p . De Cornière (2013) shows that in this case, the search engine can extract the whole profit from firms with the per-click fee. Consequently, the search engine chooses D in order to maximize firms' profits. The optimal matching accuracy for the search engine is, then, $D^{SE} \geq D^*$. The intuition behind the result is as follows. If the search engine sets $D < D^*$, the distance is strictly smaller than the consumers' reservation distance. Then, a price increase in the final good market does not lead to reduced demand. This implies that firms have an incentive to increase the price up to the reservation price of the marginal consumer, leading to a negative utility of the marginal consumer (remember that the search costs are sunk). As a consequence, consumers will not participate. Therefore, a high level of matching accuracy lowers product market competition, which dissuades consumers from using the search engine. Instead, when $D > D^*$, some consumers search more than once. This leads to lower firm prices and ensures higher consumer participation. Therefore, the search engine may find it optimal to set $D^{SE} \geq D^*$ but never $D^{SE} < D^*$. To sum up, the search engine does not have an incentive to choose a more accurate matching than advertising firms themselves.

In the empirical literature, keyword advertising and targeting technologies on the Internet have also attracted considerable attention. For example, to determine the effectiveness of keyword advertising, Ghose and Yang (2009) use a panel data set of several hundred keywords from a nationwide retailer that advertises on Google. They find that click-through and conversion rates fall in the keyword rank. However, this is not necessarily true for profitability. In particular, keywords in middle positions are often more profitable than those at the very top of a search engine's results page. Interestingly, Ghose and Yang (2009) find that the effect of retailer-specific information in a keyword is very different from brand-specific information. Whereas retailer-specific information leads to an increase in conversion rates of up to 50.6%, brand-specific information leads to a decrease of 44.2%. Similar patterns are observed for click-through rates. It could be of interest to extend the model outlined above (or related frameworks) to allow for many search results with different ranks to explain this empirical evidence.

Other studies also analyze the influence of the position (or rank) on click-through and conversion rates and provide an analysis distinguishing between different advertising effectiveness measures. For

⁵⁶ Dellarocas (2012) provides an in-depth analysis of the implications of pay-per-click pricing on final consumer prices. He shows that performance-based advertising, such as pay-per-click pricing, leads to double marginalization. As a result, consumer prices are higher than with per-impression pricing.

example, Agarwal, Hosonagar and Smith (2011) also find a positive effect on the click-through rate but demonstrate that the conversion rate is often higher for middle-ranked positions. Rutz and Trusov (2011) provide

a model and an empirical analysis of the relation between click-through-rates and conversion rates, whereas Rutz and Bucklin (2011) demonstrate positive spillovers between generic and sponsored search. That is, a generic search often generates a subsequent sponsored search.

In a recent study, Blake, Nosko and Tadelis (2014) also measure the effectiveness of keyword advertising, explicitly distinguishing between brand and non-brand keywords. When a user types in a brand such as “Macys” or “eBay” as a query in a search engine, it is very likely that the user is already familiar with the brand. In response to the user’s query, the search engines displays paid ads at the top of the search results, and the brand pays the per-click price given that the user clicks on this query. However, the user would have found the brand’s site almost surely through organic search. Blake, Nosko and Tadelis (2014) test this with data from eBay. They halted advertising for eBay-brand-related queries on the search engines Yahoo! and MSN for some time and found that 99.5 percent of traffic from the paid link was immediately captured by traffic from the organic link. Hence, substitution between paid and unpaid traffic for eBay is almost complete.⁵⁷ Is this also true for non-brand keywords? To answer this question, Blake, Nosko and Tadelis (2014) conducted another natural experiment by stopping eBay advertising via paid links in designated areas for 60 days. In addition, they segmented consumer groups into those who are frequent eBay visitors and those who are not. They found that paid links did not have a statistically significant effect on the first group since users of this group are already familiar with eBay. However, there was a significant increase in newly registered users and purchases in the second group due to exposure of paid links. This supports the informative view of advertising.

Several empirical studies focus on how targeted advertising interacts with other forms of advertising. Goldfarb and Tucker (2011b) conduct a large-scale field experiment exposing individuals to two different forms of online advertisements. The first is a contextually targeted ad, such as a banner ad for a new computer displayed on a site devoted to computing and technology. The other is a highly-visible ad, which users might consider obtrusive. For example, an ad is considered obtrusive if the ad is part of an in-stream audio or video, if it is a pop-up window, or if it automatically (non-user initiated) plays audio or video, among other characteristics. Goldfarb and Tucker (2011b) find that the effect of targeted ads *alone* (without obtrusiveness) and the effect of obtrusiveness *alone* (without targeting) have a positive influence on the effectiveness of advertising. However, both strategies *in combination* nullify this effect and are ineffective. An explanation for this can be that consumers perceive themselves to be manipulated, which reduces their purchase intentions. In particular, when exposed to targeted ads that are obtrusive, consumers may have privacy concerns. Goldfarb and Tucker (2011b) find evidence that supports this view.

More generally, advertising across different types of media and consumer online behavior are connected. Joo, Wilbur, Cowgill, and Zhu (2013) study how advertising in an offline medium (television)

⁵⁷ A natural explanation is that eBay is a well-known brand and is, therefore, highly listed in the organic search results.

affects consumers' online searches and, thus, search engine advertising. They consider TV advertisements of financial services and analyze how these commercials affect consumer search behavior. They find a significantly positive effect. For example, a few hours after being exposed to a TV ad for a particular brand, searchers have a stronger tendency to enter branded keywords instead of generic keywords.⁵⁸

Overall, this suggests that, while online and offline advertising are substitutes, offline advertising stimulates online product search. Further research, both theoretical and empirical, could be fruitful to establish a solid pattern that links the influences of advertising in one medium to consumer behavior in another.

6. Media platforms matching advertising to users

Internet media facilitate the targeting of ads to specific consumers. Traditional media provide tailored offers such that consumers self-select into particular programming and content. Advertisers then benefit from the correlation of consumer tastes with media content and with advertised products. Clearly, such tailoring strategies are also available on the Internet and were analyzed in the previous section. A novel feature of advertising on the Internet is the wealth of personal data available to data providers, which allows the matching of advertising to consumer tastes on media platforms irrespective of the media content that is consumed.⁵⁹ While this wealth of data raises serious privacy and data protection issues (not analyzed in this paper), it also affects the way media platforms operate. In addition, since consumers mostly visit multiple sites, excessive advertising beyond what consumers can digest also arises on the Internet.

6.1 Tracking and personalized offers

The Internet has opened new ways to track consumers by placing cookies. Cookies are small pieces of data sent from a website, which track the user's activities. To the extent that previous user behaviour allows inferences on users' current tastes, it becomes possible to, at least partly, avoid wasteful impressions. Google explicitly writes in its information to users: "We use cookies to make advertising more engaging to users and more valuable to publishers and advertisers." Google then provides a more detailed explanation on the use of cookies: "Some common applications of cookies are to select advertising based on what's relevant to a user; to improve reporting of campaign performance; and to avoid showing ads the user has already seen." While perfect public tracking would, in particular, allow websites to best match advertising to users, in many markets, media platforms may not share tracking information. Thus, tracking is often imperfect.

Tracking may allow for the segmentation of consumers according to some broad categories without fully personalizing the targeting of ads. This segmenting of the consumer pool may be based on past

⁵⁸ Rutz and Bucklin (2012) find a similar result for online advertising and online search.

⁵⁹ We are not claiming that the tailoring of advertising is a completely new phenomenon. For instance, advertisers may use personal information when sending out coupons by mail.

purchases (Malthouse and Elsner, 2006).⁶⁰ This information helps to increase the likelihood that the advertiser's product and the consumer's taste match.

Targeting can be based on personal characteristics. Some of the theoretical models presented below are based on this idea. Advertising has been shown to be more effective when it is targeted to particular consumers using consumer browsing behaviour (Chen, Pavlov, and Canny, 2009) or using inferred or observed demographics as consumer characteristics (Joshi, Bagherjeiran, and Ratnaparkhi, 2011).⁶¹ Thus, the empirical literature indicates that tracking can increase ad effectiveness. A more challenging question is to uncover the impact of tracking on industry outcomes.

Beyond the collection of information from cookies, the matching of advertising to users may rely on information provided by database marketing companies. Marwick (2014) provides some information on the second-largest company in the industry, Acxiom. According to Marwick (2014), Acxiom has 23,000 computer servers and processes more than 50 trillion data transactions per year, keeping records on hundreds of million U.S. residents. Data include 200 million mobile profiles, information gathered from publicly available records (such as home valuations and vehicle ownership), information about online behavior (1.1 billion browser cookies, information on browser advertising, and other information), as well as data from customer surveys and offline buying behavior. On average, for each U.S. resident, Acxiom keeps about 1500 pieces of data. Thus, Acxiom has a wealth of information that it can sell to interested parties, in particular with the aim to better match advertising or services to user tastes.

While on traditional ad-financed media, the user pays with her attention, on Internet media, the user pays not only with her attention, but also with her personal data. Thus, websites including Internet media may make revenues even if they neither charge users nor carry any advertising. They can accomplish this by opening a third source of revenues – selling user information.

A number of theoretical efforts help in understanding the forces at play when media platforms track users or rely on third-party information in their effort to best match advertising to users. The model presented at the end of this subsection explicitly includes the sale of user data for the purpose of targeting.

A media platform may provide tracking information about consumers to advertisers. Doing so, allows advertisers to bid for ads conditional on the information they receive. When advertising space is scarce, advertisers operating in such an environment internalize that in case of tracking their bids will only be successful if they provide better matches to consumers than other advertisers. As a consequence, advertisers set higher prices with tracking information than without. While tracking improves average match quality, leading to higher prices and, thus, larger industry profits, it also reduces the share of industry profits that can be extracted by the platform.

⁶⁰ Other empirical work segments consumers according to their cognitive style (Hauser, Urban, Liberali, and Braun, 2009).

⁶¹This empirical work combines tailoring and tracking, as it combines consumer characteristics with content matching.

De Corniere and De Nijs (2014) formalize this trade-off and investigate the platform's incentives to install a tracking technology. Here, through a second-price auction, a monopoly media platform sells a single advertising slot to n advertisers. This slot gives exclusive access to the consumer. Thus, sellers act as monopolists in the product market. The timing of the model is as follows: First, the platform decides whether to install a tracking technology. Second, advertisers simultaneously set the product price p_i , $i = 1, \dots, n$. Third, the consumer type is revealed to advertisers under tracking; it remains unknown otherwise. Fourth, advertisers simultaneously place bids for the advertising slot conditional on the information they received. The consumer is matched to the winning advertiser.

A consumer is of type $(\theta_1, \dots, \theta_n)$, where θ_i is i.i.d. across products and distributed according to F with density function f on $[\bar{\theta}, \bar{\theta}]$. Type θ_i for product i gives rise to a demand function $D(p_i; \theta_i)$. A higher type is assumed to be associated with larger demand for the respective product (e.g., a larger probability to buy the product); i.e., $D(p_i; \theta_i) > D(p_i; \theta'_i)$ if and only if $\theta_i > \theta'_i$. The profit of an advertiser selling to a consumer is $\pi_i(p_i, \theta_i) = (p_i - c)D(p_i; \theta_i)$. Absent tracking, if an advertiser's bid is successful, its expected profit gross of the advertising cost is $\int_{\bar{\theta}}^{\bar{\theta}} \pi_i(p_i, \theta_i) f(\theta_i) d\theta_i$. The profit-maximizing product price p^{NT} solves

$$\int_{\bar{\theta}}^{\bar{\theta}} \frac{\partial \pi_i(p_i, \theta_i)}{\partial p_i} f(\theta_i) d\theta_i = 0$$

in p_i . Since advertisers are homogeneous at the bidding stage, the media platform can extract the full expected industry profit $\int_{\bar{\theta}}^{\bar{\theta}} \pi_i(p^{NT}, \theta_i) f(\theta_i) d\theta_i$.

With tracking, it can be shown that the product price p^T solves

$$\int_{\bar{\theta}}^{\bar{\theta}} \frac{\partial \pi_i(p_i, \theta_i)}{\partial p_i} F^{n-1}(\theta_i) f(\theta_i) d\theta_i = 0$$

in p_i , since advertiser i wins the auction if and only if θ_i is larger than θ_j , $j \neq i$. As the firm is unlikely to win the auction when θ_i is small, it will set a higher price at the pricing stage than without tracking, $p^T > p^{NT}$. The price under tracking is increasing in the number of advertisers. Consequently, tracking results in a better match between advertiser and consumer and increases industry profits. With tracking, advertisers obtain a positive information rent and, thus, a strictly positive share of industry profits.

When deciding whether to install the tracking technology, the media platform faces the trade-off between increasing efficiency (and industry profits) and rent-extraction; such a trade-off also obtains in Ganuza (2004). As the number of advertisers turns to infinity, the product price p^T turns to the monopoly price of a firm facing a consumer with type $\bar{\theta}$ and, thus, the information rent of advertisers disappears. Hence, for the number of advertisers sufficiently large, the media platform installs the tracking technology and shares the consumer information with advertisers.

If the platform sells multiple advertising slots through a uniform price auction and advertisers sell independent products, the equilibrium product price is shown to be decreasing in the number of

advertising slots. In addition to the standard price-quantity trade-off, an increase in the number of advertising slots renders winning a slot in the auction less informative about the expected elasticity of demand. When the number of slots is sufficiently large, de Corniere and De Nijs (2014) show that the platform chooses not to install the tracking technology since the losing bidder who determines the ad price in the auction, tends to receive a rather bad signal with tracking.

Johnson (2013) also explores the effects of the tracking technology on advertiser profits and consumer surplus. While he does not include media platforms in his model, his analysis is useful in obtaining insights about the role of tracking when consumers can block advertising.

In his model, advertising creates an opportunity for advertisers and consumers to form a match. There is a mass 1 of advertisers and a mass 1 of consumers. For each advertiser-consumer pair, the probability of such a match is ϕ , which is distributed i.i.d. across all pairs according to some distribution function F with positive density of f on $[0,1]$. A match generates a surplus Λ for the advertiser and $1 - \Lambda$ for the consumer. Advertisers offer totally differentiated products.

The advertiser learns about the match probability with probability ψ and does not learn otherwise. In this model, improved tracking corresponds to a larger value of ψ . Thus, the probability of a match is $\psi\phi + (1 - \psi)E\phi$ if the consumer sees the ad.

Advertisers incur a cost of $\kappa > 0$ for sending an ad. Consumers have the possibility of blocking an ad with probability ζ . Hence, an advertiser decides to advertise to a consumer with signal ϕ if

$$\kappa \leq (1 - \zeta)\Lambda[\psi\phi + (1 - \psi)E[\phi]].$$

A firm that sends an ad to consumers with signal ϕ will send an ad to all consumers with larger signals. If h is the mass of ads sent by an advertiser, we have $\phi(h) = F^{-1}(1 - h)$ as the signal of the marginal consumer. An advertiser's profit is, then,

$$-\kappa h + (1 - \zeta)\Lambda \left[\psi \int_{\phi(h)}^1 x dF(x) + (1 - \psi)hE[\phi] \right].$$

All consumers are exposed to the same number of ads if they decide not to block them. Each ad they receive generates a nuisance γ from being exposed to it. However, advertising leads to consumption, which enters the consumer's utility, as well. Thus, the expected utility from receiving h ads is

$$v = -\gamma h + (1 - \Lambda) \left[\psi \int_{\phi(h)}^1 x dF(x) + (1 - \psi)hE[\phi] \right].$$

If consumers block ads, they receive an outside utility u_0 , which is distributed according to G , a continuously differentiable distribution function with support $(-\infty, 0]$. Taking h as given, each consumer compares her expected utility from receiving ads to the outside utility of blocking. Hence, the fraction of blocking consumers is $\zeta = 1 - G(v)$. Consider the game in which, simultaneously, consumers decide whether to block and advertisers decide how many ads to send. A pair (ζ^*, h^*) constitutes an equilibrium of this game. In this model, the cost-benefit ratios for advertisers and consumers, κ/Λ and $\phi/(1 - \Lambda)$, play a decisive role. They reflect the cost of an ad relative to the benefit of a successful match.

Johnson (2013) shows that the second-best advertising level (when consumers are free to block advertising) is smaller than the equilibrium level h^* if and only if the cost-benefit ratio on the consumer side is larger than the cost-benefit ratio on the advertiser side. In equilibrium, advertisers are indifferent about whether to place the marginal ad. Since the marginal probability of trade is the same for advertisers and consumers, intuitively, there is socially insufficient advertising if consumers place a higher value on ads than advertisers. If the reverse holds, advertisers post too many ads.⁶²

Improved tracking – i.e., an increase in ψ – keeping blocking decisions unchanged, leads to more advertising if the surplus derived from the marginal ad exceeds the surplus derived from the unconditional average ad. If advertisers increase advertising due to improved tracking, this must mean that they gain more from the marginal ad than from the unconditional average ad. This must continue to hold if consumers adjust their blocking decision. Johnson (2013) shows that improved tracking increases advertiser profits in equilibrium, even though this may imply more blocking. Whether consumers gain or lose from improved tracking is ambiguous.⁶³

Depending on the information available to advertisers, when there are multiple platforms, advertisers may waste impressions by hitting the same consumer more than once even if tracking is perfect on each platform. If however platforms share cookie information – a practice called cookie matching – advertisers can avoid such multiple exposures.

Athey, Calvano and Gans (2014) consider a market with two platforms that perfectly track consumers on their own platform, but may not observe the exposure of consumers to advertising on the other platform.⁶⁴ They investigate the impact of this lack of cookie matching on market outcomes and media platforms' profits. In their model, as, e.g., in Ambrus, Calvano, and Reisinger (2014), some consumers exclusively consume media content of one platform, while the others consume both (see Section 4.1). Thus, there are exclusive consumers and overlapping consumers.

Similar to previous models, media platforms provide access to consumers. To focus on advertisers' behavior, suppose first that media platforms do not make any decision and that ad prices clear the market for ads, for given ad levels for each platform. Consider a continuum of advertisers who are heterogeneous with respect to the profit per consumer they derive when successfully contacting a consumer. Their behavior determines the demand for advertising. The advertiser value per consumer is distributed between 0 and some upper bound. Advertisers with a high value per consumer have a

⁶² The results by van Zandt (2004) and Anderson and de Palma (2009) can also be interpreted as showing the possibility of socially excessive advertising. In their models, socially excessive advertising may arise because lower-value ads crowd out higher-value ads, an issue we return to in Section 6.2. In contrast, in Johnson (2013), a larger ad level encourages consumers to block ads.

⁶³ If advertisers have to pay a media website to place ads (an issue not considered by Johnson, 2013), the website can manage the ad level of advertisers. Suppose that the media website charges a price per ad. This price enters the advertiser's profit function as part of its cost κ . A welfare-maximizing single media website (which cannot directly control blocking) will then implement the second-best optimal advertising level. A profit-maximizing, monopolistic ad-financed media website will also internalize the effect of ad levels on blocking. However, it will typically not implement the second best. The effect of the tracking technology on total and consumer surplus in such a media market has yet to be explored.

⁶⁴ Ghosh, Mahdian, McAfee, and Vassilvitskii (2012) provide a related analysis, which also shows that in some cases platforms prefer to share cookie information, whereas in others they do not.

stronger incentive to contact consumers than advertisers with lower value. If there were no overlapping consumers (and perfect tracking on each platform), there would exist a marginal advertiser such that every advertiser with a lower value per consumer would not advertise, whereas all advertisers above this threshold would deliver each impression to a distinct consumer. Thus, no impression would be wasted and advertising would be delivered efficiently.

Demand for advertising depends on the tracking technology. If tracking is perfect on each platform but there is no cookie matching, multi-homing advertisers waste some impressions as they sometimes show the same ad to switchers twice. This waste, together with advertiser heterogeneity, implies sorting of advertisers: low-type advertisers single-home (and miss some consumers), while high-type advertisers (who have a higher opportunity cost of not informing consumers) multi-home.

This establishes the main insight. With perfect tracking across platforms, the number of impressions would map one-to-one into the number of consumers reached by an advertiser. On the contrary, with the above imperfect tracking, some overlapping consumers will see the same ad twice and so their attention is wasted. By increasing the number of overlapping consumers waste becomes more prevalent under imperfect tracking and the value of the advertising inventory is further degraded.

The analysis can be extended to allow for the platforms simultaneous choosing advertising inventory. In the presence of overlapping consumers, platforms become essentially Cournot competitors. Athey, Calvano, and Gans (2014) show that, in equilibrium, an increase in the fraction of overlapping consumers leads to higher advertising inventories and, in turn, lower equilibrium advertising prices.⁶⁵

Taylor (2013b) provides a different perspective on the role of tracking. In his model, media platforms choose content quality, taking into account that it increases the likelihood that a consumer does not switch to another platform. The incentive to invest in content quality is affected by the tracking technology available to the platforms. Taylor's (2013b) model has a number of features different from those of the three previous models. First, a key ingredient is product market competition among advertisers;⁶⁶ second, consumers are uninformed about content quality before visiting a website and, therefore, access media platforms at random. Thus, Taylor (2013b) focuses on how content quality affects consumer behavior *after* consumers have clicked on the media website. Key to his results is the interaction between product market competition among advertisers and the consumer decision of whether to switch to another media platform (endogenous multi-homing).

In Taylor's (2013b) model, two ad-financed media websites provide content to a large number of consumers of measure D . Consumers enjoy media consumption but incur a cost c for visiting a website.

⁶⁵ Due to the heterogeneity of advertisers, the analysis is rather intricate and best-response functions are non-monotone.

⁶⁶ Most theoretical papers on media economics postulate that advertisers have monopoly power over consumers in the product market. As discussed in Section 4.1, an exception is Gal-Or and Dukes (2003). In their model, as well as in Taylor's (2013b) model, intense product market competition creates incentives for media platforms and advertisers to sign exclusivity contracts of the form that the media platform agrees not to carry ads from competing advertisers.

The two media platforms choose content quality $q_i \in [0,1]$ and incur cost $k(q_i)$.⁶⁷ Here, content quality measures the probability that a consumer is satisfied with the content and does not move to the other website for this topic. Content quality is a search good; thus, consumers cannot assess content quality before actually visiting the website. Consumers visit websites sequentially, recall ads and make consumption choices after their media consumption.

Websites will, nevertheless, choose positive quality so as to retain the consumer's attention. Why? Websites bundle advertising to content and sell ads to advertisers. More specifically, the website places an ad together with the content offer. Each consumer has a particular consumption interest (e.g., interest in a particular product category), and an advertiser matches this interest if it offers a product in this category. The surplus generated by a successful match between advertiser and consumer is normalized to 1. Advertisers are assumed to offer non-differentiated products within a category. Therefore, if a consumer happens to see an ad from another advertiser within the same category, firms are Bertrand competitors; if this consumer were identified, and advertisers made personalized offers, the consumer would obtain the total surplus of the match, and the advertiser would make zero surplus in this interaction. Therefore, a website will carry only one ad of a given product category since it wants to raise revenues from selling ads; each website extracts the full advertiser profit.

The tracking technology determines the probability of a successful match; it allows the website to identify the product category of interest with probability ϕ and delivers any of the other categories with the remaining probability. If a consumer visits both websites, she is exposed to the advertising on the other website. For simplicity, suppose that this other websites draws on a different set of advertisers and, therefore, that a consumer will never see the same ad twice. If the consumer sees another ad in the product category of interest, she can choose between the two offers and will choose the offer at the lower price.

As long as some, but not all, consumers visit both websites (i.e., in any interior solution), there are consumers who observe one successful match and others who observe two. Thus, advertisers randomize over price. Each advertiser's equilibrium profit is equal to its monopoly price 1 times the likelihood that it provides the only match. Here, content quality helps, as it increases the probability that consumers do not visit multiple websites and, thus, prevents multiple exposures of advertising.

As Taylor (2013b) shows, equilibrium content quality can be calculated as

$$q^* = (k')^{-1}\left(\frac{\phi D}{2}\right).$$

Better tracking is valuable even in a monopoly context, as it increases the probability of a match between advertiser and consumer. Everything else equal, this increases the incentive to invest in content quality. The tracking technology matters also for product market competition since better tracking makes it more likely that a consumer who visits both websites encounters two matches. Hence, better tracking makes product market competition more intense. In this case, the website would extract lower

⁶⁷ The cost function satisfies the appropriate properties so as to ensure an interior solution characterized by a first-order condition.

rents from advertisers. To reduce the probability of multiple exposures to ads within the same product category, a website has to increase its content quality. Due to this competition effect, websites may actually overinvest in content quality compared to the welfare-maximizing solution. As Taylor (2013) illustrates, websites may actually suffer from improved tracking, as they are compelled to invest (excessively) in content quality.

Tracking is made feasible by data providers that handle large amounts of data that they obtain from placing cookies. Bergemann and Bonatti (2014) explore the interaction between data providers and advertisers. Advertisers face consumers with heterogeneous match value $v \in V$. Through cookies, they obtain precise information on consumers' value. This allows advertisers to segment the consumer side into two groups: a target group about which it collects detailed information and to which it makes personalized advertising offers; and an anonymous outsider group in which all consumers receive the same level of advertising as everybody else in this group.

If an advertiser were fully informed, he would reach a consumer with probability ϕ and then extract the full surplus v of the match between his product and the consumer. Thus, he would make revenue $v\phi$. To reach consumers with probability ϕ , he has to place $a(\phi)$ ads on the publisher's platform, which is assumed to charge p per ad. Hence, a fully informed advertiser makes profit $v\phi - pa(\phi)$. The advertiser's problem is that he knows only about some prior distribution of v and, thus, may not be able to appropriate the full surplus.

To better extract surplus, the advertiser may acquire data about the individual consumer. Such data provide a signal about the consumer's match value. This is where cookies come into play. A cookie, bought at a price P_c for a subset of types $W \subset V$, allows the advertiser to identify a consumer v . Hence, for all $v \in W$, the demand for advertising space is $v = pa'(\phi^*(v))$, where $\phi^*(v)$ is the full information demand for advertising space. For all $v \notin W$, the advertiser updates beliefs and acquires ad space according to $E[v|v \notin W] = pa'(\bar{\phi})$. Clearly, $\bar{\phi} = \phi^*(E[v|v \notin W])$. Figure 9 illustrates the market environment.

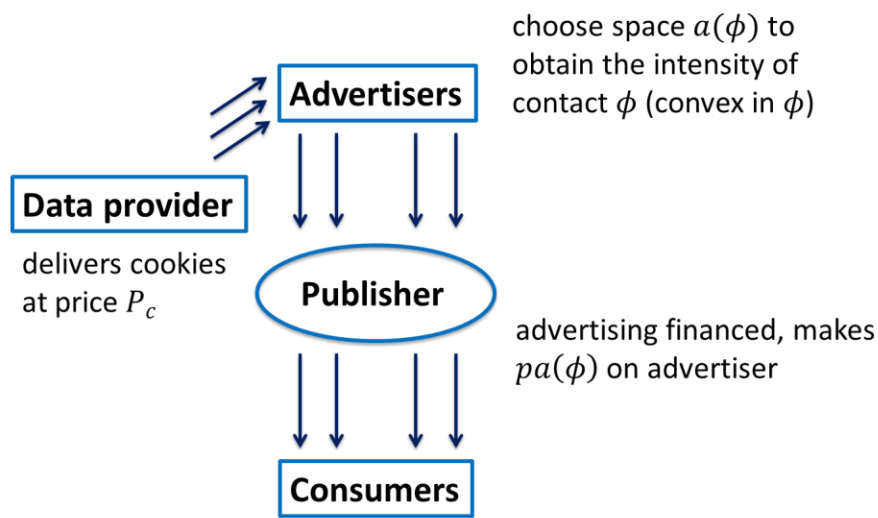


Figure 9: Internet advertising based on cookies in Bergemann and Bonatti (2014)

A fully informed advertiser makes profit $v\phi^*(v) - pa(\phi^*(v))$, which is convex in v . By contrast, absent any information, the advertiser's profit would be $v\phi^*(E[v]) - pa(\phi^*(E[v]))$, which is linear in v . An uninformed advertiser advertises too much to low-value consumers and too little to high-value consumers. Thus, the advertiser has an incentive to buy cookies.

Bergemann and Bonatti (2014) characterize the demand for cookies for a given price P_c . An advertiser may buy from a single interval that includes either the lowest-value or the highest-value consumer. Alternatively, he may want to buy from two intervals, one of which includes the lowest-value consumer and the other the highest-value consumer. Bergemann and Bonatti (2014) then determine the optimal pricing by the data provider. The data provider may limit the amount of data being bought in equilibrium, as the data provider maximizes its profits. Hence, consumers benefit if the data provider enjoys some market power because, absent market power of the data provider, advertisers would become fully informed about the match value.

To put Bergemann and Bonatti (2014) into broader perspective, we note that the publisher may be vertically integrated with the data provider. In this case, the publisher has three potential sources of revenues: First, the publisher can charge consumers directly for its content services. Second, the publisher can charge advertisers for offering advertising space (to the extent that consumers dislike advertising, this constitutes an indirect charge to consumers). These two revenue sources are well established in media economics. Third, Bergemann and Bonatti (2014) formalize that the publisher may also offer advertisers data services for which it can charge a fee. As this allows advertisers to better extract surplus from consumers, providing this information constitutes another indirect charge to consumers.

Tracking allows advertisers to use a sophisticated targeting strategy. While the above models provide relevant insights into the role of the tracking technology, they are not embedded in a dynamic environment. In particular, consumers may search within a product category on a particular website but do not close their session with a purchase. Possible reasons are that they are still in doubt or that they decide that the offer is dominated by the outside option, which may include the possibility of searching again in the future. Due to tracking, advertisers become aware of the identity of such consumers. They may, therefore, "redouble" their efforts and retarget these consumers via advertising on this website or specific offerings on another website when consumers are visiting a website that is affiliated with the ad network selling ads to advertisers. For instance, Facebook has introduced retargeted ads in its users' newsfeed. With dynamic retargeting, the ad network identifies people with the help of the individual cookie profile and recalls the exact product a consumer has looked at before. With this information, it targets the consumer by displaying the same or a related product offered by the firm visited before. This is in contrast to generic retargeting, which uses cookie information only to select the firm of which an ad is shown (see Lambrecht and Tucker, 2013). It has been claimed that dynamic retargeting is four times more effective than generic retargeting and six times more effective than generic advertising using banner ads (Criteo, 2010).

Lambrecht and Tucker (2013) provide a detailed assessment of the effectiveness of dynamic retargeting, based on a field experiment with data from a travel website that sells hotel accommodations and vacation packages. In the field experiment, consumers were randomly assigned generic and dynamic

retargeted ads when they visited an external website affiliated with the ad network. The generic ad showed a generic brand ad for the travel website, while the dynamic retargeted ad showed the hotel the consumer had looked at before, plus a few similar offerings. Perhaps surprisingly, the authors find that, on average, dynamic retargeting is not more effective, where effectiveness is measured by the probability that the consumer makes a purchase on the travel website within a specified time interval.

A possible explanation for the failure of dynamic retargeting vis-à-vis generic retargeting is that consumers may have had only a loose idea of what they were looking for (as argued by Lambrecht and Tucker, 2013). Thus, a generic ad may be more effective since consumers may have figured out that the specific offer they previously looked at was actually not what they wanted. However, such behavior seems less likely if consumers dedicate quite some effort on their search and make further inquiries. In this case, it appears likely that consumers have well-specified ideas of what they liked and may be inclined to revisit their previous searches. Lambrecht and Tucker (2013) try to proxy this by identifying consumers who used a review site. The empirical result supports the view that dynamic retargeting is more effective for those consumers. These findings suggest a rather complex effect of targeting on purchase intent. It suggests that simple and universal messages of how (re)targeting will affect market structure may be too much to expect.

6.2 Advertising congestion and limited attention

Advertisers post ads to leave an impression on consumers and to make a profit by selling a product or service or by being able to sell at a higher price. To achieve this, advertisers have to overcome several hurdles. First, consumers must not block advertising; second, consumers have to remember the ad (remember the advertised product, product features, or the ad experience as a complement) when making a purchase decision; and third, advertisers must be able to make money from such a consumer. We touched upon the first hurdle in Section 6.1 in the context of targeting (Johnson, 2013). The third hurdle arises with competing advertisers in a Bertrand world, where an ad received by a consumer who also received a competing ad is essentially a wasted impression (Taylor, 2013b). The second issue is one of limited attention, which we did not elaborate on in the context of advertising. Advertising congestion can be seen as a mismatch between advertisers and consumers since high-value advertisers may not capture consumers' attention and, thus, are not matched, whereas some lower-value advertisers do manage to gain the consumer's attention.

Anderson and de Palma (2009) formalize information congestion, postulating that consumer attention spans are limited (for a formal presentation of the model in a different context, see Section 3.2). Under "open access" to attention (for example, through billboards or bulk mail), attention is a common property resource to which all consumers have access, which, if excessively exploited, results in congestion in equilibrium. By contrast, a monopoly gatekeeper prices out congestion.

Anderson and Peitz (2014a) show how this approach can be integrated into a model of competing media platforms. Here, we illustrate their setting, following Anderson, Foros, Kind, and Peitz (2012). To see how advertising congestion changes the nature of media competition, consider, first, the situation with an invariant amount of time spent by a representative user on each of the n media platforms. The idea here is that consumers surf the web, spend more time on higher-quality websites and less on lower-quality ones. Content quality then maps into usage time. Suppose that advertisers are totally differentiated.

Absent product market competition, advertisers are monopolists vis-à-vis consumers. Suppose, furthermore, that advertisers extract the full expected surplus from consumers.

Consumers access pages of the different websites in random order. If they visit more pages of a website, they will be exposed to more ads from that website. Websites are assumed to benefit from industry-wide perfect tracking. If the value of the marginal advertiser is larger than the expected value of a repeated impression of a given ad (which holds if congestion is not too severe), all consumers will see an ad only once. Ads will be placed randomly. Even though advertising does not affect media consumption, the website will not post an unlimited number of ads. This is so because it cannot discriminate between different types of advertisers and, thus, has to lower the ad price as it takes in more advertisers.

Each media website decides how many ads to place at an initial stage since this is mostly a question of how to design the website or how to structure the bundle of content and advertising. Let σ_i be the amount of time spent on website $i = 1, \dots, n$, and let σ_0 be the time spent not watching (the outside option), and normalize the total time available to 1. If website i shows a_i ads and a consumer spends all her time on this site, the consumer's total exposure is a_i . To make a match between an advertiser and any given consumer, this consumer must be exposed to the corresponding ad and she must recall the ad. Each consumer will see $\Gamma = \sum_{i=1}^n a_i \sigma_i$ ads in total. However, if the (fixed) attention span φ of a consumer is less than this number – i.e., $\varphi < \Gamma$ – some ads will not be matched. We rank advertisers in decreasing order of willingness-to-pay to contact prospective customers. Hence, the a -th advertiser is willing to pay $p(a)$ to expose the consumer to the ad and break into her attention span. With congestion, the willingness-to-pay reduces to $p(a)\varphi/\Gamma$. With a_i ads on platform i , the ad price is the willingness-to-pay of the marginal advertiser – i.e., $p(a_i)\varphi/\Gamma$. Thus, website i maximizes its profit

$$\sigma_i a_i \frac{p(a_i)\varphi}{\Gamma} = A(a_i) \frac{\sigma_i \varphi}{\Gamma}$$

with respect to the ad space a_i , where $R(a_i)$ is the revenue per ad per viewer. Strategic interaction between websites arises because of advertising congestion and the fact that consumers have a limited attention span for all ads combined across different websites. Thus, each website has access to the common property resource, which is the attention span of each consumer. Because of the free-rider problem, advertising in a market with strictly more than one website may result in congestion. In the monopoly case, the result by Anderson and de Palma (2009) applies, and the monopoly website always sets its ad level such that congestion does not arise. If the attention span is sufficiently large, each website i chooses a_i so as to maximize revenue per ad per viewer – i.e., a_i^* solves $A'(a_i) = 0$ and all websites choose the same ad level. With congestion, when all websites offer the same content quality, the first-order condition of profit maximization becomes

$$A'(a) \frac{\varphi}{\Gamma} - A(a) \frac{\varphi}{\Gamma^2} \sigma = 0$$

after using symmetry – i.e., $\sigma = \sigma_i$ and $a = a_i$. This implicitly defines the equilibrium advertising level a^* . Since total advertising is $\Gamma = n\sigma a$, the equilibrium ad level is given by the solution to

$$\frac{aA'(a)}{A(a)} = \frac{1}{n}.$$

If $p(a)$ is logconcave, the left-hand side is decreasing in a . This implies that a larger number of websites leads to an increase in the advertising level of each website and, thus, to more-severe congestion. The intuition is that with more alternative websites around, each website internalizes to a smaller extent the negative effect that an increase in its ad level has on the resulting advertising price. In an asymmetric market, Anderson and Peitz (2014a) show that the larger website chooses a lower ad level a_i . This is intuitive, as a larger platform internalizes the congestion externality to a larger extent since more inframarginal ad slots are affected by the price decrease resulting from a larger ad volume.⁶⁸ The general lesson that emerges from the ad congestion model is that congestion generates interdependence among websites due to price effects on the advertiser side. To the extent that consumers consult a large number of ad-financed websites, the analysis suggests that the congestion problem is likely to be severe.

7. Conclusion

In this paper, we surveyed the literature on media economics as it applies to the Internet. A lot of progress has been made in understanding how media companies react to the challenges brought about by the Internet and how they adjust their business models or invent new ones. The theoretical and empirical studies show the potential for market failures on the Internet (e.g., on search engines or via targeting) and, therefore, inform firms and regulators. Naturally, our survey is restricted to recent developments, without any claim of completeness. However, due to technological progress, new revenue opportunities for websites are likely to arise soon. They are likely to provide new possibilities for innovative business models and create new regulatory and competition policy issues.

We want to point out a number of current phenomena that are not well understood and call for further research. First, we outlined at several points in this paper that the Internet affects the diversity of content. Although content of very different topics and expertise is easily accessible, the Internet is prone to information overload, and so users need to select the content. As a consequence of this selection process, users may browse websites that are pre-selected on the basis of most-read websites. Therefore, an interesting question is whether the Internet broadens or stifles diversity. Second, an important question is whether the Internet is more inclusive than traditional media. For example, the data in Section 2 showed that the Internet is more popular among young people than among older ones. In addition, in more-developed countries, the percentage of Internet usage tends to be higher than in less-developed ones. This implies that the number of recipients of news differs between age groups and regions. This may lead to some individuals or regions lagging behind, which may have adverse effects on economic growth.⁶⁹ For example, poorer and rural users tend to have slower connections and, thus, less opportunity to articulate their opinions. Third, we focused on the content generated by professional

⁶⁸ Anderson and Peitz (2014a) extend their analysis to include advertising as a nuisance. This makes the advertising market two-sided, as consumers now care are about the number of ads carried on each platform.

⁶⁹ For example, Czernich, Falck, Kretschmer and Wössmann (2011), using a panel data set from OECD countries, find that a 10% increase in broadband penetration rates leads to an increase in per-capita economic growth of around 1-1.5%.

websites with the intent of making profits. However, as outlined in Section 4.3, these websites increasingly compete with individuals, such as bloggers or Facebook friends, for the scarce attention of users. The activities of these individuals are not (necessarily) driven by monetary incentives, but, for example, by social concerns such as attention received from their followers or friends. An interesting question is how interaction and competition between “amateurs” and “professionals” play out. In particular, the presence of amateurs may change the business models of for-profit websites.

Finally, we note that many of the theories we presented focused on a particular problem of Internet media. Due to the complexity and the many layers of the problems, there is no unifying framework. However, many of the issues we discussed interact with each other. For example, the user behavior of switching more or switching less between websites affects the emergence and profitability of news-aggregators. Similarly, the possibility of tracking users in a better or worse way has profound consequences on advertisers’ willingness-to-pay and, therefore, on the profitability of Internet media platforms and search engines. Removing or refraining from net neutrality restrictions may change the advertisers’ reservation prices for different content, with consequences for tailoring. While some of these interactions have been addressed in the literature, there remain many potentially interesting links to be explored in further research.

References

- Agarwal, A., Hosonagar, K., Smith, M.D. (2011). "Location, Location, Location: An Analysis of Profitability of Position in Online Advertising Markets". *Journal of Marketing Research* 48, 1057-1073.
- Ambrus, A., Calvano, E., Reisinger, M. (2014). "Either-or-Both Competition: A Two-Sided Theory of Advertising with Overlapping Viewerships". Unpublished Manuscript.
- Ambrus, A., Reisinger, M. (2006). "Exclusive versus Overlapping Viewers in Media Markets". Unpublished Manuscript.
- Anderson, C. (2006). "The Long Tail". Hyperion, New York.
- Anderson, S.P., Coate, S. (2005). "Market Provision of Broadcasting: A Welfare Analysis". *Review of Economic Studies* 72, 947-972.
- Anderson, S.P., De Palma, A. (2009). "Information congestion". *RAND Journal of Economics* 40, 688-709.
- Anderson, S.P., Foros, Ø., Kind, H. (2010). "Hotelling Competition with Multi-Purchasing." CESifo Working Paper 3096.
- Anderson, S.P., Foros, Ø., Kind, H. (2014). "Competition for Advertisers and for Viewers in Media Markets". Unpublished Manuscript.
- Anderson, S.P., Foros, Ø., Kind, H.J., Peitz, M. (2012). "Media Market Concentration, Advertising Levels, and Ad Prices". *International Journal of Industrial Organization* 30, 321-325.
- Anderson, S.P., Neven, D.J. (1989). "Market Efficiency with Combinable Products". *European Economic Review* 33, 707-719.
- Anderson, S.P., Peitz, M. (2014a). "Advertising Congestion in Media Markets". Unpublished Manuscript.
- Anderson, S.P., Peitz, M. (2014b). "Media Seesaws: Winners and Losers in Platform Markets". Unpublished Manuscript.
- Armstrong, M. (2006). "Competition in Two-Sided Markets". *RAND Journal of Economics* 37, 668-691.
- Athey, S., Calvano, E., Gans, J. (2014). "The Impact of the Internet on Advertising Markets for News Media". Unpublished Manuscript.
- Athey, S., Ellison, G. (2011). "Position Auctions with Consumer Search". *Quarterly Journal of Economics* 126, 1213-1270.
- Athey, S., Gans, J. (2010). "The Impact of Targeting Technology on Advertising Markets and Media Competition". *American Economic Review Papers and Proceedings* 100, 608-613.

- Bagwell, K. (2007). "The Economic Analysis of Advertising". In: Armstrong, M., Porter, R. (Eds.), *Handbook of Industrial Organization*, vol. 3. North Holland, Amsterdam.
- Banerjee, A., Chandrasekhar, A.G., Duflo, E., Jackson, M.O. (2014). "Gossip: Identifying Central Individuals in a Social Network". Unpublished Manuscript.
- Becker, G., Murphy, K.M. (1993). "A Simple Theory of Advertising as a Good or Bad". *Quarterly Journal of Economics* 108, 942-964.
- Belleflamme, P., Peitz, M. (2010). "Industrial Organization: Markets and Strategies". Cambridge University Press, Cambridge.
- Belleflamme, P., Peitz, M. (2014). "Price Information in Two-Sided Markets". Unpublished Manuscript.
- Bergemann, D., Bonatti, A. (2011). "Targeting in Advertising Markets: Implications for Offline versus Online Media". *RAND Journal of Economics* 42, 417-443.
- Bergemann, D., Bonatti, A. (2014). "Selling Cookies". *American Economic Journal: Microeconomics*, forthcoming.
- Blake, T., Nosko, C., Tadelis, S. (2014). "Consumer Heterogeneity and Paid Search Effectiveness: A Large Scale Field Experiment". *Econometrica*, forthcoming.
- Börger, T., Cox, I., Pesendorfer, M., Petricek, V. (2013). "Equilibrium Bids in Sponsored Search Auctions: Theory and Evidence". *American Economic Journal: Microeconomics* 5, 163-187.
- Bourreau, M., Kourandi, F., Valletti, T. (2014). "Net Neutrality with Competing Internet Platforms". *Journal of Industrial Economics*, forthcoming.
- Burguet, R., Caminal, R., Ellman, M. (2014). "In Google we Trust?". Unpublished Manuscript.
- Butters, G.R. (1977). Equilibrium Distributions of Sales and Advertising Prices. *Review of Economic Studies* 44, 465-491.
- Cai, H., Chen, Y., Fang, H. (2009). "Observational Learning: Evidence from a Randomized Natural Field Experiment". *American Economic Review* 99, 864-882.
- Calvano, E., Jullien, B. (2012). "Issues in On-line Advertising and Competition Policy: A Two-sided Market Perspective". In: Harrington, J. E., Katsoulacos, Y. (Eds.), *Recent Advances in the Analysis of Competition Policy and Regulation*, Edward Elgar Publishing, Cheltenham.
- Chen, J., Staellert, J. (2014). "An Economic Analysis of Online Advertising Using Behavioral Targeting". *MIS Quarterly*, forthcoming.
- Chen, Y., He, C. (2011). "Paid Placement: Advertising and Search on the Internet." *Economic Journal* 121, F309-F328.

- Chen, Y., Narasimhan, C., Zhang, Z.J. (2001). "Individual Marketing with Imperfect Targetability". *Marketing Science* 20, 23-41.
- Chen, Y. Pavlov, D., Canny, J. (2009). "Large-scale behavioral targeting". In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 209-218. ACM, New York, USA.
- Chen, Y., Wang, Q., Xie, J. (2011). "Online Social Interactions: A Natural Experiment on Word of Mouth versus Observational Learning". *Journal of Marketing Research* 48, 238-254.
- Choi, J.P., Jeon, D.S., Kim, B.C. (2013). "Asymmetric Neutrality Regulation and Innovation at the Edges: Fixed vs. Mobile Networks". NET Institute Working Paper 13-24.
- Choi, J.P., Kim, B.C. (2010). "Net Neutrality and Investment Incentives". *RAND Journal of Economics* 41, 446-471.
- Criteo (2010). "Targeting & Retargeting Interview with Criteo". Report, August 26, available at <http://behavioraltargeting.biz>, last accessed October 16, 2014.
- Czernich, N., Falck, O., Kretschmer, T., Wössmann, L. (2011). "Broadband Infrastructure and Economic Growth". *Economic Journal* 121, 505-532.
- de Cornière, A. (2013). "Search Advertising". Unpublished Manuscript.
- de Cornière, A., De Nijs, R. (2014). "Online Advertising and Privacy". *RAND Journal of Economics*, forthcoming.
- de Cornière, A., Taylor, G. (2014). "Integration and Search Engine Bias". *RAND Journal of Economics* 45, 576-597.
- Dellarocas, C. (2012). "Double Marginalization in Performance-Based Advertising: Implications and Solutions". *Management Science* 58, 1178-1195.
- Dellarocas, C., Rand, W., Katona, Z. (2013). *Media, Aggregators and the Link Economy: Strategic Link Formation in Content Networks*. *Management Science* 59, 2360-2379.
- Dukes, A., Gal-Or, E. (2003). "Negotiations and Exclusivity Contracts for Advertising". *Marketing Science* 22, 222-245.
- Economides, N., Hermalin, B.E. (2012). "The Economics of Network Neutrality". *RAND Journal of Economics* 43, 602-629.
- Edelman, B., Lai, Z. (2013). "Exclusive Preferential Placement as Search Diversion: Evidence from Flight Search". Unpublished Manuscript.
- Edelman, B., Ostrovsky, M., Schwarz, M. (2007). "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars' Worth of Keywords". *American Economic Review* 97, 242-259.

- Eliasz, K., Spiegel, R. (2011). "A Simple Model of Search Engine Pricing". *Economic Journal* 121, F329-F339.
- Esteban, L., Gil, A., Hernandez, J. (2001). "Informative Advertising and Optimal Targeting in a Monopoly". *Journal of Industrial Economics* 49, 161-180.
- Evans, D.S. (2008). "The Economics of the Online Advertising Industry". *Review of Network Economics* 7, 359-391.
- Evans, D.S. (2009). "The Online Advertising Industry: Economics, Evolution, and Privacy". *Journal of Economic Perspectives* 23, 37-60.
- Gabszewicz, J.J., Laussel, D., Sonnac, N. (2004). "Programming and Advertising Competition in the Broadcasting Industry". *Journal of Economics and Management Strategy* 13, 657-669.
- Galeotti, A., Moraga-Gonzalez, J.L. (2008). "Segmentation, Advertising, and Prices". *International Journal of Industrial Organization* 26, 1106-1119.
- Gal-Or, E., Dukes, A. (2003). "Minimum Differentiation in Commercial Media Markets". *Journal of Economics and Management Strategy* 12, 291-325.
- Gal-Or, E., Gal-Or, M. (2005). "Customized Advertising via a Common Media Distributor." *Marketing Science* 24, 241-253.
- Ganuza, J.-J. (2004). "Ignorance Promotes Competition: An Auction Model of Endogenous Private Valuations". *RAND Journal of Economics* 35, 583-598.
- Gentzkow, M., Shapiro, J.M., Sinkinson, M. (2014). "Competition and Ideological Diversity: Historical Evidence from US Newspapers". *American Economic Review* 104, 3073-3114.
- Ghose, A., Yang, S. (2009). "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets". *Management Science* 55, 1605-1622.
- Ghosh, A., Mahdian, M., McAfee, R.P., Vassilvitskii, S. (2012). "To Match or not to Match: Economics of Cookie Matching in Online Advertising". *Proceedings of the 13th ACM Conference on Electronic Commerce (EC'12)*.
- Goldfarb, A. (2014). "What is different about online advertising?" *Review of Industrial Organization* 44, 115-129.
- Goldfarb, A., Tucker, C. (2011a). "Search Engine Advertising: Channel Substitution when Pricing Ads to Context". *Management Science* 57, 458-470.
- Goldfarb, A., Tucker, C. (2011b). "Online Display Advertising: Targeting and Obtrusiveness". *Marketing Science* 30, 389-404.
- Hagi, A., Halaburda, H. (2014). "Information and Two-Sided Platform Profits". *International Journal of Industrial Organization* 34, 25-35.

Hauser, J. R., Urban, G. L., Liberali, G., Braun, M. (2009). "Website Morphing". *Marketing Science* 28, 202-223.

Hoernig, S., Valletti, T.M. (2007). "Mixing Goods with Two-Part Tariffs". *European Economic Review* 51, 1733-1750.

Hotelling, H. (1929). "Stability in Competition". *Economic Journal* 39, 41-57.

Hurkens, S., López, Á.L. (2014). "Mobile Termination, Network Externalities and Consumer Expectations". *Economic Journal*, forthcoming.

IAB (2009). "IAB Social Advertising Best Practices". Interactive Advertising Bureau.

Iyer, G., Soberman, D., Villas-Boas, J.M. (2005). "The Targeting of Advertising". *Marketing Science* 24, 461-476.

Jackson, M.O. (2010). "Social and Economic Networks". Princeton University Press.

Jackson, M.O., Wolinsky, A. (1996). "A Strategic Model of Social and Economic Networks". *Journal of Economic Theory* 71, 44-74.

Jeon, D.-S., Nasr, E.N. (2013). "News Aggregators and Competition among Newspapers on the Internet". Unpublished Manuscript.

Joshi, A., Bagherjeiran, A., Ratnaparkhi, A. (2011). "User Demographic and Behavioral Targeting for Content Match Advertising. In 5th International Workshop on Data Mining and Audience Intelligence for Advertising (ADKDD 2011)". San Diego, CA.

Joo, M., Wilbur, K.C., Cowgill, B., Zhu, Y. (2014). "Television Advertising and Online Search". *Management Science* 60, 56-73.

Johnson, J. P. (2013). "Targeted Advertising and Advertising Avoidance". *RAND Journal of Economics* 44, 128-144.

Johnson, J.P., Myatt, D.P. (2006). "On the Simple Economics of Advertising, Marketing, and Product Design". *American Economic Review* 96, 756-784.

Johnson, G.A., Lewis, R.A., Reiley, D.H. (2014). "Location, Location, Location: Repetition and Proximity Increase Advertising Effectiveness". Unpublished Manuscript.

Katona, Z., Knee, J., Sarvary, M. (2013). "Agenda Chasing and Contests Among News Providers". Unpublished Manuscript.

Katona, Z., Sarvary, M. (2008). "Network Formation and the Structure of the Commercial World Wide Web". *Marketing Science* 27, 764-778.

Katz, M.L., Shapiro, C. (1985). "Network Externalities, Competition, and Compatibility". *American Economic Review* 75, 424-440.

- Katz, M.L., Shapiro, C. (1994). "Systems Competition and Network Effects". *Journal of Economic Perspectives* 8, 93-115.
- Kaye, B.K., Medoff, N.J. (2000). "Just A Click Away: Advertising on the Internet". Allyn and Bacon, Boston.
- Kempe, D., Mahdian, M. (2008). "A Cascade Model for Advertising in Sponsored Search". *Proceedings of the 4th International Workshop on Internet and Network Economics (WINE)*.
- Krämer, J., Wiewiorra, L. (2012). "Network Neutrality and Congestion Sensitive Content Providers: Implications for Content Variety, Broadband Investment, and Regulation". *Information Systems Research* 23, 1303-1321.
- Krämer, J., Wiewiorra, L., Weinhardt, C. (2013). "Net Neutrality: A Progress Report". *Telecommunications Policy* 37, 794-813.
- Lambrecht, A., Tucker, C. (2013). "When Does Retargeting Work? Information Specificity in Online Advertising". *Journal of Marketing Research* 50, 561-5767.
- Levin, J., Milgrom, P. (2010). "Online Advertising: Heterogeneity and Conflation in Market Design". *American Economic Review Papers and Proceedings* 100, 603-607.
- Lewis, R.A., Nguyen, D. (2014). "A Samsung Ad for the iPad? Display Advertising's Spillovers to Search." Unpublished Manuscript.
- Lewis, R.A., Reiley, D.H. (2014). "Online Ads and Online Sales: Measuring the Effects of Retail Advertising via a Controlled Experiment on Yahoo!". *Quantitative Marketing and Economics* 12, 235-266.
- Liaukonyte, J., Teixeira, T., Wilbur, K.C. (2014). "Television Advertising and Online Shopping". Unpublished Manuscript.
- Malthouse, E., Elsner, R. (2006). "Customisation with Cross-Basis Sub-Segmentation". *Journal of Database Marketing & Customer Strategy Management* 14, 40-50.
- Marwick, A. E. (2014). "How Your Data Are Deeply Mined". In: *New York Review of Books*, January 9, 2014.
- Milgrom, P. , Roberts, J. (1986). "Price and Advertising Signals of Product Quality". *Journal of Political Economy* 94, 796-821.
- Nelson, P. (1974). "Advertising as Information". *Journal of Political Economy* 82, 729-754.
- OECD (2014). "Connected Television: Convergence and Emerging Business Models." *OECD Digital Economy Papers*, No. 231. OECD Publishing.
- Oestreicher-Singer, G., Sundararajan, A. (2012). "The Visible Hand? Demand Effects of Recommendation Networks in Electronic Markets". *Management Science* 58, 1963-1981.

- Pan, S., Yang, H. (2014). "Targeted Advertising on Competing Platforms". Unpublished Manuscript.
- Peitz, M., Schuett, F. (2014). "Net Neutrality and the Inflation of Traffic". Unpublished Manuscript.
- Peitz, M., Valletti, T.M. (2008). "Content and Advertising in the Media: Pay-TV versus Free-to-Air". *International Journal of Industrial Organization* 26, 949-965.
- PEW Research Center (2013a). "State of the News Media". Report, available at <http://stateofthemedias.org/>, last accessed March 6, 2014.
- PEW Research Center (2013b). "Social Media Update". Report. December 30, 2013.
- PEW Research Center (2013c). "News Use Across Social Media Platforms". Report. November 14, 2013.
- PEW Research Center (2013d). "Photo and Video Sharing Grow Online". Report. October 28, 2013.
- PEW Research Center (2013e). "Civic Engagement in the Digital Age". Report. April 25, 2013.
- PwC (2013). "IAB Internet Advertising Revenue Report, 2012 full year results". Report. April 2013.
- Rochet, J.-C., Tirole, J. (2003). "Platform Competition in Two-Sided Markets". *Journal of the European Economic Association* 1, 990-1029.
- Roy, S. (2000). "Strategic Segmentation of a Market". *International Journal of Industrial Organization* 18, 1279-1290.
- Rutt, J. (2011). "Aggregators and the News Industry: Charging for Access to Content". NET Institute Working Paper 11-19.
- Rutt, J. (2012). "Targeted Advertising and Media Market Competition". Unpublished Manuscript.
- Rutz, O.J., Bucklin, R. (2011). "From Generic to Branded: A Model of Spillover Dynamics in Paid Search Advertising". *Journal of Marketing Research* 48, 87-102.
- Rutz, O.J., Bucklin, R. (2012). "Does Banner Advertising affect Browsing for Brands? Clickstream Choice Model says Yes, for Some". *Quantitative Marketing and Economics* 10, 231-257.
- Rutz, O.J., Trusov, M. (2011). "Zooming in on Paid Search Ads—A Consumer-Level Model Calibrated on Aggregated Data". *Marketing Science* 30, 789–800.
- Sahni, N. (2013). "Advertising Spillovers: Field-Experiment Evidence and Implications for Returns from Advertising". Unpublished Manuscript.
- Salaverría, R. (2005). "An Immature Medium: Strengths and Weaknesses of Online Newspapers on September 11". *Gazette* 67, 69-86.
- Sandvine (2014). "Global Internet Phenomena Report. 2 H 2013". Report.

- Seamans, R., Zhu, F. (2014). "Responses to Entry in Multi-Sided Markets: The Impact of Craigslist on Local Newspapers". *Management Science* 60, 476-493.
- Sun, M., Zhu, F. (2013). "Ad Revenue and Content Commercialization: Evidence from Blogs". *Management Science* 59, 2314-2331.
- Tarantino, E. (2013). "A Simple Model of Vertical Search Engines Foreclosure". *Telecommunications Policy* 37, 1-12.
- Taylor, G. (2013a). "Search Quality and Revenue Cannibalisation by Competing Search Engines". *Journal of Economics and Management Strategy* 22, 445-467.
- Taylor, G. (2013b). "Attention Retention: Targeted Advertising and the Ex Post Role of Media Content". Unpublished manuscript.
- Tucker, C. (2012). "Social Advertising". Unpublished Manuscript.
- Tucker, C., Zhang, J. (2011). "How does Popularity Information Affect Choices: A Field Experiment". *Management Science* 57, 828-842.
- van Zandt, T. (2004). "Information Overload in a Network of Targeted Communication". *RAND Journal of Economics* 35, 542-560.
- Varian, H. (2007). "Positioning Auctions". *International Journal of Industrial Organization* 25, 1163-1178.
- White, A. (2013). "Search Engines: Left Side Quality versus Right Side Profits". *International Journal of Industrial Organization* 31, 690-701.
- Wilbur, K.C., Zhu, Y. (2009). "Click Fraud". *Marketing Science* 28, 293-308.
- Xu, L., Chen, J., Winston, A. (2010). "Oligopolistic Pricing with Online Search". *Journal of Management Information Systems* 27, 111-141.
- Xu, L., Chen, J., Winston, A. (2011). "Price Competition and Endogenous Valuation in Search Advertising". *Journal of Marketing Research* 48, 566-586.
- Xu, L., Chen, J., Winston, A. (2012). "Effects of the Presence of Organic Listing in Search Advertising". *Information System Research* 23, 1284-1302.
- Zhang, J. (2010). "The Sound of Silence: Observational Learning in the U.S. Kidney Market". *Marketing Science* 29, 315-335.