Three essays in economics of online advertising

Martin Quinn

▶ To cite this version:

2018-ENST-0040

EDITE - ED 130

Doctorat ParisTech

THÈSE

pour obtenir le grade de docteur délivré par

TELECOM ParisTech

Spécialité « Economie »

présentée et soutenue publiquement par

Martin QUINN

le 16 Juillet 2018

Three Essays in

Economics of Online Advertising

Directeur de thèse : **Patrick WAELBROECK**

Jury

M. Paul BELLEFLAMME, Professeur, CORE, Université Catholique de Louvain	Rapporteur
M. David BOUNIE, Professeur, SES, Telecom Paristech	Examinateur
Mme Grazia CECERE, Professeure, LITEM, IMT Business School	Rapporteur
M. Regis RENAULT, Professeur, THEMA, Université de Cergy-Pontoise	Examinateur
M. Patrick WAELBROECK, Professeur, SES, Telecom Paristech	Directeur de Thèse

TELECOM ParisTech

école de l'Institut Mines-Télécom - membre de ParisTech

46 rue Barrault 75013 Paris - (+33) 1 45 81 77 77 - www.telecom-paristech.fr

T H È S E

- Abstract -

The thesis presents an economics analysis of three new subjects related to the economics of online advertising using both empirical and theoretical methodologies.

Chapter 1 introduces the context of the thesis as well as the main results developed in the following chapters.

In **Chapter 2**, the thesis shows through a theoretical model how transparency related to online advertising purchase modifies market equilibria. In a context of advertising avoidance, the chapter shows that introducing such technologies can affect welfare in various ways.

Chapter 3 tackles the economic relevance of profiling technologies that allow websites to adapt their advertising level to user's sensitivity. The chapter builds a theoretical model to draw key implications on how such technologies modify advertising intensity and welfare.

With **Chapter 4**, the thesis highlights how a privacy policy that lets users opt-out from behavioral targeting has economic implications on the online advertising market. The chapter presents a new computational methodology that estimates the effect of such opt-outs on prices of ads sold at auctions. The results differ from previous work and show that the impact may strongly depends on website's characteristics as well as other selling channels attractiveness.

The thesis concludes in **Chapter 5** and points out how key points developed across the chapters carry economic significance as they partly explain how prices and demand for advertising spaces are shaped. It also draws attention on how these new subjects stimulate relevant future research for economists.

For french readers, a summary is available in **Chapter 6**.

Acknowledgements

There is not enough time to aknowledge all the people who helped me during my PhD to the extent they deserve. I hope people will undertand the brevity of this part.

First of all, i would like to express my extreme gratitude to my doctorate advisor Prof. Patrick Waelbroeck for giving me the opportunity to start a research career. I greatly benefited from his guidance during these three years and gained unvaluable insight under his supervision. I also thank him for giving me the freedom to pursue my own research. Secondly, i would like to express my deepest appreciation to Prof. David Bounie for his helpful discussions, valuable suggestions and wise advices. Thirdly, I am grateful to Yann Balgobin, Mattias Mano and Antoine Dubus for the hard work and unvaluable brainstormings. Fourthly, I also would like to thank the Chair VPIP for indispensable financial support and its members for the insightful comments and for trusting me to conduct research. Fifthly and foremost, I would like to thank the members of the jury for accepting to review the manuscripts and participate to the defense of this thesis.

I had the chance to pursue my PhD in a great environment, surrounded by an awesome research team and helpful people. To that, i thank the SES department of Telecom Paristech and all the people whose path i may have crossed. I am appreciative of the help from researchers in University Paris Sud. I particularly show my respect to Fabrice Leguel who passed me on his passion for digital economics during my second year of Master at Telecom Paristech. I also thank members of the CESP for their great insights that appeared to be an important part of the thesis developement.

At last but not least, i would like to dedicate this thesis to my family and friends. This is particularly directed to Charlotte, as she has been the co-author of my life for many chapters, and to my parents and brothers who know how much i owe them.

Declaration

I declare that the thesis has been composed by myself and that it has not been submitted for any other degree or professional qualification. I confirm that the work submitted is my own, except where work which has formed part of jointly-authored publications has been included. My contribution and those of the other authors to this work have been explicitly indicated below. I confirm that appropriate credit has been given within this thesis where reference has been made to the work of others.

The work presented in Chapter 2 has benefited from the co-authorship of Prof. David Bounie and Valérie Morisson. The work presented in Chapter 3 is a joint work with Prof. David Bounie and Antoine Dubus. Finally, Chapter 4 is an interdisciplinary work coauthored with Dr. Vincent Toubiana.

Contents

Chapter 1

Introduction

Since the seminal work of Stigler (1961), advertising has been considered an important subject of research for economists. The approach developed by Stigler (1961) and later by Nelson (1974) was among the first to perform an economic analysis of advertising as *informative*, while advertising had been mostly considered as *persuasive* in the literature. Stigler (1961) considers in his analysis that advertising is an important pillar of information transmission that stimulates more efficiency by informing consumers of the different competitive options on the market. As pointed out by Bagwell (2007), a large body of literature has been built upon *informative* advertising [1], hence providing the basis to understanding the economics of advertising, especially on media markets coordinating the interplay between advertisers and audiences.

However, much of advertising spending is now done on Internet Media. Online advertising revenue has kept growing during the last ten years, establishing at $72B in 2016 according to the IAB (internet advertising revenue report, 2016). Moreover, Adage, 2017 points out that online ad-spending has surpassed TV's for the first time, rising to number one investment in advertising. Such figures have stimulated a new body of literature on the economic stakes of online advertising. According to Peitz and Reisinger (2015), technological progress has given firms unprecedented opportunities to inform Internet users about their product, using new techniques such as targeting. First, online media better match advertising to contents. A specific example is found in *keyword advertising*, where advertisers can choose to show their ads next to appropriate keywords. Second, by leveraging personal information of Internet users, online media are able to better match advertising to their preferences. This can be illustrated by *retargeting adver-*

[1] Parallel to the development of alternative views. See also Renault (2015) for a more complete survey on the role of advertising in markets.

tising where advertisers use past Internet browsing history of consumers to display appropriate ads. Altogether, Goldfarb (2014) argues that advertising in online media exhibits a substantial difference with respect to classical advertising: the cost of targeting is reduced.

Such benefits generate strong investments in the online advertising industry. Indeed, advertisers are better able to segment their online audience and at a lower cost compared to traditional media. However, such wealth fosters a state of continuous innovation, which consequently raises questions. We argue in the thesis that these new challenges faced by online advertising markets importantly affect demand and prices for advertising spaces, and therefore radically change market equilibria. To illustrate this argument, the thesis addresses three new economic concerns on the online advertising market. The first one relates to the lack of transparency on online ads quality, and how it bears economic impact on the online advertising market. Secondly, the thesis analyzes how new profiling algorithms used by platforms to adapt the number of ads to each user's preferences is changing websites behavior. Finally, I develop a new methodology to economically assess the impact of a privacy policy that prevents advertisers from targeting users according to their past online behavior. To an extent, the chapter illustrates how consumer trust can have economic impact on markets.

The thesis contributes to the classical literature on online advertising picturing media markets as two-sided (Rochet and Tirole, 2002; Anderson and Gabszewicz, 2006). Platforms (here content providers) attract Internet users (side 1) and show them a quantity of advertising (the advertisers are therefore side 2). Insuring a good coordination of the market creates incentives for both Internet users and advertisers to join the content provider's website. The recent economic literature on online advertising portrays Internet users perceiving advertising as a nuisance, i.e. exhibiting negative network externalities from the amount of advertising displayed on the website (Anderson and Gans, 2011; Anderson and Coate, 2005). Therefore the more ads there are on a website, the less Internet users will visit it, even if these ads may have informative properties. Hence, a challenge arises for the content provider as he faces a tradeoff between attracting more advertisers while repelling more consumers, or setting a lower number of ads, hence charging fewer advertisers while attracting more consumers.

Each chapter analyzes the economic impact of new challenges the online advertising market faces. They show that such challenges, anchored in the two-sided nature of online advertising markets exhibit sophisticated economic mechanisms. Both theoretical and empirical methodologies are developed across the thesis to address these new challenges.

The first issue addressed by the thesis is related to the lack of transparency regarding advertising purchases on online media markets. On the one hand, the introduction of complex automated mechanisms of advertising sales known as *programmatic* has increased efficiency in time and slot management according to Econsultancy Programmatic Branding report, 2015. On the other hand, according to The Guardian, March 2017 and Marketingweek, March 2017 a setback from programmatic can be found in market transparency. As advertising sales through automated mechanisms is replacing negotiations between humans, advertisers have had trouble verifying the environment in which ads are delivered. More generally, assessing whether the billions of ads delivered each day are fulfilling advertisers expectations is appearing to be quite a challenge for the market.

Three components of these expectations can be analyzed. First, advertisers have been complaining about purchasing advertising that is not seen by any Internet users. Indeed, different analyses such as comScore, 2013 have highlighted that around half of "display" ads[2] are not seen by Internet users because they are not within the boundaries of the *active view* (i.e. what part of the page the screen is showing). Second, other studies are pointing out that even if ads are viewable, they are mostly seen by robots and not by humans. Report from the Association Of National Advertisers, 2016 and Adloox, 2017 says that advertising bots fraud induced a loss between \$7.2 and \$12 billion for the industry in 2016. These issues represent a huge loss for advertisers who have been paying to display advertising, and are currently finding out that they are losing a lot revenue. Finally, brand safety issues have been recently shaking the online advertising industry. According to Business insider, December 2017 some advertisers running their campaigns on YouTube recently found that their ads were running next to inappropriate content (such as terrorist or racist videos). Backlash from the markets ensued and forced YouTube to publicly apologize.

These three illustrations underline how uncertainty regarding the context in which advertising spaces are served plays an important role on advertiser' willingness to pay. Allowing advertisers to monitor online ad quality impact may change market equilibria as it patches the lack of trust between brands and websites. Nevertheless, assessing the overall impact of the introduction of such technology is not easy. On the one hand it may reduce information asymmetry on the

[2] Ads displayed on websites or apps

market. On the other hand, such technology may end up degrading Internet users experience when visiting websites, who may consequently prefer to avoid ads.

Chapter 2 presents a theoretical model that underlines how the impossibility to control online ads environment prevents content providers from committing to a certain level of quality (Business Insider, January 2017).

The chapter focuses on the viewability of advertising: are ads viewable by Internet users when served on a specific webpage? This aspect of advertising appears to be very significant for the economics of online ads, especially in the context of branding campaigns where ads are bought *per impression* (CPM) and not *per click*[3], as advertisers pay ads when "served" (i.e. displayed) on the website. According to Comscore, Q1 2017, the ad viewability has not changed much in 4 years, since only around 50% ads served are effectively seen by Internet users in 2017. The lack of viewability is considered as a big issue since advertisers pay for ads that are not displayed to Internet users.

To overcome this hurdle, the market already developed technologies assessing the viewability level of ads. Using geometrical measures, advertisers recover the location of ads within the webpage and are able to assess how long and how much of their ad has been displayed in the active view. To foster the adoption of such technology, the Internet Advertising Bureau (IAB) and the Media Rating Council (MRC) have introduced viewability standards, counting as "viewed" any standard ad of which at least 50% of the pixels is displayed in the active view of the webpage for at least one second (IAB Measurement guidelines, 2014). [4]

The chapter analyzes how introducing such technology may impact industry profits and social welfare. Firstly, Internet users tend to see more ads when technologies allowing advertisers to verify viewability are introduced. In this case, the content providers are not constrained anymore and are able to commit to a specific level of viewability of ads. Moreover, in a context where advertising is perceived as a nuisance by Internet users, the introduction of viewability technology is beneficial for social welfare if the nuisance cost of advertising is not too high. Indeed, introducing a viewability technology increases the advertising intensity level, which consequently

[3]Advertisers distinguish two types of campaigns: branding campaigns whose goal is to have brand coverage (usually paid to be displayed), and performance campaigns whose goal is to transform into online purchase (usually paid once consumers have clicked the ad).

[4]For non-standard banners, viewability measurement criteria may change. For example advertising videos are counted as viewed when at least 50% of the pixels has been displayed in the active view of the webpage for at least two seconds.

reduces consumer surplus.

Secondly, introducing viewability technologies may create incentives for Internet users to avoid ads using for example ad-blockers. As introducing such technologies (weakly) increases the number of ads Internet users see when visiting a website, they are more likely to use an ad-blocker as the advertising nuisance is higher.[5] The chapter shows that introducing such technologies reduces uncertainty on the market by improving transparency on advertising quality, but puts websites under pressure. On the one hand, publishers' efficiency is now monitored by advertisers during campaigns, which induces them to increase the number of ads to show to Internet users. On the other hand, displaying more ads may create incentives for Internet users to avoid ads using ad-blockers. Such practice is highly monitored by the online advertising industry as it damages content providers and advertisers profits, which brings us to the next chapter.

The second concern analyzed in the thesis addresses the impact of advertising nuisance on media markets. According to a WordStream study, only 0.9% of Facebook displayed ads are being clicked on by users. Besides being uninterested in advertising, Internet users simply prefer to avoid them. As Yougov study, 2016 points out, 61% of adults surveyed dislike advertising. Consequently, consumers are massively installing advertising avoidance technologies (AATs) (PageFair, 2016) exhibiting the fact that online advertising is mostly perceived negatively.

Internet users dislike advertising - or at least advertising the way it is practiced by content providers, according to several studies such as Yougov study, 2016 or Hubspot research, July 2016. For this reason, more and more Internet users adopt Advertising Avoidance Technologies (AATs) - ad-blockers - such as Adblock Plus, when browsing online. PageFair, 2016 points out that at least 380 million mobile devices and 236 million desktop devices were blocking ads in December 2016, representing 11% of the global Internet population. Moreover, browsers such as Brave browser, Firefox Focus or the next version of Chrome browser launching in 2018 (which have around 60% market share according to NetMarketShare 2017) have built-in adblockers according to Digiday, April 2016, which increase even more the reach of advertising avoidance.

The two-sided nature of ad-financed media markets is questioned as Internet users choose to avoid ads. Indeed, as Internet users avoid ads, there is no profit generated on the advertisers' side. Consequently, content providers are starting to look for new ways to generate revenue. One of the solutions for websites has been to improve the match between advertising content

[5] This is mainly due to the two-sided market nature of ad-financed media markets.

and Internet users. Such solution may improve user experience and induce Internet users to stop avoiding ads. For example, it can determine the current interests of the Internet users and display related advertising. In this vein, new technologies provide websites with the capacity to entirely customize experience leveraging users' personal information. The type and number of content can therefore be tailored to each Internet users visiting the website. Tracking technologies are exploited to capture Internet users' traces when browsing online. These traces can take different forms: they can be the results of browsing a sport website, watching a comedy video or visiting a brand page on a social media. Using these traces, algorithms understand the preferences of a specific Internet user and to tailor their browsing experience. For example, an algorithm understanding that a specific Internet user likes golf and is not willing to see an advertising video more than 8 seconds, will adapt the website in order to match the user preferences and insure that he chooses not to avoid advertising.

Such practice may relate to first-degree price discrimination where firms can price their products to the valuation of each user. In this situation, firms can show the maximum amount of ads that a user is willing to tolerate. With these new tools, the browsing experience of an Internet user visiting a website can be completely optimized. But these technologies may have a broader impact as they change the amount of advertising on the market, which is considered as a proxy for product demand and competition by economists. Moreover, the efficiency of the technology uses Internet users online past traces, which may not always be available.

Chapter 3 addresses advertising avoidance issues and analyzes how the use of profiling technologies to counter them engenders economic implications. More precisely, considering that Internet users are heterogeneous with respect to their preferences to advertising, the chapter analyzes the implications of introducing a technology that can capture this sensitivity and customizes advertising accordingly.

The chapter studies through a theoretical model how introducing a profiling technology that can tailor advertising experience to not only the tastes of Internet users, but also their distastes for advertising, may change market equilibria. I consider two types of users: some consumers have a *high taste* for ads, while a complementary set exhibits a *low taste* for ads. We assume that low taste users are more easily annoyed with advertising, hence more likely to avoid ads. The profiling technology classifies users according to a technology that produces a signal on their ad sensitivity. The veracity of the signal is considered exogenous and can be imperfect.

Firstly, the chapter shows that a *perfectly efficient technology* (a technology that makes no mistakes classifying users) is always adopted by the website as it increases his profits on the high taste users (showing them more ads) and attract more low taste ones (less ads are shown to them, fostering users not to install ad-blockers), increasing therefore the total number of ads served on the market. As we find that such technology is always profit increasing for websites, it is not always the case for the total welfare. Indeed, we find that the total welfare always increases with the introduction of a perfectly efficient profiling technology only if the website is facing a high taste audience.

Secondly, the analysis gets more complex when 1) the website uses a perfectly efficient technology and is facing an audience with a low taste for ads and 2) when the website adopts a *imperfectly efficient technology* (a technology that is efficient enough to be adopted by the website while still making mistakes when classifying users). Moreover, the chapter considers the efficiency of the technology as exogenous, whereas it may be related to investment in statistical method as well as the quality and volume of available user data. We study the latter in the following chapter of the thesis.

The third contribution of the thesis is related to the use of consumers' personal information to optimize tailored advertising. A 2015 Pew Internet Research study points out that 93% of those surveyed in the report say that being in control of *who* gets their information is important. Such concerns stress the need to regulate firm tracking practices. On the one hand, data-driven industries such as online advertising have flourished in leveraging users' data and tracking technologies. On the other hand, such technology may hurt Internet users who want to protect their privacy.

To balance both effects, regulators have considered different privacy policies to take into account privacy concern and business practices. For example, US regulators have implemented an *opt-out* policy, where advertisers can track users by default but these same users can opt-out from this tracking. Simultaneously, forthcoming European regulation intends to implement an *opt-in* policy. Such regulation prevents advertisers from tracking Internet users, unless an explicit consent is collected.

Privacy policies may have an important economic impact on the online advertising industry, as advertisers integrate user information in their strategy to serve ads. The fact that advertisers may not be able to tailor ads to Internet users past behavior, have two implications. Firstly, it may lower the probability of click as the match between the ad and the user is of lower quality. Secondly, it may simply foster Internet users to avoid ads. Altogether, an opt-out from behavioral targeting is supposed to decrease the willingness to pay of advertisers.

Chapter 4 assesses how introducing an opt-out privacy policy impacts ad prices in a context of online advertising auctions. Previous works have already established that an opt-out policy may decrease ad prices and by extent advertisers and websites revenue.

In the chapter, I introduce a new computational methodology that captures ad prices before and after a user opted-out. More precisely, I create two groups of bots[6] with exactly the same characteristics. One group is the *control* group and the other one is the *treatment* group. I establish a list of websites where ad prices can be recovered and send the two groups to visit this same list twice in a row. Between the two visits, bots in the treatment group perform an opt-out using the AdChoice website[7], while bots in the control group remain inactive. From that experiment, we are able to create a database with prices of ads seen by bots from the control and the treatment groups, before and after the treatment takes pace. We repeat this experiment for 51 days and perform a difference-in-difference econometric analysis that assesses the impact of an opt-out from behavioral targeting on ad-prices. The paper exhibits three important features.

Firstly, the methodology precisely recovers prices before and after a specific user behavior, which is not possible in previous work from the literature. In this chapter, we test whether an opt-out from behavioral targeting impacts advertising prices. However, such methodology can be easily extended to test the impact of other policies affecting users' behavior. Secondly, the chapter finds that an opt-out from behavioral targeting weakly increases prices of online advertising sold at auctions. More precisely, we find that once bots have performed an opt-out on the AdChoice website, we see ads sold at a higher price. We contrast this result as we show how such effect may depend on the type of website where the auction takes place. This can be explained by the fact that our bots may be of low value for advertisers. Hence when advertisers don't have any information, they may choose to bid at a higher price. Thirdly, we find that an opt-out privacy policy significantly reduces the number of ads sold at auction for one user during an experiment (a day in this case). This is important as it highlights a competition between advertising selling mechanisms. This also may question the significance of our results as we

[6]Bots are computer programs that can behave according to a set of rules given by the creator.
[7]AdChoice is a service that allows users to opt-out from behavioral targeting for all advertisers at once

do not account for potential auction that did not take place due to a decrease in advertisers' participation. We offer comments and a solution in the conclusion of the chapter

Such study may be particularly relevant as the forthcoming European privacy regulation GDPR, as well as the next ePrivacy directive aim at implementing an opt-in privacy policy. Such policy requires for websites and advertisers to obtain clear consent before leveraging users' data to tailor ads. Future potential application of the methodology developed in the chapter may be applied to understanding economic effect of opt-in policies.

Large investments in advertising technology have driven innovation and market to thrive, solving old problems while creating new issues. The thesis analyzes the economic impact of these new issues rising from new online advertising technologies in **Chapters 2, 3 and 4**, while interrogating future stakes of the market in **Chapter 5**.

Bibliography

Anderson, Simon and Jean Gabszewicz, "Chapter 18 The Media and Advertising: A Tale of Two-Sided Markets," in Victor A. Ginsburg and David Throsby, eds., *Handbook of the Economics of Art and Culture*, Vol. 1, Elsevier, 2006, pp. 567 – 614.

Anderson, Simon P. and Joshua S. Gans, "Platform Siphoning: Ad-Avoidance and Media Content," *American Economic Journal: Microeconomics*, November 2011, *3* (4), 1–34.

_ **and Stephen Coate**, "Market Provision of Broadcasting: A Welfare Analysis," *The Review of Economic Studies*, 2005, *72* (4), 947–972.

Bagwell, Kyle, in Mark Armstrong and Robert Porter, eds., *The Economic Analysis of Advertising*, 1 ed., Vol. 3, Elsevier, 2007, chapter 28, pp. 1701–1844.

Goldfarb, Avi, "What is different about online advertising?," *Review of Industrial Organization*, 2014, *44* (2), 115–129.

Nelson, Phillip, "Advertising as Information," *Journal of Political Economy*, 1974, *82* (4), 729–754.

Peitz, Martin and Markus Reisinger, "Chapter 10 - The Economics of Internet Media," in Simon P. Anderson, Joel Waldfogel, and David StrÃűmberg, eds., *Handbook of Media Economics*, Vol. 1 of *Handbook of Media Economics*, North-Holland, 2015, pp. 445 – 530.

Renault, RÃl'gis, "Chapter 4 - Advertising in Markets," in Simon P. Anderson, Joel Waldfogel, and David StrÃűmberg, eds., *Handbook of Media Economics*, Vol. 1 of *Handbook of Media Economics*, North-Holland, 2015, pp. 121 – 204.

Rochet, Jean-Charles and Jean Tirole, "Cooperation among Competitors: Some Economics of Payment Card Associations," *The RAND Journal of Economics*, 2002, *33* (4), 549–570.

Stigler, George J, "The economics of information," *Journal of political economy*, 1961, *69* (3), 213–225.

Chapter 2

The Economics of Online Advertising Viewability

2.1 Introduction

Digital medias are attracting more and more advertising spending. According to eMarketer, the UK became in 2015 the first country in the world where digital media took a 50% share of advertising spending.[1] In the US, also according to eMarketer's forecast, online ad spending will surpass TV ad spending in 2017 for a total of about $77 billion, driven mostly by mobile that accounts for more than 50 per cent of total online digital spending (eMarketer, Digital Ad Spending to Surpass TV Next Year, 2016).

The rapid development of mobile usage alone does not explain the growth of online advertising spending. Programmatic advertising and the ability to collect data on consumers and ad impressions[2] allow advertisers to automate the buying and selling of ads and to achieve an effective personalized targeting of audiences. They are therefore in a better position with respect to the TV and print media to estimate how successful a particular ad is in driving a purchase decision or in raising brand awareness over time.

However, the promises of online advertising in the case of branding campaigns that rely on serving millions of ads to Internet users are today challenged. Indeed, the promises rest on the assumption that the served ad impressions are *viewable* by Internet users, i.e. "contained in the viewable space of the browser window, on an in-focus browser tab, based on pre-established cri-

[1]The US online advertising spending will amount to $77 billion (eMarketer, "UK to Achieve World First as Half of Media Ad Spend Goes Digital," 2015).

[2]The display of an ad in a page view is called an ad impression.

teria such as the percent of ad pixels within the viewable space and the length of time the ad is in the viewable space of the browser" (Internet Advertising Bureau Europe, Viewable Impressions, 2015). Viewable in this context simply means that Internet users have the opportunity to see the ad, regardless of whether they have actually seen it.

This simple assumption is however challenged by companies such as Google, comScore, Nielsen, etc, that daily analyze billions of impressions from campaigns over thousands of publishers: most of served impressions are actually never seen by Internet users. A well-known commented statistic released by comScore in 2013 indicates for example that half of the publishers' inventory is not seen by Internet users.[3] In 2016, as the Section 2.2 of this paper will show, the proportion of ads being seen by people in most of the countries around the world is still relatively low, between 40% and 50%.[4] The most popular social network website Facebook that attracts the major part of ad investments is also subject to criticisms: "Facebook ads are far less viewable than people [advertisers] were expecting" (Business Insider, December 28, 2016).

Ad viewability became therefore in recent months one of the top priorities on the agenda of advertisers (Wall Street Journal, "It's How Long Ads Are Viewed That Really Matters", February 4, 2016).[5] The concerns are perfectly understandable. In the case of branding campaigns, advertisers pay for ads by the number of impressions that a publisher has *served* (this trading currency is called "Cost-Per-Mille" (CPM)).[6] But as half of ads purchased by advertisers are never seen by Internet users, they potentially waste half of their budgets every time they pay for display ads. For the specific case of the UK for example, Meetric estimates that advertisers wasted over £600m on non-viewable ads in 2016 (Meetrics, Benchmark, 2017). Consequently, more and

[3] Different reasons explain why ad impressions are not viewed by Internet users. Firstly, the browsing behavior encompasses many possibilities to avoid the sight of an ad such as scrolling the page, resizing the window, using an ad blocker, etc. Secondly, publishers may adjust the viewability of ads to preserve user experience. In Section 2.2 of this paper, we review some of the main factors that could explain the low level of ad viewability.

[4] Ad viewability is not a new issue in media but because of the size of online markets, the problem has definitely grown and becomes a serious threat for the advertising industry. For print media, the likelihood that a reader actually sees an ad on a given page is not precise (except with QR codes). Regarding television, a commercial is supposed to be seen as soon as there is a person in a room with a TV set on. The measure is not perfect (people walking out of the room during commercial breaks, fast-forwarding through recorded ads, etc.), but the opportunity for exposure exist.

[5] During the last edition of the annual Digital & Social Media Conference in 2016, the CEO of the US Association of National Advertisers touched on media's "Big Four" concerns: ad blocking, ad fraud, media transparency and viewability/measurement.

[6] Performance-based online campaigns use another well-known metrics which is the Click-Through-Rate, i.e. the percentage of ad impressions that have led to clicks. In this paper, we only focuses on branding-based online campaigns.

more advertisers demand to pay only for *viewable* impressions and not for served impressions. A new trading currency is therefore emerging: the viewable CPM (vCPM) that prices ads by the number of impressions that can be viewed by Internet users, instead of just being served.

The adoption of technologies to measure the viewability of ads served by publishers, and subsequently the shift towards a new trading currency, may entail serious changes in the economics of online advertising. To begin with, publishers need to redesign their websites[7] to make ads more viewable and satisfy advertisers' requests to remain competitive. A large part of the current inventory with very low viewability could therefore not be sold anymore, or at a lower rate, which should decrease the revenue streams of publishers. But as publishers may have less inventory to sell, some rates might also increase (premium inventory), affecting in turn advertisers competition for high ad viewability.[8] Furthermore, as websites are redesigned to enlarge the amount of space for viewable ads to the detriment of editorial contents, the audience could shrink and the price of ads drops accordingly. The economic impact of introducing a viewability technology is therefore split between the loss in user experience and the gain in revenues from online advertising.

The objective of this paper is to analyze how the growing adoption of an ad viewability technology can affect the economics of online advertising. We construct a two-sided market model where a monopoly publisher displays an editorial content to attract Internet users on one side and advertisers on the other side. Inspired from Anderson and Coate (2005), the publisher is only financed by advertisers to display ads that are perceived by users as a nuisance. We compare two situations. In the first situation, advertisers do not have a technology to measure the viewability of ads on the publisher website. They just anticipate a global level of ad viewability. In the second situation, advertisers have a technology to precisely measure ad viewability.

Firstly, our benchmark model underlines that the use of a technology to measure ad viewability tends to increase the number of viewable ads, and publisher and advertisers' profits, but in return degrades user experience. Hence, the adoption of ad viewability improves the situation of the advertising industry (publisher and advertisers) to the prejudice of Internet users. Overall, the total welfare can be greater with ad viewability when the content quality offered by the monopoly

[7]The Guardian is a typical example. The British media website redesigned his website to render more visible ad placements (The Guardian, Optimising for viewability, 2016).

[8]According to Quantcast in 2016, inventory with viewability above 75 per cent can be up to two times more expensive than the average (Quantcast, The Road to Viewability, 2016).

publisher is relatively high and the nuisance cost of ads is not too high. This is even more relevant as the introduction of a viewability technology increases incentives for the publisher to invest in content quality.

Secondly, we extend the benchmark model to ad-blockers. In our set-up, the drop in user experience due to a higher ad viewability can be restored by adopting ad-blockers. In this case, the publisher is constrained by both sides of the market: on the one hand, it must increase the number of viewable ads to satisfy advertisers' requests, and on the other hand, it needs to lower the number of viewable ads to discourage people from installing ad-blockers. We show that the publisher is more likely to be contrained by ad-blocker when a viewability technology is introduced, as it increases disutility from advertising.

Thirdly, we build a model of "bottleneck competition", where two publishers compete for users attention on a Hotelling line. Our model shows that in the case of the two publishers being asymetric regarding their viewability capacities (i.e the maximum viewability a publisher can set for an ad), introducing a viewability technology allows to restore competition.

This paper contributes to the economics and management literature on online advertising on two points. Firstly, this paper provides the first comprehensive economic analysis of the thorny issue of ad viewability largely debated in the advertising industry but absent from the academic research. Secondly, our paper enriches the literature on Internet media (Peitz and Reisinger, 2016), and online ad effectiveness (Goldfarb and Tucker, 2011; Manchanda et al., 2006; Lambrecht and Tucker, 2013; Goldfarb and Tucker, 2015). In previous contributions, consumers like ads when they are targeted (de Cornière, 2016; Johnson, 2013), or dislike ads when they are too much intrusive (Ghose and Yang, 2009; Agarwal et al., 2011; Rutz and Bucklin, 2011; Blake et al., 2015), resulting respectively in a higher or lower demand of Internet users. But as Section 2.2 of this paper will show, there are many reasons for which targeted or intrusive ads are never seen by Internet users, regardless of whether they like ads or not. Taking ad viewability into account is therefore crucial as ads that are not or partially seen are still paid by advertisers and do not have any chance to reach consumers and to be effective.

The remainder of the paper is structured as follows. In Section 2.2, we first define the concept of ad viewability and provide some market insights. In Section 2.3, we review the academic literature. In Section 2.4 we present the setup of the and the results of the benchmark model in a monopolistic competition. In Section 2.5, we calculate the impact of ad viewability on total welfare. In Section 2.6, we extend the model to account for ad-blockers and competition. Section 2.7 presents our conclusion.

2.2 Ad Viewability: Definition and Market Insights

Online advertising requires the Internet to deliver marketing messages to promote a brand to consumers, to sign up for membership or to make purchases. To do so, marketers can use many types of ads (or creatives) such as banners, videos, etc., on desktop (personal computer) and mobile environments.

Different participants are involved in online advertising such as the publisher who places ads into his online content, the advertiser, who provides the ads to be displayed on the publisher's website, and potentially many other intermediaries (ad networks, data management platforms, media agencies, etc.). With the recent development of advertising technologies (adtech), publishers and advertisers manage less and less manually the ads on websites. Ads are served automatically by ad servers. To measure how often impressions are delivered to Internet users, publishers, advertisers and ad servers mostly use tags, a piece of HTML or JavaScript code placed on each creative to provide a complete view of campaign delivery. The tags are usually provided by a viewability vendor.

The mission of a viewability vendor is to measure the number of served and viewed impressions. The number of served impressions is just the number of tagged impressions. But not all served impressions are necessarily measured by vendors because of network failures and invalid (non-human) traffic issues.[9] For example, some ads can be tagged but not correctly delivered or fraudulently served to spiders and bots to manipulate legitimate ad serving. As a consequence, a second measure named the "number of measured impressions" is important to consider as it cleans up invalid traffic and non-served impressions. Finally, ads can be correctly served and measured but not seen by users for several reasons. For example, the ad can be served below the fold (i.e. outside the viewable browser space) far down at the bottom of a web page. Consequently, "a served ad impression can be classified as a viewable impression if the ad is contained

[9] One of the largest studies was conducted and published in December 2014 by the Association of National Advertisers (ANA) in the US and an online fraud detection firm, White Ops. According to the numbers, 11% of display and 23% of video impressions were bot-driven.

in the viewable space of the browser window, on an in focus browser tab, based on pre-established criteria such as the percent of ad pixels within the viewable space and the length of time the ad is in the viewable space of the browser" (Media Rating Council (MRC), Viewable Ad Impression Measurement Guidelines (Dekstop), 2014). The rate of ad viewability is therefore the ratio of the number of viewable impressions over the number of measured impressions.

The pre-established criteria mentioned in the quotation above have been formally defined for different ad formats by the MRC in 2014 and 2016. A display ad is considered viewable when 50 per cent of an ad's pixel are in view on the screen (on an in-focus browser tab on the viewable space of the browser page) for a minimum of one continuous second. This standard is valid for most banners but has been extended for large ad size banners: a viewable impression may be counted if 30 per cent are in view for a minimum of one continuous second. Regarding videos, it is required that 50 per cent of an ad's pixel are in view on the screen and that two continuous seconds of the video are played. Finally, regarding mobile ads, the MRC has issued its first set of guidelines last April 2016 and recommends to treat smartphone (excluding apps) and desktop ads the same: 50 per cent of an ad's pixel are in view on the screen for a minimum of one continuous second.[10]

Since 2012, numerous studies conducted by viewability vendors have measured the viewability of publishers' ad inventories. All studies conclude that a significant proportion of delivered ad impressions are never visible to the end user, resulting in relatively low viewability rates. comScore has been the first viewability vendor to conduct such analysis over thousands of campaigns spanning a mix of global advertisers who ran their ads across a variety of publisher sites and ad networks from May 2012 through February 2013. The key finding was that 54 per cent of display ads do not have the opportunity to be seen by a consumer (comScore, Viewability Benchmarks Show Many Ads Are Not In-View but Rates Vary by Publisher, 2013). Since this first and well commented statistic, other studies have confirmed this finding even if significant increases have been observed in countries like France more recently: +7.4 points between Q4 2015 and Q1 2016, and +13.1 points in one year (Integral Ad Science report: Q1 2016 International Media Quality Report). In addition, high viewability inventories are relatively rare. Quantcast for

[10] As indicated, the MRC standards value a one second impression the same as a five second impression. As a consequence, alternative trading currencies such as the 'Cost Per Hour' are being experimented by large publishers such as The Financial Times to value ad exposure time as a key dimension (Sanghvi, 2015).

example finds that "there is a very limited supply of very high viewability inventory, with viewability above 80% constitutes just two to three percent of all RTB inventory in Europe (Quantcast, Viewability: What Smart Marketers Need to Know, 2016)." In the specific case of videos, Google conducted in 2015 a study of the video advertising platforms, including Google, DoubleClick, and YouTube (Google, Are Your Video Ads Making an Impression?, 2015). He finds that 54 per cent of the videos are viewable on the web across desktop, mobile and tablets (not including YouTube).

Ad viewability also varies significantly across countries. According to Meetrics, another viewability vendor, the viewability rate for digital display ads in France stood at 65 per cent in Q4 2015, compared with 50 per cent in the UK, the lowest viewability rate than for any other country in Europe tracked by the firm (Meetrics, Viewability Benchmark, 2016). In the case of videos, Google also reports that viewability drastically varies between countries from 86 per cent for Russia for example to 54 per cent for the United States (Google, Are Your Video Ads Making an Impression?, 2015).

The gaps in viewability between countries may be due to several reasons. To begin with, the technologies used by viewability vendors differ in many ways as they do not use the same technologies to control for invalid traffic issues for example. In 2016, the syndicate of internet sales (SRI) as well as the association of media agencies (UDECAM) in France commissioned the CESP to review and compare eight different viewability measurement solutions, namely Adloox, Adledge, comScore, Integral Ad Science, Meetrics, MOAT, and two tools natively implemented in platforms (AppNexus and Google). Based on tests made by four major viewability vendors about ads placed above the fold on well-known websites,[11] CESP has reported discrepancies about the average rate of viewability between the four actors up to 36 percentage points (CESP, Mission visibilité de la publicité digitale, 2016).

Next, the publishers' strategies about the placement of ads within webpages and websites may considerably affect the viewability of ads. This is one of the first conclusions drawn by comScore early in 2013. Regardless of the publisher type, the reports emphasized that it is important to evaluate the individual publisher or network on its own merits. "The wide viewability range suggests that regardless of the publisher type, there are some members of the sell-side of the

[11] Above the fold is the upper half of the front page of a webpage.

market who are delivering very strong in-view rates and others who are falling short on their intention to deliver valuable ad inventory to advertisers." For example, for premium websites having an average CPM above USD 5.00$ and 100,000 in monthly ad revenue, the viewability ranged between 10 and 80 per cent.

Ads placed at different page depths are therefore central for viewability as ads have different likelihoods of being viewed by users. Traditionally, above the fold (ATF) has been considered as the top location by advertisers because the ad is supposed to be directly viewable in a browser window when the page first loads. But recent studies tend to refute this common belief. Quantcast, in 2015, shows for example that "ATF is a poor proxy for viewability, with one exchange at only 44% viewability rate on ATF inventory" (Quantcast Report: High Viewability Expectations Can Harm Campaign Performance). Several reasons may explain this result: first, users quickly scroll down to reach their desired destination, and second hyperlinks do not always link to the top of a page. In this respect, the format of the ad may help publishers and advertisers to attract consumers.

Viewability also varies between static banner ads and dynamic rich media. Sizmek Research analyzed in 2016 viewable data from more than 240 billion measured impressions from more than 840,000 ads and 120,000 campaigns served in 74 countries and six continents to more than 22,000 publishers and 43 programmatic partners from January 1 to December 31, 2014 (Sizmek, Viewability Benchmarks, 2016). Sizmek notes that "flash rich media ads were 18% more likely to be seen than standard banners. This was most pronounced in North America, where rich media was 29% more likely to be seen than standard banners." This finding is also confirmed by an Adform study in 2015 that found that "rich media display ads performed especially well in the UK, with such ads on its platform seeing a 71.2 per cent viewability rate. In particular, those ads saw some of the best engagement rates and times compared with other countries covered in the study." Moat analytics finally also confirms that bigger ads are more likely to be seen than smaller ones.

Ad viewability also varies by device. For example, the same report by Sizmek Research shows that mobile is generally more viewable than desktop sizes to the exception of Flash rich media. The gaps are substantial and up to 31 points of percentage for HTML5 standards banners in desktop and mobile environments. This finding is also corroborated by Google about video ads. The study reports that video ads are significantly more viewable on mobile (83 per cent, excluding apps) and tablet (81 per cent) than on desktop (53 per cent) (Google, Are Your Video Ads Making an Impression?, 2015).

To conclude this section, it is worth noting that the lack of ad viewability is not only a big waste of budgets for advertisers as indicated in Introduction, but also a matter of brand and sales impact. For example, a recent study conducted by comScore, Millward Brown & Kantar Worldpanel in April 2016 finds evidence that ads in view for longer periods increase both awareness and purchase intent metrics compared to those in view for less time (comScore, Millward Brown & Kantar Worldpanel, 2016, How Delivery and Brand/Sales Effectiveness Can Drive Digital Advertising ROI in a Cross-Media World). Understanding the issue of ad viewability is therefore a key challenge for advertisers, but also for publishers who need to demonstrate the quality of their inventory.

2.3 Related Literature

Very few academic papers have been devoted to the issue of ad viewability even though this topic is of special importance for the advertising industry. To the best of our knowledge, two papers only have been published in computer science and only one in advertising research.

The objective of the two papers in computer science is to better understand and improve ad viewability. For example, Wang et al. (2015) propose a model supposed to better predict the viewability of any given scroll depth for a user-page pair compared with other systems. In particular, they identify two features such as user geo-location and device type that have a significant impact on the maximum scroll depth. Zhang et al. (2015) investigate what percentage of viewable pixels and length of exposure time may encourage users' ad recall. They find that 75% of the ad's pixels being shown at least two seconds in the active page insure the ad to be seen by users.

Regarding research in advertising, Flosi et al. (2013) use a 2-million-person panel and census server data (cookie data) provided by comScore in 2013 to understand the extent to which ads are delivered to the right target audience. Several empirical generalizations are proposed from the study findings about cookie-related issues, viewability, geo-targeting, and non-human traffic (fraud). For the authors, viewability is a critical component of campaign validation. Several findings are commented. Firstly, the authors find that "on average, 30 percent to 37 percent of all served advertising impressions in the United States, Europe, and Canada were never actually viewable by the end user." Secondly, viewability rates vary significantly across sites and campaigns and, finally, the prices of ads are not correlated to viewability rates. This last finding may be today surprising but it is worth noting that, at the time of this study, the technologies to measure ad viewability had not yet been adopted by advertisers and publishers.

Despite the lack of research dedicated to ad viewability, online advertising has been largely studied in economics and management science (Anderson and Coate, 2005), especially with the popularity of two-sided markets (Caillaud and Jullien, 2001, 2003; Rochet and Tirole, 2002; Anderson and Gabszewicz, 2006). An excellent survey on the evolution of the online advertising business is provided by Evans (2009). The author describes how online advertising has transformed media businesses and allowed pure Internet players to compete with traditional firms. New technologies emerged allowing a better match between advertisers and consumers, transforming in turn online advertising into a reliable source of revenue. A more recent contribution by Anderson and Jullien (2016) surveys recent models of advertising in media markets developed around the concept of two-sided markets. Our paper does not directly contribute to the theory of two-sided markets but simply relies on this powerful tool to understand how publishers manage to coordinate the two sides of the market, Internet users and advertisers.

Our paper pertains more precisely to the relatively new literature on the effectiveness of online advertising in management and marketing science.[12] Manchanda et al. (2006) measure the impact of banner advertising on purchasing patterns on the Internet. The results show that the number of exposures, the number of websites, and the number of webpages all have a positive effect on repeat purchase probabilities, whereas the number of unique creatives has a negative effect. Goldfarb and Tucker (2011) explore the factors that influence the effectiveness of online advertising. They find that matching an ad to website content and increasing an ad's obtrusiveness independently increase purchase intent. However, in combination, these two strategies are ineffective. Lambrecht and Tucker (2013) measure and compare the effectiveness of dynamic retargeting (information from internal browsing data from consumers who previously visited the advertisers' website) to simple generic brand ads. They find that dynamic retargeting is less

[12] Ad effectiveness is a great concern for marketers with advances in technologies; see for example Ghosh and Stock (2010) for a case related to television with the digital video recorder.

effective than generic ads. Goldfarb and Tucker (2015) examine how the memorability of banner advertising changed with the introduction of new standard formats. They find evidence that for most ads, ad effectiveness falls as the use of standard formats rises. Finally, Andrews et al. (2016) investigate how hyper-contextual targeting with physical crowdedness, i.e. the degree of population density per unit area, may affect consumer response to mobile ads. Based on a sample of mobile phone users that mobile operators can target in subway trains, they find that commuters in crowded subway trains are about twice as likely to respond to a mobile offer by making a purchase vis-à-vis those in non-crowded trains.

Our paper particularly contributes to this literature by examining a further dimension of ad effectiveness: ad viewability. Viewability is a crucial component of ad effectiveness as an ad that is not seen or only partially does not have any chance to reach consumers. In brief, the higher the ad viewability, the higher consumer attention and ad recall. In the aforementioned studies, banner ads and other ad formats are supposed to be always viewed by Internet users. This implicit assumption is contradicted by numerous studies documented in Section 2.2.

2.4 A Theoretical Model of Online Ad Viewability

2.4.1 Model Setup

We build a model picturing a media market in which Internet users visit a website to consume content and see ads. Our model involves three types of agents: Internet users, advertisers, and a monopoly publisher.

The publisher offers an exogenous editorial content of quality q, and manages its website to attract Internet users on one side and advertisers on the other side. This is therefore a classical two-sided market in which two groups of agents interact through a platform. We assume, for the sake of simplicity, that the publisher is only financed by advertising (and not by subscription)[13]. Advertisers pay therefore the publisher to display ads and attract consumers that are interested in their products. Advertisers are concerned about paying for ads that are *seen* by users and not just *served*, as non-viewable ads do not promote the visibility of the products company.

A novelty of the model is that ad viewability is a publisher's decision variable. Indeed, as

[13] We do not consider subscription in this paper. However such set-up is discussed in the conclusion

reported in Section 2.2, the location of ads within webpages and a website considerably affects the viewability of ads, and a publisher, such as the recent case of the Guardian, can design its website so as to increase the viewability of ads (or not).

We analyze two situations. In the first situation, advertisers do not have a technology to measure the viewability of ads on the publisher website. In this case, the publisher cannot *commit* to a specific level of ad viewability, in the sense that advertisers have no technology to verify the publisher decision. Therefore, they just anticipate a global level of ad viewability. In the second situation, advertisers have a technology to measure ad viewability. Hence, the publisher can *commit* to a specific level of ad viewability that can be easily verified by advertisers. We can therefore compare the impact of the adoption of a technology to measure ad viewability on the demands and profits of Internet users, the publisher and advertisers.

Before analyzing the two situations, we describe in more details the preferences and objectives of Internet users, advertisers and the publisher, as well as the timing of the game.

The publisher. To maximize its profits, the publisher chooses the number of ads a to be displayed on the website but also the level of ad viewability b. We assume $a \in [0, \overline{a}]$, $\overline{a}$ the highest number of possible ads displayed on the website, and $b \in [\underline{b}, \overline{b}]$ with $0 \leq \underline{b} < \overline{b} \leq 1$. We assume here that websites are constrained regarding their choices in advertising, as the lenghts of articles or the network capabitilites limit to the number of ads one website can set.

The profit function of the publisher takes the following form:

$$\Pi = R(ba)N(ba), \tag{2.1}$$

where $R(ba)$ is the revenue function for setting a level a of ads with viewability b on the website and $N(ba)$ is the number of Internet users choosing to visit the publisher website when they see the ads. Following Anderson and Gans (2011) and Anderson and Jullien (2016), we state that $R(ba) = r(ba)ba$, where $r(ab)b$ denotes the willingness-to-pay of the ath advertiser for a number a of ads with viewability level b, and assume that the corresponding marginal revenue $r(ba)$ is decreasing in ba in standard fashion. Therefore, the willingness-to-pay of advertisers is concave in the level of ad viewability b. Indeed, advertisers may want to pay for viewable ads until a certain point where it would strengthen too much the competition for attention.

As discussed in Section 2.2, without viewability technology, advertisers have no idea about the level of ad viewability. In this case, we assume that all advertisers have rational expectations about ad viewability on the publisher website. Hence, we denote by b_e the estimated level of ad viewability, with $\underline{b} \leq b_e \leq \bar{b}$. However, we do assume that Internet users can always observe the true level of viewability b. We develop on this assumption later.

Considering the different cases, the profit functions of the publisher with and without viewability technology are:

$$\Pi = \begin{cases} R(b_e a)N(ba) \text{ without viewability technology} \\ \\ R(ba)N(ba) \text{ with viewability technology.} \end{cases} \tag{2.2}$$

Internet users. Internet users choose to visit (or not) the publisher website. They are uniformly distributed over a unit length with respect to their preferences θ for the editorial content of the publisher. Internet users earn a utility of 0 from not visiting the website. If they choose to visit the website, their utility can be defined as follows:

$$U = \theta q - \gamma \kappa(ba), \tag{2.3}$$

with q the quality of editorial content and $\gamma\kappa(ab)$ the nuisance cost of ads. We first assume that the utility of Internet users increases with the quality of editorial content q, and second that users perceive ads as a nuisance, driving the nuisance parameter to be positive $\gamma > 0$, and the advertising perception function $\kappa(ab)$ to be convex in ba ($\kappa'(ba) > 0$ and $\kappa''(ba) \geq 0$ for any $ba > 0$).[14]

No matter the situation (with or without viewability technology), we assume that Internet users can always observe the true level of viewability b. This is due to the fact that Internet users actually visit websites while advertisers buy campaigns without being able to visit all website pages.

Timing. The timing of the game is in two stages:

1. The monopoly publisher designs its website and chooses the level of ad viewability as well as the number of ads to be displayed on its website.

[14] We use a standard advertising disutility paradigm adopted in papers such as (Zhang and Sarvary, 2015; Dukes and Gal-Or, 2003).

2. If advertisers have a viewability technology, they pay the website according to the viewability level of ads b. If they have no technology, they pay the website according to their estimation of the viewability level b_e.[15]
3. Internet users choose to visit or not the publisher website(s) and see ads.

Step 1 and 2 are solved together. In the following, we look for the subgame perfect equilibrium, and solve the game by backward induction.

2.4.2 Equilibria

The publisher displays ads and offers an online content to attract advertisers and Internet users.

Stage 3

At stage 3, Internet users visit (or not) the publisher website. No matter $k \in \{v, nv\}$, the demand of Internet users is:

$$N(ba) = 1 - \frac{\gamma\kappa(ba)}{q} \tag{2.4}$$

As $\kappa(ba)$ increases in ba, eq. (2.4) shows that the number of Internet users visiting the publisher website decreases with the nuisance cost of ads $\gamma\kappa(ba)$.

Stage 1 and 2

The publisher maximizes its profits by choosing its advertising intensity. We first focus on the case without viewability. As advertisers cannot observe the true level of viewability and expect a level of viewability b_e, we solve the game using Bayesian perfect Nash equilibrium. We denote the solution of the such game with the subscript nv.

Inserting Eq. (2.4) into Eq. (2.2), the publisher profit function at the equilibrium of stage 1 without viewability technology is $\Pi = R(b_e a)(1 - \frac{\gamma\kappa(ba)}{q})$. As the publisher is constrained on the level of advertising a ($a \in [0, \overline{a}]$) and the level of viewability b ($\underline{b} \leq b \leq \overline{b}$), he solves the following program:

[15] In the advertising industry, advertisers can ask for compensation when they measure a lower advertising viewability that previously declared by the publisher when the ad is sold. Hence they pay advertising before ads are displayed to users.

$$\begin{aligned}
&\max_{b,a} \Pi = R(b_e a)(1 - \frac{\gamma\kappa(ba)}{q}), \\
&\text{subject to } 0 \leq a \leq \bar{a} \text{ and } \underline{b} \leq b \leq \bar{b}
\end{aligned} \tag{2.5}$$

Imposing rational expectation for advertisers and solving the program of Eq. 2.5 gives us the following solution (b^*_{nv}, a^*_{nv}) such that:

$$(b^*_{nv}, a^*_{nv}) = \begin{cases} (\underline{b}, \widetilde{a}) \text{ if } \gamma > \frac{qR'_a(\bar{a}\underline{b})}{\kappa'_a(\bar{a}\underline{b})R(\bar{a}\underline{b}) + \kappa(\bar{a}\underline{b})R'_a(\bar{a}\underline{b})}, \\ \\ (\underline{b}, \bar{a}) \text{ if } \gamma \leq \frac{qR'_a(\bar{a}\underline{b})}{\kappa'_a(\bar{a}\underline{b})R(\bar{a}\underline{b}) + \kappa(\bar{a}\underline{b})R'_a(\bar{a}\underline{b})} \end{cases} \text{[16]} \tag{2.6}$$

where $\widetilde{a_{nv}}$ satisfies[17]:

$$\frac{R'_a(b_e\widetilde{a})}{R(b_e\widetilde{a})} = \frac{\gamma\kappa'_a(b\widetilde{a})}{q - \gamma\kappa(b\widetilde{a})}\text{[18]}. \tag{2.7}$$

Advertisers pay the website where ads are displayed. If there is a viewability technology, advertisers always pay the price indexed on the viewability level of ads, as they can verify it using the technology. If no viewability technology is available, advertisers always pay the price of an ad by considering its estimated level of viewability b_e and not b_{nv}, as they have no means of knowing the actual level of ad viewability. In this scenario, the publisher cannot commit to a specific level of viewability technology. Hence, he has a dominating strategy that would be to set $b^*_{nv} = \underline{b}$. We find therefore that the publisher sets the level of ad viewability at its minimum.

Proposition 1: When there is no viewability technology, the publisher chooses the lowest level of viewability $b^*_{nv} = \underline{b}$ while advertisers expect the lowest level of viewability $b_e = \underline{b}$.

Proof. Since a higher level of ad viewability decreases the demand of users, and since advertisers have no possibilities to determine the actual level of ad viewability, it is optimal for the publisher

[16] Proof can be found in Appendix.

[17] This can be written using elasticities such that $\epsilon_R(a^*_{nv}, b_e) = \epsilon_N(a^*_{nv}, b)$. Following Anderson and Gans (2011), we define $\epsilon_R = \frac{\partial R}{\partial a}\frac{a}{R}$ as the per consumer elasticity of advertising revenue and $\epsilon_N = \frac{\partial N}{\partial a}\frac{a}{N}$ as the elasticity to advertising related to consumer demand. The optimal local advertising intensity is characterized by the equality between advertisers and Internet users elasticities (respectively ϵ_R and ϵ_N) with respect to a. We find that for a high content quality q or a low level of advertising nuisance $\gamma\kappa(.)$, the optimal level of advertisers may be above $\bar{a}$, driving the publisher to set $a^* = \bar{a}$. Hence, the maximum optimal advertising intensity when there is no viewability technology is $\underline{b}\bar{a}$.

[18] $\kappa'_a = \frac{\partial\kappa}{\partial a}$ and $R'_a = \frac{\partial R}{\partial a}$.

to set the viewability level at its minimum. Such lack of commitment power also impacts advertisers expectations. as we assume rational expectations, they anticipate that the dominating strategy of the publisher would be to set $b^*_{nv} = \underline{b}$ which drives their expectation to be $b_e = \underline{b}$. $\square$

Proposition 1 states that the publisher sets the level of ad viewability at its minimum. Such situation underlines how the lack of comitment power on the publisher side degrades the viewability level of ads on a website. Applying the same reasoning, the publisher maximizes its profit with respect to b and a under the constraint $0 < A \leq \bar{a}\bar{b}$. However, we can simplify this two variables maximization problem to a one variable in considering $A = ba$. Hence, we have the folowing program:

$$\begin{aligned} &\max_{A} \Pi = R(A)(1 - \frac{\gamma\kappa(A)}{q}), \\ &\text{subject to } 0 \leq A \leq \bar{a}\bar{b} \end{aligned} \tag{2.8}$$

Solving the program of Eq. 2.8 gives us the optimal advertising level set by the publisher $\widetilde{A}$ such that:

$$\frac{R'_A(\widetilde{A})}{R(\widetilde{A})} = \frac{\gamma\kappa'_A(\widetilde{A})}{q - \gamma\kappa(\widetilde{A})},^{19} \tag{2.9}$$

In this case, the level of viewable ads A^*_v chosen by the publisher is uniquely defined and bounded by 0 and $\bar{b}\bar{a}$.

$$A^*_v = \begin{cases} \widetilde{A} \text{ if } \gamma > \frac{qR'_A(\bar{b}\bar{a})}{\kappa'_A(\bar{b}\bar{a})R(\bar{b}\bar{a}) + \kappa(\bar{b}\bar{a})R'_A(\bar{b}\bar{a})}, \\ \bar{a}\bar{b} \text{ if } \gamma < \frac{qR'_A(\bar{b}\bar{a})}{\kappa'_A(\bar{b}\bar{a})R(\bar{b}\bar{a}) + \kappa(\bar{b}\bar{a})R'_A(\bar{b}\bar{a})}. \end{cases} \tag{2.10}$$

When γ is large, the publisher faces an interior solution problem. He can set any b between $\underline{b}$ and $\bar{b}$ and choose the level of ads a accordingly (or vice-versa). Similarly to the case without viewability, when γ is small, the publisher will face a corner solution problem and will set $A^*_v = \bar{b}\bar{a}$. Figure 2.1 shows how advertising level a and viewability level b behave at the interior equilibrium, i.e when $A^*_v = \widetilde{A}$.

[19] $\kappa'_A = \frac{\partial\kappa}{\partial A}$ and $R'_A = \frac{\partial R}{\partial A}$. Proof can be found in Appendix.

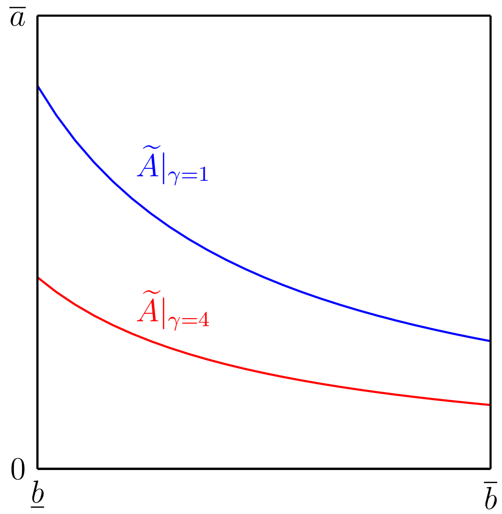

Figure 2.1: Interior level of advertising with viewability technology when $r(a) = 1 - a$, $\kappa(a) = a$, and $q = 2$

We compare the optimal advertising intensity chosen by the publisher when there is or not a viewability technology. We find that when there is a viewability technology, the advertising intensity is equal or greater than the one when there is no viewability technology $b_v^* a_v^* \geq b_{nv}^* a_{nv}^*$.

Proposition 2: Let $\dot{\gamma} = \frac{qR'_{a_{nv}}(\overline{a}\underline{b})}{\kappa'_{a_{nv}}(\overline{a}\underline{b})R(\overline{a}\underline{b})+\kappa(\overline{a}\underline{b})R'_{a_{nv}}(\overline{a}\underline{b})}$ and $A_k = b_k^* a_k^*$ with $k \in \{v, nv\}$:

(i) $A_v^* = A_{nv}^*$ when $\gamma \geq \dot{\gamma}$

(ii) $A_v^* > A_{nv}^*$ when $\gamma < \dot{\gamma}$.

Proof. See Proof in Appendix. □

Proposition 2 states that a publisher is more likely to set the level of advertising intensity to its maximum when there is a viewability technology. This finding can be disentangled in two cases. Firstly, when the advertising nuisance γ is high, the Internet users are very sensitive to the nuisance cost of ads. Hence, the publisher has no incentive to attract a lot of advertisers as it would repel too many Internet users from visiting the publisher website, whether there is a viewability technology or not. Overall, the advertising intensity is equivalent with or without viewability: when there is no viewability, the publisher will offer a low level of ad viewability and will adapt the number of ads accordingly; when there is a viewability technology, the publisher will combine different levels of ads and viewability, as long as they match the optimal level of advertising intensity.

Secondly, when the advertising nuisance γ is low, the demand of Internet users is less sensitive (elastic) to advertising. Hence, the publisher attracts a lot of Internet users even if it sets a high advertising intensity. When there is no viewability technology, the publisher is bounded by technical limitations ($a^*_{nv} < \overline{a}$) and the fact that he cannot commit to a specific level of viewability technology (leading to $b^*_{nv} = b_e = \underline{b}$ see Proposition 1). Therefore, the optimal advertising intensity cannot exceed $\overline{a}\underline{b}$ when there is no viewability technology. When there is a viewability technology, the publisher is simply bounded by a technical limitation of advertising intensity as the number of ads a^*_v cannot exceed $\overline{a}$ and the viewability level b_v is upper bounded by $\overline{b}$. Hence, the optimal advertising intensity cannot exceed $\overline{a}\overline{b}$ when there is viewability technology. As $\overline{a}\underline{b} < \overline{a}\overline{b}$, the publisher may be able to practice its optimal level of advertising intensity (or the maximum upper limit) when there is viewability, whereas it is bounded by $\underline{a}\overline{b}$ when there is no viewability technology.

We therefore find the following publisher profits at the equilibrium:

$$\Pi^*_{nv} = \begin{cases} R(\underline{b}\widetilde{a})(1 - \frac{\gamma\kappa(\underline{b}\widetilde{a})}{q}) \text{ if } \gamma > \frac{qR'_{a_{nv}}(\underline{b}\overline{a})}{\kappa'_{a_{nv}}(\underline{b}\overline{a})R(\underline{b}\overline{a})+\kappa(\underline{b}\overline{a})R'_{a_{nv}}(\underline{b}\overline{a})}, \\ R(\underline{b}\overline{a})(1 - \frac{\gamma\kappa(\underline{b}\overline{a})}{q}) \text{ if } \gamma \leq \frac{qR'_{a_{nv}}(\underline{b}\overline{a})}{\kappa'_{a_{nv}}(\underline{b}\overline{a})R(\underline{b}\overline{a})+\kappa(\underline{b}\overline{a})R'_{a_{nv}}(\underline{b}\overline{a})}, \end{cases} \tag{2.11}$$

$$\Pi^*_{v} = \begin{cases} R(\widetilde{A})(1 - \frac{\gamma\kappa(\widetilde{A})}{q}) \text{ if } \gamma > \frac{qR'_{A_v}(\overline{b}\overline{a})}{\kappa'_{b_v a_v}(\overline{b}\overline{a})R(\overline{b}\overline{a})+\kappa(\overline{b}\overline{a})R'_{A_v}(\overline{b}\overline{a})}, \\ R(\overline{b}\overline{a})(1 - \frac{\gamma\kappa(\overline{b}\overline{a})}{q}) \text{ if } \gamma \leq \frac{qR'_{A_v}(\overline{b}\overline{a})}{\kappa'_{b_v a_v}(\overline{b}\overline{a})R(\overline{b}\overline{a})+\kappa(\overline{b}\overline{a})R'_{A_v}(\overline{b}\overline{a})}, \end{cases} \tag{2.12}$$

When visiting the website, Internet users consume a content of quality q and see ads incurring a cost γ. However, the website may have different incentives to increase quality q whether there is a viewability technology or not on the market. Using the envelop theorem, we assess how the profit change with respect to the content quality q which lead to the following proposition:

Proposition 3: Let $\dot{\gamma} = \frac{qR'_{a_{nv}}(\overline{a}\underline{b})}{\kappa'_{a_{nv}}(\overline{a}\underline{b})R(\overline{a}\underline{b})+\kappa(\overline{a}\underline{b})R'_{a_{nv}}(\overline{a}\underline{b})}$:

(i) $\frac{\partial\Pi_v}{\partial q} = \frac{\partial\Pi_{nv}}{\partial q} > 0$ when $\gamma \geq \dot{\gamma}$

(ii) $\frac{\partial\Pi_v}{\partial q} > \frac{\partial\Pi_{nv}}{\partial q} > 0$ when $\gamma < \dot{\gamma}$.

Proof. Using the envelop theorem we have $\frac{\partial \Pi_i}{\partial q} = \frac{R(A_k^*)\gamma\kappa(A_k^*)}{q^2}$ with $A_k = b_k^* a_k^*$ and $k \in \{v, nv\}$. Hence, $\frac{\partial \Pi_v}{\partial q} > \frac{\partial \Pi_{nv}}{\partial q}$ only if $R(A_v^*)\kappa(A_v^*) > R(A_{nv}^*)\gamma\kappa(A_{nv}^*)$ which is always the case as $R(A_v^*) > R(A_{nv}^*)$ and $\kappa(A_v^*) > \kappa(A_{nv}^*)$ (from Proposition 2). □

Proposition 3 underlines that a publisher has (weakly) higher incentives of investing in content quality q when there is a viewability technology available. We notice that profit is increasing with content quality no matter the presence of viewability technology. This is due to the fact that a higher content quality q compensates the negative effect the nuisance cost of ads can exercise on users demand. Hence, investing in content quality q is even more efficient when the disutility from ads is high, which may happen in presence of viewability technology as the publisher sets a (weakly) higher level of advertising (see Proposition 2).

2.5 Welfare Analysis of Ad Viewability

We determine in this section whether the introduction of a viewability technology is profitable for the market, i.e. for Internet users, advertisers and the publisher. We calculate and compare the total welfare with and without viewability technology, denoted respectively by W_v^* and W_{nv}^*. The total welfare is the sum of the surplus of Internet users (S_v^{u*} and S_{nv}^{u*}), the surplus of advertisers (S_v^{a*} and S_{nv}^{a*}), and the profits of the publisher (Π_v^* and Π_{nv}^*). We consider S_i^{u*} and S_i^{a*}, where $i \in \{v, nv\}$ and $S_i^{u*} = \int_{\frac{\gamma b_i^* a_i^*}{q}}^{1} \theta q - \gamma b_i^* a_i^* \mathrm{d}\theta$ and $S_i^{a*} = N^* \int_0^{b_i^* a_i^*} (r(ba) - r(b_i^* a_i^*))\mathrm{d}a$. To keep the analysis as simple as possible, we assume $R(ba) = r(ba)ba = (1 - ba)ba$ and $\kappa(ba) = ba$,

We provide welfare calculation in Appendix, and focus in the sequel on total welfare. We find that total welfare is not the same depending on the presence of viewability technology.

Proposition 4: Let $R(ba) = r(ba)ba = (1 - ba)ba$ and $\kappa(ba) = ba$, which drives $\dot{\gamma} = \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}$ and $q_w \equiv \frac{\gamma(\underline{b}\overline{a}^2 + (b_v^* a_v^* + \gamma - 2)(\underline{b}\overline{a} + b_v^* a_v^*))}{2(\gamma-1) + \underline{b}\overline{a} + b_v^* a_v^*}$:

(i) $W_{nv}^* = W_v^*$ when $\gamma > \dot{\gamma}$,

(ii) $W_v^* > W_{nv}^*$ when $\gamma \leq \dot{\gamma}$ and $q > q_w$,

Proof. See Proof in Appendix. □

Proposition 4, illustrated in Figure 4.2, shows that the total welfare is greater with viewability technology when the quality of the content is high ($q > q_w$) and the nuisance cost of ads is low (as $\frac{\partial q_w}{\partial \gamma} > 0$).[20] To understand this result, we propose to analyze how the content quality (q) and the nuisance cost of ads (γ) impact the change in total welfare. To begin with, let us denote by $\Delta W = W_v^* - W_{nv}^*$, the change in total welfare, which is simply the difference between the total welfare with and without viewability technology. Likewise, let us denote by $\Delta \mathrm{S}^u$ the change in Internet user surplus (with $\Delta \mathrm{S}^u = \mathrm{S}_v^{u*} - \mathrm{S}_{nv}^{u}{}^*$), by $\Delta \mathrm{S}^a$ the change in advertisers' surplus (with $\Delta \mathrm{S}^{a*} = \mathrm{S}_v^{a*} - \mathrm{S}_{nv}^{a}{}^*$), and by $\Delta \Pi$ the change in publisher profits (with $\Delta \Pi = \Pi_v^* - \Pi_{nv}^*$).[21]

To summarize, with viewability technology, a large nuisance cost γ allows the publisher to practice its optimal advertising intensity ($\gamma > \dot{\gamma}$), which does not affect neither the surplus of Internet users nor the profits of the industry ($\Delta \mathrm{S}^u = \Delta \mathrm{S}^a = \Delta \Pi = 0$), leaving the total welfare unchanged ($\Delta W = 0 \equiv W_{nv}^* = W_v^*$).

However, when the advertising nuisance softens ($\gamma < \dot{\gamma}$), Internet users are worse off whereas industry profits are higher. Users' disutility from ads has therefore a mixed effect on total welfare, and the final impact depends on the content quality. More precisely, the total welfare is greater with viewability technology only if the nuisance cost of ads γ is not too high and/or the content quality q is not too low. In this latter case, the industry profits are always greater than the loss of Internet users: $\Delta W > 0 \equiv W_v^* > W_{nv}^*$. This result holds for any γ such that $\gamma < \dot{\gamma}$. This is due to the fact that the publisher sets the maximum level of advertising intensity when there is no viewability technology (i.e. $b_{nv}^* a_{nv}^* = \underline{b}\overline{a}$). In this case, a higher nuisance cost of ads γ increases the respective gains of the publisher and advertisers ($\frac{\partial \Delta S_a}{\partial \gamma}, \frac{\partial \Delta \Pi}{\partial \gamma} < 0$), while having an ambiguous effect on the loss of Internet users. Overall, when a viewability technology is adopted, a high content quality q and/or a low nuisance cost γ drive the loss of Internet user surplus to be lower than the industry profits.

[20] See the calculation in Appendix.

[21] We detail the calculation of $\Delta \mathrm{S}^u$, $\Delta \mathrm{S}^a$ and $\Delta \Pi$ in Appendix.

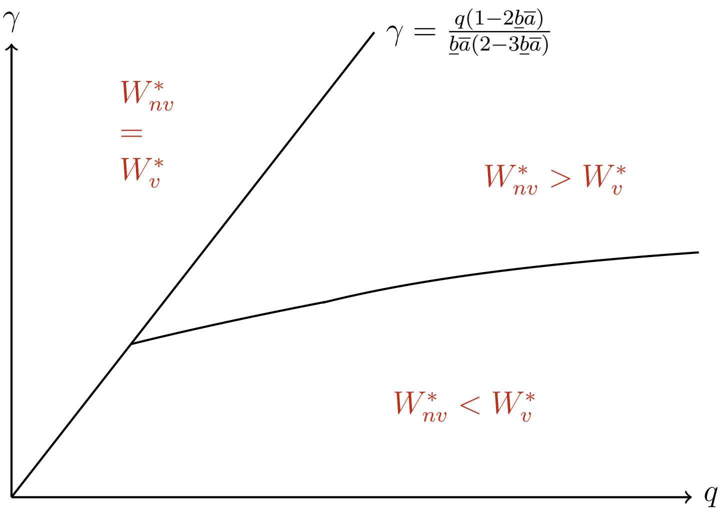

Figure 2.2: $W^*_{nv} - W^*_v$ as a Function of Content Quality (q) and Nuisance Cost of Ads (γ)

2.6 Extensions

2.6.1 Competition

In this extension we assume a "bottleneck competition" between two publishers we denote by 1 and 2. In this set-up, Internet users visits only one website (they singlehome) while advertisers contract on both websites independantly (they multihome). We model such competition in setting each platform at the extreme of a hotelling line of lenght 1 and assume that users are uniformly distributed according to their preferences x such that:

$$U_i = q - \gamma\kappa(ba) - \tau|x - \ell_i|, \tag{2.13}$$

Where τ is the transportation cost (which proxies the competition intensity between platforms), ℓ_i is the location of platform $i \in \{1, 2\}$ and x the preference of the user on the hotelling line. We assume platform 1 is located at $\ell_1 = 0$ while platform 2 is located at $\ell_2 = 1$. Both publishers are symmetric with respect to their maximum number of ads and choose their level of advertising $a_i \in [0, \overline{a}]$. Conversely, we assume that publishers are asymetric with respect to their maximum level of viewability (i.e viewability capability) and that they choose their viewability level $b_i \in [0, \overline{b_i}]$. Without loss of generality, we make the hypothesis that publisher 1 has higher viewability capabilities than publisher 2, i.e $\overline{b_1} > \overline{b_2}$.[22]

[22]This may happen because of website design or lenght of contents. We could also assume that publishers have

Assuming rational expectation for advertisers, we analyze the choice of both publishers regarding advertising and viewability, with and without viewability technology. Without technology we find that publishers optimal choice in number of ads and viewability is symmetric and satisfies the following with $i \in \{1,2\}$:

$$(b^*_{i,nv}, a^*_{i,nv}) = \begin{cases} (\underline{b}, \widetilde{a_c}) \text{ if } \gamma > \frac{\tau R'_a(\overline{a}\underline{b})}{\kappa'_a(\overline{a}\underline{b})R(\overline{a}\underline{b})}, \\ (\underline{b}, \overline{a}) \text{ if } \gamma \leq \frac{\tau R'_a(\overline{a}\underline{b})}{\kappa'_a(\overline{a}\underline{b})R(\overline{a}\underline{b})} \end{cases} \tag{2.14}$$

where $\widetilde{a_c}$ satisfies:

$$\frac{R'_a(b_e\widetilde{a_c})}{R(b_e\widetilde{a_c})} = \frac{\gamma\kappa'_a(b\widetilde{a_c})}{\tau}\text{[23]}. \tag{2.15}$$

Result in Eq. 2.16 is reminiscent of the monopoly case: viewability level is set to the lowest ($\underline{b}$) as platforms cannot commit to any higher level, while the number of ads can either be a local ($\widetilde{a_c}$) or a corner solution ($\overline{a}$) depending on the nuisance cost of ads γ.

Introducing a viewability technology when there is asymetric competition produces different results. As in the monopoly case, publisher $i \in \{1,2\}$ maximizes its profit with respect to $A_i = b_i a_i$. Solving the corresponding Lagrangean program, gives us the following solution

$$A^*_{i,v} = \begin{cases} \widetilde{A_i} \text{ if } \gamma > \frac{\tau R'_a(\overline{a}\underline{b})}{\kappa'_a(\overline{a}\underline{b})R(\overline{a}\underline{b})}, \\ \overline{a}\overline{b_i} \text{ if } \gamma \leq \frac{\tau R'_a(\overline{a}\underline{b})}{\kappa'_a(\overline{a}\underline{b})R(\overline{a}\underline{b})} \end{cases} \tag{2.16}$$

where $\widetilde{A_i}$ satisfies:

$$\frac{R'_{A_i}(\widetilde{A_i})}{R(\widetilde{A_c})} = \frac{\gamma\kappa'_{A_i}(\widetilde{A_i})}{\tau - \gamma(\kappa(\widetilde{A_i}) - \kappa(\widetilde{A_j}))},\text{[24]} \tag{2.17}$$

different advertising capabilities such that $\overline{a}$ is different for each platform. However, only assuming asymetry regarding viewability capabilities brings the same results.

[23]Where $\kappa'_a = \frac{\partial\kappa}{\partial a}$ and $R'_a = \frac{\partial R}{\partial a}$. The equilibrium must be symmetric. From Anderson and Gans (2011), if $A_1 > A_2$, publisher 1 serve ads to less Internet users than $\frac{1}{2}$. In this case, the Internet users elasticity is higher for publisher 1 than the one for publisher 2. This implies that the advertisers' elasticity would be larger for publisher 1 than the one for publisher 2 . However, since advertisers elasticity is decreasing, it means that $A_1 < A_2$, which contradicts the former.

[24]$\kappa' = \frac{\partial\kappa}{\partial A_i}$ and $R'_{A_i} = \frac{\partial R}{\partial A_i}$. Same argument of (2.15) apllies. However, as publishers have different maximum viewability level, we may have assymetric ad level.

As publishers are asymetric regarding their viewability capabilities, we have the following lemma:

Lemma 1: Let $\gamma_2 = \frac{\tau R_a'(\bar{a}\overline{b_2})}{\kappa_a'(\bar{a}\overline{b_2})R(\bar{a}\overline{b_2})}$ and $A^*_{i,v} = b^*_{i,v}a^*_{i,v}$ with $i \in \{1,2\}$:

(i) If $\gamma > \gamma_2$, $A^*_{1,v} = A^*_{2,v}$

(ii) If $\gamma < \gamma_2$, $A^*_{1,v} > A^*_{2,v}$

Proof. See Proof in Appendix. □

Lemma 1 is explained by the asymetry between platforms regarding their viewability capabilities, i.e $\overline{b_1} > \overline{b_2}$. Indeed, when nuisance cost of ads is high ($\gamma > \gamma_2$), both publishers choose the same advertising intensity $A^*_{1,v} = A^*_{2,v} = \widetilde{A_i}$. However, when nuisance cost of ads is low, competing publishers want to increase their advertising intensity. If γ is low such that $\gamma < \gamma_2$, publisher 2 will increase its advertising intensity and choose a corner solution defined by $\overline{b_2}\bar{a}$, while publisher 1 can set a higher level of advertising due to its higher viewability capability.

Such result allows us to compare advertising intensity whether there is or not viewability technology.

Proposition 5: Let $\gamma_2 = \frac{\tau R_a'(\bar{a}\overline{b_2})}{\kappa_a'(\bar{a}\overline{b_2})R(\bar{a}\overline{b_2})}$, $\underline{\gamma} = \frac{\tau R_a'(\bar{a}\underline{b})}{\kappa_a'(\bar{a}\underline{b})R(\bar{a}\underline{b})}$ and $A_{i,k} = b^*_{i,k}a^*_{i,k}$ with $i \in \{1,2\}$ and $k \in \{v, nv\}$:

(i) If $\gamma > \underline{\gamma}$, $A^*_{1,v} = A^*_{2,v} = A^*_{1,nv} = A^*_{2,nv}$

(ii) If $\gamma_2 < \gamma < \underline{\gamma}$, $A^*_{1,v} = A_{2,v} > A^*_{1,nv} = A^*_{2,nv}$

(ii) If $\gamma_2 > \gamma$, $A^*_{1,v} > A^*_{2,v} > A^*_{1,nv} = A^*_{2,nv}$

Proof. See Proof in Appendix. □

Proposition 5 shows that introducing a viewability technology has important effects on advertising competition. Firstly, results from proposition 2 still holds: a viewability technology weakly increase the level of advertising on both platforms. This is due to advertisers not expecting the lowest level of viewability when there is technology. Secondly, it also underlines that competition intensity τ can change the effect of introducing viewability technology. Indeed, in case of weak competition between platforms (high τ), both publishers have incentives to increase their advertising levels. In this case, introducing a viewability technology has a stronger effect as it incures a higher increase in ad level. Thirdly, Proposition 5 underlines how introducing such technology allows the publisher with a higher viewability capacity (publisher 1) to set a higher advertising level than publisher 2. Such case is interesting. Without viewability technology, both publishers set the same level of advertising, due to rational expectation from advertisers. However, introducing a viewability technology allows them to commit to a specific level of viewability and increase their advertising level. In the case of very low advertising nuisance ($\gamma_2 > \gamma$), viewability technology drives publisher 2 to set its advertising level to is maximum. As publisher 1 has a higher maximum level of viewability than publisher 2, he can set a higher level of advertising.

2.6.2 Ad-viewability and ad-blockers

Up to now, we assumed that Internet users cannot avoid seeing ads. However, more and more people are using ad-blockers to skip viewable ads. Ad viewability and ad-blockers are therefore intimately related.[25] A recent survey conducted by the IAB with C3Research finds for example that 26% of desktop users block ads online (IAB, Who Blocks Ads, Why, and How to Win Them Back, 2016). And their impact is significant. Research firm Ovum estimates that publishers lost $24 billion in revenue globally in 2015 due to ad blocking (Wall Street Journal, "New York Times Readies Ad-Free Digital Subscription Model," June 20, 2016).

In this extension, we introduce the possibility for consumers to block ads so as to get rid of the nuisance costs of viewable ads when visiting a monopoly publisher.[26] Installing an ad-blocker is costly for users (c), as they need to search for and to install the software.[27] We assume in the sequel that users choose to install an ad-blocker as soon as its cost c is lower or equal than the nuisance cost of viewable ads $\gamma * b * a$ when visiting the publisher website. This implies that they observe the advertising intensity before deciding to use an ad-blocker or not. By definition,

[25] An article entitled "Solving for Viewability Might Be a Reason People are Ad Blocking" published in Digiday UK on November 11, 2016, discusses the interactions between ad-blockers and ad-viewability.

[26] Ad-avoidance technologies have been largely studied (Anderson and Gans, 2011; Johnson, 2013). In line with these studies, we consider that ad-avoidance technologies involve consumers reducing the negative impact of ads. However, we only consider passive blocking from Internet users, as they do not expect a number of ads on the website before visiting it.

[27] We can also interpret c as the minimum user experience threshold required by the ad-blocker to display an ad on a website even if an ad-blocker is installed. For example, the software *Ad-blockplus* allows publisher to display ads that comply with specific criteria to be shown to users.

the installation of an ad-blocker prevents ad servers to serve ads, which delivers zero revenue for the publisher. To make positive profits, the publisher has therefore to lower the number of viewable ads to discourage people from installing ad-blockers. Consequently, ad-blockers are introduced in the model as a constraint on the maximization problem of the publisher when it sets the number of viewable ads in stage 1 of the previous game.

We study the introduction of ad-blockers and analyze its impact on total welfare. We assume that when users visit the website, they have the choice between seeing ads and incuring an advertising cost of $\gamma\kappa(ab)$ or blocking ads and facing a cost c. The publisher earns regular profit when users decide to watch ads and zero profit when they block ads. Depending on the introduction of viewability, the maximization program of the publisher is the same as in Eqs. (2.5) and (2.8), except that he is now constrained on its level of advertising $\gamma\kappa(ab) \leq c$ to make positive profits.

We know that following Eq. (2.6), without technology, the expression $\gamma\kappa(ba)$ can take on two values, $\gamma\kappa(b^*a^*)$ or $\gamma\kappa(\underline{b}\overline{a})$. Solving the langrangean programs of Eq. (2.5) with an additional constraint ($\gamma\kappa(ab) \leq c$), we show that two cases can be considered. Firstly, if $c > \min\left(\gamma\kappa(b^*a^*), \gamma\kappa(\underline{b}\overline{a})\right)$, the cost to install an ad-blocker is higher than the cost of seeing viewable ads for Internet users. Hence the publisher is not constrained and can set the optimal number of viewable ads without considering ad-blockers. Secondly, if $c \leq \min\left(\gamma\kappa(b^*a^*), \gamma\kappa(\underline{b}\overline{a})\right)$, the cost to install an ad-blocker is lower than the nuisance cost of viewable ads for users. Consequently, the publisher is forced to lower the number of viewable ads to prevent users from avoiding ads. To keep the analysis straightforward, we only focus on the second case (we however consider all the cases when computing total welfare).

When the publisher is constrained by ad-blockers, the optimal level of ad viewability is equivalent to the one without ad-blockers $b^*_{nv} = \underline{b}$. However, the number of ads chosen by the publisher is constrained by the presence of ad-blockers, driving the optimal advertising intensity chosen by the publisher to satify $\gamma\kappa(a^*_{nv}\underline{b}) = c$. In this case, the publisher is less constrained when the cost to use ad-blockers increases and can display a higher number of viewable ads. In the end, the publisher must however provide a high enough user experience to prevent Internet users from installing ad-blockers.

The optimal choice of the publisher with viewability technology is subject to the same constraint from ad-blockers. We solve the langreagan program of Eq. (2.8) but with an additional constraint ($\gamma\kappa(ba) \leq c$). Following the same reasoning as before, we conduct the analysis when $c \leq \min\left(\gamma\kappa(b_v^* a_v^*), \gamma\kappa(\overline{ba})\right)$. We find that the publisher in presence of a viewability technology and ad-blockers must limit the number of viewable ads to encourage Internet users to visit the website in satisfying $\gamma\kappa(b_v^* a_v^*) = c$.

When the publisher is constrained in both situations (i.e. when $\gamma\kappa(b_{nv}^* a_{nv}^*) > c$), he will naturally set the same advertising level, hence driving $b_v^* a_v^* = b_{nv}^* a_{nv}^*$. The publisher therefore internalizes the cost of installing the software and improves the user experience by lowering the number of viewable ads. However, this does not happen all the time, and we find that ad-blockers are more constraining for the publisher with viewability technology.

Corollary 1: The publisher is more likely to be constrained by an ad-blocker when there is viewability technology as $\kappa(b_v a_v{}^*) \geq \kappa(b_{nv}^* a_{nv}^*)$.

Proof. This directly stems from Proposition 2 as $b_v^* a_v^* \geq b_{nv}^* a_{nv}^*$ □

Corollary 1 underlines that the publisher can be constrained by adblockers only if viewability technology is on the market. Hence, ad-blockers should have an heterogeneous effect on users surplus and market profits, depending on the presence of viewability technology.

2.7 Conclusion

Ad technologies provide new opportunities to reach consumers and improve ad effectiveness. In this paper, we studied one dimension of ad effectiveness: ad viewability. The latter offers to advertisers a greater chance to determine whether ads are seen by Internet users. In this respect, ad viewability introduces more transparency between the publisher and advertisers in a context of serious doubts on digital ads (Wall Street Journal, "Doubts Rise on Digital Ads," September 24, 2016),[28] but at the same time, puts pressure on publishers to enhance their viewability performance.

Following this idea, we studied in this paper how the introduction of ad viewability changes the economics of online advertising. We basically show that ad viewability affects the way a

[28] Facebook admitted to have overestimated by up to 80% the average time people spent watching video ads on its platform. This story is not unique. Twitter (Business Insider France, "Twitter's Video Ad Metric Inflation Came at a Terrible Time," December 27, 2016) or Dentsu also acknowledged numerous cases of overcharging, amounting to at least \$2.3m (Financial Times, "Ad Scandal Puts Dentsu's Credibility on the Line," September 27, 2016).

publisher prices ads, which in turn affects its profits, the demand of advertisers, and user experience.

In a context without ad viewability, the optimal number of viewable ads displayed by the publisher is always the lowest at the equilibrium for two reasons. Firstly, a low level of ad viewability preserves user experience. Secondly, advertisers purchase impressions based on their estimated level of ad viewability, and not on the actual level of ad viewability (that they do not know). This mechanism is central: as advertisers cannot verify the level of ad viewability, the publisher has an incentive to offer the lowest level of ad viewability to preserve user experience.

This situation completely changes with a viewability technology. In this case, advertisers can verify the actual level of ad viewability, which allows the publisher to be trusted. The publisher may want therefore to raise the level of ad viewability to charge higher prices and to increase its profits. However, Internet users are not always ready to accept a higher nuisance cost of ads and may leave the market.

In the end, the optimal level of ad viewability offered by the publisher depends on both the quality of the content and the nuisance cost of ads. When the content quality is low, the publisher cannot raise its advertising intensity to make higher profits and the total welfare is not enhanced with viewability technology. By contrast, when the content quality is higher, Internet users are worse off because the publisher can raise its ad intensity, inflating in turn the profits of the industry. In that case, the welfare analysis shows that the market of online advertising can be better off with viewability technology provided that the nuisance cost of ads is not too high and the content quality is not too low. This is all the more relevant as the publisher has incentives to invest more heaviliy in content quality when viewability technology is introduced.

However, to preserve their user experience, Internet users can block ads by installing ad-blockers. Depending on the cost of installing ad-blockers, the publisher is constrained as it cannot degrade user experience by increasing too much the level of ad viewability. The publisher is therefore pressurized from both sides of the market: advertisers demand more viewable ads whereas in the same time Internet users require to preserve their online experience. Extending the baseline model, we find that when the cost to use ad-blockers is lower than the nuisance cost of viewable ads, the publisher is forced to reduce the viewability of ads to account for user experience, whether or not there is a viewability technology on the market.

Finally, we model competition between publishers that are asymetric regarding their viewability capacities. Our model underlines that without viewability technology, this asymetry has no impact on the market, as advertisers' expectations foster publishers to set their viewability to the lowest lovel. However, publishers can commit to a higher level of viewability technology once viewability technology is introduced. Such technology allows the publisher with the higest viewability capacity to take advantage of such asymetry in setting a higher level of advertising intensity, which was not possible without viewability technology.

This study can be extended in several directions. Firstly, we modeled the advertising industry as a business that rewards quantity over quality, meaning that the publisher derives revenues from the number of ads sold to advertisers, regardless of their quality. This business model is of course dominant on the market. However, due to a drop in ad revenues, many publishers refuse to add more and more ads to offset their losses and preserve user experience. They are therefore pushing advertisers to promote new ad formats to connect with their audience in a non-intrusive way. Native ads are precisely considered as the future of marketing strategy.[29] Native ads have the look and feel of the content of a website on which they are displayed, and hence do not look like simple ads. They are supposed to have higher levels of engagement than traditional non-native ads: native ads were found to deliver a 9% higher lift in brand affinity than banners (Sharethrough, Behind How Native Ads Work, 2016). A nice extension of the model would consist of accounting for the quality of the ad format, as a better ad quality may preserve user experience (even if it is more visible).

Secondly, we also assumed in the model that the demand generated on the publisher website was always valid traffic, and that fraud did not exist. But fraud is a serious concern in the advertising industry: "The World Federation of Advertisers [...] estimates that between 10 and 30% of online advertising impressions are never seen by consumers because of fraud, and forecasts that marketers could lose as much as $50bn a year by 2025 unless they take radical action." (Financial Times, "Digital advertising: Brands versus bots", July 18, 2016.) Fraud can take many forms. Unscrupulous publishers may purchase fake web traffic to inflate the price of ads. Likewise, fraud can be generated by computer programs, or "bots", that simulate users' web browsing be-

[29] According to Enders Analysis, spending on native advertising in Europe jumped by a third in 2015 alone (eMarketer, "Native Advertising in Western Europe: Paid Content Placements Gain Fans Throughout the Region," 2016).

havior. Including this dimension in the analysis would be interesting as ad viewability is affected by fraudulent traffic.[30]

Finally, in a context of increasing programmatic sales, all users do not have the same value for advertisers. This means that publishers may manage viewability levels differently depending on the value of users.

[30] US democratic senators have called on the Federal Trade Commission to protect consumers from digital advertising fraud, including potential regulation of reform of ad exchanges (Multichannel, Democratic Senators Say Digital Ad Fraud Rampant, Nov 7, 2016).

Bibliography

Agarwal, Ashish, Kartik Hosanagar, and Michael D Smith, "Location, Location, Location: An Analysis of Profitability of Position in Online Advertising Markets," *Journal of Marketing Research*, 2011, *48* (6), 1057–1073.

Anderson, Simon and Bruno Jullien, "The advertising-financed business model in two-sided media markets," Technical Report TSE 632, Toulouse School of Economics 2016.

__ **and Jean Gabszewicz**, "Chapter 18 The Media and Advertising: A Tale of Two-Sided Markets," in Victor A. Ginsburg and David Throsby, eds., *Handbook of the Economics of Art and Culture*, Vol. 1, Elsevier, 2006, pp. 567 – 614.

Anderson, Simon P. and Joshua S. Gans, "Platform Siphoning: Ad-Avoidance and Media Content," *American Economic Journal: Microeconomics*, 2011, *3* (4), 1–34.

__ **and Stephen Coate**, "Market Provision of Broadcasting: A Welfare Analysis," *The Review of Economic Studies*, 2005, *72* (4), 947–972.

Andrews, Michelle, Xueming Luo, Zheng Fang, and Anindya Ghose, "Mobile Ad Effectiveness: Hyper-Contextual Targeting with Crowdedness," *Marketing Science*, 2016, *35* (2), 218–233.

Blake, Thomas, Chris Nosko, and Steven Tadelis, "Consumer Heterogeneity and Paid Search Effectiveness: A Large-Scale Field Experiment," *Econometrica*, 2015, *83* (1), 155–174.

Caillaud, Bernard and Bruno Jullien, "Competing Cybermediaries," *European Economic Review*, 2001, *45*, 797–808.

__ **and** __ , "Chicken and Egg: Competition among Intermediation Service Providers," *The RAND Journal of Economics*, 2003, *34* (2), 309–328.

de Cornière, Alexandre, "Search Advertising," *American Economic Journal: Microeconomics*, 2016, *8* (3), 156–188.

Dukes, Anthony and Esther Gal-Or, "Negotiations and Exclusivity Contracts for Advertising," *Marketing Science*, 2003, *22* (2), 222–245.

Evans, David S., "The Online Advertising Industry: Economics, Evolution, and Privacy," *Journal of Economic Perspectives*, September 2009, *23* (3), 37–60.

Flosi, Stephanie, Gian Fulgoni, and Andrea Vollman, "If an Advertisement Runs Online And No One Sees It, Is It Still an Ad?," *Journal of Advertising Research*, June 2013, *53* (2), 192–199.

Ghose, Anindya and Sha Yang, "An Empirical Analysis of Search Engine Advertising: Sponsored Search in Electronic Markets," *Management Science*, 2009, *55* (10), 1605–1622.

Ghosh, Bikram and Axel Stock, "Advertising Effectiveness, Digital Video Recorders, and Product Market Competition," *Marketing Science*, 2010, *29* (4), 639–649.

Goldfarb, Avi and Catherine Tucker, "Online Display Advertising: Targeting and Obtrusiveness," *Marketing Science*, 2011, *30* (3), 389–404.

_ and _ , "Standardization and the Effectiveness of Online Advertising," *Management Science*, 2015, *61* (11), 2707–2719.

Johnson, Justin P., "Targeted Advertising and Advertising Avoidance," *The RAND Journal of Economics*, 2013, *44* (1), 128–144.

Lambrecht, Anja and Catherine Tucker, "When Does Retargeting Work? Information Specificity in Online Advertising," *Journal of Marketing Research*, 2013, *50* (5), 561–576.

Manchanda, Puneet, Jean-Pierre Dubé, Khim Yong Goh, and Pradeep K. Chintagunta, "The Effect of Banner Advertising on Internet Purchasing," *Journal of Marketing Research*, 2006, *43* (1), 98–108.

Peitz, Martin and Markus Reisinger, "The Economics of Internet Media," in "Handbook of Media Economics, vol. 1A," Elsevier, 2016.

Rochet, Jean-Charles and Jean Tirole, "Cooperation among Competitors: Some Economics of Payment Card Associations," *The RAND Journal of Economics*, 2002, *33* (4), 549–570.

Rutz, Oliver J and Randolph E Bucklin, "From Generic to Branded: A Model of Spillover in Paid Search Advertising," *Journal of Marketing Research*, 2011, *48* (1), 87–102.

Sanghvi, Nikul, "Cost Per Hour. Using a Time-Based Currency for Digital Advertising," Technical Report, Analytics Consultant 2015.

Wang, Chong, Achir Kalra, Cristian Borcea, and Yi Chen, "Viewability Prediction for Online Display Ads," in "Proceedings of the 24th ACM International on Conference on Information and Knowledge Management" CIKM '15 ACM New York, NY, USA 2015, pp. 413–422.

Zhang, Kaifu and Miklos Sarvary, "Differentiation with User-Generated Content," *Management Science*, 2015, *61* (4), 898–914.

Zhang, Weinan, Ye Pan, Tianxiong Zhou, and Jun Wang, "An Empirical Study on Display Ad Impression Viewability Measurements," *arXiv:1505.05788 [cs]*, May 2015. arXiv: 1505.05788.

2.8 Appendix

Ad level computation.

Ad Level without Viewability Technology

We have the following Langreagean program to solve:

$$\max_{b,a} \mathcal{L}(b,a,\lambda_1,\lambda_2,\lambda_3,\lambda_4) = R(b_e a)(1-\frac{\gamma\kappa(ba)}{q})-\lambda_1(b-\overline{b})+\lambda_2(b-\underline{b})-\lambda_3(a-\overline{a})+\lambda_4 a, \quad (2.18)$$

Solving the Langrangean program gives us the following Khun-Tucker conditions

$$(A) \quad R'_a(b_e a)(1-\frac{\gamma\kappa(ba)}{q}) - R(b_e a)\frac{\gamma}{q}\kappa'_a(ba) - \lambda_3 + \lambda_4 = 0$$

$$(B) \quad -R(b_e a)\frac{\gamma}{q}\kappa'_b(ba) - \lambda_1 + \lambda_2 = 0$$

$$(C) \quad b \leq \overline{b}, \lambda_1 \geq 0 \text{ and } \lambda_1(b-\overline{b}) = 0$$

$$(D) \quad b \geq \underline{b}, \lambda_2 \geq 0 \text{ and } \lambda_2(b-\underline{b}) = 0$$

$$(E) \quad a \leq \overline{a}, \lambda_3 \geq 0 \text{ and } \lambda_3(a-\overline{a}) = 0$$

$$(F) \quad a \geq 0, \lambda_4 \geq 0 \text{ and } \lambda_4 a = 0$$

Firstly, we see that $\lambda_4 = 0$ because $a = 0$ contradict condition (A) and (E). For conditions (C), (D) and (E) to be feasible at the same time, we need that $\lambda_1 > 0$ if $\lambda_2 = 0$ or $\lambda_1 = 0$ if $\lambda_2 > 0$. Hence solving (C), (D) and (E) gives us the following cases:

- Firstly, we see that $b = \overline{b}$ is never a solution. Indeed, in that case $\lambda_2 = 0$. Replacing in condition (B) gives $\lambda_1 < 0$ which is not possible. Hence, λ_1 is always equal to 0 to satisfy (C).

- $\lambda_1 = \lambda_2 = \lambda_3 = \lambda_4 = 0$. Replacing in condition (A) gives $a = \widetilde{a_{nv}}$, with $\widetilde{a_{nv}}$ satisfying $\frac{R'_a(b_e\widetilde{a_{nv}})}{R(b_e\widetilde{a_{nv}})} = \frac{\gamma\kappa'_a(\underline{b}\widetilde{a_{nv}})}{q-\gamma\kappa(\underline{b}\widetilde{a_{nv}})}$. Replacing in condition (B) gives $R(b_e a)\frac{\gamma}{q}\kappa'_b(ba) = 0$. This is only satisfied when $a = 0$, which contradict condition (A). Hence, this is never a solution.

- $b = \underline{b}$, $\lambda_1 = 0$, $a = \bar{a}$. Replacing in condition (B) gives $\lambda_2 = R(b_e a)\frac{\gamma}{q}\kappa'_b(\underline{b}a)$ which is greater than 0. Replacing in condition (A) gives $\lambda_3 = R'_a(b_e\bar{a})(1-\frac{\gamma\kappa(\underline{b}\bar{a})}{q}) - R(b_e\bar{a})\frac{\gamma}{q}\kappa'_a(\underline{b}\bar{a})$ which is greater than zero only when $R'_a(b_e\bar{a})(1-\frac{\gamma\kappa(\underline{b}\bar{a})}{q}) > R(b_e\bar{a})\frac{\gamma}{q}\kappa'_a(\underline{b}\bar{a})$. Hence, when $R'_a(b_e\bar{a})(1-\frac{\gamma\kappa(\underline{b}\bar{a})}{q}) > R(b_e\bar{a})\frac{\gamma}{q}\kappa'_a(\underline{b}\bar{a})$, this is a potential solution.

- $b = \underline{b}$, $\lambda_1 = 0$, $\lambda_3 = 0$ Replacing in condition (B) gives $\lambda_2 = R(b_e a)\frac{\gamma}{q}\kappa'_b(\underline{b}a)$ which is greater than 0. Replacing in condition (A) gives $a = \widetilde{a_{nv}}$, with $\widetilde{a_{nv}}$ satisfying $\frac{R'_a(b_e\widetilde{a_{nv}})}{R(b_e\widetilde{a_{nv}})} = \frac{\gamma\kappa'_a(\underline{b}\widetilde{a_{nv}})}{q-\gamma\kappa(\underline{b}\widetilde{a_{nv}})}$ and $\widetilde{a_{nv}} < \bar{a}$

When $R'_a(b_e\bar{a})(1-\frac{\gamma\kappa(\underline{b}\bar{a})}{q}) < R(b_e\bar{a})\frac{\gamma}{q}\kappa'_a(\underline{b}\bar{a})$, the only solution is $b^*_{nv} = \underline{b}$ and $a^*_{nv} = \widetilde{a_{nv}}$, with $\widetilde{a_{nv}}$ satisfying $\frac{R'_a(b_e\widetilde{a_{nv}})}{R(b_e\widetilde{a_{nv}})} = \frac{\gamma\kappa'_a(\underline{b}\widetilde{a_{nv}})}{q-\gamma\kappa(\underline{b}\widetilde{a_{nv}})}$.

When $R'_a(b_e\bar{a})(1-\frac{\gamma\kappa(\underline{b}\bar{a})}{q}) > R(b_e\bar{a})\frac{\gamma}{q}\kappa'_a(\underline{b}\bar{a})$, we have $\widetilde{a_{nv}}$ is greater than $\bar{a}$. In this case, the only solution is $b^*_{nv} = \underline{b}$ and $a^*_{nv} = \bar{a}$. Hence , assuming rational expectation from advertisers $b_e = \underline{b}$, we have:

$$(b^*_{nv}, a^*_{nv}) = \begin{cases} (\underline{b}, \widetilde{a_{nv}}) \text{ if } \gamma > \frac{qR'_{a_{nv}}(\bar{a}\underline{b})}{\kappa'_{a_{nv}}(\bar{a}\underline{b})R(\bar{a}\underline{b})+\kappa(\bar{a}\underline{b})R'_{a_{nv}}(\bar{a}\underline{b})}, \\ \\ (\underline{b}, \bar{a}) \text{ if } \gamma \leq \frac{qR'_{a_{nv}}(\bar{a}\underline{b})}{\kappa'_{a_{nv}}(\bar{a}\underline{b})R(\bar{a}\underline{b})+\kappa(\bar{a}\underline{b})R'_{a_{nv}}(\bar{a}\underline{b})}. \end{cases} \tag{2.19}$$

However, we still have to show that $\widetilde{a_{nv}}$ is a unique maximum.

We know that at the equilibrium $\widetilde{a_{nv}}$ satisfies:

$$\frac{R'_a(b_e\widetilde{a_{nv}})}{R(b_e\widetilde{a_{nv}})} = \frac{\gamma\kappa'_a(\underline{b}\widetilde{a_{nv}})}{q-\gamma\kappa(\underline{b}\widetilde{a_{nv}})} \tag{2.20}$$

and can be rewritten using elasticities:

$$\epsilon_R = \epsilon_N. \tag{2.21}$$

As Anderson and Gans (2011), we want to show that $\frac{\partial\epsilon_R}{\partial a} < \frac{\partial\epsilon_N}{\partial a}$ for any a satisfying Eq. 2.20.

Considering that $\epsilon_R = \epsilon_N = \epsilon$, we have to prove that:

$$\frac{\partial\epsilon_a}{\partial a} = \frac{ar''}{r} - \frac{(1-\epsilon)(2-\epsilon)}{a} < \frac{\gamma a\kappa''(a)}{q-\gamma\kappa(a)} + \frac{\epsilon(1+\epsilon)}{a} = \frac{\partial\epsilon_N}{\partial a}, \tag{2.22}$$

which gives:

$$\frac{a^2 r''}{r} < \frac{\gamma a^2 \kappa''(a)}{q - \gamma\kappa(a)} + 2(1 - \epsilon + \epsilon^2). \tag{2.23}$$

As we know that $\kappa''(a) \geq 0$ and $\epsilon < 1$, we know that $\frac{\gamma A^2 \kappa''(a)}{q-\gamma\kappa(a)} + 2(1-\epsilon+\epsilon^2) < 2\epsilon^2$, and we simply have to prove that $\frac{a^2 r''}{r} > 2\epsilon^2$, which is always the case for log-concavity of r in a as demonstrated in Anderson and Gans (2011).

Ad Level with Viewability Technology

We perform the same exercise when there is viewability technology.

$$\max_A \mathcal{L}(A, \lambda_1, \lambda_2) = R(A)(1 - \frac{\gamma\kappa(A)}{q}) - \lambda_1(A - \overline{a}\overline{b}) + \lambda_2 A, \tag{2.24}$$

Solving the Langrangean program gives us the following Khun-Tucker conditions:

$$(A) \quad R'_A(A)(1 - \frac{\gamma\kappa(A)}{q}) - R(A)\frac{\gamma}{q}\kappa'_A(A) - \lambda_1 + \lambda_2 = 0$$

$$(B) \quad A \leq \overline{a}\overline{b}, \lambda_1 \geq 0 \text{ and } \lambda_1(A - \overline{a}\overline{b}) = 0$$

$$(C) \quad a \geq 0, \lambda_2 \geq 0 \text{ and } \lambda_2 a = 0$$

We have the following cases:

- $A = \overline{a}\overline{b}$ and $\lambda_2 = 0$. In this case, $\lambda_1 = R'_A(\overline{a}\overline{b})(1 - \frac{\gamma\kappa(\overline{a}\overline{b})}{q}) - R(\overline{a}\overline{b})\frac{\gamma}{q}\kappa'_A(\overline{a}\overline{b})$. Hence, this case is a potential solution if $R'_A(\overline{a}\overline{b})(1 - \frac{\gamma\kappa(\overline{a}\overline{b})}{q}) \geq R(\overline{a}\overline{b})\frac{\gamma}{q}\kappa'_A(\overline{a}\overline{b})$.

- $\lambda_1 = \lambda_2 = 0$. In this case $A = \widetilde{A}$ such that $\frac{R'_A(\widetilde{A})\widetilde{A}}{R(\widetilde{A})} = \frac{\widetilde{A}\gamma\kappa'_A(\widetilde{A})}{q-\gamma\kappa(\widetilde{A})}$ and $\widetilde{A} < \overline{a}\overline{b}$. This a potential solution,

- $\lambda_2 > 0$ which never occurs because it would contradicts condition (A) and (B).

When $R'_A(\overline{a}\overline{b})(1 - \frac{\gamma\kappa(\overline{a}\overline{b})}{q}) < R(\overline{a}\overline{b})\frac{\gamma}{q}\kappa'_A(\overline{a}\overline{b})$, there is only one solution available which is $A^*_v = \widetilde{A}$ with $\lambda_1 = 0$.

If $R'_A(\overline{a}\overline{b})(1 - \frac{\gamma\kappa(\overline{a}\overline{b})}{q}) \geq R(\overline{a}\overline{b})\frac{\gamma}{q}\kappa'_A(\overline{a}\overline{b})$, the n we have $\widetilde{A} > \overline{a}\overline{b}$. In this case, the only potential solution is $A^*_v = \overline{a}\overline{b}$ and $\lambda_1 = R'_A(\overline{a}\overline{b})(1 - \frac{\gamma\kappa(\overline{a}\overline{b})}{q}) - R(\overline{a}\overline{b})\frac{\gamma}{q}\kappa'_A(\overline{a}\overline{b})$.

We know that at the equilibrium:

$$\frac{\partial R(\widetilde{A})}{\partial A}\frac{\widetilde{A}}{R(\widetilde{A})}=\frac{\partial N(A)}{\partial A}\frac{\widetilde{A}}{N(\widetilde{A})}, \tag{2.25}$$

which write equivalently in:

$$\frac{R'_A(\widetilde{A})\widetilde{A}}{R(\widetilde{A})}=\frac{\widetilde{A}\gamma\kappa'_A(\widetilde{A})}{q-\gamma\kappa(\widetilde{A})}. \tag{2.26}$$ [31]

and can be rewritten using elasticities:

$$\epsilon_A=\epsilon_N. \tag{2.27}$$

As Anderson and Gans (2011), we want to show that $\frac{\partial\epsilon_A}{\partial A}<\frac{\partial\epsilon_N}{\partial A}$ for any A satisfying Eq. 2.26.

Considering that $\epsilon_A=\epsilon_N=\epsilon$, we have to prove that:

$$\frac{\partial\epsilon_A}{\partial A}=\frac{Ar''}{r}-\frac{(1-\epsilon)(2-\epsilon)}{A}<\frac{\gamma A\kappa''(A)}{q-\gamma\kappa(A)}+\frac{\epsilon(1+\epsilon)}{A}=\frac{\partial\epsilon_N}{\partial A}, \tag{2.28}$$

which gives:

$$\frac{A^2r''}{r}<\frac{\gamma A^2\kappa''(A)}{q-\gamma\kappa(A)}+2(1-\epsilon+\epsilon^2). \tag{2.29}$$

As we know that $\kappa''(A)\geq 0$ and $\epsilon<1$, we know that $\frac{\gamma A^2\kappa''(A)}{q-\gamma\kappa(A)}+2(1-\epsilon+\epsilon^2)<2\epsilon^2$, and we simply have to prove that $\frac{A^2r''}{r}>2\epsilon^2$, which is always the case for log-concavity of r in A as demonstrated in Anderson and Gans (2011).

Ad Level with Viewability Technology in presence of adblockers

We have the following lagrangean program:

$$\max_{b,a}\ \mathcal{L}(b,a,\lambda_1,\lambda_2,\lambda_3,\lambda_3,\lambda_4,\lambda_5)=R(b_ea)(1-\frac{\gamma\kappa(ba)}{q})-\lambda_1(b-\bar{b})+\lambda_2(b-\underline{b})-\lambda_3(a-\bar{a})+\lambda_4a-\lambda_5(\gamma\kappa(ba)-c \tag{2.30}$$

which gives us the following Kuhn-Tucker conditions:

$$(A)\quad R'_a(b_ea)(1-\frac{\gamma\kappa(ba)}{q})-R(b_ea)\frac{\gamma}{q}\kappa'_a(ba)-\lambda_3+\lambda_4-\lambda_5\gamma\kappa'_a(ba)=0$$

[31] $\kappa'_A=\frac{\partial\kappa}{\partial A}$ and $R'_A=\frac{\partial R}{\partial A}$.

$$(B) \quad -R(b_e a)\frac{\gamma}{q}\kappa'_b(ba) - \lambda_1 + \lambda_2 - \lambda_5\gamma\kappa'_b(ba) = 0$$

$$(C) \quad b \leq \bar{b}, \lambda_1 \geq 0 \text{ and } \lambda_1(b - \bar{b}) = 0$$

$$(D) \quad b \geq \underline{b}, \lambda_2 \geq 0 \text{ and } \lambda_2(b - \underline{b}) = 0$$

$$(E) \quad a \leq \bar{a}, \lambda_3 \geq 0 \text{ and } \lambda_3(a - \bar{a}) = 0$$

$$(F) \quad a \geq 0, \lambda_4 \geq 0 \text{ and } \lambda_4 a = 0$$

$$(G) \quad \gamma\kappa(ba) \leq c, \lambda_5 \geq 0 \text{ and } \lambda_5(\gamma\kappa(ba) - c) = 0$$

- If $\lambda_5 = 0$, under the condition $\gamma\kappa(ba) \leq c$, potential solutions remains unchanged as we are in the case where the publisher is not constrained by adblockers in its maximization program

- If $\lambda_5 > 0$, the only potential solution is ba such that $\kappa(ba) = c$, $b = \underline{b}$, $\lambda_1 = 0$, $\lambda_3 = \lambda_4 = 0$ and $\lambda_5 = \frac{R'_a(b_e a)(1-\frac{\gamma\kappa(ba)}{q})}{\gamma\kappa'_a(ba)} - \frac{R(b_e a)}{q}$ and $\lambda_2 = R(b_e a)\frac{\gamma}{q}\kappa'_b(ba) + \gamma\kappa'_b(ba)(\frac{R'_a(b_e a)(1-\frac{\gamma\kappa(ba)}{q})}{\gamma\kappa'_a(ba)} - \frac{R(b_e a)}{q}) > 0$. Hence, to be a potential solution, we need $\frac{R'_a(b_e a)(1-\frac{\gamma\kappa(ba)}{q})}{\gamma\kappa'_a(ba)} > \frac{R(b_e a)}{q}$

When $\frac{R'_a(b_e a)(1-\frac{\gamma\kappa(ba)}{q})}{\gamma\kappa'_a(ba)} < \frac{R(b_e a)}{q}$, the set of solution remains unchanged, as the adblocker is not contraining enough.

When $\frac{R'_a(b_e a)(1-\frac{\gamma\kappa(ba)}{q})}{\gamma\kappa'_a(ba)} > \frac{R(b_e a)}{q}$, we have $\gamma\kappa(ba) > c$. In this case, the only potential solution is when $\kappa(ba) = c$ and $b = \underline{b}$.

Ad Level with Viewability Technology in presence of adblockers

We have the following lagrangean program:

$$\max_A \mathcal{L}(b, a, \lambda_1, \lambda_2, \lambda_3) = R(A)(1 - \frac{\gamma\kappa(A)}{q}) - \lambda_1(A - \bar{b}\bar{a}) + \lambda_2 A - \lambda_3(\gamma\kappa(A) - c), \quad (2.31)$$

Solving the Langrangean program gives us the following Khun-Tucker conditions:

$$(A) \quad R'_A(A)(1 - \frac{\gamma\kappa(A)}{q}) - R(A)\frac{\gamma}{q}\kappa'_A(A) - \lambda_1 + \lambda_2 - \lambda_3\gamma\kappa'_A(A) = 0$$

$$(B) \quad A \leq \bar{a}\bar{b}, \lambda_1 \geq 0 \text{ and } \lambda_1(A - \bar{a}\bar{b}) = 0$$

$$(C) \quad a \geq 0, \lambda_2 \geq 0 \text{ and } \lambda_2 a = 0$$

$$(D) \quad \gamma\kappa(A) \leq c, \lambda_3 \geq 0 \text{ and } \lambda_3(\gamma\kappa(A) - c) = 0$$

We have the following potential solutions:

We have the following cases:

- $\lambda_3 = 0$ if $\gamma\kappa(A) < c$ This gives us the same set of solution as before.
- $\gamma\kappa(A) = c$. In this case, the only solution is $\lambda_1 = \lambda_2 = 0$, with $\lambda_3 = \frac{R'_A(A)(1-\frac{\gamma\kappa(A)}{q})}{\gamma\kappa'_A(A)} - \frac{R(A)}{q}$. Hence, this a solution only if $\frac{R'_A(A)(1-\frac{\gamma\kappa(A)}{q})}{\gamma\kappa'_A(A)} > \frac{R(A)}{q}$.

When $\frac{R'_A(A)(1-\frac{\gamma\kappa(A)}{q})}{\gamma\kappa'_A(A)} < \frac{R(A)}{q}$, the publisher is not contrained and the set of solutions remains the same.

When $\frac{R'_A(A)(1-\frac{\gamma\kappa(A)}{q})}{\gamma\kappa'_A(A)} > \frac{R(A)}{q}$, we have $\gamma\kappa(A) < c$, hence, the publisher has only one solution with $\gamma\kappa(A) = c$.

Proof of Proposition 2.

Proof. We want to prove that $b_v^* a_v^* = b_{nv}^* a_{nv}^*$ when the publisher is able to practice a local solution in both cases - which is equivalent to $\widetilde{A} = \underline{b}\widetilde{a}$. From Eqs. 2.7 and 2.17, we have the two following relation that should be equivalent:

$$\frac{R'_A(\widetilde{A})}{R(\widetilde{A})} = \frac{\gamma\kappa'_A(\widetilde{A})}{q - \gamma\kappa(\widetilde{A})}. \tag{2.32}$$

$$\frac{R'_a(\underline{b}\widetilde{a})}{R(\underline{b}\widetilde{a})} = \frac{\gamma\kappa'_a(\underline{b}\widetilde{a})}{q - \gamma\kappa(\underline{b}\widetilde{a})} \tag{2.33}$$

Replacing $\widetilde{A}$ by $\underline{b}\widetilde{a}$ in Eq. (2.32), we have

$$\frac{R'_A(\underline{b}\widetilde{a})}{R(\underline{b}\widetilde{a})} = \frac{\gamma\kappa'_A(\underline{b}\widetilde{a})}{q - \gamma\kappa(\underline{b}\widetilde{a})}. \tag{2.34}$$

To show equivalence between Eq. (2.32) and (2.33), we have to show equivalence when computing the derivatives of R and κ with respect to A and a. We then compute total differential functions of R and κ:

$$\frac{\partial R}{\partial a} = \frac{\partial R}{\partial A}\frac{\partial A}{\partial a} + \frac{\partial R}{\partial a} = \frac{\partial R}{\partial A} b,$$

$$\frac{\partial \kappa}{\partial a} = \frac{\partial \kappa}{\partial A}\frac{\partial A}{\partial a} + \frac{\partial \kappa}{\partial a} = \frac{\partial \kappa}{\partial A} b.$$

Replacing the latter results in Eq. (2.34):

$$\frac{R'_a(\underline{b}\widetilde{a})}{\underline{b}R(\underline{b}\widetilde{a})} = \frac{\gamma \kappa'_a(\underline{b}\widetilde{a})}{\underline{b}(q - \gamma\kappa(\underline{b}\widetilde{a}))}. \tag{2.35}$$

Which simplifies to the Eq. (2.33). □

Welfare Analysis of Ad Viewability.

We compute total welfare in considering $r(ba) = 1 - ba$ and $\kappa(ba) = ba$, which gives us the following optimal advertising levels:

$$b^*_{nv}a^*_{nv} = \begin{cases} \frac{q+\gamma-\sqrt{q^2-\gamma q+\gamma^2}}{3\gamma} \text{ if } \gamma > \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}, \\ \underline{b}\overline{a} \text{ if } \gamma \leq \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}, \end{cases} \tag{2.36}$$

$$b^*_{v}a^*_{v} = \begin{cases} \frac{q+\gamma-\sqrt{q^2-\gamma q+\gamma^2}}{3\gamma} \text{ if } \gamma > \frac{q(1-2\overline{b}\overline{a})}{\overline{b}\overline{a}(2-3\overline{b}\overline{a})}, \\ \overline{b}\overline{a} \text{ if } \gamma \leq \frac{q(1-2\overline{b}\overline{a})}{\overline{b}\overline{a}(2-3\overline{b}\overline{a})}, \end{cases} \tag{2.37}$$

We then compute Internet users' surplus, Advertisers' profits, Publisher's profits and total welfare.

Surplus of Internet Users

Internet users do not pay to visit the publisher website, and therefore the surplus only depends on user experience. The surplus of Internet users without and with viewability technology are respectively:

$$\mathrm{S}^{u*}_{nv} = \int_{\frac{\gamma b^*_{nv}a^*_{nv}}{q}}^{1} \theta q - \gamma b^*_{nv}a^*_{nv} \mathrm{d}\theta = \begin{cases} \frac{(q-\gamma b^*_{nv}a^*_{nv})^2}{2q} \text{ if } \gamma > \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}, \\ \frac{(q-\gamma \underline{b}\overline{a})^2}{2q} \text{ if } \gamma \leq \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}, \end{cases} \tag{2.38}$$

and,

$$\mathrm{S}^{u*}_{v} = \int_{\frac{\gamma b^*_{v}a^*_{v}}{q}}^{1} \theta q - \gamma b^*_{v}a^*_{v} \mathrm{d}\theta = \begin{cases} \frac{(q-\gamma b^*_{v}a^*_{v})^2}{2q} \text{ if } \gamma > \frac{q(1-2\overline{b}\overline{a})}{\overline{b}\overline{a}(2-3\overline{b}\overline{a})}, \\ \frac{(q-\gamma \overline{b}\overline{a})^2}{2q} \text{ if } \gamma \leq \frac{q(1-2\overline{b}\overline{a})}{\overline{b}\overline{a}(2-3\overline{b}\overline{a})}, \end{cases} \tag{2.39}$$

Overall, we see that the surplus of Internet users may decrease when introducing viewability technology. Indeed, when $\gamma > \frac{q(1-2\underline{b}\bar{a})}{\underline{b}\bar{a}(2-3\underline{b}\bar{a})}$, the introduction of viewability has no impact on Internet users surplus as ${\mathrm{S}^{u}_{nv}}^{*} = \mathrm{S}^{u*}_{v} = \mathrm{S}^{u*}$. If $\frac{q(1-2\underline{b}\bar{a})}{\underline{b}\bar{a}(2-3\underline{b}\bar{a})} \geq \gamma > \frac{q(1-2\bar{b}\bar{a})}{\bar{b}\bar{a}(2-3\bar{b}\bar{a})}$, therefore ${\mathrm{S}^{u}_{nv}}^{*} = \frac{(q-\gamma\underline{b}\bar{a})^{2}*}{2q} > \frac{(q-\gamma b^{*}_{v}a^{*}_{v})^{2}}{2q} = \mathrm{S}^{u*}_{v}$. When $\gamma > \frac{q(1-2\underline{b}\bar{a})}{\underline{b}\bar{a}(2-3\underline{b}\bar{a})}$, ${\mathrm{S}^{u}_{nv}}^{*} = \frac{(q-\gamma\underline{b}\bar{a})^{2}*}{2q} > \frac{(q-\gamma\bar{b}\bar{a})^{2}}{2q} = \mathrm{S}^{u*}_{v}$.

This is very intuitive as the advertising nuisance, which decreases the Internet users' utility, is greater with viewability.

Using Eqs. (2.38) and (2.39), we are able to compute the difference in user surplus from the introduction of a viewability technology.

$$\Delta\mathrm{S}^{u} = \begin{cases} 0 \text{ if } \gamma > \frac{q(1-2\underline{b}\bar{a})}{\underline{b}\bar{a}(2-3\underline{b}\bar{a})}, \\ -\frac{\gamma(b^{*}_{v}a^{*}_{v}-\underline{b}\bar{a})(2q-\gamma(\underline{b}\bar{a}+b^{*}_{v}a^{*}_{v}))}{2q} \text{ if } \gamma \leq \frac{q(1-2\underline{b}\bar{a})}{\underline{b}\bar{a}(2-3\underline{b}\bar{a})}. \end{cases} \tag{2.40}$$

Surplus of Advertisers

The surplus of advertisers without viewability technology, defined in Eq. (2.41),

$$\mathrm{S}^{a*}_{nv} = N^{*}_{nv}\int_{0}^{a^{*}_{nv}}(r(b^{*}_{nv}a) - r(b^{*}_{nv}a^{*}_{nv})\mathrm{d}a = \begin{cases} \frac{(b^{*}_{nv}a^{*}_{nv})^{2}}{2}\frac{q-\gamma b^{*}_{nv}a^{*}_{nv}}{q} \text{ if } \gamma > \frac{q(1-2\underline{b}\bar{a})}{\underline{b}\bar{a}(2-3\underline{b}\bar{a})}, \\ \frac{\underline{b}\bar{a}^{2}}{2}\frac{q-\gamma\underline{b}\bar{a}}{q} \text{ if } \gamma \leq \frac{q(1-2\underline{b}\bar{a})}{\underline{b}\bar{a}(2-3\underline{b}\bar{a})}. \end{cases} \tag{2.41}$$

However, with viewability technology, the price of ads adjusts to the actual level of ad viewability set up by the publisher:

$$\mathrm{S}^{a*}_{v} = N^{*}_{v}\int_{0}^{a^{*}_{v}}(r(b^{*}_{v}a) - r(b^{*}_{v}a^{*}_{v})\mathrm{d}a = \begin{cases} \frac{(b^{*}_{v}a^{*}_{v})^{2}}{2}\frac{q-\gamma b^{*}_{v}a^{*}_{v}}{q} \text{ if } \gamma > \frac{q(1-2\bar{b}\bar{a})}{\bar{b}\bar{a}(2-3\bar{b}\bar{a})}, \\ \frac{\bar{b}\bar{a}^{2}}{2}\frac{q-\gamma\bar{b}\bar{a}}{q} \text{ if } \gamma \leq \frac{q(1-2\bar{b}\bar{a})}{\bar{b}\bar{a}(2-3\bar{b}\bar{a})}. \end{cases} \tag{2.42}$$

We find that the surplus of advertisers is higher with viewability technology only if the level of advertising nuisance is low. Indeed, the level of advertising intensity that maximizes advertisers surplus is $ba = \frac{2q}{3\gamma}$, which is always higher than the optimal level of advertising set up by the publisher for $q, \gamma > 0$. Therefore the higher the advertising intensity, the higher the surplus of advertisers.

Hence, when $\gamma > \frac{q(1-2\underline{b}\bar{a})}{\underline{b}\bar{a}(2-3\underline{b}\bar{a})}$, the introduction of viewability has no impact on advertisers surplus as ${\mathrm{S}^{a}_{nv}}^{*} = \mathrm{S}^{a*}_{v} = \mathrm{S}^{a*}$. If $\frac{q(1-2\underline{b}\bar{a})}{\underline{b}\bar{a}(2-3\underline{b}\bar{a})} \geq \gamma > \frac{q(1-2\bar{b}\bar{a})}{\bar{b}\bar{a}(2-3\bar{b}\bar{a})}$, therefore ${\mathrm{S}^{a}_{nv}}^{*} = \frac{\underline{b}\bar{a}^{2}(q-\gamma\underline{b}\bar{a})}{2q} < \frac{(b^{*}_{v}a^{*}_{v})^{2}(q-\gamma b^{*}_{v}a^{*}_{v})}{2q} = \mathrm{S}^{a*}_{v}$. When $\gamma > \frac{q(1-2\underline{b}\bar{a})}{\underline{b}\bar{a}(2-3\underline{b}\bar{a})}$, ${\mathrm{S}^{u}_{nv}}^{*} = \frac{\underline{b}\bar{a}^{2}(q-\gamma\underline{b}\bar{a})}{2q} < \frac{\bar{b}\bar{a}^{2}(q-\gamma\bar{b}\bar{a})}{2q} = \mathrm{S}^{a*}_{v}$.

Advertisers are better off with viewability technology. This is due to the fact that the introduction of ad viewability fosters the publisher to practice an equal or higher level of advertising intensity. Advertisers' optimal level of advertising intensity is higher than the level of advertising intensity practiced by the publisher. Therefore they will prefer a higher level of advertising intensity, which is offered with ad viewability.

Using Eqs. (2.41) and (2.42), the change in advertisers' surplus from the introduction of the viewability technology is:

$$\Delta \mathrm{S}^a = \begin{cases} 0 \text{ if } \gamma > \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}, \\ \\ \frac{(b_v^* a_v^* - \underline{b}\overline{a})(q(\underline{b}\overline{a} + b_v^* a_v^*) - \gamma(\underline{b}\overline{a}^2 + \underline{b}\overline{a} b_v^* a_v^* + (b_v^* a_v^*)^2))}{2q} \text{ if } \gamma \leq \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}. \end{cases} \tag{2.43}$$

Profits of the Publisher

The profits of the publisher are defined as:

$$\Pi_{nv}^* = \begin{cases} (1 - b_{nv}^* a_{nv}^*) A^* (1 - \frac{\gamma b_{nv}^* a_{nv}^*}{q}) \text{ if } \gamma > \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}, \\ \\ (1 - \underline{b}\overline{a})\underline{b}\overline{a}(1 - \frac{\gamma \underline{b}\overline{a}}{q}) \text{ if } \gamma \leq \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}, \end{cases} \tag{2.44}$$

and

$$\Pi_v^* = \begin{cases} (1 - b_v^* a_v^*) b_v^* a_v^* (1 - \frac{\gamma b_v^* a_v^*}{q}) \text{ if } \gamma > \frac{q(1-2\overline{b}\overline{a})}{\overline{b}\overline{a}(2-3\overline{b}\overline{a})}, \\ \\ (1 - \overline{b}\overline{a})\overline{b}\overline{a}(1 - \frac{\gamma \overline{b}\overline{a}}{q}) \text{ if } \gamma \leq \frac{q(1-2\overline{b}\overline{a})}{\overline{b}\overline{a}(2-3\overline{b}\overline{a})}. \end{cases} \tag{2.45}$$

The publisher makes higher profits with ad viewability as it practices a corner solution without ad viewability. Indeed, the publisher always prefer a non-constrained profit than a constrained one. Therefore, it will always be better with ad viewability when $\gamma < \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}$.

Using Eqs. (2.44) and (2.45), the change in publisher profits from the introduction of the viewability technology is:

$$\Delta \Pi = \begin{cases} 0 \text{ if } \gamma > \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}, \\ \\ \frac{(1-b_v^* a_v^*) b_v^* a_v^* (q - \gamma b_v^* a_v^*) - (1-\underline{b}\overline{a})\underline{b}\overline{a}(q - \gamma \underline{b}\overline{a})}{q} \text{ if } \gamma \leq \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}. \end{cases} \tag{2.46}$$

Total Welfare

The total welfare without and with viewability technology are respectively:

$$W^*_{nv} = \begin{cases} \frac{(q-\gamma b^*_{nv}a^*_{nv})}{2q}(q + b^*a^*(2 - b^*_{nv}a^*_{nv} - \gamma)) \text{ if } \gamma > \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}, \\ \frac{(q-\gamma \underline{b}\overline{a})}{2q}(q + \underline{b}\overline{a}(2 - \underline{b}\overline{a} - \gamma)) \text{ if } \gamma \leq \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}, \end{cases} \tag{2.47}$$

and,

$$W^*_{v} = \begin{cases} \frac{(q-\gamma b^*_{v}a^*_{v})}{2q}(q + b^*_{v}a^*_{v}(2 - b^*_{v}a^*_{v} - \gamma)) \text{ if } \gamma > \frac{q(1-2\overline{b}\overline{a})}{\overline{b}\overline{a}(2-3\overline{b}\overline{a})}, \\ \frac{(q-\gamma \overline{b}\overline{a})}{2q}(q + \overline{b}\overline{a}(2 - \overline{b}\overline{a} - \gamma)) \text{ if } \gamma \leq \frac{q(1-2\overline{b}\overline{a})}{\overline{b}\overline{a}(2-3\overline{b}\overline{a})}, \end{cases} \tag{2.48}$$

Proof of Proposition 4.

Proof. We look at the value of ΔW^* with respect to q and γ.

Firstly, when $\gamma > \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}$, the introduction of viewability has no impact on total welfare as $W^*_{nv} = W^*_v = W^*$.

Secondly, we analyze the case $\gamma < \frac{q(1-2\underline{b}\overline{a})}{\underline{b}\overline{a}(2-3\underline{b}\overline{a})}$. In this case we compare the value of $b^*_v a^*_v$ and $\underline{b}\overline{a}$ with respect to q. We find that $W^*_v - W^*_{nv} > 0$ only when $q(2\gamma + b^*_v a^*_v + \underline{b}\overline{a} - 2) < (b^*_v a^*_v + \underline{b}\overline{a})\gamma^2 + (b^*_v a^{*2}_v + (\underline{b}\overline{a} - 2)b^*_v a^*_v + \underline{b}\overline{a}^2 - 2\underline{b}\overline{a})\gamma$.

- If $\gamma > 1 - \frac{\underline{b}\overline{a} + b^*_v a^*_v}{2}$, as$b^*_v a^*_v > \underline{b}\overline{a}$, we always have $W^*_{nv} > W^*_v$.
- If $\gamma < 1 - \frac{\underline{b}\overline{a} + b^*_v a^*_v}{2}$, $W^*_{nv} < W^*_v$ only if $q > \frac{(b^*_v a^*_v + \underline{b}\overline{a})\gamma^2 + ((b^*_v a^*_v)^2 + (\underline{b}\overline{a} - 2)b^*_v a^*_v + \underline{b}\overline{a}^2 - 2\underline{b}\overline{a})\gamma}{2\gamma + b^*_v a^*_v + \underline{b}\overline{a} - 2} \equiv q_w$

We want to know the sign of $\frac{\partial q_w}{\partial \gamma}$. We find that $\frac{\partial q_w}{\partial \gamma} > 0 \Leftrightarrow \frac{\gamma(2-b^*_v a^*_v)b^*_v a^*_v + \underline{b}\overline{a}(\underline{b}\overline{a}-2)}{(2(\gamma-1)+\underline{b}\overline{a}+(b^*_v a^*_v)^2} + \frac{\underline{b}\overline{a}^2 - (\underline{b}\overline{a} + b^*_v a^*_v)(2-\gamma-b^*_v a^*_v)}{2(\gamma-1)+\underline{b}\overline{a}+b^*_v a^*_v} + \frac{\partial (b^*_v a^*_v)}{\partial \gamma}\gamma\big(\underline{b}\overline{a} + 2b^*_v a^*_v - 2 + \gamma\big)\big(\underline{b}\overline{a} + b^*_v a^*_v + 2(\gamma - 1)\big) - 2\big(\underline{b}\overline{a}^2 - (b^*_v a^*_v + \underline{b}\overline{a})(2 - \gamma - b^*_v a^*_v)\big) > 0$. As we focus on the case $\gamma < 1 - \frac{\underline{b}\overline{a} + b^*_v a^*_v}{2}$, the derivative is always positive, whether $b^*_v a^*_v$ depends on γ or not. □

Proof of Lemma 1.

Proof. Let $\gamma_2 = \frac{\tau R'_a(\overline{a}\overline{b_2})}{\kappa'_a(\overline{a}\overline{b_2})R(\overline{a}\overline{b_2})}$ denote the level of ad nuisance cost such that:

$$A^*_{2,v} = \begin{cases} \widetilde{A_2} \text{ if } \gamma > \gamma_2, \\ \overline{a}\overline{b_2} \text{ if } \gamma \leq \gamma_2 \end{cases} \tag{2.49}$$

with $\frac{R'_{A_2}(\widetilde{A_2})}{R(\widetilde{A_2})} = \frac{\gamma\kappa'_{A_2}(\widetilde{A_2})}{\tau-\gamma(\kappa(\widetilde{A_2})-\kappa(\widetilde{A_1}))}$.

Hence if $\gamma > \gamma_2$, both publisher pratice the local solution and we have $A^*_{1,v} = A^*_{2,v} = \widetilde{A_i}$.

Let $\gamma_1 = \frac{\tau(\kappa(\overline{a}\overline{b_1})-\kappa(\overline{a}\overline{b_2}))R'_a(\overline{a}\overline{b_1})}{\kappa'_a(\overline{a}\overline{b_2})R(\overline{a}\overline{b_1})}$. If $\gamma_1 < \gamma < \gamma_2$, publisher 1 sets $A^*_{1,v} = \widetilde{A_1}$ such that $\frac{R'_{A_2}(\widetilde{A_1})}{R(\widetilde{A_1})} = \frac{\gamma\kappa'_{A_1}(\widetilde{A_1})}{\tau-\gamma(\kappa(\widetilde{A_1})-\kappa(A^*_{2,v}))}$ while $A^*_{2,v} = \overline{a}\overline{b_2}$. Finally, if $\gamma < \gamma_1$, we have $A^*_{1,v} = \overline{a}\overline{b_1}$ and $A^*_{2,v} = \overline{a}\overline{b_2}$

□

Proof of Proposition 5.

Proof. Let $\gamma_2 = \frac{\tau R'_a(\overline{a}\overline{b_2})}{\kappa'_a(\overline{a}\overline{b_2})R(\overline{a}\overline{b_2})}$, $\underline{\gamma} = \frac{\tau R'_a(\overline{a}\underline{b})}{\kappa'_a(\overline{a}\underline{b})R(\overline{a}\underline{b})}$ and $A_{i,k} = b^*_{i,k}a^*_{i,k}$ with $i \in \{1, 2\}$ and $k \in \{v, nv\}$.

We know that $a^*_{1,nv} = a^*_{2,nv} = \widetilde{a}_c$ when $\gamma > \underline{\gamma}$ and that $a^*_{1,nv} > a^*_{2,nv}$ when $\gamma < \underline{\gamma}$.

We know from previous Lemma that $A^*_{1,v} = A^*_{2,v} = A_i$ when $\gamma > \gamma_2$ and that $A^*_{1,v} > A^*_{2,v}$ when $\gamma < \gamma_2$.

By the same mecanism of the Proof of Proposition 2, it is easy to demonstrate that $\widetilde{a}_c\underline{b} = \widetilde{A_i}$ when $\gamma > \underline{\gamma}$. As we have $\underline{\gamma} > \gamma_2$, when $\gamma > \underline{\gamma}$, both publishers always choose the corner solution no matter the introduction of viewability technology. We have therefore $A_{1,v} = A_{1,v} = A_{1,nv} = A_{2,nv}$.

However, when $\underline{\gamma} > \gamma > \gamma_2$, both publishers are symmetrically constrained on the maximum number of ads they can practice when there is no viewability technology $A_{1,nv} = A_{2,nv} = \overline{a}\underline{b}$, while no publisher is contrained when there is viewability technology $A_{1,v} = A_{2,v} = \widetilde{A_{i,v}}$. This is due to rational anticipations of advertisers $b_e = \underline{b}$ in the absence of viewability technology. Hence we have $A_{1,v} = A_{2,v} > A_{1,nv} = A_{1,nv}$.

When $\gamma_1 < \gamma < \gamma_2$, both publishers are again constrained when there is no viewability technology $A_{1,nv} = A_{2,nv} = \overline{a}\underline{b}$. However, only publisher 2 is constrained when there is viewability technology $A_{2,v} = \overline{a}\overline{b_2}$, while publisher 1 can practice the corner solution $A_{1,v} = \widetilde{A_i}$. In this case, we have $A_{1,v} > A_{2,v} > A_{1,nv} = A_{1,nv}$.

Finally, when $\gamma_1 > \gamma$, all publishers are constrained, no matter the presence of viewability technology. Whithout viewability technology, all publishers are contrained by $\overline{a}\underline{b}$. However, with viewability technology, publisher i is contrained by $\overline{a}\overline{b_i}$. As $\overline{b_1} > \overline{b_2}$, we have $A_{1,v} > A_{2,v} > A_{1,nv} = A_{1,nv}$. □

Chapter 3

Targeting Advertising Preferences

3.1 Introduction

Advertising is the main mechanism to finance media content such as TV, radio or magazines. It also constitutes most of media platforms income on the Internet (Shiller et al., 2017). For instance, in the third quarter of 2016, U.S. advertisers invested $17.6 billion in digital advertising according to the Internet Advertising Bureau,[1] a 20 percent increase over the same time period in 2015.

However, Internet users tolerate less and less online ads that degrade their online experience (Manchanda et al., 2006; Goldfarb and Tucker, 2011), and as a response, they block ads using ad-avoidance technologies (AATs).[2] In 2016, 26.3% of US online users were using an ad-blocker (69.8 million Americans), a jump of 34.4% over last year.[3] AATs can be installed on web browser, and forthcoming versions of Chrome and Firefox will even directly integrate them by default.[4] The ad-blocking feature will filter out automatically ads such as pop-ups and auto-playing video.

The growing adoption of AATs by a significant part of online users confirms that advertising is perceived as a nuisance. However, an even larger proportion of consumers still continue to visit websites and click on ads that are related or targeted to their personal tastes and interests. Taste for advertising, or equivalently, ad preferences therefore vary across online users.

[1] IAB, Highest Third Quarter Spending on Record, December 28, 2016.
[2] Adblock Plus or uBlock are examples of popular AATs on the Internet.
[3] eMarketer, US Ad Blocking to Jump by Double Digits This Year, June 21, 2016.
[4] ArsTechnica, Report: Google will add an ad blocker to all versions of Chrome Web browser, April 19, 2017.

Online content providers may adopt various strategies to account for users taste for ads. For example, Tag (2009) analyzes a monopoly platform that offers either an ad-free pure subscription option or an ad-only program to discriminate consumers. The author shows that the subscription option induces a higher level of advertising for those remaining on the ad-only option. In the end, the aggregate consumer surplus falls, whereas the advertiser surplus rises through lower ad prices. Anderson and Gans (2011) extend the possibility for consumers to choose a costly ad-avoidance technology (such as TiVo) to strip out the ad nuisance in place of a subscription option. They confirm the results of Tag (2009), and also show that the adoption of AATs may reduce program quality.

The platform strategies analyzed in Tag (2009) and Anderson and Gans (2011) rely on the idea that ads are only perceived as a nuisance, and that consumers' advertising preferences are not directly observable (they are exogenous). By choosing a program, online users self-select and reveal afterward their taste for advertising. However, recent profiling technologies are allowing platforms to infer the taste for advertising of online users. On the basis of past behavior (click on ads, number of ads seen, etc.), or by associating consumers that have similar profiles, platforms can determine whether a user is more or less ad-sensitive, and tailor accordingly the level of advertising to prevent him adopting AATs. For example, digital platforms are using profiling technologies to limit the number of ads seen by consumer, the frequency of exposure of a consumer to a given ad, and ads to products a consumer has already purchased. As illustrations, Facebook and Snapchat limit the number of ads users can see in newsfeeds[5] or to be appeared between friends' snaps and stories.[6] This technique is called "frequency capping" (Buchbinder et al., 2011).

This article analyzes the impact of the use of a profiling technology on platforms' strategies, the surplus of consumers and advertisers, and the volume of ads served on the market. We develop a model in line with Anderson and Coate (2005), Tag (2009) and Anderson and Gans (2011), where a monopoly platform infers consumers' ad preferences using a profiling technology. To measure the impact of such technology, we base the model on three key features. Firstly, advertising is not necessarily perceived as a nuisance, as it may benefit Internet users to see some ads. Secondly, online users are heterogeneous in their taste for ads. Two types of consumers are

[5] TechCrunch, How Facebook News Feed Works, Sep 6, 2016.

[6] Adweek, Snapchat Is Letting More Brands Run Ads Between Friends' Stories, August 10, 2016.

distinguished: there is a proportion of consumers who have a *low taste* (γ_l) for advertising and a proportion of consumers who have a *high taste* (γ_h). Thirdly, we allow the platform to classify users according to their taste for ads using a profiling technology which precision is exogenous and can range between no identification and perfect identification. However, the profiling technology is not always *perfectly efficient*, i.e. it does not perfectly identify the users taste for ads. Indeed, the technology may be influenced by the provision of users' personal information as well as by the sophistication of the technology itself. In this case, consumers γ_l and γ_h may respectively see a level of advertising ill-adapted to their advertising preferences, hence adopting AATs.

Our model shows that the use of a profiling technology by a platform has strategic implications. Firstly, we highlight that the profiling technology may not always be used by the platform. If the technology is not efficient enough - or there is no profiling technology -, the platform cannot precisely infer the advertising preferences of each online user. It will therefore set a unique level of advertising for all users on the site, whether they like advertising or not. But this strategy may foster users to avoid ads. Conversely, the platform uses the profiling technology only if it generates greater revenue. This situation arises when the technology is *efficient* such that the probability of correctly predicting users taste for advertising is sufficiently high. In this case, the technology allows the platform to tailor the levels of advertising to the taste of users, thus increasing its profits. However, one important feature of our model considers that an efficient profiling technology is not always flawless. Indeed, a *perfect* profiling technology always perfectly classifies Internet users according to their type. This would happen in presence of a perfectly trained and up-to-date profiling technology. However, an efficient but *imperfect* technology makes mistakes, even if it generate higher profits for the platform. This is more likely to happen as the profiling technology may face barriers when predicting users' type. Hence, the technology may classify a consumer as enjoying advertising when he actually hates it, and *vice versa*. Indeed, even if the use of the technology increases the platform's revenue, it may not properly tailor the level of advertising to the taste of each type of user. This can have big implications on the platform profits. For example, a user who is wrongly classified as hating advertising will see a lower number of ads, thus reducing the profit opportunities for the platform. On the contrary, a user who is wrongly classified as enjoying advertising will see too many ads, and will have incentives to avoid ads.

Secondly, the use of an efficient profiling technology by the platform changes the total number of ads served on the market. We find that when the platform uses a perfect profiling technology, the total number of ads served to Internet users is always higher than without technology. This is intuitive as the technology allows the platform to serve more ads to users who have a high taste for advertising, while serving an appropriate level of ads to users who were avoiding ads before. Conversely, we contrast this result when the platform uses an efficient but imperfect technology. In this case, if the platform faces an audience that is hard to attract , more ads will be served with technology. Conversely, if the platform addresses an audience that have high taste for advertising, the total number of ads served is higher with technology only if this technology is highly efficient.

Thirdly, we show that introducing a perfect profiling technology is always welfare increasing when the platform is facing and audience that is cheap to attract, as the gains for advertisers and the platform offset the potential loss in Internet users' surplus. However, the analysis become more complex when 1) the technology is perfect but the platform is facing an audience that dislike ads very much or 2) the technology is imperfect. In both case, the impact of the technology on welfare relies on Internet user surplus and advertisers profits. Firstly, the impact of a profiling technology on Internet users surplus depends on whether users with a high taste for advertising prefers to see many ads (being correctly classified) or few ads (being misclassified). Secondly, we show that advertisers may not benefit from the introduction of a profiling technology, while the platform does. This situation happens when the technology is sufficiently efficient to be implemented by the platform but not efficient enough to increase advertisers' profits.

The chapter is organized as follows. The next section describes the model. Section 3.3 solves the case where the platform does not use a profiling technology. Section 3.4 introduces the profiling technology. Section 3.5 analyzes the aggregate number of ads served on the market. Section 3.7 summarizes the key findings and strategic implications of the chapter, and concludes.

3.2 Description of the Model

Our model deals with a monopolist publisher (also called platform) that delivers content and displays ads on a website to online users (also named consumers), and sells advertising space to advertisers. The platform therefore manages its website to attract online users on one side and advertisers on the other.

3.2.1 Online Users

Online users visit the website and receive utility:

$$U = \begin{cases} 1+\gamma(a), \text{ if choosing to visit and see ads,} \\ \\ \theta, \text{ if choosing to visit without seeing ads,} \end{cases} \tag{3.1}$$

and 0 otherwise. A consumer receives therefore utility by viewing an editorial content of quality q, normalized here to 1. If she chooses to visit the website and see ads, she receives $\gamma(a)$, where a is the number of ads (or advertisers), and γ the taste for advertising. $\gamma(a)$ is assumed to be continuous, differentiable, strictly concave in a, with $\gamma(a=0)=0$. $\gamma(a)$ has two notable properties. First, when no ads is displayed on the website, the consumer only enjoys the quality of editorial content. Second, we hypothesize that a low number of ads is enjoyable and provides utility for consumers; targeted advertising for instance provides useful information about product firms. However, too many ads displayed on the website, even though they are targeted, will degrade the user experience and provide disutility. Advertising is perceived in this case as a nuisance. This is crucial in generating demand for AATs. Indeed, when advertising is perceived as a nuisance, the consumer can choose to view content and earn utility θ for using AATs. To ensure that, we therefore also assume that $\gamma(a)$ admits a unique positive maximum such that the Internet user enjoys an additional ad until a certain point where it generates more nuisance than utility.

Consumers are heterogeneous in their taste for advertising. To keep things as simple as possible, we consider two populations. First, there is a proportion β of consumers who have a *low taste* for ads: γ_l. By contrast, there is a proportion $1-\beta$ of consumers who has a *high taste* for ads: γ_h. By definition, consumers of type γ_l would prefer to see a lower number of ads compared to *high taste* consumers γ_h. From the properties on $\gamma(a)$, we make two hypothesis. Firstly, we assume that γ_l is dominated by γ_h such that $\gamma_h(a) \geq \gamma_l(a)$, $\forall a > 0$. This corresponds to the fact that consumers who have a *high taste* for ads are willing to watch more ads than those have a *low taste*. Secondly, we assume that the value of a that maximizes $\gamma_l(a)$ is lower than the one that maximizes $\gamma_h(a)$. Hence consumers who have a *high taste* for ads ideally prefer to watch more ads than *low taste* ones.

To summarize, the utility function of a consumer of type $i \in \{l, h\}$ is:

$$U_i = \begin{cases} 1 + \gamma_i(a), \text{ if choosing to visit and see ads,} \\ \\ \theta, \text{ if choosing to visit and avoid ads,} \end{cases} \tag{3.2}$$

From Eq. 3.2, it is straightforward to see that when the cost to block ads is high, i.e. when $\theta < 0$, the choice of the consumer is, whatever his taste for advertising, either to consume with ads or not to visit the website. In the sequel, we will only consider the cases where the cost of using AATs is low enough, i.e. when $\theta > 0$. In other words, Internet users have only the choice between consuming and seeing ads or consuming without ads and using AATs. To ensure that Internet users do not always avoid ads, we assume that when there is zero ads ($a = 0$), Internet users always get a higher utility from consuming content than from trying to avoid them, which translates by $\theta < 1$.

This assumption is clearly in line with recent announces from Google and other large Internet companies. Google for example is about to roll out an ad-blocking version of his Chrome web browser (Wall Street Journal, Google Plans Ad-Blocking Feature in Popular Chrome Browser, April 19, 2017).[7] The ad-blocking feature, which could be switched on by default, would filter out automatically unacceptable ads such as pop-ups and auto-playing video that degrade online user experience. As the feature of ad-blocking would be installed by default, the cost of blocking ads would be dramatically low.

3.2.2 Advertisers

Following Anderson and Coate (2005) and Anderson and Gans (2011), we assume that a single ad suffices to reach all consumers on the platform, and so an advertiser decides to place an ad as long as the profit per consumer is no smaller than the price paid for an ad per consumer reached. We also rank advertisers from highest to lowest willingness to pay per consumer and derive the advertiser inverse demand curve $r(a)$. The corresponding total advertising revenue earned per individual, $R(a) = r(a)a$, where $r(a)$ is concave, and $r'(a) < 0$ when $r(a) > 0$, making $R(a)$ concave in a.[8]

[7] In the U.S., Chrome has nearly 47.5% of the browser market across all platforms, according to online analytics provider StatCounter.

[8] These properties are similar to the ones in Anderson and Gans (2011).

3.2.3 Platform

The platform delivers content to Internet users and displays ads on his website. The platform is only financed by advertising (and not by subscription). His profit function corresponds to the revenue per Internet user $R(a)$ times the number N of Internet users:

$$\Pi = r(a)aN(a) = R(a)N(a). \tag{3.3}$$

As $R(a)$ is concave in a (because $r(a)$ is), the platform is interesting to get on board enough advertisers to be profitable, but up to a certain level, too many advertisers lower its revenues (at a certain point, $r(a)$ is too low as the additional attracted advertiser expects a lower impact of its ad). We notice that $N(a)$ is the discrete function picturing user demand. Depending on the level of advertising a chosen by the platform, it can be equal to 0, $1-\beta$ or 1.[9]

3.3 Baseline Model Without Profiling Technology

In the baseline case, the platform has no profiling technology to find out the consumer taste for advertising. We introduce the profiling technology in Section 3.4.

The timing of the game in this baseline case is in two stages. In stage 1, the platform chooses the number of advertisers a (or equivalently, the number of ads to be displayed), and in stage 2, consumers choose to visit the website. In other words, when visiting the website, consumers discover the level of advertising and its possible nuisance, and choose to see or not ads using AATs. We solve the game by backward induction.

3.3.1 Stage 2

The users decide to view ads or to avoid them using AATs. Consumers of type γ_l and γ_h choose to visit the website depending of their respective taste for advertising. Their respective utilities

[9] By assumption $1+\gamma_h(a) > 1+\gamma_l(a)\ \forall a$, and $N(a)$ can be written as:

$$N(a) = \begin{cases} 0, \text{ if } \theta > 1+\gamma_h(a), \\ 1-\beta, \text{ if } 1+\gamma_h(a) > \theta > 1+\gamma_l(a), \\ 1, \text{ if } 1+\gamma_l(a) > \theta. \end{cases} \tag{3.4}$$

are defined in Eq. (3.2). Let define m_l and m_h as the solutions for which $1 + \gamma_l(m_l) = \theta$ and $1 + \gamma_h(m_h) = \theta$. These respective solutions can be interpreted as the maximum levels of advertising that users γ_l and γ_h are willing to accept to visit a website.

If the cost of AATs decreases for example, consumers γ_l and γ_h are both less ready to see ads. However, consumers γ_l and γ_h respond differently to a change in AATs. For any given θ, consumers γ_l who have a low taste for ads to advertising, will accept to see less ads than consumers γ_h, who have a high taste for advertising. As consumers γ_l are more sensitive to the nuisance cost of ads than consumers γ_h, the former will be more likely to use AATs when its utility from avoidance ads θ increases.

3.3.2 Stage 1

As the platform cannot discriminate consumers in the baseline case with a profiling technology, it can choose either to sell advertising spaces to a low number of advertisers, serving possibly consumers γ_l and γ_h, or to sell spaces to a high number of advertisers and reach only consumers γ_h who have a high taste for ads. As setting a too high advertising level repel all consumers from visiting the platform, the constrained optimization problem is:

$$\begin{aligned} &\max_a \; \Pi = r(a)aN = R(a)N. \\ &\text{subject to } a \leq m_h \end{aligned} \tag{3.5}$$

The strict concavity of the profit function guarantees the existence and uniqueness of an optimal level of advertising noted $\hat{a}$.[10] The latter may be lower than the maximum levels of advertising m_l and m_h that consumers γ_l and γ_h are willing to accept to visit the website. In this case, the ad-avoidance technology does not exert any competitive pressure on the platform, and there is an interior equilibrium where ads are served to both types of users. This is examined in Proposition 1.

Proposition 1: When advertising avoidance is expensive (which corresponds to $1 + \gamma_l(\hat{a}) \geq \theta$), interior equilibrium exists where the platform sets a level $\hat{a}$ ads to be served to both types of users.

Proposition 1 contrasts the analysis of Tag (2009) and Anderson and Gans (2011). When Internet users have a high taste for ads, ads are perceived as a benefit and not as a nuisance by

[10] Where it satisfies the following condition $r(\hat{a}) = -\hat{a}\frac{\partial r(\hat{a})}{\partial a}$.

both types of users; consumers derive a positive utility from seeing ads, and can be defined as ad-lovers. In this case, the platform can choose to display few ads to both types of users to maximize its profits. The level of advertising does not therefore depend on AATs. This case is displayed in Figure 3.1: when setting $a^* = \hat{a}$, the platform earns $\Pi^* = R(a^* = \hat{a}) = \hat{R}$, and both types of consumers γ_l and γ_h visit the website ($N = 1$) as the optimal level of advertising is lower than the ones which would exclude both types of consumers from visiting the website, i.e. $\hat{a} < m_l < m_h$.

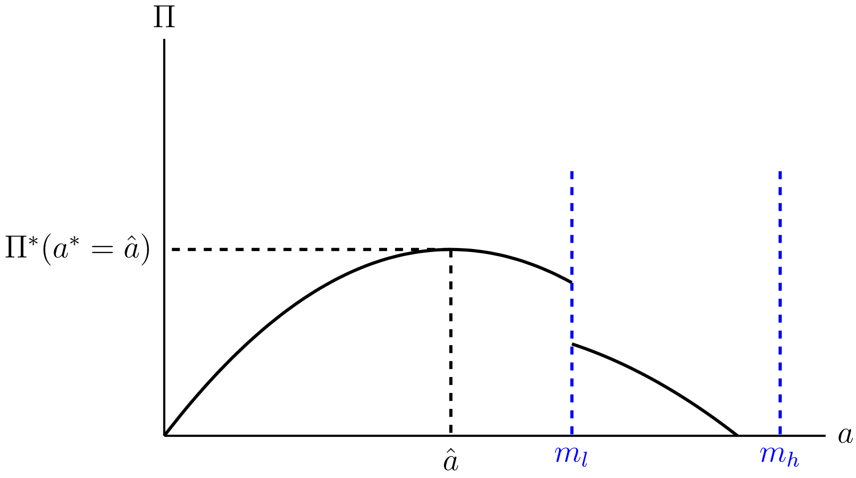

Figure 3.1: Platform's profits and choice of advertising level when $1 + \gamma_l(\hat{a}) \geq \theta$

However, advertising avoidance may be less costly. Indeed, $\hat{a}$ may be greater than the maximum levels of advertising m_l and/or m_h that consumers γ_l and γ_h are willing to accept to visit the website, forcing the platform to lower its advertising level to prevent Internet users from avoiding ads. For example, when $\min(\hat{a}, m_l) = m_l$ (which corresponds to $1 + \gamma_l(\hat{a}) < \theta$), if the platform sets $a^* = \hat{a}$, it cannot reach consumers who are strongly ad-sensitive. Two main cases are thus possible.

In the first case, when $\hat{a} < m_h$ (which corresponds to $\theta \leq 1 + \gamma_h(\hat{a})$), the platform can only reach consumers who have a high taste for ads when setting $a^* = \hat{a}$. As a consequence, the platform has the choice between setting $a^* = m_l$ and serving all the users ($N = 1$), or setting $a^* = \hat{a}$ and reaching only high taste consumers ($N = 1 - \beta$). We refer to this case in the sequel as a case where advertising avoidance is *relatively costly*. In this case, high taste users remain *ad-lovers* as they always get positive utility from advertising, regardless the choice of the platform.

Conversely, users who have a low taste for ads are *ad-neutral* when $a^* = m_l$, as they receive similar benefit from advertising than with AATs, and *ad-averse* when $a^* = \hat{a}$, as they all prefer AATs and not see ads. Figure 3.2 (left) illustrates the case where both types of consumers visit the website and the profits of the platform are higher for $a^* = m_l$. However, setting $a^* = \hat{a}$ as it is displayed in Figure 3.2 (right) would be preferable for the platform with the low taste users on board.

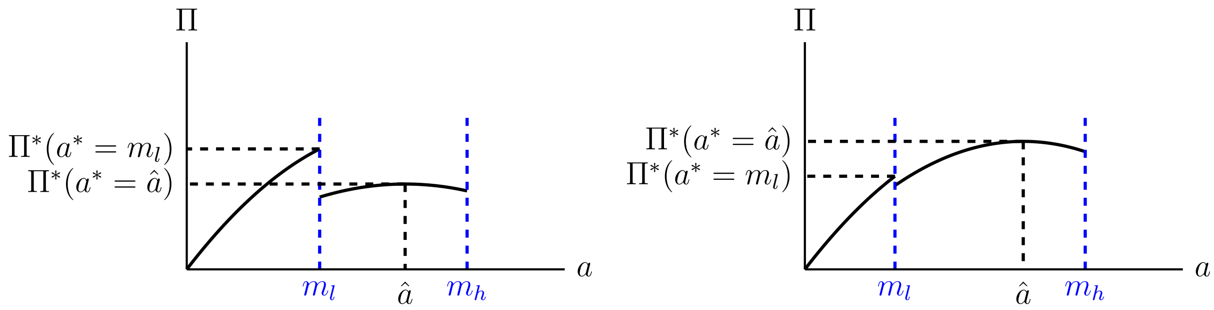

Figure 3.2: Platform's profits and choice of advertising level when $1 + \gamma_l(\hat{a}) < \theta \leq 1 + \gamma_h(\hat{a})$

In the second case, when $\hat{a} > m_h$ (which corresponds to $1 + \gamma_h(\hat{a}) < \theta$), both types of consumers are not interested in visiting the website when setting $a^* = \hat{a}$, as they get a higher utility in avoiding ads. The platform has the choice between setting $a^* = m_l$, and reach both types of consumers ($N = 1$), or setting $a^* = m_h$ and serve again only high taste consumers ($N = 1 - \beta$). By reference to the previous strategy, we refer to this case as when avoiding ads is *cheap*. In this case, users who have a high taste for ads are now ad-neutral as they are indifferent between seeing ads or using AATs. However, users with a low taste for ads are ad-averse as they all prefer AATs and not see ads. Moreover, if the strategy of the platform is $a^* = m_l$, it follows the same logic as in the previous case: high taste users are ad-lovers while low taste users are ad-neutral. The choice faced by the platform is illustrated in Figure 3.3. In Figure 3.3 (left), setting a level of advertising $a^* = m_h$ to target users who have a high taste for ads is less profitable for the platform than setting $a^* = m_l$, as a lower number of advertisers create incentives both to watch ads for both types of consumers.

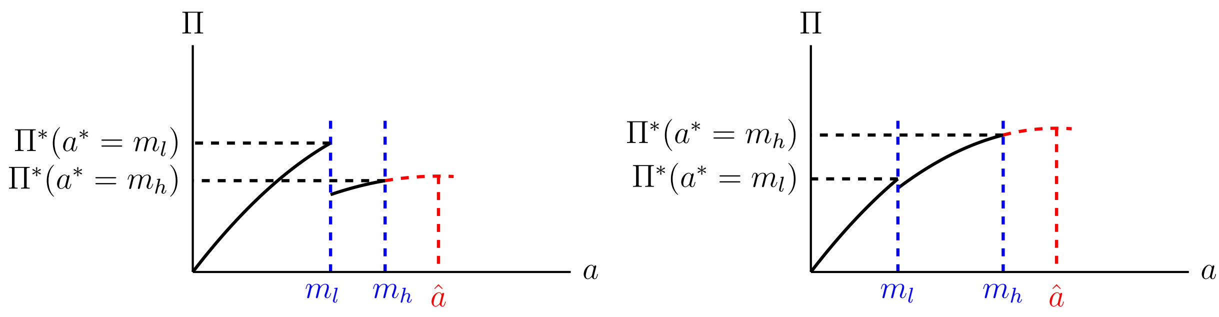

Figure 3.3: Platform's profits and choice of advertising level when $1 + \gamma_h(\hat{a}) < \theta$

To summarize, when advertising avoidance is relatively costly or even cheaper on the level of advertising, the platform can choose to serve either both types of users or only users who have a high taste for ads. The profits of the platform are in these cases:

$$\Pi^* = \begin{cases} R_l, & \text{if both types of consumers visit the website,} \\ \\ (1-\beta)R_h, & \text{if the only consumers with a high taste for ads visit the website,} \end{cases} \tag{3.6}$$

where $R_l = R(a^* = m_l)$. R_h represents either the revenues from setting the optimal level of ads $\hat{a}$ (implying $R_h = R(a^* = \hat{a})$) when users with a high taste for ads have a higher utility when viewing ads rather than avoiding ads ($\hat{a} < m_h$), or the revenues from setting the optimal level of ads m_h (implying $R_h = R(a^* = m_h)$), when the low taste users have a higher utility when avoiding ads rather than viewing advertising ($\hat{a} > m_h$). The revenue per user when the platform chooses to focus only on consumers who have a high taste for ads can therefore be written as $R_h = R(a^* = \min(\hat{a}, m_h))$. To keep the notations as simple as possible, we will use R_l and R_h in the rest of the chapter.

Proposition 2 summarizes the different cases regarding the choices of advertising levels.

Proposition 2: When avoiding ads is *relatively costly* ($\hat{a} > m_l$), two cases are possible. First, when $\frac{R_l}{R_h} > 1 - \beta$, the platform chooses the level of advertising $a^* = m_l$, and displays ads to users γ_l and γ_h. Second, when $\frac{R_l}{R_h} < 1 - \beta$, the platform chooses $a^* = \min(\hat{a}, m_h)$ and only displays ads to consumers γ_h. See Proof of Proposition 2 in Appendix 3.8.

Proposition 2 states that when the proportion of low taste users is high in the population ($\frac{R_l}{R_h} > 1 - \beta$), the platform may find profitable to set the level of advertising so as to attract not only low taste users but also high taste ones. By contrast, when the proportion of users who have a low taste for ads is low ($\frac{R_l}{R_h} < 1 - \beta$), it is most profitable to only target users with a high taste for advertising..

3.4 Equilibrium with Profiling Technology

In the baseline case, the platform does not have a profiling technology to find out the consumer preferences for advertising, and sets an equal level of advertising for both types of consumers. In this section, the platform can now use a profiling technology to tailor a level of advertising to each type of consumer γ_l and γ_h, and prevent consumers from adopting AATs.

The profiling technology exploits programmatic techniques, which are methods for buying and selling online ads in real time. Real time means that the entire process takes only a few milliseconds to complete, before a web page is loaded by a consumer. Typically, a consumer visits a website. Once connected, the technology gathers personal information about users with the help of cookies, and produces a signal that can be of two types, s_h or s_l.[11] A signal s_l (resp. s_h) means that the technology correctly classifies a user of type γ_l (resp. γ_h) with probability δ. Without loss of generality, we assume that $\delta \in [\frac{1}{2}, 1]$; $\delta = 1$ means that the technology is perfectly efficient and always classifies correctly Internet users, and $\delta = \frac{1}{2}$ indicates that the technology brings no information, and correctly classifies Internet users with probability $\frac{1}{2}$.

Once the classification has been made, the platform is able to display the content and the level of advertising adapted to the probable type of consumer: for example, the level of advertising ${a_l}^T$ will be displayed to consumers classified as having a low taste for ads (γ_l), and the level of advertising ${a_h}^T$ to the ones classified as consumers with a high taste for ads (γ_h). If the profiling technology is efficient, i.e. if $\delta = 1$, the platform perfectly discriminates the types of consumers, and tailors the level of advertising ${a_l}^T$ and ${a_h}^T$ to users γ_l and γ_h, who do not use in turn AATs. On the other hand, if the technology does not perfectly identify the types of users $\delta \in [\frac{1}{2}, 1[$, consumers γ_l and γ_h may see a level of advertising unsuitable to their advertising preferences, and adopt in reaction AATs.

The use of a profiling technology does not change the timing structure for the order of moves except that the platform now choses the number of advertisers a in a context where the profiling

[11] The profiling technology uses machine-learning methods to construct signals.

technology is effective in stage 2. In stage 1, the platform sets the advertising levels ${a_l}^T$ and ${a_h}^T$ to be displayed to consumers classified as γ_l and γ_h. In stage 2, consumers choose to visit the website, and decide to avoid ads if the level of advertising is not adapted. We solve the game by backward induction.

3.4.1 Stage 2

Following the baseline case, the users decide to view ads or to avoid them using AATs. Depending on their taste, an internet user of type $i \in \{l, h\}$ exhibit the following utility function:

$$U_i = \begin{cases} 1 + \gamma_i({a_i}^T), \text{ if choosing to visit with ad level } {a_i}^T, \\ 1 + \gamma_i({a_{-i}}^T), \text{ if choosing to visit with ad level } {a_{-i}}^T, \\ \theta, \text{ if choosing to visit and avoid ads,} \end{cases} \tag{3.7}$$

3.4.2 Stage 1

Description of the profiling technology

Following the Bayes' rule, we can calculate the expected profits of the platform. The probability to receive the signal s_l knowing that the Internet user is of type γ_l is equal to δ, and can be written as $\mathbb{P}(s_l|\gamma_l) = \delta$. The signal s_l is received with probability $\mathbb{P}(s_l) = \beta\delta + (1-\beta)(1-\delta)$ and s_h with probability $\mathbb{P}(s_h) = \delta(1-\beta) + \beta(1-\delta)$. Upon receiving the signal s, the platform then knows that this signal is true with probability:

$$\mathbb{P}(\gamma_l|s_l) = \frac{\mathbb{P}(s_l|\gamma_l)\mathbb{P}(\gamma_l)}{\mathbb{P}(s_l)} = \frac{\delta\beta}{\delta\beta + (1-\delta)(1-\beta)},$$

and

$$\mathbb{P}(\gamma_h|s_h) = \frac{\mathbb{P}(s_h|\gamma_h)\mathbb{P}(\gamma_h)}{\mathbb{P}(s_h)} = \frac{\delta(1-\beta)}{\delta(1-\beta) + (1-\delta)\beta}.$$ [12]

[12] By definition, $\mathbb{P}(\gamma_h|s_l) \equiv 1 - \mathbb{P}(\gamma_l|s_l)$, and $\mathbb{P}(\gamma_l|s_h) \equiv 1 - \mathbb{P}(\gamma_h|s_h)$.

Applying the Bayes rule, the expected profits of the platform are:

$$\begin{aligned}\max_{a_l{}^T, a_h{}^T} \mathbb{E}\left(\Pi^T\right) &= \mathbb{P}(s_l)\left(\mathbb{P}(\gamma_h|s_l) + \mathbb{P}(\gamma_l|s_l)\right) R(a_l{}^T) + \mathbb{P}(s_h)\left(\mathbb{P}(\gamma_h|s_h) + \mathbb{P}(\gamma_l|s_h)\right) R(a_h{}^T) \\ &= [\beta\delta + (1-\beta)(1-\delta)]R(a_l{}^T) + [\beta(1-\delta) + \delta(1-\beta)]R(a_h{}^T)\end{aligned} \tag{3.8}$$

Advertising choice of the platform with profiling technology

Similarly to the baseline case, when avoiding ads is expensive, that is when seeing advertising provide utility for both types of users ($\hat{a} < m_l$), the optimal amount of ads targeted for consumers γ_l and γ_h is exactly the same. The profiling technology is not therefore useful in this case. This result is summarized in Proposition 3.

Proposition 3: When avoiding ads is *expensive* ($\hat{a} < m_l$), the platform always sets its interior level of advertising which is equal for all users, i.e. $a_l{}^{T*} = a_h{}^{T*} = \hat{a}$. The profiling technology is not useful in this case.

Proposition 3 is equivalent to Proposition 1, and the same intuitions apply. A more interesting case arises when advertising avoidance is affordable, which drives the platform to be constrained. In this case, the profiling technology can be helpful as it discriminates consumers in setting two distinct levels of advertising.

Focusing on the case where the publisher is constrained $\hat{a} > m_l$, we notice that expected profits in Eq. (3.8) are made of four terms: two terms related to the probability of receiving the signal s_l, and two terms related to the probability of receiving the signal s_h. First, upon receiving the signal s_l, the platform will set the level of advertising $a_h{}^{T*} = m_l$ to attract users γ_l. The resulting profits are also composed of two terms. The first term, $\beta\delta R(a_h{}^{T^*})$, is the revenue related to the successful classification of users γ_l, and the second term, $(1-\beta)(1-\delta)R(a_l{}^{T*})$, is the revenue from the wrong classification of users γ_h. In both cases, users γ_l and γ_h will choose to see ads without AATs.

Second, upon receiving the signal s_h, the platform will set the level of advertising $a_l{}^{T*} = \min(\hat{a}, m_h)$ to attract users γ_h. The resulting profits are also come from the proportion of successfully classified users γ_h, $\delta(1-\beta)R(a_l{}^{T*})$, and the wrong classification of users γ_l, $\beta(1-\delta)R(a_l{}^{T*})$. In this case, only users γ_h will choose to see ads without AATs. Indeed, users γ_l receive a higher utility when using AATs than from seeing ads.

Lemma 1: When avoiding ads is *relatively costly* ($\hat{a} > m_l$), the use of the profiling technology allows the platform to discriminate consumers γ_l and γ_h: a proportion $\delta\beta + (1-\delta)(1-\beta)$ of users classified as γ_l sees a level of advertising ${a_l}^{T*} = m_l$, while a proportion $\delta(1-\beta)+(1-\delta)\beta$ classified as γ_h sees a level of advertising ${a_h}^{T*} = \min(\hat{a}, m_h)$.

In setting the optimal level of advertising, the strategy of the platform will depend on the probability δ to correctly classify users γ_l and γ_h. The expected profits of the platform at the equilibrium can be written as:

$$\mathbb{E}\left(\Pi^T\right)^* = \delta\beta R_l + (1-\beta)(1-\delta)R_l + \delta(1-\beta)R_h.^{13} \tag{3.9}$$

Technological choice of the platform

Given the efficiency δ of the profiling technology, the platform may or may not adopt it. We therefore compare the platform's profits with and without profiling technology to underline the condition of adoption of the profiling technology by the platform.

When $\hat{a} < m_l$, Proposition 3 already demonstrates that the platform has no use for the profiling technology.

Conversely, when $\hat{a} > m_l$, comparing profits in Eq. (3.9) to Eq. (3.6) allows us to obtain the minimum efficiency value δ for which it is always profitable for the platform to use the profiling technology.

Proposition 4: When avoiding ads is *relatively costly* ($\hat{a} > m_l$), the platform always uses the technology if $\frac{R_l}{R_h} \geq 1-\beta$ and $\delta \in [\delta_l, 1]$, or $\frac{R_l}{R_h} \leq 1-\beta$ and $\delta \in [\delta_h, 1]$, with $\delta_l = \frac{\beta R_l}{R_l\beta+(1-\beta)(R_h-R_l)}$ and $\delta_h = \frac{(1-\beta)(R_h-R_l)}{R_l\beta+(1-\beta)(R_h-R_l)}$. See Proof of Proposition 4 in Appendix 3.8.

Proposition 4 states that the platform will use the profiling technology when the probability of correctly classifying the types of users is sufficiently high, i.e. when $\delta > \max(\delta_l, \delta_h)$.[14] Indeed, when the profiling technology is deficient ($\delta < \max(\delta_l, \delta_h)$), the platform does not properly classify a part of Internet users, resulting into lower profits. For example, when users γ_l are classified as γ_h, their utility decreases and they choose to avoid ads by adopting AATs, which induces in a drop in profits for the platform.

[13] We recall that $R_l = R(m_l)$ and $R_h = R(\min(\hat{a}, m_h))$.

[14] When $\frac{R_l}{R_h} \leq 1-\beta$ we have $\delta_h \leq \frac{1}{2} \leq \delta_l \leq 1$, and when $\frac{R_l}{R_h} \geq 1-\beta$ we have $\delta_l \leq \frac{1}{2} \leq \delta_h \leq 1$. Therefore, we know from Proposition 4 that the platform will use the profiling technology if $\delta > \max(\delta_l, \delta_h)$.

3.5 Some considerations about the volume of ads in equilibrium

The use of a profiling technology affects the level of ads served to consumers with respect to a situation where the platform does not use a profiling technology. In this section, we determine more precisely how the number of ads served at equilibrium changes when the platform uses the profiling technology. Do consumers see more ads when a platform uses a profiling technology?

We restrict the analysis to the cases where the platform may use the technology.[15] Lemma 2 presents two different cases.

Lemma 2: When avoiding ads is *relatively costly* ($\hat{a} > m_l$) and the platform chooses to use the technology ($\delta > \max(\delta_h, \delta_l)$), it affects Internet users in two ways.

- Firstly, when $\frac{R_l}{R_h} \geq 1 - \beta$, a proportion $\delta\beta + (1-\delta)(1-\beta)$ of consumers sees the same level of ads with or without technology, whereas a proportion $\delta(1-\beta) + (1-\delta)\beta$ of consumers sees more ads. The level of ads increases so much for the proportion $(1-\delta)\beta$ of consumers that they choose to adopt AATs when they would choose to see ads without technology.

- Secondly, when $\frac{R_l}{R_h} < 1 - \beta$, a proportion $\delta(1-\beta) + (1-\delta)\beta$ of consumers sees the same level of ads as without technology, whereas a proportion $\delta\beta + (1-\delta)(1-\beta)$ sees less ads. The level of ads decreases so much for the proportion $\delta\beta$ of consumers that they prefer to see ads when they would prefer to avoid them without technology.

Lemma 2 shows that introducing a profiling technology greatly changes the level of ads consumers may see with respect to a situation without profiling technology.

Firstly, Lemma 2 shows that if the platform does not use a profiling technology and $\frac{R_l}{R_h} \geq 1 - \beta$, it will set a low level of ads $a^* = m_l$ and will attract all Internet users. However, when using the profiling technology, the platform sets two levels of advertising ${a_l}^{T*} = m_l$ and ${a_h}^{T*} = \min(\hat{a}, m_h)$ when receiving the signals s_l and s_h. The profiling technology does not however always correctly classify consumers, which may encourage some to adopt AATs. This

[15] We indeed established that when all Internet users are not ad-sensitive ($\min(\hat{a}, m_l) = \hat{a}$), the platform does not use the technology.

is exactly the case when the platform tailors a level of ads ${a_l}^{T*}$ for users γ_l but classifies some of them as γ_h. In this latter case, it is preferable for the proportion $(1-\delta)\beta$ of these users to adopt AATs. The number of consumers who chooses to see ads therefore decreases with respect to the baseline case.

Second, Lemma 2 also shows that if the platform does not use a profiling technology and $\frac{R_l}{R_h} < 1-\beta$, it may want to set a higher level of advertising $a^* = \min(\hat{a}, m_h)$ to get higher profits in attracting less Internet users γ_h. In the case with profiling technology, both levels of advertising ${a_l}^T$ and ${a_h}^T$ are still available. As the platform sets ${a_l}^{T*} = m_l$ when receiving s_l, the level of advertising decreases for respective proportions $\delta\beta$ and $(1-\delta)(1-\beta)$ of users correctly and incorrectly classified as γ_l. In this case, the proportion of users $\delta\beta$ correctly identified as γ_l chooses to see ads when they would choose to adopt AATs without profiling technology. The use of a profiling technology allows the platform in this case to increase the number of consumers who choose to see ads with respect to the baseline case.

To evaluate in the end whether the levels of ads with profiling technology ${a_l}^{T*}$ and ${a_h}^{T*}$ differ from the one without technology ${a_l}^{T*}$, we need to calculate how many ads are served to consumers. Indeed, we cannot simply compare the levels of ads with or without technology, as different proportions of consumers can be targeted by two distinct levels of advertising. Basically, as an illustration, we have to compare a case without profiling technology where the platform sells 10 ads to reach 100 consumers to a case where the platform can use a profiling technology and sell on the one hand 8 ads to reach 75 consumers γ_l, and 10 ads to reach 25 consumers γ_h on the other.

We denote by $V = Na^*$ and $V^T = N_{|s=s_l}{a_l}^{T*} + N_{|s=s_h}{a_h}^{T*}$, the total number of ads served by the platform without and with profiling technology. We have:

$$V = \begin{cases} \hat{a} \text{ if } \hat{a} < m_l, \\ \\ m_l \text{ if } \hat{a} > m_l \text{ and } \frac{R_l}{R_h} \geq 1-\beta, \\ \\ (1-\beta)\min(\hat{a}, m_h) \text{ if } \hat{a} > m_l \text{ and } \frac{R_l}{R_h} < 1-\beta, \end{cases}$$

and,

$$V^T = \begin{cases} V \text{ if } \hat{a} < m_l \text{ or } \hat{a} > m_l \text{ and } \delta < \max(\delta_h, \delta_l), \\ \\ m_l(\delta\beta + (1-\beta)(1-\delta)) + \delta(1-\beta)\min(\hat{a}, m_h) \text{ if } \hat{a} > m_l \text{ and } \delta > \max(\delta_h, \delta_l). \end{cases}$$

When online users are not very ad-sensitive, the platform will not use the technology as all users still visit the website when setting its favorite level of advertising $a^* = \hat{a}$. When at least a proportion of users have a low taste for ads ($\min(\hat{a}, m_l) = m_l$), it is clear that when the profiling technology does not correctly classify users, i.e. when the technology is deficient ($\delta < \max(\delta_h, \delta_l)$), the platform does not use it, and the situation remains unchanged: $V^T = V$. However, when the technology is efficient ($\delta > \max(\delta_l, \delta_h)$), the platform uses the profiling technology to tailor the level of ads to both types of consumers. This situation has to be compared with the two possible strategies available for the platform when it does not have a profiling technology, i.e. when $V = m_l$ or $V = (1-\beta)\min(\hat{a}, m_h)$. We establish the following proposition:

Proposition 5: When avoiding ads is *relatively costly* ($\hat{a} > m_l$) and the platform chooses to use the technology ($\delta > \max(\delta_h, \delta_l)$), the total number of ads served by the platform to Internet users is higher $V^T > V$ when $\frac{R_l}{R_h} \geq 1-\beta$. If $\frac{R_l}{R_h} < 1-\beta$, $V^T > V$ if $\delta \in [\delta_h^V, 1]$, with $\delta_h^V = \frac{(1-\beta)(\min(\hat{a}, m_h) - m_l)}{(1-\beta)(\min(\hat{a}, m_h) - m_l) + \beta m_l}$.

See Proof of Proposition 5 in Appendix 3.8.

Proposition 5 shows that the platform will serve more ads in total when $\frac{R_l}{R_h} \geq 1-\beta$, while it is not always the case when $\frac{R_l}{R_h} < 1-\beta$.

Indeed, according to Proposition 2, introducing a technology when $\frac{R_l}{R_h} \geq 1-\beta$ decreases the number of Internet users seeing as. On the one hand, the platform attracts all Internet users without technology, while it misclassifies a proportion of low taste users when using one. On the other hand, the technology allows the platform to set a higher level of ads to a proportion of correctly classified high taste users. All in all, we find that the higher level of ads shown to correctly classified users with a high taste for ads always offsets the decrease in the number of Internet users seeing advertising.

Conversely, according to Proposition 2, introducing a technology when $\frac{R_l}{R_h} < 1-\beta$ allows the platform to attract more users. On the one hand, correctly classified users with a low taste for ads choose not to avoid advertising and see a low level of ads. On the other hand, misclassified low taste users see less ads than before. All in all, we show that the higher number of Internet users seeing ads offsets the lower level of ads seen by misclassified user who have a high taste for ads only if the profiling technology does not misclassify too much.

3.6 Winners and Losers of the Profiling Technology

We analyze how the introduction of a profiling technology may affect consumers, advertisers, and the platform. In doing that, we define total welfare as the sum of the platform profits (Π), the advertisers surplus (S_a) and the users surplus (S_u). We compute the total welfare without technology (W) and the total welfare with the profiling technology (W^T) when it is good enough to be used by the platform (i.e. when $\delta > \max(\delta_l, \delta_h)$).

3.6.1 Internet user surplus

Two cases are interesting to analyze when $\hat{a} > m_l$.[16]

Lemma 3: When avoiding ads is *relatively costly* ($\hat{a} > m_l$) and the platform chooses to use the technology ($\delta > \max(\delta_h, \delta_l)$):

- If $\frac{R_l}{R_h} > (1-\beta)$, $\mathrm{S}_u^T > \mathrm{S}_u$ if $\gamma_h(\min(\hat{a}, m_h)) > \gamma_h(m_l)$.
- If $\frac{R_l}{R_h} < (1-\beta)$, $\mathrm{S}_u^T > \mathrm{S}_u$ if $\delta \neq 1$ and $\gamma_h(\min(\hat{a}, m_h)) < \gamma_h(m_l)$.

See Proof of Lemma 3 in Appendix 3.8.

Lemma 3 underlines two important results. Firstly, the impact of the profiling technology on users' surplus largely depends on the utilities associated to a correct or a misclassification. For example, an Internet users could tolerate a high advertising intensity, but may prefer a lower one, hence gaining more utility from being misclassified. Conversely, a user may want to receive more ads on a specific subject, hence generating a higher utility from being correctly classifed. Secondly, it also depends on the structure of the market, as the impact of introducing a technology depends on the advertising level set by the platform without technology.

Indeed, when $\frac{R_l}{R_h} > (1-\beta)$, introducing a profiling technology will increase consumer surplus only if users with a high taste for ads prefer to be correctly classified ($\gamma_h(\min(\hat{a}, m_h)) >$

[16]The consumer surplus will not be affected by the introduction of technology when $\hat{a} = \hat{a} < m_l$. Indeed, Proposition 3 shows that the technology will not be used by the platform in this case.

$\gamma_h(m_l)$). [17] This is due to the fact that the platform attract all Internet users with a low level of advertising, hence showing few ads to users who have a high taste for ads. Hence, using the technology allows the platform to show more ads to high taste users. The impact of the technology on user surplus therefore depends on whether those users prefer to see more ads (being correctly classified) or fewer ads (being misclassified).

The inverse situation arises when $\frac{R_l}{R_h} < (1-\beta)$ as introducing an imperfect profiling technology increases consumer surplus only if users who have a high taste for ads prefer to be misclassified ($\gamma_h(m_l) > \gamma_h(\min(\hat{a}, m_h))$). In this situation, the platform having no profiling technology chooses to strip out high taste users from their benefit while excluding users who have a low taste for ads. Hence, introducing a technology would reduce the level of ads shown to misclassified users with a high taste for ads. Consequently, user surplus increase with a profiling technology if users who have a high taste for ads prefers a lower level of ads (being misclassified) than a high level of ads (being correctly classified).

3.6.2 Advertiser surplus

Like the consumer surplus analysis, we only analyze the situation encountered when advertising avoidance is affordable ($\hat{a} > m_l$). More generally, introducing a perfect profiling technology will increase advertisers' surplus. However, an imperfect technology may have different effect as the platform have different practice when it has no profiling technology.

Lemma 4: When avoiding ads is *relatively costly* ($\hat{a} > m_l$) and the platform chooses to use the technology ($\delta > \max(\delta_h, \delta_l)$), advertisers always benefit from it $\mathrm{S}_a^T > \mathrm{S}_a$ if $\frac{R_l}{R_h} \geq 1-\beta$. If $\frac{R_l}{R_h} \leq 1-\beta$, advertisers benefit from the technology if $\delta \in [\delta_h^{s^a}, 1]$, with $\delta_h^{s^a} = \frac{(1-\beta)(s_h^a - s_l^a)}{s_l^a \beta + (1-\beta)(s_h^a - s_l^a)}$. See Proof of Lemma 4 in Appendix 3.8.

Lemma 4 highlights important results for advertisers. When $\frac{R_l}{R_h} \geq (1-\beta)$, the advertisers will gain from the profiling technology as long as it is beneficial for the platform to introduce it. Conversely, if $\frac{R_l}{R_h} < (1-\beta)$, there is an area on δ where introducing the technology is profitable for the platform but lower the advertisers surplus.

[17] As $\gamma_h(m_l) > \gamma_h(m_h)$, the profiling technology can only increase consumer surplus when advertising avoidance is affordable and $\gamma_h(m_l) < \gamma_h(\hat{a})$. This is intuitive as a correct identification in this case generates higher utility than a misclassification for Internet users who have a high taste for ads .

In this case, the platform may choose to adopt the technology even if it lowers advertisers' surplus, hence making their situation worse than without technology. This is due to the fact that an efficient profiling technology allows the platform for a better detection of low taste users (which require a lower level of ad), hence decreasing advertisers' surplus. This result exerts a strong impact on the computation of total welfare.

3.6.3 Profits of the platform

The analysis follows the same mechanisms as in Proposition 4. The platform choose to use the technology only if it generates high enough profits, i.e. if the technology is good enough. Therefore the condition for the technology to be profit increasing are the same than in Proposition 4.

3.6.4 Total welfare

We simplify the analysis in writing the industry profits as $\underline{G} = R_h + s_h^a$ and $\overline{G} = R_l + s_l^a$ and the incentives to be correctly identified for consumers who have a high taste for ads as $I^{\gamma_h} = \gamma_h(\min(\hat{a}, m_h)) - \gamma_h(m_l)$.

Proposition 6: When avoiding ads is *relatively costly* ($\hat{a} > m_l$) and the platform chooses to use the technology ($\delta > \max(\delta_h, \delta_l)$):

- If $\frac{R_l}{R_h} > (1-\beta)$, $W^T > W$ if $\delta(\beta\overline{G} + (1-\beta)(\underline{G} - \overline{G}) + (1-\beta)I^{\gamma_h}) > \beta\overline{G}$
- If $\frac{R_l}{R_h} < (1-\beta)$, $W^T > W$ if $\delta(\beta\overline{G} + (1-\beta)(\underline{G} - \overline{G}) + (1-\beta)I^{\gamma_h}) > (1-\beta)(\underline{G} - \overline{G}) + (1-\beta)I^{\gamma_h}$.See Proof of Lemma 6 in Appendix 3.8.

Proposition 6 underlines different situations. We analyze only the cases where the technology is efficient enough and have an impact on market equilibria (that is when advertising avoidance is affordable such that the platform is constrained by the nuisance of Internet users when setting its level of advertising).

A first situation arises when $\frac{R_l}{R_h} > (1-\beta)$. In that case, the industry profits - that is the platform profits and the advertisers surplus - always benefit from the introduction of an efficient technology. Two cases can arise. Firstly, if high taste users generate more utility from being correctly identified ($\gamma_h(\min(\hat{a}, m_h)) > \gamma_h(m_l)$), the introduction of technology will be welfare increasing as it improves the situation of all agents in the market. Secondly, if Internet users with a high taste for ads generate more utility from being misclassified as having a low taste for ads ($\gamma_h(\min(\hat{a}, m_h)) < \gamma_h(m_l)$), the introduction of technology will be welfare increasing if the opportunity cost of not being misclassified is not too high ($\gamma_h(\min(\hat{a}, m_h)$ close to $\gamma_h(m_l)$).

A second situation arises when $\frac{R_l}{R_h} < (1-\beta)$. In this case, the industry profit may not increase with the introduction of the technology, as platform may use the technology while it decreases advertisers' surplus. More precisely, three different situations may occur. Firstly, if the introduction of an efficient technology decreases industry surplus [18] and Internet users who have a high taste for ads earn higher utility from being correctly classified ($\gamma_h(\min(\hat{a}, m_h)) > \gamma_h(m_l)$), it will be welfare decreasing as both industry profits and Internet users surplus are lower with technology than without. Secondly, if the introduction of an efficient technology increases industry surplus, and high taste ads users earn higher utility from being misclassified ($\gamma_h(\hat{a}, m_h)) < \gamma_h(m_l)$), everyone benefits from the technology. Finally, if the technology decreases industry surplus and increases Internet user surplus or increases industry surplus and decreases Internet user surplus, the introduction of such technology may have an ambiguous effect on welfare. Results are summed up in Figure 3.4.

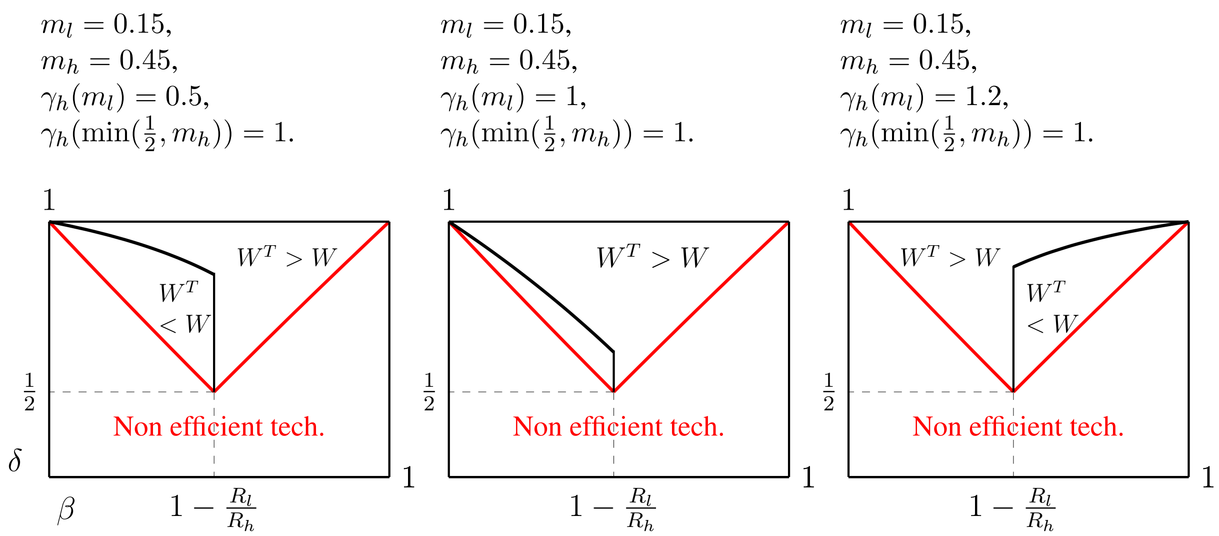

Figure 3.4: Welfare impact of technology (with $r(a) = 1 - a$) with respect to δ(technology efficiency) and β (low taste users)

[18]The technology is efficient enough to be used by the platform but sharply decreases advertisers surplus ($\underline{\delta_{s^a}} > \delta > \underline{\delta}$)

3.7 Strategic implications and conclusion

In this part, we summarize the key results and implications of the model, and concludes.

We observe that the main findings depend on whether attracting users who have a low taste for ads is profitable or not for the platform. We recall from Proposition 2 that when there is no profiling technology, the publisher chooses to set a low level of advertising and attract all Internet users if $\frac{R_l}{R_h} > 1 - \beta$ or set a high level of advertising and only attract the least sensitive Internet users if $\frac{R_l}{R_h} < 1 - \beta$.

Let assume for simplicity that the marginal revenue function $R(a)$ of the publisher is fixed, as well as the maximum advertising level users of each population are willing tolerate m_l and m_h. [19] In this case, the incentives of the platform to attract or not the entire audience depends on the proportion β of users with a low taste for ads visiting its website. Firstly, we note that when there is a high proportion of low taste user (high β), the publisher is more likely to attract all Internet users with a low level of advertising. When this is the case, we consider that the platform is facing an "expensive audience": an audience that requires to lower the number of ads to attract them, hence decreasing revenue by user. The audience is expensive in a sense that it must give up on high profits on the advertiser side to attract users. Conversely, when there is a dominant proportion of Internet users with a high taste (low β), the publisher will be more likely to set a high level of advertising and attract only high taste Internet users. In that case, the platform is facing a "cheap audience", in the sense that the audience can be attracted without renouncing to high profits on the advertisers side.

From above, we consider that the platform can encounter two types of audience: cheap or expensive. For example, illegal streaming or downloading websites such as The Pirate Bay or DpStream generally face a cheap audience, and choose to set a high level of advertising. Conversely, News websites such as The New York Times, The Washington Post or The Financial Times address a more expensive audience which is significantly more elastic to advertising levels, hence fostering them to lower the level of advertising.

The platform sets its advertising level depending on the type of audience it faces, knowing its marginal revenue function $R(a)$. Considering this typology of audiences, we are able to derive key strategic implications:

[19] which defines R_l and R_h as $R(a^* = m_l) = R_l$ and $R(a^* = m_h) = R_h$

Key Result 1:

- On a platform facing a cheap audience, the introduction of a profiling technology reduces advertising avoidance.

- Conversely, on a platform facing an expensive audience, the introduction of a profiling technology increases advertising avoidance.

This result directly stems from Proposition 2. Without profiling technology, a platform facing a cheap audience choose to focus on attracting the fringe of Internet users that is the less elastic to ads (i.e. users who have a high taste for ads). Hence, introducing a profiling technology allows the platform to attract the small proportion of users who have a low taste for ads in offering a better Internet user experience. Hence, introducing a profiling technology on DpStream would improve the user experience of a part of the audience.

Conversely, without profiling technology, a platform facing an expensive audience chooses to attract the entire audience of Internet users. Therefore, introducing a profiling technology may degrade the experience of low taste users who are misclassified, and therefore foster them to use AATs instead of watching ads. As an example, introducing a profiling technology on The New York Times would deteriorate the user experience of a part of the audience.

A subsequent analysis can be drawn on the impact of the profiling technology on welfare. Proposition 6 determines the conditions for which the introduction of such technologies is welfare increasing. Mapping the propositions to our analysis underline the following key results:

Key Result 2:

- On a platform facing a cheap audience, the introduction of an efficient profiling technology is welfare increasing if Internet users benefits more from being misclassified than correctly classified and advertisers benefit from the technology.

- Conversely, on a platform facing an expensive audience, the introduction of a profiling technology is always welfare increasing if Internet users benefits more from being correctly classified than misclassified.

Key result 2 highlights that the impact of the introduction of a profiling technologies depends on the audience of the website that adopts it. From the analysis on Proposition 6, we understand that if the profiling technology is highly efficient, its introduction is welfare enhancing no matter the type of websites. However, if the profiling technology is less efficient, the introduction of a profiling technology impacts welfare differently depending on the structure of its audience.

For example, introducing a profiling technology on a platform facing an expensive audience such as The Washington Post is welfare increasing if this very audience is composed of internet users that enjoy being correctly classified. Indeed, on the one hand, both advertisers and the platform benefit from the introduction of a profiling technology. On the other hand, the less elastic fringe of the audience is seeing more ads with profiling technology than without. Therefore, the impact of the profiling technology on total welfare obviously depends on whether the audience prefers to be correctly classified (seeing more ads) or not (seeing less ads).

Conversely, introducing the profiling technology on platform facing a cheap audience such as The Pirate Bay may have more ambiguous impact on total welfare. Firstly, the advertisers do not always enjoy the introduction of such technology when facing that type of audience, as the platform sets lower level of ads. Secondly, the less elastic fringe of the audience is seeing less ads with profiling technology than without in this case. Overall, the impact of the profiling technology on total welfare depends on whether the audience prefers to be correctly classified (seeing more ads) or not (seeing less ads) and if the advertisers extract benefit from the profiling technology. Therefore, the introduction of a profiling technology on a platform facing a cheap audience such as DpStream is ambiguous as it depends on the shape of Internet users utility as well as on the efficiency of the profiling technology.

More generally, the efficiency of the profiling technology relies on two important points. On the one hand, the profiling technology must use a good enough prediction technology (such as machine learning), which may represent an investment for the platform. On the other hand, the profiling technology tracks Internet user and collect their data such as their browsing history to classify them as averse or not to advertising. Thus, the fact that more and more Internet users tend to protect themselves using Privacy Enhancing Tools (PETs) may lower quality of the profiling technology. This last point is crucial for regulators. Forthcoming European regulations intent to better frame the use of personal information for business purpose, giving more market power to Internet user. For these reasons, one of the limit of the chapter is to consider the efficiency of the profiling technology as exogenous, while it is endogenous to the choice of Internet users and depends on the regulation. Future research may therefore consider the link between the willingness to share information of Internet users, the privacy regulation operating on the market and the efficiency of profiling technologies used by firms to serve advertising.

Bibliography

Anderson, Simon P. and Joshua S. Gans, "Platform Siphoning: Ad-Avoidance and Media Content," *American Economic Journal: Microeconomics*, November 2011, *3* (4), 1–34.

— **and Stephen Coate**, "Market Provision of Broadcasting: A Welfare Analysis," *The Review of Economic Studies*, 2005, *72* (4), 947–972.

Buchbinder, Niv, Moran Feldman, Arpita Ghosh, and Joseph (Seffi) Naor, "Frequency Capping in Online Advertising," in Frank Dehne, John Iacono, and Jörg-Rüdiger Sack, eds., *Algorithms and Data Structures: 12th International Symposium, WADS 2011, New York, NY, USA, August 15-17, 2011. Proceedings*, Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 147–158.

Goldfarb, Avi and Catherine Tucker, "Online Display Advertising: Targeting and Obtrusiveness," *Marketing Science*, 2011, *30* (3), 389–404.

Manchanda, Puneet, Jean-Pierre Dubé, Khim Yong Goh, and Pradeep K. Chintagunta, "The Effect of Banner Advertising on Internet Purchasing," *Journal of Marketing Research*, 2006, *43* (1), 98–108.

Shiller, Ben, Joel Waldfogel, and Johnny Ryan, "Will Ad Blocking Break the Internet?," Working Paper 23058, National Bureau of Economic Research January 2017.

Tag, J., "Paying to Remove Advertisements," *Information Economics and Policy*, 2009, *21* (4), 245–252.

3.8 Appendix

Proof of Proposition 2.

We compare the profits of the platform when it serves consumers γ_l and γ_h to the ones when it serves only consumers γ_h:

- $\theta \leq 1 + \gamma_h(\hat{a}(v))$: $\Pi(a^* = \hat{a}) = (1-\beta)R(a^* = \hat{a})$
- $\theta \geq 1 + \gamma_h(\hat{a}(v))$: $\Pi(a^* = m_h) = (1-\beta)R(a^* = m_h)$

The threshold $\frac{R_l}{R_h} > 1-\beta$ results directly from the comparison.

Proof of Proposition 4.

Two cases have to be analyzed.

- First case: $\frac{R_l}{R_h} > 1-\beta$.

  Two strategies are available, which gives the following profits:

$$\Pi^* = \begin{cases} \mathbb{P}(s_l)R_l + \mathbb{P}(s_h)\mathbb{P}(\gamma_h|s_h)R_h \text{ if the platform follows the signal,} \\ \\ R_l \text{ else.} \end{cases}$$

  Comparing both profits leads directly to the rule that following the signal is more profitable for $\delta \in [\frac{R_l\beta}{R_l\beta+(1-\beta)(R_h-R_l)}, 1]$. Else, ignoring it is more profitable.

- Second case: $\frac{R_l}{R_h} \leq 1-\beta$.

  Again two strategies are available, either following the signal or not:

$$\Pi^* = \begin{cases} \mathbb{P}(s_l)R_l + \mathbb{P}(s_h)\mathbb{P}(\gamma_h|s_h)R_h \text{ if the platform follows the signal,} \\ \\ (1-\beta)R_h \text{ else.} \end{cases}$$

  Comparing both profits leads directly to the rule that following the signal is more profitable for $\delta \in [\frac{(R_h-R_l)(1-\beta)}{(R_h-R_l)(1-\beta)+R_l\beta}, 1]$. Else, ignoring it is more profitable.

It is immediate to see that both values are below 1, and a direct resolution allows us to verify that they are above $\frac{1}{2}$.

Proof of Proposition 5.

- When $\frac{R_l}{R_h} \geq 1-\beta$, $V^T > V$ if $\delta \in [\delta_l^V, 1]$.
- When $\frac{R_l}{R_h} < 1-\beta$, $V^T > V$ if $\delta \in [\delta_h^V, 1]$.

With $\delta_l^V = \frac{\beta m_l}{(1-\beta)(\min(\hat{a},m_h)-m_l)+\beta m_l}$ and $\delta_h^V = \frac{(1-\beta)(\min(\hat{a},m_h)-m_l)}{(1-\beta)(\min(\hat{a},m_h)-m_l)+\beta m_l}$.

However, we know from Proposition 4 that for the technology to be used by the platform, it has to be efficient enough such that $\delta > \max(\delta_h, \delta_l)$.

We therefore compare the conditions for which the platform would use the technology and the volume of ads may increase with profiling technology. We find that as $r(a)$ is decreasing in a, $r(m_l) > r(\min(\hat{a}, m_h))$ and therefore $\delta_l^V < \overline{\delta}$ and $\delta_h^V > \underline{\delta}$.

Hence, we are able to write that :

- When $\frac{R_l}{R_h} \geq 1-\beta$, $V^T > V$
- When $\frac{R_l}{R_h} < 1-\beta$, $V^T > V$ if $\delta \in [\delta_h^V, 1]$.

with $\delta_h^V = \frac{(1-\beta)(\min(\hat{a},m_h)-m_l)}{(1-\beta)(\min(\hat{a},m_h)-m_l)+\beta m_l}$.

Proof of Lemma 3.

$$\mathrm{S}_u = \overline{NU} + \underline{NU} = \begin{cases} \beta(1+\gamma_l(\hat{a})) + (1-\beta)(1+\gamma_h(\hat{a})) \text{ if } \hat{a} < m_l \\ \beta\theta + (1-\beta)(1+\gamma_h(m_l)) \text{ if } \hat{a} > m_l \text{ and } \frac{R_l}{R_h} > (1-\beta), \\ \beta\theta + (1-\beta)(1+\gamma_h(\min(\hat{a}, m_h))) \text{ if } \hat{a} > m_l \text{ and } \frac{R_l}{R_h} < (1-\beta), \end{cases} \tag{3.10}$$

$$\mathrm{S}_u^T = \begin{cases} \beta(1+\gamma_l(\hat{a})) + (1-\beta)(1+\gamma_h(\hat{a})) \text{ if } \hat{a} < m_l \\ \beta\theta + (1-\delta)(1-\beta)(1+\gamma_h(m_l)) + \\ \delta(1-\beta)(1+\gamma_h(\min(\hat{a}, m_h)) \text{ if } \hat{a} > m_l \end{cases} \tag{3.11}$$

Proof of Lemma 4.

$$\mathrm{S}_a = N\int_0^{a^*} (r(a) - r(a^*))\mathrm{d}a = \begin{cases} s^a(\hat{a}) \text{ if } \hat{a} < m_l \\ s^a(m_l) \equiv s_l^a \text{ if } \hat{a} > m_l \text{ and } \frac{R_l}{R_h} > (1-\beta), \\ (1-\beta)s^a(\min(\hat{a}, m_h)) \equiv (1-\beta)s_h^a \text{ if } \hat{a} > m_l \text{ and } \frac{R_l}{R_h} < (1-\beta), \end{cases} \tag{3.12}$$

$$\mathrm{S}_a^T = \begin{cases} s^a(\hat{a}) \text{ if } \hat{a} < m_l \\ \\ (\delta\beta + (1-\delta)(1-\beta))s_l^a + \delta(1-\beta)s_h^a \text{ if } \hat{a} > m_l \end{cases} \tag{3.13}$$

When avoiding ads is *relatively costly* ($\hat{a} > m_l$) and the platform, the advertisers benefit from the technology $\mathrm{S}_a^T > \mathrm{S}_a$ if $\frac{R_l}{R_h} \geq 1-\beta$ and $\delta \in [\delta_l^{s^a}, 1]$, or $\frac{R_l}{R_h} \leq 1-\beta$ and $\delta \in [\delta_h^{s^a}, 1]$, with $\delta_l^{s^a} = \frac{\beta s_l^a}{s_l^a\beta + (1-\beta)(s_h^a - s_l^a)}$ and $\delta_h^{s^a} = \frac{(1-\beta)(s_h^a - s_l^a)}{s_l^a\beta + (1-\beta)(s_h^a - s_l^a)}$.

However, we know from Proposition 4 that for the technology to be used by the platform, it has to be efficient enough such that $\delta > \max(\delta_h, \delta_l)$.

We compare the conditions that insure benefits from the technology for advertisers and the platform. We find that $\frac{R_l}{R_h} \geq 1-\beta$ (resp. $\frac{R_l}{R_h} < 1-\beta$) the condition for the platform to benefit from the profiling technology overlooks the one for advertisers i.e. $\delta_l > \overline{\delta_{s^a}}$ (resp. $\delta_h > \delta_h^{s^a}$) if $\frac{R_l}{R_h} > \frac{s_l^a}{s_h^a}$ (resp. $\frac{R_l}{R_h} < \frac{s_l^a}{s_h^a}$).

We call $G(a)$ the antiderivative of $r(a)$. We therefore can rewrite the above condition in $\frac{R_l}{R_h} > \frac{\overline{G} - R_l}{\underline{G} - R_h}$ (resp. $<$), which can be simplified to $\frac{R_l}{R_h} > \frac{\overline{G}}{\underline{G}}$. Computing the derivative of G and R with respect to a, we find that $G' = r$ and $R' = ar' + r$. As $r' < 0$, $G' > R'$, which prove the condition $\frac{R_l}{R_h} > \frac{\overline{G}}{\underline{G}}$ to be always true.

We can therefore wright that if $\frac{R_l}{R_h} \geq 1-\beta$, $\mathrm{S}_a^T > \mathrm{S}_a$ and if $\frac{R_l}{R_h} < 1-\beta$, $\mathrm{S}_a^T > \mathrm{S}_a$ only if $\delta \in [\delta_h^{s^a}, 1]$ with $\delta_h^{s^a} = \frac{(1-\beta)(s_h^a - s_l^a)}{s_l^a\beta + (1-\beta)(s_h^a - s_l^a)}$.

Proof of Proposition 6.

$$W = \begin{cases} R(\hat{a}) + s^a(\hat{a}) + 1 + \beta\gamma_l(\hat{a}) + (1-\beta)\gamma_h(\hat{a}) \text{ if } \hat{a} < m_l \\ \\ R_l + s_l^a + \beta\theta + (1-\beta)(1 + \gamma_h(m_l)) \text{ if } \hat{a} > m_l \text{ and } \frac{R_l}{R_h} > (1-\beta), \\ \\ (1-\beta)(R_h + s_h^a) + \beta\theta + \\ (1-\beta)(1 + \gamma_h(\min(\hat{a}, m_h) \text{ if } \hat{a} > m_l \text{ and } \frac{R_l}{R_h} < (1-\beta), \end{cases} \tag{3.14}$$

$$W^T = \begin{cases} R(\hat{a}) + s^a(\hat{a}) + 1 + \beta\gamma_l(\hat{a}) + (1-\beta)\gamma_h(\hat{a}) \text{ if } \hat{a} < m_l \\ \\ (\delta\beta + (1-\delta)(1-\beta))(R_l + s_l^a) + \delta(1-\beta)(R_h + s_h^a) \\ +\beta\theta + (1-\delta)(1-\beta)(1 + \gamma_h(m_l)) + (1-\beta)(1 + \gamma_h(\min(\hat{a}, m_h)) \text{ if } \hat{a} > m_l \end{cases} \tag{3.15}$$

The difference between W^T and W directly gives proposition 6.

Chapter 4

Privacy protection and online advertising market: Opt-out impact on ad prices.

4.1 Introduction

Online advertising generated an annual revenue of $41.9 billion in 2017, which is 12.3% higher than in 2016, according to this IAB report. This strong growth is mainly driven by fact the use of information on users by advertisers and websites. Indeed, the same report points out that a large share of this revenue is relatable to behavioral targeting technologies. Such techniques are very attractive to advertisers as they collect user information that enables personalization of ads. For example, according to the Washington post, Facebook allows advertisers to leverage around 100 different data points when segmenting their audience. Using basic information like user location, age, gender, or more advanced characteristics such as job title, credit card type or an ongoing disease treatment, Facebook help advertisers to tailor their campaign. The power of targeting technologies also relies on automated mechanisms for selling advertising spaces at auctions. More commonly known as "programmatic selling", these mechanisms accounted for around 80% of total ads displayed in 2017 according to EMarketer. The report suggests that 86% of ads sold through programmatic in Europe in 2017 were part of a behavioral targeting strategy. Such technologies allow advertising buyers to internalize ad spaces as well as users characteristics recovered from tracking before bidding in an auction.

As big as they are today, firms tracking practices may be perceived as intrusive by Internet users, who seem to be more and more sensitive to the use of their personal information. Report from the Pew Research Center finds that 93% of surveyees declared that being in control of *who* gets their information is important. As detailed in Tucker (2012), Internet users may exert disutility from the perceived intrusiveness of personalized ads.

Such concerns have fostered regulators to provide a framework for tracking practices, hence considering different privacy policies. The US regulator have for example favored the implementation of an *opt-out* policy, that lets advertisers and websites track by default but allows Internet user to prevent it by "opting-out" from such practice. If they do opt-out, Internet users still see ads, but these ads are not tailored based on their past behavior. Conversely, forthcoming European regulation will soon implement an *opt-in* policy. Indeed, forthcoming *GDPR* regulation and more precisely the *ePrivacy* directive prevent firms from tracking Internet users unless they expressly said it so. In this case, advertisers must recover users consent to practice behavioral targeting. Finally, a *tracking ban* simply prevent advertisers and websites from tracking users. As personal information is intensively used in many industries, regulators are struggling to implement the best policy. Indeed, choosing the best regulation requires carrying out an economic assessment of each privacy policy stated above.

In this chapter, we analyze the economic impact of an *opt-out* privacy policy on online advertising auctions prices. As an opt-out modifies the available information on users, advertisers are less able to conduct tailored advertising campaign, which in turn modifies their willingness to pay for ads. Such question has already been studied by Johnson (2013). Using a proprietary dataset, the author's results predict that an opt-out would reduce publishers and advertisers revenue of respectively 3.9% and 4.6%. Other work from Johnson et al. (2017) assesses the impact of the AdChoice program - that allows users to opt-out from behavioral targeting - on ad prices. The author observes transaction data with ad prices and Internet users' behavior with respect to an opt-out. The study predict that opting-out from behavioral targeting on AdChoice would decrease the revenue by 59.2% with respect to comparable ads for users who did not opted-out. However, both abovementioned papers are unable to assess the impact of opting-out from behavioral targeting on the prices of ads seen by a single user. Indeed, recovering such data would require being able to track and identify users even after they opted-out from tracking, which is impossible for firms.

We develop a new, scalable and reproducible methodology that precisely lets us recover prices before and after an opt-out by Internet users. We carry out our experiment for 51 days in January and February 2017 and recover 6682 auction prices. Using a diff-in-diff methodology, we compute the difference in price before and after an opt-out allows us to draw the economic impact of an opt-out privacy policy on the market for online advertising spaces sold at auctions. The chapter underlines four key points.

Firstly, our methodology differs from previous works. We build a computational methodology that recovers advertising prices before and after an opt-out, which is not the case in Johnson (2013) and Johnson et al. (2017). We create bots[1] to be perceived as real Internet users, with no past history, that we divide in two groups: one control group and one treatment group. Both groups visit a list of websites twice successively. Between the two visits, a privacy behavior - i.e. an opt-out from behavioral targeting - is applied to the treatment group. Our database is therefore composed from advertising prices paid by winning bidders of online advertising auctions taking place in the first visit, second visit, and on bots from the treatment group and the control group.

Secondly, we do not find a clear relationship between opting-out from behavioral targeting and the prices of advertising spaces sold at auction. Moreover, our results indicate that websites included in specific categories see a weak increase in average advertising price when auctioning ads on users that have opt-out. This result can appear counter-intuitive. This is not in line with previous results that show a significantly lower average price when an ad is shown after an opt-out. One potential explanation lies in the fact that our bots are not valuable for advertisers. In this case, advertisers would bid higher when they are not able to recover bots characteristics (i.e. after an opt-out).

Thirdly, we contrast this result as we show that the effect of an opt-out significantly differ according to the category of the website on which the ad is shown. It contrasts with the previous results that predict a uniform decrease in prices when advertisers are unable to leverage users' information. Following our results, it appears that such effect may be conditional to websites characteristics.

Finally, we show that an opt-out from behavioral targeting significantly decreases the number of ads sold at auction. This is an important result as it underlines how a privacy policy changes competition between ad selling mechanisms. In this case, our result clearly indicates that automatic selling mechanisms are less employed by websites after an opt-out from behavioral targeting. It may also jeopardize our precedent results, as we do not account for cancelled auctions

[1]Bots are computer programs designed to perform tasks on computers. For example, we can "code" bots to visit a list of websites.

that arise in case of non-participation of advertisers (or null price), the diff-in-diff analysis may be biased. We offer a solution to control for such issue in Section 4.5.

Section 4.2 presents the literature related to the chapter. The computational methodology is presented in Section 4.3. The empirical analysis and the results are presented in Section 4.4. Section 4.5 concludes.

4.2 Related literature

Firstly, the literature on the economics of online advertising is naturally of great importance to our work. Online advertising markets are generally considered as two sided markets. Evans (2009) explains how the transformation of media markets stimulates online advertising and how data management intermediaries emerged from the need of a better match between advertisers and Internet users. The paper particularly relates to targeted advertising, as it massively uses personal information to tailor ads.

Hence, the economics of privacy related to online behavior appears important as users are increasingly hiding their personal information online. Empirical studies in the literature show that consumers have idiosyncratic valuation for their personal data. Acquisti and Grossklags (2006) presents an experiment which observes the willingness to accept to sell personal data against monetary compensation and the willingness to protect data. They find a clear cut result where the willingness to protect is sharply lower than the willingness to sell personal data, underlining the fact that user's valuation of personal information depend on the current situation. A more specific literature underlines the economic importance of privacy in a context of online advertising. Important papers such as Evans (2009), Acquisti et al. (2009) and Goldfarb and Tucker (2011) give insights on the link between privacy and online advertising.

More related to our paper, another stance of the literature analyzes the impact of new mechanisms used to sell online advertising. The existing literature has already established that preventing the use of consumers' personal information decreases advertisers' willingness to pay for an advertising space (De Corniere and De Nijs (2011) or Beales and A.Eisenach (2014)). Other papers empirically assess how privacy concern, and more precisely privacy regulations, may impact the online advertising market. Goldfarb and Tucker (2011) shows that European privacy laws[2]

[2]Authors investigate the effect of the "Privacy and Electronic Communications Directive" (2002/58/EC) on the

reduced advertising effectiveness, and therefore may represent a financial loss. All in all, above mentioned papers underline how advertisers' inability to target user segments is mechanically internalized and lowers their willingness to pay.

Finally, our research question is aligned with Johnson (2013) and Johnson et al. (2017): how an opting-out privacy policy impact the prices of advertising spaces sold at auctions? In both papers, the author analyzes online advertising in auction framework. He finds that prices are different according to privacy policies (opt-in, opt-out and tracking ban). Hence, the inability to use consumers' information for targeted ads bears economic significance for the online advertising industry. However, both papers estimate ad prices after opting-out as they cannot recover them. Our paper differs from the work of Johnson as we build a new methodology from the tools created in Olejnik et al. (2013) to conduct experiments. It allows us to recover real price distribution even after an opt-out. Hence, we are able to precisely measure the impact of opting-out from behavioral targeting on prices of advertising spaces sold at auction.

4.3 Methodology

4.3.1 Real Time Bidding auctions

The purchase of advertising through automatized auction mechanisms - known as Real Time Bidding (RTB) - has been attracting a lot of advertisers. In 2017, around 40% of display ads have been sold by RTB according to Emarketer. The same study predicts that the spending in RTB will reach $14B by the end of 2017.

In practice, RTB allows publishers to sell their advertising spaces at auction to advertisers. Such mechanism offers many advantages with respect to direct buying, such as information on sold space, pricing precision and optimization capability.[3] Among these advantages is the possibility to better trigger user data when formulating a bid. The immediate and simultaneous access to information on websites and users allows advertisers to precisely understand their valuation of displaying an ad.

We observe many intermediaries on the RTB advertising market. In the paper, a simple version of the RTB ecosystem is described where a single user visits a website and sees an ad.

online advertising industry

[3]See articles from Periscopix or Knowonlineadvertising. Many other results can be found on Internet or specialized magazines.

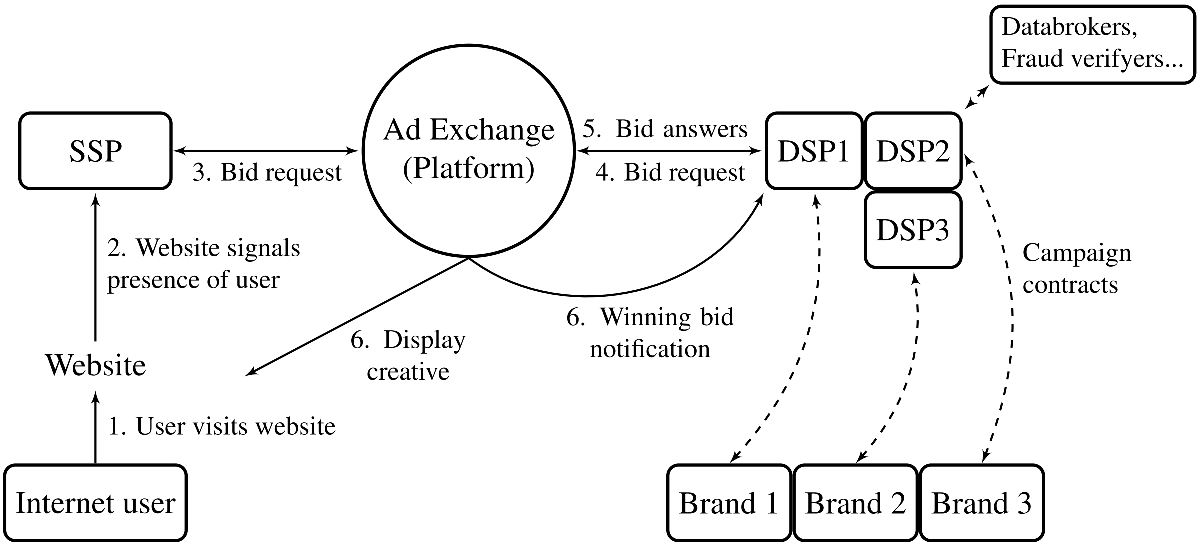

Figure 4.1: Real Time Bidding process

The website advertising spaces are managed by a Supply Side Plaform (SSP), which submits a bid request to an Ad Exchange (platform) when an Internet user connects to the website. The Ad Exchange organizes the auction in contacting potential interested bidders (advertisers). Usually denoted as Demand Side Platforms (DSP), bidders manage brands campaigns. To formulate a precise bid, a DSP mobilizes different third party actors such as databrokers which recover user information. By the time the webpage is loaded in the user's browser, the winning bidder has injected its ad. The entire RTB process takes less than $10ms$ and is summarized in Figure 4.1.

We use a flaw in the RTB process to recover advertising prices [4]. This allows us to analyze how price may vary with respect to Internet user behavior. We decompose our methodology in two steps: the computational process and the econometric analysis.

4.3.2 Recovering ad prices using a computational process

The methodology uses automated computer scripts (commonly called ”bots”) and exploits a flaw in the RTB process to recover ad prices. Automated computer scripts are built using the tools developed by the WebTAP platform project. It allows the bots to automatically visit a list of websites in impersonating a random human behavior. Each bot integrates an extension developed by Privactics (Inria) which recovers prices from auctions when the ad exchange notifies the winning

[4] This flaw has been discovered by Privactics (Inria) and highlighted in their research paper Olejnik et al. (2013)

bidder (event 3. on Figure 4.2).

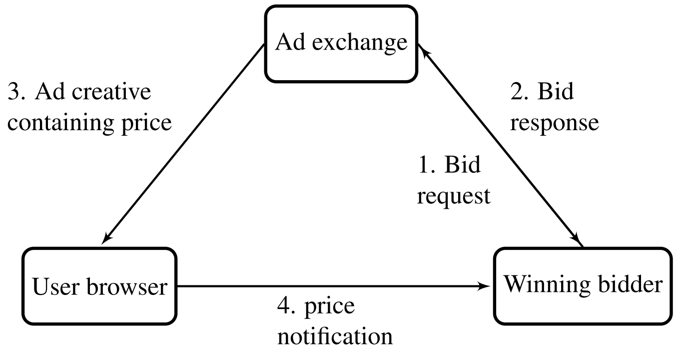

Figure 4.2: Simplified RTB process

The flaw in the RTB process lets us capture around 30% of prices that are usually encrypted. We correlate prices paid by advertisers with bot characteristics and the context in which the auction takes place. Some concerns may be raised regarding the reason why around 30% of the prices are not encrypted. This may represent a limitation as only clear-text prices are recovered. The Internet Advertising Bureau (IAB) considers that it is a matter of security: "Bidders [...] may optionally use price encryption [...] to reduce the risk of win-notifications being manufactured and falsified". This may indicate that the encryption may be related to website security reputation. In this case, one way to control for such issue would be to account for website fixed effects.

The computational methodology can be decomposed in five steps. All the process is roughly displayed in Figure 4.3:

1. Profile construction: two identical groups of bots are created. Bots are attributed with specific characteristics,

2. First period: both groups of bots visit a list of websites we denote by L. Advertising prices of period 1 are recovered,

3. Shock: one of the two groups of bots is "treated" with a shock,

4. Second period: both groups of bots visit for the second time the exact same list L of identified websites. Advertising prices of period 2 are recovered,

5. Database enrichment: using scrapping techniques, auction and environmental characteristics are recovered,

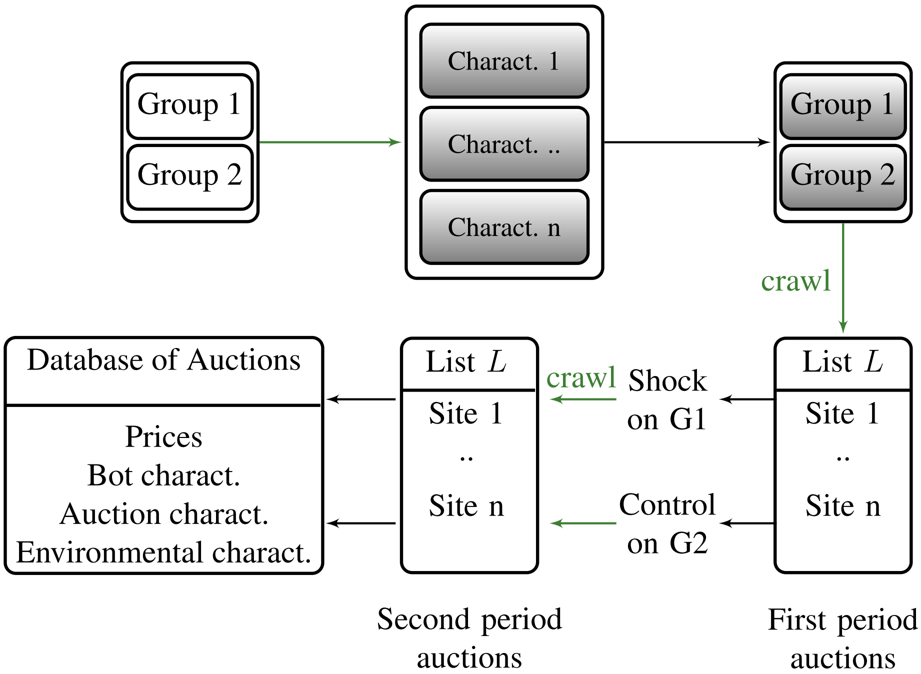

Figure 4.3: Computational process

Profile construction

The two groups of bots differ according to shocks but have the same static characteristics. Each bot has C characteristics such that each characteristic $c \in \{1, .., C\}$ can take K_c modalities such that $k_c \in \{1_c, ..K_c\}$. We create two groups of size $N = \prod_{c=1}^{C} \binom{1}{K_c}$ uniquely characterized bots. This means that each bots in one group have a unique combination of modalities from characteristics.

We use two methodologies to attribute characteristics to bots. A first option is to code the bot to visit a list of website related to a specific characteristic's modality. For example, the characteristic being the content preference, and the modality being "interested in sport", a bot can visit the top 50 website related to sport according to Alexa[5], to be identified as interested in sports. A second option is to use databrokers dashboard such as Exelate[6] which is used by

[5] Alexa is an aggregator of website metrics
[6] Exelate is a databroker that gather data and work with many actors participating in targeted advertising

many actors on the market. Exelate dashboard allows us to report characteristics such as age, gender, lifestyle and more generally personal preferences to be used by advertisers when serving advertising.

Crawl strategy and shock

In a first period, the two indistinguishable groups of bots visit a list L of websites and recover prices of ad spaces sold at auctions. After that first period, we perform the shock on one of the group, which is denoted by the "treatment" group after being shocked. The group that has not been shocked is denoted by the "control" group. In a second period, the two groups of bots visit for the second time the same list L and recover again ad prices.

Database enrichment

We also need to recover the context in which an auction takes place. For each auction we recover websites information from two online metric services: Mustat and SimilarWeb. From that tool we manage to recover the estimated traffic, SEO rank, estimated category and server's location on each websites where an auction took place. Secondly, our crawler is also designed to recover auctioned ad and environmental characteristics. We parse that information and recover the auction winner and exact date of the auction.

Using recovered ad prices, we compile a database of auction events and their corresponding features, including bots, websites and ad-slots characteristics.

4.3.3 Empirical Methodology

Our methodology does not recover the number of bidders for a specific auction. Previous works from Song (2004) and Marra (2015) offer methods to estimate auction winning bid without the number of bidders. Both methodologies perform identification from the bid distribution of each auctions. However, our methodology only collects the winning bid for each auction. In order to bypass endogeneity problem linked to the non-observation of the number of bidders, we use a diff-in-diff methodology with fixed effect in assuming that the number of contacted bidders for an auction on a specific website does not change from one period to the other. The assumption stems from the fact that websites are under contracts with ad-exchanges hosting auctions, which are also contractually linked to advertisers. We adopt the following econometric specification:

$$
\begin{aligned}
Price_{it} &= \beta_1 X_{i,t} + \beta_3 W_{i,t} + \beta_2 Z_{i,t} + \beta_4 AfterShock_t \\
&+ \beta_5 Treatment_i + \beta_6 (Treatment_i * AfterShock_t) + \epsilon_i
\end{aligned}
$$

Where i represents the auction number and t indexes the periods. $AfterShock_t$ is a dummy which takes value 1 when the auction takes place in the second period and 0 if it takes place in the first period. $Treatment_i$ is a dummy which takes value 1 when the bot is part of the treatment group. Hence, the variable $Treatment_i * AfterShock_t$ is equal to 1 if the bot watching the ad is from the treatment group and the auction takes place in the second period. Vector $Z_{i,t}$ is the set of characteristics of auction i at time t. Vectors X_i and W_i are respectively sets of characteristics of the bot and the website part of auction i.

4.4 Empirical analysis

4.4.1 Modeling an opt-out choice

We use the abovementioned methodology to model how an opt-out from behavioral targeting can impact prices of ad sold at auctions. To make the analysis simple, we assume that an experiment is composed of two groups of two bots. Within a group, the bots exhibit the same range of age located in 18-34, the same preference for fashion shopping but can be distinguished according to their gender: one bot is gendered as male and the other as female. We therefore apply the same steps as the methodology mentioned in Section 4.3.

Profile construction: age and gender characteristics are assigned to bots using Exelate dashboard. Preferences categories are attributed to each bot in visiting the top 20 websites related to fashion shopping on Alexa[7]. Both methodologies are explained is Section 4.3. Our two populations are indistinguishable regarding characteristic.

First period: in a first period, both groups of bots visit a list of websites on which our tool can recover non-encrypted prices of ads sold at auctions. The website list come from previous Inria project that developed the tools that recover advertising price. [8]

[7]The list can be found in Appendix 4.6 in table 4.9.

[8]In the paper, authors identify websites that are more likely to pass encrypted price through the bot's browser when selling ad at auctions. The list can be found in Appendix 4.6 in table 4.8.

Shock: after the first visit, bots from the treatment group perform an opt-out from behavioral targeting using Adchoice program website. The AdChoice program has been launched in 2010 by the Digital Advertising Alliance, which gather US advertising groups and consortium to foster the use "self-regulatory solutions to online consumer issues". The AdChoice program allows users to opt-out from behavioral targeting. In this case, users still see ads, but not targeted ads based on past browsing behavior.

Second period: in a second period, both groups of bots visit for the second time the exact same list of websites, in the exact same order.

Database enrichment: we enrich the data base in recovering additional auctions characteristics. Firstly, we recover contextual data from the scraper (date, winning bidder). Secondly, we recover websites characteristics from Mustat and SimilarWeb.

We repeat the experiment each day for 51 days, from the beginning of January to the end of February 2017, which gives us a database of 6082 auctions distributed over 46 websites.

4.4.2 Descriptive statistics

Bot populations' summary statistics

In Table 4.1 we summarize key variables on bot populations. We see that around 40% of the auctions in the database take place in the second period (i.e. after the shock), which may point out a lower interest for advertisers to win auction at RTB. 48% of the auctions in the database have been performed while a bot from the treatment group was browsing the corresponding website, which highlights a stable number of auctions between the two groups. However, only 15% of the auctions take place in the second visit while showing ads to a bot from the treatment group. Conversely, 31% of the auctions take place in the second period but show ads to a bot from the control group.

Overall, Table 4.1 underlines that bots tend to see less auction in the second period and particularly when they are part of the treatment group.

We also try to analyze the difference in average price between the different groups of bots (control and treatment) and the different periods (first period and second period). We consider prices formulated in CPM in euros.[9] Table 4.1 shows that the average price of ads shown to bots

[9]Hence, each price recovered must be divided by 1000, in order to find the price for a single ad.

from the treatment group is significantly higher in the second period (i.e. after the bots from the treatment group have opted-out) than in a first period. Conversely, the average price of ads is more or less similar whether the bot is from the control or the treatment group. Combining periods and bot groups, we find that average price of ads in the first period is about the same whether bots are from the control or the treatment group. However, the second period exhibits higher average price of ads for bots from the treatment group than from the control group.

Table 4.1: Average price by bot subpopulation

Population	Ave. price	S.d	Min	Max	N
All population	0.33	0.77	0.000025	12.54	6682
Before shock (t=1)	0.31	0.78	0.000025	12.54	3994
After shock (t=2)	0.36	0.76	0.000033	8.59	2688
Control	0.34	0.8	0.000033	8.59	3480
Treatment	0.32	0.74	0.000025	12.54	3202
Control#Before Shock	0.32	0.79	0.000073	8.59	1957
Control#After Shock	0.37	0.81	0.000033	9.59	1523
Treatment#Before Shock	0.31	0.77	0.000025	12.54	2037
Treatment#After Shock	0.35	0.68	0.01	8.59	1165

These simple statistics gives us a first hint on how advertising prices may vary when bots seeing ads have previously opted out from behavioral targeting. However, these first clues may just be driven by website and advertisers strategies, as well as events happening on the day of the experiment. We therefore analyze how the number of auctions and average prices change according to website, advertisers and experiment number.

Websites summary statistics

For each experiment, our two populations of bots visit a list of 46 websites [10] on which we recover clear-text advertising prices. Firstly, we notice that websites exhibit very different numbers of auctions and exhibit a very concentrated distribution. Indeed, *mobafire.com*, *thechive.com* and *hunt4freebies.com* gather more than 40% of the auctions and hence highly influence the results. Figure 4.4 displays the number of auctions for the most important websites.[11]

We analyze how the average price changes across websites. The highest average price is on *steamkitten.com*. We notice that the first three websites exhibiting the highest average prices

[10] The list can be found in Appendix 4.6 in table 4.10.
[11] The total list can be found in Appendix.

Figure 4.4: Number of auctions by website

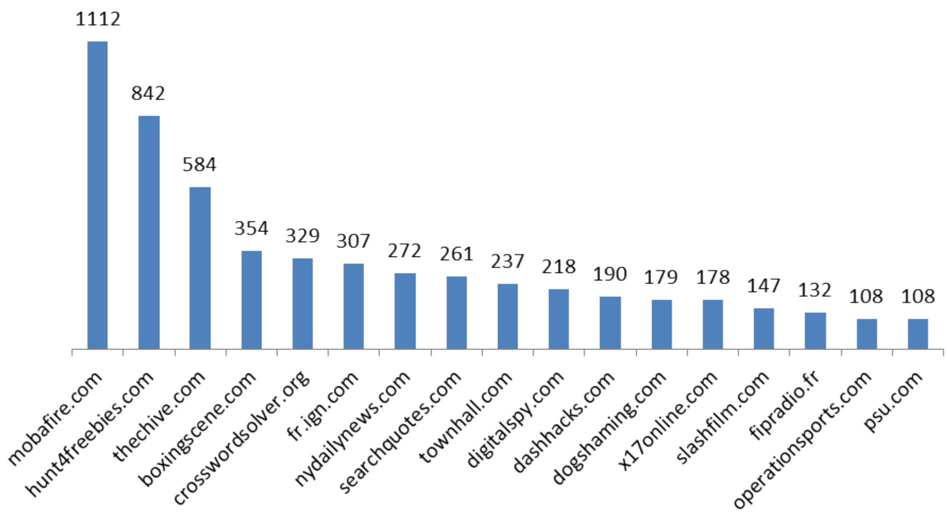

also concentrate less than 30 auctions. The fourth is *searchquotes.com* that exhibits 261 auctions (Figure 4.5).

Figure 4.5: Average price by website

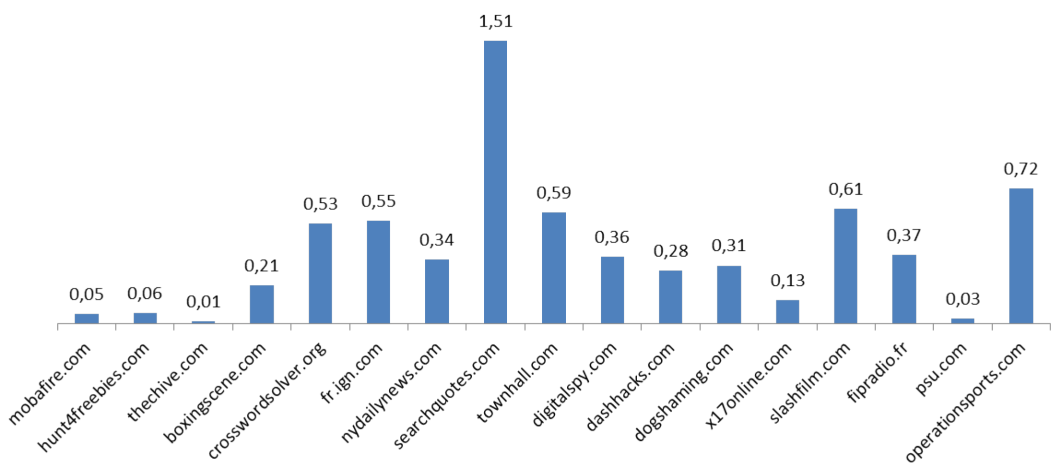

More generally, the average price does not seem to be correlated with the number of auctions in our database (Figure 4.6). It still underlines that websites having the highest number of advertising auctions are exhibiting really low advertising prices.

Finally, 4.7 shows how revenue is distributed among websites. We notice that *searchquotes.com* generates the highest revenue among all the websites.

Figure 4.6: Average price by number of auctions for each website

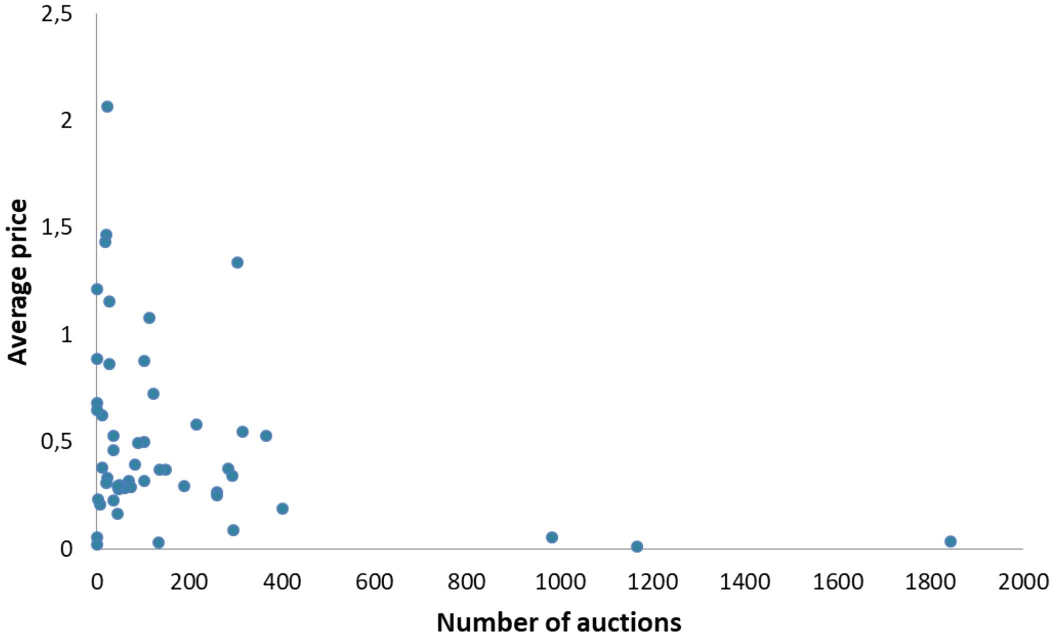

Figure 4.7: Revenue by website

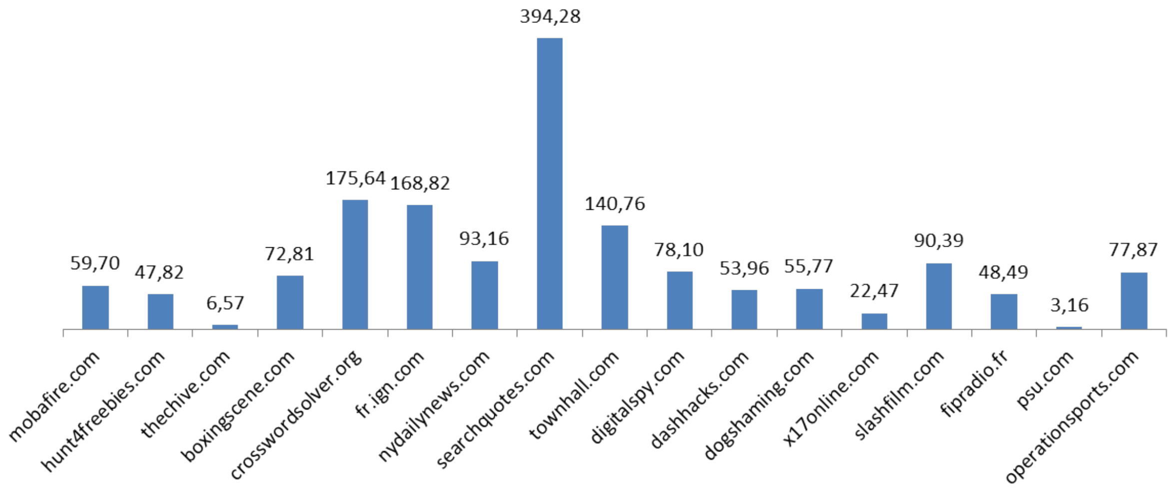

Overall, websites exhibit strong heterogeneity regarding advertising prices and number of auctions, hence suggesting that some of them may be more attractive to advertisers.

Advertisers summary statistics

Using the tools developed by INRIA-PRIVATICS, we are able to recover the name of the winning advertiser. We are not able to recover its exact valuation but only the market price (i.e. the second highest bid) paid to display an ad. We first analyze how the number of won auctions varies between advertisers. In Table 4.8, we highlight the importance of *Rfihub* in our database with almost 50% of won auction.[12]

Figure 4.8: Number of won auction by advertiser

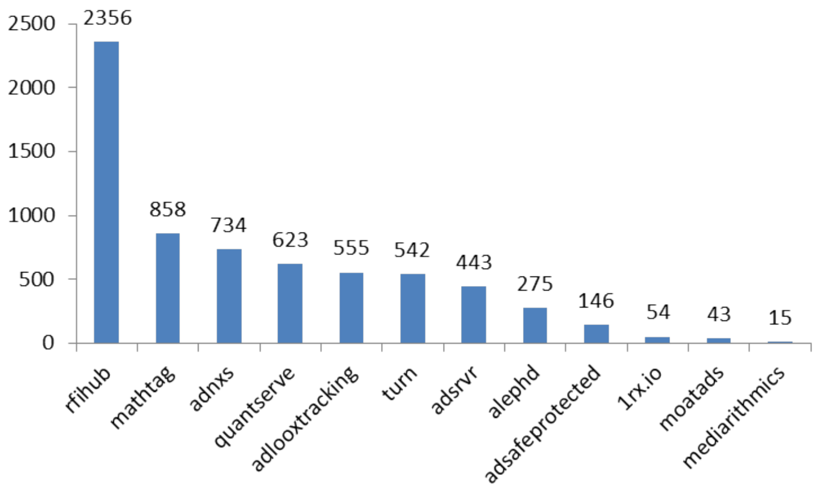

Conducting the same analysis for average price, we find that it highly varies between advertisers. Among advertisers that have won more than 15 auctions, the highest average price is paid by *Adsrvr* over only 443 auctions (Figure 4.9).[13] Confronting the number of auctions with the average price paid per advertisers brings no significant correlations (Figure 4.10).

The highest amount spent during our experiment among advertisers is illustrated in Figure 4.11: *Mathtag* and *Adsrvr* - both large international companies - spent more money on advertising than the other advertisers during the total experiment.

Experiment summary statistics

[12] *Rfihub* stands for Rocket Fuel Incorporated Hub, which is a firm belonging to Sizmek, one of the leading online marketing company in the world. Remaining actors are well established firms operating both in Europe and in the US.

[13] The complete summary statistics about winning bidder can be found in 4.6 in Table 4.11.

Figure 4.9: Average price paid by advertiser

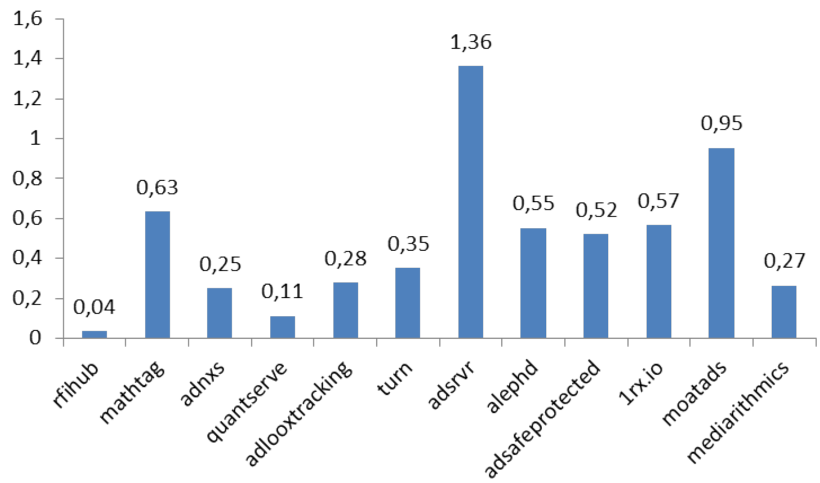

Figure 4.10: Average price by number of auctions for each advertiser

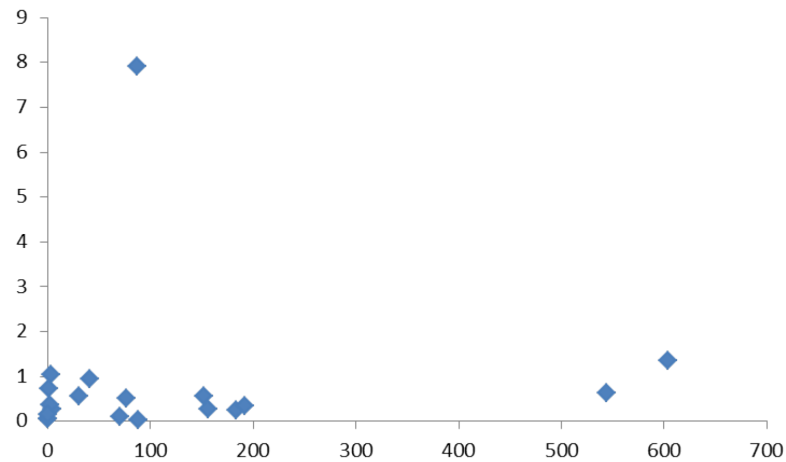

Figure 4.11: Total paid by advertiser

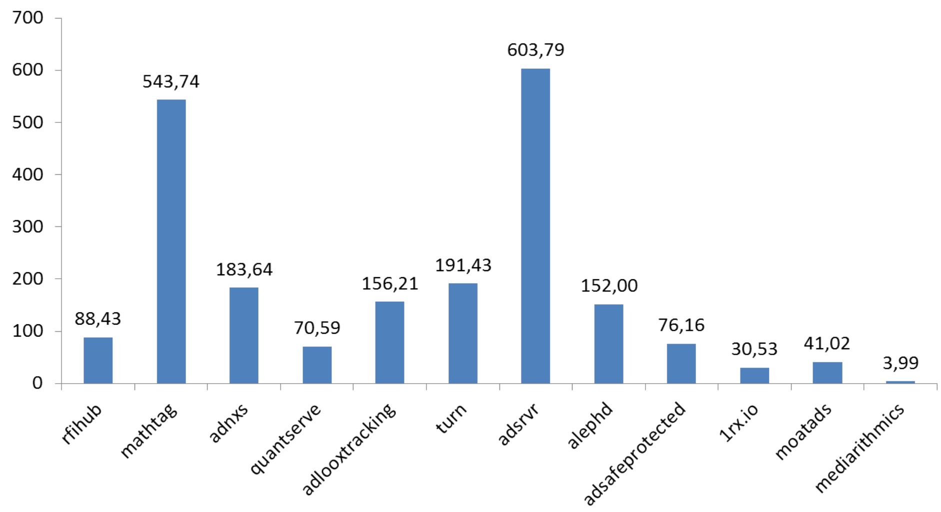

We repeat the experiment for 51 days. We therefore have 51 databases of auctions. Firstly, we notice that the number of auctions varies according to the experiment number. We find an exceptionally high number of auctions on the 49th experiment which corresponds to the 23^{rd} of February 2017 (see Figure 4.12). We don't find any particular event that day that would indicate a higher number of ads sold at auction. This could be explained by a lot of different reasons such as the launch of a new campaign for a leading advertisers. More generally, we observe significant heterogeneity between experiments regarding advertising prices.

Figure 4.13 shows that the average price for ads also varies by experiment. It also indicates different periods with different average prices. For example, we notice that from experiment 6 to 16, the average price seems lower, while it is higher for the next periods.

Analyzing the link between the number of auctions and the average price, we notice a weak negative correlation. Indeed, for a specific experiment number, the higher the number of auctions, the lower the average advertising price (see Figure 4.14).

This points out how advertising number and average price depends on many different variables. Moreover, we observe heterogeneity in generated revenue across experiments, which underline how generated revenue for websites may vary from one day to the other (see Figure 4.15).

After analysis, we notice that advertisers, websites and experiments must be taken in account in the econometric analysis. As we perform a diff-in-diff, we observe that fixed effect variables

Figure 4.12: Number of auctions by experiment

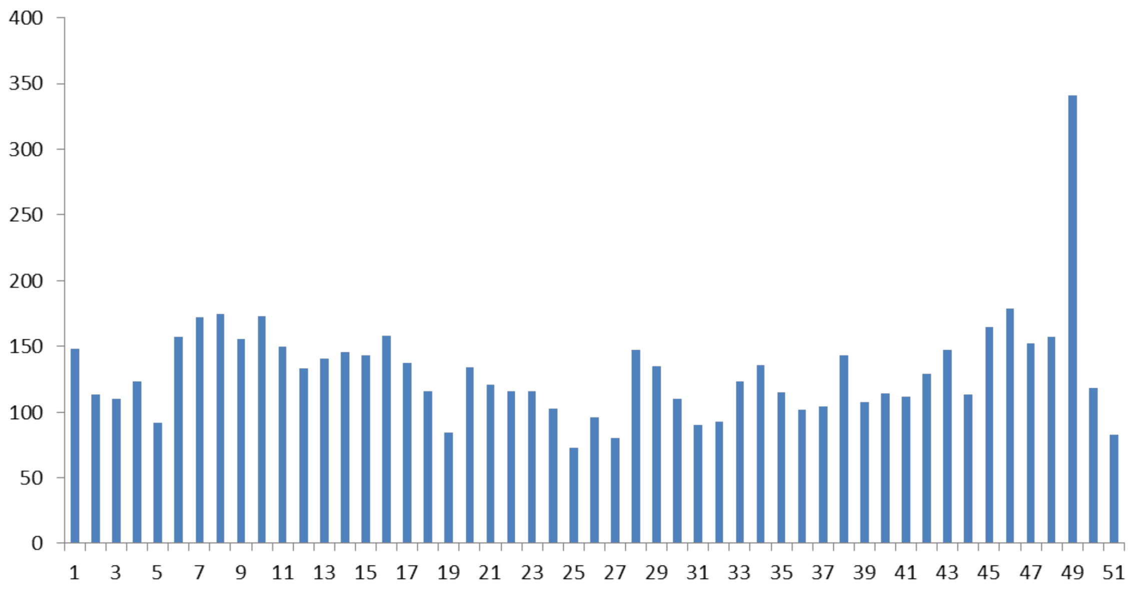

Figure 4.13: Average price paid by experiment

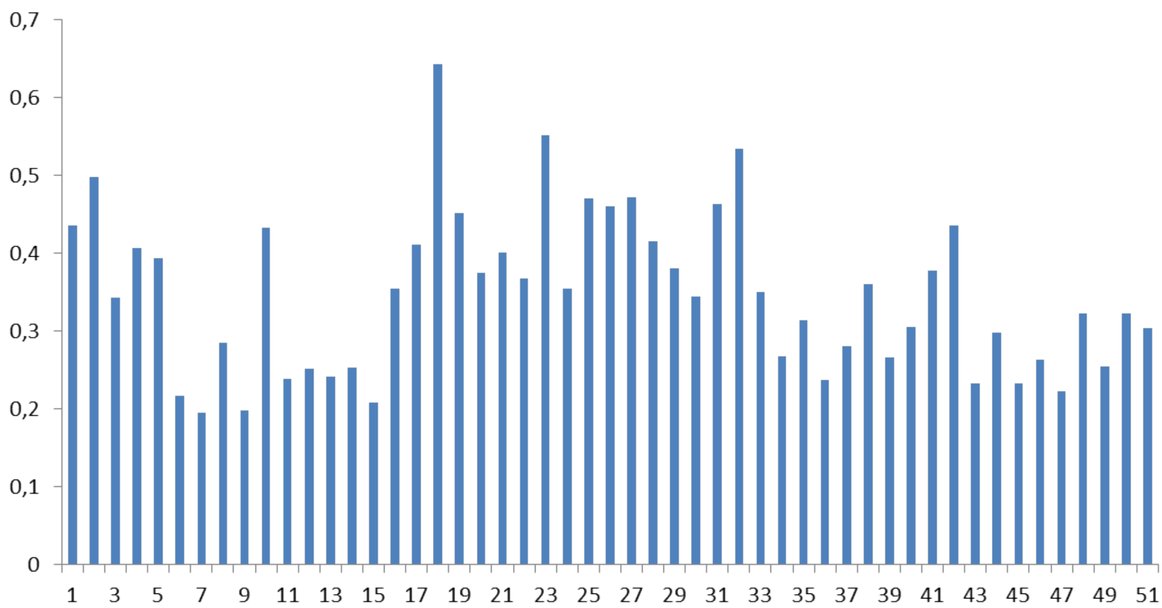

Figure 4.14: Average price by number of auctions for each experiment

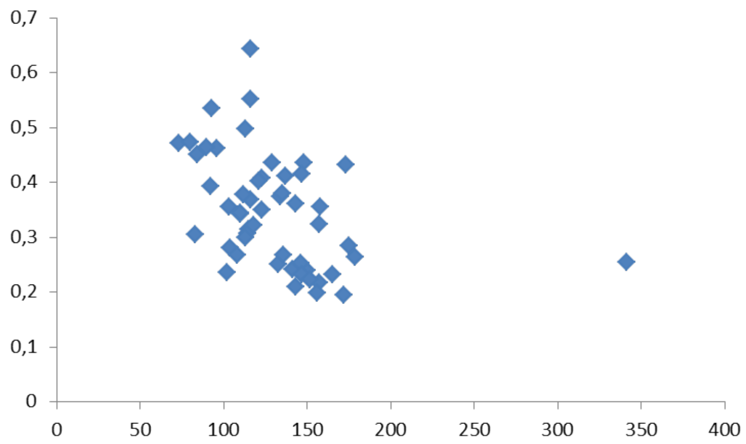

Figure 4.15: Average money transfer by experiment

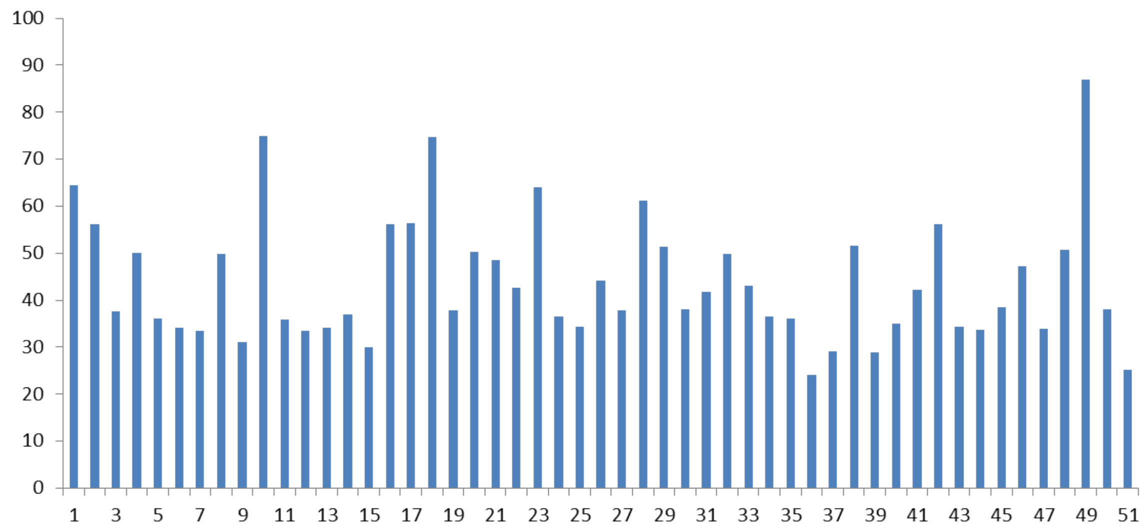

should have no effect in the analysis, as websites, advertisers, experiment and bot characteristics remains the same no matter the period. However, as the number of auctions changes between the two periods the average effect of fixed effect are not the same as well. This is very important as it underlines an inconsistency in the number of ads sold at auction between the periods.

4.4.3 Econometric analysis

The effect of an opt-out privacy policy on prices

We first run a simple regression based on the methodology described in Section 4.3. We want to measure how prices of ad spaces sold at auction are modified when users opt out from behavioral targeting. As many prices are smaller than 1, we use the log of the price. We focus on the bot group (Control/Treatment), the period (After shock/Before shock) and the interaction between the two groups. Our interest lies in the significance of the average price difference when the bot is part of the treatment group and the auction takes place after bots from the treatment group have opted-out (i.e. the variable *Treatment#After shock*).

Results are presented in Table 4.2. We find that opting out from behavioral targeting significantly increases advertising prices in a simple diff-in-diff regression without fixed effects (column (1)). However, in order to account for endogeneity from not recovering the number of bidders for each auctions, we have to control for the website where the ad is sold. Inserting fixed effects at the website level does not highlight any effect of an opt-out from behavioral targeting on prices (column (2)).

To contrast this result, we classify websites with categories recovered from SimilarWeb.[14] We create interactions between website categories and the diff-in-diff variable (*Treatment#After shock*). From regression (3) to regression (5), we find that opting out from behavioral targeting significantly increases prices of advertising spaces sold at auction. However, we show that this impact is not uniform between websites. The variable *Treatment#Aftershock#Website Type* represents the effect of the website reference category "Arts and Entertainment", which contains very general websites according to SimilarWeb. Hence, results in the table display effects of opting-out from behavioral targeting on a specific type of website relative to opting-out on a website from the "Arts and Entertainment" category. Opting out from behavioral targeting significantly

[14]The distribution of websites by category can be found in Appendix 4.6 in table 4.13. Summary statistics about website categories can be found in table 4.12

Table 4.2: Diff-diff regressions

OLS reg. Dependent variable: ln(price)	(1)	(2)	(3)	(4)	(5)
After shock (t=2)	-0.24***	-0.17	-0.19	-0.19	-0.16
	(0.08)	(0.08)	(0.13)	(0.11)	(0.08)
Treatment	-0.05	-0.01	-0.03	-0.02	-0.01
	(0.06)	(0.02)	(0.03)	(0.03)	(0.03)
Treatment#After shock	1.21***	0.16	0.6***	0.27**	0.25**
	(0.1)	(0.14)	(0.21)	(0.11)	(0.09)
Treatment#After shock#Website Type					
(Reference category is "Arts and Entertainment")					
Autos and Vehicles			-0.45***	-0.11	-0.06
			(0.09)	(0.08)	(0.06)
Computers and Electronics			-0.42***	-0.008	-0.04
			(0.15)	(0.11)	(0.11)
Food and Drink			-0.72***	-0.38***	-0.32***
			(0.15)	(0.11)	(0.08)
Games			-0.16	-0.11	-0.15
			(0.85)	(0.58)	(0.55)
News and Media			-0.39	-0.11	-0.09
			(0.27)	(0.16)	(0.18)
People and Society			- 0.49**	-0.24	-0.23*
			(0.18)	(0.16)	(0.11)
Reference			-2.1***	-0.24	-0.27
			(0.15)	(0.42)	(0.43)
Shopping			-0.27*	-0.09	-0.1
			(0.15)	(0.07)	(0.07)
Sports			-0.23	-0.009	-0.009
			(0.17)	(0.11)	(0.09)
Website Dummies		X	X	X	X
Advertisers (DSP) Dummies		X		X	X
Batch number Dummies		X			X
Constant	-2.88***				
	(0.05)				
N	6682	6682	6682	6682	6682
R^2	0.03	0.67	0.67	0.67	0.78
Within R^2		0.006	0.01	0.005	0.005

Robust standard errors in parentheses, clustered at Website, experiment and advertiser level.

*** $p<0.01$, ** $p<0.05$, * $p<0.1$

decreases prices paid for advertising space on website that falls into the "Food and Drink" or "People and Society" categories with respect to the "Arts and Entertainment" category. [15] For example, an opt-out from behavioral targeting would decrease prices on average 27%[16] more on a website from the category "Food and Drinks" relative to a website from the category "Arts and Entertainment".

Our analysis underlines that opting-out from behavioral targeting does not uniformly impact prices. To clarify such effect, we run separate analysis for website categories that exhibit more than 800 auctions. We retain the following four categories: "Arts and Entertainment" (17 websites) , "Games" (6 websites), "News and Media" (9 websites) and "Shopping" (1 websites).

Firstly, we show in Table 4.3 that the average price tends to be higher in the subgroup *Treatment#After shock* than in the subgroup *Treatment#Before shock* no matter the category of the website. Conversely, we do not find the same clear results for the *Control* group. Indeed, the average price is not always higher in the subgroup *Control#After shock* compared to subgroup *Control#Before shock*.

Table 4.3: Average price by website category and bot subpopulation

	"Arts and Entertainment"	"Games"	"News and Media"	"Shopping"
All population	0.23	0.25	0.55	0.05
Before shock (t=1)	0.19	0.27	0.48	0.05
After shock (t=2)	0.3	0.22	0.65	0.06
Control	0.19	0.24	0.54	0.05
Treatment	0.28	0.25	0.54	0.06
Control#Before Shock	0.16	0.29	0.49	0.05
Control#After Shock	0.22	0.2	0.64	0.05
Treatment#Before Shock	0.2	0.24	0.47	0.05
Treatment#After Shock	0.43	0.26	0.67	0.066

Secondly, we analyze how the number of auctions varies according to website categories. We find a lower number of auctions in the subgroup *Treatment#After shock* than in the subgroup *Treatment#Before shock* no matter the category of the website. We find that it is not always the case for the *Control* group.

[15] As "Arts and Entertainment" gather very general content websites, we may assume that other websites belonging to more specialized categories are more or less affected by an opt-out from behavioral targeting.

[16] $100(e^{-0.32} - 1) = -0.27$

Table 4.4: Number of auction by website category and bot subpopulation

	"Arts and Entertainment"	"Games"	"News and Media"	”Shopping”
All population	1672	1979	834	842
Before shock (t=1)	1066	1112	550	534
After shock (t=2)	695	867	308	2688
Control	930	1123	414	373
Treatment	831	856	420	269
Control#Before Shock	522	532	273	264
Control#After Shock	408	591	141	109
Treatment#Before Shock	544	580	277	270
Treatment#After Shock	287	276	143	199

We perform the same diff-in-diff regression for the four categories independently. Results can be found in Table 4.5. We find that an opt-out from behavioral targeting does not impact the price of ads sold at auction, except for the "Arts and Entertainment" category, which exhibits prices that are 23%[17] higher on average when advertisers can't use behavioral targeting.

A higher advertising prices after an opt-out appears counter-intuitive and is not in line with precedent results. Two potential reason may motivate this result.

Firstly, the short browsing history of our bots may reflect a low value for advertisers. In that case, advertisers lower their bids when behavioral targeting is allowed, as they know they compete to display ads to low value Internet user. Hence, when behavioral targeting is not allowed, advertisers have no information on Internet users and adjust their bids accordingly. In this case, this result illustrates how low value consumers opting-out from behavioral targeting may impact the price of ads sold at auctions. It contrasts with existing studies in underlining how having access to consumers' information lowers advertising prices if these consumers are not of interest for advertisers.

Secondly, higher prices after an opt-out may be explained by advertisers considering that only high value users hide their personal information. In that case, behavioral targeting allows advertisers to adjust their bid for each Internet users, as they can precisely compute their willingness to pay from available information. Conversely, when behavioral targeting is not allowed, advertisers are unable to leverage user information. Hence, advertisers adjust their bids with users' characteristics that can be correlated with hiding information online (which would be positive in

[17] $e^{0.21} - 1 = 0.23$

Table 4.5: Diff-diff regressions by categories

OLS reg. Dependent variable: ln(price)	"Arts and Entertainment"	"Games"	"News and Media"	”Shopping”
After shock (t=2)	-0.25	-0.33**	-0.06	-0.13
	(0.14)	(0.11)	(0.09)	(0.07)
Treatment	0.003	-0.08	0.02	0.04
	(0.04)	(0.09)	(0.04)	(0.03)
Treatment#After shock	0.21*	0.08	0.004	0.09
	(0.11)	(0.3)	(0.1)	(0.09)
Website Dummies	X	X	X	X
Advertisers (DSP) Dummies	X	X	X	X
Batch number Dummies	X	X	X	X
N	1672	1659	834	842
R^2	0.78	0.67	0.67	0.48
Within R^2	0.01	0.006	0.03	0.01

Robust standard errors in parentheses, clustered at Website, experiment and advertiser level.
*** p<0.01, ** p<0.05, * p<0.1

that case).

The first effect relates to the economics of low value Internet users when those hide their information. The second effect relates to privacy as a signal for advertisers. Overall both explanations could be tested and may be of interest for future economic research.

The effect of an opt-out privacy policy on the number of ads sold at auction

As we saw in Section 4.4.2, the number of ads vary according to websites, experiments, winning advertisers and bot characteristics. Hence, we wonder how an opt-out privacy policy may change the number of ads sold at auction. Indeed, programmatic selling mechanism may be less attractive for websites or advertisers compared to direct selling for example.

We aggregate data at the experiment, bot's type and period level. As we have 51 experiments, 4 unique bots within each experiment, and 2 periods (after and before shock) for each bot, we obtain a database of $51 * 4 * 2 = 408$ observations. Each observations represent the aggregated journey of one bot visiting the list of websites during one period in one experiment. Hence, we observe the number of ads sold at auction that each bot in each experiment have seen during a period, all websites taken together.

We see from Table 4.6 that there is less ads served in average after an opt-out from behavioral targeting. It also seems that such effect is stronger when the bots are part of the treatment group.

Table 4.6: Summary statistics on the number of auctions

	N	Ave. num. of auction	Sd	Min	Max
All population	408	16.3	8	3	67
Before shock (t=1)	204	19.5	8	4	67
After shock (t=2)	204	13.17	6.73	3	57
Control	204	17	8	4	57
Treatment	204	15.6	8	3	53
Control#Before Shock	102	19.18	9.13	4	67
Control#After Shock	102	14.93	7.18	4	57
Treatment#Before Shock	102	19.97	7.77	5	53
Treatment#After Shock	102	11.42	5.77	3	48

We perform a diff-in-diff regression on the number of auctions seen by a bot during a period within an experiment:

$$\begin{aligned} Auction\ Number_{e,t,b} &= \beta_1 X_e + \beta_2 W_t + \beta_3 W_b + \beta_4 AfterShock_{e,t} \\ &+ \beta_5 Treatment_{e,b} + \beta_6 (Treatment_{e,b} * AfterShock_{e,t}) + \epsilon_e \end{aligned}$$

Where e indexes experiments, b indexes the bots participating in the experiment and t indexes the period within an experiment. The explained variable is the number of ads sold at auction seen by a specific bot during one period of one experiment. X_e, W_t and W_b respectively represent characteristics related to experiment, period and bot. In the end, only the bot category remains valid as bots can differ according to their gender.

Results in Table 4.7 shows a significantly lower number of auctions during the second period. Moreover, this effect is doubled when the bot has opted-out from behavioral targeting between the two periods. Indeed, we notice that a bot that opted-out from behavioral targeting triggers around 4 less auctions than one that did not. This result underlines that at the experiment level, bots that have opted-out from behavioral targeting see less ads sold at auction than the others.

This may be explained by the fact that programmatic selling is very convenient to leverage personal information as pointed by the IAB 2017 report on behavioral targeting. Hence, opting-

Table 4.7: Regression on the number of auction

OLS reg. Dependent variable: Nb of auction	(1)	(2)	(3)
After shock (t=2)	-4.25***	-4.27***	-4.24***
	(1.07)	(1.07)	(0.90)
Treatment	0.78	0.78	0.78
	(1.11)	(1.11)	(0.74)
Treatment#After shock	-4.29***	-4.29***	-4.28***
	(1.44)	(1.44)	(1.43)
Female bot		X	X
Batch number Dummies			X
Constant	19.18***	17.60***	20.58***
	(0.8)	(1.79)	(4.98)
N	408	408	408
R^2	0.18	0.18	0.56
Within R^2			0.3

Robust standard errors in parentheses
*** $p<0.01$, ** $p<0.05$, * $p<0.1$

out may impact the participation of advertisers to an auction. In the case of low or no participation to an auction, websites may prefer to sell ads through another channel. Two cases may arise.

If we assume that opting-out from behavorial targeting increases prices, websites should increase the number of ads sold at auctions (or do nothing if all ads are already sold at auction).

Conversely if opting-out decreases prices, websites would benefit from favoring other sales channel and may reduce the number of ads sold at auction. In that case, the effect of opting-out from behavioral targeting on advertising prices depicted in Tables 4.2 and 4.5 may be biased. Indeed, we now have to account that after an opt-out, websites may choose direct sales over auctions as they generate lower revenue. In that case, we don't recover the full effect of an opt-out on prices as this truncation at the bottom of the auction price distribution is very likely to bias the average price over periods when using diff-in-diff regression.

This problem may be identified as a sample selection issue, as we only observe auction data when at least one advertiser have participated. One way to treat this problem would be to use a Heckman two step model that would explain how a privacy policy favoring an opt-out from behavioral targeting may impact 1) advertisers participation and 2) prices of ads separately. However, we do not recover observations that captures non-participation from advertisers. Hence estimating advertisers' participation is not possible.

Overall, our analysis displays results that help us better understand how an opt-out privacy policy may have a wide range of effects on the market for online ads sold at auction. We sum up the results of the chapter, the limitations of the methodology and the future research at stake in the next section.

4.5 Conclusion and limitation

The chapter analyzes the impact of opting out from behavioral targeting on the online advertising market. Former literature underlines how the lack of information stemming from a privacy behavior decreases targeting capabilities of advertisers, which foster them to lower their advertising prices. Another stance of the literature has already provided answers to our research question in using different methodologies: It underlines that opting out from behavioral targeting significantly decreases prices for online advertising spaces sold at auctions.

We build a computational experiment that allows us to recover prices of advertising spaces sold at auction. Conversely to previous work, our methodology can recover prices of ads even after an opt-out. Such capability enables a more precise estimation of the impact of opting out from behavioral targeting on prices. Our analysis underlines three important results.

Firstly, an opt-out privacy policy may have a weak positive impact on prices. Indeed, Internet users may be sold at a higher price when advertisers cannot leverage information about them. This may be due to two reasons. First, our bots exhibit a very short browsing history, and are maybe considered of low value by advertisers. In that case, having no information prevent advertisers from adjusting bids downward. Second, hiding information may be correlated with valuable user characteristics. In that case, advertisers may increase their bids when they can't access information about users.

Secondly, this effect appears to rely on the auction context, and especially on the type of website where the auction takes place. This makes sense as websites may have different reputation or show different type of information. In this case, auctions context appears to be more important for advertisers as they cannot leverage information about users anymore.

Thirdly, an opt-out from behavioral targeting seems to significantly decrease the number of ads sold at auction. This effect underlines a competition between advertising selling mechanisms, where websites choose the best solution to maximize their revenue. This also may jeopardize our precedent results, as we do not account for potential auction that have been substituted to other selling channel by the opt-out privacy policy.

We recover precise data on online ads sold at auction. However, some limitations arise. Firstly, we only recover 30% of auctions prices. To use such prices as representative sample, we have to make sure that the choice to encrypt prices is independent from the bid strategy of advertisers, which we are unable to properly verify yet. Secondly, we do not recover the number of bidders. While we do make assumptions on websites' behavior and use a diff-in-diff strategy, we may still encounter endogeneity issues. Thirdly, a larger set of unique bots as well as a longer website list may strengthen the results. Finally, we may have a sample selection bias based on the fact that an opt-out may affect advertising price until a point where websites may choose to not use auction to sell ads. One potential solution that could be implemented is to recover the number of ads on each webpage, hence providing us with a way to compute the proportion of ads sold through auction mechanisms. It would also allow us to better model how the demand for automatic selling mechanism depend on privacy policy.

Our methodology can be extended in various ways. A potential one would be to test how an *opt-in* privacy policy may impact prices of ads sold at auctions. This is all more important as European privacy regulations are currently implementing such privacy policy that would require website to collect Internet users' consent before making use of their personal information. This may have large impact on users' behavior as it seems to lower the cost of hiding. Correspondingly, such regulation may cover implications for advertisers that would serve less precise advertising as they operate without information. Overall, a deeper analysis of the effect of an opt-in privacy regulation on the advertising market may be needed, as it will surely change market equilibria.

Bibliography

Acquisti, Alessandro and Jens Grossklags, "What Can Behavioral Economics Teach Us About Privacy?," 2006.

__ **, Leslie John, and George Loewenstein**, "What is privacy worth," in "Workshop on Information Systems and Economics (WISE)" 2009.

Beales, J.Howard and Jeffrey A.Eisenach, "An empirical analysis of the value of the information sharing in the market for online content," *Navigant Economics*, January 2014.

Corniere, Alexandre De and Romain De Nijs, "Online advertising and privacy," Technical Report, mimeo 2011.

Evans, David S., "The online advertising industry: Economics, evolution, and privacy," *The journal of economic perspectives*, 2009.

Goldfarb, Avi and Catherine E. Tucker, "Privacy Regulation and Online Advertising," *Management Science*, 2011, *57* (1), 57–71.

Johnson, Garrett A., "The Impact of Privacy Policy on the Auction Market for Online Display Advertising," SSRN Scholarly Paper ID 2333193, Social Science Research Network, Rochester, NY May 2013.

Johnson, Garrett A, Scott Shriver, and Shaoyin Du, "Consumer Privacy Choice in Online Advertising: Who Opts Out and at What Cost to Industry?," 2017.

Marra, M, "Nonparametric identification of ascending auctions without the number of bidders," in "in" 2015.

Olejnik, Lukasz, Tran Minh-Dung, and Claude Castelluccia, "Selling Off Privacy at Auction," December 2013.

Song, U, "Nonparametric estimation of an eBay auction model with an unknown number of bidders," Technical Report 2004.

Tucker, Catherine E., "The economics of advertising and privacy," *International Journal of Industrial Organization*, 2012, *30* (3), 326 – 329. Selected Papers, European Association for Research in Industrial Economics 38th Annual Conference, Stockholm, Sweden, September 1-3, 2011.

4.6 Appendix

nowish.cheezburger.com	kidsgamesheroes.com	imnotobsessed.com	thecelebritycafe.com	cinetrafic.fr
crosswordsolver.org	mobafire.com	abovetopsecret.com	allfreefightvideos.com	voici.fr
hunt4freebies.com	sidereel.com	engadget.com	shefinds.com	babycenter.fr
techeblog.com	stuff.tv	socialitelife.com	blacksportsonline.com	sport.fr
digitalspy.co.uk	tvovermind.com	aol.ca	atlnightspots.com	bilansgratuits.fr
boxingscene.com	thechive.com	routenet.nl	starcasm.net	manageo.fr
kevinandamanda.com	autosport.com	oceanup.com	thatgrapejuice.net	butfootballclub.fr
steamykitchen.com	eurweb.com	voetbalzone.nl	ask-oracle.com	fipradio.fr
nesn.com	kitchendaily.com	bostonherald.com	striptease.keenspot.com	livefoot.fr
psu.com	cougarboard.com	livingrichwithcoupons.com	thegrio.com	funradio.fr
lucianne.com	dashhacks.com	coolinterview.com	southwales-eveningpost.co.uk	voici.fr
operationsports.com	toptenz.net	firstshowing.net	mediengestalter.info	babycenter.fr
coches.net	searchquotes.com	thecelebritycafe.com	notizieshock.it	sport.fr
pajiba.com	thetechjournal.com	independent.co.uk	songlyrics.com	bilansgratuits.fr
lolzepic.com	beliefnet.com	ps3hax.net	celebuzz.com	manageo.fr
dogshaming.com	fm-base.co.uk	theanswerbank.co.uk	danoah.com	butfootballclub.fr
geeksaresexy.net	toptenz.net	straightfromthea.com	mysavings.com	fipradio.fr
townhall.com	nydailynews.com	awkwardfamilyphotos.com	roadbikereview.com	livefoot.fr
gooddrama.net	x17online.com	uberhumor.com	vitalfootball.co.uk	funradio.fr
foxnews.com	fr.ign.com	abritel.fr	franceinfo.fr	
wickedlocal.com	cinetrafic.fr	experienceproject.com	vgchartz.com	
rasamalaysia.com	voici.fr	gamezone.com	textsfromlastnight.com	
yourtango.com	babycenter.fr	bootyarcade.com	salary.com	
dashhacks.com	sport.fr	uberhumor.com	thinkbabynames.com	
slashfilm.com	bilansgratuits.fr	cheezburger.com	southwales-eveningpost.co.uk	
dogshaming.com	manageo.fr	i4u.com	recipegirl.com	
relevantmagazine.com	butfootballclub.fr	clientsfromhell.net	visualnews.com	
mediatakeout.com	fipradio.fr	autocar.co.uk	globalsecurity.org	
dispatch.com	livefoot.fr	thefutoncritic.com	247wallst.com	
crushable.com	funradio.fr	49erswebzone.com	toyotanation.com	

Table 4.8: Crawling website list where prices are recovered

Table 4.9: Website list used to specialize bots in shopping

amazon.com
ebay.com
ikea.com
store.steampowered.com
gap.com
hm.com
nike.com
ticketmaster.com
forever21.com
autotrader.com
gamestop.com
victoriassecret.com
rakuten.com
sephora.com
jcrew.com
urbanoutfitters.com
ralphlauren.com
ae.com
mango.com
ulta.com

Table 4.10: Website summary statistics

Website	Freq.	Percent	Cum.
beliefnet.com	29	0.43	0.43
bilansgratuits.fr	13	0.19	0.63
boxingscene.com	354	5.30	5.93
butfootballclub.fr	88	1.32	7.24
cheezburger.com	58	0.87	8.11
cinetrafic.fr	13	0.19	8.31
coches.net	12	0.18	8.49
cougarboard.com	28	0.42	8.90
crosswordsolver.org	329	4.92	13.83
crushable.com	102	1.53	15.35
dashhacks.com	190	2.84	18.20
digitalspy.com	218	3.26	21.46
dispatch.com	65	0.97	22.43
dogshaming.com	179	2.68	25.11
eurweb.com	12	0.18	25.29
fipradio.fr	132	1.98	27.27
fm-base.co.uk	90	1.35	28.61
fr.ign.com	307	4.59	33.21
gooddrama.to	7	0.10	33.31
hunt4freebies.com	842	12.60	45.91
kidsgamesheroes.com	15	0.22	46.14
livefoot.fr	37	0.55	46.69
lucianne.com	65	0.97	47.67
mediatakeout.com	2	0.03	47.70
mobafire.com	1,112	16.64	64.34

Website	Freq.	Percent	Cum.
nesn.com	97	1.45	65.79
nydailynews.com	272	4.07	69.86
operationsports.com	108	1.62	71.48
pajiba.com	37	0.55	72.03
psu.com	108	1.62	73.65
rasamalaysia.com	52	0.78	74.42
relevantmagazine.com	28	0.42	74.84
searchquotes.com	261	3.91	78.75
sidereel.com	2	0.03	78.78
slashfilm.com	147	2.20	80.98
sport.fr	17	0.25	81.23
steamykitchen.com	24	0.36	81.59
stuff.tv	1	0.01	81.61
techeblog.com	52	0.78	82.39
thechive.com	584	8.74	91.13
toptenz.net	30	0.45	91.57
townhall.com	237	3.55	95.12
voici.fr	32	0.48	95.60
wickedlocal.com	28	0.42	96.02
x17online.com	178	2.66	98.68
yourtango.com	88	1.32	100.00
Total	6,682	100.00	

Table 4.11: Dsp summary statistics

Dsp	Freq.	Percent	Cum.
1rx.io	54	0.81	0.81
adlooxtracking	555	8.31	9.11
adnxs	734	10.98	20.10
adsafeprotected	146	2.18	22.28
adsrvr	443	6.63	28.91
alephd	275	4.12	33.03
basebanner	6	0.09	33.12
criteo	3	0.04	33.16
erne	2	0.03	33.19
mathtag	858	12.84	46.03
mediarithmics	15	0.22	46.26
mmtro	3	0.04	46.30
moatads	43	0.64	46.95
openx	1	0.01	46.96
quantserve	623	9.32	56.29
rfihub	2,356	35.26	91.54
tubemogul	11	0.16	91.71
turn	542	8.11	99.82
w55c	12	0.18	100.00
Total	6,682	100.00	

Table 4.12: Website categories

Categories	Freq.	Percent	Cum.
Arts and Entertainment	1761	26.35	25.02
Autos and Vehicles	12	0.18	25.20
Business and Industry	13	0.19	25.40
Computer and Electronics	190	2.84	28.24
Food and Drink	76	1.14	29.38
Games	1959	29.62	54.21
News and Media	834	12.48	68.21
People and Society	117	1.75	69.96
Reference	261	3.91	78.46
Shopping	842	12.60	91.07
Sports	597	8.93	100.00
Total	6,682	100.00	

Table 4.13: Website categories classification

Arts and Entertainment	Autos and Vehicles	Business and Industry	Computer and and Electronics
cheezburger.com cinetrafic.fr crushable.com digitalspy.com dogshaming.com eurweb.com fipradio.fr gooddrama.to mediatakeout.com pajiba.com relevantmagazine.com sidereel.com slashfilm.com thechive.com toptenz.net voici.fr x17online.com	coches.net	bilansgratuits.fr	dashhacks.com

Food and Drink	Games	News and Media	People and Society
rasamalaysia.com steamykitchen.com	crosswordsolver.org fr.ign.com kidsgamesheroes.com mobafire.com operationsports.com psu.com	dispatch.com lucianne.com nesn.com nydailynews.com sport.fr stuff.tv techeblog.com townhall.com wickedlocal.com	beliefnet.com yourtango.com

Reference	Shopping	Sports
searchquotes.com	hunt4freebies.com	boxingscene.com butfootballclub.fr cougarboard.com fm-base.co.uk livefoot.fr

Chapter 5

Conclusion

The thesis provides analysis of three new considerations that impact the online advertising markets. Chapters respectively question the role of accountability, profiling algorithms and personal information in online media markets. The conclusion focuses on the thesis results and provide future potential research question related to each chapters.

The first chapter analyzes how a technology that let advertisers verify advertising quality may modify market equilibria. It shows that without technology, the publisher cannot commit on its advertising quality and therefore under-advertises with respect to its optimum level. It also underlines that the introduction of such technology may be welfare increasing only if user disutility from advertising is not too high. Finally, the chapter reviews the link between such technology and advertising avoidance. The chapter is one of the first academic production to underlines the importance of transparency regarding online advertising purchases.

However, advertisers are concerned about market transparency on many other issues. A first example relates to online advertising fraud. According to a study by the Association of National Advertisers, 2016, "advertising bot fraud", where robots visit websites and are considered as real Internet users, represents around 40% of online traffic and have cost \$7.2 Billion to brands in 2016. In other words, content providers have been considering 40% of served ads as successfully served to an Internet user (sometimes even clicked) whereas it was in fact displayed to a robot. This type of fraud increases the uncertainty advertisers' face when purchasing ads. Bot fraud is one of the most common advertising fraud across the Internet. Other scams such as "ad-stacking", that stacks many advertising in one ad-slot, "pixel stuffing", that displays advertising in a single invisible pixel on the webpage, or "domain spoofing", where domain name of famous websites are faked to sell nonexistent ad-slot, have increasingly spread, degrading trust even more between advertisers and websites. Without fraud identifying technologies, many actors have incentives to set up online advertising fraud schemes. As pointed out by Bloomberg, October 2006 in a not so recent article, "*Google and Yahoo! at times passively profit from click fraud and, in theory, have an incentive to tolerate it*". Consequently, the market heavily invested in fraud detection technologies in the past few years. According to Mediatel, September 2017 Google is reportedly investing in invalid traffic detection in order to "*build trust in the advertising supply chain*. These technologies act like filters, identifying fake traffic from legitimate one. Using consumers panels based on the behavior of real humans (Conscore, Trusted Detection of Invalid Traffic and MRC, October 2015), detection technologies establish criterions that allows them to spot suspiciously robot behaviors. In response, more sophisticated robots have been developed in order to trick detection technologies, exhibiting hence a technological war between *inspectors* and *inspectees*. Other initiatives to clean the market from fraud have been carried out. Last year, the Interactive Advertising Bureau (IAB) introduced "ads.txt", a way to check websites characteristics and link before bidding on an ad-slot, hence insuring that they are no fake websites whose name domain have been spoofed. Despite its economic importance, online advertising fraud has been relatively unexplored by economists. Asdemir et al. (2008) is one of the only academic work to analyze to which extent advertising fraud impact the online advertising market. They underline how ad-financed search engine using an imperfect technology to identify fake traffic may end-up trimming off valid clicks. However the online advertising fraud is nowadays diversified and widespread across the Internet. Moreover, the introduction of technology to detect fraud allow advertisers to adjust their investment in real time when buying advertising on websites. All in all, as fraud and anti-fraud technologies modifies price and demand for advertising spaces, performing an economic analysis may help understand how it separately affects Internet users, advertisers and websites.

In a second chapter, the thesis analyzes how publishers are now able to understand advertising sensitivity of Internet users and tailor the number of ads to fit their preferences. The chapter details the implication of the introduction of such profiling technology on optimal advertising level, profits and welfare. It shows that depending on the audience elasticity to advertising, the profiling technology may foster or hinder Internet users from watching ads. At last, the chapter shows that the impact of such technology on welfare depends on the type of website, audience and the efficiency of the technology.

The chapter relies in the fact that profiling technologies are fueled by traces that Internet users leave on Internet (Estrada-Jiménez et al., 2017). In that sense, such technologies are therefore dependent upon the willingness of Internet users to disclose information. This appears to be relevant as many of them prefer to deter access to their personal information when browsing Internet. To do so, they use softwares such as Ghostery, Privacy Badger or uMatrix. Such behavior relates to the literature on privacy concern and advertising (see Turow et al. (2009), Tucker (2012) and Pew Research Center, Internet Users Don't like Targeted Ads, 2012). If Internet users start masking their browsing behavior, technologies may be less efficient and more likely to make mistakes. As an extension, an economic analysis would be relevant to measure the impact of user empowerment regarding personal information on the online advertising market.

More generally, as Internet users hide their traces, brands adapt their online advertising strategies, which may change how online content are financed. The different privacy policies considered by regulators try to strike a balance between protecting Internet users' right to privacy and insuring online content production. Such policies need careful economic analysis as they could radically change advertisers' willingness to pay and therefore websites profits.

In a final chapter, the thesis assesses how a privacy policy that let online users opt-out from behavioral targeting may affect prices of online ads sold at auction. The chapter present a new computational methodology that recovers advertising prices and let us estimate by how much prices change when a user opt-out from behavioral targeting. Firstly, the chapter finds that opting out from behavioral targeting may increase prices of ads sold at auctions, which is not in line with previous work from the literature. Secondly, the chapter also shows that opting-out from behavioral targeting does not uniformly impact websites, but that such effect depends on website type. These results draws important potential implications of how a privacy concern can impact prices and quantity of advertising on a website, as it provides an additional ground to assess the economic impact of an *opt-out* privacy policy.

However, recently implemented privacy regulations such as the General Data Protection Regulation and the ePrivacy directive intent to "give citizens back control over their personal data, and to simplify the regulatory environment for business" and are favoring *opt-in*. In order to leverage users' personal information, websites will have to recover users consent first in specifying the exact purpose of the collection as well as the exact data collected. Work such as Johnson (2013) tries to estimate how a privacy policy favoring an opt-in may affect prices and profits in online advertising markets. Other economic analysis must be carried out in order to assess the impact of such privacy policy and hence facilitate the choice of regulators.

The thesis present three chapters that leverage classical literature on online advertising. However, as online advertising markets quickly evolves, many new transformation may appear to be interesting for economists. For example, Hsu (2016) raises the fact the online advertising industry is largely driven by two actors: Facebook and Google. Firstly, according to Business Insider, April 2017, both actors concentrate 99% of revenue growth from digital advertising in the US in 2016 highlighting to which extent "*Google and Facebook dominate the industry presently*". Secondly, both firms are strongly present in all the online advertising ecosystem, from selling data to organizing programmatic auctions. Both points tend to underline a market concentration that may define the future of selling ads on Internet, and may therefore be relevant for economist.

Other authors underline that marketers are currently redirecting money on other media rather than wasting money on "classical" online advertising. For example, famous multiproduct firm *Proctor & Gamble* experimented that a cut of $100 millions in digital advertising spending in June 2017 had little to no impact on profits, arguing that online ads are largely ineffective (Wall Street Journal, July 2017). Such competition between media is transforming the competition landscape for consumers' attention and drive new advertising strategies to emerge. The understanding of these new strategies and how they shape the product markets would benefit from an economic analysis, as it may impact consumers surplus, welfare and competition.

Bibliography

Asdemir, Kursad, ÃŰzden Yurtseven, and Moin A. Yahya, "An Economic Model of Click Fraud in Publisher Networks," *International Journal of Electronic Commerce*, 2008, *13* (2), 61–90.

Estrada-Jiménez, José, Javier Parra-Arnau, Ana Rodríguez-Hoyos, and Jordi Forné, "Online advertising: Analysis of privacy threats and protection approaches," *Computer Communications*, 2017, *100*, 32 – 51.

Hsu, David, "The End of Digital Advertising as We Know It," Article November 2016.

Johnson, Garrett A., "The Impact of Privacy Policy on the Auction Market for Online Display Advertising," SSRN Scholarly Paper ID 2333193, Social Science Research Network, Rochester, NY May 2013.

Tucker, Catherine E., "The economics of advertising and privacy," *International Journal of Industrial Organization*, 2012, *30* (3), 326–329.

Turow, Joseph, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy, "Americans Reject Tailored Advertising and Three Activities that Enable It," Working Paper 1478214, SSRN September 2009.

Chapter 6

Résumé de la thèse

6.1 Introduction

Depuis le travail fondateur de Stigler (1961), la publicité est considérée comme un sujet de recherche important pour les économistes. L'approche développée dans ses travaux et plus tard par Nelson (1974) fut parmi les premières à effectuer une analyse économique de la publicité en tant qu'*informative*, alors généralement considérée comme *persuasive* dans la littérature. Stigler (1961) considère dans son analyse que la publicité est un pilier important de la transmission de l'information, permettant aux consommateurs d'être au fait de l'existence des différents produits sur le marché. Comme le souligne Bagwell (2007), un large pan de la littérature et de l'analyse économique de la publicité en ligne s'est construit sur l'analyse de sa propriété *informative* [1], fournissant ainsi la base nécessaire pour comprendre le rôle de la publicité, en particulier sur les marchés des médias ou se coordonne la rencontre entre annonceurs et audiences.

Aujourd'hui, une grande partie des dépenses publicitaires s'effectue sur Internet. Les recettes publicitaires en ligne ont continué de croître au cours des dix dernières années, atteignant 72 milliards de dollars en 2016, selon l'IAB, 2016. À ajouter à cela, Adage, 2017 souligne que ces dernières ont dépassé depuis 2016 celles de la télévision, devenant le premier investissement médiatique. Ces chiffres ont stimulé une nouvelle littérature sur les enjeux économiques de la publicité en ligne. Selon Peitz and Reisinger (2015), les progrès technologiques ont offert aux entreprises une occasion sans précédent d'informer les internautes de leur produit, à travers de nouvelles techniques telles que le ciblage comportementale. Cela à plusieurs conséquences.

[1] Parallèlement au développement de points de vue alternatifs. Voir aussi Renault (2015) pour une analyse plus complète sur le rôle de la publicité dans les marchés.

Premièrement, cela permet une meilleure correspondance entre publicité et contenu. La publicité par mots clés (*keyword advertising*), où les annonceurs peuvent choisir de diffuser leurs annonces en fonction de mot clefs saisis par un internaute sur un moteur de recherche, en est un parfait exemple. Deuxièmement, en exploitant les traces que les internautes laissent en ligne, les sites web sont en mesure d'afficher une publicité plus pertinente. Ceci s'illustre par le reciblage publicitaire (*retargeting*) où les annonceurs utilisent l'historique de navigation Internet des consommateurs pour afficher des annonces appropriées. Plus généralement, Goldfarb (2014) affirme que la publicité dans les médias en ligne présente une différence substantielle par rapport à la publicité classique : le coût du ciblage est réduit.

Ces nouvelles technologies génèrent des investissements importants dans le secteur de la publicité en ligne, mais génèrent également des externalités. Nous soutenons dans la thèse que ces dernières modifient radicalement les équilibres du marché. Pour illustrer cet argument, la thèse aborde trois nouvelles préoccupations économiques sur le marché de la publicité en ligne. Le premier a trait au manque de transparence sur la qualité des annonces publicitaires en ligne. Deuxièmement, la thèse analyse comment les nouveaux algorithmes de profilage utilisés par les plateformes pour adapter le nombre d'annonces aux préférences de chaque utilisateur modifient le comportement des sites web. Enfin, la thèse développe une nouvelle méthodologie permettant d'évaluer économiquement l'impact d'une régulation sur l'utilisation des données personnelles.

La thèse contribue à la littérature économique sur la publicité en ligne, analysant les marchés des médias comme des marchés bifaces (Rochet and Tirole, 2002; Anderson and Gabszewicz, 2006). Les plateformes (ici les fournisseurs de contenu) attirent les internautes (côté 1) et leur montrent une quantité de publicité (les annonceurs sont donc du côté 2). Assurer une bonne coordination du marché permet la rencontre des internautes avec les annonceurs sur le site web du fournisseur de contenu. La littérature économique récente fait l'hypothèse que les internautes considère la publicité comme une nuisance, présentant des externalités de réseau négatives provenant de la quantité de publicité affichée sur le site web. Par conséquent, plus il y a d'annonces sur un site web, moins les internautes le visitent, même si ces annonces peuvent avoir des propriétés informatives. Dans ce contexte, le fournisseur de contenu doit faire face à un défi : afficher un grand nombre de publicité tout en repoussant une partie des consommateurs, ou réduire le nombre d'annonces, générant ainsi moins de nuisance et attirant davantage de consommateurs.

Chaque chapitre fourni une analyse économique de nouveaux défis auxquels le marché de la publicité en ligne est confronté. Ils montrent que ces transformations, ancrées dans une nature biface, présentent des mécanismes économiques sophistiqués. Des méthodologies théoriques et empiriques sont développées tout au long de la thèse pour analyser comment ces transformations impactent le marché de la publicité en ligne

Le premier problème abordé dans la thèse a trait au manque de transparence des achats publicitaires en ligne. D'une part, l'introduction de mécanismes automatisés complexes de vente de publicité connus sous le nom de *programmatique* a accru l'efficacité de vente - comme indiqué par le Rapport d'Econsultancy, 2015. D'autre part, selon The Guardian, 2017 et Marketingweek, 2017, cette complexité rend difficile la vérification de la qualité des emplacements publicitaires vendus aux enchères. Ceci est d'autant plus important que des milliards d'annonces sont venus chaque jours et que les annonceurs. Être en mesure de répondre aux attentes des annonceurs - concernant la qualité des emplacements publicitaires - semble constituer un véritable défi pour le marché.

Le marché de la publicité en ligne s'est saisi de trois composantes de cette " qualité média". La première concerne la visibilité des publicités en ligne. Les annonceurs ont en effet pu acheter des publicités des publicités vues par aucun Internautes. Différentes analyses telles que celle de ComScore, 2013 ont souligné que près de la moitié des annonces "display" [2] ne sont pas vues par les internautes, ne se situant pas dans les limites de la *page active* (la partie de la page affichée par le navigateur). Deuxièmement, d'autres études soulignent que même si les publicités sont visibles, elles sont principalement vues par des robots et non par des humains. Les rapports de l'Association of National Advertisers (ANA), 2016 et de Adloox, 2017 indiquent que la fraude par les robots publicitaires est responsable d'une perte comprise entre 7,2 milliards de dollars et 12 milliards de dollars pour 2016 sur le marché de la publicité en ligne. Enfin, la notion de *brand safety* a récemment secoué le secteur de la publicité en ligne. Selon Business Insider, 2017, certains annonceurs ayant ayant conduit des campagnes publicitaires sur YouTube ont pu constater que leurs annonces étaient diffusées avant la lecture de contenus inappropriés (vidéos terroristes ou racistes, par exemple).

Ces trois illustrations - visibilité, fraude et brand safety - soulignent comment l'incertitude autour de la vente d'espaces publicitaires joue un rôle important dans la disposition à payer des

[2]les annonces affichées sur des sites web ou des applications mobiles

annonceurs. Le marché a développé des technologies afin de palier à ces problèmes. Néanmoins, il n'est pas facile d'évaluer comment l'introduction de ces technologies transforment les équilibres de marché. D'une part, cela peut inciter les annonceurs à ajuster leurs investissements. D'autre part, la mise en pratique de ces technologies pourrait dégrader l'expérience des internautes, préférant par conséquent éviter les publicités.

Le contexte de la thèse est introduit dans le **Chapitre 1**. Il y est également présenté la problèmatique ainsi que les principaux résultats.

Le **Chapitre 2** présente un modèle théorique qui analyse comment l'introduction d'une technologie permettant de vérifier la qualité des espaces publicitaires transforme le marché de la publicité en ligne.

Le chapitre se penche particulièrement sur le cas de la visibilité de la publicité. Cet aspect semble être très important pour l'économie de la publicité en ligne, en particulier dans le contexte des campagnes de *branding* où les publicités sont achetées *par impression* (CPM) et non pas *au clic* (CPC) [3], les annonceurs payant des annonces lorsque "servie"(c'est-à-dire affiché) sur le site web. Selon Comscore, T1 2017, la visibilité des annonces n'a pas beaucoup changé en 4 ans, puisque environ 50% des annonces diffusées sont effectivement vues par les internautes, ce qui pose logiquement problèmes aux annonceurs paient pour les publicités.

Pour surmonter cet obstacle, le marché a déjà mis au point des technologies évaluant le niveau de visibilité des annonces. À l'aide de mesures géométriques, les annonceurs récupèrent l'emplacement des annonces dans la page web et détermine la durée et la quantité d'affichage de leur annonce dans la vue active. Pour favoriser l'adoption de cette technologie, Internet Advertising Bureau (IAB) et le Media Rating Council (MRC) ont introduit des normes de visibilité, qualifiant de "visionnées" toute annonce "display"dont au moins 50% des pixels sont affichés dans la vue active de la page web pendant au moins une seconde (IAB Guidelines, 2014). [4]

Le chapitre analyse comment l'introduction de cette technologie impacte le marché de la publicité en ligne. Premièrement, les internautes ont tendance à voir plus d'annonces lorsque les technologies permettant aux annonceurs de vérifier la visibilité sont introduites. En ef-

[3] Les annonceurs distinguent deux types de campagnes : des campagnes de promotion de la marque (généralement payées pour être affichées) et campagnes de performance ayant pour objectif d'inciter à l'achat en ligne (généralement payées lorsque les internautes ont cliqué sur l'annonce).

[4] Pour les bannières non "display", les critères de mesure de la visibilité peuvent changer. Par exemple, les vidéos publicitaires sont considérées comme vues lorsqu'au moins 50% des pixels ont été affichés dans la vue active de la page web pendant au moins deux secondes.

fet, en présence de technologies, les fournisseurs de contenu peuvent s'engager à pratiquer un niveau spécifique de visibilité des annonces. De plus, dans un contexte où la publicité est perçue comme une nuisance par les utilisateurs d'Internet, l'introduction cette technologie est bénéfique pour le bien-être social si la nuisance publicitaire n'est pas trop élevée. Deuxièmement, comme l'introduction de ces technologies peut également inciter les internautes à éviter les annonces (en utilisant par exemple des bloqueurs de publicité). L'introduction de telles technologies augmentant (faiblement) le nombre d'annonces sur un site web et par conséquent la nuisance publicitaire, les internautes sont plus susceptibles d'utiliser un bloqueur de publicité. [5]

Le chapitre montre que l'introduction de telles technologies réduit l'incertitude sur le marché en améliorant la transparence sur la qualité de la publicité, tout en mettant les sites web sous pression. D'une part, l'efficacité des éditeurs est désormais contrôlée par les annonceurs lors des campagnes, ce qui les incite à augmenter le nombre d'annonces à diffuser aux internautes. D'autre part, l'affichage d'un plus grand nombre d'annonces peut inciter les internautes à utiliser des bloqueurs de publicité. Ce dernier point est surveillé de près par le secteur de la publicité en ligne, car elle nuit aux bénéfices des fournisseurs de contenu et des annonceurs, ce qui nous amène au chapitre suivant.

La deuxième préoccupation analysée dans la thèse traite de la nuisance publicitaire sur le marché des médiatique en ligne. Selon une étude publiée par WordStream, seuls 0,9% des annonces affichées sur Facebook sont cliquées par les utilisateurs. Au-delà de se désintéresser de la publicité, les internautes préfèrent simplement les éviter : l'étude de Yougov souligne par exemple que 61% des adultes n'aiment pas la publicité en ligne. Par conséquent, les consommateurs installent massivement des technologies permettant d'éviter les publicités en ligne (PageFair, 2016).

Les internautes n'aiment pas la publicité - ou du moins la publicité telle qu'elle est pratiquée par les fournisseurs de contenu en ligne, comme souligné par plusieurs études telles que Yougov, 2016. Pour cette raison, de plus en plus d'internautes adoptent des bloqueurs de publicités, tel que Adblock Plus, lors de leur navigation en ligne. PageFair, 2016 souligne qu'en décembre 2016, au moins 380 millions de mobiles et 236 millions d'ordinateurs étaient équipé d'un bloqueurs de publicité, soit 11% de la population Internet mondiale. En outre, des navigateurs tels que Brave, Firefox Focus ou la récente version de Chrome lancé en 2018 (qui représente environ 60% de part

[5] Ceci est principalement dû à la structure biface du marché de la publicité en ligne.

de marché selon NetMarketShare 2017) ont intégrés des bloqueurs de publicité selon Digiday, avril 2016.

Avec les bloqueurs de publicités, la nature biface du marché de la publicité en ligne est remis en cause, le site Internet ne générant aucun revenus. Par conséquent, ceux-ci commencent à rechercher de nouveaux modèles d'affaires.

Par exemple, le marché tente d'améliorer la correspondance entre le contenu publicitaire et les préférences des internautes. Une telle solution a pour but d'améliorer l'expérience utilisateur et ainsi d'inciter les internautes à ne plus éviter les annonces. Dans cette optique, de nouvelles technologies offrent aux fournisseurs de contenus la capacité de personnaliser entièrement leur expérience en exploitant les informations personnelles des internautes. Le type et le nombre de contenus peuvent donc être adaptés à chaque visiteur. Des technologies de suivi (comme les cookies) sont exploitées pour récupérer les traces des internautes. Ces traces peuvent prendre différentes formes : elles peuvent être le résultat de la lecture d'un article sur un site sportif, du visionnage d'une bande annonce de film ou de la visite d'une page de marque sur un réseau social. À l'aide de ces traces, les algorithmes infèrent les préférences des internautes et proposent alors une expérience de navigation adaptée.

Une telle pratique peut ressembler à une discrimination par les prix au premier degré. Dans cette situation, les entreprises peuvent afficher le nombre maximal d'annonces qu'un internaute est prêt à tolérer. Avec ces nouveaux outils, l'expérience de navigation peut être totalement personnalisée. Or, ils modifient également la quantité de publicités servies sur Internet, considérée comme un indicateur concurrentiel important par les économistes. De plus, l'efficacité de la technologie repose sur les traces en ligne des utilisateurs d'Internet, qui dans certains cas ne sont pas disponibles.

Le **chapitre** 3 traite des problèmes d'évitement publicitaire et analyse comment l'utilisation de technologies de profilage a des conséquences économiques importantes. Plus précisément, considérant que les internautes sont hétérogènes quant à leurs préférences en matière de publicité, le chapitre analyse les implications de l'introduction d'une technologie capable de capturer cette sensibilité et personnalise la publicité en conséquence.

Le chapitre étudie, à l'aide d'un modèle théorique, comment l'introduction d'une technologie de profilage permettant d'adapter l'expérience publicitaire aux préférences des internautes, peut modifier les équilibres de marché. Il considère deux types d'utilisateurs : certains consommateurs ont une préférence *forte* pour les publicités, tandis qu'un ensemble complémentaire présente une préférence *faible*. Nous supposons que les utilisateurs à faible préférence sont plus facilement gênés par la publicité, donc plus susceptibles de les éviter. La technologie de profilage classe les utilisateurs en fonction d'un signal sur leur préférence. La précision du signal est considérée comme exogène et peut être imparfaite. [6]

Tout d'abord, le chapitre montre qu'un fournisseur de contenu adopte toujours une technologie quand elle est *parfaite* (une technologie qui ne fait pas d'erreurs en classant les utilisateurs) car il augmente ses profits aussi bien sur les internautes ayant une préférence forte pour la publicité (en leur montrant plus d'annonces) que sur ceux avec une préférence faible (en leur montrant moins d'annonce ce qui les incite à visiter le site sans bloqueurs de publicité).

L'analyse devient plus complexe lorsque le site web adopte une technologie *imparfaite* (une technologie suffisamment efficace peut être adopté par le site web tout en faisant des erreurs lors de la classification des utilisateurs). Le chapitre considère l'efficacité de la technologie comme exogène, alors qu'elle peut être liée à l'investissement dans une méthode statistique ainsi qu'à la qualité et au volume des informations disponibles sur les internautes. Nous étudions ce dernier point dans le troisième chapitre de la thèse.

La troisième contribution de la thèse traite de l'utilisation des traces des internautes dans le but d'améliorer l'efficacité de la publicité ciblée. Dans un rapport de 2015 du Pew Internet Research center, 93% des personnes interrogées affirment qu'il est important de pouvoir contrôler *qui* obtient leurs informations personneles. Ces préoccupations soulignent la nécessité de réglementer les pratiques des entreprises sur Internet. D'une part, les industries axées sur la donnée, telles que la publicité en ligne, exploitent les informations des internautes. D'autre part, une telle pratique peut nuire à la protection de la vie privée et être perçu comme une nuisance par les citoyens

Pour équilibrer ces deux effets, les régulateurs ont envisagé différentes politiques publique. Par exemple, le régulateur américain a mis en place une stratégie d'*opt-out*, dans laquelle les annonceurs peuvent suivre les utilisateurs par défaut, alors même que ces utilisateurs peuvent choisir de ne pas être suivi. Sur ce même terrain, le règlement européen sur la protection des données (RGPD) récemment adopté a pour objectif de mettre en oeuvre une politique d'*opt-in*, ce qui empêche alors les annonceurs de suivre les internautes, à moins d'en recueillir le

[6] Si le signal est tout le temps correct, cela signifie qu'il classifiera parfaitement tous les internautes

consentement explicite.

Ces règlementations impactent largement les profits du secteur de la publicité en ligne. Ainsi, l'absence de ciblage publicitaire peut avoir deux conséquences. Tout d'abord, cela peut réduire le taux de clic, car la correspondance entre l'annonce et l'utilisateur est de qualité moindre. Deuxièmement, cela peut également inciter les internautes à éviter les publicités. Plus généralement, l'absence de possibilité de ciblage lors d'une campagne publicitaire est censée réduire la volonté à payer des annonceurs et donc les prix des publicités.

Le **chapitre 4** évalue comment l'adoption d'une politique publique favorisant l'opt-out peut avoir un effet sur les prix des publicités, dans le contexte d'une vente d'espace publicitaire dite programmatique. [7] Des travaux antérieurs ont déjà établi qu'une telle politique publique tend à réduire les prix des publicités et donc les revenus sites.

Le chapitre présente une nouvelle méthode permettant de capturer les prix des annonces avant et après l'opt-out de l'utilisateur. À partir de cette méthode, le chapitre produit trois résultats importants.

Premièrement, la méthodologie récupère précisément les prix des publicités avant et après un opt-out du ciblage publicitaire, ce qui n'était pas possible dans les travaux antérieurs présent dans la littérature. Une telle méthodologie peut facilement être étendue pour tester l'impact d'autres politiques publiques affectant le comportement des internautes. Deuxièmement, le chapitre observe qu'une absence de ciblage publicitaire augmente légèrement le prix de la publicité en ligne vendue aux enchères. Troisièmement, nous constatons également qu'une absence de ciblage publicitaire réduit significativement le nombre d'annonces vendues aux enchères au cours d'une expérience, ce qui souligne une concurrence entre les mécanismes de vente publicitaire sur les sites web.

Une telle étude s'avère particulièrement pertinente dans la mesure où le RGPD, ainsi que la prochaine directive ePrivacy, visent à mettre en oeuvre une régulation "opt-in". Dans ce cas précis, les sites web et les annonceurs doivent obtenir un consentement clair des internautes avant de tirer parti des données des utilisateurs. La méthodologie développée dans le chapitre pourrait également permettre de comprendre les effets économiques d'une régulation opt-in sur le marché de la publicité en ligne.

[7] La vente de publicité programmatique désigne un processus de vente automatisé ou un espace publicitaire est mis aux enchères. L'annonceur avec l'enchère la plus haute gagne le droit d'afficher sa publicité sur le site visité par l'internaute

Les parties suivantes résument les résultats des chapitres présentés dans cette introduction.

6.2 Chapitre 1 : Mesure de visibilité publicitaire et marché de la publicité en ligne - une analyse économique

Les dépenses publicitaires en ligne ne cessent d'augmenter. Selon eMarketer, le Royaume-Uni est devenu en 2015 le premier pays au monde où les médias numériques représentent 50% des dépenses publicitaires. [8] Aux États-Unis, également selon les prévisions d'eMarketer, avec environ 77 milliards de dollars, les dépenses publicitaires en ligne dépassent celles des publicités télévisées et ce principalement grâce au mobile qui représente plus de 50% du marché (eMarketer, 2016) .

Le développement fulgurant du marché du mobile n'explique cependant pas à lui seul la croissance des dépenses de publicité en ligne. L'avènement de la publicité programmatique et la possibilité de collecter des données précises sur les consommateurs mais aussi sur les impressions [9] permet aux annonceurs web d'automatiser l'achat et la vente de publicité et de réaliser un ciblage personnalisé et efficace. Les médias numériques deviennent ainsi plus efficaces que la télévision et la presse écrite pour évaluer l'impact d'une publicité dans la décision d'achat ou la notoriété d'une marque.

Toutefois, les promesses de la publicité en ligne sont aujourd'hui contestées. Ces promesses reposent en effet sur le postulat que les publicités diffusées sur les sites sont *vues* par les internautes, c'est-à-dire contenues dans l'espace visible de la fenêtre du navigateur, basé sur des critères préétablis tels que le pourcentage des pixels de la publicité et la durée de leur affichage dans cet espace. (IAB Europe, 2015). Visible dans ce contexte signifie simplement que les internautes ont la possibilité de voir l'annonce.

Ce postulat simple se trouve toutefois contesté par des entreprises telles que Google, comScore ou encore Nielsen qui analysent quotidiennement des milliards d'impressions de campagnes provenant de milliers d'éditeurs : une statistique publiée par comScore en 2013 indique

[8]Les dépenses publicitaires en ligne aux États-Unis s'élèvent à 77 milliards de dollars (RelevanceeMarketer, "Le Royaume-Uni réalisera une première mondiale alors que la moitié des dépenses publicitaires des médias passe au numérique, "2015).

[9]L'affichage d'une annonce dans une page est appelé une impression.

en effet que la moitié de l'inventaire des éditeurs n'est en réalité pas vue par les internautes. [10] En 2016, le taux de visibilité des annonces dans la plupart des pays du monde reste relativement faible, entre 40 et 50%. [11] Facebook, le réseau social qui attire la majeure partie des investissements publicitaires, ne fait pas exception et se trouve également sujet aux critiques (Business Insider, 2016).

La visibilité des annonces est donc devenue au cours des derniers mois l'une des problématiques majeures des annonceurs (Wall Street Journal, 2016). [12] Et leurs préoccupations sont parfaitement légitimes. Dans le cas de campagnes de stratégie de marque, les annonceurs paient les publicités en fonction du nombre d'impressions sur le site d'un éditeur (appelé "Coût par mille" (CPM)). [13] Mais comme la moitié des publicités achetées par les annonceurs ne sont jamais vues par les internautes, ceux-ci gaspillent potentiellement la moitié de leur budget. Selon les estimations de Meetric, en 2016 au Royaume-Uni, les annonceurs ont gaspillé 600 millions d'euros en annonces non vues par les internautes (Meetrics, 2017). Par conséquent, de plus en plus d'annonceurs exigent de ne payer que pour des impressions *vue* et non plus seulement pour les impressions simplement “affichées”s ur une page web. Une nouvelle monnaie d'échange émerge : le “viewable” CPM (vCPM ou CPM visible) qui tarife les annonces en fonction du nombre d'impressions vues par les utilisateurs Internet, et non simplement affichées.

L'utilisation de technologies permettant de mesurer la visibilité des annonces diffusées par les éditeurs accompagnée par l'émergence d'une nouvelle monnaie d'échange pourrait entraîner des changements structurels au sein de l'économie de la publicité en ligne. Tout d'abord les

[10] Différentes raisons expliquent pourquoi les impressions ne sont pas vues par les internautes. Les comportements de navigation présentent de nombreuses possibilités d'évitement telles que le défilement de la page, les dimensions de la fenêtre de navigation ou bien l'utilisation d'un bloqueur de publicité. Les éditeurs peuvent choisir d'ajuster la visibilité des publicités afin de préserver une bonne expérience utilisateur.

[11] La visibilité de la publicité n'est pas un problème nouveau dans les médias, mais en raison de la taille des marchés en ligne, le problème persiste, se développe même et devient une menace sérieuse pour le secteur de la publicité. Pour les médias papier, la probabilité qu'un lecteur voit réellement une annonce sur une page donnée n'est pas précise (à l'exception de l'utilisation de QR codes). En ce qui concerne la télévision, une publicité est censée être vue à partir du moment où il y a une personne dans une pièce avec une télévision allumée. La mesure n'est cependant pas parfaite : un téléspectateur peut sortir de la pièce pendant les pauses publicitaires ou bien passer les publicités si le programme est enregistré. Mais il existe une “occasion de voir”.

[12] Au cours de l'édition 2016 de l'annuelle Conférence sur les médias sociaux et numériques, le PDG de l'association américaine des annonceurs nationaux a évoqué les préoccupations des "quatre grands" médias : blocage des publicités, fraude publicitaire, transparence des médias et visibilité/mesure.

[13] Les campagnes utilisent un autre indicateur bien connu : le taux de clics, c'est-à-dire le pourcentage d'impressions ayant généré des clics. Dans ce chapitre, nous concentrons l'étude uniquement sur les campagnes en ligne basées sur de la stratégie de marque.

éditeurs doivent repenser leur site web [14] pour rendre les annonces publicitaires plus visibles et satisfaire les demandes des annonceurs. Une large partie de l'inventaire actuel avec une visibilité très faible ne pourrait donc plus être vendue, ou à un prix inférieur, ce qui devrait réduire les flux de revenus de certains éditeurs. Mais comme dans ce cas précis, ces derniers peuvent avoir moins d'inventaire à vendre, certains tarifs peuvent également augmenter (inventaire premium), ce qui transforme la concurrence entre annonceurs pour une visibilité optimale des annonces. [15] De plus, à mesure que les sites web sont réaménagés pour augmenter l'espace réservé aux annonces visibles au détriment du contenu éditorial, l'audience pourrait diminuer et le prix des annonces baisser en conséquence. L'incidence économique de l'introduction d'une technologie de visibilité est donc partagée entre la perte d'expérience utilisateur et le gain de revenus de la publicité en ligne.

Le chapitre a pour objectif d'analyser l'impact qu'ont les technologies de mesure de visibilité sur l'économie de la publicité en ligne. L'analyse se base sur un modèle de marché biface sur lequel un éditeur en monopole affiche un contenu éditorial pour attirer les internautes d'un côté, et vendre des espaces publicitaires aux annonceurs de l'autre. Dans ce modèle inspiré de Anderson and Coate (2005), l'éditeur de contenu est financé uniquement par les annonceurs pour afficher des publicités qui sont vues par les utilisateurs comme une nuisance. Nous proposons de comparer deux situations : dans un premier temps, les annonceurs ne disposent pas d'une technologie de mesure de visibilité sur le site web de l'éditeur ; ils ne peuvent ainsi qu'anticiper un niveau global de visibilité de leurs annonces. Dans un second temps, les annonceurs disposent d'une telle technologie et peuvent mesurer avec précision la visibilité de leurs publicités.

En premier lieu, le modèle souligne que l'utilisation d'une technologie de mesure de la visibilité tend à augmenter le nombre de publicités visibles et ainsi les bénéfices des éditeurs et des annonceurs ; mais en retour, cela dégrade l'expérience utilisateur. En somme, l'adoption de tels outils améliore la situation des éditeurs et des annonceurs au détriment de celle des internautes. L'introduction d'une technologie de mesure de visibilité peut augmenter le bien-être total sur le marché seulement si la nuisance des publicités est faible et/ou si la qualité du contenu proposé par l'éditeur est élevée. Par ailleurs, le modèle montre que l'éditeur réalisera de plus gros

[14] The Guardian est un exemple typique. Le site web du journal britannique a repensé son site web afin de rendre les emplacements d'annonce publicitaires mieux visibles (The Guardian, 2016)

[15] Selon Quantcast en 2016, l'inventaire avec une visibilité supérieure à 75% peut aller jusqu'à deux fois plus cher que la moyenne (Quantcast, En route vers la visibilité, 2016).

investissement dans la qualité des contenus qu'il propose en présence de technologie.

Ensuite, nous étendons le modèle aux bloqueurs de publicités. Dans notre configuration, adopter des bloqueurs de publicités représente en effet une réponse face à la baisse d'utilité des internautes. Dans ce cas, l'éditeur se trouve contraint par les deux faces du marché : d'une part il doit augmenter le nombre de publicités afin de satisfaire les demandes des annonceurs et d'autre part, il doit le réduire pour ne pas amener les internautes à installer des bloqueurs de publicités.

Le développement de l'analyse amène à la construction d'un modèle de concurrence au goulot d'étranglement, dans lequel deux éditeurs se font concurrence pour attirer l'attention des utilisateurs. Notre modèle montre que dans le cas où les deux éditeurs sont asymétriques concernant leur capacité de visibilité (c'est-à-dire le niveau de visibilité maximale qu'un éditeur peut définir pour une annonce), l'introduction d'une technologie de mesure de visibilité permet de rétablir une concurrence saine.

Ce chapitre contribue à la littérature économique sur la publicité en ligne en deux points. Premièrement, il constitue la première analyse économique de la question de la visibilité qui, si elle se trouve largement débattue dans le secteur de la publicité, reste cependant absente de la recherche universitaire. Deuxièmement, notre article enrichit la littérature sur le fonctionnement du marché des médias en ligne (Peitz and Reisinger, 2015) ainsi que celle sur l'efficacité des publicités correspondantes (Goldfarb and Tucker, 2011; Manchanda et al., 2006; Lambrecht and Tucker, 2013). Dans les contributions précédentes, les internautes préféraient les publicités lorsqu'elles étaient ciblés (de Cornière, 2016; Johnson, 2013b) et, paradoxalement, n'appréciaient pas les publicités lorsqu'elles étaient trop intrusives, entraînant des variations de demande. Or, il existe de nombreuses raisons pour lesquelles les publicités ciblées ou intrusives ne sont jamais vues par les internautes, quel que soit leur préférence pour la publicité. Il apparaît donc crucial de prendre en compte la visibilité des publicités en ligne, car celles-ci même si elles ne sont pas ou seulement partiellement visionnées sont pour autant toujours payées par les annonceurs.

6.3 Chapitre 2 : Analyse économique de l'introduction d'une technologie de discrimination par la nuisance publicitaire

La publicité constitue le mécanisme principal de financement des canaux médiatiques tels que la télévision, la radio ou la presse écrite. Il représente également l'essentiel des revenus des plateformes médiatiques sur Internet. Au troisième trimestre 2016 par exemple, les annonceurs américains ont investi 17,6 milliards de dollars dans la publicité en ligne, selon l'IAB [16], ce qui constitue une augmentation de 20% par rapport à la même période en 2015.

Si la publicité en ligne ne cesse de se développer, elle est cependant de moins en moins tolérée par les internautes qui la considèrent comme responsable de la dégradation de leur expérience en ligne (Manchanda et al., 2006; Goldfarb and Tucker, 2011). Face à cette situation, les utilisateurs ont de plus en plus recours à des technologies d'évitement de publicité (*Anti Advertising Technologies* - AAT). [17] En 2016 c'est plus de 26% des internautes américains qui utilisent un bloqueur de publicité (ce qui représente environ 69 millions d'utilisateurs), soit un bond de 34% par rapport à l'année précédente. [18] Les AAT peuvent être installés sur un navigateur web. De plus, les nouvelles versions de Chrome et de Firefox les proposent même directement par défaut. [19] La fonctionnalité de blocage de la publicité filtre automatiquement les annonces telles que les fenêtres contextuelles et la lecture automatique de vidéos.

Cette adoption croissante d'AAT par une large partie des internautes montre que la publicité est majoritairement perçue comme une nuisance. Cette observation est cependant à nuancer : une proportion encore importante de consommateurs continue de consulter des sites web et de cliquer sur des publicités. Le goût et les préférences en matière de publicité varient donc d'un utilisateur à l'autre.

Afin de tenir compte du goût des utilisateurs pour la publicité, les éditeurs de contenu en ligne adoptent des stratégies diverses. Par exemple, Tag (2009) analyse une plateforme en monopole proposant aux internautes de consommer les contenus avec de la publicité, ou en prenant une option d'abonnement sans publicité. L'auteur montre qu'introduire une option d'abonnement incite le site à augmenter le nombre de publicité par page (pour les utilisateurs qui décident de rester sur l'option publicitaire). En fin de compte, le surplus du consommateur total diminue avec l'introduction d'un abonnement, alors que le surplus de l'annonceur augmente. Anderson and Gans (2011) considère la possibilité pour les internautes d'utiliser une technologie coûteuse permettant d'éviter les publicités, ce qui supprime la nuisance publicitaire au même titre qu'un

[16] IAB, 2016.
[17] Adblock Plus ou uBlock sont des exemples d'AAT populaires sur Internet.
[18] eMarketer, 2016.
[19] ArsTechnica, Rapport : Google Vous ajouterez un bloqueur de publicité à toutes les versions du navigateur web Chrome, le 19 avril 2017.

abonnement. Ils confirment les résultats de Tag (2009) et montrent également que l'adoption d'AAT peut réduire la qualité des contenus proposés par l'éditeur.

Les stratégies de plateforme analysées dans Tag (2009) et Anderson and Gans (2011) reposent sur l'idée que les publicités ne sont perçues que comme une nuisance et que les préférences des consommateurs en matière de publicité ne sont pas directement observables. En choisissant de s'abonner ou pas, les internautes révèlent alors indirectement leur goût pour la publicité. Or les technologies de profilage récentes permettent aux plateformes de déduire le goût des utilisateurs en ligne pour la publicité. Sur la base du comportement passé (clic sur les publicités, nombre de publicités vues, etc.), les plateformes peuvent déterminer si un internaute est plus ou moins sensible à la publicité et adapter en conséquence le nombre d'annonces diffusées afin de l'inciter à ne pas éviter la publicité. En effet, les plateformes numériques utilisent des technologies de profilage pour limiter le nombre d'annonces visionnées par le consommateur, la fréquence d'exposition d'un consommateur à une annonce donnée et les annonces relatives aux produits déjà achetés. C'est ainsi que Facebook et Snapchat limitent le nombre d'annonces que dans les fils d'actualités [20] Cette technique s'appelle le "frequency capping" (Buchbinder et al., 2011).

Dans ce chapitre, nous analysons comment une technologie de profilage permettant de classifier les utilisateurs en fonction de leur préférence pour la publicité impacte les stratégies des plateformes, le surplus des internautes et des annonceurs et le volume d'annonces diffusées sur le marché. Nous développons un modèle en lien avec ceux développés dans Anderson and Coate (2005), Tag (2009) et Anderson and Gans (2011), où une plateforme en monopole déduit la sensibilité publicitaire des internautes (i.e. combien chacun d'eux est prêt à voir de publicité). Pour mesurer l'impact d'une telle technologie, nous basons le modèle sur trois caractéristiques principales. Tout d'abord, la publicité n'est pas nécessairement perçue comme une nuisance. Deuxièmement, les utilisateurs sont hétérogènes en sensibilité publicitaire et deux types sont distingués : une proportion de consommateurs avec une préférence *forte* (γ_l) pour la publicité tandis qu'une proportion correspondante présente une préférence *faible* (γ_h). Enfin, nous permettons à la plateforme de classifier les utilisateurs en fonction de leur goût pour les publicités en utilisant une technologie de profilage dont la précision est exogène et dont l'efficacité peut aller de l'absence d'identification à l'identification parfaite. En effet, la technologie de profilage

[20] TechCrunch, 2016.

n'est pas toujours *parfaite*, c'est-à-dire qu'elle n'identifie pas à 100% la sensibilité des internautes pour la publicité. On parle dans ce cas-là d'une technologie *imparfaite*. Cette dernière dépend en effet de la disponibilité des informations personnelles des utilisateurs ainsi que de la sophistication de la technologie elle-même. Dans ce cas, les internautes mal classifiés observent un niveau de publicité inadapté à leurs préférences en matière de publicité, et préfèrent par conséquent éviter la publicité.

Le modèle présenté dans le chapitre souligne également que l'utilisation d'une technologie de profilage engendre des implications stratégiques. Premièrement, la plateforme peut choisir de ne pas utiliser la technologie de profilage. En effet, si celle-ci n'est pas suffisamment efficace, la plateforme ne peut pas déduire avec précision les préférences publicitaires de chaque utilisateur en ligne. Dans le cas où cette précision est vraiment faible, la plateforme préfèrera fixer un niveau de publicité unique pour tous les utilisateurs du site, qu'ils aiment la publicité ou non. Si inversement (i. e. si la précision de la technologie et suffisamment bonne), la plateforme choisira d'implémenter la technologie de profilage dans sa stratégie publicitaire. Cette situation survient lorsque la technologie est assez précise, de sorte que la probabilité de prédire correctement les préférences des utilisateurs est suffisamment élevée. Dans ce cas, la technologie adapte alors le nombre de de publicité à chaque utilisateur, augmentant ainsi ses bénéfices. Il est alors important de souligner qu'une technologie de profilage peut être adoptée par la plateforme alors qu'elle produit un nombre conséquent d'erreurs de classification. En effet, une technologie de profilage *parfaite* classe toujours parfaitement les internautes en fonction de leur type. Cela se produirait en présence d'une technologie de profilage parfaitement formée et à jour. Cependant, une technologie efficace mais *imparfaite* commet des erreurs, même si elle génère des profits plus importants pour la plateforme. Cela est susceptible de se produire sachant qu'une technologie de classification peut être confrontée à plusieurs obstacles techniques. Ainsi, même si l'utilisation de la technologie augmente les revenus de la plateforme, elle risque de ne pas adapter correctement le niveau de publicité aux préférences de chaque internaute. Cela engendre alors des répercutions importantes sur les bénéfices de la plateforme : un utilisateur classé à tort dans la catégorie "préférence faible pour la publicité" verra un nombre de publicité inférieur à ses attentes, réduisant ainsi les possibilités de profit pour la plateforme ; à l'inverse, un utilisateur classé à tort dans la catégorie "préférence forte pour la publicité"verra un nombre trop important de publicités, ce qui l'incitera à mettre en place des stratégies d'évitement.

Ensuite, l'utilisation par la plateforme d'une technologie de profilage modifie le nombre total de publicités diffusées sur le marché. Lorsque cette technologie est parfaite, le volume de publicités affichées aux internautes se révèle toujours supérieur comparé au cas où la technologie est absente. La plateforme peut en effet diffuser davantage d'annonces publicitaires aux utilisateurs qui ont un goût prononcé pour la publicité, tout en proposant un niveau approprié aux utilisateurs qui les évitaient auparavant. Ce résultat s'inverse lorsque la plateforme utilise une technologie imparfaite.

Enfin, nous montrons qu'introduire une technologie de profilage parfaite augmente toujours le bien-être lorsque la plateforme est confrontée à une audience majoritairement composé d'internaute avec une préférence forte pour la publicité. En effet, les gains pour les annonceurs et la plateforme compensent la perte potentielle du surplus des internautes. Cependant, l'analyse devient plus complexe lorsque 1) la technologie est parfaite mais que la plateforme est confrontée à une audience majoritairement composée d'internaute avec une préférence faible pour la publicité ou 2) que la technologie est imparfaite. Dans les deux cas, l'impact de la technologie sur le bien-être social repose sur l'arbitrage entre le surplus des internautes et les bénéfices des annonceurs.

6.4 Chapitre 3 : Impact d'une régulation de la vie privée “opt-out”sur le marché de la publicité en ligne

La publicité en ligne a généré un chiffre d'affaires de 41,9 milliards de dollars en 2017, soit une hausse de 12,3% par rapport à 2016, selon ce rapport de l'IAB. Cette importante croissance s'explique notamment par le dÃľvelopement important de la publicité ciblée. En effet, le même rapport souligne qu'une grande partie de ces investissements est liée aux technologies de ciblage comportemental. Ces techniques se révèlent très attractives pour les annonceurs : elles collectent des informations sur les internautes ce qui permet une personnalisation des annonces. Par exemple, selon le Washington Post, Facebook permet aux annonceurs de tirer parti d'une multitude de données pour construire leurs campagnes. En utilisant des informations telles que la localisation de l'utilisateur, son âge et son genre mais aussi des caractéristiques plus avancées telles que l'intitulé du mÃľtier ou le type de carte de crédit, Facebook aide les annonceurs à personnaliser leur campagne. La puissance des technologies de ciblage repose également sur des mécanismes automatisés de vente d'espaces publicitaires - mobilisant notamment des enchères. Plus communément appelé "vente programmatique", ces mécanismes seraient responsable de l'affichage d' environ 80% des annoncesen 2017 selon EMarketer. Également, le rapport de l' IAB suggère que 86 % de publicités vendues par vente programmatique en Europe en 2017 incorporaient une stratégie de ciblage comportemental.

Ces pratiques peuvent être perçues comme intrusives par les internautes, de plus en plus sensibles quant à l'utilisation de leurs informations personnelles. Comme détaillé dans Tucker (2012a), les internautes peuvent être gêné face à l'intrusion perçue des publicités personnalisées. Ces préoccupations ont incité les régulateurs à fournir un cadre de suivi des pratiques, prenant ainsi en compte différentes régulations sur l'utilisation des données personnelles des internautes. Le régulateur américain a par exemple favorisé la mise en oeuvre d'une stratégie *opt-out* permettant aux annonceurs et aux sites web d'effectuer un pistage par défaut, mais permettant aux internautes de l'empêcher en "désactivant" cette pratique. S'ils choisissent de ne pas communiquer leurs données, les internautes voient toujours les annonces mais celles-ci ne sont plus personnalisées en fonction de leur comportement antérieur. À ce sujet, la règlementation européenne favorise une politique *opt-in*. En effet, le *RGPD* et plus précisément la directive *ePrivacy* empêchent les entreprises de pister les internautes à moins qu'ils ne l'aient expressément autorisé. Les annonceurs doivent donc récupérer le consentement des utilisateurs avant de pratiquer un ciblage comportemental. Enfin, un *tracking ban* empêche tout simplement les annonceurs et les sites web de pister les utilisateurs. Les informations personnelles étant utilisées de manière intensive dans de nombreux secteurs, il est compliqué pour les régulateurs de choisir la meilleur régulation. Il est alors nécessaire de procéder à une évaluation économique de chaque politique de confidentialité décrite ci-dessus.

Dans ce chapitre, nous analysons l'impact économique d'une régulation de type *opt-out* sur les prix des enchères de publicité en ligne. Johnson (2013a) a déjà étudié cette question. À l'aide d'un ensemble de données exclusif, les résultats de l'auteur prédisent qu'une option de retrait réduirait les revenus des éditeurs et des annonceurs de respectivement 3,9% et 4,6%. D'autres travaux de Johnson et al. (2017) évaluent l'impact du programme AdChoice - programme qui permet aux internet de se retirer du ciblage comportemental - sur le prix des annonces. Toutefois, ces deux articles ne sont pas en mesure d'évaluer l'impact du retrait du ciblage comportemental sur les prix des annonces du point de vue d'un internaute unique. Récupérer ces données nécessiterait en effet de pouvoir suivre et identifier les utilisateurs même après un opt-out, ce qui est impossible pour les entreprises.

Nous développons une nouvelle méthodologie qui nous permet de récupérer les prix avant et après un opt-out des internautes. Nous menons notre expérience pendant 51 jours, pendant janvier et février 2017, et récupérons 6682 prix de publicités en ligne vendues aux enchères. À l'aide d'une méthode diff-in-diff, nous calculons la différence de prix avant et après un opt-out, ce qui nous permet de souligner l'impact économique d'une régulation sur l'utilisation des données personnelles. Le chapitre permet d'éclairer quatre points clés. Toutr d'abord, la méthodologie utilisée se différencie des travaux précédents : nous construisons une méthode computationnelle qui récupère les prix de la publicité avant et après un opt-out du ciblage publicitaire, ce qui n'est pas possible dans le cas de Johnson (2013a) et Johnson et al. (2017). Nous construisons des bots [21] de manière à ce qu'ils soient perçus comme de vrais internautes. Nous les divisons en deux groupes identiques : un groupe de contrôle et un groupe de traitement. Les deux groupes visitent une liste de sites web deux fois de suite. Entre les deux visites, un comportement relatif à la vie privée - à savoir un opt-out du ciblage comportemental - est appliqué au groupe de traitement. Notre base de données est donc composée de prix payés par les annonceurs gagnants des enchères en ligne effectuées lors de la première visite, de la deuxième visite et des robots du groupe de traitement et du groupe de contrôle. Deuxièmement, nous ne trouvons pas de relation claire entre le retrait du ciblage comportemental et les prix des espaces publicitaires vendus aux enchères. De plus, nos résultats indiquent que certains sites web - inclus dans des catégories spécifiques - subissent une faible augmentation du prix publicitaire moyen lorsque le ciblage sur les utilisateurs n'est pas possible. Ce résultat peut paraître contre-intuitif : il ne correspond pas aux résultats précédents qui indiquaient un prix moyen nettement inférieur lorsqu'une annonce est affichée après un opt-out. Cela peut s'expliquer par le fait que nos robots ne sont que peu attractifs pour les annonceurs. Si c'est le cas, en l'absence d'informations sur les internautes (i.e. nos bots) les annonceurs augmenteraient leurs enchères. Troisièmement, nous contrastons ce résultat en montrant que l'effet d'un opt-out est très différent selon la catégorie du site web sur lequel l'annonce est diffusée. En effet, les résultats précédents de la littérrature estimait une baisse uniforme des prix lorsque les annonceurs sont dans l'incapacité d'exploiter les informations des utilisateurs.

[21] Les bots sont des programmes informatiques conçus pour effectuer des tâches sur des ordinateurs. Par exemple, nous pouvons "coder" les robots pour visiter une liste de sites web

Ce résultat souligne l'importance des caractéristiques contextuelle d'une publicité pour les annonceurs. Enfin, nous montrons qu'un opt-out de la publicité ciblée réduit considérablement le nombre d'annonces vendues aux enchères. Ce résultat important souligne la manière dont une régulation sur l'utilisation des données personnelles modifie la concurrence entre les mécanismes de vente. Dans notre cas, le chapitre indique clairement que les sites web utilisent moins intensivement la programmatique pour vendre leurs espaces publicitaires lorsqu'ils n'ont pas accÃÍs aux techniques de ciblage. Cela impacte également nos résultats précédents et peut biaiser notre analyse économétrique. Nous proposons une solution à ce problème dans le chapitre.

6.5 Conclusion

La thèse analyse trois nouvelles considérations importante pour le marché de la publicité en ligne. Les chapitres y interrogent respectivement le rôle de la transparence, des algorithmes de profilage et des informations personnelles. La conclusion se concentre sur les résultats de la thèse et présente des extensions potentielles pour chaque chapitre.

Le premier chapitre analyse comment une technologie permettant aux annonceurs de vérifier la visibilité des publicités peut modifier les équilibres du marché. Le chapitre souligne que sans technologie, l'éditeur ne peut pas s'engager sur la visibilité de ses espaces publicitaire et choisit donc de diminuer son volume de publicité par rapport à son niveau optimal. Il montre également que l'introduction de telles technologies peut accroître le bien-être total uniquement si la nuisance publicitaire n'est pas trop élevée. Enfin, le chapitre examine le lien entre cette technologie et l'évitement de la publicité. Ce chapitre est l'une des premières productions académiques à souligner l'importance de la transparence des achats sur le marché de la publicité en ligne.

Cependant, d'autres problÃÍmes de marché - également en lien avec la transparence - inquiètent aujourd'hui les annonceurs. Un premier exemple concerne la fraude à la publicité en ligne. Selon une étude de Association of National Advertisers, 2016, les “publicités frauduleuses”, où des robots visitent des sites web et y sont considérés comme de véritables internautes, représentent environ 40 % du trafic en ligne et aurait coûté 7,2 milliards de dollars aux marques en 2016. En d'autres termes, les fournisseurs de contenu ont considéré 40 % des annonces affichées aux internautes (parfois même cliquées), alors qu'elles étaient en réalité affichées à un robot. La fraude publicitaire mobilisant des bots est l'une des plus courantes sur Internet. D'autres escroqueries telles que l' "ad-stacking", qui empile plusieurs publicités dans un seul espace ou le "pixel stuffing", qui affiche la publicité dans un seul pixel de la page, se sont de plus en plus répandus, dégradant encore plus le lien de confiance entre les annonceurs et les sites web. Sans technologies de detection, de nombreux acteurs sont incités à mettre en place des systèmes de fraude à la publicité en ligne. Comme l'a souligné Bloomberg dans un article datant de 2006, Google et Yahoo! auraient un intérêt théorique à tolérer la fraude. Par conséquent, le marché a fortement investi dans les technologies de détection au cours des dernières années. Selon Mediatel, 2017, Google financerait la détection de trafic non valide afin de renforcer la confiance dans la chaîne logistique de la publicité. Ces technologies agissent comme des filtres, distinguant le faux trafic de celui légitime. D'autres initiatives visant à débarrasser le marché de la fraude ont vu le jour. L'année dernière, l'Interactive Advertising Bureau (IAB) a lancé "ads.txt", un moyen de vérifier les caractéristiques des sites web avant d'y diffuser une publicité, garantissant ainsi leur intÃl'gritÃl'.

Malgré son importance économique, la fraude en matière de publicité en ligne a été relativement peu explorée par les économistes. Asdemir et al. (2008) analyse comment un moteur de recherche financé par la publicité, utilisant une technologie imparfaite pour identifier un trafic factice, peut également réduire le nombre de clics valables. Or, la fraude en matière de publicité en ligne est aujourd'hui diverse. De plus, l'introduction d'une technologie de détection de fraude permet aux annonceurs d'ajuster leur investissement en temps réel lorsqu'ils achètent de la publicité en ligne. Dans l'ensemble, étant donné que les technologies de fraude et de lutte contre celle-ci modifient la demande d'espaces publicitaires, une analyse économique s'avererait utile pour comprendre comment les internautes, les annonceurs et les sites web sont impactés.

Dans un deuxième chapitre, la thèse analyse comment les éditeurs sont désormais en mesure de comprendre la sensibilité publicitaire des internautes et adapte ainsi le nombre de publicités à leurs préférences. Le chapitre détaille les conséquences de l'introduction de cette technologie de profilage sur le niveau publicitaire optimal du marché, les bénéfices des entreprises et le bien-être total.

Le chapitre s'appuie sur le fait que les technologies de profilage sont alimentées par les traces laissées par les internautes (Estrada-Jiménez et al., 2017). Ces technologies dépendent donc de leurs volontés de divulguer des informations. Or, une proportion non négligeable choisit de ne divulguer d'informations personnelles sur Internet. Pour se protéger de ce genre de pratique, les internautes utilisent des logiciels tels que Ghostery, Privacy Badger ou uMatrix. Une littérature sur la protection de la vie privée en lien avec la publicité a déjà commençé à analyser ce type de comportement (voir Turow et al. (2009) et **?**). Si les internautes masquent leur comportement de navigation, les technologies risquent d'être moins efficaces et seront donc plus susceptibles de faire des erreurs. Sur ce sujet, une analyse économique serait pertinente et permettrait de comprendre comment ces nouveaux comportements impacte le marché de la publicité en ligne.

Plus généralement, lorsque les internautes cachent leurs traces, les marques adaptent leurs modèles d'affaires, ce qui peut modifier le mode de financement du contenu en ligne. Différentes régulation tentent de trouver un équilibre entre protection du droit à la vie privée et financement de contenu en ligne. Ces dernières nécessitent une analyse économique minutieuse, car elles pourraient modifier radicalement la volonté des annonceurs Ãă payer pour de la publicité en ligne, et par extension les bénéfices des sites web.

Dans un dernier chapitre, la thèse évalue comment une rÃl'gulation sur la protection des données personnelles des internautes peut influer sur le prix des annonces en ligne vendues aux enchères. Le chapitre présente une nouvelle méthode qui permet d'estimer la variation du prix publicitaire lorsqu'on ne peut pas cibler un internaute. En premier lieu, le chapitre constate que l'absence de ciblage peut faire augmenter le prix des annonces vendues aux enchères. Deuxièmement, le chapitre montre également que l'absence de ciblage n'affecte pas les sites web de manière uniforme, mais que cet effet dépend du type de site web. Ces résultats offrent une piste de compréhension de l'impact qu'une régulation sur l'utilisation des données personnelles des internautes - notamment *opt-out* - peut avoir sur les prix et la quantité de publicité sur un site web.

Cependant, des réglementations, telles que le Réglement Général sur la Protection des Données et la future directive ePrivacy ont pour but de “redonner aux citoyens le contrôle de leurs données personnelles et de simplifier l'environnement réglementaire des entreprises”, et favorise une régulation de type *opt-in*. Afin de tirer parti des informations personnelles des utilisateurs, les sites web devront d'abord récupérer leur consentement en précisant le caractère des données collectées, le but exact de la collecte ainsi que le nom du responsable de traitement. Des travaux tels que Johnson (2013a) essaient d'estimer l'impact d'une politique de confidentialité privilégiant un opt-in sur le marché de la publicité en ligne. Une autre analyse économique pourrait venir renforcer les travaux existant, et permettrait de mieux évaluer l'impact de ces régulations sur les marchés.

La thèse présente trois chapitres qui exploitent la littérature classique sur la publicité en ligne. Cependant, de nombreuses nouvelles transformations peuvent intéresser les économistes. Par exemple, Hsu (2016) souligne le fait que le secteur de la publicité en ligne est largement dominé par deux acteurs : Facebook et Google. En effet, selon Business Insider, 2017, les deux acteurs concentrent plus de 90% de la croissance des dépenses du marché de la publicité en ligne aux États-Unis en 2016. Ensuite, les deux acteurs semblent être présents sur l'ensemble de l'écosystème de la publicité en lignec- de la vente de données à l'organisation de ventes aux enchères programmatiques. Ces deux points tendent à souligner une concentration du marché, pouvant naturelement intéresser les chercheurs en sciences économiques.

D'autres auteurs soulignent également le doute planant sur l'efficacité de la publicité en ligne. À titre d'illustration, la célèbre société multiproduits *Proctor & Gamble* a fait savoir qu'une réduction de 100 millions de dollars des dépenses publicitaires numériques en juin 2017 n'avait eu que peu, voire aucune incidence sur les bénéfices (Wall Street Journal, juillet 2017). Ainsi, les marques préfèrent se tourner vers d'autres média pour leurs campagnes publicitaires. Cette concurrence entre médias fait émerger de nouvelles stratégies. La compréhension de ces nouvelles stratégies et de la manière dont elles façonnent les marchés de produits bénéficierait d'une analyse économique.

Bibliography

Anderson, Simon and Jean Gabszewicz, "Chapter 18 The Media and Advertising: A Tale of Two-Sided Markets," in Victor A. Ginsburg and David Throsby, eds., *Handbook of the Economics of Art and Culture*, Vol. 1, Elsevier, 2006, pp. 567 – 614.

Anderson, Simon P. and Joshua S. Gans, "Platform Siphoning: Ad-Avoidance and Media Content," *American Economic Journal: Microeconomics*, November 2011, *3* (4), 1–34.

_ **and Stephen Coate**, "Market Provision of Broadcasting: A Welfare Analysis," *The Review of Economic Studies*, 2005, *72* (4), 947–972.

Asdemir, Kursad, ÃŰzden Yurtseven, and Moin A. Yahya, "An Economic Model of Click Fraud in Publisher Networks," *International Journal of Electronic Commerce*, 2008, *13* (2), 61–90.

Bagwell, Kyle, in Mark Armstrong and Robert Porter, eds., *The Economic Analysis of Advertising*, 1 ed., Vol. 3, Elsevier, 2007, chapter 28, pp. 1701–1844.

Buchbinder, Niv, Moran Feldman, Arpita Ghosh, and Joseph (Seffi) Naor, "Frequency Capping in Online Advertising," in Frank Dehne, John Iacono, and Jörg-Rüdiger Sack, eds., *Algorithms and Data Structures: 12th International Symposium, WADS 2011, New York, NY, USA, August 15-17, 2011. Proceedings*, Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 147–158.

de Cornière, Alexandre, "Search Advertising," *American Economic Journal: Microeconomics*, 2016, *8* (3), 156–188.

Estrada-Jiménez, José, Javier Parra-Arnau, Ana Rodríguez-Hoyos, and Jordi Forné, "Online advertising: Analysis of privacy threats and protection approaches," *Computer Communications*, 2017, *100*, 32 – 51.

Goldfarb, Avi, "What is different about online advertising?," *Review of Industrial Organization*, 2014, *44* (2), 115–129.

_ **and Catherine Tucker**, "Online Display Advertising: Targeting and Obtrusiveness," *Marketing Science*, 2011, *30* (3), 389–404.

Hsu, David, "The End of Digital Advertising as We Know It," Article November 2016.

Johnson, Garrett A., "The Impact of Privacy Policy on the Auction Market for Online Display Advertising," SSRN Scholarly Paper ID 2333193, Social Science Research Network, Rochester, NY May 2013.

Johnson, Garrett A, Scott Shriver, and Shaoyin Du, “Consumer Privacy Choice in Online Advertising: Who Opts Out and at What Cost to Industry?,” 2017.

Johnson, Justin P, “Targeted advertising and advertising avoidance,” *The RAND Journal of Economics*, 2013, *44* (1), 128–144.

Lambrecht, Anja and Catherine Tucker, “When Does Retargeting Work? Information Specificity in Online Advertising,” *Journal of Marketing Research*, 2013, *50* (5), 561–576.

Manchanda, Puneet, Jean-Pierre Dubé, Khim Yong Goh, and Pradeep K. Chintagunta, “The Effect of Banner Advertising on Internet Purchasing,” *Journal of Marketing Research*, 2006, *43* (1), 98–108.

Nelson, Phillip, “Advertising as Information,” *Journal of Political Economy*, 1974, *82* (4), 729–754.

Peitz, Martin and Markus Reisinger, “Chapter 10 - The Economics of Internet Media,” in Simon P. Anderson, Joel Waldfogel, and David StrÃűmberg, eds., *Handbook of Media Economics*, Vol. 1 of *Handbook of Media Economics*, North-Holland, 2015, pp. 445 – 530.

Renault, Régis, “Chapter 4 - Advertising in Markets,” in Simon P. Anderson, Joel Waldfogel, and David StrÃűmberg, eds., *Handbook of Media Economics*, Vol. 1 of *Handbook of Media Economics*, North-Holland, 2015, pp. 121 – 204.

Rochet, Jean-Charles and Jean Tirole, “Cooperation among Competitors: Some Economics of Payment Card Associations,” *The RAND Journal of Economics*, 2002, *33* (4), 549–570.

Stigler, George J, “The economics of information,” *Journal of political economy*, 1961, *69* (3), 213–225.

Tag, J., “Paying to Remove Advertisements,” *Information Economics and Policy*, 2009, *21* (4), 245–252.

Tucker, Catherine E., “The economics of advertising and privacy,” *International Journal of Industrial Organization*, 2012, *30* (3), 326 – 329. Selected Papers, European Association for Research in Industrial Economics 38th Annual Conference, Stockholm, Sweden, September 1-3, 2011.

Turow, Joseph, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy, “Americans Reject Tailored Advertising and Three Activities that Enable It,” Working Paper 1478214, SSRN September 2009.

Three Essays in Economics of Online Advertising

Martin QUINN