

Ad-blocking Games: Monetizing Online Content Under the Threat of Ad Avoidance

Nevena Vratonjic, Mohammad Hossein Manshaei, Jens Grossklags and Jean-Pierre Hubaux

Abstract Much of the Internet economy relies on online advertising for monetizing digital content: Users are expected to accept the presence of online advertisements in exchange for content being free. However, online advertisements have become a serious problem for many Internet users: while some are merely annoyed by the incessant display of distracting ads cluttering Web pages; others are highly concerned about the privacy implications – as ad providers typically track users’ behavior for ad targeting purposes. Similarly, security problems related to technologies and practices employed for online advertisement have frustrated many users. Consequently, a number of software solutions have emerged that block online ads from being downloaded and displayed on users’ screens as they browse the Web.

We focus on these advertisement avoidance technologies for online content and their economic ramifications for the monetization of websites. More specifically, our work addresses the interplay between users’ attempts to avoid commercial messages and content providers’ design of countermeasures. Our investigation is substantiated by the development of a game-theoretic model that serves as a framework usable by content providers to ponder their options to mitigate the consequences of ad avoidance techniques. We complement our analytic approach with simulation results, addressing different assumptions about user heterogeneity.

Our findings show that publishers who treat each user individually, and strategically deploy fee-financed or ad-financed monetization strategy, obtain higher revenues, compared to deploying one monetization strategy across all users. In addition, our analysis shows that understanding the distribution of users’ aversion to ads and valuation of the content is essential for publishers to make a well-informed decision.

Nevena Vratonjic

School of Computer and Communication Sciences, EPFL, Switzerland,
e-mail: nevena.vratonjic@epfl.ch

Mohammad Hossein Manshaei

Department of Electrical and Computer Engineering, Isfahan University of Technology, Iran,
e-mail: manshaei@gmail.com

Jens Grossklags

College of Information Sciences and Technology, The Pennsylvania State University, USA,
e-mail: jensg@ist.psu.edu

Jean-Pierre Hubaux

School of Computer and Communication Sciences, EPFL, Switzerland,
e-mail: jean-pierre.hubaux@epfl.ch

1 Introduction

It is difficult to produce a television documentary that is both incisive and probing when every twelve minutes one is interrupted by twelve dancing rabbits singing about toilet paper. (Rod Serling, 1997)

Consumers and content providers have a love-hate relationship with advertisements. In the area of online news sites, 81% of a surveyed consumer sample report the acceptance of the presence of online advertising in exchange for content being free. At the same time, 77% state that they would hardly ever click on these ads [36]. More significantly, across all media channels, 69% say they are “interested in products and services that would help them *skip* or *block* marketing [40].”

Each media genre is affected with its own specific advertisement circumvention challenges. During TV commercial breaks, watchers can leave the room to do small chores. Ads in video recordings can be manually skipped with fast-forwarding or are automatically marginalized with advanced functions of digital video recorders (e.g., TiVo) and VCRs [27]. This trend has accelerated with the availability of Home Theater PC systems such as Windows Media Center, SageTV Media Center and MythTV where available third-party add-ons allow consumers to conveniently skip ads (e.g., Comskip and ShowAnalyzer). In telemarketing, consumers are able to screen calls with CallerID or utilize software tools that act on their behalf (e.g., Telemarketing Blocker). Further, regulatory intervention can have a significant impact, for example, with the US Do-Not-Call list that upon registration allows consumers to opt out from unsolicited telephone marketing calls [47].

We focus on advertisement avoidance technologies for Web content and their economic ramifications. In the past few years, a number of effective software solutions have emerged of which the most prominent is perhaps the Adblock Plus third-party extension for the Firefox browser family [17, 39]. According to up-to-date statistics provided by Mozilla, Adblock Plus has been downloaded over 128 Million times since July 2006, and has an active daily user base of about 14 Million consumers. Further, it is also among the most popular add-ons for the Google Chrome browser with more than 100,000 weekly installs. Observers from the advertising business have predicted that the “importance of Adblock is its potential for extreme menace to the online-advertising business model [20]”. However, many other technology options exist to block ads.

The emergence of behavioral advertisement and the associated increase in advertisers’ incentives for user tracking, has led to what some observers call a “data collection arms race” (see, for example [10]). Most recently, Google’s proposed changes to its privacy practices that would allow for more pervasive user data aggregation have refreshed privacy concerns in consumers’ minds (see, for example [13]). And consumers object to such practices [31, 44]. However, in absence of truly effective *and* wide-spread technologies to opt-in/opt-out from tracking and the later usage of such information for advertisements, consumers only have the option to decide on their own personal mix of avoidance technologies. For example, while consensus for a powerful and broadly applicable Do Not Track mechanism is still

absent, some users may seek to disable scripting languages, Flash or cache cookies.¹ Others may use advanced privacy-enhancing technologies such as Tor just for the purpose of evading such commercially-motivated tracking. Finally, to be effective, avoidance of tracking does frequently necessitate also the blocking of the display of ads since ad campaigns almost always involve some form of campaign management tools. While this is trivially necessary to allow for ad-related payment flows, consumers cannot easily distinguish between different degrees of tracking severity.

So far, the impact of the circumvention of online tracking and advertisements has been moderated by the overall growth of the market for Internet commercials. The Interactive Advertising Bureau estimates that through the first half of 2011, online advertising in the United States totaled \$12.1 billion and has grown by 11 percent compared to the same time frame in the previous year [35]. Nevertheless, many content sites suffer from the burden of ad-blockers, in particular, if they cater to a technology-savvy audience (see, for example, [23]). The search for an adequate response to this threat has so far proven inconclusive. In particular, monetization approaches do not only have to be economically sensible, but need to be accompanied by technically sound implementations. So far, ad-block deterrence solutions have been notably absent from the marketplace, even though the cost of development and deployment of simple approaches would be very manageable. In fact, as the majority of ad-blocking solutions are based on filtering out elements whose URLs contain keywords like ad or click, omitting these keywords would make existing ad-blocking tools ineffective. In addition, existing tools cannot automatically detect URLs likely to be advertisements. Therefore, if publishers start using different keywords, ad-blocking systems would not work [39].

The stakes described in this paper are very high and are relevant beyond the discussions about the effectiveness of marketing or commercial mechanisms. In fact, the popularity of Adblock-style add-ons represents only the tip of the iceberg, as many related challenges are consuming the attention of content producers. For example, applications such as Flipboard allow users to conveniently grab pictures and articles from many different content resources to display them in a variety of user-defined formats, and ads may be left behind (or are replaced).

Our work studies in detail and in a quantitative manner the implications of a (likely to happen) growing usage of ad blocking techniques and addresses the economic justification for effective countermeasures concerning ad avoidance. To achieve that goal, we develop a game-theoretic model that takes into account the most relevant parameters, identifies different canonical options (strategies) that the content providers and the users can choose from and forecasts the most likely outcome of such situations. The models we provide rely on Subgame Perfect Nash Equilibria (SPNE) and on Perfect Bayesian Nash Equilibria (PBNE). We complement our analytic approach with simulation results by addressing different assumptions about user heterogeneity. We make "common sense" assumptions in terms of cost and show that in general, content providers are better off when they make use

¹ It is unlikely that a meaningful compromise on Do Not Track will be reached quickly. See, for example, the counter-arguments on such technology brought forward by leading content providers [6].

of a "mixed approach", namely when they simultaneously rely on fee-funded and ad-funded monetization strategies.

Closely related to our work is an economic model by Tåg [43]. Content providers decide whether to offer to users a subscription option that eliminates advertisements. By increasing the amount of advertisements to non-subscribers, the content provider can further increase the differentiation between these options. In another model, Shah accounts for ad avoidance technologies [37]. Users may invest in ad avoidance options but will still see a certain fraction of the commercials. A content provider can make use of this fact by optimally differentiating the amount of advertisements catered to the two groups (i.e., users with and without ad avoidance products). In a two-sided market model, Anderson and Gans similarly show that content providers may increase the number of advertisements to those users who do not invest in avoidance technologies [16].

In our model, we explicitly consider the limited information aspects related to advertisement avoidance and its detection. As a result, content providers must invest in detection technologies to be able to distinguish between consumers that utilize ad-blocking and those who do not engage in such activities.

Further academic works on advertisement circumvention have been undertaken in the context of "old media" from an empirical perspective (e.g., about video recording technologies) or address the problem from a legal or ethical perspective [27, 41, 45, 48].

Additional recent work has been focused on improvements of the mechanisms for ad allocations and techniques to lower the impact of manipulation by malicious actors. See, for example, research papers on ad auctions (e.g., [22, 46]) and click fraud [30, 32].

The paper is structured as follows. In Section 2, we introduce the reader to background information relevant to the problem area of ad avoidance. After briefly laying out the roadmap for our analysis in Section 3, we delve into the details of our game-theoretic models in Section 4. Our simulation results follow in Section 5, before we end with concluding remarks.

2 Background

In this section, we discuss the drivers of consumer resistance to advertisements and their propensity for ad blocking. We also review existing technologies for ad avoidance and approaches by website owners to detect blocking software.

2.1 Why do consumers block ads?

Previous research has studied a variety of ad avoidance behaviors such as eliminating, ignoring or quickly flipping past commercial messages [42]. Graphical and

auditory stimuli are frequently considered annoying or unconvincing, irrespective of the actual information content [42]. Online ads are more likely to be avoided if consumers hold expectations of a negative experience, are generally skeptical towards the advertising messages or contest their relevance [29]. Further, if a user perceives an interruption in her primary interaction objective or considers ads to clutter her workspace, marketing messages are more likely to be blocked or ignored [19]

Further, sophisticated online advertising approaches such as personalized, behavioral or targeted delivery mechanisms rely on the collection and use of data about users' Web interactions. Different studies have documented users' misgivings and privacy concerns about these practices. For example, in an interview study of 1000 adult consumers, 66% objected to tailored ads [44]. Due to the pervasiveness of these concerns, (self-)regulatory and technical proposals are under consideration, e.g., that would allow users to opt-out from such data collection practices by signing up for a Do-Not-Track list [21]. At the same time, users may attempt to block advertisements altogether when suspecting that they are triggered by the tracking of their online trails.

A survey of 1543 Adblock Plus users further evidenced that privacy and security concerns are major factors to select this application [34]. Avoiding distractions and improving website load time performance, however, are the dominating reasons. Interestingly, the lowest score of importance was given to ideological reasons. See Table 1 for the full results [34].

Table 1 Survey results: Why do consumers use Adblock Plus?

Reasons	No Opinion	Not Important	Somewhat Important	Important
Distracting animations and sound	4.3%	5.6%	15.6%	74.5%
Offensive/inappropriate ad content	8.0%	20.1%	23.3%	48.6%
Reduce page load time and bandwidth use	5.7%	10.1%	22.6%	61.6%
Missing separation between ads and content	13.2%	11.5%	27.5%	47.8%
Privacy concerns	8.3%	9.9%	27.5%	54.3%
Security concerns	8.0%	9.7%	26.1%	56.3%
Ideological reasons	20.2%	32.0%	24.2%	23.7%

2.2 What technologies are involved?

Ad-blocking tools prevent online ads from being downloaded and displayed on users' screens as they browse the Web. They can also be considered privacy tools as some forms of online tracking (e.g., via cookies) can be evaded. Typically, ad-blocking tools are available as free downloadable plug-ins and exist for several Web browsers. For example, AdblockPlus is open-source and maintained by an international community of voluntary helpers. Internet Explorer 9 includes a directly

embedded functionality primarily used for *tracking protection*, but also allows to block some unwanted content.

Ad-blocking tools rely on two mechanisms to block ads: (i) Prevent loading elements whose URLs match *filter rules* used to declare elements as ads, and (ii) Hide page elements that match a Cascading Style Sheets (CSS) selector. Users may subscribe to different community-generated filter lists or manually specify filtering rules themselves. They can also decide to allow loading of some elements of a page or to turn-off ad-blocking on specific pages or websites. However, this feature is not widely used among Adblock users [34].

Ad-blocking causes revenue loss for advertisers and ad networks but it has the most significant impact on websites whose business model is based on online advertising. The majority of websites today rely on ad revenue, whereas only a few websites have successfully implemented subscription and membership based systems for revenue. Therefore, it is understandable that site operators might want to discourage or thwart ad-blocking. In particular, a website may be able to detect the use of ad-blocking tools with a JavaScript that executes after the page has been loaded and verifies that the ads are indeed displayed. Then, the website may take one of the following countermeasures: (i) inform users about adverse effect of ad-blocking on the website and ask them to turn it off; (ii) prevent users from accessing the content unless they disable ad-blocking; (iii) embed ads in a way that ad-blocking filters cannot easily differentiate ads from content; (iv) tie the functionality of websites to the download of ad elements; and (v) offer users to pay subscription fees for ad-free content.

Both the blocking of ads and measures against blocking currently come at a very low cost. The former requires the user to install a browser plug-in and subscribe to filter lists. As for detecting ad blockers, the required JavaScript code is easily available online.

3 Analysis Overview and Assumptions

We propose a game-theoretic model of the informational consequences of consumer’s ad circumvention and website owner’s detection of these practices. In our analysis, we model the strategic interactions between a generic website W and a user U and we iteratively consider the following three cases: (i) without the presence of ad blocking and ad circumvention detection technologies; (ii) with ad blocking but no detection, and (iii) where both technologies are available to consumers and website owners, respectively.

A key assumption we make is that the website owner attempts to analyze users *individually*. A number of technologies exist to implement various forms of conditional content and ad delivery (see, for example, [26]) ranging from tailoring a website’s appearance to the type of browser and operation system in use by the consumer. Note that the individualized analysis does not necessarily translate into unique monetization strategies for each user.

Website owners can utilize two canonical types of monetization strategies in response to a particular user: either employ advertisement financed content delivery or propose a micropayment for access to content (as a representative subcase of a wider range of payment-based strategies, such as subscriptions). The consideration of micropayments for newspaper content is extremely timely. Not only has the debate about micropayment schemes for news and other digital content been fought very passionately over the last few years [28, 38]; But from an actual deployment point of view, easy-to-manage systems are now available, for example, One Pass from Google [7] or PayPal for Digital Goods [14]. And consumers seem more willing than ever to accept small charges in response for immediate content or entertainment needs [38].

We further assume that the website owner is aware of the user’s valuation of content, for example, because of the cooperation with advertisement networks, inference about the resources the user is trying to access or previous interactions. In practice, websites work on obtaining such information and use it to, for example, compute appropriate prices for their services or content (e.g., New York Times’ subscription price is based on the estimates of readers’ valuations of the content, which is set such that the current paywall system should be accepted by a certain fraction of their readership [8]). Our analysis can also be easily extended to introduce uncertainty about user’s content preferences from the content providers’ perspective.

Not all aspects about user behavior are immediately observable without sophisticated detection technologies. In particular, the website owner cannot easily deduce whether the consumer is taking advantage of ad blocking technologies. This is especially the problem in the impression-based ad revenue model, in which the website obtains ad revenue for each ad displayed to its visitors. For example, if the feedback cycle between the ad network and the content provider is not real-time then payoff consequences of ad avoidance are only realized at a later time. In the click-based ad revenue model, a website gets paid for users’ clicks that get reported to the ad network, thus perhaps enabling more direct and immediate control. The absence of signals may indicate to ad networks (and websites) a change in the user’s behavior (e.g., use of AB software). The website can mitigate this information disadvantage by investing in technologies to detect ad avoidance. In this work, we focus on impression-based ad revenue model and we note that the similar analysis can be provided for the click-based ad revenue model.

Based on these assumptions, we model each website visit as a sequential game between the two players, a website W and a user U , to highlight the informational and strategic aspects of the interactions. We represent the different cases as game trees (see Figures 1 and 2) with the notation provided in Table 5, on page 24. In each game, the players can choose from the corresponding strategy sets, S_W and S_U , and the payoffs achieved at the end of the game are represented in a format (P_W, P_U) , where P_W and P_U are total payoffs of W and U , respectively.

4 Game-theoretic Models

In this section, we introduce game-theoretic models that capture strategic interactions of a website and a user. For each model, we present analysis methodology and the obtained results.

4.1 Model 1: No Blocking and No Detection

We want to introduce the reader to our approach by first proposing a relatively trivial basic model of the interaction between website owners and consumers in which no ad-blocking or detection technologies are used by users and the website, respectively. Afterwards, we will slowly increase the complexity of the model to account for ad avoidance and countermeasures.

4.1.1 Model Setup: An Extensive Form Game with Complete Information

The content provider selects between fee-financed (e.g., micropayments) and ad-financed monetizing scheme for his content (denoted by FF and AF , respectively). If presented with a website that solicits a fee to access its content, users can elect to transmit a payment, P , or to deny payment and forfeit access, N . The website owner will either earn positive revenues from the ad impression, r_i , or from the micropayment, s . The consumer receives a benefit, b , from accessing the content and pays either the fee, s , or has a cost c due to accepting ads. The subscription fee s is determined by the content provider and it is the same for all the users, because it has been shown that price discrimination is not economically optimal for providers [15] and because of users' protest (e.g., case of Amazon [3]). Determining the optimal price is not the goal of this work, but is certainly noteworthy to explore. The cost c captures all the negative aspects of receiving ads from the users' point of view (summarized in Table 1). Figure 1 summarizes the characteristics of the basic model.

4.1.2 Analysis Methodology: Subgame Perfect Nash Equilibrium

The model in this section belongs to the class of *perfect and complete information extensive form* games. In these games, each player always knows the previous moves of all players when he has to make his move. In [24], it is proven that every finite extensive-form game of perfect information has a pure-strategy Nash equilibrium. Let us first define *Subgame Perfect Nash Equilibrium* (SPNE) that is a refinement of a Nash equilibrium used in dynamic games. In game theory, a strategy profile is a subgame perfect Nash equilibrium if it represents a Nash equilibrium of every sub-




Fig. 1 Model 1 - Game tree for the basic model with no blocking and no detection.

game of the original game. A common method for determining SPNE is *backward induction* and we apply it in our analysis.

Backward induction can be applied to any finite game of perfect information. This technique eliminates incredible equilibria and assumes that: (i) the players can reliably forecast the behavior of other players and (ii) the players believe the other players can do the same. In the game defined by Figure 1, the user knows that he is the player that has the last move. Hence, for each possible move of the website the user selects his best response. For example, if the website plays *FF*, the user concludes that with move *P* he obtains the best payoff if and only if $b > s$.

Now we consider how the website chooses his best strategy, by using backward induction. Let us assume that $b > s$. In this case, the website knows that if he selects *FF*, then the best response for the user is to play *P*, which results in the payoff of s for the website. However, if the website uses *AF*, its payoff would be equal to r_i . Hence, the best response of the website is *FF*, if $s > r_i$. In summary, if $b > s$ and $s > r_i$ strategy profile (FF, P) is the SPNE of the game in Figure 1. Table 4 summarizes all possible SPNE of the defined game.

4.1.3 Results

Following this methodology, Table 4 summarizes all possible SPNE of the defined game, considering different values of game parameters.

Table 2 SPNE of Game Model 1.

$b > s$	$s > r_i$ (FF, P)
	$s < r_i$ (AF, P)
$b < s$	(AF, N)

It follows that a website owner would only implement fee-financed revenue scheme when users' value of the provided content is sufficiently high, $b > s$, and the expected ad-revenue does not exceed fee payments, $s > r_i$.

The first condition is relatively difficult to assess for a large number of diverse users if the revenue policy cannot be set adaptively for each consumer. In contrast, the second condition allows for a more straightforward calculation - at least for an impression-based advertisement scheme. We address the impact of the heterogeneity of the user base in our simulations (Section 5).

4.2 Model 2: Blocking, Detection vs. No Detection

In the following, we extend the analysis to include consumers having the opportunity to utilize ad-blocking software and website owners to potentially respond by investing in detection technologies. The expanded game is represented in Figure 2.

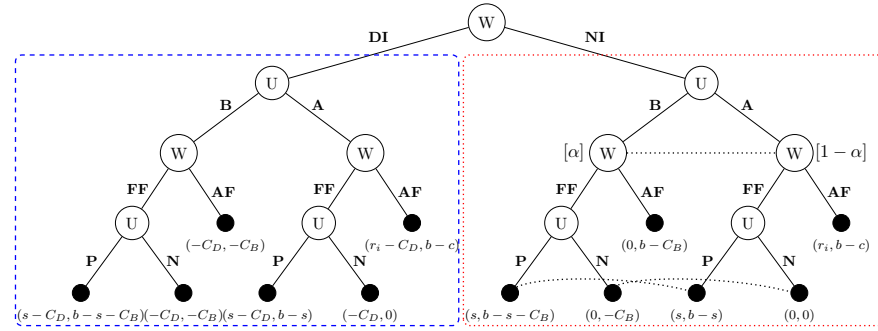


Fig. 2 Model 2 - Game tree for ad-blocking with and without detection technologies.

4.2.1 Model Setup: An Extensive Form Game with Imperfect Information

Consumers now have the option to use an ad-blocker, B , at cost C_B , or to abstain from blocking ads, A , which does not incur any direct cost. We assume that website owners are aware of the possibility of ad-blocking, but without an investment into detection technologies, NI , are not able to differentiate between individuals with and without AB software and thus hold only *imperfect* information about the user's action and its payoff consequences. In contrast, when the website is equipped with detection technologies, DI , at cost C_D , then the information barrier is resolved.

The informational consequences are easily discernible in Figure 2 by observing the dotted lines between information sets that indicate the website owner's uncertainty about the reached state in the game and the eventual outcomes. Website own-

ers have to formulate a probabilistic assessment α of the reached state of the game, following the user's decision to block ads or to abstain.

We further break down the game into two subgames concerning the website owner's decision to invest or not into detection of AB software, as highlighted by the left and right boxes in Figure 2, respectively. The analysis of the lefthand side game in Figure 2 (i.e., when website plays *DI*) is similar to the calculation of SPNE, presented in Section 4.1.2. Using the same methodology we obtain SPNE of this game and present the obtained results later in Table 4.

4.2.2 Analysis Methodology: Perfect Bayesian Nash Equilibrium

The subgame in the righthand side of the game in Figure 2 belongs to the class of *complete imperfect sequential* games, because one player does not have information about the opponent's action played in the previous stage of the game. In other words, the website owner does not know whether the user has already installed AB or not, when he wants to choose the charging mechanism for the website's content (i.e., use ad-financed or fee-financed mechanism).

In this section, we discuss the game-theoretic concept of the *Perfect Bayesian Nash Equilibrium* (PBNE) that helps us get an insight into the strategic behavior of players in such games. PBNE was developed in order to refine the Bayesian Nash equilibrium concept and remove implausible equilibria in sequential games [25]. More specifically, the concept of PBNE is defined by four Bayes requirements that eliminate unwanted subgame-perfect equilibria [33]. We discuss these requirements considering the defined game represented in the righthand side box in Figure 2.

Requirement 1: The player with the move must have a belief about which node in the information set has been reached by the play of the game. For example, in Figure 2 the website believes that the user installed AB with a probability of α .

Requirement 2: At the PBNE strategy profile, players must be sequentially rational given the players' beliefs. A strategy profile is sequentially rational if and only if the expected payoff of the player who has the move at that information set is maximal given the strategies played by all the other players. For example, in Figure 2 the website should calculate its average payoff for playing *AF* and *FF*, given its belief α and choose the strategy that maximizes its average payoff. Given website belief, the expected payoff from playing *FF* is $\alpha \times s + (1 - \alpha) \times s = s$. The expected payoff from playing *AF* is $\alpha \times 0 + (1 - \alpha) \times r_i = (1 - \alpha)r_i$. Hence if $\alpha > \frac{r_i - s}{r_i}$, the website plays *FF* to be sequentially rational.

Requirement 3: The player must update his belief at the PBNE to remove implausible equilibria of BNE on the equilibrium path. These beliefs are determined by Bayes' rule and the players' equilibrium strategies. In other words, players should first calculate the equilibrium paths of the complete perfect information game. If the calculated strategy that satisfies sequential rationality is on the equilibrium path, there is no uncertainty for the player at the PBNE (i.e., α equals 0 or 1).

Requirement 4: Finally, the belief should be updated considering the sequential rationality and players' equilibrium strategies where it is possible.

In the righthand subgame presented in Figure 2, if $b > s$, $s < r_i$, and $C_B > c$ there exists an equilibrium path of $(A|P, AF)$. Although, the user cannot play P when the website deploys AF strategy, we use $A|P$ notation to represent the full strategy profile of the user at the equilibrium path. This means that if $\alpha < \frac{r_i - s}{r_i}$, the PBNE is $(A|P, AF; \alpha = 0)$ (i.e., Requirement 3). Requiring that each player have a belief and act optimally given this belief suffices to eliminate the implausible equilibria for the belief of $0 < \alpha < \frac{r_i - s}{r_i}$. But, if $\alpha > \frac{r_i - s}{r_i}$, the PBNE is $(A|P, FF; \alpha)$, because we cannot eliminate any implausible equilibria for this strategy profile (i.e., Requirement 4). Similar calculations can be made for other cases.

4.2.3 Results

Applying this methodology, we can derive results presented in tabular fashion for the righthand side (Table 3) and the lefthand side (Table 4) of Figure 2.

If website owners do not invest in detection, we observe that ad blocking happens in two instances (see Table 3). First, when consumers do not value the content highly enough to pay a fee ($b < s$), and ad-blocking is cheap relative to ads' nuisance ($C_B < c$). Second, if website owners believe it to be unlikely that consumers block ads ($\alpha < \frac{r_i - s}{r_i}$) and blocking is cheap, then ad avoidance can persist even when users value the content sufficiently ($b > s$). In both cases, the user will exploit his information advantage to avoid ad clutter while the website owner will gain nothing through the interaction (because he mistakenly relies on ad-funding, AF).

Table 3 PBNE of submodel without detection.

	$C_B < c$	$C_B > c$
$s > r_i$	$(A P, FF; \alpha = 0)$	$(A P, FF; \alpha = 0)$
$b > s$	$(A P, FF; \alpha) \alpha > \frac{r_i - s}{r_i}$	$(A P, FF; \alpha) \alpha > \frac{r_i - s}{r_i}$
	$(B P, AF; \alpha) \alpha < \frac{r_i - s}{r_i}$	$(A P, AF; \alpha = 0)$
$b < s$	$(B N, AF; \alpha = 1)$	$(A N, AF; \alpha = 0)$

In contrast, with an investment in detection technology the website owner can partially crowd out the ill-effects of ad avoidance. He can successfully solicit a micropayment even when blocking technology is cheap as long as the user values the content sufficiently (see Table 4). However, the website owner will still not extract any benefits from a user who does not value the content highly and has access to cheap ad-blocking technology. Interestingly, the website owner is indifferent in the latter case about allowing the user to access the content freely (and with blocked ads) or not.

Importantly, the introduction of detection technology also lowers the threshold of what a user considers to be a cheap advertisement blocker, i.e., the consumer now

internalizes the cost of the expected price of the micropayment when making the assessment ($C_B < c - s$).

Table 4 SPNE of submodel with detection.

	$C_B < c - s$	$C_B > c - s$
$s > r_i$		(A P, FF)
$b > s$	$s < r_i$	(B P, FF) (A P, AF)
	$C_B < c - b$	$C_B > c - b$
$b < s$	(B N, FF)	(A N, AF)
	(B N, AF)	




Fig. 3 Case 2: $b > s$ and $s < r_i$.




Fig. 4 Case 3: $b < s$.

We are now proceeding to visualize the space of equilibria from a different perspective in Figures 3 and 4 by integrating the results of the games from the lefthand and righthand side of Figure 2. The figures show how the equilibrium strategies of the players depend on the cost of detection, C_D , and blocking, C_B , technologies, respectively. We are breaking down the results based on the equilibrium beliefs of the website owner, i.e., Figure 3 is split according to the threshold belief, $\alpha^* = \frac{r_i - s}{r_i}$. Figure 4 shows the cases where the website owner is certain about the consumer's strategies. In addition (and not visualized), for the case of high content value, $b > s$, and low ad-revenue, $s > r_i$, we also find that website owners and users select PBNE1 = $(NI|FF, A|P; \alpha = 0)$, independently of C_B and C_D .

5 Simulation approach and results

Our analysis in Section 4 provides a framework that content providers can use to determine which countermeasures concerning ad avoidance they should use to maximize the revenue. Our results show that the best response depends on the type of users that a given website serves. In this section, we illustrate how our framework can be used to determine the best response while taking into account different assumptions about user heterogeneity with respect to user perception of content and ads.

5.1 Simulation Setup

We model the application of our framework to a popular website with specific and unique content that is of high value to its visitors (e.g., Financial Times). Financial Times is a good example as it is a content provider that deploys both monetization strategies: fee-financed and ad-financed. Our game-theoretic analysis shows that the outcome of the game depends mostly on the parameters that characterize visitors of a given website: users' benefit of viewing the content, users' cost of viewing ads with the content and ad revenue that the website earns for each pageview. As discussed in Section 3, the values of per impression ad revenue and users' benefit of viewing the content are available to the stakeholders, namely websites and ad networks. It is more difficult to obtain exact values for users' cost of viewing ads and to do so, websites could perhaps position themselves with respect to the reasons users have named in the survey on why they block ads (Table 1). Depending on how much they match users' criteria, they can estimate their visitors' costs. In addition, as we will show, knowing the distribution of such a variable for which the relevant parameter is the fraction of users who use ad-blocking software (e.g., available from Firefox statistics) is sufficient for the model.

We rely on Web analytic providers, Alexa and Google’s DoubleClick Ad Planner, to obtain the data based on which we can estimate the parameter values. We use the following values in our evaluations:

1. The website receives 1 million pageviews per day, as reported by Google’s DoubleClick Ad Planner [2].
2. In the case of fee-financed content, we consider a micropayment of $s = \$0.321$ per pageview. We compute this value based on the Financial Times’ subscription fee of \$4.99 per week [12] and the 2.22 average number of pageviews per visitor per day, as reported by Alexa [1]. As explained in Section 4, the subscription fee should be the same for all users.
3. We model the impression-based ad revenue per pageview with a beta distribution represented in Figure 5 based on the estimated Cost-per-Mille (CPM) between \$1 and several tens of dollars [5]. CPM is a cost that advertisers pay for thousand impressions and thus we compute the per-pageview ad revenue as $CPM/1000$ for the considered values of CPM. We select skewed distribution as most of advertisers pay CPM in the range of couple of dollars and only a very few major advertisers pay a high CPM in the order of tens of dollars. The total ad revenue that the website can earn in our model is in the range of the reported ad revenues by the top blog websites [9] that have a similar number of daily pageviews [2].




Fig. 5 Distribution of user-generated impression-based ad revenue per pageview.

4. Benefit b (expressed in US dollars) of users viewing the content (Figure 6) is drawn from a beta distribution (in the range of values comparable to the impression-based ad revenue per pageview), such that 25% of the visitors would opt for fee-financed content (i.e., has $b > s$). This number is in compliance with 25% of Financial Times’s visitors paying for digital subscriptions [11]. In addition, for most of the websites user’s benefit is high due to users’ self-selection bias. The exact values are not necessary, the important parameter is the fraction of users accepting to pay the subscription fees.




Fig. 6 Distribution of users' benefits of viewing content per pageview.

5. We consider a population of visitors that consists of: (i) users who dislike ads and therefore use AB software and (ii) users who do not use AB software and are heterogenous in how much they like or dislike ads. We assume that users who have installed AB software associate a high cost (expressed in US dollars) to viewing online ads. The cost may even surpass the benefit they associate to viewing the content. However, users that have not installed AB software may have an arbitrary cost of viewing ads. In other words, we do not interpret the absence of AB software as users necessarily liking ads. Therefore, the parameter c that represents users' costs of viewing ads is drawn from a bimodal distribution (Figure 7), that assigns a high cost to the users with AB software (the second mode of the distribution) and an arbitrary cost to users without AB software (the first and the second mode). The values of c are in the range comparable to the impression-based ad revenue per pageview and users' valuation of the content.
6. In practice, the cost of blocking ads (C_B) corresponds to the cost of installing and maintaining a browser add-on and subscribing to filter lists that define blocking rules. At the moment, the cost (C_D) of detecting AB software on users' machines corresponds to the cost of including a specific Javascript into Web pages. Nowadays, both of these costs (expressed in US dollars) are very small and we use values of $C_B = \$0.01$ $C_D = \$0.001$ for our simulations. Note that these values represent costs per interaction and have such a low value as they are factored out on millions of users (for C_D) and a number of pageviews per day (for C_B). These costs may increase if an arms race develops between AB softwares and detection tools, as it was the case with pop-up ads and pop-up blockers [4]. We evaluate the effect of higher costs of blocking and detection later in the analysis.

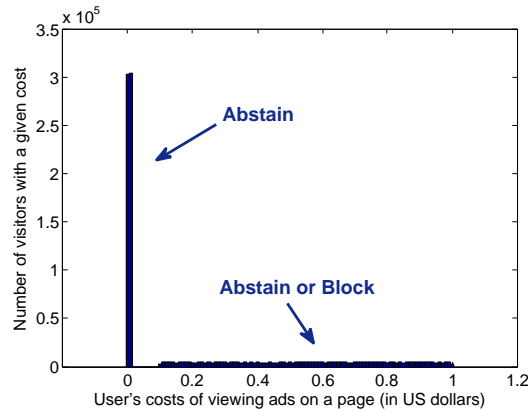


Fig. 7 Distribution of users' costs of viewing ads per pageview (Abstain - users without AB software; Block - users with AB software).

5.2 Results

We simulate the interaction between the website and the population of users, based on our game-theoretic model and parameter values described above. First, we compare the revenues that the website obtains by deploying three different monetizing strategies: (i) serving ad-financed content (AF model) to all visitors, regardless of whether they use AB software or not; (ii) serving fee-financed (FF) content, where only users for which the benefit is higher than the subscription fee are the source of revenue; (iii) game-theoretic approach (GT model) where a website chooses an appropriate strategy according to our analysis, and can either serve ad-financed or fee-financed content to different users. The fraction of users that use AB software is a variable in our simulations.

Figure 8 depicts the daily revenue of the website, for the three models, depending on the fraction of users that use AB software. We observe that the revenue that the website obtains with GT monetizing model is superior to using pure fee-financed or ad-financed models. The reasoning behind such a result is as follows. In the AF model, users with AB software do not generate ad revenue for the website, as ad impressions are blocked on their machines. The higher is the number of users with AB software, the higher is the revenue loss for the website. In the FF model, only users who value the content more than the subscription fee are willing to pay, thus the revenue is not influenced by the users who use AB, only by the number of subscriptions. FF revenue depends on the subscription fee that the website can charge, which mostly depends on the content it serves and how valuable it is to its visitors. The GT model represents a compromise between AF and FF. For users who dislike ads, but value content enough to pay subscription fees, the website will apply FF strategy. With AF, the website cannot make profit out of these users as they block ads. For users who do not dislike ads as much, the website might either use FF or




Fig. 8 Website's daily revenue (in US dollars) with different monetizing models.

AF strategy, whichever is more profitable. Thus, the GT model enables the website to take into account users' heterogeneity and maximize its profit.




Fig. 9 Fraction of visitors that generate revenue for each monetizing model.

In Figure 9 we show the fraction of users that generate profit for the website with the three monetizing models. Curve labeled *AF-Block* represents the fraction of users from which the website profits in the AF model. In this model, the ad revenue is generated only by the users without AB software. Note that nevertheless all users obtain the content. In the fee-financed model, only the users who opt to pay the subscription generate the revenue for the website and obtain the content (*FF* curve²). In the GT model, the website profits from serving ad-financed content to a fraction of users (*GT-AF* curve) and fee-financed content to another fraction of users

² Note that *FF* curve overlaps with *GT-FF* curve.

(*GT-FF* curve). Users served with ad-financed content are those who: (i) accept to view ads in exchange for free content (for which the outcome of the game is PBNE 9: $(NI|AF, A|N; \alpha = 0)$), or (ii) value the content more than they dislike ads, but not enough to pay the subscription fee for ads-free content (for which the outcome of the game is SPNE 1: $(DI|AF, A|N)$). Users who are served fee-financed content are those who: (i) dislike ads but value the content, or (ii) users who accept ads but also value the content, thus leaving the choice to the website that may decide to offer the subscription model to such users as it may be more profitable. These are the users for which the outcome of the game is PBNE 1: $(NI|FF, A|P; \alpha = 0)$. We observe that the total fraction of users that generate the revenue for the website in the GT model (*GT-Block*) is higher than in either AF or FF model.

Users who do not generate revenue and do not obtain the content in the GT model are those who dislike ads and do not value the content enough to pay subscription fees. This case corresponds to PBNE 8: $(NI|AF, B|N; \alpha = 1)$. Note that the impact of the users with AB software is smaller in the GT model, and we see that in the worst case about 27% blocks ads (and generate revenue loss for the website) compared to the 50% in the AF model. These results are in line with the results in Figure 8 and explain why the website earns more with the GT monetizing model. In the worst case, the GT revenue is around 12% higher than FF revenue and it may not seem justified to deploy GT model for that increment in the revenue. However, one major advantage of the GT model is that it maximizes the number of users who obtain the content (73% in the GT model compared to 25% in the FF model, in the worst case). We conclude that the GT model allows the website to adapt its monetizing strategy such that it maximizes the number of visitors from whom it profits, as well as its visibility or impact factor.

As discussed previously, the website may deploy a strategy of making it more difficult for AB software to filter out and block ads. In our GT model, this action can be represented with an increase in the users' cost of blocking ads and a higher investment in the detection. We simulate the effect of a higher blocking and investment costs ($C_B \in \{0.01, 0.1, 0.5, 0.7, 1\}$ and $C_D = \$0.1$) and represent the results in Figure 10. Different curves correspond to the fraction of users that the website can profit from in the GT model, considering a different cost of blocking. We observe that the fraction of users that will block ad-financed content decreases with the increase in the cost of blocking ads. As both the website and users are behaving strategically in the GT model, with the higher cost rational users deter from blocking ads and it shows that the website has a good return-on-investment with the strategy of making ad blocking more difficult.

In summary, we have illustrated how a website may use our framework in practice to determine the best monetizing strategy given the population of users it serves. We have demonstrated how a website maximizes its revenue with a strategic choice of its best response when facing users with different preferences with respect to ads and content. Such a strategic behavior allows for the website to maximize the number of users from which it can profit, as well as to apply the strategy that maximizes the profit. Users' strategic behavior allows them to maximize their utility as well, by having a choice of viewing ad-financed or fee-financed content.




Fig. 10 Fraction of visitors that generate revenue in the GT model, considering higher blocking and detection costs.

6 Conclusion

In this paper, we conduct a systematic study of the consequences of ad-blocking on the business model of content providers. We develop a framework usable by content providers to ponder their options to mitigate the consequences of ad avoidance techniques. We carefully devise and analyze a game-theoretic model of the impression-based ad revenue mechanism and illustrate with simulations the impact of different strategies under parameter assumptions motivated by real-world data. Our analysis shows that deploying a game-theoretic approach, i.e., strategically applying fee-financed or ad-financed monetization strategy, and treating each user individually yields higher revenues for publishers, compared to deploying one strategy across all users. Also, understanding the distribution of users' aversion to ads and valuation of the content is essential for publishers to make a well-informed decision. We expect that our modeling and simulation assumptions are a reasonable, but likely not a perfect fit for every situation involving content providers and ad avoiders. In future work, we intend to further explore deviations from our modeling assumptions and expand our framework to additional problem areas.

The current paper is only a first step to account for the complicated interactions between ad avoidance and content monetization. For example, a promising area for additional work is to more carefully address the impact of the negative feedback spiral caused by the presence of ad-blocking under the presence of limited information. A loss of revenue through an increase of website visitors who use ad-blocking software will frequently trigger a more aggressive pursuit of advertisement opportunities. Those might even include consumer-unfriendly affiliate marketing schemes. While this may create short-term benefits, additional consumers will depart or try to avoid these practices.

In addition, we aim to consider measures of concentration and interdependency in the ad industry. For example, a recent study shows that Google-controlled cookies

were present on 97 of the top 100 websites [18]. The same study also documents the growing intricacy of tracking attempts that will make it very difficult for users to find adequate countermeasures in absence of market (self-)regulation.

In conclusion, we expect content providers that serve a technology-minded audience to suffer most from ad avoidance technologies. And, in absence of a broad consensus between the ad and content industry, on the one side, and consumers, on the other side, the trend towards blocking of advertisement is likely to grow. Resistance to user tracking and the desire for ad blocking are tightly interwoven, even though we do not model the related long term trends in the moment, i.e., users rarely become technology-savvy ad avoiders over night. However, the potential for a significant shift in consumer behavior is large and should not be under-appreciated.³

7 Acknowledgments

We thank the anonymous reviewers and workshop participants for their valuable comments and feedback. The presentation at the workshop is partially supported by travel funding from the Volkswagen Foundation.

References

1. Alexa, The Web Information Company.
Available online at <http://www.alexa.com/siteinfo/ft.com>
2. DoubleClick Ad Planner by Google.
Available online at <https://www.google.com/adplanner/>
3. Bezos calls Amazon experiment "a mistake". Available online at <http://www.bizjournals.com/seattle/stories/2000/09/25/daily21.html> (2000)
4. Can't stop the pop-ups. Available online at http://news.cnet.com/Can%27t-stop-the-pop-ups/2100-1024_3-5226273.html (2004)
5. The Average CPM Rates Across Different Verticals. Available online at <http://www.labnol.org/internet/average-cpm-rates/11315/> (2010)
6. Arstechnica Opposition Letter. Available online at <http://static.arstechnica.com/oppositionletter.pdf> (2011)
7. Google One Pass.
Available online at <http://www.google.com/landing/onepass/> (2011)
8. NYTimes' "Fair" Prices. Available online at <http://www.mondaynote.com/2011/03/21/nytimes-%E2%80%9Cfair%E2%80%9D-prices/> (2011)
9. Top Earning Blogs. Available online at <http://onlineincometeacher.com/money/top-earning-blogs/#.Tre9pLJkxbw> (2011)
10. Data Collection Arms Race Feeds Privacy Fears. Available online at www.chicagotribune.com/business/sns-rt-us-data-collection-tre81i0ap-20120219,0,2999768.story (2012)
11. Financial Times, Digital Subscribers. Available online at <http://aboutus.ft.com/corporate-information/ft-company/#axzz1d16rUEub> (2012)

³ A 2010 study revealed that up to 40% consumers are willing to change their online behavior if advertisers were collecting data [31].

12. Financial Times, Subscription Fees. Available online at <https://registration.ft.com/signup/standard/?segid=70152&execution=els1> (2012)
13. Google Privacy Changes Must Be Stopped, Group's Lawsuit Says. Available online at <http://www.businessweek.com/news/2012-02-13/google-privacy-changes-must-be-stopped-group-s-lawsuit-says.html> (2012)
14. PayPal Merchant Services. Available online at https://merchant.paypal.com/cgi-bin/marketingweb?cmd=_render-content&content_ID=merchant/digital_goods (2012)
15. Acquisti, A., Varian, H.R.: Conditioning prices on purchase history. *Marketing Science* **24**, 367–381 (2005)
16. Anderson, S., Gans, J.: Platform siphoning: Ad-avoidance and media content. *American Economic Journal: Microeconomics* **3**(4), 1–34 (2011)
17. Aycock, J.: *Spyware and Adware (Advances in Information Security)*. Springer Verlag (2010)
18. Ayenson, M., Wambach, D.J., Soltani, A., Good, N., Hoofnagle, C.J.: *Flash Cookies and Privacy II: Now with HTML5 and ETag Respawning*. World Wide Web Internet And Web Information Systems (2011)
19. Cho, C., Cheon, H.: Why do people avoid advertising on the internet? *Journal of Advertising* **33**(4), 89–97 (2004)
20. Cohen, N.: Whiting out the ads, but at what cost? *The New York Times* (2007)
21. Commission, F.T.: Protecting consumer privacy in an era of rapid change: A proposed framework for businesses and policymakers. Preliminary FTC Staff Report. (2010)
22. Edelman, B., Ostrovosky, M., Schwarz, M.: Internet advertising and the generalized second-price auction: Selling billions of dollars worth of keywords. *American Economic Review* **97**(1), 242–259 (2007)
23. Fisher, K.: Why ad blocking is devastating to the sites you love. *Ars Technica* (2010)
24. Fudenberg, D., Levine, D.: Subgame-perfect equilibria of finite- and infinite-horizon games. *Journal of Economic Theory* **31**(2), 251–268 (1983)
25. Fudenberg, D., Tirole, J.: Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory* **53**(2), 236–260 (1991)
26. Ha, S.: An intelligent system for personalized advertising on the internet. In: K. Bauknecht, M. Bichler, B. Pröll (eds.) *E-commerce and Web technologies (LNCS 3182)*, pp. 21–30. Springer Verlag (2004)
27. Haskins, J.: Commercial skipping technology and the new market dynamic: The relevance of antitrust law to an emerging technology. *Duke Law & Technology Review* (6) (2009)
28. Isaacson, W.: How to save your newspaper. *TIME Magazine* (2009)
29. Kelly, L., Kerr, G., Drennan, J.: Avoidance of advertising in social networking sites: The teenage perspective. *Journal of Interactive Advertising* **10**(2), 16–27 (2010)
30. Kshetri, N.: The economics of click fraud. *IEEE Security & Privacy* pp. 45–53 (2010)
31. McDonald, A., Cranor, L.: Americans' attitudes about internet behavioral advertising practices. In: E. Al-Shaer, K. Frikken (eds.) *WPES*, pp. 63–72 (2010)
32. Mungamuru, B., Weis, S.: Competition and fraud in online advertising markets. In: G. Tsudik (ed.) *Financial Cryptography, LNCS*, vol. 5143, pp. 187–191. Springer (2008)
33. Okada, A.: Perfect bayesian equilibrium and sequential equilibrium. In: J. Cochran, L. Cox, P. Keskinocak, J. Kharoufeh, C. Smith (eds.) *Wiley Encyclopedia of Operations Research and Management Science*. John Wiley & Sons (2010)
34. Palant, W.: Adblock Plus User Survey. Available online at <http://adblockplus.org/blog/adblock-plus-user-survey-results-part-2> (2011)
35. Press, A.: Goldman Sachs forecasts online ad slowdown in 2012. Available from multiple news sources. (2011)
36. Rainie, L., Purcell, K.: *State of the news media 2010: Online economics and consumer attitudes*. Report produced by the Pew Internet Project and the Pew Research Center's Project for Excellence in Journalism. (2010)
37. Shah, S.: *Ad-Skipping and Time-Shifting: A Theoretical Examination of the Digital Video Recorder* (2011). Working paper. University of Virginia.

38. Sindik, A., Graybeal, G.: Newspaper micropayments and millennial generation acceptance: A brand loyalty perspective. *Journal of Media Business Studies* **8**(1), 10–20 (2011)
39. Singh, A., Potdar, V.: Blocking online advertising - A state of the art. In: Proceedings of the 2009 IEEE International Conference on Industrial Technology (ICIT) (2009)
40. Smith, W.: Consumer Resistance to Marketing Reaches All-Time High, Marketing Productivity Plummetts, According to Yankelovich Study (Yankelovich President J. Walker Smith Addressed Topic at AAAA Conference Today) (2004)
41. Snow, N.: The TiVo question: Does skipping commercials violate copyright law? *Syracuse Law Review* **56**(1), 27–84 (2005)
42. Speck, P., Elliott, M.: Predictors of advertising avoidance in print and broadcast media. *Journal of Advertising* **26**(3), 61–76 (1997)
43. Tåg, J.: Paying to remove advertisements. *Information Economics and Policy* **22**(4), 245–252 (2009)
44. Turow, J., King, J., Hoofnagle, C., Bleakley, A., Hennessy, M.: Americans reject tailored advertising and three activities that enable it. Available at University of Pennsylvania Scholarly Commons. (2009)
45. Vallade, J.: AdBlock Plus and the legal implications of online commercial-skipping. *Rutgers Law Review* **61**(3), 823–853 (2009)
46. Varian, H.: Position auctions. *International Journal of Industrial Organization* **25**(6), 1163–1178 (2007)
47. Varian, H., Wallenberg, F., Woroch, G.: The demographics of the do-not-call list. *IEEE Security & Privacy* **3**(1), 34–39 (2005)
48. Wilbur, K.: A two-sided, empirical model of television advertising and viewing markets. *Marketing Science* **27**(3), 356–378 (2008)

Appendix

Table 5 Symbols for the game-theoretic models.

Symbol	Definition
b	User's "benefit" of viewing content
c	User's "cost" of viewing ads
s	Subscription fee
r_i	Ad Network's per-impression ad revenue
C_B	Cost of using AB software
C_D	Cost of detecting AB software
α	Belief about the reached information set
P_W	Website's total payoff
P_U	User's total payoff
AF	Ad-financed content
FF	Fee-financed content (micropayments)
DI	Invest in detection of AB softwares
NI	No investment in detection of AB softwares
B	Block ads
A	Abstain from blocking ads
P	Pay subscription
N	Not pay subscription
$(x y)$	First (x) and second (y) action of a player
(x,y)	Strategy profile: (first mover,second mover)