informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–23
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

Ad Blocking

Aleksandr Gritckevich,[a] Zsolt Katona,[b] Miklos Sarvary[a]

[a] Columbia Business School, Columbia University, New York, New York 10027; [b] Haas School of Business, University of California, Berkeley, Berkeley, California 94720

Contact: agritckevich20@gsb.columbia.edu (AG); zskatona@haas.berkeley.edu, https://orcid.org/0000-0001-8411-6040 (ZK); miklos.sarvary@columbia.edu (MS)

Received: March 8, 2021
Accepted: April 5, 2021
Published Online in Articles in Advance: October 25, 2021

https://doi.org/10.1287/mnsc.2021.4106

Abstract. In recent years, ad blocking has become a significant threat to advertising-supported content. Adblockers typically negotiate with publishers, allowing some ads to go through in return for a payment, a practice called (partial) *whitelisting* in the industry. Ad blocking has a direct positive effect on consumers by reducing advertising intensity. On the other hand, the practice clearly hurts publishers and reduces their incentives to invest in content quality. Lower content quality, in turn has an indirect negative effect on consumers. This paper builds an analytic model to explore the net impact of ad blocking on consumers, how it depends on various market characteristics, and how uniformly it affects consumers. The results show that under a broad set of market conditions, total consumer surplus and even total welfare decline under ad blocking. Whereas some consumers are always better off with an ad blocker, for the average consumer, the impact of quality decline is larger than that of ad reduction. The analysis highlights the detrimental role of ad blockers' current revenue model—in which value is created for the consumers but it is captured from publishers—in decreasing quality, consumer surplus, and total welfare. Analyzing the impact of varying levels of negotiation power between the ad blocker and publisher reveals that full negotiation power is not preferred by the ad blocker. A lower negotiation power allows the ad blocker to commit to less value extraction from the publisher, thereby leading to higher content quality. Additional model extensions show that the main results are robust. In the case of multiple publishers with different levels of competition between them, the strong negative effect of ad blocking on quality holds.

History: Accepted by Juanjuan Zhang, marketing.
Supplemental Material: The web appendix is available at https://doi.org/10.1287/mnsc.2021.4106.

Keywords: advertising • game theory • media competition • extortion

1. Introduction

Ad blocking, the practice of installing a third-party software or app to block advertisements from loading on visited websites, has grown steadily since 2013 to reach hundreds of millions of internet users worldwide. In the United States alone, eMarketer estimates that in 2020, approximately 26.4% of U.S. internet users installed ad blockers on their connected devices (Reyes 2020). Although this proportion is stabilizing, eMarketer notes that it still represents 76.1 million users and threatens about a quarter of publishers' (dollar) 142 billion total U.S. online advertising revenues. The use of ad blockers is more popular in some other geographies (e.g. in the Asia-Pacific or Eastern Europe), but recent studies indicate that in developed countries, usage rates are stabilizing. For example, in the United Kingdom, the Association for Online Publishing found that by the end of 2019, ad blocking rates actually fell slightly compared with the previous year to 20.6% (Perrin 2019). Importantly, in all countries, a significant number of consumers never use ad block software even though these applications are nominally free. Clearly, the use of ad blockers is not without cost to consumers.[1]

Ad blocking's benefits to consumers are obvious: ad blocking significantly reduces the intensity of advertising, which is a direct transfer of surplus to consumers. This direct positive effect is certainly one factor behind the public's generally favorable attitude toward ad blockers and may also explain ad blockers' favorable treatment by regulators. Indeed, in their public relations, ad blockers often cultivate a Robin Hood image, claiming to represent a consumer movement against intrusive online advertising.

In contrast, the practice is a major threat to content providers (publishers), whose revenue model is based on advertising. Publishers complain to be victims of "extortion" and liken ad blocking to highway robbery (Katona and Sarvary 2018). To respond to the threat to their revenues, some publishers (e.g., Axel Springer, a large German publishing group, the *Washington Post*, or the United Kingdom's ITV and Channel 4) have denied content to their readers who blocked ads, encouraging them to stop doing so or to purchase a subscription (Vasagar 2015). These early strategies provided mixed results. In the case of Springer, rather than paying for a subscription, most users switched off their ad blocker, but some simply left the site. It is also clear that such strategies are unlikely to work for lesser-known publishers who do not have enough resources to set up and manage a subscription system.

Publishers have also tried to bring ad blockers to court. For example, Eyeo, the German owner of the world's most popular ad blocker, Adblock Plus, faced (and won) half a dozen lawsuits in the past years. The question even reached the Supreme Court in Germany, who ruled ad blocking to be legal.[2] It is clear that neither the courts nor regulators seem to take the side of publishers in their fight with ad blockers.

There is no doubt that ad blocking hurts publisher profits and, since it reduces the returns to investing in quality content, indirectly, it results in lower content quality (Shiller et al. 2018). Putting aside the externalities associated with a general decline in, say, news quality, even a narrow focus on consumer welfare should ask: does the reduced intensity of advertising compensate consumers for the decline in content quality? Asked differently, is the current regulatory stance, which unambiguously supports ad blocking in its present form, a wise approach to regulation?

The purpose of this paper is to shed light on these questions and examine the welfare implications of the current practice of ad blocking with a special focus on consumer welfare. Given regulators' and courts' unambiguously positive attitude toward ad blockers, it is important to examine if indeed ad blocking increases, or is at least neutral to welfare. Specifically, the paper builds an analytic model to explore if and when the direct positive effect of ad reduction on consumers outweighs the indirect negative effect of lower content quality and how total welfare changes in the presence of ad blocking. Surprisingly, we find that under a broad set of market conditions, total consumer surplus and even total welfare decline under ad blocking. Whereas some consumers are always better off in the presence of ad blocking, for the average consumer, the negative impact of quality decline is larger than the positive effect of ad reduction.[3]

Importantly, beyond welfare implications, our analyses also highlight the role of ad blockers' unusual revenue model in sustaining the industry ecosystem. Indeed, although ad blockers create value for consumers, they almost never charge consumers for their services. No doubt, a reason behind this approach is related to the relative technical difficulty of collecting small fees from a large and fragmented customer base, many of whom are used to and expect free digital products. Instead, ad blockers negotiate with publishers for a fraction of their ad revenues in exchange for a limited amount of, presumably higher quality, ads to pass the blockade. This practice, which may affect up to 30% of the publishers' ad revenues is known as (partial) *whitelisting* in the trade (Maheshwari 2016).[4]

In deciding the right amount of advertising volumes allowed, a key tradeoff for ad blockers is how attractive they remain to consumers and how much surplus they can extract from publishers. The lower the ad volume, the more consumers are likely to download the ad blocker app, which increases the ad blocker's leverage over the publisher. However, a lower advertising volume also means that the potential surplus extracted from a publisher is also lower.[5] Making ad blocking software attractive to consumers is a key consideration. Recent trends of ad blocker usage indicate that the proportion of consumers using them (though different across countries) has stabilized. Importantly, the majority of consumers still do not use ad blockers.

Our model takes into account these complex strategic interactions in the ecosystem of publishers, ad blockers, and consumers. In particular, we pay special attention to consumer heterogeneity given the nonuniform behavior of consumers in using ad blockers and the known differences across consumers in their sensitivity to advertising and content quality. Our analysis considers a single ad blocker and one publisher that relies on an advertising revenue model. The publisher makes a costly investment in content quality. Consumers are assumed to be heterogeneous in three dimensions: (i) their sensitivity to quality, (ii) their disutility from advertising, and (iii) their cost of downloading the ad blocker app. We show that all three of these consumer characteristics play a critical role in understanding the complex dynamics of ad blocking. The ad blocker sets an upper limit to the ad volume seen by consumers who downloaded its app. Furthermore, we assume that it can also reduce ad volume for these consumers to zero. This threat allows the ad blocker to extract surplus from the publisher.

The solution of this simple model is consistent with the anecdotal evidence observed in the industry. In equilibrium, the ad blocker does not block all ads and collects a fee from the publisher, which is reminiscent of the practice of whitelisting in the industry. A fraction of consumers—those who (i) are sensitive to advertising, (ii) have high valuation for quality, and (iii) have a relatively low cost of downloading—use the ad blocker's app. We find that consumer heterogeneity is critical in explaining in which way consumer welfare changes in the presence of ad blocking. In equilibrium, ad blocking always increases the surplus of some consumers; in fact, some consumers only consume content in the presence of an ad blocker (i.e., the ad blocker may expand the market). However, unless consumer heterogeneity is very limited both in advertising sensitivity and the valuation of quality, overall consumer welfare is lower with ad blocking. In the more realistic case of heterogeneous consumers, the publisher decreases its quality substantially enough to erase the benefit of lower ad exposure for most consumers. As a result, not only is the publisher worse off in the presence of ad blocking, but total consumer surplus and overall welfare also decline. If consumers are homogeneous in both advertising disutility and the appreciation of quality, then overall consumer welfare can increase under ad blocking. In this case, without the ad blocker, the publisher can extract almost all surplus from consumers. Although ad blocking still reduces quality, the effect on consumers is small as they did not have much surplus to begin with. In contrast, advertising reduced by the ad blocker provides a substantial benefit.

Overall, our first result suggests that the revenue model where ad blockers charge publishers is more akin to the highway robbery narrative than the Robin Hood image that ad blockers promote. It is true that ad blockers transfer some surplus from publishers to consumers, but they also collect substantial rents in the process. In an ecosystem with a standard publisher-consumer relationship where both sides walk away with healthy surpluses, the entry of an ad blocker can be highly detrimental to publishers and, by extension, to consumers. The situation becomes even worse if ad blockers provide a full ad blocking option to consumers. Hence, we conclude that, in contrast to prevailing regulatory practice, ad blocking is typically not an efficient market solution to regulate the amount of advertising shown on websites.

To further substantiate this conclusion, we calculate welfare measures for a hypothetical model where ad blockers directly charge consumers for ad blocking in a traditional fee for service revenue model. We find that under realistic market conditions with heterogeneous consumers, this revenue model is better for publishers as well as consumers than the current practice of publisher extorsion. In the case of a heterogeneous consumer base, more surplus is left with the publisher, which limits the decline of content quality that, in turn, also increases average consumer surplus. We conclude that restricting ad blockers' value extraction practices would be a relatively simple way to unambiguously increase publishers' and consumers' welfare.

In the second part of the paper, we analyze several variations in our setting to explore important institutional details. First, we explore the case when the publisher has more (and varying levels of) power in its negotiation with ad blockers. This is a relevant practical consideration given the variety of investments publishers consider to protect themselves against ad blocking.[6] The results show that even at very high levels of publisher power, ad blocking usually has a considerable negative effect on content quality and welfare. In fact, surprisingly, we find that if the ad blocker could choose its negotiation power at no cost (and commit to it), the ad blocker would exercise restraint by renouncing to extracting all potential surplus from the publisher. This suggests that an additional avenue for a more efficient regulation may be to help the coordination of incentives between publishers and ad blockers for a better consumer (and general) welfare outcome.

Our final step is to discuss a variety of robustness checks that are based on additional analyses documented in the web appendix. Specifically, for the more common case of heterogeneous consumers, we extend the base model to consider multiple publishers with different levels of competition between them. Our main results concerning the practice of whitelisting, publishers' quality choice, consumer surplus, and general welfare remain similar, although we also highlight a variety of mechanisms through which an ad blocker common to multiple publishers may affect (typically increase) competition between them.

Finally, we look at the case of a transaction platform offering ad blocking on its own. This last setup is relevant for multiple reasons. First, more and more content consumption happens on platforms where consumers and publishers interact (e.g., iOS, Google, or Facebook). Second, such platforms have tremendous power in regulating these interactions. We show that the platform tends to internalize competition between the ad blocker and the publisher, leading to better welfare outcomes. However, unless the platform can fully commit to the maximum level of advertising before the publisher sets quality, the platform cannot provide ad blocking services that would improve on our benchmark, that is, the case without any ad blocking.

The rest of the paper is organized as follows. In the next section, we survey the relevant literature. Section 3 describes and analyzes the main model featuring a single publisher and an ad blocker that may offer various levels of ad blocking to consumers. This section also compares ad blockers' current business practices to a more conventional revenue model. Section 4 explores the ad blocker-publisher relationship and negotiations in more detail. Section 5 discusses extensions to gauge the model's robustness and discusses limitations. The paper ends with concluding remarks. To improve readability, all proofs are relegated to the appendix and some extensions are detailed in a separate web appendix.

2. Relevant Literature

A fundamental driving force behind ad blocking is that most consumers dislike ads. They try to mitigate the disutility caused by avoiding ads, which raises a number of important questions. Anderson and Gans (2011) study the impact of ad avoidance technologies and their effects on content providers. They find that the reaction to ad avoidance leads to increased ad clutter and reduced content quality and decreases welfare. In another paper, Johnson (2013) considers how better targeted ads interact with ad avoidance, showing that consumers may underutilize ad avoidance tools. Eventually, consumers and content providers end up in a game where each player uses some methods to avoid or push through ads (Vratonjic et al. 2013).

On the empirical side, Goldstein et al. (2014) examine the costs of placing ads that are annoying. They find that such ads can cause users to abandon the website and negatively affect the process of consuming the site's content, potentially costing more than the revenue from the ads. Shiller et al. (2018) make the first attempt at thoroughly documenting the impact of ad blocking on online publishers. Using web traffic and ranking data for a broad set of sites, they show that, as a larger proportion of consumers use ad blockers, publisher sites decline in their rankings (e.g., their rank becomes larger) and this effect is larger for worse-ranked sites. They conclude that "ad blocking poses a substantial threat to the ad-supported web (in Abstract)." Our theoretical results are largely consistent with this general empirical finding.

More recent theoretical work in marketing challenges this general view. In particular, in a model with exogenous ad blocking, Despotakis et al. (2021) show that under some conditions, ad blocking can benefit competing publishing platforms as they allow them to discriminate between consumers who have high and low ad sensitivity, respectively.[7] They also find that under some conditions, content quality and overall welfare can also increase. Our model is fundamentally different from their setup in that it assumes the ad blocker to be an active, profit-maximizing agent. Our results attenuate some of the results by Despotakis et al. (2021) and reverse others. On the one hand, we do find that some consumers benefit from the presence of ad blockers and, for fixed quality, the average consumer is better off. On the other hand, we show that quality is lower in the presence of an active ad blocker, and this leads to lower consumer and overall welfare in equilibrium. We therefore argue that it is critical to model ad blockers as active agents.

On the conceptual front, our paper relates to several streams of literature. The base model of publishers providing content to consumers in return for showing advertisements is reminiscent of typical two-sided markets (Rochet and Tirole 2006), although for simplicity, we assume exogenous advertising prices. We need this simplification because with an ad blocker, the model extends to a three-sided platform. On the consumer side, we assume a standard utility model with heterogeneity at multiple levels. In particular, in terms of consumers' valuation for content quality, we adopt the well-accepted approach of the vertical differentiation literature (see Shaked and Sutton 1982, Moorthy 1988). Furthermore, the ad blocking ecosystem exhibits characteristics of double marginalization, well-documented in the literature (Spengler 1950, Bresnahan and Reiss 1985).

At the high level, ad blocking shows similarities with online music and software piracy, as a third party (in our case the ad blocker) allows users to consume content without proper compensation (advertising eyeballs in case of ad blocking). Piracy has been extensively studied in the marketing and economics literature. Prominent work includes Jain (2008), who investigates digital piracy from a competitive perspective, showing that weaker copyright protection may be beneficial to firms selling content because it softens price competition by allowing price-sensitive consumers to pirate the content. Danaher et al. (2010) show how reducing the availability of content on paid channels increases the demand for the pirated version of the same content, whereas Vernik et al. (2011) look at digital rights management, finding that piracy might decrease even if usage restrictions are removed from digital products.

A distinct, but fairly small body of literature shows interesting resemblances to our research questions. As ad blockers essentially hold content publishers "hostage," there is a clear connection to the economic analysis of extortion. Several important details of ad blocking relate to the mechanics underlying extortion. For example, Konrad and Skaperdas (1997) examine the credibility of threats related to extortion. They find that an equilibrium with extortion only exists if there is a large enough number of victims. Choi and Thum (2004) investigate the dynamics of repeated extortion. In their setting, government officials demand bribes in a repeated fashion, which shows similarities to ad blocking, and they also find that a pure strategy equilibrium does not exist. In a fascinating setting, Olken and Barron (2009) study extortion in the field by analyzing the bribes demanded from truckers by corrupt officials. The findings demonstrate how multiple layers of extortion interact with higher amounts of bribes closer to the end of the journey. Overall, our contribution lies in uncovering the dynamics of the strategic interaction between ad blockers, content publishers, and consumers. Besides shedding light on some new industry practices, the paper examines the impact of ad blocking on content quality, consumer surplus, and overall welfare.

3. Base Model

Our base model has three types of active players: (i) a publisher, (ii) consumers, and (iii) an ad blocker. The price of advertising (impression) is assumed to be exogenous and constant, denoted p_a.[8] The publisher chooses the quality level of its content q at cost $c(q) = cq^2$ and collects revenue through advertising by choosing the level of advertising volume, A.[9]

The ad blocker offers a tool (e.g., an app or browser plug-in) to consumers free of charge that limits the advertising on the publisher's site. We denote the amount of advertising the ad blocker lets through by $\mathcal{V}$. We assume that consumers incur cost κ when downloading (using) the ad blocker app, which varies across consumers. The ad blocker commits to a maximum advertising level V, such that $\mathcal{V} \le V$. Setting a lower V results in more consumers downloading the app. For the consumers who download the app, the ad blocker can also decide to block all ads by setting $\mathcal{V} = 0$ if it wants to. The threat of committing to $\mathcal{V} = 0$ allows the ad blocker to extract surplus from the publisher. The basic tradeoff faced by the ad blocker is to commit to a low enough V such that many consumers download the app, but not too low so that it can still extract sufficient surplus from the publisher. It is important to highlight that this setup is usually the default setting of large ad blockers (e.g. the dominant ad blocker, Adblock Plus). However, in practice, ad blockers also provide consumers with the option to fully block all advertising at some additional (hassle) costs. In Section 3.3, we analyze this practically relevant modification and show that, beyond reinforcing our main results, it also provides key insights with respect to the dynamics of potential competition between ad blockers.

A unit mass of consumers are distributed in a three-dimensional cuboid with their total number normalized to 1.[10] Consumers get utility from visiting the publisher's site in the amount of

$$U = \max(\vartheta q - \gamma A, \vartheta q - \gamma \mathcal{V} - \kappa, 0),$$

where $\vartheta \sim U(t, 1)$ is consumers' sensitivity to quality, $\gamma \sim U(r, 1)$ measures their sensitivity to advertising, and $\kappa \sim U(0, \bar{\kappa})$ is consumers' cost for downloading the ad blocker app. We assume that ϑ, γ, and κ are independently distributed from each other. We also assume that $0 \le t \le 1, 0 < r \le 1$ and that $\bar{\kappa}$ is large enough.[11] The restriction on $\bar{\kappa}$ is similar to assuming that—in line with practice—there is a sufficiently high proportion of consumers who never download the ad blocker app.

The publisher's profit depends on whether consumers choose to visit its site and how much advertising they see. The publisher's profit is

$$D_A A p_a + D_{\mathcal{V}} \mathcal{V} p_a - cq^2 - T,$$

where D_A is the number of consumers who read without using the ad blocker; $D_{\mathcal{V}}$ is the number that reads with the ad blocker installed; and T denotes a payment the publisher makes to the ad blocker. This payment represents the rent extracted by the ad blocker from the publisher as the ad blocker can threaten to lower the ads to zero for consumers who download the app. Since the payment may be decided through negotiations, the details of the negotiation process have a high impact on its outcome. In the basic model, we assume that the ad blocker moves first and makes a take-it-or-leave-it offer for T combined with a promise of a low $\mathcal{V}$. We assume that the publisher always accepts when indifferent. We will see that in this setup, the ad blocker extracts all advertising revenue generated from consumers who downloaded the app. In Section 4, we explore the case when the ad blocker cannot extract all surplus from the publisher (for example, when the publisher has substantial negotiation power or when it can introduce a subscription model).

The timing of the game is the following: First, the publisher chooses the level of quality q. Then the ad blocker sets V, the maximum advertising level when using the ad blocker. This timing reflects the fact that investing in quality (e.g., a brand or a large editorial staff) is a much larger commitment than setting the level of advertising, V. Given V, consumers decide whether to download the ad blocker app. Next, the ad blocker makes a take-it-or-leave-it offer and the publisher accepts or rejects. Then the ad blocker can set $\mathcal{V} \le V$ and finally, the publisher decides on the volume of advertising, A. Given $\mathcal{V}$ and A, consumers decide on content consumption and firms' profits are realized.

3.1. Benchmark: No Adblocker

We first examine the case without the presence of an ad blocker as a benchmark. For simplicity, we present a brief analysis for $t = 0$ and relegate $0 < t \le 1$ to the appendix. Consumers visit the publisher's website when $\vartheta q - \gamma A \ge 0$. Without the ad blocker, κ is irrelevant, hence we can derive the consumer demand by considering only ϑ and γ. Figure 1 illustrates consumer choice by representing zero utility lines on the (γ, ϑ) graph. Consumer demand is the area above the line corresponding to a particular scenario, that is, $A \le q$ or $q \le A$. There are three cases depending on the values of q and A:

Case 1: $A \le q$. Demand equals the area of trapezoid (abdc) multiplied with density $\frac{1}{1-r}$:

$$D(q, A) = \frac{(1-r)\left(1 - \frac{A}{q} + 1 - \frac{A}{q} r\right)}{2(1-r)} = 1 - \frac{A}{2q}(1 + r).$$

Case 2: $rA \le q \le A$. Demand equals the area of triangle (def) multiplied with density $\frac{1}{1-r}$:

$$D(q, A) = \frac{\left(1 - \frac{A}{q} r\right)\left(\frac{q}{A} - r\right)}{2(1-r)} = \frac{(q - Ar)^2}{2Aq(1-r)}.$$

Figure 1. Consumer Demand as a Function of ϑ and γ for Different Values of A and q with $r > 0$ and $t = 0$

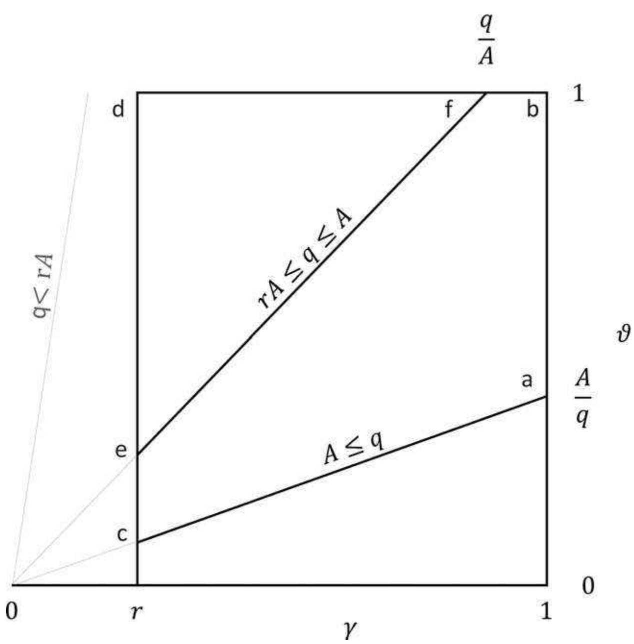

Note. Demand is the area below the $\gamma = \frac{q}{A}\vartheta$ line.

Case 3: $q < rA$. In this case, there are no consumers visiting the publisher, hence $D(q, A) = 0$.

Summarizing the three cases, consumer demand faced by the publisher is:

$$D(q, A) = \begin{cases} 1 - \dfrac{A}{2q}(1 + r) & \text{if } A \le q, \\ \dfrac{(q - Ar)^2}{2Aq(1 - r)} & \text{if } rA \le q \le A, \\ 0 & \text{if } q < rA. \end{cases} \tag{1}$$

Given this demand, for a given quality, the publisher maximizes the following objective function:

$$\max_A p_a A D(q, A) - cq^2,$$

where p_a is the price of an ad impression. We consider the two relevant cases with positive revenue:

Case 1: $A \le q$. The publisher maximizes

$$p_a A\left(1 - \frac{A}{2q}(1 + r)\right) - cq^2 \Rightarrow \max_A \quad \text{subject to } A \le q.$$

This is a concave second degree equation in A. The first-order condition provides the following solution for A as $A^\star = \frac{q}{1+r}$ (the constraint is always satisfied):

Case 2: $rA \le q < A$. Here, the publisher maximizes

$$p_a A\frac{(q - Ar)^2}{2Aq(1 - r)} - cq^2 \Rightarrow \max_A \quad s.t.\ q < A \le \frac{q}{r}.$$

The derivative of the objective function is always negative and implies that the optimal choice of advertising is $A_2^\star = q$ and, hence, Case 2 is never relevant. Conducting the same analysis for $0 < t \le 1$ reveals the general solution for the optimal advertising level as:

$$A^\star = \begin{cases} \dfrac{q}{(1 + r)} & \text{if } 0 \le t \le \dfrac{r}{1 + r} \\ q(1 - r(1 - t)) & \text{if } \dfrac{r}{1 + r} \le t \le 1. \end{cases} \tag{2}$$

Solving for the optimal quality provides the equilibrium that we summarize next.

Lemma 1. *Without the presence of the ad blocker, the publisher chooses the following quality and level of advertising:*

$$q^\star = q_{NA}^\star = \begin{cases} \dfrac{p_a}{4c(1 + r)(1 - t)} \\ \dfrac{p_a(1 - r(1 - t) + t)}{4c} \end{cases} \text{ and}$$

$$A^\star = \begin{cases} \dfrac{q^\star}{(1 + r)} & \text{if } 0 \le t \le \dfrac{r}{1 + r} \\ q^\star(1 - r(1 - t)) & \text{if } \dfrac{r}{1 + r} \le t \le 1. \end{cases}$$

Both quality and advertising decrease in r and c and increase in t and p_a.

As expected, a higher cost reduces the quality investment, whereas a higher price at which ads can be sold increases it. Consumer preferences also have an impact. Equilibrium quality is higher when consumers value quality more or are less sensitive to advertising.

3.2. Analysis with Ad Blocker

In the presence of the ad blocker there are three new stages added to the game. The first is when the ad blocker sets the maximum advertising it lets through, V. The second is when consumers make a decision to download the ad blocker at cost κ, and finally, when the ad blocker and the publisher negotiate.

Solving the game entails backward induction. Details of the proofs are relegated to the appendix; here we walk through the main steps and illustrate the intuition for $t = 0$. The last stage of setting the advertising is similar to the benchmark case and, given our assumption that $\bar{\kappa}$ is large enough we get that $A^* = \frac{q}{1+r}$ when t is small. The stage before that is the negotiation phase. The ad blocker's main leverage here is the set of consumers who downloaded its app. The app limits advertising to a level of V for these consumers, but the ad blocker can reduce that to 0. Note that a rejected offer by the publisher makes the ad blocker indifferent in setting $\mathcal{V}$ anywhere within the $[0, V]$ interval. This serves as a threat and allows the ad blocker to extract some of the publisher's revenue.

Given the base model's assumption that the ad blocker starts the negotiation with a take-it-or-leave-it offer and the publisher accepts if indifferent, the ad blocker is able to extort a payment T equal to the entire advertising revenue the publisher generates from consumers who downloaded the ad blocker by threatening to block all ads. Hence, the size of the segment that downloads the ad blocker becomes crucial and therein lies the key tradeoff: as the ad blocker reduces V, more consumers download the ad blocker, but at the same time its revenue, T, decreases as the amount of ads V decreases.

It is important to realize that consumers are rational and they expect a deal between the ad blocker and the publisher. Hence when they contemplate downloading the app, they know that it does not provide full ad blocking, rather it limits advertising to V.[12] Figure 2 illustrates the download decision for consumers with a given κ cost when $t = 0$. There are two constraints of interest. The vertical line $\gamma = \frac{\kappa(1+r)}{q-V(1+r)}$ represents the consumers who are indifferent between visiting the publisher's site while being exposed to advertising A^* or downloading the ad blocker, and being exposed to only V level of advertising. Hence, consumers to the right of this line are potential downloaders, whereas consumers to the left never download. But we also have to make sure that consumers who download the ad blocker get positive utility (i.e., they want to visit the site with V level of advertising). The indifference curve is given by the $\vartheta = \frac{\kappa+\gamma V}{q}$ line. Consumers below this line never download, whereas consumers above may. Hence, exactly consumers in the (abhg) trapezoid download (and use) the ad blocker. As the (cg) line represents the threshold for consumers to visit the publisher's site in the absence of an ad blocker, it is clear that consumers in the (cdhg) trapezoid continue to visit the site without the ad blocker. It is interesting to see that there is a segment of consumers (afg) who do not visit the site without the ad blocker but do visit with it.

Figure 2 illustrates the tradeoff clearly. As V decreases, the (abhg) trapezoid grows, but the amount of rent to be extracted decreases, as it is exactly V for each consumer in this segment. The ad blocker hence needs to find the optimal V that balances these two forces. The figure shows our reasoning for a given κ, but to find the optimal V, we need to integrate over all possible values of κ and solve

$$\max_V T(q, V) = \max_V V \cdot \int_0^{\bar{\kappa}} S_V(\kappa, q, V) d\kappa,$$

where $S_V(\kappa, q, V)$ is the area of the segment that downloads the ad blocker. Deriving the integral in the proof yields the following proposition.

Proposition 1. *For given quality, $V^\star = qF(r,t) \le A^\star/2$, where $F()$ is defined in the appendix. The function $F(r, t)$ is decreasing in r and increasing in t. Furthermore, $F(r,0) \xrightarrow{r\to 0} 1 - \frac{1}{\sqrt{3}}$ and $F(r,1) = \frac{1}{2}$ for any $r > 0$.*

This result captures the key idea behind how the ad blocking ecosystem works. The ad blocker makes money by posing a threat to the publisher: "pay us or we will completely block advertising." Since ad blockers typically do not charge consumers, this is their main source of revenue. In order for the threat to be real, the ad blocker needs downloads. It entices consumers to download the app by filtering out at least some ads, that is, limiting the total amount of disutility from advertising. The consumer segment that downloads consists of people who tend to be sensitive to advertising, appreciate quality, and have a low downloading cost.

It is important to point out that in equilibrium, the ad blocker does not completely block ads. This is consistent with the practice of whitelisting that we observe in reality. Most ad blockers let "good" ads through and only block the intrusive ones, essentially limiting the amount of advertising. As our model demonstrates, whitelisting is a central element of the ad blocking business model, in which the ad blocker extracts rent from publishers.

Despite these whitelisting practices, the ad blocker still generates value for consumers. The maximum advertising is always less than half of what consumers without an ad blocker encounter. The proposition also shows that consumer heterogeneity matters. The less heterogeneity there is in sensitivity to ads (increasing r), the lower $V^\star/A^\star$ will fall. On the contrary, if consumers are less heterogeneous with respect to their valuation of quality (increasing t), the more ads they will see. At the extreme, when all consumers have the same valuation for quality, exactly half of the ads are blocked.[13]

Overall, if publisher quality is fixed, the presence of the ad blocker clearly benefits consumers.[14] Not only by making it less painful for existing visitors to consume the publisher's content, but by making it worthwhile for new visitors to come to the publisher's site. These latter are consumers who would not visit the site without the ad blocker.

This brings us to the very first stage of the game. Up until now, we have treated the publisher's quality as given. However, many argue (Shiller et al. 2018) that ad blockers endanger the entire advertising-supported publishing business model. If the ad blocker can extort publishers and extract advertising revenues, the publisher will have no incentive to invest in content quality, which may ultimately hurt consumers. Our next result speaks to this issue.

Figure 2. (Color online) Download Decision of Consumers with Cost κ in the (γ, ϑ) Space When the Ad Blocker Sets Maximum Ad Volume to V Assuming $t = 0$

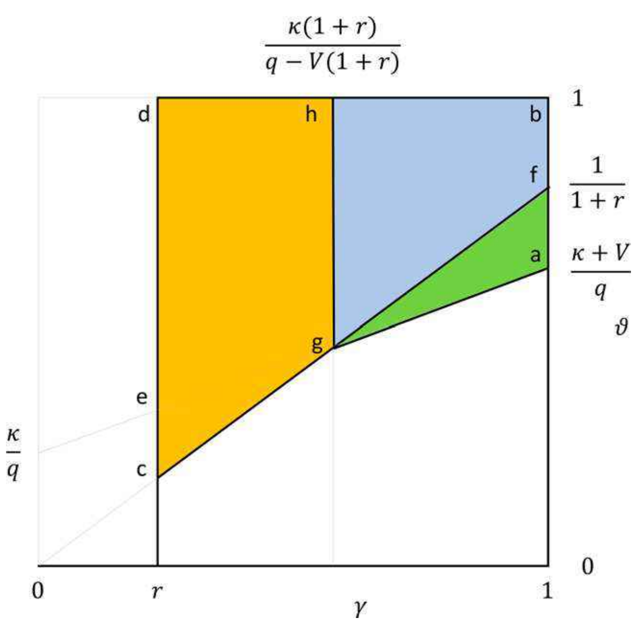

Proposition 2. *There exists $G(r,t) > 0$, such that the optimal quality is*

$$q^\star = q^\star_{AB} = \begin{cases} \left(\frac{4c(1+r)(1-t)}{p_a} + \frac{G(r,t)}{\bar{\kappa}}\right)^{-1} & \text{if } 0 \le t \le \dfrac{r}{1+r} \\ \left(\frac{4c}{p_a(1-r(1-t)+t)} + \frac{G(r,t)}{\bar{\kappa}}\right)^{-1} & \text{if } \dfrac{r}{1+r} \le t \le 1. \end{cases}$$

We have $q^\star_{AB} < q^\star_{NA}$ and $q^\star_{AB}$ is increasing in $\bar{\kappa}$ and converges to $q^\star_{NA}$ as $\bar{\kappa} \to \infty$. The publisher's profit is also lower than in the benchmark case and increases with $\bar{\kappa}$.

The results confirm our intuition that quality decreases in the presence of the ad blocker. This is driven by the reduction in advertising revenue the publisher can keep. At the extreme, if $\bar{\kappa}$ were zero, that is, all consumers could download the ad blocker without any nuisance cost, the publisher would lose all incentive to invest in quality and the market would collapse. What protects the publisher's revenues is that downloading the ad blocker is costly. The more consumers with high download costs, the less the publisher has to worry. To better illustrate how the download decision is made by different consumers, we next examine the size of the segment that downloads.

Corollary 1. *The size of the consumer segment that downloads the ad blocking software is increasing in the ad price, p_a, and decreasing in the download cost, $\bar{\kappa}$, converging to zero as $\bar{\kappa} \to \infty$.*

The corollary is in line with some of our main intuition. A higher ad price leads to a higher quality level, which increases advertising levels. However, the gap between the advertising levels with and without the ad blocker also increases. As a result, the benefit from downloading the ad blocker goes up with quality if the download costs are fixed. Conversely, when download costs increase, fewer consumers download the ad blocker. This protects the publisher from the ad blocker's extortion threat yielding higher quality.

Clearly, the ad blocker has an unambiguous negative effect on the publisher by extracting rent through the downloader segment. The effect on consumers is less clear. On one hand, content quality becomes lower in the presence of the ad blocker, especially when the download costs are not very high. However, the consumers who downloaded the ad blocker see significantly less advertising. In fact, we can show that there are segments of consumers whose net benefit from the presence of the ad blocker is positive.

Corollary 2. *Unless $r = t = 1$, there are consumers who are strictly worse off in the presence of an ad blocker than in the benchmark case with no ad blocker. There always are consumers who are better off. In particular, there is a segment of consumers that only visits the publisher's site and derive positive utility in the presence of the ad blocker.*

This latter segment consists of consumers who are sensitive to advertising, but value quality moderately. They only visit the publisher's site if the ad blocker is available. Importantly, this segment would not consume any content and therefore would obtain zero utility without the ad blocker. In the presence of the ad blocker, they get positive utility even though quality decreases. Note that this particular segment is only part of the consumers who are better off with an ad blocker. Consumers with very low download cost and high ad sensitivity are also better off with the ad blocker, especially if most other consumers have high download costs and, thus, quality levels are not reduced.

Overall, ad blocking has two opposing effects on consumers. The direct effect increases consumer utility by removing some ads and the associated disutility. The indirect, quality reduction effect reduces the utility consumers obtain from visiting the site, which is proportional to the content quality. Depending on the heterogeneity in the consumer population, either effect can dominate and total consumer surplus may increase or decrease. The following proposition compares total consumer surplus in the benchmark case with no ad blocking (CS_{NA}) with that in the presence of an ad blocker (CS_{AB}).

Proposition 3. *If either r or t is sufficiently low, then $CS_{AB} < CS_{NA}$. If both r and t are sufficiently high, then $CS_{AB} > CS_{NA}$ as long as κ is high enough.*

The main determinant of whether total consumer surplus declines or increases with the entry of the ad blocker is the extent of consumer heterogeneity. The reason is that the impact of the quality reduction effect on consumers is determined by two factors: the extent of the quality reduction and the surplus consumers can extract relative to quality. The latter is strongly dependent on consumer heterogeneity. Consider the case when both r and t are close to 1. In this case, the publisher is able to extract all surplus from consumers by setting a high advertising level. Hence, in the case of no ad blocking, consumers walk away with almost no surplus. The entry of the ad blocker changes the equation, because those with low download cost benefit end up with a positive surplus. Although the ad blocker's entry does reduce quality levels, this does not hurt consumers as their surplus was close to zero to start with. Hence total consumer surplus increases.[15]

On the other hand, when consumer heterogeneity is sufficiently high (either r or t is low enough), consumers enjoy a good positive surplus in the absence of the ad blocker. With its entry, consumers are hurt, because the extent to which the ad blocker's rent-seeking behavior reduces quality investment and content production counters its beneficial impact on advertising disutility.

In our setting, consumer heterogeneity has to be fairly limited for the positive effects to dominate. Figure 3 depicts the region in the (r, t) parameter space where consumer surplus is higher than without an ad blocker. The thresholds are close to 1. When $r = 1$, total consumer surplus increases if and only if $t > 4/5$ as long as κ is high enough. When $t = 1$, the threshold for r is $r > (\sqrt{141} - 1)/14$.

We next investigate how the total welfare of consumers, the publisher, and the ad blocker (TW_{AB}) compares to the total welfare in the benchmark, where TW_{NA} is the sum of consumer and publisher surplus.

Corollary 3. *We have $TW_{AB} < TW_{NA}$, unless the conditions for $CS_{AB} > CS_{NA}$ hold and both p_a and c are sufficiently low.*

The results show that total welfare can only increase if consumer surplus increases. The reason is that the entry of the ad blocker lowers the value that the publisher generates. Since $q^{AB} < q^{NA}$, the total surplus divided between the ad blocker and the publisher is lower than that of the publisher in the benchmark case. Only consumer surplus can flip the inequality. When consumer surplus is higher with the ad blocker than in the benchmark case, total welfare can also go up, but only if the increase in consumer surplus is high enough compared with the decrease due to lower quality. For the latter reduction to be low, both p_a and c have to be low. The price p_a has to be low so that the rate at which advertising space is converted into revenue is low relative to the consumer surplus, and c has to be low so that quality stays generally high enough for the consumer surplus increase to be substantial. Overall, we see that total welfare decreases with the entry of the ad blocker unless a number of (quite demanding) conditions align.

3.3. Full Ad Blocking as an Option

In this section, we examine the case where, besides opting for partial ad blocking, consumers have the option to use full ad blocking. This is a relevant case in practice, because in mainstream ad blocking software, though the default option is partial ad blocking, a full ad blocking option is also available if the consumer is willing to incur some extra cost.[16]

To study this scenario, we modify our basic model as follows. When consumers make the download decision, they now have three options: (i) not downloading the ad blocker, (ii) downloading the ad blocker with some ads still displayed (V) at cost κ as before, and (iii) downloading the ad blocker and using it with full ad blocking at cost $\kappa(1 + \mu)$, where $\mu \geq 0$ is a fixed parameter. For simplicity, we also set $t = 0$. Note that this is a conservative assumption as total consumer surplus is lower with the ad blocker for this parameter range. We investigate whether adding the full ad blocking option changes that. All other aspects of the model are the same as in our basic model.

Figure 3. Parameter Values in the (r, t) Space Where Total Consumer Surplus Is Higher with the Ad Blocker (Shaded Area) Than in the Benchmark Case as Long as κ Is High Enough

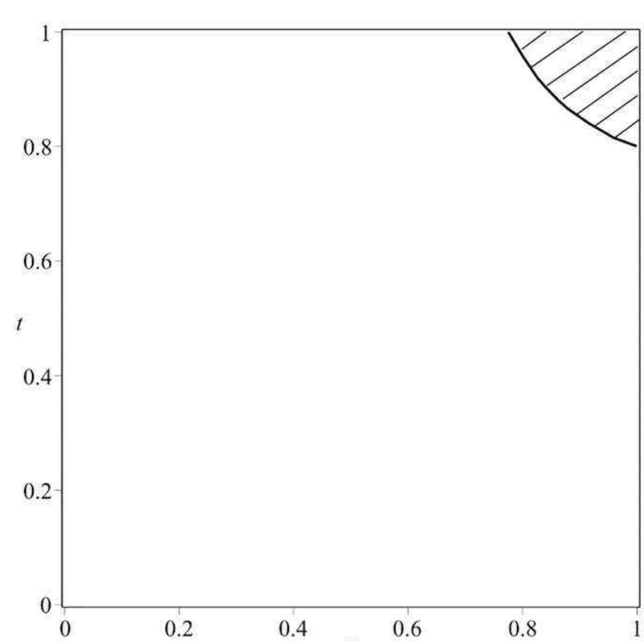

The main difference in the analysis surfaces when we examine the decision of the ad blocker to set V. As before, V has to be low enough to entice consumers to download the ad blocker, but at the same time high enough to create advertising revenue for extortion potential. The novel force in our modified model is that there is a full ad blocking option which, even though at a higher cost, presents a potentially more attractive option for consumers. Moreover, those who choose the full ad blocking option will not contribute to the advertising revenue and thus will not increase the ad blocker's potential revenue source. Our first result concerns the new equilibrium V.

Lemma 2. *For given quality q, we have $V^\star = qF_\mu(r)$, where $F_\mu(.)$ is defined in the appendix; $F_\mu(r) < F(r,0)$; and $F_\mu(r)$ is decreasing in r and increasing in μ. Furthermore,*

$$F_0(r) = 0, \quad F_\mu(r) \overset{\mu\to\infty}{\longrightarrow} F(r,0) \quad \text{for} \quad \forall\ r < 1;$$
$$F_\mu(r) \overset{r\to 0}{\longrightarrow} \frac{\mu}{1+\mu}\Big(1-\sqrt{1/3}\Big), \quad \text{and} \quad F_\mu(1) = 0.$$

The results are similar to those we obtain from our main model in Proposition 1, but there are a number of notable differences. Whereas V^* is still proportional to quality, q, the ratio V^*/q is smaller. Furthermore, the lower the extra cost of using the ad blocker to completely dispose of ads, μ, the lower the multiplier: as it becomes easier for users to get rid of all ads, the more the ad blocker is forced to let through only a small amount of ads. At the other extreme, when the extra cost of full ad blocking approaches infinity, we get a V that is identical to our basic results as long as $r < 1$. Figure 4 illustrates the consumer choice for a given κ. As before, consumers with low ad sensitivity (cdhg trapezoid) do not download the ad block app. The difference is that those who download it are split into two segments. Moderately ad-sensitive consumers in the (ghik) trapezoid only incur the κ cost and download it for partial ad blocking. Highly ad-sensitive consumers in the (kibj) rectangle incur the extra cost and use it for full ad blocking. This latter segment is the one that forces the ad blocker to lower V further than before, especially if μ is small. Interestingly, when r approaches 1 the ad blocker finds itself in a tough spot as it does not have enough heterogeneity in ad sensitivity to differentiate the partial ad blocking option from the full ad blocking option. Therefore, it is forced to set a low V reaching $V = 0$ as r approaches 1 (i.e., when there is no consumer heterogeneity in advertising disutility). In sum, the extortion-based business model is no longer feasible without heterogeneity in ad sensitivity and the ad blocker faces similar problems with extracting revenue as in Section 3.4 even if μ is high. Overall, we see that the possibility of full ad blocking reduces the advertising level that the ad blocker allows through. We next examine how the publisher's quality choice is affected by μ.

Proposition 4. *The optimal quality is $q^\star_\mu = p_a\left(4c(1+r)+\frac{2p_a(r^2+4r+1)(1-(1+r)F_\mu(r))}{3\bar{\kappa}(1+r)^2}\right)^{-1}$. Quality $q^\star_\mu$ is increasing in μ and $q^\star_\mu \overset{\mu\to\infty}{\longrightarrow} q^\star$. As before, $q^\star_\mu$ is decreasing in r and increasing in $\bar{\kappa}$ and converges to the benchmark level as $\bar{\kappa}\to\infty$.*

Again, the results are similar to those in our main model, but the more stringent blocking of ads has a negative effect on the publisher. With the lower ad levels, the ad blocker is more attractive and more consumers download it, leaving the publisher with lower ad revenues for a given quality level. In turn, the incentives to invest in quality are further diminished, resulting in a lower equilibrium quality. The lower the cost of full ad blocking, the lower the quality.

Corollary 4. *For $t = 0$, total consumer surplus and total welfare is always lower with the full ad blocking option than in the benchmark with no ad blocker for any μ value.*

Even though some of the very ad-sensitive consumers benefit from full ad blocking, overall consumer surplus is lower than without an ad blocker for $t = 0$, just as in Proposition 3. No matter how easy it is to use the ad blocker to fully block ads (even as μ becomes 0), consumers overall will always be worse off than without an ad blocker. The reduction in quality is just too severe and the gains from blocked ads cannot counteract the disutility resulting from lower quality. Although we restricted this section to $t = 0$ for parsimony, it is instructive to check what happens when $t = 1$. Analyzing this case shows that a lower μ does, in fact, improve overall consumer surplus. There is a critical r above which consumer surplus is higher than in the benchmark case. This result parallels our previous finding that in the absence of consumer heterogeneity, ad blocking increases consumer surplus and overall welfare. In Proposition 3, $r > (\sqrt{141}-1)/14 \approx 0.78$ ensured that consumers are better off. Here, as μ approaches 0, the threshold changes to $r > (\sqrt{33}-1)/8 \approx 0.59$. Overall, in this case, a lower μ does benefit consumers, but there is a wide range of r, t parameter values for which consumers are worse off in the presence of an ad blocker even if $\mu = 0$.

3.4. Evaluating the Ad Blocking Business Model

The original benchmark we established was a world without ad blocking. We showed how the entry of an ad blocker reduces quality and generally decreases consumer surplus and welfare in the practically relevant scenario when consumers exhibit some level of heterogeneity. The main reason is that the ad blocker extorts too much surplus from the publisher, thereby strongly disincentivizing content quality investment. Consumers do not benefit much from the reduced advertising as the ad blocker lets through a fair amount of ads to be able to put pressure on the publisher by wielding the threat of completely blocking ads. This equilibrium allows for the market to maintain a certain level of content and advertising, but reduces overall welfare and consumer surplus. The situation is even worse if ad blockers also offer a full ad blocking option, as is often the case in practice.

There is an important regulatory debate whether this type of business model, often labeled extortion, should be allowed or not. In some European countries, these questions reached the supreme courts and in most cases the practice has been deemed legal.[17] Our model allows us to speak to this issue, by comparing the practice to a world where ad blocking exists, but the ad blocker is not allowed to extract a fee from the publisher, who is hurt by ad blocking in the first place. The logical and more conventional alternative revenue model consists of the ad blocker charging consumers, for whom ad blocking creates value. We have mentioned before that ad blockers may shy away from charging consumers for practical reasons (i.e., the difficulty of collecting revenue from a large fragmented customer base with a generally negative attitude toward paying for digital products). Our analysis that follows shows that the reason behind ad blockers' choice of an unconventional business model could be more fundamental, driven by the potential surplus they generate this way.

In implementing the ad blockers' new revenue model, we need to slightly alter the timing of the game. As before, the first stage is comprised of the publisher setting its quality q, followed by the ad blocker setting V. At this point we add a step where the ad blocker sets its price, p_b for consumers. Upon observing both V and p_b, consumers decide whether to download the ad blocker app. We remove the stage where the ad blocker asks the publisher for a payment, and the last stage is where the publisher sets its advertising level, consumers make their choice to visit the site or not, and profits are realized.

Next we compare this charge consumers (CC) case to our full model with the ad blocker charging the publisher as is consistent with current practice (we continue to denote this case by AB) and our original benchmark of no ad blocking (NA). For simplicity, we focus on two extreme cases that exhibit starkly different overall outcomes in the main model where the ad blocker extorts the publisher. The first case represents the lower left corner of Figure 3, where both r and t are small and, therefore, consumer heterogeneity is maximal. This case is closest in relevance for practice. We also examine another case represented by the upper right corner of the figure where both r and t are high and thus heterogeneity is very low.

Proposition 5. *When r and t are sufficiently small, the equilibrium outcome in the case of charging consumers falls between the benchmark with no ad blocker and the main model of charging publishers with respect to quality, publisher payoff, consumer surplus, and total welfare: $q^\star_{AB} < q^\star_{CC} < q^\star_{NA}$, $\Pi^{Pub\star}_{AB} < \Pi^{Pub\star}_{CC} < \Pi^{Pub\star}_{NA}$, $CS^\star_{AB} < CS^\star_{CC} < CS^\star_{NA}$, and $TW^\star_{AB} < TW^\star_{CC} < TW^\star_{NA}$. Furthermore, the ad blocker has lower profits when charging consumers: $\Pi^{Adb\star}_{AB} > \Pi^{Adb\star}_{CC}$.*

When $r = t = 1$, we have $q^\star_{CC} = q^\star_{AB} < q^\star_{NA}$, $\Pi^{Pub\star}_{CC} = \Pi^{Pub\star}_{AB} < \Pi^{Pub\star}_{NA}$, $TW^\star_{CC} = TW^\star_{AB} < TW^\star_{NA}$, and $CS^\star_{CC} = CS^\star_{AB} > CS^\star_{NA}$.

The results paint an interesting picture of how the business model affects profits and welfare. When consumer heterogeneity is high, surprisingly, both the publisher and consumers are better off when the ad blocker charges the consumers compared with when it charges the publisher. Furthermore, total welfare is also higher in this case, but still lower than without the presence of the ad blocker. Hence, ad blocking is still harmful, but less so than in the case when the ad blocker extorts the publisher. At the same time, the ad blocker is worse off with the traditional business model of charging consumers. This analysis suggests that it is not a surprise that we do not see many ad blockers employing such a business model in reality.

Before we describe the intuition behind the results, we note that although the decision to set $V > 0$ is still an option when charging consumers, the ad blocker does not exercise this option for trivial reasons. The more ads the software blocks, the more valuable it is allowing the ad blocker to collect more revenue. Hence, in equilibrium V is always 0 when the ad blocker charges consumers. Instead of adjusting V, the ad blocker uses the price p_b to balance the revenue and demand for its software. Setting a higher p_b is similar to setting a higher V as both decrease the amount of consumers who download the ad blocker, but both increase the revenue that can be extracted from each consumer. Here, the revenue is extracted directly through a price, whereas in our main model, the revenue for each downloader is extracted indirectly by collecting it from the publisher.

Charging a single price p_b allows for less discrimination than setting V, as the latter affects consumers with different ad sensitivities to a different degree. As a consequence, charging consumers directly will lead to lower overall demand for the software and less revenue for the ad blocker. On the flip side, fewer downloaders benefit the publisher as it can collect more ad

Figure 4. (Color online) Download Decision of Consumers in the (γ, ϑ) Space When the Ad Blocker Sets Maximum Ad Volume to V

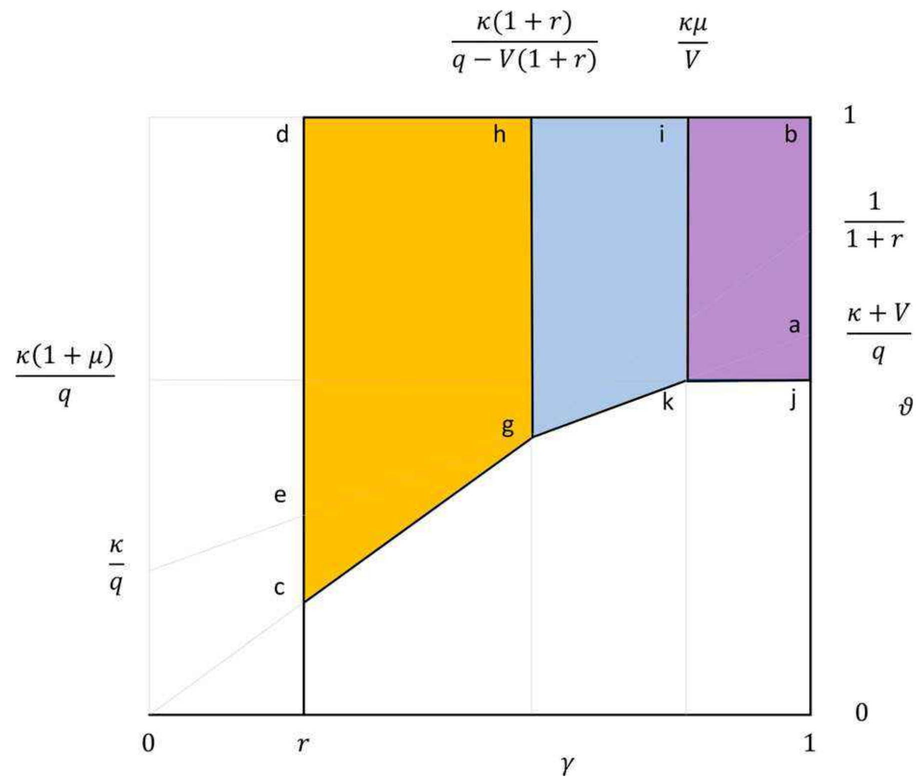

Note. The cost of downloading and using the ad blocker with a V level of advertising is κ, whereas using it for full blocking costs an additional $\mu\kappa$.

revenue, which can then support a higher level of quality investment. Higher qualities lead to a higher total welfare.

Surprisingly, consumers are better off when they are charged directly. This happens for two reasons. First, a higher quality always benefits consumers. The other reason is that the ad blocker is able to extract less surplus from consumers when charging them directly. Even though no surplus is extracted directly from consumers when the ad blocker charges the publishers, a positive V causes consumers to lose utility by having to endure some advertising. When they are charged for the ad blocker directly, they do pay a price, but at least there is full ad blocking.

When consumer heterogeneity is very low, the dynamics are different. In the extreme case of $r = t = 1$, when consumers are only heterogeneous with respect to κ, charging the publisher and consumers results in very similar outcomes. The optimal p_b is $q/2$, which closely corresponds to a $V = q/2$ in our main model. Hence, the direct price the consumers pay for purchasing the software is the same as the indirect price they pay in our main model, by enduring some advertising. As a result, the segment of downloaders is exactly the same in both cases and hence the publisher's incentives are also the same in both cases, leading to the same quality, publisher profit, consumer welfare, and total welfare.

Overall, we see that charging consumers makes a difference when consumer heterogeneity is significant. Such an ad blocking business model is less harmful to the ad-supported publishing industry and, importantly, it benefits consumers as well. However, ad blockers typically do not want to use such a model as it hinders their ability to extract surplus and leads to lower revenues. Although in practice there may be a variety of other factors that make it difficult to collect revenues from consumers, it is clear that ad blockers have no incentive to do so even if these factors were mute.

4. Publisher-Ad Blocker Relations

So far, we have considered all negotiation power to be in the hands of the ad blocker who presents a take-it-or-leave-it offer to the publisher. In practice, this assumption may apply to the case of small, powerless publishers. However, it is reasonable to assume that larger publishers have more negotiation power. The most obvious way a publisher can avoid ad blockers is by erecting a paywall around its content; we formally explore this case in a web appendix to show that, indeed, if the publisher can switch to a subscription model, it is partially shielded from the ad blocker's extorsion. But publishers may have many other (costly) means to avoid ad blockers. For example, one option that often emerges in discussions about ad blocking is technology that is able to prevent ad blocking (so-called "ad blocker blockers").[18]

Although ad blockers' business model relies on extracting rent from publishers, there is no institutionalized system to determine and collect these payments. Hence, bilateral negotiations between the ad blocker and the publishers and certain actions that each side can take play a crucial role. In this section, we explore this important issue.

Mathematically, we operationalize the negotiating power of a publisher as the share of the surplus from consumers who downloaded the ad blocking app that the publisher can keep. We define the publisher's share as $1-\alpha$, with the ad blocker's share being α. We first study the case of an exogenous α. For simplicity, we assume $t=0$ throughout this section.

Several aspects of the analysis remain the same as in the base model. For a fixed V and q, the publisher chooses $A=\frac{q}{1+r}$ given our assumption that $\bar{\kappa}$ is high enough. The ad blocker's problem of setting V also remains the same as its payoff is simply multiplied by α, but otherwise remains unchanged. Hence $V^\star(\alpha)=V^\star=qF(r,0)$ remains the same. However, the publisher's quality choice problem is different since the payoff function now includes some portion of the ad revenue from the downloader segment as well. The analysis provides the following result.

Proposition 6. *The optimal quality is*

$$q^\star(\alpha)=p_a\left(4c(1+r)+\frac{\begin{array}{c}2p_a[(r^2+4r+1)(1-(1+r)F(r))\\-(1-\alpha)H(r)]\end{array}}{3\,\bar{\kappa}(1+r)^2}\right)^{-1},$$

where $H(r)>0$ only depends on r. Both $q^\star(\alpha)$ and the publisher's profit are decreasing in α, r and increasing in $\bar{\kappa}$.

Overall, from the publisher's perspective, having more negotiation power acts similarly to having a lower cost. The ad blocker sets V at the same percentage of A regardless of α giving the same incentive to download the ad blocker, resulting in the exact same downloader population if quality were fixed. However, being able to collect and keep some portion of the ad revenue from the downloaders is a boost and results in a higher quality investment, which also benefits consumers.

Corollary 5. *Consumer surplus is decreasing in α. When $t=0$ and $\bar{\kappa}$ is sufficiently high, then $CS^\star(\alpha)>CS^\star_{NA}$ if and only if both r and α are sufficiently small.*

Since consumer behavior is only affected by α through the increased q, consumers are actually better off with a stronger publisher and weaker ad blocker. So much so that when α becomes small enough, consumer surplus can be higher than in the benchmark with no ad blocker. However, this can only happen when r is small enough. Numerically, the maximum thresholds for consumers to be better off are at about 0.12 for α and 0.009 for r. If either r or α are larger than these thresholds then consumers are worse off than without the ad blocker.

Note that there is a slight difference in the outcome when $t=1$. Consumers are still better off as α decreases, but the pattern reverses in r. Consumers are better off than in the benchmark case if and only if r is large. Recall that in Proposition 3, $r>(\sqrt{141}-1)/14\approx 0.78$ ensured that consumers are better off. Here, as α approaches 0, the threshold changes to $r>(\sqrt{33}-1)/\,8\approx 0.59$. Interestingly, this is exactly the same threshold as what we find in Section 3.3 with the full ad blocking option as μ approaches 0. Hence, when $t=1$, the outcome is equally good for consumers if they can easily use the ad blocker for full ad blocking or if there is limited ad blocking, but the publisher keeps all the revenue from the ad blocker. Note, however, that the quality is higher in the latter case.

Overall, lowering α is beneficial to both the publisher and the consumers. However, for a wide range of parameters, consumers are still worse off than in the benchmark case with no ad blocker. The reason is that the quality reduction effect dominates: even though the ad blocker does not capture much surplus because of its low negotiation power, ad blocking can simply act as a suboptimal limit on ads (from both the publisher's and the consumers' perspectives), which reduces quality. This is similar to what happens when a platform limits advertising for publishers on its site, but there, the limit will be higher because the platform maximizes the total surplus in the entire ecosystem as opposed to only the ad blocker's profit.[19]

4.1. Ad Blocker Restraint

Thus far, we have treated α as an exogenous variable. It is likely that both the publisher and the ad blocker can take certain actions to influence α and prepare for the negotiation. It is plausible that both sides can invest in certain costly actions to tilt the scale, such as legal representation, better outside options, and technology. As we have seen in Proposition 6, the publisher is certainly better off with a lower α and would potentially spend money to make it lower. Does that mean that the ad blocker is always interested in the opposite, that is, a higher α? Here, we examine this question by assuming that the ad blocker can freely choose an α at the beginning of the game, but it has to (and can) commit to that α throughout the game.

Proposition 7. *The publisher chooses an optimal $\alpha^* > 0$, which is increasing in c, $\bar{\kappa}$ and decreasing in p_a. The optimal $\alpha^* < 1$ if and only if $\bar{\kappa}c/p_a$ is sufficiently low.*

Thus, we get the interesting result that the ad blocker wants to restrain itself from extracting all surplus from the publisher. The intuition is that even though a better negotiation power allows it to get a higher share of the pie, the pie itself decreases substantially due to the reduced incentive of the publisher to invest in quality. The ad blocker's payoff turns out to be proportional to $q^2\alpha$. Clearly, setting a very low α cannot be optimal as the ad blocker would get only a small share. But increasing α is only beneficial up to a point because quality decreases. Whether that point is more or less than $\alpha = 1$ depends on how sensitive the publisher is to α. The sensitivity depends on a combination of three variables: the download cost, the quality cost, and the ad price. An increase in the first two make the publisher more sensitive, whereas an increase in ad prices makes it less sensitive. Hence, in a situation where either costs are relatively low, or when ad prices are high, the optimal α is less than one and the ad blocker wants to show restraint.

Whether the ad blocker can commit to an upper limit on α depends on the circumstances. A plausible scenario is when both the ad blocker and the publisher can invest in technology that ex post increases their relative negotiation power. In such a game (that is, unfortunately, intractable), it is easier for the ad blocker to commit by not investing, for example. Similarly, the ad blocker can voluntarily subject itself to a regulator, who can enforce ex post its commitment to an $\alpha < 1$.

5. Model Extensions and Robustness Checks

In this section, we examine extensions to our base setup to speak to relevant market contexts and address a few limiting assumptions. First, we consider multiple publishers, which may compete to a varying degree. Next, we explore the increasingly common situation where publishers and consumers interact on a platform, which has broad powers to regulate (ban or promote, even provide) ad blocking. Finally, we discuss what happens if we relax various model assumptions. The detailed analyses and results of these model extensions are in the web appendix. Here, we only discuss the key findings.

5.1. Multiple Publishers

We study two publishers that may have different costs of investing in quality. Importantly, we assume that the ad blocker cannot discriminate between publishers and sets a common V for both, which is consistent with practice and is also needed if the ad blocker wants consumers to develop rational expectations about the proportion of ads blocked when they download the ad blocker.[20] To vary the level of competition, we first consider two noncompeting publishers with overlapping consumers, that is, consumers who visit either or both publishers. Then, we consider competition between two publishers where all consumers consider both publishers but visit at most one of them. For parsimony, we assume $t = 0$ throughout this section.

5.1.1. Noncompeting Publishers. Consumers visit a particular publisher if it provides positive utility. There is no publisher competition because consumers do not compare the utility they get from the two publishers. Publishers may have different costs of quality and, as a result, will choose different quality levels. The rest of the model setup is identical to the basic model. All actions that publishers take are simultaneous with each other.

Even without competition, there are two sources of strategic interaction between the publishers. First, the ad blocker sets a common V, which is a function of both publisher qualities. Second, consumers in their overlapping consumer segment make the ad blocker download decision considering both publishers.

We find that the level of asymmetry between the players is important in explaining the strategic interaction between them and the qualities they set. Regardless of whether this asymmetry is high or low, the presence of another player has a negative effect on a publisher, but the mechanisms are slightly different. In the highly asymmetric case, although the ad blocker focuses on the high-quality player, it also extracts rent from the low-quality publisher, reducing V compared with the basic model with just the high-quality player. This suppresses the high-quality as a lower V generally leads to more downloads of the ad blocker and generally hurts publishers. The low-quality player is also hurt. Even though V is higher than it would be with only the low-quality player in the market, V is low enough that it encourages consumers to download the ad blocker with the purpose of using it for the high-quality site as well. As a result, a larger number of consumers download the ad blocker than if the low-quality player were alone, which allows the ad blocker to extract more surplus from even the low-quality player than if it were alone.

In the low asymmetry case, consumers use the ad blocker for both sites and consequently the ad blocker uses roughly the average quality to set V. At the extreme, if the two qualities are equal, the ad blocker will react similarly as if there were a single player. But an important difference from the basic model is that the consumer benefit from a single download doubles: for the same download cost κ, a consumer gets ad blocking on two sites. The larger downloading segment suppresses publisher profits and leads to lower quality investments.

5.1.2. Competing Publishers. When publishers are substitutes, each consumer has capacity to visit at most one publisher's site, choosing the one that provides higher utility. To reach an analytical solution, we use the convention of the literature (Shaked and Sutton 1982) that quality investment is costless with a maximum quality of 1. Thus, here, there is no inherent heterogeneity in quality between publishers, although in equilibrium, quality choices differ. This feature allows us to isolate the strategic effect that the ad blocker might have on competing publishers.[21] The rest of the model construction is identical to our previous setup.

In this case, a key insight is that the ad blocker's V falls between the two firms' advertising levels. This is because the ad blocker intends to undercut the ad level on the high-quality site, but not on the low-quality site. As before, this hurts the high-quality site, because advertising is reduced and the ad blocker is able to extract more revenue. But interestingly, it also hurts the low-quality site because a segment of consumers switches to the high-quality site. These consumers previously could not afford the high-quality site because of the higher advertising, but now with the ads limited by the ad blocker, they can enjoy the higher quality content. Furthermore, similarly to our basic model, there are some consumers who did not visit any site before, but now they do and they immediately jump to the high-quality content.

The second point brings forward an interesting strategic effect. Given our simplified cost structure, the higher quality publisher always chooses the maximum quality of 1 to differentiate. The interesting question is how the lower quality changes with the ad blocker. We find that it increases, therefore leading to less differentiation. The intuition relies on the notion that ad blocking makes consumers switch away from the low-quality site. The overall magnitude of switching and its impact on revenue depends on the quality differential. The biggest loss from switching happens at a quality level that is lower than the equilibrium quality without an ad blocker. The impact diminishes as the lower quality approaches the higher quality, because the overall pool of consumers that could switch decreases. Combining this effect with the basic forces that drive differentiation, we get that there is still differentiation, but less than without the ad blocker.

In this case, the effect of the ad blocker on consumers is mostly positive. One reason is that without quality costs, we do not observe a reduction in quality. In fact, with less differentiation the lower quality increases. Furthermore, ad blocking directly increases consumer utility and as we noted, there is a segment of consumers that switches from the low-quality to the high-quality publisher and there is a segment that moves from no consumption to visiting a site. In addition to these segments, the reduced differentiation even benefits consumers who are not using the ad blocker as competition leads to lower ad levels. It is important to note that if we add back quality costs, ad blocking will have a primary effect of suppressing quality investment as in the base model. However, the key finding of this analysis is that competition between publishers may alleviate some of those negative factors.

5.2. Platform Regulating Ad Blocking

In this section, we will consider the case of a platform that does ad blocking itself instead of a third-party ad blocker. This setup is relevant for multiple reasons. First, more and more content consumption happens on platforms where consumers and publishers interact (e.g., iOS, Google, or Facebook). For example, Google's Chrome browser offers to block certain ads. Second, multisided platforms have tremendous power in "regulating" the interactions between the different sides and can potentially internalize competition between the ad blocker and the publisher, leading to better welfare outcomes. Indeed, the platform is interested in maximizing the total surplus generated by the transactions whereas a third-party ad blocker is only interested in maximizing its own profit. In this sense, the incentives of the platform and the ad blocker are not necessarily aligned, leading to different outcomes. This tension shows in practice, where ad blockers are criticizing Google, for example, for not blocking enough ads. At the same time, Google is attempting to prevent ad blocking software from operating properly in its system. As we see a trend of greater platform control, especially in the mobile content ecosystem, it is important to understand how different entities can and should block or limit ads.

To speak to this issue, we analyze a model in the web appendix in which the platform itself does the ad blocking with the objective to maximize total welfare. Interestingly, we find that in such a setup, total surplus is always lower compared with the no ad blocker (benchmark) case. Even though for a fixed value of q, the platform indeed sets V such that the total surplus is higher than in the benchmark, the platform's choice of V lowers the publisher's profit for a given q. This implies that the equilibrium q in the benchmark model is higher than in the model where the platform is the ad blocker. Moreover, the total surplus in the latter case can be sometimes even lower than with a third-party ad blocker. Depending on the value of advertising, an ad blocker may be less harmful to quality levels than the platform itself. Despite the drop in quality, consumer surplus is mostly positively affected. Although the combination of reduced advertising and reduced quality hurts total surplus, the former factor often dominates when considering consumers' payoffs.

5.3. Additional Robustness Checks

We conduct various other robustness checks in the web appendix to address some limitations of our model. We start by alleviating the assumption that the price of advertising is exogenous. Solving the model with a fully endogenous p_a is intractable but we can analyze a quasi-endogenous case, where advertising price is an increasing (concave) function of quality. We find that the structure of the analysis and the results are similar to the main model's.

We also explore the case when the publisher can switch to a subscription revenue model, which essentially puts a lower bound on the publisher's profit. In turn, this provides more negotiation power to the publisher. Although this limits the decline in the quality level, our main findings concerning consumer surplus and welfare remain valid.

Throughout the analysis we consider a monopolist ad blocker. This assumption is driven by the fact that, whereas one can observe a large number of ad blockers, the market is highly concentrated.[22] [22]Indeed, although the cost of entry is relatively low, reliably operating ad blocking requires scale. As such, new entrants typically shut down because they cannot generate revenues. Nevertheless, it is relevant to ask how the presence of competing entrants, often offering full ad blocking, affects our results. Fortunately, the version of our model in which the ad blocker also provides a full ad blocking option (see Section 3.3) is essentially the same as this setup as long as full ad blockers represent higher downloading costs to consumers, which is safe to assume. Furthermore, in the web appendix, we also explore the possibility that two symmetric ad blockers compete. If such competition only concerns the limit of advertising, V, then the dynamics are similar to Bertrand competition. However, if ad blockers compete in V as well as in consumers' downloading cost, a differentiated outcome with positive profits is possible, with one ad blocker offering lower V at a higher downloading cost, while the other does the opposite.

6. Concluding Remarks

We have developed a model to evaluate the impact of ad blocking on publishers' content quality and, in turn, how the decline in quality affects consumers compared with the beneficial effect of reduced ad exposure. Consistent with practice, we assumed that the ad blocker uses the unconventional revenue model, in which it charges the publisher in exchange for allowing a certain amount of ads through its blockade, a practice often qualified as extorsion by publishers. Our analysis reveals that this practice, although beneficial to ad blockers, is typically an inefficient way to limit consumers' advertising exposure that ends up leading to lower overall consumer surplus. By directly charging publishers—whose business is damaged by ad blocking in the first place—the ad blocker forces them to reduce quality investment to such a degree that the resulting lower content quality has more negative effect on consumers than the positive impact of lower advertising exposure due to ad blocking. We show that the conventional revenue model of directly charging consumers would alleviate this problem, although it is not in the interest of ad blockers. This insight may provide some guidance to policymakers who, at present seem to largely support ad blockers' current, so-called whitelisting business practice.

We have explored a number of model extensions to address important institutional details and describe additional business practices. These include the possibility for ad blockers to offer a full ad blocking option, publishers' capability to introduce a subscription revenue model and the possibility of negotiation between the ad blocker and publisher. We have also found that the model is robust to various structural factors, such as the existence of multiple publishers, quasi-endogenous advertising prices, ad blocker competition, and the case where a transaction platform provides the ad blocking service.

The surge of ad blocking in 2016 created a small crisis among publishers, predicting the end of advertising-supported content creation. Today, although ad blocking has stabilized, it still represents a loss of almost 25% of total digital advertising revenues for publishers. Understanding its drivers and its impact on welfare is crucial for the media industry and policymakers. The present paper intends to be a first step in this direction.

Acknowledgments

The authors thank participants at the 2016 and 2019 *Marketing Science* conferences and at UC Davis' marketing seminar for their helpful comments and suggestion.

Appendix

Proof of Lemma 1. The paper already covers the case of $t = 0$. The exact same analysis applies to $0 < t < \frac{r}{1+r}$. When $\frac{r}{1+r} \le t \le 1$, the steps are similar.

Summarizing the three cases, consumer demand faced by the publisher is:

$$D(q, A) = \begin{cases} 1 & \text{if } 0 \le A \le qt \\ 1 - \dfrac{(A - tq)^2}{2Aq(1-r)(1-t)} & \text{if } qt \le A \le q \\ \dfrac{q(1+t) - 2Ar}{2A(1-r)} & \text{if } q \le A \le t\dfrac{q}{r} \\ \dfrac{(q - Ar)^2}{2Aq(1-r)(1-t)} & \text{if } tq/r \le A \le q/r \\ 0 & \text{if } q/r \le A. \end{cases} \tag{A.1}$$

The objective function $p_a AD(q,A)$ is increasing in A in the first case and decreasing in A in the last three. In the second case, we maximize

$$\max_A \left(p_a A - \frac{p_a (A - tq)^2}{2q(1-r)(1-t)} - cq^2 \right),$$

which yields $A^\star = (1 - r(1-t))q$. Substituting into the profit function results in a quadratic function of q:

$$\max_q \left(\frac{1 - r(1-t) + t}{2} \cdot p_a q - cq^2 \right),$$

Maximizing gives the stated formula. □

Proof of Proposition 1. First, we determine the optimal A in the last stage of the game, given V and the set of consumers who download the ad blocker app. If the publisher optimizes only for consumers with a high enough κ who never download the app, the optimal A is the same as in the benchmark case, that is, $A^* = \max\left(\frac{q}{1+r}, q(1 - r(1-t))\right)$. Setting a different level of advertising would result in a discrete downward jump in revenue for the segment that never downloads. Therefore, our assumption of a high enough $\bar{\kappa}$ ensures that the mass of consumers with high enough κ to never download the app is high enough for the overall optimum A to remain the same.[23]

In the negotiation stage, the ad blocker first makes a take-it-or-leave it offer. As long as this offer results in a nonnegative payoff for the publisher, the publisher will accept. Therefore, the ad blocker always has an incentive to increase the demanded payment until it extracts all possible surplus. When this payment reaches the maximum, the publisher becomes indifferent and our tie-braking assumption dictates that the publisher accepts. We discuss relaxing the tie-braking rule and the assumption around the power balance of negotiations in Section 4. Since the publisher will always accept the ad blocker's offer, consumers will expect an advertising level of V when making the decision to download. Consumers who have a cost κ will download, if and only if their $\gamma > \frac{(1+r)\kappa}{q - V(1+r)}$ and $\vartheta > \frac{\kappa}{q} + \gamma \frac{V}{q}$. The first constraint ensures that a consumer is better off with an ad level of V instead of $A = \frac{q}{1+r}$ and the second constraint guarantees a positive utility. When $\kappa > \max\left(\frac{q}{1+r}, q(1 - r(1-t))\right) - V$, no consumer downloads the ad blocker. When $\kappa < \frac{qr}{1+r} - rV$, all consumers who visit the publisher's site download the ad blocker. When κ is between the two, we have a binding $\hat{\gamma}$ cutoff with $r < \hat{\gamma} = \frac{(1+r)\kappa}{q - V(1+r)} < 1$.

To determine the optimal V, we need to calculate how much surplus the ad blocker can extract during the negotiations. This surplus is equal to the difference between the revenue of the publisher under light ad blocking (V) and full ad blocking. This difference is simply the product of Vp_a and the amount of consumers who visit the site and download the ad blocker. We begin with the case of $t \le \frac{r}{1+r}$ when $A^\star = \frac{q}{1+r}$ and we also assume $V \ge \frac{qt}{r}$. For $\kappa < \frac{qr}{1+r} - rV$, the size of the segment is all consumers with $\vartheta > \frac{\kappa}{q} + \gamma \frac{V}{q}$, that is,

$$S_{V1}(\kappa, q, V, r) = \frac{1}{1-t} \left(1 - \frac{\kappa}{q} - \frac{(1+r)V}{2q} \right). \tag{A.2}$$

For $\frac{qr}{1+r} - rV < \kappa < \frac{q}{1+r} - V$, the $\gamma > \frac{(1+r)\kappa}{q - V(1+r)}$ constraint is also binding and the size of the segment becomes

$$S_{V2}(\kappa, q, V, r) = \frac{1}{(1-r)(1-t)} \left(1 - \frac{(1+r)\kappa}{q - V(1+r)} \right) \times \left(1 - (1/2)\frac{\kappa}{q - V(1+r)} - (1/2)\frac{\kappa + V}{q} \right). \tag{A.3}$$

To determine the objective function of the ad blocker, that is, the total surplus it can extract, we have to integrate over κ for each segment:

$$\begin{aligned} T(q,V) = T_1(q,V) &= p_a V \int_0^{\frac{qr}{1+r} - Vr} \frac{1}{\bar{\kappa}} S_{V1}(\kappa, q, V, 0) d\kappa \\ &\quad + V p_a \int_{\frac{qr}{1+r} - Vr}^{\frac{q}{1+r} - V} \frac{1}{\bar{\kappa}} S_{V2}(\kappa, q, V, 0) d\kappa = \\ &= p_a V \frac{(q - V(1+r))(q(2 + 5r + 2r^2) - V(1+r)(1 + r + r^2))}{6q(1+r)^2(1-t)\bar{\kappa}} \end{aligned}. \tag{A.4}$$

There are two more cases we need to examine if $t \le \frac{r}{1+r}$. When $\frac{qt}{r} \ge V \ge qt$, we get

$$T_2(q,V) = p_a \frac{\begin{array}{l} V^2(1+r)^2(V - 3q(1 + r^2(1-t))) \\ +Vq^2(2 + 3r - 3r^2 + 3rt^2 - 2r^3 + 6r^2t^2 + 3r^3t^2) \\ -q^3t^3(1+r)^2 \end{array}}{6q(1+r)^2(1-r)(1-t)\bar{\kappa}},$$

and when $qt \ge V \ge 0$, we have

$$T_3(q,V) = p_a V \frac{q(r + (2 - 3t^2)(1+r)^2) - 3V(1+r)^3(1-t)}{6(1+r)^2(1-t)\bar{\kappa}}.$$

All three of these functions are either cubic or quadratic in V, hence we can maximize them in V. Depending on the value of t, the optimal V can fall in either of the three segments. When $t \ge \frac{r}{1+r}$, we have $A^\star = q(1 - r(1-t))$ and we can similarly integrate to get the ad blocker's payoff functions as

$$T_4(q,V) = p_a \frac{\begin{array}{l} (q(1 - r(1-t)) - V)(V(1 - r(1-t)) \\ (q(2 - rt - 3r^2(1-t)) - V) - q^2t^3) \end{array}}{6q(1-r)(1-t)(1 - r(1-t))\bar{\kappa}},$$

when $v \ge tq$. If $tq \ge V \ge 0$, then

$$T_5(q,V) = p_a V \frac{\begin{array}{l} q(3r^3t^2 - 6r^3t + 2r^2t^2 + 3r^3 + 2r^2t + 2rt^2 \\ -4r^2 + 2rt - t^2 - r + 2t + 2) - 3V(1+r)(1 - r(1-t)) \end{array}}{6(1 - r(1-t))\bar{\kappa}}.$$

Again, these are cubic and quadratic in V. Overall maximization yields that $V^* = qF(r,t)$, where

$$F(r,t) = \begin{cases} \text{i)} & \dfrac{1+r}{1+r+r^2} - \dfrac{\sqrt{1+5r+9r^2+5r^3+r^4}}{\sqrt{3}(1+r)(1+r+r^2)} \\ \text{ii)} & \dfrac{3(1+r)(1-r^2(1-t) - \sqrt{\begin{array}{l}3(1-r)(1-3r^5t^2+6r^5t-9r^4t^2-3r^5+18r^4t\\ -12r^3t^2-9r^4+18r^3t-9r^2t^2-6r^3+6r^2t\\ -3rt^2+4r^2+4r)\end{array}}}{3(1+r)} \\ \text{iii)} & \dfrac{2+5r-3t^2+2r^2-6rt^3-3r^2t^2}{6(1+r)^3(1-t)} \\ \text{iv)} & \dfrac{3(1-r(1-t))(1-r^2(1-t)) - \sqrt{\begin{array}{l}3(1-r(1-t)(1-r)(1-t)\\ (3r^4t^2-6r^4t+3r^4+3r^3t+r^2t^2-3r^3\\ +r^2t+rt^2-2r^2+rt+t^2+r+t+1)\end{array}}}{3(1-r(1-t))} \\ \text{v)} & \dfrac{\begin{array}{l}3r^3t^2-6r^3t+2r^2t^2+3r^3+2r^2t\\ +2rt^2-4r^2+2rt-t^2-r+2t+2\end{array}}{6(1+r)(1-r(1-t))}. \end{cases}$$

$$\text{if} \begin{cases} \text{i)} & 0 \le t \le r\dfrac{3(1+r)^2 - \sqrt{3r^4+15r^3+27r^2+15r+3}}{3(1+r)(1+r+r^2)} \\ \text{ii)} & r\dfrac{3(1+r)^2 - \sqrt{3r^4+15r^3+27r^2+15r+3}}{3(1+r)(1+r+r^2)} \le t \le \min\left(\dfrac{r}{1+r}, \dfrac{3(1+r)^2 - \sqrt{9r^4+24r^3+18r^2+9r+3}}{3(1+r)(1+2r)}\right) \\ \text{iii)} & \dfrac{3(1+r)^2 - \sqrt{9r^4+24r^3+18r^2+9r+3}}{3(1+r)(1+2r)} \le t \le \dfrac{r}{1+r} \\ \text{iv)} & \dfrac{r}{1+r} \le t \le \dfrac{(3r^2-r-2+\sqrt{3(1+r)(3r+2)})(1-r)}{1+4r+4r^2-3r^3}, \\ \text{v)} & \max\left(\dfrac{r}{1+r}, \dfrac{(3r^2-r-2+\sqrt{3(1+r)(3r+2)})(1-r)}{1+4r+4r^2-3r^3}\right) \le t \le 1. \quad \square \end{cases} \tag{A.5}$$

Proof of Proposition 2. In the first stage, the publisher maximizes its profit given V. The publisher's revenue is exclusively from consumers who did not download the ad blocker as all the surplus generated by the downloaders is extracted by the ad blocker. We begin with the case of $t \le \frac{r}{1+r}$. We calculate the lost revenue compared with the benchmark case of no ad blocker. The lost revenue for consumers in $\frac{qr}{1+r} - Vr < \kappa < \frac{q}{1+r} - V$ is

$$\frac{qp_a}{(1+r)(1-r)(1-t)}\left(1 - \frac{\kappa}{2(q-V(1+r))} - \frac{1}{2(1+r)}\right)\left(1 - \frac{\kappa(1+r)}{(q-V(1+r))}\right). \tag{A.6}$$

For $\kappa > \frac{q}{1+r} - V$, the lost revenue is 0, whereas for $\kappa < \frac{qr}{1+r} - Vr$, it is $\frac{qp_a}{2(1+r)(1-t)}$. Integrating over κ yields

$$\frac{qp_a}{(1+r)(1-t)\bar{\kappa}}\left(\int_{\frac{qr}{1+r}-Vr}^{\frac{q}{1+r}-V} \frac{1}{1-r}\left(1 - \frac{\kappa}{2(q-V(1+r))} - \frac{1}{2(1+r)}\right)\left(1 - \frac{\kappa(1+r)}{(q-V(1+r))}\right)d\kappa + \int_0^{\frac{qr}{1+r}-Vr} 1 d\kappa\right) = \frac{(r^2+4r+1)(q-(1+r)V)qp_a}{6\bar{\kappa}(1+r)^3(1-t).} \tag{A.7}$$

Substituting the equilibrium V and subtracting from the benchmark, we get the objective function of

$$\frac{qp_a}{2(1+r)(1-t)} - \frac{(r^2+4r+1)(1-(1+r)F(r))q^2p_a}{6\bar{\kappa}(1+r)^3(1-t) - cq^2.}$$

Since this is a simple quadratic function in q, the optimal quality is

$$q^* = \left(\frac{4c(1+r)(1-t)}{p_a} + \frac{2(r^2+4r+1)(1-(1+r)F(r))}{3\bar{\kappa}(1+r)^2}\right)^{-1},$$

which is clearly an increasing function of $\bar{\kappa}$ and converges to $\frac{p_a}{4c(1+r)(1-t)}$ as $\bar{\kappa} \to \infty$. For $t \ge \frac{r}{1+r}$, similar calculations yield the objective function as

$$\frac{qp_a(1-r(1-t)+t)}{2} - \frac{\begin{array}{c}(3r+1)(1-r)^2+\\ ((1+r+t+rt+4r^2+r^2t-6r^3+3r^3t)t)\\ (1-r(1-t)-F(r))q^2p_a\end{array}}{6\bar{\kappa}(1-r(1-t)) - cq^2.}$$

Maximizing in q yields the formula stated in the proposition where

$$G(r,t) = \begin{cases} \dfrac{2(r^2+4r+1)(1-(1+r)F(r,t))}{3(1+r)^2} & \text{if} \quad t \le \dfrac{r}{1+r} \\ \dfrac{\begin{array}{l}2((3r+1)(1-r)^2+(1+r+t+rt+4r^2\\ +r^2t-6r^3+3r^3t)t)(1-r(1-t)-F(r,t))\end{array}}{3(1-r(1-t))(1-r(1-t)+t)} & \text{if} \quad \dfrac{r}{1+r} \le t \end{cases}. \tag{A.8}$$

Given the quadratic profit function, it is easy to calculate the equilibrium profit, yielding the comparative statics results:

$$\Pi^*_{PUB} = \begin{cases} \dfrac{q^*p_a}{2(1+r)(1-t)} & \text{if} \quad t \le \dfrac{r}{1+r} \\ \dfrac{q^*p_a(1-r(1-t)+t)}{4} & \text{if} \quad \dfrac{r}{1+r} \le t. \end{cases} \quad \square$$

Proof of Corollary 1. In (A.4), and the following equations, we have essentially calculated the size of the downloader segment, which is $T(q,V)/(Vp_a)$. Plugging in $V^* = q^*F(r)$ yields that in equilibrium the size of this segment is a linear function of $\frac{q^*}{\bar{\kappa}}$ and its multiplier only depends on r and t. Since q^* is an increasing function of p_a, so is the segment size. Furthermore, $\frac{q^*}{\bar{\kappa}}$ is a decreasing function of $\bar{\kappa}$. $\square$

Proof of Corollary 2. Consumers who do not download the ad blocker are worse off, because the ad blocker causes the publisher quality to decline, reducing their surplus. In particular, a consumer with $\gamma = r$ and $\vartheta = 1$ has a surplus of $q - rA = q(\min\left(\frac{1}{1+r}, 1 - r + r^2(1-t)\right)) > 0$ if $r < 1$. This is surplus clearly goes down as the ad blocker enters and quality decreases. On the other hand, there are consumers who only visit the publisher's site if there is an ad blocker. Consumers who are below the $\vartheta = \gamma\max\left(\frac{1}{1+r}, 1 - r(1-t)\right)$ line on the (ϑ, γ) space do not visit the publisher's site if there is no ad blocking and hence get zero utility in the absence of the ad blocker. However, in the presence of an ad blocker, consumers with ϑ in the range $\gamma\max\left(\frac{1}{1+r}, 1-r(1-t)\right) > \vartheta > \frac{\kappa+\gamma V}{q}$ visit the site get positive utility if their κ is low enough.

Proof of Proposition 3. Total consumer utility can be calculated as

$$CS_{AB} = q^{\star}C_0(r,t) + (q^{\star})^2\frac{C_{\Delta}(r,t)}{\bar{\kappa}}. \tag{A.9}$$

The first component is the benchmark consumer surplus. The second component is positive and represents the additional surplus consumers enjoy from ad blocking when quality is fixed. The benchmark consumer surplus multiplier is

$$C_0(r,t) = \begin{cases} \dfrac{1+r+r^2}{6(1+r)^2} & \text{if} \quad t \le \dfrac{r}{1+r} \\ \dfrac{(1-t)^2(1+r+r^2) + 3((t-r(1-t)) \\ (1-r+r^2(1-t))}{6(1-r(1-t))} & \text{if} \quad \dfrac{r}{1+r} \le t. \end{cases}$$

The second component is a longer formula, but we can calculate it for certain parameter regions:

$$C_{\Delta}(r,0) = \frac{(2(r+1) - F(r,t)(1+r^2))(1-F(r,t)(1+r))^2}{24(1+r)}.$$

$$C_{\Delta}(r,1) = \frac{(1+r+r^2)(1-F(r,t))^2}{6}.$$

$$C_{\Delta}(0,t) = \begin{cases} \dfrac{(F(r,t)(2-F(r,t)) - t^4)(1-F(r,t))^2}{24F(r,t)(1-t)} & \text{if} \quad t \le \sqrt{6}-2 \\ \dfrac{2(1+t+t^2-t^3) \\ +F(r,t)(4(F(r,t)-5-5t+t^2+t^3)}{24} & \text{if} \quad \sqrt{6}-2 \le t \end{cases}$$

$$C_{\Delta}(1,t) = \begin{cases} \dfrac{(2-F(r,t))(1-2F(r,t))^2}{24} & \text{if} \quad t \le \dfrac{4-\sqrt{7}}{6} \\ \dfrac{4(1-t)(3F(r,t)(F(r,t)-1-t)+t^2) \\ -3F(r,t)-2+4t^2}{24(1-t)} & \text{if} \quad \dfrac{4-\sqrt{7}}{6} \le t \le \dfrac{r}{1+r} \\ \dfrac{(t-F(r,t))^2}{2} & \text{if} \quad \dfrac{r}{1+r} \le t \end{cases}$$

To compare with CS_{NA}, we examine CS_{AB} as a function of $\bar{\kappa}$. As $q^{\star}$ converges to $q^{\star}_{NA}$ when $\bar{\kappa} \to \infty$, we get the same for CS_{AB}. What we need to determine is whether it converges increasingly or decreasingly. Plugging in the formula for $q^{\star}$ from Proposition 2 and differentiating with respect to $\bar{\kappa}$, we find that there are two distinct cases. The derivative is always positive for small values of $\bar{\kappa}$. However, depending on whether $C_{\Delta}(r,t) > C_0(r,t)G(r,t)$ or not, the derivative either crosses into negative at some threshold of $\bar{\kappa}$ or stays positive for all $\bar{\kappa}$. Hence, consumer surplus is increasing for all $\bar{\kappa}$ values if $C_{\Delta}(r,t) \le C_0(r,t)G(r,t)$ or it is first increasing and then decreasing after reaching a maximum if $C_{\Delta}(r,t) > C_0(r,t)G(r,t)$. The function $G(r, t)$ is given in (A.8) and, as we can see, the condition only depends on r, t. Figure 3 presents the contour line for $C_{\Delta}(r,t) = C_0(r,t)G(r,t)$, which is the solution of a sixth degree polynomial, hence it has no closed form solution. For the purposes of the corollary, all we need to show is that $C_{\Delta}(0,t) \le C_0(0,t)G(0,t)$ for any t, that $C_{\Delta}(r,0) \le C_0(r,0)G(r,0)$ for any r, and that $C_{\Delta}(1,1) \le C_0(1,1)G(1,1)$. With the formulas given earlier, this is straightforward to verify. Since both sides of the inequality change continuously, we get that the left side is smaller when either r or t is low and that the left side is larger when both r and t are high, yielding the stated result. □

Proof of Corollary 3. To see that total welfare is lower (even if we include the ad blocker's surplus), we need to establish that the sum of the publisher surplus and the ad blocker surplus is lower in the presence of the ad blocker. There are two effects. First, the ad blocker extracts some of the surplus generated by the publisher. Keeping the quality fixed, this is a zero-sum game. The second effect of the ad blocker is that quality decreases. But with a lower quality, the publisher surplus decreases. The combination of the two effects reduces the sum of the publisher and ad blocker surpluses. Clearly, if consumer surplus is also reduced by the ad blocker's entry, the total welfare decreases as well.

When $CS_{AB} > CS_{NA}$, we need to check whether the increase can counteract the decrease in the sum of the publisher and ad blocker's surplus. Price p_a plays a crucial role here as both the publisher's and the ad blocker's surplus are scaled by it. Quality $q^{\star}$ is also dependent on p_a, but only as a function of p_a/c. Hence if p_a is reduced while keeping p_a/c constant, the sum of the losses of the publisher and the ad blocker can be reduced compared with the gains in consumer surplus. On the contraery, if either p_a or c is too high, total welfare is reduced. □

Proof of Lemma 2. We determine the optimal V just as in the proof of Proposition 1 by first calculating the amount of people who download the ad blocker and use it for partial ad blocking. In the typical case depicted in Figure 4, this segment is in the (ghik) trapezoid with a size of

$$S_{V2}(\kappa,\mu,q,V,r) = \frac{1}{1-r}\left(\frac{\mu\kappa}{V} - \frac{(1+r)\kappa}{q-V(1+r)}\right)\left(1 - \frac{\kappa}{2(q-V(1+r))} - \frac{(1+\mu)\kappa}{2q}\right).$$

This formula applies for $\frac{qr}{1+r} - rV < \kappa < \frac{V}{\mu}$. For smaller $\kappa < \frac{qr}{1+r} - rV$, we have

$$S_{V1}(\kappa,\mu,q,V,r) = \frac{1}{1-r}\left(\frac{\mu\kappa}{V} - r\right)\left(1 - \frac{(1+\mu)\kappa}{2q} - \frac{\kappa+rV}{2q}\right).$$

representing the (deki) trapezoid. And for larger $\frac{V}{\mu} < \kappa < \frac{q}{1+r} - V$, we obtain

$$S_{V3}(\kappa,\mu,q,V,r) = \frac{1}{1-r}\left(1 - \frac{(1+r)\kappa}{q-V(1+r)}\right)\left(1 - \frac{\kappa}{2(q-V(1+r))} - \frac{\kappa+V}{2q}\right).$$

for the (abik) trapezoid. Assuming $r < 1$ and integrating over κ yields the publisher's payoff function:

$$\begin{aligned} T(q,V) &= Vp_a\int_0^{\frac{qr}{1+r}-Vr} 1\bar{\kappa}S_{V1}(\kappa,q,V,0)d\kappa + \\ &+ Vp_a\int_{\frac{qr}{1+r}-Vr}^{\frac{V}{\mu}} 1\bar{\kappa}S_{V2}(\kappa,q,V,0)d\kappa + Vp_a\int_{\frac{V}{\mu}}^{\frac{q}{1+r}-V} 1\bar{\kappa}S_{V3}(\kappa,q,V,0)d\kappa = \\ &= Vp_a\left(\frac{(q-V(1+r))(q(2+5r+2r^2) - V(1+r)(1+r+r^2))}{6q(1+r)^2\bar{\kappa}} - \frac{3\mu qV - (1+2\mu)V^2}{6q\mu^2(1-r)\bar{\kappa}}\right), \end{aligned}$$

if $V < \frac{\mu q}{(1+\mu)(1+r)}$, and 0 otherwise. Maximizing $T(q, V)$ yields

$$V^* = q\mu\left(\frac{1-\mu-\mu r^2}{1-2\mu-\mu^2-\mu^2 r^3} - \frac{\sqrt{\begin{array}{l}\mu^2(1-r)^2(1+5r+9r^2+5r^3+r^4)\\ +2\mu(1-r)(3r^3+7r^2+4r+1)\\ +2r^3+6r^2+3r+1\end{array}}}{\sqrt{3}(1+r)(1-2\mu-\mu^2-\mu^2 r^3)}\right) = qF_\mu(r).$$

Analyzing $F_\mu(r)$ at different values yields the stated results. □

Proof of Proposition 4. The analysis is identical to that in the proof of Proposition 2, but with $F_\mu(r)$ instead of $F(r)$. □

Proof of Corollary 4. Total consumer utility can be calculated as $CS_\mu = q^\star C_{\mu 0}(r,0) + (q^\star)^2 \frac{C_{\mu\Delta}(r,0)}{\bar{\kappa}}$ where

$$C_0(r,0) = \frac{1+r+r^2}{6(1+r)^2} \quad \text{and} \quad C_{\mu\Delta}(r,0) = \frac{(2(r+1)-F_\mu(r)(1+r^2))(1-F_\mu(r)(1+r))^2}{24(1+r)}.$$

Using $G_\mu(r,0) = \frac{2(r^2+4r+1)(1-(1+r)F_\mu(r))}{3(1+r)^2}$, we can see that $C_{\mu\Delta}(r,0) < C_0(r,0)G_\mu(r,0)$, yielding that consumer surplus increases in $\bar{\kappa}$ as it approaches the benchmark. We have shown in the proof of Corollary 3 that for $t = 0$ the sum of the publisher's and the ad blocker's surplus is lower in the full model than in the benchmark case. Both the publisher's and the ad blocker's surplus are lower in this version of the model with the full ad blocking option as the quality is lower and the ad blocker collects a lower revenue per consumer from a smaller segment.

When $t = 1$, we have

$$F^\star_\mu(r,1) = \frac{\mu}{2(1+\mu)},\ G_\mu(r,1) = \frac{1+r}{2}\cdot\frac{2+\mu}{1+\mu},\ C_{\mu 0}(r,1) = \frac{1-r}{2},\ C_{\mu\Delta}(r,1) = \frac{(4+\mu)(1+r+r^2)}{24(1+\mu)}.$$

Comparing $C_{\mu\Delta}(r,1)$ with $C_0(r,1)G_\mu(r,1)$ reveals that the left-hand side is larger if and only if $r > \hat{r}(\mu) = \frac{\sqrt{141\mu^2+552\mu+528}-4-\mu}{14\mu+32}$, where $\hat{r}(\mu)$ is increasing in μ. □

Proof of Proposition 5. As in the main model, the advertising levels are not affected by the ad blocker because there are enough consumers with high download cost who do not use the ad blocker and thus ad levels are the same as given in Lemma 1 as a funciton of q. Specifically, for $0 \le t \le \frac{r}{r+1}$, we have $A = \frac{q}{1+r}$. Given the ad levels and p_b, it is straightforward to calculate the number of downloaders, which yields the optimal price. Assuming $0 \le t \le \frac{r}{r+1}$, and $r \le \frac{\sqrt{57}-7}{2}$ we get

$$p_b^* = \frac{9(1+r)+1-\sqrt{81r^2+84r+36}}{16(1+r)}q.$$

Given this price, consumers with $\gamma > \frac{\kappa+p_b^*}{q}(1+r)$ and $\vartheta > \frac{\kappa+p_b^*}{q}$ will download and use the ad blocker. Hence, the publisher will only be able to generate revenue on the remaining consumers. The publisher's payoff can be calculated as

$$\frac{qp_a}{2(1+r)(1-t)} - \frac{(3-r-6r^2)(3-r)^2 q^2 p_a}{384\bar{\kappa}(1+r)^2(1-r)(1-t)-cq^2},$$

as long as $r < \frac{\sqrt{57}-7}{\sqrt{57}-5}$. Since this is a simple quadratic function in q, the optimal quality is

$$q^* = \left(\frac{4c(1+r)(1-t)}{p_a} + \frac{(3-r-6r^2)(3-r)^2}{96\bar{\kappa}(1+r)^2(1-r)(1-t)}\right)^{-1}.$$

When $t = 0$, this yields $q^*_{CC} \xrightarrow{r\to 0} \left(\frac{4c}{p_a} + \frac{9}{32\bar{\kappa}}\right)^{-1} < q^*_{NA} = \frac{p_a}{4c}$. Calculating $q^*_{AB} \xrightarrow{r\to 0} \left(\frac{4c}{p_a} + \frac{2\sqrt{3}}{9\bar{\kappa}}\right)^{-1}$ yields $q^*_{AB} < q^*_{CC}$ for small values of r, specifically for $r < \frac{\sqrt{57}-7}{\sqrt{57}-5}$. The publisher's profit is then simply $\Pi^{pub\star}_{CC} \xrightarrow{r\to 0} p_a^2\left(16c + \frac{9p_a}{8\bar{\kappa}}\right)^{-1} < \Pi^{pub\star}_{NA} = \frac{p_a^2}{16c}$, and $\Pi^{pub\star}_{AB} \xrightarrow{r\to 0} p_a^2\left(16c + \frac{8\sqrt{3}}{9\bar{\kappa}}\right)^{-1}$ in the two cases, yielding $\Pi^{pub\star}_{AB} < \Pi^{pub\star}_{CC} < \Pi^{pub\star}_{NA}$ for small r values. Next, we calculate the ad blocker's profit as a function of q for the same range of small r values:

$$\Pi^{Adb}_{CC} = \frac{\left(9(1+r)+1-\sqrt{81r^2+84r+36}\right)\left(67r^2+124r-20+(21r+18)\sqrt{81r^2+84r+36}\right)}{12288\bar{\kappa}(1+r)^3(1-t)q^2}.$$

We then have $\Pi^{Adb}_{CC} \xrightarrow{r\to 0} \frac{11}{384\bar{\kappa}}(1-t)q^2$ and from our previous analysis, we have $\Pi^{Adb}_{AB} \xrightarrow{r\to 0} \frac{\sqrt{3}}{18\bar{\kappa}}(1-t)q^2$. Thus, for the same quality level, charging consumers leads to more than three times lower profit. However, the quality is slightly higher when charging consumers, but $\frac{q^\star_{CC}}{q^\star_{AB}} < \frac{64\sqrt{3}}{81} < 1.369$ for small r, hence $\Pi^{Adb\star}_{CC} < \Pi^{Adb\star}_{AB}$. Moving on to the consumer surplus, we start from (A.9) to establish that

$$CS_{AB} \xrightarrow{r\to 0} q^\star_{AB}C_0(0,t) + (q^\star_{AB})^2\frac{3+\sqrt{3}}{216\bar{\kappa}},$$

where $CS_{NA} = q^\star_{NA}C_0(r,t)$. We can similarly calculate

$$CS_{CC} \xrightarrow{r\to 0} q^\star_{CC}C_0(0,t) + (q^\star_{CC})^2\frac{27}{1024\bar{\kappa}},$$

and a simple comparison yields $CS^\star_{AB} < CS^\star_{CC} < CS^\star_{NA}$ for small r values. Adding the resulting consumer suprlus to publisher and ad blocker profits further yields $TW^\star_{AB} < TW^\star_{CC} < TW^\star_{NA}$.

When $r = t = 1$, the analysis becomes simpler as consumers are only heterogenous with respect to the download costs. As before, the publisher sets the advertising to q, regardless of who downloads the ad blocker. Then any consumer for whom $\kappa + p_b < q$ will download and use the ad blocker. This yields a profit function of $p_b\frac{q-p_b}{\bar{\kappa}}$ for the ad blocker with an optimal price of $p_b^\star = q/2$. Recall that this is the same value as $V^\star = q/2$ in the main model. As a result in equilibrium, the exact same consumers will download the ad blocker regardless of whether the publisher or consumers are charged. The publisher only makes money on consumers who do not download the ad blocker, hence publisher profits will be the same as well. Given the equal cost to the downloading consumers—whether through a direct price or nonzero advertising—consumer surplus will be the same, as well as ad blocker profits and also total welfare.

Proof of Proposition 6. We have calculated the size of the segment who does not download the ad blocker in (A.7) as

$$SNDL(r) = \frac{1}{2} - \frac{(r^2+4r+1)(1-(1+r)F(r))q}{6\bar{\kappa}(1+r)^2}.$$

The publisher gets $\frac{qp_a}{1+r}$ for each consumer in this segment. The size of the segment that downloads was calculated in (A.4). After substituting $V = qF(r)$, we get

$$SDL(r) = \frac{q(1-(1+r)F(r))(2+5r+2r^2-(1+r)(1+r+r^2)F(r))}{6(1+r)^2\bar{\kappa}}.$$

The publisher gets $(1-\alpha)Vp_a = (1-\alpha)F(r)qp_a$ for each consumer in this segment. Hence the publisher maximizes

$$\frac{qp_a}{1+r}SNDL(r) + (1-\alpha)F(r)qp_aSDL(r) - cq^2$$

in q, yielding

$$q^\star(\alpha) = \left(\frac{\frac{4c(1+r)}{p_a} + \frac{2(1-(1+r)F(r))}{3\bar{\kappa}(1+r)}}{r^2+4r+1-(1-\alpha)(1+r)F(r) \cdot \frac{(2+5r+2r^2-(1+r)(1+r+r^2)F(r))}{1+r}}\right)^{-1},$$

which is the same as stated in the proposition with $H(r) = (1+r)F(r)(1-(1+r)F(r))(2+5r+2r^2-(1+r)(1+r+r^2)F(r))$. Since $H(r) > 0$ is increasing, we get the comparative statics results. □

Proof of Corollary 5. We use the same steps as in the proof of Proposition 3 and check whether $C_\Delta() > C_0()G()$. The particular formulas for $t = 0$ are

$$C_0(r,0) = \frac{1+r+r^2}{6(1+r)^2},$$

$$C_\Delta(r,0) = \frac{(2(r+1)-F(r,0)(1+r^2))(1-F(r,0)(1+r))^2}{24(1+r)},$$

$$G_\alpha(r,0) = \frac{2(1-(1+r)F(r,0))\Big(r^2+4r+1-(1-\alpha)(1+r)F(r,0)\,(2+5r+2r^2-(1+r)(1+r+r^2)F(r,0))\Big)}{3(1+r)^2},$$

We find that when $r = \alpha = 0$, we get $C_\Delta(r,0) > C_0(r,0)G_0()$. For positive α and r values, the inequality quickly reverses. When $t = 1$, we have

$$C_0(r,1) = \frac{1-r}{2},\quad C_\Delta(r,1) = \frac{1+r+r^2}{24},\quad G_\alpha(r,1) = \frac{(1+r)(1+\alpha)}{4},$$

yielding that consumer surplus can go up if and only if $1+r+r^2 > 3(1-r)(1+r)(1+\alpha)$. □

Proof of Proposition 7. The ad blocker maximizes $(1-\alpha)F(r)qp_aSDL(r)$ in α. This function takes the form of $const \cdot q^\star(\alpha)^2\alpha$, where $q^\star(\alpha) = const' \cdot (D+\alpha)^{-1}$ with

$$D = \frac{(c\bar{\kappa}/p_a)\cdot 2(1+r)^3 + (r^2+4r+1)(1-(1+r)F(r)) - H(r)}{H(r)},$$

and the constasnt not depending on α. Maximizing $(D+\alpha)^{-2}\alpha$ yields a single maximum at $\alpha = \min(D,1)$. Given the expression for D, we see that it is increasing in $(c\bar{\kappa}/p_a)$. Furthermore, setting $c\bar{\kappa}/p_a = 0$ gives $0 < D < 1$, and since D is linear in $(c\bar{\kappa}/p_a)$, we get that there is a single threshold where $D = 1$. □

Endnotes

[1] Consumers using ad blockers face various nuisance costs (e.g., loss of content, frequent requests for upgrades). Some consumers feel uncomfortable toward publishers when using ad blockers whereas others face access difficulties (e.g., many corporate users cannot freely download browser extensions on their computers).

[2] See, for example, Reuters (2018).

[3] Our analysis is strictly based on economic surplus and does not take into account other negative externalities related to the reduction of content quality (e.g., its consequences for democracy).

[4] The term *whitelisting* is also used in other contexts (e.g., when consumers allow some publishers' advertising to pass the ad blocker). In this paper, we consistently refer to whitelisting as the practice by ad blockers of allowing a certain amount of qualified advertising from specific publishers to pass the blockade, often in return for payment.

[5] Some ad blockers may also allow consumers to block all ads, although this option is discouraged and usually comes with additional costs to consumers. Nevertheless, some consumers do opt for this option.

[6] We have already mentioned vigorous legal action but many publishers have also invested in so-called ad blocker blockers or introduced a subscription revenue model (a case we explicitly analyze).

[7] See also Chen and Liu (2021) on how ad blocking may help publishers in a context where advertising serves as a signaling device.

[8] This is a reasonable assumption given the low proportion of people who use ad blockers and the vast inventory of online advertising available. Also, in the quasi-endogenous case (explored in the web appendix) when p_a is a function of the publisher's content quality, we find that our key findings remain similar.

[9] In the web appendix, we also analyze the case when the publisher directly charges consumers in a subscription revenue model.

[10] We find that all three dimensions of consumer heterogeneity are necessary to explain the relevant dynamics of ad blocking. For additional clarity, the web appendix presents the analyses of cases when consumer heterogeneity is only along two dimensions.

[11] Setting r strictly above 0 is a purely technical assumption to avoid the case where advertising can be set to infinity for $r = 0$ consumers. Parameters r and t allow us to gauge the impact of the level of consumer heterogeneity on ad blocking.

[12] In its default setting, Adblock Plus, the largest ad blocker, explicitly tells consumers that some, so-called acceptable ads will pass the blockade. With a few more clicks, consumers can also opt to block all ads, a case we will explicitly model further later.

[13] It is worth noting that if we change r and t, the mean sensitivity levels also change, but $V^\star/A^\star$ is invariant to the scaling of the means.

[14] Note that advertising levels do not change with the entry of the ad blocker and this is driven by a large enough mass of consumers with high download costs. In the web appendix, we analyze a setting where all consumers have the same moderate level of download cost and we find that advertising levels may slightly increase with the entry of an ad blocker.

[15] We thank one of the associate editors for drawing our attention to this case.

[16] Beyond the somewhat trivial cost of a few additional clicks, many consumers experience a psychological cost for hurting publishers by blocking acceptable ads.

[17] See Reuters (2018).

[18] Publishers can also coordinate legal action against a powerful ad blocker and have done so in European countries

[19] We analyze this specific case in a web appendix.

[20] For completeness, in the web appendix, we also analyze the case when the ad blocker can discriminate between publishers.

[21] If we use the same cost structure as in the basic model with different costs, we have to resort to numerical analysis to calculate the equilibrium qualities.

[22] Adblock Plus, the most popular ad blocker, is estimated to represent about 60% of the desktop browser market, for example, with no comparable second player (Maheshwari 2016).

[23] We discuss the case of a fixed, potentially lower κ in the web appendix.

References

Anderson SP, Gans JS (2011) Platform siphoning: Ad-avoidance and media content. *Amer. Econom. J. Microeconomics* 3(4):1–34.

Bresnahan TF, Reiss PC (1985) Dealer and manufacturer margins. *RAND J. Econom.* 16(2):253–268.

Chen Y, Liu Q (2021) Signaling through advertising when ad can be blocked. *Marketing Sci.* Forthcoming.

Choi JP, Thum M (2004) The economics of repeated extortion. *RAND J. Econom.* 35(2):203–223.

Danaher B, Dhanasobhon S, Smith MD, Telang R (2010) Converting pirates without cannibalizing purchasers: The impact of digital distribution on physical sales and internet piracy. *Marketing Sci.* 29(6):1138–1151.

Despotakis S, Ravi R, Srinivasan K (2021) The beneficial effects of ad blockers. *Management Sci.* 67(4):2096–2125.

Goldstein DG, Suri S, McAfee P, Ekstrand-Abueg M, Diaz F (2014) The economic and cognitive costs of annoying display advertisements. *J. Marketing Res.* 51(6):742–752.

Jain S (2008) Digital piracy: A competitive analysis. *Marketing Sci.* 27(4):610–626.

Johnson JP (2013) Targeted advertising and advertising avoidance. *RAND J. Econom.* 44(1):128–144.

Katona Z, Sarvary M (2018) Eyeo's Adblock Plus: Consumer movement or advertising toll booth? *Berkeley Hass Case Series* (March 31), https://cases.haas.berkeley.edu/2018/04/adblock/.

Konrad KA, Skaperdas S (1997) Credible threats in extortion. *J. Econom. Behav. Organ.* 33(1):23–39.

Maheshwari S (2016) Adblock Plus, created to protect users from ads, instead opens the door. *New York Times* (September 18). https://www.nytimes.com/2016/09/19/business/media/adblock-plus-created-to-protect-users-from-ads-opens-the-door.html.

Moorthy KS (1988) Product and price competition in a duopoly. *Marketing Sci.* 7(7):141–168.

Olken BA, Barron P (2009) The simple economics of extortion: Evidence from trucking in Aceh. *J. Political Econom.* 117(3):417–452.

Perrin N (2019) Consumer attitudes on marketing 2019. *eMarketer* (August 29), https://www.emarketer.com/content/consumer-attitudes-on-marketing-2019.

Reuters (2018) German Supreme Court rules ad blockers legal, in defeat for Springer. *Reuters* (April 19), https://www.reuters.com/article/us-germany-trial-adblocking/german-supreme-court-rules-ad-blockers-legal-in-defeat-for-springer-idUSKBN1HQ277.

Reyes MS (2020) Consumers would still rather protect their data than get personalized ads. *eMarketer* (November). Accessed November 25, 2020, https://content-na1.emarketer.com/consumers-would-still-rather-protect-their-data-than-personalized-ads.

Rochet J-C, Tirole J (2006) Two-sided markets: A progress report. *RAND J. Econom.* 37(3):645–667.

Shaked A, Sutton J (1982) Relaxing price competition through product differentiation. *Rev. Econom. Stud.* 49(1):3–13.

Shiller B, Waldfogel J, Ryan J (2018) The effect of ad blocking on website traffic and quality. *RAND J. Econom.* 49(1):43–63.

Spengler JJ (1950) Vertical integration and antitrust policy. *J. Political Econom.* 58(4):347–352.

Vasagar J (2015) Axel Springer winning fight against adblockers. *Financial Times* (November 4), https://www.ft.com/content/0e58d68a-8305-11e5-9dc0-186bb1146746.

Vernik DA, Purohit D, Desai PS (2011) Music downloads and the flip side of digital rights management. *Marketing Sci.* 30(6):1011–1027.

Vratonjic N, Manshaei MH, Grossklags J, Hubaux J-P (2013) Ad-blocking games: Monetizing online content under the threat of ad avoidance. Böhme R, ed. *The Economics of Information Security and Privacy* (Springer, Berlin, Heidelberg), 49–73.