THE EFFECT OF PIRACY WEBSITE BLOCKING ON CONSUMER BEHAVIOR

Brett Danaher[‡]
bdanaher@wellesley.edu

Michael D. Smith[†]
mds@cmu.edu

Rahul Telang[†]
rtelang@andrew.cmu.edu

This Version: March 2018

Acknowledgements: This research was conducted as part of Carnegie Mellon University's Initiative for Digital Entertainment Analytics (IDEA), which receives unrestricted (gift) funding from the Motion Picture Association of America. Danaher acknowledges support from an NBER Economics of Digitization research grant. This research was conducted independently without any oversight or editorial control. The authors presented an earlier version of this paper at the December 2014 Workshop on Information Systems and Economics and to seminar participants at the University of Arizona's Eller School of Management and thank participants for their helpful feedback. All findings and errors are entirely our own.

‡ Argyros School of Business and Economics, Chapman University, Orange, CA 92868. Corresponding author.

† School of Information Systems and Management, Heinz College, Carnegie Mellon University, Pittsburgh, PA, 15213.

THE EFFECT OF PIRACY WEBSITE BLOCKING ON CONSUMER BEHAVIOR

ABSTRACT

Understanding the relationship between copyright policy and consumer behavior is an increasingly important topic for both rights holders and policymakers. In this paper we study how consumer behavior changes when Internet Service Providers are required to block access to major piracy websites. We do this in the context of three court-ordered events affecting consumers in the UK: A blocking order directed at The Pirate Bay in May 2012, a blocking order directed at 19 major piracy sites in November 2013, and a blocking order aimed at 52 different piracy sites in 2014.

Our results show that blocking a single site—The Pirate Bay—only caused a small reduction in total piracy and no increase in usage of legal sites. Instead, consumers seemed to turn to other piracy sites or Virtual Private Networks that allowed them to circumvent the block. In contrast, blocking 19 different major piracy sites caused a meaningful reduction in total piracy and subsequently led former users of the blocked sites to increase their usage of paid legal streaming sites such as Netflix by 11% on average. Similarly, blocking 52 sites in 2014 caused treated users to increase their usage of legal subscription sites by 10% and legal ad-supported streaming sites by 11.5%. These results are heterogeneous across groups such that users who have not yet formed a strong tie to either legal or piracy channels are the most likely to be impacted by the blocks.

Keywords: *Piracy, regulation, digital distribution, motion picture industry, natural experiment.*

1. Introduction

One of the most important challenges facing the media industries today is whether and how copyright policy should be adapted to the realities of the digital age. The invention and subsequent adoption of filesharing technologies[1] have eroded the strength of copyright law across many countries. In the ten years following the introduction of Napster in 1999, worldwide revenues from recorded music fell by 50% (IFPI 2010), and in the four years after the introduction of BitTorrent, home video sales declined in the film industry by 27% (Zentner 2010). The vast majority of the academic literature has found that digital piracy causes a significant reduction in sales of music and motion picture content (see Danaher et al. 2014b for a review of this literature). Though the literature on piracy and the supply of creative works is somewhat inconclusive, there exists some evidence that diminished revenues from piracy have the potential to lead to a decrease in the quantity and quality of films that are produced (Telang and Waldfogel 2014, Danaher and Smith 2016). Thus it is important, not only from a business perspective but also from a social welfare perspective, to understand how to design and enforce copyright policy in an age of filesharing technologies.

Accordingly, there is tremendous interest in evaluating the impact of antipiracy legislation on consumer behavior and market outcomes. Several papers in the literature examine the impact of antipiracy interventions on legal consumption (e.g., Adermon and Liang 2014, Danaher et al. 2014a, Danaher and Smith 2014. Aguiar et. al. 2018). However, our study is unique in several ways. Notably, we are the first to study the specific impact of piracy website blocking on

[1] As is customary in the economics and information systems literature, we use the terms filesharing and piracy interchangeably. When we use these terms we are referring collectively to all of the major forms of Internet media piracy including BitTorrent and other peer-to-peer protocols, direct cyberlocker downloads, and illegal streaming sites.

legal consumption.[2] As well, prior studies on antipiracy enforcements have each focused on just one action – we are the first to study the impact of several interventions of the same type but of varying strength, showing that weak interventions may have no impact on legal consumption while stronger interventions of the same nature can have a meaningful impact. Finally, ours is the first study of which we are aware yield insight into which pirates are most affected by antipiracy measures.

Evaluating the impact of website blocking on consumer behavior is important from a policy perspective. Unlike shutting down entire sites (such as the shutdown of Megaupload.com analyzed in Danaher and Smith 2014), website blocking is a strategy whereby governments or courts order Internet Service Providers within a country to not resolve domain names pertaining to a website that has been shown to facilitate illegal copyright infringement. This could include piracy cyberlockers, BitTorrent tracker sites (which do not host actual content but rather index the "tracker" files that filesharers require in order to download a media file through the BitTorrent protocol), or unauthorized media streaming sites.[3] As a legal matter, ISP-level blocking is easier to implement than full site shutdowns, and it has gained wide use in recent years as an anti-piracy strategy.[4] It is currently being debated as potential policy in Canada. Some countries, such as Australia, have explicitly stated that the effectiveness of website blocking will be

[2] Although one study exists on the blocking of The Pirate Bay (Poort et al. 2014), this study looks only at the effect of blocking on total piracy levels and does not explore whether legal channels benefitted from the block or whether VPNs were used to circumvent it.
[3] To be specific, it is more common that blocks are ordered against sites that link to or stream from cyberlocker content, rather than blocking the cyberlockers themselves.
[4] These countries include Australia, Argentina, Austria, Belgium, Chile, Denmark, Finland, France, Germany, Greece, Iceland, India, Indonesia, Ireland, Italy, Malaysia, Norway, Portugal, Russia, Saudi Arabia, Singapore, South Korea, Spain, Turkey, and the UK. (http://www2.itif.org/2016-website-blocking.pdf)

reviewed to determine whether policies should be changed, and our research provides guidance to these countries as to whether and how website blocking can be effective.[5]

Website blocks may also have a different impact from complete site shutdowns (e.g., Megaupload) because with a website block the content is still available on the servers of the blocked sites and there are a number of ways in which consumers and suppliers of pirated content may circumvent the block to obtain access to the infringing content. This leaves consumers with a choice between (a) finding ways to circumvent the blocks, (b) finding other sites to access pirated content, (c) increasing their use of legal channels, or (d) simply decreasing their consumption of the media in question. By studying three different examples of website blocking, our data allow us to gain insight into the conditions where website blocking is and is not effective at changing consumer behavior, something we have not seen in any of the prior work on antipiracy interventions. Further, by using individual level data rather than industry sales data, we are also able to evaluate which types of consumers are most impacted by website blocking.

In this paper we study three specific periods of website blocking orders granted by the UK High Court. The first was directed to The Pirate Bay in May 2012, the second was directed to 19 different major filesharing websites during October and November 2013,[6] and the third was directed to 52 remaining filesharing websites in November 2014. For the first two waves, we obtain panel data on aggregate groups of consumers surrounding each block, allowing us to determine the differential effects of blocking one major piracy site verses blocking many. For the third wave of blocks, we were able to obtain individual level panel data, allowing us to verify

---

[5] https://www.theregister.co.uk/2017/01/30/australia_to_review_effectiveness_of_isps_copyrightdefending_website_blocks/

[6] Actually, 28 sites were ordered blocked during this period of time. However, 9 of them were music-only piracy sites, and this paper focuses on video content, which was accessible through only 19 of these sites. Thus, from this point on we will refer to the 19 site blocks in October-November 2013.

some of the assumptions we made in studying the first two waves of blocks, and also allowing us to gain greater insight into the specific pirates who were most affected by the blocks.

These data show that blocking The Pirate Bay, one of the largest BitTorrent sites in the UK, caused only a small decrease in total piracy and no increase in the adoption of legal distribution services for digital movies and television. The data suggest that former Pirate Bay users merely switched to unblocked "proxy" sites that mirrored the contents of The Pirate Bay or dispersed to other filesharing websites to consume media illegally.

However, our data suggest that when 19 major piracy websites were simultaneously blocked in October-November 2013, there was a strong decrease in total piracy levels that caused users of the blocked sites to increase their usage of paid legal streaming sites by 11%. Later, the blocking of 52 remaining piracy sites in 2014 had a significant and similar impact on legal consumption. Together our results show website blocking may have a significant impact on legal consumption when multiple sites are blocked at once and when legal digital services are well-developed and convenient. Notably, more moderate media consumers are most impacted by the blocks, as opposed to heavy pirates or heavy users of legal channels. We discuss the explanations for these results and their policy implications in the conclusion of this paper.

2. Background on the Film Industry and Website Blocking

The film industry is a significant force in the world economy, with $38.6 billion in total theatrical revenue in 2016.[7] However, the advent of the BitTorrent filesharing protocol in 2003 led to a rapid spread of Internet movie piracy, and several studies (cited and discussed in section III) have causally linked this widespread piracy with significant lost revenues in motion picture revenue in all major sales channels.

[7] http://variety.com/2017/film/news/box-office-record-china-1202013961/

The industry has reacted to this threat by changing their distribution strategies in a variety of ways. For example, Danaher and Waldfogel (2012) show that since the advent of BitTorrent, movie studios have steadily decreased the windows between the US box office premiere of a movie and the international premieres. Similarly, Danaher et al. (2010) and Danaher et al. (2015) demonstrate that making content available on legal digital channels, such as iTunes and Hulu, can reduce the incidence of piracy for that content as some consumers switch from piracy to legal consumption. Zhang (2016) demonstrates that music labels increased legal song and album sales by removing Digital Rights Management and making digital download content more convenient. In addition to changing their business strategies in an attempt to make legal consumption more attractive than piracy, the film and television industries have also attempted to make pirated content less attractive than legal consumption by supporting various government antipiracy interventions such as the shutdown of Megaupload.com and Megavideo.com.

One of the more common antipiracy methods in recent times has been piracy website blocking, a strategy that has been attempted in over 25 countries to date.[8] For example, the UK has used website blocking to fight piracy since October 2011 when British Telecom and five other UK ISPs[9] were ordered by the High Court to block their customers from accessing Newzbin2, an indexing site for binary files posted to the Usenet. Following the Newzbin2 precedent, as of April 2015, over 125 copyright infringing sites were subject to court-ordered blocks in the UK.

Website blocking of this sort may be an attractive alternative strategy to graduated response laws and site seizures because, unlike graduated response laws it does not involve the legal and regulatory overhead necessary to adjudicate copyright claims against individuals, and

[8] http://www2.itif.org/2016-website-blocking.pdf, page 12.
[9] Specifically the ISPs Everything Everywhere, Sky, TalkTalk, Telefónica and Virgin Media.

unlike site seizures it does not involve cross-country cooperation for non-domestic websites. Instead, website blocking involves implementing requirements for domestic ISPs to not resolve domain names that have been shown to facilitate access to copyright infringing content.

Our present analysis concerns three waves of UK blocks that occurred in 2012, 2013 and 2014, respectively. Specifically, in April 2012 five major UK ISPs were ordered by the court to block access to The Pirate Bay, a major website for indexing the tracker files necessary to gain access to pirated media files through BitTorrent.[10] The Pirate Bay reportedly had 3.7 million users in the UK, and reportedly made about $3 million in October 2011 alone from advertising revenues.[11] Later, in October and November 2013, these five ISPs were ordered to block access to 19 additional piracy websites that provided access to copyrighted video content. Finally, in November 2014 they were again ordered to block access to a total of 52 additional piracy sites.

These orders, as well as other instances of mandated piracy website blocking around the world, have been met with controversy, as some claimed that this was opening the door to censorship of content on the Internet. This paper does not attempt to evaluate such claims. Rather, our purpose is to understand the impact of different levels of piracy website blocking on user behavior, and to extend the prior literature on antipiracy efforts by using a granular consumer level dataset to determine which consumers respond most heavily to blocks.

As we noted above, from a theoretical perspective, website blocking may have a different impact than site seizures because, given that the site is still operational and "connected" to the Internet and that the hosted content is still available, sophisticated users are able to find ways around the ISP-level block through the use of Virtual Private Network services or proxy server sites. For example, if a court orders an ISP to block access to a particular domain, say ThePi-

[10] British Telecomm, the sixth major ISP, was subsequently ordered to block The Pirate Bay in June 2012.
[11] http://www.theguardian.com/technology/2012/apr/30/british-isps-block-pirate-bay

rateBay.com, operators of the blocked website may set up a "proxy server" at a different domain that links users to the same content on the blocked site—for example, ThePirateBay.se. Even if the ISPs are ordered to block all future incarnations of the site in question (as is the case in the UK), there may still be some time between the introduction of a new domain and the ISPs recognition of it as a proxy to a blocked site. Thus, website blocking has been compared to the game "whac-a-mole,"[12] implying that it will be ineffective at increasing legal consumption as authorities or ISPs will be unable to keep up with agile piracy websites that are able to move domains and set up proxy servers more quickly than authorities can order those domains blocked.[13] In addition, because website blocking only impacts ISPs in the country in which the block was ordered, pirates can use Virtual Private Network (VPN) services to bypass the block by appearing to be connecting from a different country.

Nonetheless, investing the time and money involved in finding new domains (and knowing whether to trust them) or purchasing and learning how to use VPN services may come at some cost to the user. In this regard, prior research has demonstrated that actions that make legal content more attractive to users or that make illegal content less attractive to users can cause pirates to switch to paid content (see Danaher et al. 2014b for a summary of such studies). Thus, website blocking may be effective in changing consumer behavior if the potential workarounds have a sufficiently high level of inconvenience or sufficiently high learning costs.

Given the widespread adoption of website blocking and the theoretical ambiguity of their impact on consumer behavior, the question of whether and how piracy website blocking can be effective is of great practical and policy importance. We explore this below in the context of the

---

[12] See, for example, Nick Bilton's New York Times August 2012 editorial titled "Internet Pirates Will Always Win." (http://www.nytimes.com/2012/08/05/sunday-review/internet-pirates-will-always-win.html)

[13] See for example http://www.theguardian.com/world/2014/sep/10/blocking-copyright-infringing-websites-derided-whacking-moles

three UK blocking events described above: the blocking of The Pirate Bay in May 2012, the near-simultaneous blocking of nineteen major piracy sites in October-November 2013, and the blocking of fifty two piracy sites in November 2014.

3. Literature Review

Our study fits into several streams of the academic literature. First, there is a significant body of work on the relationship between piracy and sales of video content, including Rob and Waldfogel (2006), Smith and Telang (2010), Danaher et al. (2010), Zentner (2012), Danaher and Waldfogel (2012), and Ma et al. (2014).[14] The vast majority of this literature finds evidence of sales displacement caused by piracy across a variety of media types, including the consumption of television content, DVDs, and box office attendance.

Second, scholars in the information system and economics disciplines have begun to ask how government anti-piracy interventions can impact consumer behavior and revenues from legal media markets. Bhattacharjee et al. (2006) found that the RIAA's highly publicized law-suits against music pirates had a significant negative impact on the availability of pirated content, though a substantial amount of infringing content remained available even after the lawsuits. Danaher et al. (2014a) found that the French graduated response anti-piracy law "HADOPI" increased digital music sales for the major labels by around 25%. Danaher and Smith (2014) found that the shutdown of the popular piracy cyberlocker Megaupload.com increased digital movie revenues by 6-8%. Adermon and Liang (2014) demonstrated that the Swedish IPRED directive increased total music sales by 36% after being passed, but that sales reverted back to original levels within 6 months, possibly due to a lack of enforcement. Aguiar et. al. (2016) find that the shutdown of a single German piracy linking site, Kino.tv, had only a short-lived impact on levels

[14] We refer the interested reader to Danaher et al. (2014b) for a review of this literature.

of legal and illegal consumption. Peukert, Claussen, and Kretschmer (2017) find that the shutdown of Megaupload led to a decrease in sales of smaller, independent films. Finally, in perhaps the closest study to our own, Poort et al. (2014) used survey data to study the impact of the Dutch courts' order to Internet Service Providers (ISPs) to block Dutch access to The Pirate Bay and related sites, finding little impact on total piracy activity.

Our study contributes to the literature in several ways. First, we study website blocking rather than shutdowns or demand-side antipiracy policies, which is important given the widespread use of this strategy and the debate over its effectiveness. Our study extends the literature by analyzing the impact of website blocking on both piracy activity and visits to legal channels. Second, we study three different instances of website blocking whereas prior studies of antipiracy policies were largely confined to one instance. Importantly, our findings corroborate those of Poort et al. (2014) for the blocking of The Pirate Bay and Aguiar et al. (2018) for the shutdown of Kino.tv, but contrast them during the multi-site blocks, allowing us to draw inferences as to when and how website blocking may be effective and painting a more holistic picture than prior research has. Finally, by delving further into the individual level characteristics in our dataset, we provide the first evidence of which we are aware as to which pirates are most influenced by antipiracy policies like website blocking.

4. Data

We obtained data from an anonymous Internet consumer panel tracking company, which we refer to as PanelTrack in this paper.[15]

2012 and 2013 Blocks

[15] Despite their requirement to remain anonymous in our study, this tracking company is one of several leaders in the field and their data has been used in other peer reviewed papers to study the behavior of consumers on the Internet.

For the first two natural experiments, PanelTrack could not provide us data at the consumer level, and so provided us with aggregate data for groups of consumers defined based on observed behavior. We requested that PanelTrack define the groups by sorting consumers based on their pre-block usage of the blocked sites. For the study of The Pirate Bay block, consumers were sorted into ten different groups based on their total number of visits to The Pirate Bay during February and March 2012, given that the block occurred in May 2012. Similarly, for studying the blocking of 19 sites in November 2013, consumers were sorted into ten different groups based on their total visits to any of the 19 different blocked sites during August and September 2013. Thus, for each event that we study, we have ten different consumer segments, each of which we observe for seven months surrounding their respective blocks. For The Pirate Bay, we observe each segment from February through August 2012, and for the 19-site block, we observe each segment from August 2013 until February 2014. Consistent with industry practices (such as in determining TV audience numbers through ratings), PanelTrack scales their sample to respresent the population by multiplying each user's data by a scaling coefficient to project the sample to the population.

For each month-segment, we observe the following outcome variables: visits to the blocked sites, visits to mirrors of the blocked sites, visits to other unblocked torrent sites, visits to cyberlockers and streaming piracy sites, visits to VPN sites, and visits to paid legal streaming sites (such as Netflix and Viewster). Thus we can observe how each consumer segment changes their behaviors over time, both before and after the blocks.

2014 Blocks

Because PanelTrack only provided us with data on ten segments of consumers for the first two waves of blocks, there may be some concern about a lack of precision in our estimates.

We were able to obtain from PanelTrack actual unscaled individual level data to analyze the third wave of blocks. Specifically, we obtained monthly site visit data for a panel of 28,314 UK Internet users from August 2014 to February 2015, which includes three months before the blocks, the month of the blocks (November 2014), and three months after the blocks. Paneltrack again provided us with visits to blocked piracy sites and visits to unblocked piracy sites, but did not break down the unblocked piracy sites visits into cyberlockers vs. torrent sites. However because PanelTrack had updated some of their tracking methods, we were able to break down visits to legal sites into two categories: subscription streaming sites (like Netflix) and free ad-supported sites (like the BBC's iPlayer or Channel 5's "Demand 5"). [16]

Descriptive Statistics

For The Pirate Bay block, Table 1 provides mean visits to each site during February, March, and April (the months before the block) for each of the consumer segments.

TABLE 1 – DESCRIPTIVE STATISTICS (PIRATEBAY BLOCK)

Consumer Segment	% of Sample in Group	Pre-block Pirate Bay Visits Per User	Unblocked Torrent Sites (1000's)	Piracy Cyberlockers (1000's)	VPN Sites (1000's)	Paid Legal Streaming (1000's)
0	N/A	0	59,735	51,362	1,507	7,094
1	9%	1	1,167	949	27	79
2	12%	2.5	1,737	1,113	26	97
3	13%	5	1,570	718	19	52
4	11%	8.2	1,421	603	13	179
5	10%	13.4	2,122	687	14	189
6	11%	20.8	1,568	675	23	155
7	10%	36	1,367	554	19	85
8	10%	67.9	1,721	493	36	65
9	14%	226.3	2,907	559	21	96

[16] Unfortunately, the largest a la carte purchasing store is the iTunes store. But this is used most often as an application rather than a website, and many things other than intent to purchase may cause this app to open. So PanelTrack could not accurately track a la carte purchasing activity or even intent to purchase.

Note that we report the percent of the treated sample in each treated segment and that these are percentages are relatively equal across segments. However, PanelTrack indicated that around 90% of the total sample was in the control group, as most users were not accessing the Pirate Bay in the months before the block. Although piracy sites are substitutes for one another, the heaviest users of The Pirate Bay were also disproportionately heavy users of other torrent sites, even before the block. On the other hand, heavier users of The Pirate Bay (a torrent site) were lighter users of cyberlocker sites, which may imply that pirates tend to stick to a particular protocol/method for filesharing. Notably, the heaviest users of legal streaming sites are actually the mid-tier users of The Pirate Bay. Legal streaming sites were somewhat nascent at this time and had relatively low visits compared to piracy.

Table 2 reports the same statistics but for the consumer segments that we used to study the blocking of nineteen sites in October/November 2013.

TABLE 2 – DESCRIPTIVE STATISTICS (19-SITE BLOCK)

Consumer Segment	% of Sample in Group	Pre-block Visits/User to Blocked Sites	Unblocked Torrent Sites (1000's)	Piracy Cyberlockers (1000's)	VPN Sites (1000's)	Paid Legal Streaming (1000's)
0	NA	0	26,452	25,744	1,555	53,863
1	24%	1.0	589	490	9	1,488
2	13%	2.0	394	343	8	695
3	9%	3.0	454	368	3	771
4	7%	4.0	208	217	45	323
5	9%	5.4	272	542	29	479
6	10%	8.2	486	494	10	614
7	9%	13.2	651	673	11	607
8	10%	23.8	624	753	23	422
9	9%	66.4	719	1,927	23	956

Again we report the percent of treated users in each treated segment. In this case, the control segment makes up about 95% of the sample, indicating that most internet users in the sample were not users of these 19 piracy sites in the months before the blocks. The heavier users of the

19 blocked sites were also heavier users of other torrent sites. However, in this case, they were also heavier users of cyberlocker piracy sites, which may be because the 19 blocked sites included 8 non-torrent sites. Finally, visits to paid legal streaming sites are much higher for each of the segments than during the Pirate Bay blocks. This is likely due to the increased diffusion of the largest legal streaming sites (such as Netflix). One can compare totals between Table 1 and Table 2 despite the different sample sizes because the samples were projected to the population.

In the case of the 2014 blocking of 52-sites, although our data were obtained (and will be analyzed) at an individual level, in Table 3 we have broken the data into ten segments based on visits to the blocked during August and September of 2014. However, these data cannot easily be compared to Tables 2 and 3 because they were generated from raw, unscaled data (the figures are lower because they have not been projected to population values).

TABLE 3 – DESCRIPTIVE STATISTICS (52-SITE BLOCK)

Consumer Segment	Users in Segment	Pre-block Visits/User to Blocked Sites	Unblocked Piracy Site Visits	Legal Ad-Supported Visits	Legal Subscription Visits	VPN Visits
0	N/A	0	135,294	61,967	57,475	4,854
1	31%	1	28,903	6,610	7,692	390
2	14%	2	15,735	2,346	3,322	147
3	8%	3	13,209	2,286	1,871	166
4	6%	4	9,915	1,119	1,301	18
5	8%	5.4	12,821	1,590	1,968	229
6	9%	8.3	17,932	2,389	2,666	71
7	7%	13.2	21,635	1,999	3,446	524
8	9%	23.8	40,912	3,448	3,018	115
9	8%	78.6	96,732	3,178	2,496	28

In these data, the control group makes up about 90% of the sample, but we again report the percent of users in each segment as a percentage of total treated users. We note that visits to legal subscription sites has continued to grow, due in part to increased diffusion of Netflix. With our ability to separately identify visits to ad-supported free legal sites in these data, we are able to observe that such visits are similar in magnitude to subscription websites. Finally, because 52 different sites were blocked in 2014, the percent of users visiting at least one blocked site before the blocks is higher (only 90% in the control group compared to 95%).

Blocking Effectiveness

It is worth noting that the data from PanelTrack do show that all three blocking injunctions were effective in drastically reducing traffic to the blocked sites. Total visits to The Pirate Bay across all treated groups dropped by nearly 90% in the 3 months after the block as compared to the three months before.[17] Total visits by the treated groups to the 19 sites after the November blocks dropped by 83%, and total visits to the 53 sites blocked in 2014 dropped by 90% after the blocks were implemented.

One might ask why the drop was not 100% if the sites were blocked. One possible reason is that although we dropped the month the blocks were ordered from the analysis, we do not know exactly when each ISP implemented the blocks. Thus, it is possible that some Internet users had access to some of the blocked sites even in the early part of the post period. Additionally, to the degree that users circumvented the blocks by using VPNs or similar measures, their visits may continue to show up in our data as visits to the blocked sites even after the block was enforced. Finally, smaller ISP's in the UK were not required to participate in the blocks. Nonetheless, it is clear that the blocking injunctions caused major decreases in total visits to the blocked sites from all consumer segments, and thus these events constitute meaningful experiments with which to determine the impact of website blocking on consumer behavior.

[17] We ignore May itself, since the blocks occurred mid-month.

In the next section, we present our empirical model to analyze these experiments and determine their causal effect on consumer behavior.

5. Empirical Model and Results

5.1 2012 Pirate Bay Block

We first turn our attention to the blocking of The Pirate Bay in 2012. Recognizing that changes in outcome variables, such as use of paid streaming channels or use of other piracy sites, might change over time for reasons other than the block, we identify the causal impact of the block by comparing treated users (those who used The Pirate Bay before the block) with "control" users (those who did not).[18] We also divide treated users into nine different groups based on their number of visits to The Pirate Bay two months before the block. We call this variable for each user group the 'treatment intensity' variable as it serves as a proxy for the intensity of treatment that the block had on that group. Our identification relies on asking whether treated users change their visitation to paid legal viewing sites (or other potential outcomes) more than control users do, as well as examining how the pattern of visitation changes across different levels of treatment intensity. Users have of course quasi- "self selected" into these groups by choosing their on levels of usage of the blocked sites before they were blocked, and this choice is likely correlated with other static characteristics of the user, such as their usage of legal services. However, it is plausible that pre-block visits to blocked sites (treatment intensity) is not correlated with month to month trends in usage of other sites in the absence of a treatment, and we can partly test this identifying assumption by asking whether treatment intensity is correlated with month-to-month changes in our outcome variables prior to the blocks (we show that it is not).

[18] More precisely, we consider a user a control user if they did not use The Pirate Bay in the two months prior to the block, i.e. the users in group 0. Some of these users may have been rare users who then would have made some use of The Pirate Bay after the block (and thus were partly treated by the block). If this is the case, our results will be conservative since the control group may have been lightly impacted by the block.

To estimate the effect of blocking access to the Pirate Bay on visits to paid streaming sites, visits to other torrent sites, visits to cyberlocker piracy sites, and visits to VPN sites, we run the following model:

$$LnVisits_{jt} = \beta_0 + \beta_1 After_{\tau} + \beta_2 TreatIntensity_j * After_{\tau} + \mu_j + \varepsilon_{jt} \tag{1}$$

where $LnVisits_{jt}$ indicates the natural log of visits (to whatever category of sites we are examining) made by consumer group j during period t. $After_t$ is a dummy variable equal to one if the observations is for the post period. By including this variable, we control for differences between the pre-block period and the post-block period that would, on average, affect all segments evenly, such as any outside factors which increase or decrease the appeal of streaming services, VPN's, or piracy (for example, an increase in the quantity or quality of content offered on legal services). $Treatintensity_j$ indicates the number of visits that the average consumer in group j made to The Pirate Bay during March 2012. Finally, μ_j is a vector of group fixed effects and ε_{jt} is an idiosyncratic error term. In this model, β_2 is the variable of interest and, under the assumption that each group's trend after the block would have been uncorrelated with that group's treatment intensity, it indicates the causal impact of the block on visits to sites in the outcome group in question (e.g. paid legal streaming sites).

Table 4 below shows the results from models (1) and (2) for each of four outcome variables.

TABLE 4 – EFFECT OF 2012 PIRATE BAY BLOCK ON SITE VISITS

	Paid Streaming	Other Torrent	Cyberlockers	VPNs
After Block	-0.623+	-0.242*	-0.370+	-1.009**
	(0.302)	(0.042)	(0.168)	(0.362)
TreatIntensity * After Block	0.001	0.002*	0.000	0.010+
	(0.003)	-0.001	(0.002)	(0.005)
Constant	11.949*	14.688*	13.866*	10.377*
	(0.185)	(0.026)	(0.103)	(0.221)
Observations	20	20	20	20
Consumer groups	10	10	10	10
R-squared	0.853	0.994	0.951	0.822

Robust standard errors in parentheses
p-values calculated based on a t distribution with 8 degrees freedom (# groups - 2)
+ significant at 10%; ** significant at 5%; * significant at 1%

The first column of Table 4 examine the impact of the block on visits to paid legal streaming channels We note that β_2 is close to zero and statistically indistinguishable from zero. Thus we are unable to clearly detect any increase in usage of paid legal streaming sites. In contrast, β_2 for other torrent sites (the second column) is positive and statistically significant at a 99% confidence level, indicating an increase in the use of other unblocked torrent sites caused by the block. The coefficient in the third column for cyberlockers is effectively zero, indicating that users of The Pirate Bay did not turn to cyberlockers as a piracy alternative after the block. Finally, in the fourth column, β_2 is measured as 0.01 (at 90% confidence), indicating an increasing use of VPN's to circumvent the block. The constant in this column is low and so the change in levels for VPN usage may not be large, but the percent change in VPN use caused by the block is the largest of any of the outcome variables. Further inspection of the data reveals that this coefficient is heavily driven by the highest treatment intensity treatment group—people who used The Pirate

Bay over 200 times in the two months before the block. One might speculate that since these users had by far the strongest preferences for The Pirate Bay, they had the strongest incentive to find a way around the block rather than turning to legal sources for content or even other torrent sites.

In short, our regression results suggest that blocking The Pirate Bay in May 2012 caused users to gravitate toward other piracy sites or to use VPN's to circumvent the block, but we see no indication of an increase in paid legal sources of video content from blocking just The Pirate Bay.

5.2 19-Site Block in November 2013

We now turn our attention to the second event in our study: the blocking of 19 different video piracy websites within a 30-day period between October and November of 2013. We estimate the empirical model using OLS regression, and report the results in Table 5.

TABLE 5 – EFFECT OF 2013 BLOCKS ON SITE VISITS

	Paid Streaming	Other Torrent	Cyberlockers	VPNs
After Block	-0.027	-0.342*	-0.377*	0.201
	(0.061)	(0.093)	(0.102)	(0.263)
TreatIntensity * After Block	0.007**	0.006	-0.006	0.027**
	(0.003)	(0.004)	(0.004)	(0.011)
Constant	13.815*	13.436*	13.575*	9.997*
	(0.036)	(0.055)	(0.060)	(0.155)
Observations	20	20	20	20
Consumer groups	10	10	10	10
R-squared	0.993	0.983	0.981	0.915

Robust standard errors in parentheses
p-values calculated based on a t distribution with 8 degrees freedom (# groups - 2)
+ significant at 10%; ** significant at 5%; * significant at 1%

In this case, the results are noticeably different than those for The Pirate Bay block. First, there is a statistically significant increase in use of paid legal streaming sites, which if we assume that the segments should have trended similarly if not for the block, can be attributed to a shift toward legal channels caused by the block. Specifically, an individual who made 10 visits to the blocked sites in the month before the block increased his visits to legal streaming sites by 7% more than an individual who didn't use the blocked sites.

In the second and third columns we observe no statistically significant changes in visits to other unblocked torrent sites or to cyberlockers sites, although the point estimates indicate some shifting of blocked consumers to unblocked torrent sites. Finally, we observe that the blocks caused a statistically significant increase in use of VPN services, much like The Pirate Bay block did.

In claiming that it was the blocks that caused the increase in legal consumption, we rely on the identifying assumption that, in the absence of the blocks, there would have been no correlation between pre-block usage of the blocked sites and month-to-month trends in legal consumption. This assumption could be problematic given that users have self-selected into the ten groups of varying treatment intensity. We estimate the following more granular model in order to test this assumption:

$$LnVisits_{jt} = \beta_0 + \beta_1 month_t + \beta_2 TreatIntensity_j * month_t + \mu_j + \varepsilon_{jt} \qquad (2)$$

This model is very similar to (1) except that instead of t indexing either the pre- or post-blocking period, it indexes month. And instead of including an after-treatment dummy, it includes a vector of dummy variables for each month after the first month, and interacts these variables with treatment intensity. Here, β_2 is the coefficient of interest, as it indicates the degree to which treatment intensity moderates month-to-month trends in legal subscription visits. If our identifying assumption is correct, then we would expect β_2 to be statistically insignificant before the treatment. To examine this identifying assumption, we estimate model (2) and then in Figure 1 we plot β_2 below for each month, along with its 95% confidence interval.

FIGURE 1 – MODERATION OF LEGAL STREAMING BY TREATMENT INTENSITY DURING 2013 BLOCKS

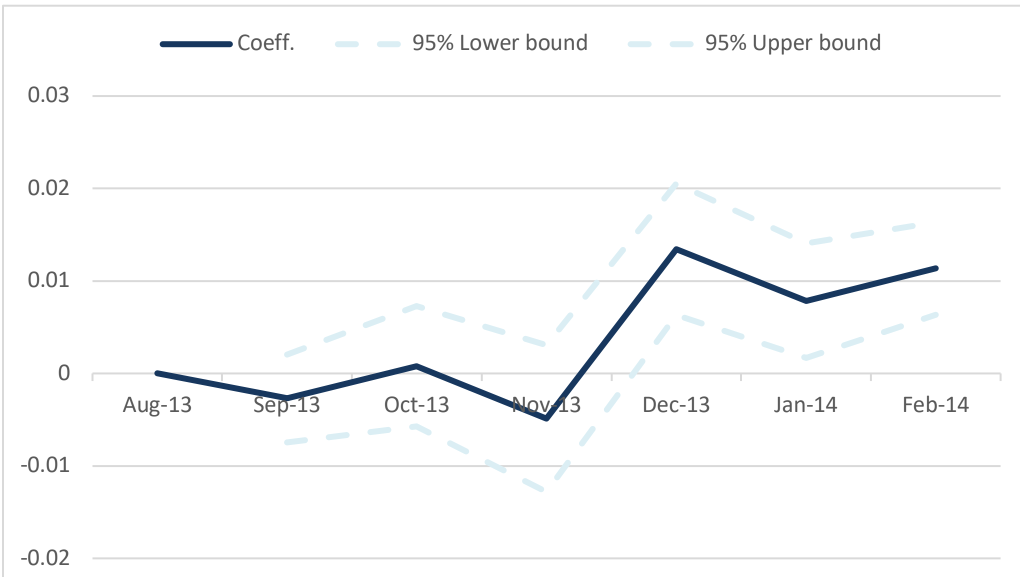

In Figure 1 there is almost no correlation between month-to-month trends in usage of legal subscription sites and pre-block visits to the blocked sites until after the blocks were implemented in November 2013. Afterward, there is a strong and persistent positive correlation. Although the blocks were ordered in November 2013, it appears as if there is no positive impact until December. The most plausible explanation for this for this delay is that it took the ISP's a few weeks to implement the blocks, leading to little or no impact in November (indeed, we still see a number of visits to blocked sites in November in the data). Thus our prior estimates, which included November in the post-treatment period, are likely conservative estimates of the true effect of the blocks. Nonetheless, the graph clearly shows that although pre-treatment trends are uncorrelated with treatment intensity, after the blocks the segments that were more heavily affect by the blocks increase their visits to paid legal subscription sites more than segments that were less affected.

Though we showed that pre-existing trends were not correlated with treatment intensity in the months before the blocks, one might worry with so few clusters (10 segments) whether our results are driven by one outlier segment. We re-estimated all of the regressions in Table 5 iteratively, each time dropping one segment from the data. All results held in sign and most held in significance, though in a few cases the significance fell to $p<0.1$, likely due to the loss of data.[19] In short, our results are robust to removal of any one group.

5.3 52-Site Block in November 2014

Finally, we turn our attention to the last wave of blocks in our study, the blocking of 52 of the largest remaining piracy sites in November 2014. Recall that for this wave of blocks, we were able to negotiate with PanelTrack to obtain raw individual level monthly visits to each group of sites we studied. As a result, instead of having ten different experimental groups, we can treat each individual as a unique experiment, asking whether individuals who used the blocked sites more before the blocks changed their behavior more than individuals who used them less or not at all. Because we are no longer aggregating individuals into groups, there will be more noise in the data, but now we will have over 28,000 unique observations before and after the blocks rather than just 10.

Having individual level data introduces several issues. First, 3,694 individuals in our sample did not consume any video content. In other words, they did not visit any blocked or unblocked piracy sites, nor did they visit any legal consumption sites. We drop such individuals

[19] We also performed this exercise for the 2012 Pirate Bay block with the same result.

because they cannot inform the question we are studying. Second, our data on visits are count data. When performing the analysis at the segment level, counts were high enough that they could reasonably be analyzed as continuous data. But at the individual level, visits are often low enough that OLS is unlikely to fit the data well. Even Poisson regression is unlikely to fit well due to overdispersion in visits to the outcome sites and the fact that the most common value for visits is zero. Thus we select negative binomial regression as the most appropriate specification to estimate our model.

We estimate model (1) on our individual dataset for the 2014 blocks, with two differences. First, in this case j indexes the individual rather than the segment. Second, as discussed, we estimate a negative binomial model rather than OLS on the log of the dependent variable.

TABLE 6 – EFFECT OF NOVEMBER 2014 BLOCKS ON SITE VISITS (INDIVIDUAL DATA)[20]

	Unblocked Piracy	VPNs	Legal Ad-supported	Legal Subscription
After Block	-1.588**	-0.933**	-1.023**	-0.975**
	(0.016)	(0.087)	(0.016)	(0.021)
TreatIntensity * After Block	-0.000	0.003	0.005**	0.006**
	(0.000)	(0.002)	(0.001)	(0.001)
Constant	-0.101**	-0.738**	-0.291**	-0.760**
	(0.019)	(0.170)	(0.023)	(0.029)
Observations	34,034	1,198	26,876	17,934
Individuals	17,017	599	13,438	8,967
Wald chi^2	10,936	117	3,956	2,241

robust standard errors in parentheses
+ significant at 10%; * significant at 5%; ** significant at 1%

Due to the link function used in negative binomial (and Poisson) regression, the coefficients in Table 6 can be interpreted as the change in the natural log of the dependent variable for a 1-unit increase in the independent variable. In the first column, we observe no increase in usage of unblocked piracy sites as a result of the blocks. In the second column, we observe that the blocks caused an increase in usage of VPN sites, though the effect is statistically insignificant, possibly due to the small sample size of 599 individuals. Relatively few individuals used VPN's at any point in our data.

In the third column and fourth columns we observe that the blocks caused users to increase their visits to legal ad-supported viewing sites and legal subscription sites. For each additional pre-block visit to the blocked sites, a user increased her visits to legal subscription sites by

[20] Note that negative binomial fixed effects models do not estimate a true fixed effect, and thus can technically estimate a coefficient for "TreatIntensity". All findings remain similar in sign and significance if we include this term in the model.

0.6 % after the blocks and her visits to legal ad-supported free streaming sites by 0.5%. Note that our ability to analyze these data at the individual level has increased the statistical precision in our estimates—the standard errors around these estimates are smaller than for the aggregated data from the prior two experiments and we can reject the null hypothesis that treatment intensity does not correlate with post-block changes in legal consumption at alpha=.01.[21] We note that these estimates of the impact on legal consumption are very similar to the one estimated by for the effect of the November 2013 blocks—thus a similar impact was achieved, though it was achieved by blocking over twice as many sites.

We again test the identifying assumption—that an individual's month to month changes in legal viewing would be uncorrelated with treatment intensity in the absence of the treatment—by estimating model (2) for both legal subscription visits and legal ad-supported visits, graphing the coefficient of interest in Figures 2 and 3 below.

[21] We also estimated each of these models using Poisson regression, which can estimate a true fixed effect. Results were the same in sign and significance, but we believe the negative binomial model produces more accurate estimates of the treatment effect.

FIGURE 2 – MODERATION OF LEGAL AD-SUPPORTED VIEWING BY TREATMENT INTENSITY DURING 2014 BLOCKS

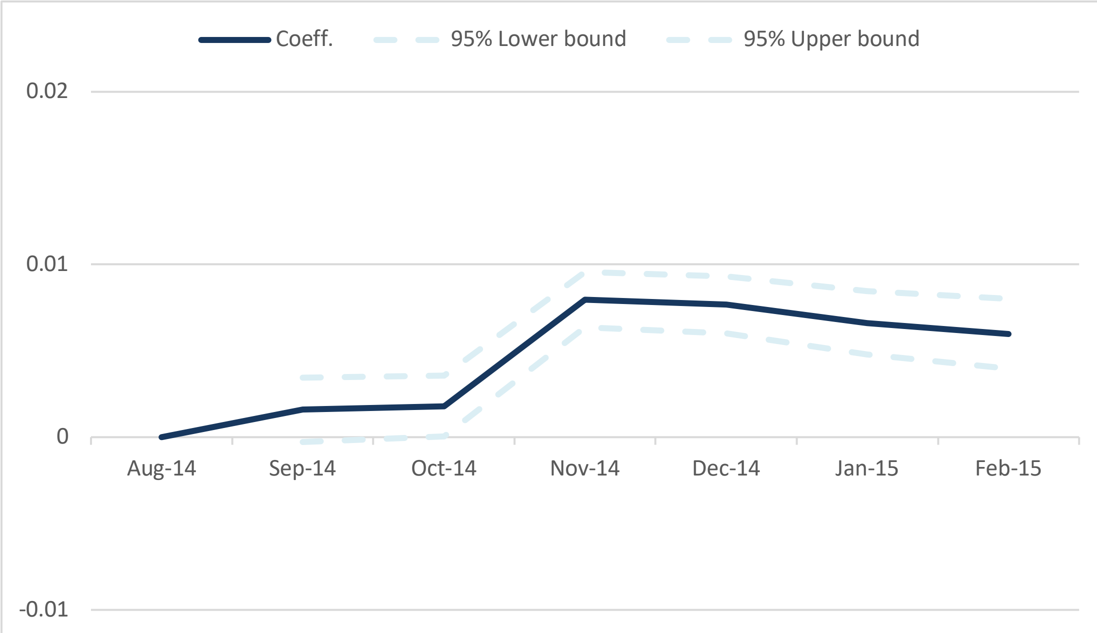

FIGURE 3 – MODERATION OF LEGAL STREAMING BY TREATMENT INTENSITY DURING 2014 BLOCKS

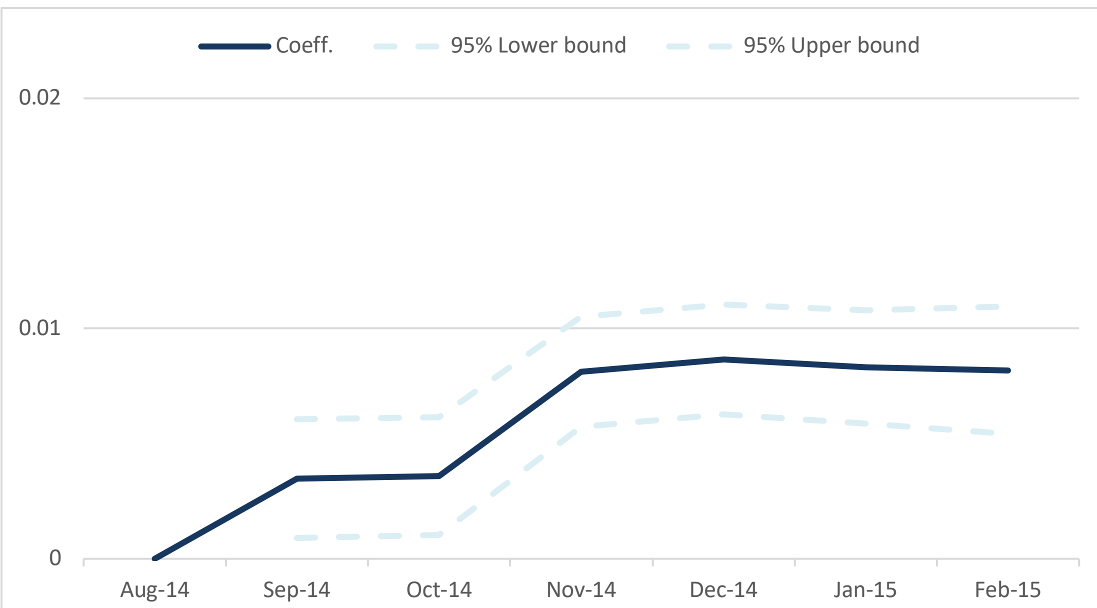

In Figure 2 it is clear that, prior to the blocks in November 2014, there was little to no correlation between pre-block visits to blocked sites and month-to-month changes in visits to ad-supported sites. Immediately following the blocks, we observe a positive and statistically significant correlation. This supports our identifying assumption that month-to-month viewing trends are not correlated with treatment intensity in the absence of a treatment. Figure 3 is slightly more complicated. It does appear as if heavier users of the blocked sites increased their usage of legal subscription sites more from August to September 2014 than did lighter users. However, no correlation was observed from September to October. In November, when the treatment occurred, there was a statistically significant correlation between the increase in legal subscription views and treatment intensity, and this correlation persevered for the following three months. We note that in the post period, after experiencing an increase in subscription streaming immediately following the blocks, heavier users of the blocked sites were not growing their subscription streaming any faster than lighter users. Thus we believe that the most logical conclusion is that, in spite of the one anomaly from August to September 2014, the 2014 blocks appear to have caused an increase in visits to subscription streaming sites.

Importantly, we find relatively similar results if we group the users into ten segments based on their various degrees of treatment intensity (i.e. the approach that we were forced to use for the first two waves due to limitations in the data provided by PanelTrack). This lends support to the notion that using the segment level data in the first two experiments yielded results similar to what we would have observed if we could have obtained individual level data for these periods.

5.4 Heterogenous Impact of Website Blocking

One benefit of having individual level data for the third wave of blocks is to verify the validity of the less granular approach used to analyze the 2012 and 2013 waves of blocks. Another benefit, however, is that it allows us to ask whether different types of individuals are impacted differently by the blocks. Unfortunately, we do not have demographic information on each individual in our study. However, we can take the data we do have on individuals before the blocks—total legal visits (ad-supported or subscription-based) and total piracy visits (blocked or unblocked sites)—and cluster users based on these pre-treatment variables. We ran k-means clustering on these data and found five clusters of individuals in our dataset based on their pre-treatment activity. These clusters are described in Table 7.

TABLE 7 – AVERAGE VISITS PER USER IN EACH CLUSTER (PRE-2014 BLOCKS)

Cluster	# of Individuals	Blocked Piracy	Unblocked Piracy	Legal Ad-Supported	Legal Subscription
1	161	4.1	19.4	100.9	117.7
2	1,314	28.2	122.3	5.3	4.0
3	151	125.2	477.0	11.5	5.0
4	22,004	1.2	6.6	1.6	1.3
5	990	2.7	13.1	26.6	33.0

The clusters in Table 7 are relatively easy to interpret. Cluster 1 is a small cluster of individuals who make very heavy use of legal channels with much smaller use of piracy. Clusters 2 and 3 are heavy pirates (both blocked and unblocked sites) with cluster 3 being the most extreme – both clusters make some limited use of legal channels. Cluster 4 is by far the largest cluster and are the lightest consumers of content, consuming less content through all of the channels. Cluster 5 includes consumers who consume more content than cluster 4, but who tend to favor legal channels over illegal ones.

Having identified these clusters, we estimate model (1) again on each cluster separately, using the individual level data and negative binomial regression. We consider legal ad supported visits as the outcome variable in Table 8, and legal subscription visits as the outcome in Table 9.

TABLE 8 – IMPACT OF 2014 BLOCKS ON LEGAL AD-SUPPORTED VIEWING, BY CLUSTER

Cluster	1	2	3	4	5
After Block	-1.060**	-0.478**	0.045	-1.099**	-1.198**
	(0.132)	(0.064)	(0.177)	(0.019)	(0.058)
TreatIntensity * After Block	-0.004	0.002	-0.001	0.030**	0.005
	(0.009)	(0.001)	(0.001)	(0.003)	(0.007)
Constant	-0.049	-0.393**	-0.234	-0.325**	0.040
	(0.155)	(0.069)	(0.173)	(0.027)	(0.070)
Observations	292	1,852	238	22,790	1,704
Individuals	146	926	119	11,395	852
Wald chi^2	72	72	3	3,462	468

robust standard errors in parentheses
+ significant at 10%; * significant at 5%; ** significant at 1%

TABLE 9 – IMPACT OF 2014 BLOCKS ON LEGAL SUBSCRIPTION VIEWING, BY CLUSTER

Cluster	1	2	3	4	5
After Block	-1.272**	-0.214*	0.490+	-1.041**	-1.319**
	(0.138)	(0.090)	(0.276)	(0.024)	(0.066)
TreatIntensity * After Block	-0.007	0.001	-0.002	0.031**	0.007
	(0.009)	(0.002)	(0.002)	(0.005)	(0.006)
Constant	0.278+	-1.197**	-1.280**	-0.813**	-0.214**
	(0.168)	(0.099)	(0.241)	(0.034)	(0.076)
Observations	268	1,158	140	14,860	1,508
Individuals	134	579	70	7,430	754
Wald chi^2	98	7	3	1,918	429

robust standard errors in parentheses
+ significant at 10%; * significant at 5%; ** significant at 1%

The interaction term is the term of interest in both Table 8 and Table 9. First, we note that our prior estimates of a positive impact of the blocks on legal subscription viewing and legal ad supported viewing are driven entirely by cluster 4 (no other clusters exhibit a positive and statistically significant effect). Of course, it is not surprising that the results we estimated for the total sample are similar in sign and significance to those we estimated for cluster 4, as cluster 4 is comprised of a large percent of the sample. But it is interesting that no other cluster exhibits a causal increase in legal viewing from the blocks. Cluster 4 are the lightest media consumers—in other words, neither heavy pirates nor heavy legal consumers of media increased their legal consumption as a result of piracy website blocking. Not only are the estimates of the impact of the blocks for the other four clusters statistically insignificant, but they are also smaller by an order of magnitude. Said another way, our study shows that the users whose behavior is mostly likely to be changed by website blocking are those users who are not strongly committed to a mode of consumption—either legal or piracy.

Thus we observe that while the majority of Internet users consume only a moderate amount of video content and can be influenced to turn toward legal channels by blocking piracy websites, neither the heaviest users of pirate sites nor the heaviest users of legal sites do not appear to turn toward legal channels when piracy sites are blocked. Nonetheless, because moderate media consumers make up the bulk of the population, blocking multiple sites at once appears to increase legal consumption for the population of media consumers. In the final section, we calculate the magnitude of these effects.

6. Discussion

While the use of supply side antipiracy actions has increased in recent years as a tool in the fight against intellectual property theft, there are few studies that have empirically analyzed their effectiveness in changing user behavior, and none that we are aware of that have analyzed the impact of actions with heterogeneous intensity, or the impact of blocks across heterogeneous groups of users. Moreover, the most popular form of supply side antipiracy, website blocking, has gone largely unstudied. Our study seeks to fill this gap in the literature, using data provided by a panel tracking company to analyze the impact of three website blocking events in the UK: The blocking of The Pirate Bay in May 2012, the blocking of 19 additional sites in October and November of 2013, and the blocking of another 52 sites in November 2014

Our results suggest that blocking The Pirate Bay in May 2012 led to an increase in the usage of other unblocked torrent sites and of VPN sites, causing only a small decrease in total piracy and having no statistical impact on legal entertainment sites. However, the blocking of 19 different sites in November of 2013, and 52 different sites in November 2014 led to significant decreases in total piracy and caused a statistically significant increase in the usage of legal streaming sites. This pattern between pre-block piracy visits and post-block changes in legal visits was not observed in the months before the block, suggesting that this impact is causally related to the block.

A natural question to ask is: what was the total impact of the blocks on the use of legal streaming sites. In 2013, for each segment, we can also start with the total post-block visits to paid legal streaming sites for that group and determine what the counterfactual visits would have been if the blocks had not happened (which is equivalent to estimating what would have happened if the treatment intensity variable were 0). The difference between observed visits and counterfactual visits in the post-period is the causal impact of the block on visits for that segment. This leads to the conclusion that treated users (i.e., users who made at least one visit to the blocked sites before they were blocked) increased their visits to subscription streaming sites by

11% as a causal effect of the blocks. These results were significant at the 95% confidence level. We did not have data on ad-supported site visits for this wave of blocks.

Turning to the 2014 blocks, we can use the same approach to make back-of-the-envelope estimates of the total effect on legal consumption, only we perform the analysis at the individual level rather than the segment level. Doing this, we find that this wave of blocks caused treated users to increase their visits to free legal ad-supported sites by 11.5% and legal paid subscription sites by 10%. These estimates were significant at the 99% confidence level.[22]

Notably, the blocks did cause some users to increase visits to VPN's during all waves of blocks. However, VPN use was not common in general, and this was far from the dominant effect of waves 2 and 3 of blocking.

One might ask why the second two waves of blocks caused an increase in paid legal consumption but the Pirate Bay block in May 2012 did not. One potential explanation is that legal services were more developed by 2013 than they were in 2012. However, this cannot explain the entire difference. Even in 2012 there were a number of consumers using legal subscription services in addition to piracy, and these users did not increase their visits to these services as a result of the Pirate Bay block. Moreover, subscription services were more widely known in 2014 than 2013, but the 2014 blocks caused no larger increase in use of these sites than 2013. Instead, we offer another explanation—when only one site is blocked, even a relatively popular one, it remains relatively easy for former users of that site to find a reliable substitute piracy site. But

[22] This larger significance is expected given that PanelTrack was able to provide us with individual level data rather than aggregated consumer segments. It may however seem counterintuitive that the percent increase for ad supported streaming was slightly larger than that for subscription streaming, despite the fact that the regression coefficient of interest for subscription sites was slightly higher than that for ad-supported. This is explained by the fact that it was more common for heavy users of ad-supported sites to be heavy users of the blocked sites (and thus get a larger legal increase) than it was for heavy users of subscription sites. On average, pirates in our data prefer free (ad supported) legal sites to paid (subscription) legal sites.

when 19 or more sites are blocked, finding an alternative requires substantially more effort. This explanation is consistent with a theoretical model developed by Dey et. al (2017), who demonstrated that supply-side antipiracy enforcement will only be effective at increasing legal consumption if it sufficiently increases the cost of piracy.

This result is important. Several prior studies (cited in our literature review) have shown that supply-side antipiracy enforcement efforts that have blocked or shut down a single site have not been effective at decreasing the supply of pirated content or increasing legal consumption. Our results are consistent with these findings, but also show that when enough sites are blocked, it can sufficiently increase the cost of piracy such that some users—especially moderate media consumers—will shift from piracy to legal consumption. In addition, our results imply that piracy does indeed displace usage of legal paid streaming sites and of free ad-supported sites, despite the relative convenience and low cost of such sites.

There are several limitations to this study. First, we were only able to study paid legal subscription streaming in the first and second wave of blocks as well as ad-supported streaming in the third wave. Limitations in PanelTrack's data did not allow us to study other forms of legal consumption during the timeframe covered by these data. Second, because the blocks were not implemented by all ISP's and may have been implemented by participating ISP's some time into the first month of our "after" periods, our results may underestimate the true effect of website blocking on legal consumption. Third, because the size of the consistent panel (users tracked for the entire time of the study) would decrease exponentially when increasing the months of the study, we were only able to observe the first three months after each wave of blocks. Although we observed persistent effects, it remains possible that increases in legal consumption caused by the blocks fade over time as consumers identify alternate piracy sites. Finally, we are not able to fully estimate the social welfare implications of these blocks because our data do not allow us to estimate the value of the impacts (just their relative sizes) or the costs of implementing the blocks, and because we have no data on the impact of increased profitability on industry output. Future work should focus on these issues to obtain a better understanding of the broader impacts of site blocking and other anti-piracy measures.

References

Adermon, A., C.Y. Liang. 2014. Piracy and Music Sales: The Effect of an Anti-Piracy Law. *Journal of Economics Behavior and Organization*. 105, pp 90-106.

Aguiar L, Peukert C., and Claussen J. 2018. Catch Me If You Can: Effectiveness and Consequences of Online Copyright Enforcement. Forthcoming, *Information Systems Research.*

Bertrand, M, E. Duflo, S. Mullainathan. 2004. How Much Should We Trust Differences-in-Differences Estimates? *The Quarterly Journal of Economics*, 119(1) 249-275.

Bhattacharjee, S., R. Gopal, K. Lertwachara, J. Marsden. 2006. Impact of Legal Threats on Online Music Sharing Activity: An Analysis of Music Industry Legal Actions. Journal of Law and Economics, 49(1) 91–114.

Danaher, B., S. Dhanasobhon, M.D. Smith, R. Telang. 2010. Converting Pirates without Cannibalizing Purchasers: The Impact of Digital Distribution on Physical Sales and Internet Piracy. *Marketing Science*, 29(6) 1138-1151.

Danaher, B., J. Waldfogel. 2012. Reel Piracy: The Effect of Online Movie Piracy on Film Box Office Sales. Working Paper. Wellesley College. Boston, Massachusetts.

Danaher, Brett, Michael D. Smith. 2014. Gone in 60 Seconds: The Impact of the Megaupload Shutdown on Movie Sales. *International Journal of Industrial Organization*. 33 1-8.

Danaher, Brett, Michael D. Smith, Rahul Telang, Siwen Chen. 2014a. The Effect of Graduated Response Anti-Piracy Laws on Music Sales: Evidence from an Event Study in France. *Journal of Industrial Economics*. 62(3) 541-553.

Danaher, Brett, Michael D. Smith, Rahul Telang. 2014b. Piracy and Copyright Enforcement Mechanisms, Lerner and Stern, eds. Innovation Policy and the Economy, Volume 14, Chapter 2 (pp. 31-67), National Bureau of Economic Research, University of Chicago Press, Chicago, Illinois.

Danaher, Brett, Samita Dhanasobhon, Michael D. Smith, Rahul Telang. 2015. Understanding Media Markets in the Digital Age: Economics and Methodology, Greenstein, Goldfarb, and Tucker, eds. *Economics of Digitization: An Agenda*. Chapter 13 (pp. 385-406). University of Chicago Press, Chicago, Illinois.

Dey, Debabrata, Antino Kim, Atanu Lahiri. 2016. Online Piracy and the 'Longer Arm' of Enforcement. Forthcoming, *Management Science*.

IFPI. 2010. IFPI Response to Commission Green Paper on Creative and Cultural Industries. July 2010.

Ma, Liye, Alan Montgomery, Michael D. Smith, Param Singh. 2014. The Effect of Pre-Release Movie Piracy on Box Office Revenue. *Information Systems Research*. **25**(3) 590-603.

Peukert C., Claussen J, and Kretschmer, T. 2017. Piracy and Movie Revenues: Evidence from Megaupload." *International Journal of Industrial Organization* 52, 188-215

Poort, J., Leenheer, J, Ham, JVD, and Dumitru, C. 2014. Baywatch: Two Approaches to Measure the Effects of Blocking Access to The Pirate Bay. *Telecommunications Policy*. 38(1) 383-392

Rob, R., J. Waldfogel. 2006. Piracy on the High C's: Music Downloading, Sales Displacement, and Social Welfare in a Sample of College Students. *Journal of Law and Economics*, 49(1) 29-62

Smith, M.D., R. Telang. 2010. Piracy or Promotion? The impact of broadband Internet penetration on DVD sales. *Information Economics and Policy*. 22 289-298.

Telang, R. and Waldfogel, J. 2014. "Piracy and New Production Creation: A Bollywood Story." Forthcoming, *Information Economics and Policy*.

Zentner, A. 2010. Measuring the Impact of File Sharing on the Movie Industry: An Empirical Analysis Using a Panel of Countries. Working Paper. University of Texas at Dallas, Dallas, Texas. (Available from http://ssrn.com/abstract=1792615)

Zentner, A. 2012. Measuring the Impact of File Sharing on the Movie Industry: An Empirical Analysis Using a Panel of Countries. Working Paper, University of Texas at Dallas, Dallas, Texas. (Available from http://ssrn.com/abstract=1792615)

Zhang, Laurina. 2016. Intellectual Property Strategy and the Long Tail: Evidence from the Recorded Music Industry. *Management Science*. Articles in Advance. (Available from https://doi.org/10.1287/mnsc.2016.2562)