

THE EFFECT OF PIRACY WEBSITE BLOCKING ON CONSUMER BEHAVIOR

Brett Danaher[‡]
bdanaher@wellesley.edu

Michael D. Smith[†]
mds@andrew.cmu.edu

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

This Version: May 2015

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 this paper at the December 2014 Workshop on Information Systems and Economics and thank attendees for their helpful feedback. All findings and errors are entirely our own.

[‡] Department of Economics, Wellesley College, Wellesley, MA, 02481

[†] 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 participants in digital media markets. In this paper we seek to 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 two court-ordered events affecting consumers in the UK: The blocking order directed at The Pirate Bay in May 2012, and blocking orders directed at 19 major piracy sites in October and November 2013.

Our results show that blocking The Pirate Bay had little impact on consumption through legal channels — instead, consumers seemed to turn to other piracy sites, Pirate Bay “mirror” sites, or Virtual Private Networks that allowed them to circumvent the block. In contrast, blocking 19 different major piracy sites caused users of those sites to increase their usage of paid legal streaming sites such as Netflix by 12% on average. The lightest users of the blocked sites (and thus the users least affected by the blocks, other than the control group) increased their clicks on paid streaming sites by 3.5% while the heaviest users of the blocked sites increased their paid streaming clicks by 23.6%, strengthening the causal interpretation of the results.

Our results suggest that website blocking requires persistent blocking of a number of piracy sites in order to effectively migrate pirates to legal channels.

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 technologies¹ 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. 2014 for a review of this literature). The recent literature suggests that in the film industry, 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). 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 legislations. Several papers exist that examine the impact of antipiracy interventions on legal consumption (Adermon and Liang 2014, Danaher et. al 2014, Danaher and Smith 2014), but each of these use aggregate market data and thus cannot capture insights into how consumers choose whether to circumvent such legislation or why they increase their legal consumption if they do. Our study is the first of which we are aware to use a consumer-level dataset to understand the various ways in which Internet users react to instances of anti-piracy legislation and provide insights into why.

¹ As is customary in the economics and information systems literature, we use the terms filesharing and piracy interchangeably and without intending to impose any connotation. As well, when we say filesharing, we are referring to all of the major protocols of Internet media piracy – peer-to-peer programs, BitTorrent, using cyberlockers for piracy, and illegal streaming sites.

Specifically, we analyze the effect of piracy website blocking, a type of legislation that has not yet been addressed in the literature.² Unlike shutting down entire sites (such as the shutdown of Megaupload.com), website blocking is a strategy whereby governments or courts order Internet Service Providers within a country to simply block users' access 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 However, the effectiveness of website blocking may be different from a complete site shutdown, like Megaupload, because 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. Thus, website blocks are empirically interesting to study as consumers have a choice between finding ways to circumvent the blocks, finding other sites to access pirated content, increasing their use of legal channels, or simply decreasing their consumption of the media in question. Our data allow us to gain insight into how consumers make this choice, and thus what types of antipiracy interventions are more likely to increase legal consumption.

In this paper we study two specific periods of website blocking orders granted by the UK High Court: the first directed to The Pirate Bay in May 2012, and second directed to nineteen different major filesharing websites during October and November 2013.⁴ Our analysis uses a

² Though one study exists on the blocking of The Pirate Bay (Poort et al. 2014), this study looks only at the effect on total piracy levels and does not explore how consumer behavior changed following the block nor whether legal channels benefitted from the block.

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 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.

novel datasets that allow us to study these events with arguably greater precision than prior work. Unlike prior studies that used market data to examine the impact of anti-piracy interventions on sales alone, we obtained panel data on the actual behavior of a large group of Internet users in the UK. As such, we are able to determine the effect of these website blocks on not only legal purchasing activity, but also dispersion to other unblocked piracy websites and on the use of technologies that can circumvent the blocks.

These data show that the blocking of The Pirate Bay, one of the largest BitTorrent sites in the UK caused 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, unrelated filesharing websites to consume media illegally.⁵

However, our data suggest that when nineteen major piracy websites were simultaneously blocked in October-November 2013, the results were different. Though we observe that this caused some consumers to disperse to other piracy sites or to adopt technologies that allow circumvention of the block, we also find that these blocks caused users of the blocked sites to increase their usage of paid legal streaming sites by 12%. The lightest users of the blocked sites (and thus the users least affected by the blocks, other than the control group) increased their clicks on paid streaming sites by 3.5% while the heaviest users of the blocked sites increased their paid streaming clicks by 23.6%. Thus, our results show website blocking may have a signif-

⁵ In March 2015, Aguiar, Peukert, and Claussen reported a similar result in the context of shutting down a major piracy linking site in Germany, kino.to. Their methodology was similar to our own and their results are largely consistent with our Pirate Bay result, as they found that while prior users of kino.to decreased their total piracy to a degree, many of them substituted to other piracy sites and the event only caused a modest increase in usage of paid licensed services.

ificant impact on legal consumption when multiple sites are blocked at once. We discuss the explanations for and implications of these results in the conclusion of this paper.

Finally, because we observe data on behavior of groups of consumers rather than aggregate market data, we are able to shed light on which types of consumers are most responsive to anti-piracy interventions and discuss what this means for copyright reform.

The paper proceeds as follows: Section II presents a short background on the motion picture industry and website blocking. Section III describes academic literature related to our study. Section IV describes our data and section V outlines our empirical model and presents our results. Finally, section VI concludes our study with potential explanations for our findings and with implications for firms and policymakers.

2. Background on the Film Industry and Website Blocking

The film industry is a significant force in the world economy, with \$35.9 billion in total theatrical revenues in 2013.⁶ 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 the box office, home theater, and digital film markets.

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,

⁶ http://www.mpaa.org/wp-content/uploads/2014/03/MPAA-Theatrical-Market-Statistics-2013_032514-v2.pdf

can reduce the incidence of piracy for that content as some consumers switch from piracy to legal consumption. 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 anti-piracy interventions such as the shutdown of Megaupload.com and Megavideo.com.

Recently many governments and courts worldwide have adopted “website blocking” as an additional anti-piracy approach. For example, the UK has used website blocking to fight piracy since October 2011 when British Telecom and five other UK ISPs⁷ 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 because unlike graduated response laws it does not involve the legal and regulatory overhead necessary to adjudicate copyright claims against individuals, and unlike site seizures it does not involve cross-country cooperation for non-domestic websites. Instead, website blocking involves requirements for domestic ISPs to block access to domain names or IP addresses that have been shown to facilitate access to copyright infringing content.

Our present analysis concerns UK blocks that occurred in 2012 and 2013. 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.⁸ The Pirate Bay reportedly had 3.7 million users in the UK, and the record labels claimed that this site made about \$3 million in October 2011 alone from advertising reve-

⁷ Specifically the ISPs Everything Everywhere, Sky, TalkTalk, Telefónica and Virgin Media.

⁸ BT was subsequently ordered to block The Pirate Bay in June 2012.

nues.⁹ Later, in October and November 2013, these five ISPs were ordered to block access to 19 piracy websites that provided access to copyrighted video content.

These orders, as well as other instances of mandated piracy website blocking around the world, were initially met with some degree of controversy, as detractors 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 piracy website blocking in a more complete way than prior studies on anti-piracy interventions through the use of a more granular consumer level dataset.

This impact is theoretically ambiguous. Website blocking may be less effective than site seizures or takedowns because, given that the site is still operational and “connected” to the Internet and that the hosted content is still available, technically sophisticated users may be able to find ways around the ISP-level block through Virtual Private Network services or through proxy server sites. For example, if a court orders an ISP to block access to a particular domain, say ThePirateBay.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,”¹⁰ 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

⁹ <http://www.theguardian.com/technology/2012/apr/30/british-isps-block-pirate-bay>

¹⁰ 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>)

or more quickly than ISPs can detect and then carry out orders that require them to block access to any future proxies.¹¹ In addition, because website blocking only blocks users from 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 make illegal content less attractive to users can convert pirates into paid consumers (see Danaher et al. 2014 for a summary of such studies). Thus, in spite of potential workarounds, website blocking may still be effective in changing consumer behavior if potential workarounds have a sufficient level of inconvenience or sufficiently high learning costs. We explore this hypothesis below in the context of the two UK blocking events described above: the blocking of The Pirate Bay in May 2012, and the near-simultaneous blocking of nineteen major piracy sites in October-November 2013.

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).¹² 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.

¹¹ See for example <http://www.theguardian.com/world/2014/sep/10/blocking-copyright-infringing-websites-derided-whacking-moles> and <http://www.nytimes.com/2012/08/05/sunday-review/internet-pirates-will-always-win.html>

¹² We refer the interested reader to Danaher et al. (2014) for a thorough review of this literature.

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. (2008) found that the RIAA's highly publicized lawsuits 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. (2014) 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. In a yet unpublished study, Peukert, Claussen, and Kretschmer (2014) suggest 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. They find little impact on total piracy activity.

Our study contributes to the literature in several ways. First, we look at the consumer behavior at the micro segment level which has been missing in prior literature. As we will see, consumer respond differently and hence treating consumers as aggregated homogenous block misses many important nuances. Second, in the literature only the Poort et al. paper studies website blocking activities. Unlike their paper, we study not only piracy activity but changes in legal consumption and how pirates choose to continue pirating after the website is blocked. Third, we study both the blocking of a single site (The Pirate Bay) and the subsequent nearly simultaneous

blocking of nineteen different popular piracy sites. Our findings corroborate those of Poort et al. for the blocking of The Pirate Bay but contrast them during the nineteen website block, allowing us to draw inferences as to when and how website blocking may be effective. Fourth, because we find heterogeneous effects across consumer groups, we are able to provide insight into which consumers are more likely to increase legal consumption after piracy websites are blocked. Finally, due to our novel use of Internet consumer panel tracking data, we are able to utilize an identification strategy that might be employed to study a number of other natural experiments regarding behaviors on the Internet.

4. Data

Unlike previous studies analyzing the impact of antipiracy interventions which relied on aggregate market data to analyze only the impact on legal activity, our analysis utilizes a novel Internet user dataset to better understand how consumers react to such interventions. We obtained data from an anonymous Internet consumer panel tracking company, which we refer to as PanelTrack in this paper.¹³ While PanelTrack could not provide us data at the consumer level, they were able to provide 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 studying the blocking of The Pirate Bay, consumers were sorted into ten different groups based on their total number of visits to The Pirate Bay during March and April, the two months before the block. For studying the blocking of nineteen sites in October/November 2013, consumers were sorted into ten different groups based on their total visits during August and September to any of the nineteen different blocked sites. Thus, for

¹³ 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.

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 nineteen site block, we observe each segment from August 2013 until February 2014.

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.

Table 1 provides mean visits to each site during February, March, and April (the months before the block) for each of the consumer segments that we used to study the initial Pirate Bay block.

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	10%	1	1,167	949	27	79
2	13%	2.5	1,737	1,113	26	97
3	13%	5	1,570	718	19	52
4	10%	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	11%	36	1,367	554	19	85
8	11%	67.9	1,721	493	36	65
9	11%	226.3	2,907	559	21	96

Note that since we were not given the number of users in the control group, it is difficult to make comparisons to it, though we know that it is significantly larger than any of the treat-

ment groups (the majority of people were non-users of The Pirate Bay). However, each of the treatment groups make up relatively similar portions of the rest of the sample, so we can observe several important facts. First, though 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. Interestingly, the heaviest users of legal streaming sites are actually the mid-tier users of The Pirate Bay.

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	Unknown	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, while it is unfortunate that we do not know the actual size of each group, and thus cannot directly compare the numbers in this table to those of Table 1, we do note that the “pre-block visits to blocked sites” column is per user and thus can be used as a point of comparison — and in general, the means here are lower than those in Table 1. Though 19 sites were blocked, The Pirate Bay was a particularly large source of pirated content and these 19 sites, even in ag-

gregate, do not appear to be as intensely accessed by each user. Again, 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 (1 cyberlocker and 7 piracy link/streaming sites). Finally, although we cannot compare numbers across tables without knowledge of group size, it is likely that the higher visits to legal streaming sites in Table 2 reflect the fact that paid legal streaming sites like Netflix had higher adoption levels by October 2013 than they had in May 2012.

It is worth noting that the data from PanelTrack do show that both blocking injunctions – The Pirate Bay and the eighteen site block – were effective in drastically reducing traffic to the blocked sites. Total visits to The Pirate Bay (the main site, not proxies) across all treated groups dropped by 80% in the 4 months after the block as compared to before. Total visits by the treated groups to the nineteen sites after the October 2013 blocks dropped by 75%. One might ask why the drop was not 100% if the sites were blocked. The primary reason is that in each case, we consider the month during which the blocks occurred as being in the “after” period in order to be conservative in our estimates in the empirical section. Thus, for as many as four potential weeks in the after period some of the blocked sites may still have been available to users, depending on exactly when during that month the user’s ISP blocked each site. Additionally, to the degree that users circumvented the blocks by using a VPN or similar measure, this may continue to show up in our data as visits to the blocked sites even after the block was enforced. Finally, some smaller ISP’s in the UK may not have participated in the blocks. Nonetheless, it is quite 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.

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

The Pirate Block in May 2012: 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).¹⁴ We also divide treated users into nine different groups based on their number of visits to The Pirate Bay in the two months before the block — thus, we will 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.

Before turning to regressions to make causal inference, it makes sense to visualize the data in a way that best represents our identification strategy. Figure 3 shows a scatterplot where the x-axis represents the intensity of treatment (pre-treatment Pirate Bay visits) for each consumer group and the y-axis represents the percent change in visits to paid legal streaming sites. Figure 4 then shows a similar scatterplot except that the y-axis represents the percent change in visits to other, unblocked torrent sites.

¹⁴ 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.

FIGURE 3: TREATMENT INTENSITY VS. CHANGE IN VISITS TO PAID LEGAL STREAMING SITES





FIGURE 4: TREATMENT INTENSITY VS. CHANGE IN VISITS TO UNBLOCKED TORRENT SITES



Note the difference between these graphs. In Figure 3, while there is significant variance in the change in visits to paid legal streaming sites across consumer groups, there is no clear relationship between the change in visits and the measure of treatment intensity (we formally test this statement later in the regression analysis). Some of the groups with the lowest intensity of treatment have some of the largest increases in visits to paid streaming sites. However, in Figure 4, a more clear pattern emerges: Generally, the greater the intensity of treatment, the larger the increase in visits to other, unblocked torrent site after The Pirate Bay was blocked. These two scatterplots demonstrate the idea behind our methodology, applying a difference-in-difference estimate to our outcome variables where one of the differences is pre-treatment intensity of usage of the blocked site in question.

Having visually demonstrated the identification strategy for some of the data, we can now formally test the hypotheses of whether the block of The Pirate Bay causally impacted visits to paid streaming sites, visits to other torrent sites, visits to cyberlocker piracy sites, and visits to VPN sites.

To do this we run the following model:

$$LnVisits_{jt} = \beta_0 + \beta_1 After_t + \beta_2 TreatIntensity_j * After_t + \mu_j + \varepsilon_{jt} \quad (1)$$

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 month is May, June, July, or August and $treatintensity_j$ indicates the number of visits that the average consumer in group j made to The Pirate Bay during April and March of 2012. Finally, U_j is a vector of group fixed effects and e_{jt} is an idiosyncratic error term. In this model, B_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).

This difference-in-difference approach should be a well-identified estimation strategy for the impact of the block on visits to other piracy channels or VPN's. However, we note that even though the highest users of The Pirate Bay were the most heavily treated by the block, they were also the heaviest users of other unblocked sites (see Table 1), which means that they had the least cost of finding alternate piracy sources after the block. As a result, they may have a lower probability of shifting their consumption to legal sites. The fact that they are the most affected by the block must dominate this secondary factor (because if you aren't impacted by the block, your probability of shifting activity to legal channels is irrelevant). But this factor implies that although the impact of the block (if it exists) should increase with treatment intensity, it may do so nonlinearly, as the highest treatment groups may only receive a small marginal impact from the block. Thus, for regressions involving the impact of the block on visits to legal sites, we also include a squared term for the variable of interest.

$$LnVisits_{jt} = \beta_0 + \beta_1 After_{\tau} + \beta_2 TreatIntensity_j * After_{\tau} + \beta_3 TreatIntensity_j^2 * After_{\tau} + \mu_j + \varepsilon_{jt} \quad (2)$$

Table 3 below shows the results from models (1) and (2) for each of four outcome variables.

TABLE 3 – EFFECT OF PIRATE BAY BLOCK ON SITE VISITS

	Paid Streaming	Paid Streaming	Other Torrent	Cyberlockers	VPNs
After Block	-0.623+	-0.639	-0.242*	-0.370+	-1.009**
	(0.302)	(0.398)	(0.042)	(0.168)	(0.362)
TreatIntensity * After Block	0.001	0.001	0.002*	0.000	0.010+
	(0.003)	(0.019)	-0.001	(0.002)	(0.005)
TreatIntensity ² * After Block		0.000			
		(0.000)			
Constant	11.949*	11.949*	14.688*	13.866*	10.377*
	(0.185)	(0.198)	(0.026)	(0.103)	(0.221)
Observations	20	20	20	20	20
Consumer groups	10	10	10	10	10
R-squared	0.853	0.854	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 two columns of Table 3 examine the impact of the block on visits to paid legal streaming channels, with the first column forcing a linear relationship and the second allowing for a quadratic curve. We note that B_2 and B_3 are 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, B_2 for other torrent sites (the third 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 fourth 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 fifth column, B_2 is measured as 0.01 (at 90% confidence), indicating a strong turn to 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 – the 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 any increase in paid legal sources of video content from blocking just this site.

TABLE 4 – EFFECT OF OCT/NOV BLOCKS ON SITE VISITS

	Paid Streaming	Paid Streaming	Other Torrent	Cyberlockers	VPNs
After Block	-0.027 (0.061)	-0.126+ (0.062)	-0.342* (0.093)	-0.377* (0.102)	0.201 (0.263)
TreatIntensity * After Block	0.006** (0.003)	0.022** (0.007)	0.006 (0.004)	-0.006 (0.004)	0.027** (0.011)
TreatIntensity ² * After Block		-0.0003** (0.0001)			
Constant	13.815* (0.036)	13.815* (0.028)	13.436* (0.055)	13.575* (0.060)	9.997* (0.155)
Observations	20	20	20	20	20
Consumer groups	10	10	10	10	10
R-squared	0.993	0.996	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%

19-Site Block in October/November 2013: We now turn our attention to the second event in our study: the blocking of 19 different video piracy websites within a thirty day period be-

tween October and November of 2013. We again estimate the empirical model using OLS regression and report the results in Table 4.

In this case, the results are noticeably different than those for the single block above. 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. An individual who made 10 visits to the blocked sites in the month before the block subsequently increased his visits to legal streaming sites by 19% more than an individual who didn't use the blocked sites, based on the second column.¹⁵ But the impact is non-linear – an individual who made 40 visits to the blocked sites in the month before the sites were blocked increased his visits to legal streaming sites about 49% more than a non-user of the blocked sites.¹⁶ This diminishing marginal increase in treatment intensity is consistent with our hypothesis that heavier users of the blocked sites have a lower cost to find alternate piracy sites after a block, and thus are less likely to shift consumption toward legal channels.

Indeed, Table 2 demonstrates that the heaviest users of the blocked sites were also disproportionately the heaviest users of other unblocked torrent sites as well as cyberlocker sites, and so even though they were most impacted by the blocks they were somewhat more resistant to converting activity toward legal consumption. However, note that within the range of treatment intensity (pre-block visits to blocked sites) in the data, the relationship between treatment intensity and post-block visits to legal sites never turns negative, indicating that all levels of treatment groups increased their legal consumption by more than the control group.

¹⁵ We describe how we arrive at this calculation in the text just before Table 6 below.

¹⁶ Note that even in the first column, which forces a linear relationship between treatment intensity and change in visits to paid streaming sites, the coefficient of interest is still positive and statistically significant. However, it is clear from the higher t-values and higher R^2 that a quadratic model is a better fit, and so we use this model for our calculations.

In the third and fourth columns we measure no statistically significant changes in visits to other unblocked torrent sites or to cyberlockers sites, though 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.

TABLE 5 – PLACEBO TEST AUGUST TO SEPTEMBER 2013

	Paid Streaming	Paid Streaming	Other Torrent	Cyberlockers	VPNs
After Block	-0.256*	-0.2160+	-0.226*	-0.314*	-0.405**
	(0.078)	(0.106)	(0.070)	(0.069)	(0.142)
TreatIntensity * After Block	0.0003	-0.006	-0.003	-0.003	-0.002
	(0.003)	(0.015)	(0.002)	(0.003)	(0.006)
TreatIntensity ² * After Block		0.0001			
		(-0.0002)			
Constant	13.921*	13.921*	13.626*	13.800*	10.176*
	(0.046)	(0.048)	(0.040)	(0.041)	(0.084)
Observations	20	20	20	20	20
Consumer groups	10	10	10	10	10
R-squared	0.989	0.988	0.991	0.991	0.973

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%

Measuring an increase in visits to paid legal streaming sites in this case is important, and deserves further exploration. Our coefficients should only be interpreted as causal impacts if each segment’s trends would have been uncorrelated with treatment intensity if not for the block. Thus, before the block, we should observe no correlation in the month-to-month trend in visits and the treatment intensity variable. In Table 5 we run the same model, except that we consider only August in the pre period and only September in the post, thus “pretending” as if the block

happened at the end of August. This is essentially a placebo test: a test to show that pre-existing trends are uncorrelated with treatment intensity.

The placebo test yields the expected results: not a single one of the B_2 's are statistically significant at even the 90% confidence level. Further, the coefficients are generally much smaller than those measured in Table 4 (studying the true period of the 19 blocks). Thus, at least in the two months before the blocks occurred, there was no relationship between each group's time trend in visits to each category of sites and that group's treatment intensity. The correlation appears only after the blocks happened, lending considerable strength to the causal interpretation of the results.

Table 6 – Estimated Causal Increase in Visits to Paid Legal Streaming Sites

Consumer Segment	Pre-block Visits/User to Blocked Sites	Causal Increase in Visits to Legal Sites
0	0	0.0%
1	1.0	2.2%
2	2.0	4.4%
3	3.0	6.5%
4	4.0	8.7%
5	5.4	11.7%
6	8.2	17.5%
7	13.2	26.8%
8	23.8	42.4%
9	66.4	14.8%

If we consider the estimates from the second column in Table 4 to measure the causal influence of the blocks on each segment's usage of paid legal streaming sites, we can then say that the blocks increased the natural log of visits to legal sites for a segment by $0.022x - 0.0003x^2$ where x is the treatment intensity (or pre-block visits to blocked sites per user) of the segment. Thus, the causal percent increase in visits to legal sites for a segment is $[\text{EXP}(0.022x - 0.0003x^2)]$

– 1] * 100. Table 6 shows these estimated causal increases in legal streaming site visits for each segment.

These figures are consistent with what we observed in the raw data – the greater the treatment intensity, the larger the increase in visits to paid legal streaming sites after the block. The exception is the heaviest group of pirates (segment 9), who still increase paid streaming more than the control group, but not as much as some of the less intense users of the blocked sites (segments 6, 7, and 8). We discuss this result in the following section.

6. Discussion

While the use of website blocking has increased in recent years as a tool in the fight against intellectual property theft, ours is the first study we are aware of that analyzes their effectiveness in changing consumer behavior. We use data provided by a panel tracking company to analyze the impact of two website blocking events in the UK: The blocking of The Pirate Bay in May 2012 and the blocking of 19 additional sites in October and November of 2013.

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, but had no statistical impact on legal entertainment sites. However, the blocking of 19 different sites in October and November of 2013 caused a statistically significant increase in the usage of legal streaming sites: The more a consumer visited the blocked sites in the months leading up to the block, the more likely they were to visit legal sites after the block. This pattern 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. Table 6 showed the percent increase in legal streaming for each consumer seg-

ment. 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.

If we combine groups 1, 2, and 3 (the lightest users of the blocked sites), our model suggests that the blocks caused users in these groups to increase their clicks on paid legal streaming sites by 3.5% on average. Combining groups 7, 8, and 9 (the heaviest users of the blocked sites), we find that the blocks caused them to increase their use of paid legal streaming sites by 23.6%. If we consider all users of the blocked sites together, we estimate that the blocks caused them to increase clicks to paid streaming sites by 12%.

We also note that these estimates are conservative in a variety of ways. Recall that after the blocks occurred we only see a 75% decrease in total visits to blocked sites, rather than a 100% decrease. There are several reasons for this. First, although we include all of November in the “post” period, we do not know when in November each ISP executed the court-ordered blocks. Thus, part of our post period includes observations before the sites were blocked by some ISPs, causing our model to underestimate the true effect of the blocks. Second, even if we could correct for this timing issue, there would still be use of the blocked sites among subscribers to smaller ISPs that did not participate in the blocking program. Third, some subscribers of ISPs that did participate in the blocking program adopted VPNs to circumvent the block. To be conservative, we do not attempt to correct our estimates for non-participating ISPs or VPN-based circumvention.

Notably, the blocks did cause some users to increase visits to VPN's and may have caused dispersion of piracy to other torrent sites, though this latter finding was not measured with statistical significance. But while some people may circumvent the blocks, it is clear from the data that in addition to reducing visits to blocked sites, site blocking caused some former pirates to migrate their consumption toward legal channels. Importantly, because we are examining groups of individuals, we cannot determine the degree to which this increase comes from new consumers turning to these legal channels or increased usage of these legal channels by existing users. This remains an interesting question for future research.

One might ask why blocking 19 sites at once caused an increase in visits to legal sites while blocking a single site did not, especially given that the single blocked site was one of the most popular piracy sites ever in the UK. Although we cannot rule out alternative explanations, we believe the most likely explanation is that when a single site is blocked, many pirates will know at least one other good site, but when many sites are blocked, the cost to find another reliable site is higher for all but the most active pirates. This explanation is supported both by the non-linear (and decreasingly positive) relationship between treatment intensity and increases in legal consumption, and by the fact that the heaviest users of the 19 blocked sites were more likely to use other piracy sites even before the blocks, implying that it was easier for them to find alternate sites after the blocks.

In addition to the finding that a one website block was ineffective at increasing legal consumption while 19 site blocks did increase legal consumption, our results demonstrate that piracy does indeed displace usage of legal paid streaming sites, despite the relative convenience and low cost of such sites. This reinforces results from earlier studies which suggest that while mak-

ing legal content more attractive can turn some pirates into legal consumers, it is more effective when accompanied by policies that make illegal content less attractive.

There are several limitations to this study. First, we only study a one site blocking injunction and injunctions for blocking nineteen sites. Further study of additional blocking actions (including different numbers of sites) would be valuable to verify the conclusion that the number of sites blocked is a meaningful moderator of the impact. Second, while we have shown that the heaviest pirates are more reticent to turn to legal channels (as observed in the non-linearity of the impact across segments), we can only suggest potential explanations for this and future work should delve further into the reason for this diminishing marginal impact. Third, the increase we observe in use of paid streaming sites cannot be broken down into new users or increased usage from existing users, and separating these two effects may have important managerial implications, as converting an individual from piracy toward a new legal streaming service may have different implications than causing a pirate to increase usage of a service to which he was already subscribed. Fourth, we have used paid legal streaming sites in the UK as a proxy for increasing legal consumption behavior. Other legal channels exist in the UK such as paid digital downloads (on iTunes, for example) or DVD/Blu-Ray purchases, and understanding the impact of site blocks on these channels would be of interest to managers and policymakers as well. Data on such behavior were not available to us for this particular study. 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.
- 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 91–114.
- 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, 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, Brett, Michael D. Smith, Rahul Telang. 2014. 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.
- Danaher, B., J. Waldfogel. 2012. Reel Piracy: The Effect of Online Movie Piracy on Film Box Office Sales. Working Paper. Wellesley College. Boston, Massachusetts.
- Fowler, Geoffrey, D. Barrett, and S. Schechner. 2012. “U.S. Shuts Offshore File-Share ‘Locker’.” *Wall Street Journal*, 20 January.
- 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. 2013. Piracy and Movie Revenues: Evidence from Megaupload: A Tale of the Long Tail? Working Paper, University of Zurich. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2176246
- Aguiar L, Peukert C., and Claussen J. 2015. Online Copyright Enforcement, Consumer Behavior, and Market Structure. Working Paper. Available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2604197.

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." Working Paper, Carnegie Mellon University.

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.

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>)

**Appendix A – List of Piracy Sites (Allowing Access to Pirated Video Content)
Blocked in October-November 2013**

1. YIFY-torrents
2. FTVO
3. vodly.to
4. primewire.ag
5. watchfreemovies.ch
6. 1337X
7. Bitsnoop
8. Extratorrent
9. Monova
10. Torrentcrazy
11. Torrentdownloads
12. Torrenthound
13. Torrentreactor
14. Torrentz
15. Filecrop
16. Filestube
17. Rapidlibrary
18. solarmovie.so
19. tubeplus.me