The Effect of Piracy Website Blocking on Consumer Behavior

Abstract

In this paper we ask whether antipiracy enforcement interventions that aim to make copyright-infringing content more difficult to access can decrease piracy and increase legitimate consumption. We do this in the context of three court-ordered events affecting consumers in the UK: We first study Internet Service Providers' blocking of 53 major piracy sites in 2014 and we then study two smaller waves of blocking – the blocking of 19 piracy sites in 2013 and the blocking of The Pirate Bay in 2012.

We show that blocking 53 sites in 2014 caused treated users to decrease piracy and to increase their usage of legal subscription sites by 7-12%. Similarly, we found that the 2013 block of 19 different piracy sites caused users to increase visits to legal sites by 8%. However, blocking a single dominant site in 2012—The Pirate Bay—caused no increase in usage of legal sites, but it did cause users to increase visits to other unblocked piracy sites and VPN sites. This suggests that to increase legal IP use when faced with a dominant piracy channel, the optimal policy response must block multiple channels of access to pirated content, a distinction that the current literature has not made clear.

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

THE EFFECT OF PIRACY WEBSITE BLOCKING ON CONSUMER BEHAVIOR

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 demand-side antipiracy interventions (enforcement actions against consumers) on legal consumption (e.g., Bhattacharjee et. al. 2006, Adermon and Liang 2014, Danaher et al. 2014a)

[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.

and demonstrate that such policies can be effective at increasing legal sales. However, the literature is divided on whether supply-side interventions (enforcement actions against websites and protocols that facilitate piracy) are effective at reducing piracy and increasing legitimate consumption. Danaher and Smith (2014) analyze the shutdown of Megaupload.com, the world's largest piracy cyberlocker[2] which housed 25 petabytes of copyright infringing content. They show that this action increased legal digital revenues from Hollywood film content by 6.5-8.5%. However, Poort et. al (2014) and Aguiar et. al. (2018) each study antipiracy interventions that targeted individual popular filesharing sites and find no lasting impact on piracy levels or legitimate consumption.

Given this apparent tension in empirical research on supply-side antipiracy enforcement, we ask two important research questions. First, we examine the court ordered blocking of 53 major piracy websites in the United Kingdom in November 2014 and ask what effect this had on users' illegal and legal consumption behaviors. We exploit this natural experiment by using a panel dataset on the behaviors of 24,620 UK Internet users before and after the blocks. Although all individuals were blocked from accessing the sites, the "bite" of this treatment varied from user to user. An individual who used the blocked sites frequently was more heavily shocked than an individual who was only an infrequent user of the blocked sites, and both of these individuals were more heavily shocked than non-users of the blocked sites. We therefore implement a generalized version of the difference in difference model, asking whether an individual's treatment intensity (pre-block usage of the blocked sites) was correlated with her pre-post change in visits to legal media sites and/or alternate unblocked piracy sites.

[2] A cyberlocker is a cloud site or server that provides file storing and sharing. In the context of piracy, cyberlockers are repositories of illegal content, whereas other types of sites merely provide links or tracker files that link to pirated content stored elsewhere.

We find that these blocks caused meaningful decreases in total piracy as well as a 7-12% increase in usage of paid legal streaming sites among users affected by the blocks. This motivates our second research question – we ask why some supply-side antipiracy interventions are effective and some are not. We exploit two earlier natural experiments in the UK – the blocking of The Pirate Bay, a dominant piracy site, in May 2012 and the blocking of 19 major piracy sites in 2013. We obtain panel data on aggregated behaviors of a number of different groups of consumers before and after the blocks where the groups differ primarily in terms of the treatment intensity of the blocks – some groups were making light use of the blocked sites before they were blocked while others were using them more heavily. Using our generalized difference-in-difference model, we demonstrate that the blocking of 19 sites in November 2013 also caused a decrease in piracy leading to an 8% increase in consumption from legal sites. However, when we study the blocking of the Pirate Bay in 2012, we find results extremely similar to Aguiar et. al. (2018) – the block caused at most a one month increase in legal consumption which immediately returned to pre-block levels. Users of the Pirate Bay then increased their usage of other unblocked piracy sites and Virtual Private Networks (VPNs) to circumvent the block.

Our study contributes to the academic literature in that we are the first to study multiple instances of the same type of intervention at varying degrees of strength. As such, we show that differences in the effectiveness of various antipiracy interventions may be explained by the strength of the intervention, a finding that is consistent with analytical results from Dey et. al. (2018). Our results validate the findings of prior researchers showing that blocking access to content through a single popular site is ineffective at reducing piracy or increasing legitimate consumption, but we demonstrate that the effect of supply-side antipiracy enforcement is more nuanced, as blocking a number of sites at once causes meaningful increases in legal consumption.

From a policy perspective, we are the first to evaluate specifically the impact of piracy website blocking on legal consumption.[3] Unlike shutting down entire sites (as in Danaher and Smith 2014 or Aguiar et. al. 2018), 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 P2P filesharers require in order to download a media file through the BitTorrent protocol), or unauthorized media streaming "link" sites (sites that conveniently provide links to stream illegal content hosted on other sites). 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 and in Japan. Some countries, such as Australia, have explicitly stated that the effectiveness of website blocking will be 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] Thus, our work has practical importance for policymakers trying to decide whether to turn to piracy website blocking as a solution to copyright violations.

[3] Reis et. al. (2018) have concurrently undertaken a similar study on piracy website blocking and legitimate consumption in Portugal, This study is not yet published.
[4] Some countries implementing site blocking 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)
[5] https://www.theregister.co.uk/2017/01/30/australia_to_review_effectiveness_of_isps_copyrightdefending_website_blocks/

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 the same 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.

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.[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 3) have causally linked this widespread piracy with significant losses in motion picture revenue in all major sales channels.

One of the more common antipiracy methods in recent years has been piracy website blocking, a strategy that has been attempted in over 25 countries to date. For example, the UK has used website blocking to fight piracy since October 2011 when British Telecom and five other UK ISPs were ordered by the High Court to block their customers from accessing Newzbin2, an indexing site for pirated content posted to the Usenet. Following the Newzbin2

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

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. 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. Such methods are distinct from packet filtering, which inspects transmitted content at the packet level, and thus website blocking is more easily circumventable.

Our present analysis concerns three waves of UK blocks that occurred in 2012, 2013 and 2014, respectively. Specifically, in April 2012 six major UK ISPs were ordered by the courts 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.[7] 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.[8] Later, in November 2013, these six 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 53 additional piracy sites.

Website blocking may have a different impact than site seizures because, while blocked, the site is still operational and may serve users outside of the country where the blocking occurs.

[7] Specifically the ISPs Everything Everywhere, Sky, TalkTalk, Telefónica and Virgin Media were ordered to block access in April resulting in the block occurring in May. The sixth ISP, British Telecomm, implemented the blocks in June.
[8] http://www.theguardian.com/technology/2012/apr/30/british-isps-block-pirate-bay

Sophisticated domestic users are able to circumvent the ISP-level block through the use of Virtual Private Network services[9] or 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. 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.[11] Indeed, Aguiar et. al. (2018) find evidence that after Kino.to was shut down, a number of new piracy sites and proxy sites appeared in its place and traffic gravitated to these sites, leading to the unintended consequence of more dispersed piracy.

Nonetheless, users may bear a time and monetary cost to finding new mirror/proxy sites and purchasing VPN access, or they may be reticent to trust downloading files from new piracy sites that they do not know. Prior research shows that making legal content more attractive or illegal content less attractive can induce 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 switching or learning costs are high.

[9] By using a VPN, a user can appear to be attempting to access a blocked site from another country, and thus the request to the site will resolve.

[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)

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

3. Existing Literature

There is a significant body of work on the relationship between piracy and sales of video content, including Rob and Waldfogel (2006), Liebowitz (2008), Smith and Telang (2010), Danaher et al. (2010), Zentner (2012), and Ma et al. (2014).[12] The majority of this literature finds evidence of sales displacement caused by piracy across a variety of media types, though there are several prominent dissenting studies, including Oberholzer-Gee and Strumpf (2007) and Aguiar and Martins (2016).

There have also been a number of studies on whether antipiracy enforcement actions by government can influence pirates to turn to legal channels of consumption. Studies on the effects of demand-side antipiracy interventions targeting consumers of pirated content show that these actions tend to be effective in reducing pirated content and/or increasing legitimate consumption (Bhattacharjee et. al. 2006, Danaher et. al. 2014a, Adermon and Liang 2014), though McKenzie (2017) shows no impact of such interventions on box office revenues. In general, the reaction to demand-side interventions has been to view them as draconian and political taste for such enforcement activity has diminished. (Danaher et. al. 2017)

However, our study analyzes a supply side anti-piracy enforcement, defined as an enforcement action against the sites and protocols that facilitate media piracy. This literature is divided. Danaher and Smith (2014) find that the shutdown of the popular piracy cyberlocker Megaupload.com increased digital movie revenues of Hollywood films by 6-8%, but Peukert, Claussen, and Kretschmer (2017) find that this shutdown led to a decrease in the box office of smaller, independent films. Notably, because Megaupload was a cyberlocker, the shutdown of this site led to the removal of a vast amount of pirated content from the Internet. In contrast,

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

Poort et. al. (2014) find that when Dutch courts ordered ISPs to block The Pirate Bay, it caused only a small, transient decrease in illegal downloading activity. Because this was a blocking action rather than a shutdown, no content was actually removed from the Internet and it remains very possible that Dutch consumers could find ways to access the torrent trackers that The Pirate Bay had linked. Finally, Aguiar et. al. (2018) find that the shutdown of a single major German piracy linking site, Kino.to, had only a short-lived impact on illegal downloads and no impact on legal consumption in Germany. Although this was a site shutdown, Kino.to was not a cyberlocker and therefore hosted very little illegal content – instead, it was a site that conveniently provided a centralized set of links to dispersed external sources of pirated content. This shutdown, therefore, had far more in common with "blocking" access to content than it did with the shutdown of a major cyberlocker like Megaupload. Thus, there is reasonable debate over whether supply-side antipiracy enforcement is actually effective, and in particular the current evidence on blocking or making access to content less convenient suggests that it does not reduce piracy or increase legal consumption as consumers continue to access the same amount of illegal content through use of other sites.

Our study helps to resolve this tension in the existing literature. In demonstrating that the blocking of The Pirate Bay caused only a short, statistically insignificant increase in usage of legal sites, later reversed by a statistically significant increase in usage of unblocked piracy sites, we confirm the important findings of Poort et. al. (2014) and Aguiar et. al. (2018) who both studied only actions against a single dominant site. However, because we are the first to also analyze the blocking of multiple sites at once, we provide evidence that the full story is more nuanced – blocking a number of popular illegal sites at once can cause a significant increase in legal consumption. This provides an empirical test for the theoretical result of Dey et. al. (2018), who suggest that supply-side antipiracy enforcement can be effective, but only when the enforcement makes piracy significantly inconvenient.

Why should we consider blocking 53 piracy sites a stronger enforcement action than blocking one dominant piracy site? Prior research demonstrates the presence of switching costs between Internet sites/portals and suggests that such costs may be lower when consumers are already aware of the other sites (Chen and Hitt 2002, Goldfarb 2006), and this appears to be empirically true for piracy sites as well (Aguiar et. al. 2018). If one site is blocked, most or perhaps even all users are likely to be aware of and perhaps even have proficiency with another alternate piracy site. On the other hand, if 19 or 53 sites are blocked, some users of those sites may have little or no awareness of any remaining piracy sites. In such cases, switching costs would involve discovering another site, deciding whether to trust it (piracy sites outside of a trusted few are notorious for the spread of malware), and learning how to properly use it. Thus blocking 53 (or 19) sites is more likely to lead to higher switching costs than simply blocking one popular site, even in a case where the one site served as many users as the 53 combined. Adopting a legal site, on the other hand, involves very little switching cost as such sites are well-advertised, easy to use, and largely trusted. The implication is that it is possible that blocking one very large piracy site could be ineffective at changing user behavior while blocking a number of prominent piracy sites could induce users to increase their legal consumption.

4. Data

We obtained data from an anonymous Internet consumer panel tracking company, which we refer to as PanelTrack in this paper.[13] PanelTrack offers individuals compensation to participate in their panel, and then subsequently installs software that monitors a user's PC Internet activity unnoticeably in the background for as long as the user remains in the panel.[14]

We use each wave of blocks as a natural experiment affecting piracy at the blocked sites to determine how consumers respond. They may increase usage of remaining unblocked piracy sites, increase usage of legal sites, circumvent the blocks by using VPNs, or simply stop consuming the media that they had been pirating.

2014 Blocks

We use a panel of 24,620 UK Internet users, covering the time period from August 2014 to February 2015, to study the impact of the 53 site blocks in November 2014. This panel includes each individual's monthly visits to: 1) blocked piracy sites, 2) unblocked piracy sites, 3) VPN sites; and 4) paid legal streaming sites like Netflix or LoveFilm (See Appendix A for a list of the blocked sites in each wave of blocks and Appendix B for an explanation of how we determined the sets of legal subscription websites, unblocked piracy sites, and VPN sites). The panel is unbalanced – a number of these users joined the sample after our study began or left before it ended. As such, we observe 67,098 user-months. Our difference-in-difference approach can be

[13] Because our study is about piracy, PanelTrack required that the company remain anonymous. However, 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.
[14] Though this observation occurs in the background, we cannot rule out that the sample is biased due to Hawthorne effects. However, these data are the standard in the entertainment industry as well as others for learning about Internet consumer behavior over time. Because we study changes within users before and after the blocks, our methodology helps to difference out any degree to which users behave differently as a result of participating in the panel.

effectively applied to unbalanced panels, but we also show in appendix D that all of our 2014 results hold when estimated using only the balanced panel of individuals observed in all seven months.

Table 1 shows some descriptive statistics for the panel of UK Internet users before and after the blocks in November 2014. We present average monthly site visits, prior to and after the blocks, to the categories of sites considered as our outcome variables of interest. We exclude observations from November from this table, as the blocks were in the process of being implemented, and thus November is considered "partially" treated. We define treatment intensity as the average monthly visits to pre-blocked sites that were subsequently blocked in November 2014. We see that the November 2014 blocks were effective at reducing visits to blocked sites. Visits to blocked sites dropped by 88% from the three months before the blocks to the 3 months after (not counting the month of the blocks).[15] It also appears that visits to unblocked sites decreased and visits to paid sites increased. However, we will investigate this more rigorously using a difference-in-difference analysis to control for the time trends underlying the blocks that may be driving the changes.

[15] There are several reasons why the drop may not be 100%. First, not all ISPs were forced to comply with the blocks, only the six largest ones. Second, users on VPNs could still access the blocked sites and PanelTrack's machine-side software would still detect such visits. Finally, it may be that some ISPs did not fully operationalize blocks to all of the sites by the beginning of our post-period.

Table 1: 2014 Summary Statistics

Treatment Intensity (average monthly pre-block blocked site visits)	Blocked Sites		Unblocked Sites		VPN Sites		Paid Streaming Sites		N
	Pre	Post	Pre	Post	Pre	Post	Pre	Post	
0-1	0.0	0.1	3.7	4.5	0.1	0.1	1.5	2.2	19525
1-5	1.9	0.5	9.9	7.9	0.1	0.1	1.4	2.2	3323
5-10	6.8	1.3	20.6	11.8	0.0	0.1	2.4	2.7	798
10-50	20.2	3.6	47.0	24.9	0.1	0.2	2.6	3.1	852
50+	84.7	15.6	179.2	62.1	0.1	0.2	1.0	3.2	122

Notes: This table shows average monthly site visits by site category and treatment intensity. Treatment intensity is measured by average monthly visits, prior to the block, to sites that would eventually be blocked. N shows number of users within each bucket of treatment intensity over which averages are calculated.

2012 and 2013 Blocks

Due to a change in policy surrounding privacy concerns, PanelTrack would not release individualized data for the 2012 and 2013 waves of blocks. We instead received monthly data with individuals aggregated by groups stratified by users' pre-block usage of websites that were subsequently blocked. Thus, instead of measuring the treatment intensity (the "bite" of the blocks) on individual users by how many visits each user made to blocked sites prior to the blocks, for these waves we observe aggregate group behavior. The treatment intensity "bite" of the blocks for each group of users is defined as that group's overall average monthly visits to the blocked sites in the three months before the blocks. For example, during the 2012 Pirate Bay Block, one of the groups had averaged only 1 visit per user per month to The Pirate Bay before it was blocked, while the group with the heaviest Pirate Bay users averaged 230 visits per user per month to that site in the months before it was blocked.[16] Thus for the 2012 wave and for the 2013 wave we have separate datasets, each of which observes aggregate visits for 10 different consumer groups for the seven months surrounding each wave of blocks. Importantly, because we only observe behavior at the group level and cannot observe individuals entering and exit observations, we required that the groups be formed only of individuals who were observed during all seven months, and thus the groups were created from a balanced panel of users.

Again because of privacy concerns, PanelTrack would not reveal the number of individuals comprising each group in our aggregate group data. They provided visit counts scaled to the UK population of internet households by using sample weights to expand individual data to the population sample. We were assured these sample weights were designed such that the sample is representative of the population of UK Internet users.[17] Thus, what we observe for each group during each month is the aggregate population-scaled visits for that group during that month to each of the site types in question (legal, illegal, VPNs). Because we also know the number of projected users in each group, we can divide scaled visits by projected users to determine the average number of visits per user per month. Importantly, we confirmed with PanelTrack that each group of users in both datasets is comprised of at least two hundred raw users, and thus the data points are generated by a relatively large sample within each group.

A positive feature of our 2012 and 2013 data is that we were able to separate the unblocked piracy sites into two categories of piracy sites – unblocked torrent sites, and unblocked cyberlocker sites – and obtain each groups' aggregate monthly visits to each of these sub-types

[16] Exact details of how PanelTrack sorted their users into these ten groups are in Appendix C. This explains why there is no "control group" per se, though such a concern does not prohibit drawing inferences from the generalized version of the difference-in-difference model with a continuous treatment intensity variable.
[17] Although we would have preferred raw data for these two waves of blocks, PanelTrack's practice of scaling the data is consistent with industry practices such as scaling television ratings data to determine population audience sizes.

of piracy sites.[18] This allows us to measure whether any increase or decrease in unblocked piracy was disproportionately driven by a particular piracy protocol.

Table 2 provides average monthly visits per user by type of site during the pre-period (February, March, and April 2012) and post period (June, July, and August 2012). We exclude May 2012 from this table as the block was in the process of being implemented this month and we consider it "partially" treated.

Table 2 – Avg Monthly Visits Per User Before and After 2012 Pirate Bay Block

	The Pirate Bay		Other Torrent Sites		Cyberlocker Sites		VPN Sites		Paid Streaming Sites	
Group	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
1	**0.8**	0.1	4.9	3.7	4.0	1.6	0.1	0.0	0.3	0.3
2	**2.0**	0.2	11.2	9.9	7.2	4.2	0.2	0.1	0.6	0.3
3	**2.1**	0.4	3.5	3.1	3.0	3.9	0.2	0.2	0.4	0.5
4	**4.2**	0.4	13.9	10.8	6.4	3.5	0.2	0.1	0.5	0.5
5	**6.8**	0.5	16.4	11.9	7.0	6.5	0.1	0.2	2.1	1.0
6	**12.8**	1.5	40.4	29.1	13.1	4.9	0.3	0.2	3.6	0.3
7	**17.1**	2.6	25.6	21.2	11.0	14.9	0.4	0.2	2.5	0.9
8	**38.5**	2.1	32.2	28.9	13.1	7.6	0.4	0.3	2.0	1.3
9	**55.0**	4.5	42.9	50.8	12.3	11.0	0.9	0.8	1.6	2.1
10	**231.2**	11.7	80.2	102.6	15.4	12.0	0.6	1.7	2.7	2.0

The first column in Table 2 (in bold) indicates each group's average monthly pre-block visits to The Pirate Bay, and thus it is our measure of treatment intensity for that group. Clearly there is dispersion across groups in usage of The Pirate Bay, indicating that the bite of the treatment was different for each group. Visits to blocked sites drop by 80-95% across the various groups, indicating an effective block. Heavy users of The Pirate Bay were also heavier users of other torrent and cyberlocker sites. Visits to pirate sites appear more common than visits to paid streaming sites. This may be because during this period, a number of paid streaming sites were in

[18] Piracy linking sites (sites that conveniently provide links to content hosted on other sites) were less common during this time, but those that were present we categorized with cyberlockers as they mostly linked to content on cyberlockers.

their infancy (Netflix, for example, launched in January 2012) and thus was not yet widely adopted.

Table 3 – Avg Monthly Visits Per User Before and After November 2013 Blocks

	Blocked Sites		**Other Torrent Sites**		**Cyberlocker Sites**		**VPN Sites**		**Paid Streaming Sites**	
Group	**Pre**	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
1	**1.4**	0.3	1.2	1.1	0.7	0.5	0.0	0.0	2.0	2.1
2	**2.6**	0.4	1.7	1.0	1.1	0.7	0.0	0.0	2.2	2.4
3	**3.4**	0.4	2.3	1.6	1.5	0.5	0.0	0.0	3.2	2.0
4	**3.9**	0.5	1.6	1.7	1.3	0.6	0.3	0.0	1.9	2.1
5	**5.1**	0.9	2.0	1.9	1.5	1.1	0.3	0.1	3.1	4.1
6	**5.6**	0.7	1.5	1.4	2.2	0.9	0.1	0.1	1.9	1.8
7	**7.5**	0.8	2.4	2.9	1.9	1.4	0.0	0.1	2.4	2.3
8	**12.5**	1.3	3.2	3.1	2.8	1.2	0.0	0.3	2.4	3.2
9	**20.0**	2.8	4.0	4.6	3.6	1.7	0.1	0.4	2.0	3.2
10	**51.5**	5.0	6.1	6.8	8.8	2.6	0.1	0.7	3.9	6.8

Table 3 reports the summary statistics for the data surrounding the 19 site blocks in November 2013. Again, the first column is average monthly pre-block visits to the blocked sites, and thus indicates the bite of the treatment on each group. All groups decrease their usage of blocked sites by 80-90%. Visits to paid streaming sites appear higher in 2013 than they were in 2012, likely due to increased adoption levels of these services. Visits to unblocked piracy sites appears lower, but this is likely because 19 major piracy sites (as well as a number of their mirror sites) are included in the blocked sites rather than just one, leaving less piracy remaining in unblocked sites..

Because visits to blocked sites drop by 80-90% in each wave of blocks, we have clear discrete shocks to piracy at the blocked sites. 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 November 2014 Blocking of 53 Major Piracy Sites

We first turn our attention to how blocking 53 major piracy sites in November 2014 affected consumer use of legal and illegal media channels. Because changes in outcome variables, such as use of paid streaming channels, might change over time for reasons other than the block, we employ a generalized version of the difference-in-difference model using a continuous treatment variable. This is a common method when the treatment being studied is not binary but rather varies in intensity, which is the case with our data. Here the treatment variable is a measure of the "bite" of the treatment on each affected user. We define each user's treatment intensity as proportional to their average monthly visits to the 53 blocked sites before the blocks were enacted. Our logic is that users who visited these blocked sites more before they were blocked were more impacted by the treatment than users who visited them less.

In line with prior use of this generalized difference-in-difference model, we identify the causal effect of the blocks by comparing individuals' pre-post changes in the outcome variables of interest with those individuals' treatment intensity. We acknowledge that individuals are self-selecting into different measures of treatment intensity (based on their tendency to visit the blocked sites). We control for time-invariant differences across users by including individual fixed effects in our model. Our approach relies on the assumption that a user's month to month changes in the outcome variable would be uncorrelated with treatment intensity in the absence of the treatment, which is the parallel trends assumption. Because we observe the individual for three months before the blocks, we can partly test this assumption by testing whether there is a correlation during the pre-period.

We estimate a model of the form

$$Visits_{it} = \beta_0 + \beta_1 * month_t + \beta_2 TreatmentIntensity_i + \beta_3 TreatmentIntensity_i \cdot month_t + \mu_i + \epsilon_{it} \quad (1)$$

Where $Visits_{it}$ indicates the number of website visits by individual i in month t to the set of sites in question (paid legal subscription sites, unblocked piracy sites, or VPN sites). $Month_t$ is a vector of month fixed effects and $TreatmentIntensity_i$ is the average number of monthly visits prior to the block made by individual i. μ_i is an individual fixed effect and e_{it} is the error term. Our outcome variables are visits to sites and should be modeled as count data as they cannot be non-integers or less than zero. We estimate equation (1) using a negative binomial fixed effects regression because of this concern, and prefer it over Poisson because of the issue of over-dispersion in our data. We use the negative binomial fixed effects model as developed by Hausman et. al. (1984).[19] We note that negative binomial fixed effects models do not estimate a true fixed effect due to the incidental parameters problem, and thus can technically estimate a coefficient for $TreatmentIntensity_i$. However, we demonstrate in Appendix D that our results are robust in sign and significance to estimation using Poisson as well as using OLS with log of visits plus one as the outcome variable (both of which estimate a true fixed effect).

The chief coefficient of interest is β_3, which indicates the degree to which treatment intensity is correlated with individuals' month-to-month changes in site visits of interest. Under the parallel trends assumption, we expect β_3 to be statistically indistinguishable from 0 during the months before the treatment. Then, after the treatment, β_3 gives the causal effect of the blocks on site visits. Figure 1 below shows the estimation of equation (1) plotting β_3 and its

[19] Note that negative binomial fixed effects models do not estimate a true fixed effect, and thus can technically estimate a coefficient for "TreatIntensity". As already discussed, in Appendix D we demonstrate that our results are the same in sign and significance and similar in magnitude if we estimate using a Poisson fixed effects regression which estimates the standard conditional likelihood fixed effect.

standard errors over time. We observe that the 2014 website blocks caused a decrease in unblocked piracy site visits, an inconclusive impact on VPN site visits, and a persistent increase in legal subscription site visits. The parallel trends assumption appears to hold for VPN sites visits and legal subscription site visits. Thought it does not appear to hold for unblocked piracy site visits, the discrete jump in the post period is much larger than the slight pre-existing downward trend.

Figure 1: Effect of 2014 Blocks on Outcomes

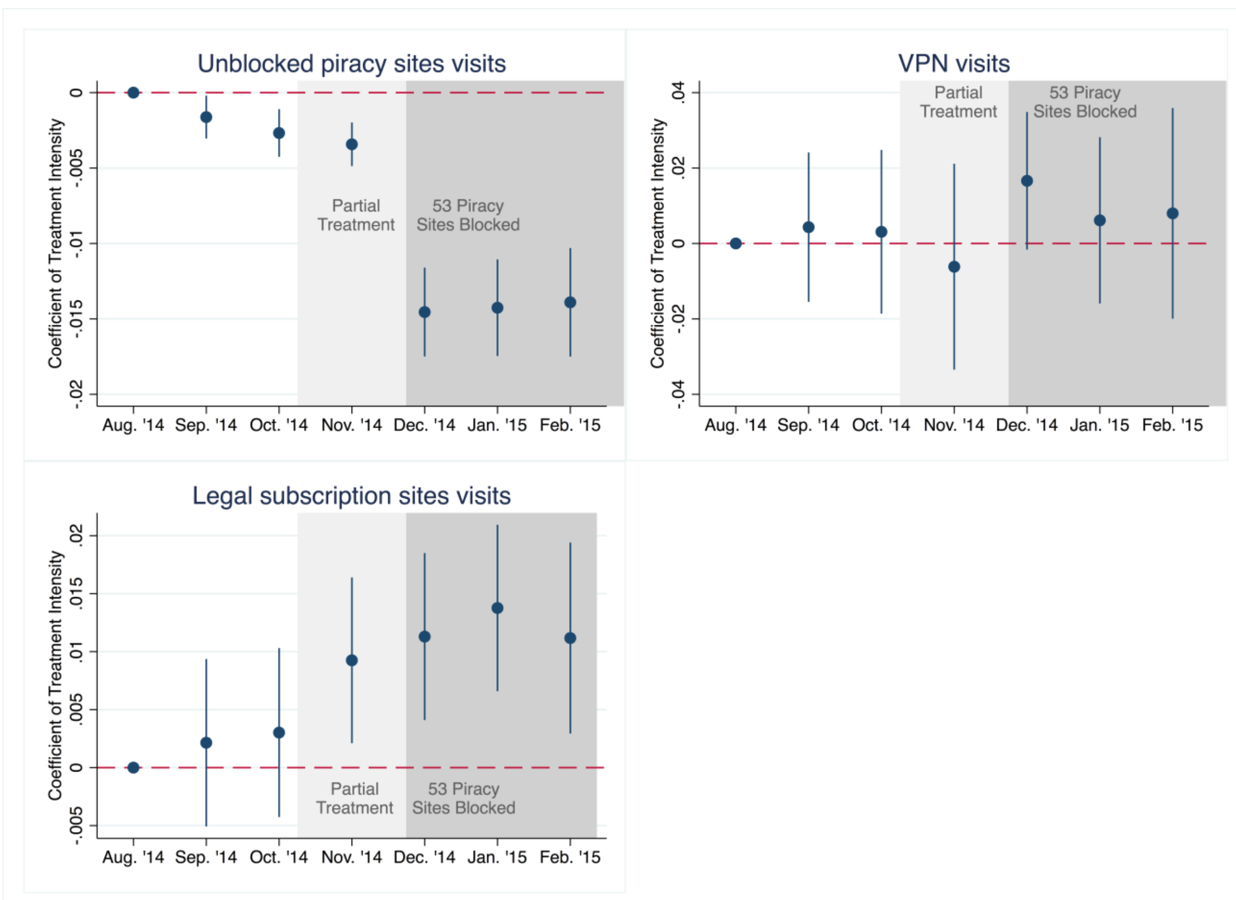

The natural next step is to estimate the size of these effects. To do so, we estimate the following model:

$$Visits_{it} = \beta_0 + \beta_1 Post_t + \beta_2 TreatmentIntensity_i + \beta_3 Post_t \cdot TreatmentIntensity_i + \beta_4 \cdot PartialTreatment_t + \beta_5 Partial_t \cdot TreatmentIntensity_i + \mu_i + \epsilon_{it} \quad (2)$$

Here we have replaced the month dummies with an indicator variable for the "partial treatment" period (November 2014) as it is a partial treatment month and then an indicator for the post period, equal to 1 for the months of December, January, and February. Under the identifying assumption, the interaction of treatment intensity and the partial treatment indicator represents the impact of the blocks on the outcome variable in the month they were being implemented and the interaction of treatment intensity with the post dummy represents the effect of the blocks during the following three months.

Table 4: Estimated Impact of 53 Site Block in November 2014 on User Site Visits

	(1)	(2)	(3)
Dependent variable:	**Unblocked torrent sites**	**VPN sites**	**Legal subscription sites**
Post treatment	-0.250***	-0.0995	0.130***
	(0.0113)	(0.0828)	(0.0178)
Treatment intensity	0.00701***	-0.0266**	-0.0180***
	(0.000608)	(0.00815)	(0.00230)
Post X treatment intensity	-0.0125***	0.00999	0.0104***
	(0.000969)	(0.00610)	(0.00229)
During treatment	-0.196***	-0.152	-0.0194
	(0.0132)	(0.111)	(0.0234)
During X treatment intensity	-0.00173***	-0.00851	0.00743**
	(0.000175)	(0.0122)	(0.00288)
Constant	0.369***	-0.0637	-0.271***
	(0.0115)	(0.113)	(0.0188)
Individual FE?	Y	Y	Y
N	46733	2546	29685
Individuals	11847	556	7095
Log-Likelihood	-88399.33	1706.17	-36720.16

Notes: Standard errors are shown in parentheses and clustered by user. + p<0.10 * p<0.05 ** p<0.01 *** p<0.001.

In the columns above, we see from the coefficient on the post treatment dummies that piracy at unblocked sites was generally decreasing (for users with 0 treatment intensity) during this time while traffic to legal subscription sites was increasing. The coefficient on Treatment intensity is estimated due to the aforementioned caveats associated with negative binomial fixed effects models – it implies that heavier users of the blocked sites made lighter use of paid streaming sites in the pre period (consistent with descriptive statistics in Table 1). While the interactions between the partial treatment dummy and treatment intensity may be interesting, we focus on analysis of the effect of the blocks once they were fully implemented – this is represented by the interaction between the post dummy and treatment intensity. We see that it is negative and significant for unblocked piracy sites, positive but insignificant for VPN sites, and positive and significant for paid legal streaming sites. In column 3, the coefficient on the interaction term is 0.0104 – this implies that an individual who visited the blocked sites one more time before the blocks increased his usage of paid legal streaming sites by 1.04% more than he otherwise would have had the blocks not affected him (i.e. were his treatment intensity zero).

It is clear from Table 1 that the panel is unbalanced as some individuals were not observed during some months. Though fixed effects models are robust to unbalanced panels, one might worry whether the months in which individuals are not observed are somehow selected with bias – for example, if individuals periodically choose to be unobserved due to their behaviors. As a check on this, in Table A2 in Appendix D we re-estimate our results using only a strictly balanced panel of individuals who are observed in all seven months of the dataset. Our results remain similar in sign and significance. [20] However, in the balanced panel, the coefficient

[20] Sign and significance are also robust to estimating our model using Poisson regression or OLS with log(visits + 1) as the outcome variable, though the magnitude of the results differ in the latter as a result of improperly fitting small count data using a linear model. See Table A3 and A4 in Appendix D.

on the interaction between post and treatment intensity for legal subscription sites is .0169. This indicates that a user with one additional pre-block visit to blocked sites increases usage of paid legal streaming by 1.69% more after the blocks then she otherwise would have. While the balanced panel may be preferred for the reason stated above, the unbalanced panel has a much larger number of observations. As well, the balanced panel may select for users who are more consistent consumers of media overall if users in the panel sometimes drop out of observation specifically when they are not visiting any sites. Because the unbalanced panel and the balanced panel each have advantages, we consider these estimates as indicating the range of possible effects of the blocks.

One concern with our identification strategy is that high treatment intensity users, though most affected by the blocks, might be the least likely to turn to legal channels due to their affinity for piracy. If this were true, it would bias our result towards zero, making it harder to find an effect. However, this is not what we observe, as more heavily treated user decrease their use of unblocked piracy sites, and disproportionately turn towards legal sites.

Because our results contrast those of Poort et. al (2014) and Aguiar et. al. (2018) we ask a second research question. We ask whether our findings can be explained by the increased deterrent effect of blocking 53 sites rather than fewer sites. To do so, we examine the blocking of 19 sites in November 2013 and also a single site in May 2012.

5.2 November 2013 Blocking of 19 Major Piracy Sites

Recall that because of privacy concerns PanelTrack would only release monthly data aggregated into consumer groups for 2012 and 2013. While using aggregate grouped data is clearly inferior to using individual data, our analysis is able to recover the impact of website blocks on legal and illegal media usage with the appropriate inferential statistics. We rely on the estimator and inference suggested by Donald and Lang (2007), who discuss methods for correcting for common group errors when the treatment is assigned at the group level such as in our data. The estimator they recommend as most efficient in many circumstances is the "between-group estimator", which is in fact a regression of grouped means against group average outcomes. The approach relies on the fact that each data point is based on the aggregated behavior of a sufficiently large underlying group and thus is measured with greater precision than if each observation were generated by one individual. This is precisely the data for the website blocks in 2012 and 2013 we have at our disposal. Inference for our estimators is given by a t_{G-2} distribution, where G indicates the number of groups.[21]

We estimate the following model:

$$LnVisits_{jt} = \gamma_0 + \gamma_1 month_t + \gamma_2 TreatmentIntensity_j \cdot month_t + \mu_j + \epsilon_{jt} \qquad (3)$$

where all terms are the same as in (1) except that the j subscript now denotes a group of consumers (as opposed to the i indexing the individual). Because our data are aggregated across groups, the resulting visits are large enough that we can estimate OLS – however, we log these data because visits are right skewed and because we expect trends across the groups to be comparable on a relative (percent) basis. γ_2 is the coefficient of interest – as a test of the parallel trends assumption, we ask if γ_2 is 0 for all months before the blocks. We plot all γ_2 coefficients below for the various outcome variables.

[21] Note this distribution is more conservative than using a t-distribution with $t_{M_1+M_2+\cdots+G-2}$ degrees of freedom that would be used were we to estimate pooled OLS. This point is made by lecture notes by Jeffrey Wooldridge (2007).

Figure 2: Effect of November 2013 Blocks on Outcomes

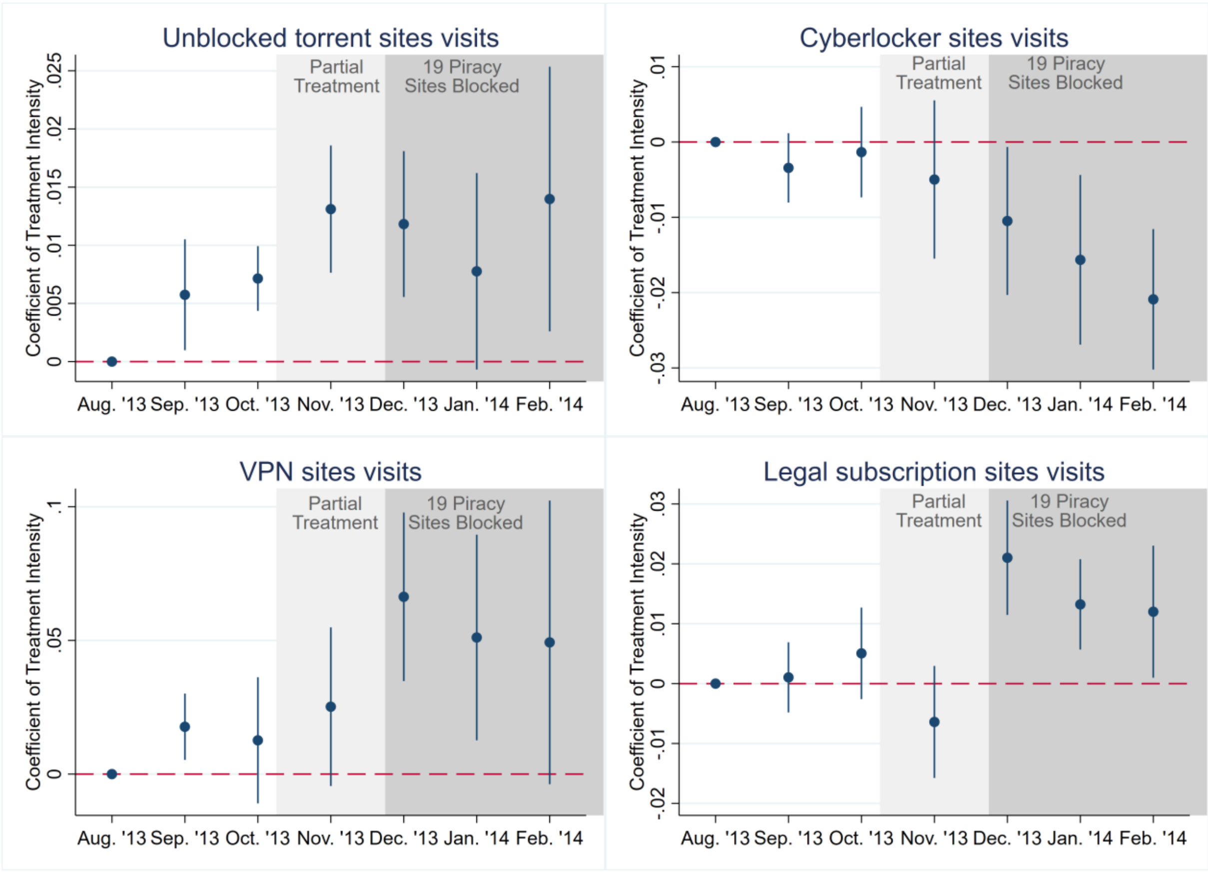

In the post period, it appears as if visits to unblocked cyberlocker piracy sites decreased as a result of the blocks while visits to legal subscription sites increased – both of these results are consistent with our results from 2014. It appears as if VPN visits as well as visits to unblocked torrent sites increased as a result of the November 2013 blocks.

While the parallel trends assumption holds for cyberlocker visits, for legal subscription, and (almost) for VPN site visits, it fails for unblocked torrent site visits. Heavier users of the blocked sites appeared to grow their usage of unblocked torrent sites more in the pre-period than lighter users. We are not certain why this is the case – it is possible that because the court case ordering the blocks actually occurred during October 2013, some users of pirate sites may have had advance knowledge of the blocks and started to rely more other sites. Some of the unblocked torrent sites may in fact be proxy sites for the blocked sites. Either way, any results for the effect of the 2013 wave of blocks for visits to unblocked torrent sites must be taken with caution as they may not be causal.

To measure the overall effect of the November 2013 blocks on the outcome variables and to determine statistical significance, we estimate the following model:

$$\begin{aligned} LnVisits_{jt} = \gamma_0 + \gamma_1 Post_t + \gamma_2 TreatmentIntensity_j \cdot Post_t + \gamma_3 PartialTreatment_t \\ + \gamma_4 Partial_t \cdot TreatmentIntensity_i + \mu_j + \epsilon_{jt} \end{aligned} \tag{4}$$

Model (4) is similar to (2) except that j indexes each group and the outcome variable is logged due to our ability to estimate using OLS.

Table 5: Estimated Impact of 19 Site Block in November 2013 on User Site Visits

	(1)	(2)	(3)	(4)
Dependent variable:	**Unblocked torrent sites**	**Cyberlockers**	**VPN sites**	**Legal subscription sites**
Post Treatment	-0.149	-0.527**	-0.0540	-0.0339
	(0.0934)	(0.121)	(0.418)	(0.0974)
Post X treatment intensity	0.00689+	-0.0141**	0.0454**	0.0134**
	(0.00324)	(0.00305)	(0.0140)	(0.00359)
Partial Treatment	0.0295	-0.214	0.456+	-0.00699
	(0.0697)	(0.122)	(0.225)	(0.0691)
Partial X treatment intensity	0.00881**	-0.00339	0.0151	-0.00843*
	(0.00186)	(0.00343)	(0.0109)	(0.00364)
Constant	13.74***	13.55***	9.927***	13.78***
	(0.0318)	(0.0472)	(0.147)	(0.0380)
User Group FE?	Y	Y	Y	Y
N	70	70	70	70
User Groups	10	10	10	10
Adjusted R2	.125	.633	.221	.255

Notes: Standard errors are shown in parentheses and clustered by user group. + p<0.10 * p<0.05 ** p<0.01 *** p<0.001

In Table 5, we see from the Post Treatment dummy that all of the outcome variables were decreasing over time, though the decreases are relatively small for VPN sites and legal subscription sites. The coefficients of interest are those on the Post x treatment intensity interaction term. Here, we see an increase usage of unblocked torrent sites (significant at alpha = 0.1), but we know from the time plot in Figure 2 that this may be an extension of pre-existing trends. We observe a causal decrease in visits to unblocked cyberlocker sites, a causal increase in visits to VPN sites, and a causal increase in visits to legal subscription sites. Thus, our results indicate that, like the blocking of 53 sites in November 2014, the blocking of 19 sites did drive some users to paid legal streaming sites and reduced at least some forms of piracy. An individual who made one more visit per month to blocked sites during the pre-period increased her monthly visits to legal subscription sites 1.34% more than she would have if not for the blocks.

We followed Donald and Lang (2007) in computing p-values when outcome variables are aggregate data from large groups and believe this to sufficiently correct for any downward bias in standard errors. However, because our number of clusters is small, we also impute even more conservative p-values using the wild cluster bootstrap approach (Cameron et. al. 2008). These p-values on the coefficients of interest can be found in Table A7 in Appendix D – there the coefficient for visits to paid legal subscription has a p-value of 0.08.

Both the 2013 and 2014 waves of blocks involved the blocking of a number of major piracy sites. Next, we ask whether the blocking of one major piracy site – an experiment more akin to those in Poort et. al. (2014) and Aguiar et. al. (2018) – demonstrates similar outcomes or produces a different set of results.

5.3 May 2012 Blocking of The Pirate Bay

The data we obtained from PanelTrack to study the blocking of The Pirate Bay in 2012 are similar to the data from 2013 – we observe outcomes by consumer group by month, generated from a balance panel of consumers observed in all months. Again, PanelTrack sorted consumers into groups based on pre-block usage of the blocked site – in this case, The Pirate Bay.

We estimate model (3) for each of the outcome variables and plot the coefficients of interest for the models in figure 3.

Figure 3: Effect of May 2012 Pirate Bay Block on Outcomes

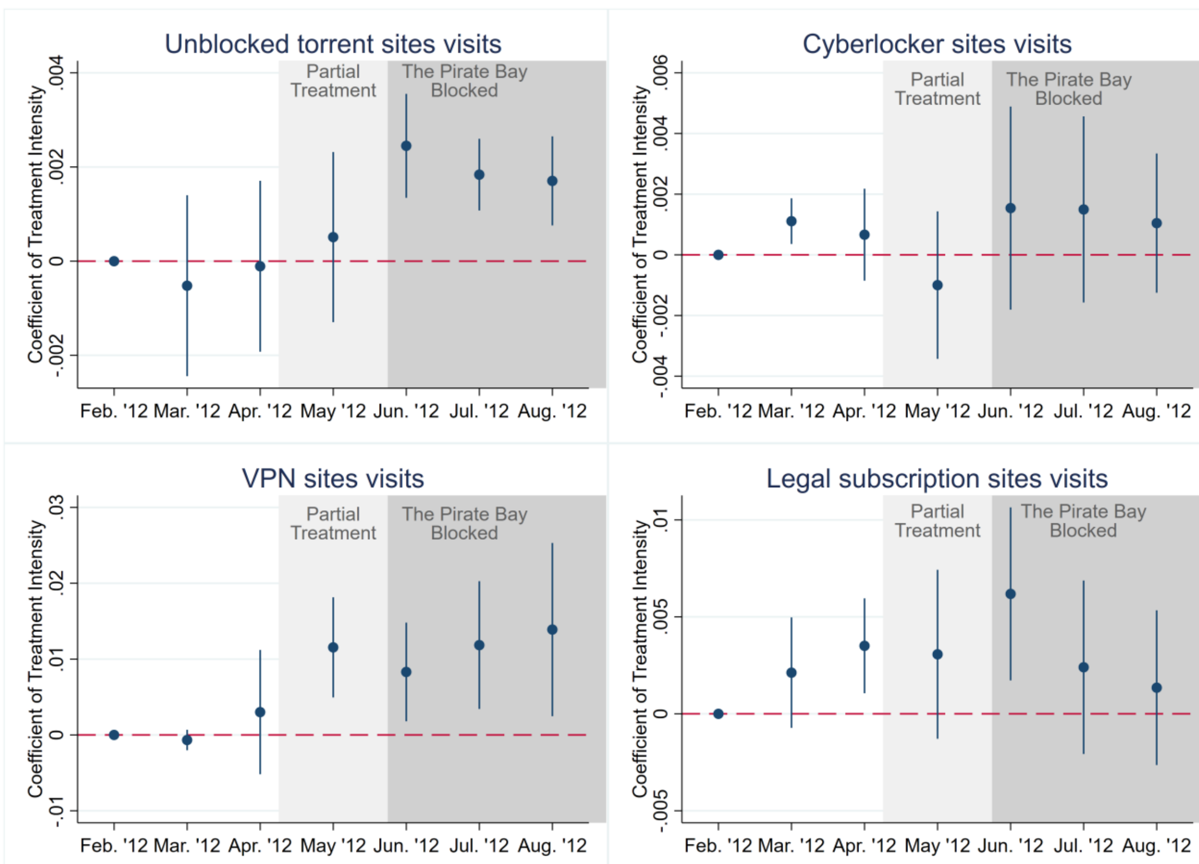

The 2012 blocking of the Pirate Bay appears to cause an increase in visits to unblocked torrent sites as well as visits to VPN sites. We observe no clear effect on visits to cyberlockers and while we may see an increase in usage of legal subscription sites in the month after the blocks, it disappears by the second and third months after the blocks. The parallel trends assumption appears to hold for visits to unblocked torrent sites, VPN sites, and (nearly) for unblocked cyberlockers. However, the parallel trends assumption fails for visits to legal subscription sites as treatment intensity appears positively correlated with changes in visits to subscription sites. There is a compelling explanation for this fact – one of the paid streaming sites, Netflix, was introduced to the UK in January 2012, and it quickly became popular due to the fame of the brand. During the initial adoption period, we argue that people who were pirating a lot of content (relative to people who were pirating little) are more likely to have an initial interest in Netflix and therefore subscribe. This would explain the elevated γ_2 coefficients in March and April. The direction of the pre-existing trend would actually suggest that the correlation should have increased in the post period, and instead we see it remain flat or decrease other than in June 2012. Thus we conclude that blocking The Pirate Bay caused no lasting increase in paid legal consumption and at most a temporary one month increase.

We estimate (4) for each of the outcome variables and present the results below[22].

[22] As with the 2013 blocks, we present in Table A8 in Appendix D the same estimates but with standard errors estimated using the wild cluster bootstrap approach. The increase in visits to unblocked torrent sites remains significant with a p-value of 0.036.

Table 6: Estimated Impact of The Pirate Bay Block in May 2012 on User Site Visits

	(1)	(2)	(3)	(4)
Dependent variable:	**Unblocked torrent sites**	**Cyberlockers**	**VPN sites**	**Legal subscription sites**
Post	-0.207**	-0.388+	-0.962+	-0.576+
	(0.0434)	(0.173)	(0.508)	(0.292)
Post X treatment intensity	0.00221***	0.000769	0.0106**	0.00143
	(0.000363)	(0.000935)	(0.00261)	(0.00148)
Partial treatment	-0.312**	-0.340+	-1.085*	-0.707+
	(0.0811)	(0.160)	(0.392)	(0.322)
Partial X treatment intensity	0.000720	-0.00159+	0.0108**	0.00119
	(0.000436)	(0.000835)	(0.00233)	(0.00166)
Constant	14.67***	13.85***	9.897***	11.90***
	(0.0216)	(0.0785)	(0.217)	(0.141)
User Group FE?	Y	Y	Y	Y
N	70	70	70	70
User Groups	10	10	10	10
Adjusted R2	0.274	0.288	0.050	0.203

Notes: Standard errors are shown in parentheses and clustered by user group. + p<0.10 * p<0.05 ** p<0.01 *** p<0.001

In Table 6, the post treatment coefficient indicates that most of the response variables are decreasing during this time in general. This is not surprising as May indicates the end of many television seasons and thus interest in digital video content tends to decline each year. However, the coefficients of interest are those for Post x treatment intensity. We observe a statistically significant increase in usage of other unblocked torrent sites such that a person visiting The Pirate Bay one more time during the pre-period increased her usage of other torrent sites 0.22% more than she would have after the blocks if she had not been using The Pirate Bay. We observe a statistically significant increase in usage of VPN sites, indicating that some heavy users of the Pirate Bay turned to using a VPN to circumvent the blocking of the site. However, the economic significance of this may be small as the constant term here is small – visits to VPN sites in the data are relatively low. Finally, the coefficient for paid legal streaming sites is positive but small and insignificant. Typically, this might lead to an inconclusive interpretation – is the increase positive or 0? From Figure 3 we know that any increase in the post period is driven entirely by the first month, after which it disappears. And we also know that even this first month effect may be the result of a pre-existing trend. Thus, like Aguiar et. al. (2018) we find no lasting causal effect of the May 2012 blocking of The Pirate Bay on legitimate consumption.

5.4 Summary of Empirical Results

In summary, we found that the 2014 blocking of 53 major piracy sites not only decreased visits to the blocked sites but also caused a decrease in usage of other unblocked piracy sites. As a result, we observe that it causally increased usage of paid legal streaming sites, implying that supply-side antipiracy enforcement can be effective in turning users of illegal piracy channels toward paid legal consumption. In November 2013 when 19 major piracy sites were blocked, our evidence is inconclusive as to the effect on visits to unblocked torrent sites but we observe a causal decrease in visits to unblocked cyberlocker sites. We also observe a statistically significant increase in usage of paid legal streaming sites. Because our results contrast those from prior work that studied the blocking / shutting down of one site, we also examined the May 2012 blocking of The Pirate Bay. There, our results were largely consistent with the literature – we found that consumers increased their usage of alternate unblocked piracy sites or VPN sites to circumvent the blocks, and we found no evidence of a lasting increase in usage of paid legal sites. In the next and final section, we provide some interpretations and implications of our results, and we estimate the magnitude of the of causal increases in legal consumption.

6. Discussion

While the use of supply side antipiracy actions has increased greatly in recent years as a tool in the fight against intellectual property theft, there are only a few studies that have empirically analyzed their effectiveness in changing user behavior. These existing studies each only examined the impact of a single intervention against a single site, and evidence was mixed as to whether the interventions were effective. By analyzing the blocking of a single major piracy site in the UK in 2012, we confirm the findings from previous research that removing access to content through a single site is ineffective in increasing legal consumption because pirates simply find other ways to access pirated content. But we demonstrate that the effect of supply-side anti-piracy policies are more nuanced – removing access to content through a number of the most popular sites causes decreases in piracy and increases in usage of paid legal channels. We suggest that differences in the effectiveness of various supply-side interventions are likely due to differences in how costly or inconvenient they make piracy, a finding that is consistent with theory from Dey et. al. (2018) but which had not yet been empirically tested.

One objection to the causal interpretation of our results might be that legal services could have started advertising their services more heavily around the time of the blocks in 2014 and 2013. First, our difference-in-difference model is able to capture common time trends through the time fixed effects, so this would only be a concern if legal services could somehow target high piracy individuals more than low piracy individuals with such advertisements. Second, we observe a lack of pre-existing time differential time trends, and so this counter explanation is only relevant if legal services started targeting heavier users of the blocked sites (and not lighter users) with increased advertising, and they did so exactly at the timing of both the 2013 and the 2014 blocks. The discrete jumps in legal consumption following each of these blocks were not followed by continuing upward trends, and so the timing of the correlation between treatment intensity and discrete changes in legal consumption is telling. While no quasi-experimental claims of causality are ever 100% perfect, we suggest that alternative explanations for our results are unlikely.

6.1 Possible Mechanisms

While it is easy to understand why blocking a number of sites at once might cause enough inconvenience to push some consumers toward legal channels, it is not clear why these blocks would cause decreases in piracy at other unblocked sites. Yet this is what we observed for unblocked piracy cyberlockers in 2013 and for unblocked piracy sites in general in 2014. We propose several possible explanations. First, seeing a number of the most popular sites blocked at the same time may have caused an overall chilling effect for pirates – an individual who was making use of the blocked sites and some of the unblocked sites may have elected to abandon piracy once it was clear that legal actions against piracy were strengthening (indeed, Danaher et. al. (2014) found that the HADOPI law in France began to affect consumer behavior when it became salient to French consumers, even before it actually went into effect). Second, we know that some users of blocked sites increased their user of legal subscription sites as a result of the blocks. Legal subscription sites have an upfront fixed cost and so once these pirates paid the legal subscription price, they may have found piracy at unblocked sites less appealing than all of the new legal content available to them at zero marginal cost. Finally, it could be that some of the unblocked sites were partly dependent on users being able to access the blocked sites in order to function properly (for example, this could be true for piracy sites that linked to content on the blocked sites). Our data do not allow us to determine the mechanism causing the decrease in usage of unblocked piracy sites, but the hypotheses above may be interest areas for future research.

6.2 Aggregate Impact of Website Blocking

While the effect of the 2013 and 2014 waves of blocks on legal channels were statistically significant, it is important to ask whether they were economically significant. In 2014, we start with each individual's observed post-treatment visits to paid legal subscription sites. We estimate their counterfactual post-treatment visits to these subscription sites by predicting what they would have been if treatment intensity were zero (our estimate of the counterfactual, or what they would have been had the individual not been affected by the blocks). We aggregate the difference across all individuals between observed visits to legal sites and counterfactual visits to determine the total causal uplift in visits to legal subscription sites and divide this by the total counterfactual visits to get the overall percent increase. If we use the coefficient estimate from the unbalanced panel estimates in Table 4 (0.0104), we find that users of the blocked sites in 2014 increased their usage of legal subscription sites by 7% relative to what they would have done in the absence of the blocks. If we use the coefficient estimate from the balanced panel estimates in Table A2 (0.0169), we find that this increase was 12%. As both the balanced and unbalanced panels have strengths and weaknesses (discussed in Section 5), we suggest that the effect of the 2014 blocks on the usage of legal sites by treated users was somewhere between 7% and 12%.

Performing the same analysis for the 2013 blocks (but at the group level rather than the individual), we find that on average the blocks caused treated users to increase their visits to paid subscription sites by 8% relative to what they would have done if not for the blocks. Thus both the 2014 wave and 2013 wave appear to have had similar impacts on legal consumption.

There are several limitations to this study. First, we were only able to study legal consumption of media through paid legal subscription sites. Users may consume media legally in other ways, such as by digital purchase/rental, physical purchase/rental, or legal free ad-supported viewing channels. Because PanelTrack observes clickstream data but not actual e-commerce, we cannot infer a la carte purchases or rentals (people visit a site like Amazon.com for many reasons other than purchasing movies or television).[23] Second, because the blocks were not implemented by all ISP's and may only have been perfectly implemented by participating ISP's some time into the "after" period, our results may underestimate the true effect of website blocking on legal consumption. Third, we only observe three months after each wave of blocks, and thus we do not know how long our measured impacts lasted. Although the effects on legal subscription visits appeared persistent in our data, it remains possible that increases in legal consumption caused by the blocks fade over time as consumers eventually identify and grow to trust 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.

Given the accumulated evidence, how should policymakers view supply-side interventions to curb illegal piracy? We consider by analogy the Greek myth of the Hydra, the mythical, multi-headed beast. The Hydra is one of most difficult animals to kill in Greek mythology. Decapitating any single one of its heads only results in several more growing back to replace it, an

[23] PanelTrack does also track when a user opens the iTunes application on their computer, a common channel for purchasing digital media. However, many things (including plugging in one's devices) cause an app like this to open and so we chose not to purchase these data from PanelTrack.

excellent analogy for our results and those of prior researchers. It is only when a sword is plunged into its heart that it dies. Removing the source of the pirated content stored in cyberlockers and linked to by many other sites is akin to stabbing the Hydra in the heart (and akin to shutting down Megaupload.com); this is effective but may not always be feasible. Blocking a single site is akin to decapitating only one of the Hydra's heads. The result will only be a more diffuse network of piracy sites, with no curb on pirating activity. Blocking multiple sites at once is akin to decapitating several of the Hydra's heads. With the network of sites significantly disrupted, this could possibly be a mortal wounding. We have shown that users' behavior is sufficiently disrupted and switching costs are high enough such that some increase the use of legal channels, and reduce illegal ones.

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 and Martins, B. 2016. Digital Music Consumption on the Internet: Evidence from Clickstream Data. *Information Economics and Policy*, 34, pp. 27-43.

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.

Cameron, C., Gelbach, J., and Miller, D. 2008. Bootstrap-Based Improvements for Inference with Clustered Standard Errors. Review of Economics and Statistics. 90(3), pp. 414-427.

Chen, P., and Hitt, L. M. 2002. Measuring Switching Costs and the Determinants of Customer Retention in Internet-Enabled Businesses: A Study of the Online Brokerage Industry." *Information Systems Research*, 13 (3), pp. 255-274.

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.

Danaher, Brett, Michael D. Smith, and Rahul Telang. 2017. Copyright Enforcement in the Digital Age: Empirical Evidence and Policy Implications. *Communications of the ACM*, Vol 60(2), pp 68-75.

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

Donald, S.G. and Lang, K., 2007. Inference with difference-in-differences and other panel data. *The Review of Economics and Statistics*, 89(2), pp.221-233.

Goldfarb, A. 2006. State dependence at internet portals. *Journal of Economics & Management Strategy*. 15 (2), 317{352.

Hausman, J., Hall, BH., and Griliches, Z. 1984. "Econometric models for count data with an application to the patents-R&D relationship." *Econometrica*, 52: 909-938

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

Liebowitz, S. 2008. Testing File-Sharing's Impact on Music Album Sales in Cities. Management Science. 54(4).

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.

McKenzie, J. (2017). Graduated Response Policies to Digital Piracy: Do They Increase Box Office Revenues of Movies?" Information Economics and Policy, 38, 1-36.36

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*.

Wooldridge, Jeffrey. "What's New in Econometrics? Lecture 10 Difference-in-Differences Estimation". NBER Summer Institute, 2007. Available at: http://www.nber.org/WNE/Slides7-31-07/slides_10_diffindiffs.pdf

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)

Appendix A: List of Blocked Sites in Each Wave

2012 Block Wave – The Pirate Bay Site Only
thepiratebay.se

2013 Block Wave – 19 Unique Video Piracy Sites
In 2013, the following 19 sites were blocked. A number of known mirror sites with similar domains (different suffixes) were blocked along with these. For example, while torrentz.eu was ordered blocked, ISP's were also ordered to block sites such as torrentz.net when it was verified that they contained the same content.

1337x.org	rapidlibrary.com	torrentz.eu
bitsnoop.com	solarmovie.us	torrentz.eu
Extratorrent.cc	torrentcrazy.com	tubeplus.me
filecrop.com	torrentdownloads.me	vodly.com
filestube.com	torrenthound.com	watchfreemovies.com
monova.org	torrentreactor.net	yify-torrents.com
primewire.net		

2014 Block Wave – 53 Unique Video Piracy Sites Plus Known Mirrors

nowtorrents.com	btloft.com	iwannawatch.to
torrentdb.li	picktorrent.com	warez-bb.org
watchseries.to	seedpeer.me	icefilms.info
heroturko.me	torlock.com	Tehparadox.com
torrentbytes.net	torrentbit.net	scnsrc.me
seventorrents.org	torrentdownload.ws	rapidmoviez.com
tormovies.org	torrentexpress.net	isohunt.to
bts.to	torrentfunk.com	torrentz.pro
limetorrents.com	torrentproject.com	torrentbutler.eu
torrentus.eu	torrentroom.com	iptorrents.com
vitorrent.org	torrents.net	sumotorrent.sx
movie25.cm	torrentz.cd	torrentday.com
iwatchonline.to	torrentzap.com	torrenting.com
losmovies.com	watchseries.lt	bitsoup.me
torrents.fm	stream-tv.me	yourbittorrent.com
bittorrent.am	watchserieshd.eu	demonoid.ph
btdigg.org	cucirca.eu	torrent.cd
vertor.eu	rarbg.com	

Appendix B – Creation of Illegal and Legal Site Lists

While the list of blocked sites in each wave were publicly available based on court orders and summarized (with citations) on Wikipedia[24], an important question is how we created the lists of legal sites and other unblocked piracy sites that we provided to PanelTrack in order for them to provide clickstream visits.

The list of legal sites for each wave of blocks was relatively easy to put together based on a combination of Internet research for available legal services as well as conversations with film and television industry contacts.

However, the list of piracy websites was more difficult since such sites do not necessarily advertise themselves. And it is important that we capture the majority of piracy activity in this list or else it could be that pirates thwarted from the blocked sites turn to other unblocked sites and we would not observe it. As such, we went through a multi-step process to determine the set of unblocked piracy sites available in the UK during each wave of blocks.

1. We collected lists of potential piracy sites from various sources, including
    a. The City of London Police's "infringing website" list, available at https://www.cityoflondon.police.uk/advice-and-support/fraud-and-economic-crime/pipcu/Pages/Operation-creative.aspx
    b. The Google Transparency Report provides lists of websites with removal requests due to copyright infringement, available here: https://www.cityoflondon.police.uk/advice-and-support/fraud-and-economic-crime/pipcu/Pages/Operation-creative.aspx
    c. The MPAA's "Online Notorious Market Report", provided to us directly by the MPAA.
2. We then also asked PanelTrack for their list of piracy sites. PanelTrack categorizes many of the sites that show up in their clickstream data and one of the categories is piracy.
3. After merging #1 and #2, we connected to a UK VPN (in order to appear to be accessing websites from the UK) and attempted to connect to each site to determine if it truly was a site mostly dedicated to piracy, and for 2012 and 2013 to determine whether it was a torrent/P2P site or a cyberlocker/link site. We removed any sites that were not piracy sites, or that were clearly dedicated only to music, anime, adult, games, or eBook content since the blocked sites and legal sites that we studied were largely dedicated to standard television and film content.
4. Finally, we provided the resulting lists to PanelTrack (for each wave, this list was hundreds of sites) who then confirmed based on their data that most of the piracy visits were concentrated within a small number of these sites. We therefore think it unlikely that we missed major relevant piracy sites.

[24] Accessible at https://en.wikipedia.org/wiki/List_of_websites_blocked_in_the_United_Kingdom#Court_ordered_implementations_targeting_copyright_and_trademark_infringement

Appendix C: Explanation of Consumer Group Formations in 2012 and 2013

As described in the data section, our datasets for 2012 and 2013 include aggregated observations for groups of consumers (each month) rather than individual level data. While this is not ideal, at the time it was mandated by PanelTrack due to issues of privacy and the topic of our research (piracy).

If PanelTrack had randomly assigned consumers in their panel into groups, we would expect very little variation in treatment intensity (pre-blocked visits to blocked sites) across group due to the central limit theorem. Thus, we asked that PanelTrack ensure variation in treatment intensity by bucketing consumers into bins based on their pre-block usage of blocked sites in the three months before the blocks. Instead, PanelTrack bucketed consumers based on their visits to blocked sites in the first month of the study (for example, for the 2012 data, consumers were placed into groups based on February 2012 visits to thepiratebay.se). We consider average visits to blocked sites in the three months before the blocks to be a more accurate representation of a consumer's pre-block behavior than just one month's worth of visits, and so we use the former as our measure of treatment intensity for each group.

One consequence of this is that there is no group with 0 treatment intensity, because some consumers who didn't make visits to the blocked sites in the first month of the panel still made some in the second and third months, leading to a treatment intensity greater than 0. We note that even were there a group with 0 treatment intensity, this group would not truly be a pure "control" group as it remains possible that some individuals who did not visit blocked sites before the blocks would have attempted to afterward, and thus would have been treated by the blocks. Pre-block usage of blocked sites is merely a measure of the "bite" of the treatment on each group, and a zero-value group is unnecessary for identification in this generalized form of the difference-in-differences model.

Appendix D: Supporting Tables and Figures

Table A1: Monthly Impact of November 2014 Site Blocks

	(1)	(2)	(3)
Dependent variable:	**Unblocked torrent sites**	**VPN sites**	**Legal subscription sites**
Month=0	0	0	0
	(.)	(.)	(.)
Month=1	-0.0104	-0.0709**	-0.00425
	(0.0498)	(0.0265)	(0.192)
Month=2	0.0156	-0.00328	-0.144
	(0.0496)	(0.0262)	(0.206)
Month=3	0.0123	-0.192***	-0.258
	(0.0488)	(0.0278)	(0.211)
Month=4	0.180***	-0.253***	-0.0687
	(0.0467)	(0.0285)	(0.197)
Month=5	0.301***	-0.250***	0.0675
	(0.0455)	(0.0284)	(0.192)
Month=6	0.113*	-0.343***	-0.488*
	(0.0446)	(0.0270)	(0.207)
Treatment intensity	-0.0195***	0.0165***	-0.0402*
	(0.00483)	(0.00145)	(0.0203)
Month=0 # Treatment intensity	0	0	0
	(.)	(.)	(.)
Month=1 # Treatment intensity	-0.00536	-0.000796	0.0135
	(0.00676)	(0.00140)	(0.0222)
Month=2 # Treatment intensity	-0.00224	-0.00180	-0.0201
	(0.00642)	(0.00142)	(0.0379)
Month=3 # Treatment intensity	0.0108+	-0.00735***	0.0136
	(0.00560)	(0.00160)	(0.0232)
Month=4 # Treatment intensity	0.0145**	-0.0128***	0.0245
	(0.00528)	(0.00188)	(0.0218)
Month=5 # Treatment intensity	0.0120*	-0.0146***	0.0242
	(0.00529)	(0.00193)	(0.0211)
Month=6 # Treatment intensity	0.0122*	-0.0147***	0.0259
	(0.00547)	(0.00197)	(0.0215)
Constant	-0.432***	0.417***	-0.138
	(0.0401)	(0.0242)	(0.183)
Individual FE?	Y	Y	Y
N	11152	15369	1128
Individuals	1837	2549	176

Notes: Standard errors are shown in parentheses and clustered by user. + $p<0.10$ * $p<0.05$ ** $p<0.01$ *** $p<0.001$

Figure A1: 2014 Blocks - Estimates of Model (1) Using Balanced Panel

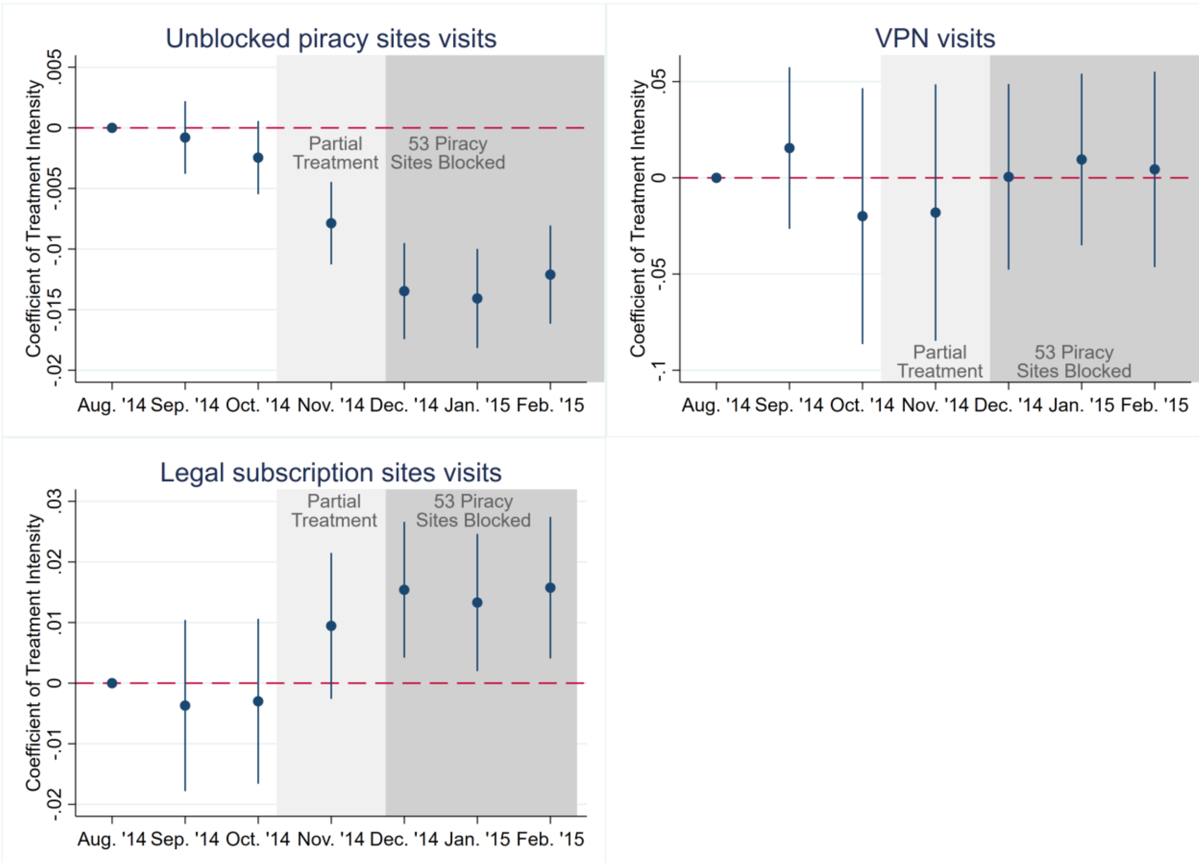

Table A2: 2014 Blocks - Negative Binomial Estimates of Model (2) Using Balanced Panel

	(1)	(2)	(3)
Dependent variable:	**Unblocked torrent sites**	**VPN sites**	**Legal subscription sites**
Post treatment	-0.267***	-0.0452	0.155***
	(0.0201)	(0.140)	(0.0338)
Partial treatment	0.131***	-0.0533	-0.0902+
	(0.0288)	(0.220)	(0.0470)
Partial X treatment intensity	0.00533**	-0.0226	-0.00536
	(0.00171)	(0.0307)	(0.00423)
Treatment intensity	0.0137***	-0.0477+	-0.0247***
	(0.00134)	(0.0253)	(0.00363)
Post X treatment intensity	-0.0121***	0.00197	0.0169***
	(0.00121)	(0.0137)	(0.00350)
Constant	0.487***	-0.0678	-0.310***
	(0.0239)	(0.193)	(0.0378)
Individual FE?	Y	Y	Y
N	9478	826	7133
Individuals	1354	118	1019
Log-Likelihood	-25326.04	-571.30	-11136.01

Notes: Standard errors are shown in parentheses and clustered by user. + p<0.10 * p<0.05 ** p<0.01 *** p<0.001

Table A3: 2014 Blocks - OLS Estimates of Model (2) With Log(Visits + 1) as Outcome

	(1)	(2)	(3)
Dependent variable:	**Unblocked torrent sites**	**VPN sites**	**Legal subscription sites**
Post treatment	-0.175***	-0.00188	0.0466***
	(0.0100)	(0.00159)	(0.00803)
During treatment	-0.129***	-0.00170	-0.00493
	(0.0116)	(0.00127)	(0.00712)
During X treatment intensity	-0.00808**	-0.000134	0.00103+
	(0.00254)	(0.000120)	(0.000612)
Post X treatment intensity	-0.0229***	0.000372	0.00257**
	(0.00209)	(0.000278)	(0.000890)
Constant	1.334***	0.0182***	0.426***
	(0.00302)	(0.000471)	(0.00247)
Individual FEs?	Y	Y	Y
N	67098	67098	67098
Individuals	24620	24620	24620
r2	0.0386	0.000337	0.00260

Notes: Standard errors are shown in parentheses and clustered by user. + p<0.10 * p<0.05 ** p<0.01 *** p<0.001

Table A3: 2014 Blocks – Poisson Estimates of Model (2)

Dependent variable:	(1) **Unblocked torrent sites**	(2) **VPN sites**	(3) **Legal subscription sites**
Post treatment	-0.262***	-0.284***	-0.0213**
	(0.00390)	(0.0314)	(0.00678)
During treatment	-0.153***	-0.0995*	-0.0509***
	(0.00425)	(0.0414)	(0.00868)
During X treatment intensity	-0.00208***	-0.0296**	0.00297**
	(0.0000563)	(0.0111)	(0.00107)
Post X treatment intensity	-0.00766***	0.0187***	0.00817***
	(0.000146)	(0.00432)	(0.000795)
Individual FEs?	Y	Y	Y
N	46733	2546	29685
Individuals	11847	556	7095
Log-likelihood	-178247.95	-2846.39	-68904.83

Notes: Standard errors are shown in parentheses and clustered by user. + p<0.10 * p<0.05 ** p<0.01 *** p<0.001

Table A5: Monthly Impact of November 2013 Site Blocks

	(1)	(2)	(3)	(4)
Dependent variable:	**Unblocked torrent sites**	**Cyberlockers**	**VPN sites**	**Legal subscription sites**
month=1	0	0	0	0
	(.)	(.)	(.)	(.)
month=2	-0.253**	-0.311**	-0.533*	-0.283*
	(0.0766)	(0.0808)	(0.174)	(0.0944)
month=3	-0.310***	-0.371**	-0.407	-0.138
	(0.0436)	(0.0939)	(0.255)	(0.123)
month=4	-0.158	-0.441*	0.143	-0.147
	(0.0937)	(0.163)	(0.242)	(0.115)
month=5	-0.211*	-0.683**	-0.681	-0.270
	(0.0897)	(0.153)	(0.398)	(0.155)
month=6	-0.273*	-0.634**	-0.208	0.0751
	(0.120)	(0.193)	(0.503)	(0.112)
month=7	-0.527**	-0.947***	-0.212	-0.327*
	(0.135)	(0.165)	(0.722)	(0.129)
month=1 # treatintensity	0	0	0	0
	(.)	(.)	(.)	(.)
month=2 # treatintensity	0.00575*	-0.00344	0.0177*	0.00105
	(0.00211)	(0.00204)	(0.00548)	(0.00259)
month=3 # treatintensity	0.00714***	-0.00134	0.0126	0.00506
	(0.00123)	(0.00265)	(0.0104)	(0.00338)
month=4 # treatintensity	0.0131***	-0.00498	0.0252+	-0.00640
	(0.00242)	(0.00464)	(0.0131)	(0.00413)
month=5 # treatintensity	0.0118**	-0.0105*	0.0663**	0.0210***
	(0.00277)	(0.00434)	(0.0139)	(0.00422)
month=6 # treatintensity	0.00776+	-0.0156*	0.0511*	0.0132**
	(0.00373)	(0.00498)	(0.0170)	(0.00333)
month=7 # treatintensity	0.0140*	-0.0209***	0.0493+	0.0120*
	(0.00503)	(0.00412)	(0.0235)	(0.00487)
Constant	13.88***	13.80***	10.13***	13.90***
	(0.0405)	(0.0751)	(0.193)	(0.0615)
User Group FEs?	Y	Y	Y	Y
N	70	70	70	70
Individuals	10	10	10	10
Adjusted R2	.350	.783	.166	.513

Notes: Standard errors are shown in parentheses and clustered by user. + $p<0.10$ * $p<0.05$ ** $p<0.01$ *** $p<0.001$

Table A6: Monthly Impact of May 2012 The Pirate Bay Block

	(1)	(2)	(3)	(4)
Dependent variable:	**Unblocked torrent sites**	**Cyberlockers**	**VPN sites**	**Legal subscription sites**
month=1	0	0	0	0
	(.)	(.)	(.)	(.)
month=2	-0.160	-0.219***	-0.941***	-0.237
	(0.100)	(0.0451)	(0.114)	(0.181)
month=3	-0.232+	-0.156	-2.549**	-0.237
	(0.103)	(0.0985)	(0.590)	(0.205)
month=4	-0.442**	-0.465*	-2.248**	-0.865*
	(0.107)	(0.193)	(0.490)	(0.380)
month=5	-0.342***	-0.587*	-1.493*	-0.874+
	(0.0644)	(0.241)	(0.496)	(0.396)
month=6	-0.288***	-0.495+	-2.084**	-0.835+
	(0.0565)	(0.261)	(0.588)	(0.393)
month=7	-0.385**	-0.455*	-2.799*	-0.495
	(0.0840)	(0.175)	(1.106)	(0.341)
month=1 # treatintensity	0	0	0	0
	(.)	(.)	(.)	(.)
month=2 # treatintensity	-0.000522	0.00111**	-0.000671	0.00212
	(0.000849)	(0.000333)	(0.000596)	(0.00126)
month=3 # treatintensity	-0.000108	0.000664	0.00301	0.00351*
	(0.000802)	(0.000670)	(0.00362)	(0.00108)
month=4 # treatintensity	0.000510	-0.000999	0.0115**	0.00307
	(0.000799)	(0.00107)	(0.00292)	(0.00193)
month=5 # treatintensity	0.00245***	0.00154	0.00831*	0.00618*
	(0.000488)	(0.00148)	(0.00288)	(0.00197)
month=6 # treatintensity	0.00184***	0.00150	0.0118*	0.00240
	(0.000337)	(0.00136)	(0.00372)	(0.00198)
month=7 # treatintensity	0.00170**	0.00104	0.0139*	0.00135
	(0.000418)	(0.00101)	(0.00505)	(0.00176)
Constant	14.81***	13.96***	11.03***	11.98***
	(0.0448)	(0.105)	(0.256)	(0.184)
User Group FEs?	Y	Y	Y	Y
N	70	70	70	70
Individuals	10	10	10	10
Adjusted R2	.319	.231	.234	.169

Notes: Standard errors are shown in parentheses and clustered by user. + p<0.10 * p<0.05 ** p<0.01 *** p<0.001

Table A7: 2013 Estimates of Model (4) Using Wild Cluster Bootstrap Standard Errors

	(1)	(2)	(3)	(4)
Dependent variable:	**Unblocked torrent sites**	**Cyberlockers**	**VPN sites**	**Legal subscription sites**
Post X treatment intensity	0.00689	-0.0141	0.0454	0.0134+
	[0.205]	[0.188]	[0.132]	[0.082]

Notes: All other variables from model (4) were estimated but suppressed. Table contains only the estimates for the coefficient of interest. P-values computed using wild cluster bootstrapping of standard errors are displayed in brackets below the estimates.

Table A7: 2012 Estimates of Model (4) Using Wild Cluster Bootstrap Standard Errors

	(1)	(2)	(3)	(4)
Dependent variable:	**Unblocked torrent sites**	**Cyberlockers**	**VPN sites**	**Legal subscription sites**
Post X treatment intensity	0.00221*	0.000769	0.0106	0.00143
	[0.036]	[0.469]	[0.193]	[0.355]

Notes: All other variables from model (4) were estimated but suppressed. Table contains only the estimates for the coefficient of interest. P-values computed using wild cluster bootstrapping of standard errors are displayed in brackets below the estimates.