


Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/authorsrights>

Contents lists available at [ScienceDirect](#)

Telecommunications Policy

URL: www.elsevier.com/locate/telpol

Baywatch: Two approaches to measure the effects of blocking access to The Pirate Bay



Joost Poort^{a,*}, Jorna Leenheer^b, Jeroen van der Ham^c, Cosmin Dumitru^c

^a Institute for Information Law, University of Amsterdam, The Netherlands

^b CentERdata, Tilburg University, The Netherlands

^c System and Network Engineering, University of Amsterdam, The Netherlands

ARTICLE INFO

Available online 3 February 2014

Keywords:

Unauthorised file sharing

Piracy

p2p

BitTorrent monitoring

Blocking access

The Pirate Bay

Online copyright enforcement

ABSTRACT

In the fight against unauthorised sharing of copyright protected material, Dutch Internet Service Providers have been summoned by courts to block their subscribers' access to The Pirate Bay and related sites. This paper studies the effectiveness of this approach towards online copyright enforcement, using both a consumer survey and a newly developed non-infringing technology for BitTorrent monitoring. While a small group of respondents download less from illegal sources or claim to have stopped doing so, no impact is found on the percentage of the Dutch population downloading from illegal sources. Slight changes are found on the distribution of Dutch peers, but these seem related to the awareness raised by blocking rather than the blocking itself.

© 2013 Elsevier Ltd. All rights reserved.

1. Introduction

In early 2012, Dutch rights holders claimed two potentially important legal victories in their fight against unauthorised sharing of copyright protected material on the Internet, also known as online piracy. As from February 2012, two large Internet Service Providers (ISPs) were ordered by the [Court of The Hague \(2012a\)](#), a lower Dutch court, to block access to The Pirate Bay (TPB) website and a list of subdomains and mirror sites. In a second ruling in May 2012, the same court ordered four other Dutch ISPs to block access to TPB within ten days ([Court of The Hague, 2012b](#)). Both rulings combined imply that more than 90% of Dutch Internet subscribers cannot access TPB directly through their ISP. Both rulings are currently under appeal from the ISPs. They are part of a manifold of legal actions against TPB in Sweden, Germany and other countries. According to the Dutch court, TPB is currently the world's largest index site for BitTorrent files and as such an important platform for online piracy. Other legal efforts to take down this site have failed so far. The court considered blocking access to TPB for all subscribers of these ISPs proportional, as an estimated 90% to 95% of the material offered via this site is illegal while legally offered material can also be obtained through other sites. If this situation were to change, withdrawal of the ruling could be ordered ([Court of The Hague, 2012a](#), nr. 4.27–4.29).

The effectiveness of the requested measures was an important issue in both lawsuits. Rights holders' representatives presented evidence from Italy and Denmark that blocking access to TPB had significantly reduced its number of unique visitors, despite the claim by the defendants that the intervention is easily circumvented, for instance by making use of *virtual hosting* or an anonymous web proxy provider ([Court of The Hague, 2012a](#), nr. 4.34–4.36). From an economic

* Correspondence to: Kloveniersburgwal 48, 1012 CX Amsterdam, The Netherlands. Tel.: +31 205 253 324.

E-mail addresses: poort@uva.nl (J. Poort), j.leenheer@uvt.nl (J. Leenheer), vdham@uva.nl (J. van der Ham).

perspective, however, the relevant question is not whether blocking access to TPB decreased the number of visitors to this website, but what the effect is on online copyright infringement as a whole.

Blocking access to TPB may affect unauthorised file sharing through various mechanisms. First, the mere announcement of the intervention may discourage downloading because of anticipation or awareness that file sharing is not appreciated by rights holders¹ and may deprive authors from their income (*awareness effect*). Second, after the actual blocking has become effective, illegal content may become less attractive, because it is more difficult to find. This raises transaction costs, which may have an instant negative *blocking effect* on file sharing. After these immediate effects, two opposing mid- or long-term effects could occur: On the one hand, consumers may be triggered by the intervention to use legal alternatives and continue using them (*conversion effect*), on the other hand, they may learn how to circumvent the blocking, and new illegal sources may be launched, causing file sharing to increase again (*relapse effect*).

This paper assesses the effects of blocking access to TPB on unauthorised file sharing and the use of legal channels, using two empirical approaches next to each other: (1) Consumer surveys among representative consumer samples (two measurements) and (2) BitTorrent monitoring on selections of torrents (three measurements). BitTorrent monitoring concerns an innovative data collection technique that directly measures BitTorrent participation by monitoring the distribution of peers for a sample of torrent files, without participating in the file sharing. The first consumer survey and BitTorrent Monitoring took place three months after the first ruling; the second monitoring after four months. The second survey took place six months after the second ruling and ten months after the first. Finally, the second monitoring took place approximately one year after the first ruling (see Fig. 1).

Both methods combined lead to the conclusion that while a small group of respondents downloads less from illegal sources or has stopped doing so altogether and a significant but small effect is found on the distribution of Dutch peers, there is no impact on the percentage of the Dutch population downloading from illegal sources. As such, this paper contributes to the literature on the effectiveness of online copyright enforcement. Moreover, it provides a novel and non-infringing technology for BitTorrent monitoring.

The rest of this paper is structured as follows. A short technical introduction to the BitTorrent file sharing mechanism, and an overview of the emerging literature on copyright enforcement and on BitTorrent monitoring are provided in Section 2. The design and results of the consumer surveys are presented in Section 3, and BitTorrent monitoring in Section 4. The conclusions are summarised in Section 5.

2. Background and literature

2.1. The BitTorrent file sharing mechanism

The BitTorrent protocol is a peer-to-peer protocol in which peers cooperate in distributing content over the Internet. A *peer* is a programme running on a computing node that participates in downloading and uploading content. This content is divided into *blocks* of data, which are exchanged between peers and together form the complete content. A *swarm* is a set of peers sharing a single set of files, a *torrent file*. This torrent file describes the relevant metadata of the content being distributed to support the BitTorrent protocol.

Trackers are used to bootstrap and accelerate BitTorrent swarms; they participate in swarms by keeping track of all participants and provide a peer with information on other peers in the swarm. A peer can discover other peers through *peer exchange*, that is sharing known peers with connected peers.

The initial version of the BitTorrent protocol (Cohen, 2008) used torrent files to describe content. Later versions have added another layer of distribution by storing the BitTorrent files in a *Distributed Hash Table* (DHT) storage network created by all global peers. A so-called *magnet link* can then be used to address content in this DHT network, which provides the contents of a torrent file, and several participating peers. The main advantage of magnet links is that they remove the need for a central node, the tracker, but often magnet links also contain pointers to trackers to improve the peer discovery process.

2.2. Effectiveness of measures against unauthorised file sharing

Since early this century, a substantial empirical literature emerged on the effects of unauthorised file sharing on the sales of entertainment products. Early contributions focused on the music industry, later some studies appeared on movies. In their literature review, Smith and Telang (2012) conclude that a large majority of these studies find a negative effect of unauthorised file sharing on sales.

Over the years, the entertainment industry has pursued a variety of strategies to combat unauthorised file sharing. Some concern their own supply, for instance the use of Digital Rights Management (DRM) technology to prevent users from sharing legally acquired content. For the music industry this strategy proved to be counterproductive and was abandoned (Sinha, Machado, & Sellman, 2010; Vernik, Purohit, & Desai, 2011), while for audio-visual products, e-books and games the

¹ In the Netherlands, downloading from illegal sources is allowed under the private copying exception, but uploading is illegal. By default BitTorrent clients download as well as upload content.

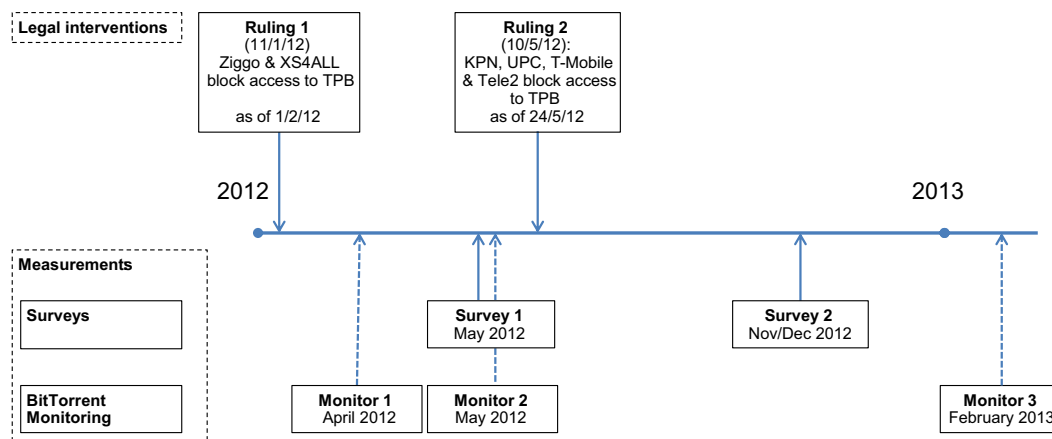


Fig. 1. Timeline of legal interventions and measurements.

use of DRM is still common. Another strategy is to offer legal digital alternatives. [Danaher, Dhanasobhon, Smith, and Telang \(2010\)](#) study the effect of the removal of NBC content from the iTunes store in December 2007 and its restoration in September 2008, on BitTorrent piracy and DVD sales on Amazon. They associate the removal with an 11.4% increase in piracy of this content, twice the legal digital sales prior to removal. After the content was restored, no significant effects on DVD sales were found, nor on piracy levels. A more controversial strategy involves the pollution or poisoning of illegal file sharing networks with useless decoys ([Christin, Weigend, & Chuang, 2005](#)).

Blocking access to TPB, the object of this study, stands in a tradition of *legal* actions against file sharing. These can be distinguished in action against individual file sharers, the demand-side of the illegal market, and actions against the supply-side, platforms that accommodate unauthorised file sharing.

2.2.1. Legal action against individual file sharers

In June 2003, the Recording Industry Association of America (RIAA) initiated a series of lawsuits against individual file sharers. [Bhattacharjee, Gopal, Lertwachara, and Marsden \(2006\)](#) tracked the online file sharing behaviour of over 2000 individuals. They found that in reaction to these lawsuits, the majority of substantial file sharers decreased the number of files shared typically by 90% and small time file sharers typically to a third. However, the individuals who continued unauthorised file sharing increased their activity again after a court ruling that made it harder for the RIAA to request the names of file sharers from ISPs. Furthermore, the authors note that individuals may have gone off the radar, using more covert file sharing technologies.

[Adermon and Liang \(2011\)](#) study the effects of the implementation of the Intellectual Property Rights Enforcement Directive (IPRED) in Sweden on music and movie sales. This European directive, implemented on 1 April 2009, substantially increased the risk of being caught and prosecuted for online file sharing. The authors have found an 18% drop in Internet traffic during the six months following the implementation. Using difference-in-difference analysis with Finland and Norway as controls, they conclude that the implementation led to an increase in the sale of physical music by 27% and digital music by 48%. No significant effects were found on cinema visits or DVD sales. On the other hand, it was also shown in the study that “the reform effects more or less disappeared after six months except for digital music sales” – the aforementioned *relapse effect*. They also report the outcome of two consumer surveys on file sharing. In 2009, 23% of the respondents stated they had stopped using file sharing sites as a result of the new legislation, 37% used file sharing sites less ($N=429$). In 2010, 52% stated they used file sharing sites less for downloading music than the year before ($N=1060$). From this group who reported to download less than the year before, 56% mentioned Spotify as the reason for this, while 34% mentioned the IPRED, and 25% “better legal services”.

[Danaher, Smith, Telang, and Chen \(in press\)](#) study the effect of the French HADOPI legislation on digital sales in the iTunes store. Under this “three strikes” legislation, implemented in October 2009, infringers caught first receive a warning. When caught again, they get a second warning, and after this, suspension of their Internet connection may be ordered. Using a difference-in-difference approach comparing French data with other countries, the authors have found a positive effect on song and album sales at iTunes of 22.5% and 25% respectively (*conversion effect*). However, it is impossible to disentangle the effects of the actual legislation and the education campaigns accompanying the introduction of HADOPI (*awareness effect*). Most of the effect seems to have arisen before the (amended) legislation was finally accepted by the Constitutional Council and diminished since then.

2.2.2. Legal action against platforms that accommodate file sharing

A different strategy is directed towards platforms that accommodate file sharing, the supply side of the illegal market. Blocking access to TPB stands in this tradition. An early victory of right holders against this supply side was the shutdown of Napster in July 2001. However, Napster was soon succeeded by alternative platforms such as KaZaA and BitTorrent clients that decentralise the file sharing process. The bootstrapping of the process occurs at sites such as TPB. An alternative

technology is provided by cyberlockers (or *one-click hosters*), cloud services where individuals can store copyright protected content anonymously for others to download.

In January 2012, Megaupload, the most popular cyberlocker, was shut down. [Danaher and Smith \(2013\)](#) study the effects of this natural experiment on unauthorised file sharing and legal online movie rentals and purchases. They analyse cross-country variation in the use of Megaupload before and the change in legal sales after the shutdown. No relation is found between the penetration of Megaupload and the digital sales *prior* to the shutdown. However, a significant positive relationship is found between this penetration and the sales change *after* the shutdown (*blocking effect*). For each additional 1% of pre-shutdown penetration, the post-shutdown sales increased by an extra 2.5–3.8%. The absence of a relation between Megaupload penetration and digital sales prior to shutdown suggests that the effect of the shutdown is temporary and lasts until consumers have found their ways to alternative suppliers of illegal video content. [Peukert, Claussen, and Kretschmer \(2013\)](#) also study the effect of the Megaupload shutdown and have found a negative effect of the shutdown on box office revenues for smaller and mid-range movies. Apparently, only large blockbusters benefit from the shutdown of Megaupload, whereas smaller movies may benefit more from file sharing through word-of-mouth in social networks.

[Lauinger et al. \(2013\)](#) also study the effect of legal actions against cyberlockers, such as removing certain content. They have found that such actions are a nuisance to the users of cyberlockers but that their effect on overall availability of content and on file sharing activity is limited. They conclude that cyberlockers “are probably most vulnerable to antipiracy measures targeted at removing external sources of revenue. Indexing sites may be less affected, especially those that are less driven by (and reliant on) monetary gain” ([Lauinger et al., 2013, p. 12](#)).

In sum, this review of the literature shows that legal actions against file sharers and platforms for unauthorised file sharing often have immediate effects (*awareness* and *blocking effects*) which disappear after typically six months, as illegal supply and demand find other places to meet (*relapse effect*). This is congruent with the conclusion drawn by [Cammaerts, Mansell, and Meng \(2013\)](#), that “[t]argeting individual internet users is not likely to reverse the trend toward an online sharing culture.” Subsequently, they stress the importance of independent evidence for copyright policy. This study on the effect of blocking access to TPB at several points in time during the first year after the intervention adds to this literature and body of evidence.

2.3. BitTorrent monitoring

There is a large body of research on monitoring BitTorrent and other peer-to-peer networks. Many studies are focused on detecting monitors and escaping detection from these monitors. [Piatek, Kohno, and Krishnamurthy \(2008\)](#) describe a reverse-engineering approach to BitTorrent monitoring by copyright holders attempting to identify infringing users. They have found that this monitoring has become more systematic, yet not conclusive. In their experiments, they are able to inject false information, which is then served with complaints about copyright infringement. Furthermore, blacklists used by the BitTorrent community at the time were not effective in identifying these monitors.

[Toro and Chothia \(2009\)](#) wrote a BitTorrent monitor for examining the behaviour of peers participating in swarms. This can then be used to classify peers heuristically, so that suspicious peers showing deviant behaviour can be identified and thus be avoided.

[Bauer, McCoy, Grunwald, and Sicker \(2009\)](#) note that passive monitors often produce false positives, and that active monitoring of a BitTorrent swarm is much more effective. They have created a tool, BitStalker, which probes participating peers, exchanges a block of data and then requests a peer exchange. This allows for monitoring a BitTorrent swarm in a way that is robust against trackers providing false data and also verifies whether peers are actively participating. [Jünemann, Andelfinger, Dinger, and Hartenstein \(2010\)](#) and [Wolchok and Halderman \(2010\)](#) monitor the Distributed Hash Table (DHT) storage network. The BitMON tool created by Jünemann et al. monitor the participants in the DHT network, and allows tracking their behaviour and the stability of the DHT network. Wolchok and Halderman instead crawl the DHT to discover the stored torrent files but also identify peers downloading these files. This is an indirect way of identifying BitTorrent participants.

[Hoßfeld et al. \(2010\)](#) use the test bed distributed by PlanetLab to monitor BitTorrent swarms, with the objective of identifying how much the performance can be improved by adjusting the BitTorrent distribution protocol by leveraging distance in the network in forming the overlay network. They show that it is possible for most swarms to identify almost all of the participating IP addresses.

[Kryczka, Cuevas, Guerrero, Azcorra, and Cuevas \(2011\)](#) classify many different BitTorrent monitoring techniques: Portal, tracker and peer crawling, but also a custom client/plugin. They identify the possibilities of these techniques and identify a custom client/plugin as the best method for gathering information about peers. [Chothia, Cova, Novakovic, and Toro \(2013\)](#) classify monitoring techniques as direct and indirect, equivalent to the active and passive techniques mentioned earlier. They observe that both techniques are used to identify infringing peers.

3. Consumer survey

To study the effects of blocking access to TPB, two surveys were held among representative samples of the Dutch population aged 16 years and older. Combining both surveys yields a multi-period measurement of reactions to blocking access expected among individuals, and of their (self-reported) reaction after three, six and ten months (see [Fig. 1](#)).

3.1. Sample and response

Both surveys were conducted in the CentERpanel, a representative online household panel. In contrast to most commercial panels, CentERpanel is not an access panel: households can only participate on invitation. The CentERpanel recruits participants with the argument that they support scientific and societal research without commercial purposes. Household selection is done through random sampling from Dutch address registers (probability sample), households without internet access receive necessary devices and support from CentERdata. As such the panel is highly representative on both observed and unobserved characteristics.² The panel aims to keep participants attached on a permanent basis (since 1990), but some panel attrition renders periodic panel recruitment necessary. In fact, recruitments occurred between both survey measurements, during which a wave of new participants were recruited and inactive members were dropped.

A total of 2009 people fully completed the first questionnaire, yielding a response rate of 64.4% (3118 members had been invited). The second questionnaire revealed a response of 2422, a response rate of 78.4%. 1692 panel members (54.3% of the first sample) participated in both surveys. The measurements are treated as two independent cross-sections in order to keep the information of one-time participants. The first sample consists of 55% men, 38% have a college degree, 40% live in a highly urbanised area. For the second sample the percentages are 53% men, 41% with a college degree, and 38% live in a highly urbanised area. As file sharing differs strongly among age groups and young age groups are somewhat under-represented, all data were weighted by age on seven different age groups.

3.2. Results

3.2.1. Market developments

The first measurement reveals that 27.8% of Dutch consumers purchased music in a physical format (CD, LP) in the preceding six months (Table 1). For 63.1% buying physical music formats was longer than six months ago, whereas 9.1% never did so. Overall, 51.7% obtained music from a legal source in the preceding half year: In a physical format (27.8%), as paid download or streaming (14.2%), and/or as free download or streaming from a legal source (33.2%). Finally, 18.3% downloaded music from an illegal source such as TPB in the preceding six months.

The second measurement shows that purchasing music in physical formats increased slightly ($\chi^2 = 11.8$; $p = 0.01$; $df = 3$) to 30.4% in the preceding six months. This increase is unlikely to be caused by the blocking of TPB, given the fact that acquiring music from illegal sources remained constant ($\chi^2 = 4.0$; $p = 0.26$; $df = 3$). Paid downloading remained stable as well ($p > 0.05$), whereas free downloading and streaming from a legal source decreased somewhat ($\chi^2 = 21.1$; $p < 0.001$; $df = 3$).

Data on downloading and streaming for music, films and series, books, and games is provided in Table 2. Downloading music from an illegal source is most common, closely followed by downloading films and series. The majority of Dutch consumers has never downloaded any of the content types from an illegal source (58.7% in the second measurement). Whereas downloading music was practically equal between both measurements ($\chi^2 = 4.0$; $p = 0.26$; $df = 3$), downloading films and series, games and books increased somewhat ($ps < 0.0$).

3.2.2. The effects of blocking access to The Pirate Bay

The self-reports of behavioural changes in reaction to blocking access to TPB by the two sets of ISPs are described in Table 3.³ During the first measurement, subscribers to UPC, KPN, Tele2 and T-Mobile were not confronted with the blocking yet and were asked about their expected reaction. More than half of the downloaders (56.1%) expected to keep their downloading rate unchanged; 28.8% expected to decrease their downloading, either by downloading less (21.7%) or by quitting completely (7.1%); 15.2% expected to download more.

After the blocking had become effective, consumers were asked about the actual impact of the blocking. The reported behaviour three months after the blocking differed significantly from the expected change reported before the blocking ($\chi^2 = 12.1$; $p = 0.007$; $df = 3$): The percentage of downloaders that did not change their downloading behaviour was higher (71.4%) than initially expected (56.1%). The percentage that stopped downloading was slightly higher than initially expected (8.0% versus 7.1%), but the share of customers that downloaded less was substantially lower (14.9% vs. 21.7%). On the other hand, fewer consumers increased downloading (5.7%) than they previously expected (15.2%).

There is no significant difference between the reported reaction three and six months after the blocking ($\chi^2 = 0.6$; $p = 0.91$; $df = 3$), nor between six and ten months after the blocking ($\chi^2 = 0.7$; $p = 0.87$; $df = 3$). Thus, an immediate effect of the blocking is found that is smaller than the expected effect prior to the blocking, and this effect does not change within this time span. It is important to realise that the majority of customers was not affected by the blocking, simply because they were no downloaders at the time of the blocking. Overall, 4–6% of all consumers have decreased their downloading.⁴

² For representativeness figures, see: <http://www.centerdata.nl/en/about-centerdata/what-we-do/data-collection/centerpanel/centerpanel-representativity-figures-may>.

³ Because the focus is on the developments in these subsamples, unweighted observations of those who were downloaders at the time of the blocking are used.

⁴ Approximately 25% of consumers downloaded from illegal sources in the preceding six months, of which 20–25% decreased downloading in reaction to the blocking.

Table 1
Purchasing, downloading and streaming music.

Last time	Physical format (CD/LP) (1)	Downloading and streaming from a legal source		Downloading and streaming from an illegal source (4)	All channels (1 to 4)	Total legal (1 to 3)
		Paid (2)	Free (3)			
May 2012 (N=2009)						
Past 6 months	27.8%	14.2%	33.2%	18.3%	53.6%	51.7%
6–12 months	12.2%	2.9%	3.3%	3.4%	9.4%	9.0%
> A year ago	50.9%	11.7%	11.6%	12.4%	30.6%	32.5%
Never	9.1%	71.2%	51.9%	65.9%	6.4%	6.7%
November–December 2012 (N=2422)						
Past 6 months	30.4%	14.8%	31.5%	18.2%	55.3%	53.3%
6–12 months	12.3%	3.5%	4.2%	3.5%	9.8%	9.7%
> A year ago	49.7%	12.6%	14.0%	13.7%	29.0%	30.9%
Never	7.7%	69.0%	50.4%	64.6%	5.9%	6.1%
Comparison between measurements:						
χ^2	11.8*	6.9	21.1*	4.0	4.7	5.7
(p-Value)	(0.01)	(0.07)	(< 0.001)	(0.26)	(0.20)	(0.13)

* Significant change between measurements ($p < 0.05$).

Table 2
Downloading and streaming from illegal sources.

	Music	Films and series	Books	Games	Total
May 2012 (N=2009)					
Past 6 months	18.3%	16.8%	5.1%	4.4%	24.0%
6–12 months	3.4%	1.5%	1.2%	1.9%	3.2%
> A year ago	12.4%	5.6%	2.3%	7.1%	12.0%
Never	65.9%	76.0%	91.4%	86.7%	60.8%
November–December 2012 (N=2422)					
Past 6 months	18.2%	17.8%	8.5%	6.4%	24.5%
6–12 months	3.5%	2.1%	1.7%	1.7%	3.4%
> A year ago	13.7%	8.2%	3.2%	8.7%	13.4%
Never	64.6%	72.0%	86.6%	83.2%	58.7%
Comparison between measurements:					
χ^2	4.0	41.6*	74.6*	34.7*	6.3
(p-Value)	(0.26)	(< 0.001)	(< 0.001)	(< 0.001)	(0.10)

* Significant change between measurements ($p < 0.05$).

Table 3
Reaction or expected reaction to blocking access to The Pirate Bay of customers downloading from illegal sources at the time of blocking (two measurements, split sample).

	UPC, KPN, Tele2 and T-Mobile (expected reaction, $t=0$)*	Ziggo and XS4ALL (reaction $t=3$)**	UPC, KPN, Tele2 and T-Mobile (reaction $t=6$)*	Ziggo and XS4ALL (reaction $t=10$)**
Stop	7.1%	8.0%	9.2%	8.4%
Less	21.7%	14.9%	14.5%	15.3%
Just as much	56.1%	71.4%	70.2%	71.8%
More	15.2%	5.7%	6.1%	4.6%
N	198	262	228	131
Comparison with previous measurement:				
χ^2		12.1	0.6	0.7
(p-Value)		(0.05)	(0.91)	(0.87)

* Part of 1st measurement.
** Part of 2nd measurement.

Downloading from illegal sources has not decreased since the interventions. This is confirmed in Table 4. In fact, both for UPC, KPN, Tele2 and T-Mobile ($\chi^2=43.6$; $p < 0.001$; $df=6$) and for Ziggo and XS4all ($\chi^2=942.8$; $p < 0.001$; $df=6$) the percentage of consumers downloading in the preceding six months increased. For the former it increased from 15.7% just before the blocking to 18.4% six months after. For the latter it increased from 22.5% three months to 25.2% ten months after

Table 4

Downloading and streaming from illegal sources per blocking situation (two measurements, split sample).

	UPC, KPN, Tele2 and T-Mobile		Ziggo and XS4ALL	
	No blocking, $t=0$	Blocking, $t=6$	Blocking, $t=3$	Blocking, $t=10$
Past 6 months	15.7%	18.4%	22.5%	25.2%
< week	6.0%	8.1%	7.8%	11.3%
Week-month	3.7%	4.2%	6.8%	4.1%
1–3 months	3.7%	3.4%	6.0%	5.4%
3–6 months	2.3%	2.7%	1.9%	4.4%
6–12 months	2.8%	3.3%	3.0%	3.5%
> A year ago	9.5%	12.0%	13.2%	15.1%
Never	72.1%	66.4%	61.3%	56.3%
χ^2	43.6		942.8	
(p -value)	< 0.001		< 0.001	

the intervention. For both sets of ISPs, the percentage of customers downloading very recently (preceding week or month) also increased. Thus, though a small share of downloaders reports a decrease in their downloading activities after the blocking, this effect is not reflected in the overall numbers. A likely explanation is that there are also new consumers who have started downloading from illegal sources, since the percentage of consumers that has never downloaded decreased over the measurements.

4. BitTorrent monitoring

4.1. Monitoring tools used and torrent samples

The initial monitoring (April and May 2012) (Van der Ham et al., 2012) started as an ad hoc way to chart effects of the first blocking, by Ziggo and XS4ALL, in order to come up with measurements shortly after this intervention (see Fig. 1). A programmable interface of the popular Transmission client was used. By means of a script, a torrent magnet link was added programmatically, and then every minute the list of peers the client was interacting with was requested and stored. The default limit of peers to interact with was raised to 1024 (the maximum allowed) to record as many peers as possible. This methodology recorded activity on several magnet links at different times during a few days.

The above method of recording peers yielded a list of IP addresses for each of the torrents. To convert this list to usable information, first the Team Cymru IP-to-ASN mapping service (<http://www.team-cymru.org/Services/ip-to-asn.html>) was used to record which ISP the IP address came from. This service has combined all the IP address registrations from the Internet Registries. Unfortunately, the country data from these registries is not always accurate. Many ISPs also have IP subnets registered as EU. To pin down the location further, the MaxMind Free GeoIP database was used (<http://www.maxmind.com/en/country>). In cases of conflicting results, the latter was preferred.

In April 2012, 59 Dutch spoken or subtitled torrents were selected, which yielded a total of 12,942 Dutch peers. In May 2012, 19 Dutch spoken or subtitled torrents were selected which yielded 2566 Dutch peers.

The monitor described above is an active client, from which data are exported. The methodology provides a valid insight into the BitTorrent activity but could be improved on its effectiveness to record peer activity. Therefore, a new monitor was designed from scratch for the third measurement using Python and the libtorrent library. This library (<http://www.rasterbar.com/products/libtorrent/>) implements the BitTorrent protocol and is used in many popular BitTorrent clients. The new monitor uses the library to appear as an active client but is configured such that it does not download or upload any content. The monitor joins the torrent swarm and records activity, it requests a new set of peers from the tracker as often as allowed and records all these IP addresses. The monitor does not exhibit any suspicious behaviour as defined by Toro and Chothia (2009), because it only monitors a maximum of ten torrent swarms simultaneously and behaves like a regular client.

The above monitor is a stand-alone process, which submits all its recorded peers to a database server, where they are stored and processed. Each peer record contains the IP address, the torrent it was recorded in, and the time it was recorded. During February 2013, the server and three monitors ran at different locations on the Internet and recorded activity in ten torrent swarms over a period of two weeks. After analysing the records with the Team Cymru IP-to-ASN mapping and MaxMind GeoIP database, 98,807 Dutch peers were obtained from ten Dutch spoken or subtitled torrents. As such, the new monitor proved to be much more effective than the former one.

4.2. Results

The Dutch peers were attributed to a total of 133 ISPs (108 of which recorded less than 50 peers over all measurements combined). While the court rulings primarily affect only six ISPs, these account for over 90% of Dutch residential broadband subscribers and of all Dutch peers recorded. Table 5 presents the percentage of Dutch peers for each ISP of interest for the three consecutive measurements. Although differences in the distributions of peers are statistically significant, changes are

Table 5

Distribution of peers amongst ISPs affected by the 1st or 2nd ruling.

ISP	April 2012	May 2012	February 2013
Total Dutch peers	12,942	2566	98,807
<i>Ruling 1: blocking since 1-2-2012</i>	29.7%	33.2%	31.7%
Ziggo	27.6%	31.5%	29.3%
XS4ALL	2.1%	1.7%	2.4%
<i>Ruling 2: blocking since 24-5-2012</i>	63.9%	61.1%	60.3%
KPN	29.2%	25.9%	30.3%
UPC	21.8%	23.9%	19.1%
Tele2	7.2%	6.3%	5.9%
T-Mobile	5.7%	5.0%	5.0%
<i>Other Dutch ISPs</i>	6.4%	5.7%	8.1%
Comparison with previous measurement:			
χ^2 ($df=2$)		15.7	1075.8
(p -Value)		(< 0.001)	(< 0.001)
Total peers ISPs of interest	12,118	2420	90,852
% Of all Dutch peers	93.6%	94.3%	91.9%
Total residential broadband subscriptions (million)	6.47	6.48	6.61
<i>Market shares:</i>			
Ziggo	26.3%	26.4%	26.6%
XS4ALL	3.7%	3.7%	3.5%
KPN	37.0%	37.0%	37.4%
UPC	15.3%	15.3%	15.6%
Tele2	6.6%	6.6%	5.9%
T-Mobile	4.4%	4.3%	4.1%
File sharing propensity ISPs Ruling 1	0.99	1.11	1.05
File sharing propensity ISPs Ruling 2	1.01	0.97	0.96

Note: Figures on market shares for residential broadband based on linear interpolation of Albrecht (2013). For consistency, all figures for KPN include 1.8% market share estimate of fibre subsidiaries taken over in Q4 2012.

small, which implies limited effects of the intervention on BitTorrent file sharing. The percentage of peers associated with the ISPs affected by the *first ruling* increased by 3.5%-point between April and May 2012 while the percentage affected with the *second ruling* decreased by 1.8%-point. The changes were statistically significant ($\chi^2 = 15.7$; $p < 0.001$; $df = 2$). Since the second ruling was not yet enacted by that time, this increase cannot be the result of the aforementioned *blocking effect*. It is more likely the result of an *awareness effect* on subscribers of the ISPs affected by the second ruling or of a *relapse effect* for subscribers of Ziggo.

Between May 2012 and February 2013, the percentage of peers affected by the first ruling decreased, albeit not back to the April 2012 level, while the percentage of peers affected by the second ruling decreased only slightly ($\chi^2 = 1075.8$; $p < 0.001$; $df = 2$). The latter indicates that the actual *blocking effect* adds little to the earlier *awareness effect*.

The status of the remaining smaller ISPs after the second ruling complicates this analysis. Some have also blocked access to TPB at some point even though the rulings do not explicitly apply to them: they may want to avoid prosecution. Moreover, the awareness effect may also occur for the subscribers of these ISPs. Hence, this group cannot be seen as a completely *untreated* control group. This can be resolved by only comparing the ISPs addressed in either ruling and gives similar results.

Finally, changes in ISPs' market share for residential broadband between the three measurements could influence the outcomes. To control for such market dynamics, the second half of Table 5 give the relevant residential broadband market shares as well as total market size. Combining this with the distribution of Dutch peers, a metric can be developed for the file sharing propensity of subscribers per ISP as: share of total peers/share in broadband market. For the entire market, this propensity is 1 by definition, but it differs between ISP according to the socio-demographic composition of their client base. Changes over the short time frame studied are most likely related to the court rulings. For the ISPs in ruling 1, the pattern for this metric is identical to that for the percentage of Dutch peers. For the ISP in ruling 2, the decrease between the last two measurements is negligible, despite the blocking that was ordered in between. This suggests that the *awareness effect* dominates any *blocking effect*.

5. Conclusions and recommendations

Following rulings from a Dutch court, the major Dutch Internet Service Providers have blocked access to The Pirate Bay (TPB) since February/May 2012, and more than 85% of Dutch Internet subscribers can no longer (directly) access this popular website facilitating the unauthorised exchange of copyright protected material. This study contributes to the emerging literature on the effectiveness of online copyright enforcement by measuring the effects of these interventions on

downloading from illegal sources, possibly in favour of legal channels. Two complementary empirical methods (yielding five measurements) for assessing the effects of this intervention on downloading from illegal sources are presented in the paper. No strong indications are found for any structural effects.

Two consecutive consumer surveys provide insight into consumers' reactions to the intervention after three, six and ten months, as well as the reaction they expect shortly before blocking. The intervention can only affect consumers who download or intend to download from illegal sources, 27–28% over the past year. For this segment of the population, it is found that a large majority (70–72%) is non-responsive to blocking access to TPB. This is significantly more than consumers expect prior to the blocking. About half of those who report a response to the intervention state they download less, while a third state they stopped downloading altogether. The rest claim to download more as a result of the intervention.

This would suggest a small negative *blocking effect* of the intervention on the percentage of the population downloading from illegal sources. However, no such effect is found. Instead, the percentage downloading films & series, games and books from illegal sources in the preceding six months increased between May and November/December 2012, while downloading music from illegal sources remained constant. This implies that any behavioural change in response to blocking access to TPB has had no lasting net impact on the overall number of downloaders from illegal sources, as new consumers have started downloading from illegal sources and people learn to circumvent the blocking while new illegal sources may be launched, causing file sharing to increase again (*relapse effect*).

These findings are corroborated by the second, complementary method presented: BitTorrent monitoring. BitTorrent monitoring measures observed rather than reported behaviour, but with the short-coming that it cannot observe consumers circumventing the blocking by downloading via VPN connections or from newsgroups and cyberlockers. BitTorrent monitoring reveals only small changes in the distribution of Dutch peers over the different ISPs for the three measurements, which implies very limited effects of the intervention on BitTorrent file sharing.

For the small changes observed, it is not fully possible to disentangle the different and opposing effects of the *blocking* itself, *awareness* of the intervention, *conversion* to legal alternatives induced by the blocking, and a *relapse* as a result of circumvention or the launch of new file sharing platforms. The fact that consumers report a significantly smaller response to the intervention than they expect in advance, indicates that the *awareness effect* wears off quickly. Furthermore, the increase between the two surveys of the percentage of ISP subscribers admitting to having downloaded from illegal sources in the preceding period indicate a *relapse effect*: After a small initial *awareness* or *blocking effect*, the market moves back towards the earlier equilibrium, with no or only very small structural effects.

These results are in line with a tendency found in the literature that any effects of legal action against file sharing often fade out after a period of typically six months, as the initial *awareness effect* wears off and illegal supply and demand find other places to meet. Probably, the required ICT knowledge to circumvent the blocking is no more advanced than the knowledge required to download from illegal sources. Hence, targeting individual file sharers and blocking access to file sharing platforms seem relatively ineffective to reduce unauthorised file sharing, while such measures bear a risk of alienating customers from the content industries and giving them incentives to adopt covert technologies such as dark nets, IP-spoofing and VPN. These interventions also threaten the transparency of the Internet, effectively introducing censorship. In France, such considerations as well as the costs of enforcement have led the government to temper the HADOPI-sanctions (see [Cammaerts et al., 2013](#) for a brief discussion).

Since there are no indications for structural effects of the interventions, it is unlikely that the increased use of legal channels and decrease in file sharing observed for music over a longer time span of four years ([Poort & Leenheer, 2012](#)) can be attributed to the intervention (*conversion effect*). If that were the case, a similar pattern would be expected for films and series, and books, while the opposite is found. A more likely explanation is the development of successful and comprehensive legal business models for downloading and streaming music. Therefore, policy makers and the content industry had best focus on removing any legal or practical obstacles for comprehensive and attractive legal online models, not only for music but also for films, series and books, instead of combating unauthorised file sharing. Researchers could support this by studying the dynamics between the adequacy of legal supply and file sharing.

Acknowledgements

The survey presented was part of a larger research project conducted at the initiative and under the authority of IViR and CentERdata, with partial financial support from The Netherlands Ministry of Education, Culture and Science, Ziggo, KPN, XS4ALL, DELTA, CAIW, the Royal Dutch Book Trade Association (KVB) and SOS Internet. Funding sources had no involvement in the analysis.

References

- Adermon, A., & Liang, C. (2011). Piracy, music, and movies: A natural experiment. *IFN working paper no. 854*. Retrieved August 27, 2013 from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1752224.
- Albrecht, K. (2013). *Dutch Broadband Q3 2013. The Netherlands: Telecompaper*, Houten.
- Bauer, K., McCoy, D., Grunwald, D., & Sicker, D. (2009). Bitstalker: Accurately and efficiently monitoring BitTorrent traffic. In *Proceedings of the First IEEE international workshop on information forensics and security*. Retrieved from <http://www.mit.edu/~ke23793/papers/bauer-wifs09.pdf>.

- Bhattacharjee, S., Gopal, R. D., Lertwachara, K., & Marsden, J. R. (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.
- Cammaerts, B., Mansell, R., & Meng, B. (2013). Copyright & creation: A case for promoting inclusive online sharing. *Media Policy Brief*, 9. The London School of Economics and Political Science, Department of Media and Communications, London, UK. Retrieved December 7, 2013 from <http://www.lse.ac.uk/media@lse/documents/MPP/LSE-MPP-Policy-Brief-9-Copyright-and-Creation.pdf>.
- Chothia, T., Cova, M., Novakovic, C., & Toro, C. (2013). The unbearable lightness of monitoring: Direct monitoring in BitTorrent. *Security and privacy in communication networks. Lecture notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, 106, 185–202 (Springer Berlin/Heidelberg, Germany).
- Christin, N., Weigend, A. S., & Chuang, J. (2005). Content availability, pollution and poisoning in file sharing peer-to-peer networks. In *Proceedings of the 6th ACM Conference on Electronic Commerce* (pp. 68–69–77), New York. Retrieved from <http://p2pecon.berkeley.edu/pub/CWC-EC05.pdf>.
- Cohen, B. (2008). *The BitTorrent Protocol Specification*. Retrieved December 6, 2013 from http://www.bittorrent.org/beps/bep_0003.html.
- Court of The Hague. (2012a). LJN: BV0549; 374634/HA ZA 10-3184. Retrieved from <http://deeplink.rechtspraak.nl/uitspraak?id=ECLI:NL:RBSGR:2012:BV0549>.
- Court of The Hague. (2012b). LJN: BW5387; 413085/KG ZA 12-156. Retrieved from <http://deeplink.rechtspraak.nl/uitspraak?id=ECLI:NL:RBSGR:2012:BW5387>.
- Danaher, B., Dhanasobhon, S., Smith, M. D., & Telang, R. (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., & Smith, M. D. (2013, September 14). Gone in 60 s: The impact of the megaupload shutdown on movie sales. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2229349.
- Danaher, B., Smith, M., Telang, R., & Chen, S. The effect of graduated response anti-piracy laws on music sales: Evidence from an event study in France. *Journal of Industrial Economics* (in press), <http://www.jindc.org/article/effect-graduated-response-anti-piracy-laws-music-sales-evidence-event-study-france>.
- Ham, J. van der, Rood, H., Dumitru, C., Koning, R., Sijm, N. & De Laat, C. (2012). *Review en Herhaling BREIN steekproeven 7-9 april 2012*, Technical report UVA-SNE-2012-01, Retrieved from <http://staff.science.uva.nl/~vdham/research/publications/dutchpirate.pdf>.
- Hoßfeld, T., Hock, D., Oechsner, S., Lehrieder, F., Despotovic, Z., & Kellerer, W. (2010). Measurement of BitTorrent Swarms and their AS Topologies. *University of Würzburg institute of computer science research report series*. Retrieved from <http://www3.informatik.uni-wuerzburg.de/TR/tr464.pdf>.
- Jünemann, K., Andelfinger, P., Dinger, J., & Hartenstein, H. (2010). Bitmon: A tool for automated monitoring of the BitTorrent DHT. In *Proceedings of the 2010 IEEE Tenth International Conference on Peer-to-Peer Computing (P2P)* (pp. 1–2). Retrieved from ieeexplore.ieee.org/xpls/abs_all.jsp?arnumber=5569978.
- Kryczka, M., Cuevas, R., Guerrero, C., Azcorra, A., & Cuevas, A. (2011). Measuring the bittorrent ecosystem: Techniques, tips, and tricks. *IEEE Communications Magazine*, 49(9), 144–152.
- Lauinger, T., Szydłowski, M., Onarlioglu, K., Wondracek, G., Kirda, E., & Kruegel, C. (2013). Clickonomics: Determining the effect of anti-piracy measures for one-click hosting. In *Proceedings of NDSS Symposium 2013*, San Diego, CA. Retrieved from http://www.internetsociety.org/sites/default/files/07_1_0.pdf.
- Peukert, C., Claussen, J., & Kretschmer, T. (2013, August 20). *Piracy and movie revenues: Evidence from Megaupload*. A tale of a long tail? Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2176246.
- Piatek, M., Kohno, T., & Krishnamurthy, A. (2008). *Challenges and directions for monitoring P2P file sharing networks, or, why my printer received a DMCA takedown notice*. Retrieved from http://www.usenix.org/event/hotsec08/tech/full_papers/piatek/piatek.html.
- Poort, J. & Leenheer, J. (2012). File sharing 2©12. Downloading from illegal sources in The Netherlands. Institute for Information Law & CentERdata, Amsterdam/Tilburg. Retrieved from http://www.ivir.nl/publications/poort/Filesharing_2012.pdf.
- Sinha, R. K., Machado, F. S., & Sellman, C. (2010). Don't think twice, it's all right: Music piracy and pricing in a DRM-free environment. *Journal of Marketing*, 74(2), 40–54.
- Smith, M. D., & Telang, R. (2012, August 19). *Assessing the academic literature regarding the impact of media piracy on sales*. Retrieved from http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2132153.
- Toro, C. A. G., & Chothia, T. (2009). *Analysis of BitTorrent Peers' Behavior and Monitoring Trends (Category: Law and security)*. Retrieved from http://www.kaspersky.com/images/camilo_andrzej_gonzalez_toro-10-75858.pdf.
- Vernik, D. A., Purohit, D., & Desai, P. S. (2011). Music downloads and the flip side of digital rights management. *Marketing Science*, 30(6), 1011–1027.
- Wolchok, S., & Halderman, J. (2010). Crawling BitTorrent DHTs for fun and profit. In *Proceedings of WOOT (Washington, DC, USA, 2010)*. Retrieved from http://www.usenix.org/event/woot10/tech/full_papers/Wolchok.pdf.
- <http://www.centerdata.nl/en/about-centerdata/what-we-do/data-collection/centerpanel/centerpanel-representativity-figures-may> Accessed 12.12.13.
- <http://www.maxmind.com/en/country> Accessed 06.04.13.
- <http://www.rasterbar.com/products/libtorrent/> Accessed 06.04.13.
- <http://www.team-cymru.org/Services/ip-to-asn.html> Accessed 07.04.13.