

Batikas, Michail; Claussen, Jörg; Peukert, Christian

Conference Paper

Follow The Money: Piracy and Online Advertising

28th European Regional Conference of the International Telecommunications Society (ITS):
"Competition and Regulation in the Information Age", Passau, Germany, 30th July - 2nd
August, 2017

Provided in Cooperation with:

International Telecommunications Society (ITS)

Suggested Citation: Batikas, Michail; Claussen, Jörg; Peukert, Christian (2017) : Follow The Money: Piracy and Online Advertising, 28th European Regional Conference of the International Telecommunications Society (ITS): "Competition and Regulation in the Information Age", Passau, Germany, 30th July - 2nd August, 2017, International Telecommunications Society (ITS), Calgary

This Version is available at:

<https://hdl.handle.net/10419/169448>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Follow The Money: Piracy and Online Advertising*

Michail Batikas¹, Jörg Claussen¹, and
Christian Peukert²

¹LMU Munich

²University of Zurich

July 1, 2017

Abstract

Online copyright enforcement, in the form of either direct action against the supply-side (via website shutdowns) or the demand-side (via individual lawsuits against users), has not been very effective in reducing piracy. Regulators have therefore put forward the so called “follow the money” approach. Because the main source of revenue for infringing websites often comes from online advertising, the idea is that cutting access to advertisers could lower the financial incentives for website owners.

In this paper, we aim to provide systematic evidence on the effectiveness of such a policy. We collect data on the advertising services associated with a large number of piracy websites and corresponding set of legitimate “placebo” websites before and after the first steps of a self-regulatory effort in the European Union went in place. Preliminary results suggest that advertising services indeed reduce their activities on piracy websites, however only those that are most likely to be directly affected by the regulation. We further provide evidence that some advertising services that did not cater to the piracy market before, start to do so – perhaps as a strategic response.

Keywords: Piracy, Copyright Enforcement, Online Advertising

JEL No.: L82, M37, D83

*Batikas (corresponding author): m.batikas@lmu.de, Claussen: j.claussen@lmu.de, Peukert: christian.peukert@uzh.ch.

1 Introduction

Advertising is an important institution in multi-sided markets. On the one hand, it enables firms to inform and persuade consumers about their products, on the other, it generates incentives for the creation of content. The classic example is the newspaper market, where advertisers fund journalism by effectively subsidizing the price charged to readers. It is well understood that in this market, as well as in many other media markets, the interests of advertisers and consumers are not necessarily aligned, which potentially affects the type of content created by publishers.

However, there can be important externalities that make advertising-funded content overall welfare decreasing, although it creates a positive surplus for all directly involved parties (advertisers, consumers and publishers). In the context of websites that offer access to large volumes of unlicensed entertainment content (online piracy), negative externalities shape up from potentially decreased long run incentives to create entertainment content. Recently, the European Commission has put forward a self-regulatory approach. The idea is that the online advertising industry will eventually promise to ban certain publishers from their platforms. A first significant step towards this happened in October 2016, when industry representatives agreed on “the basic principles of an agreement”.

In this paper, we aim to empirically evaluate the effects of such efforts. We collect data on the advertising services associated with a large number of piracy websites and corresponding set of legitimate “placebo” websites. We then compare the underlying market structure of the advertising industry and its dynamics. Exploiting exogenous variation from the timing of the self-regulation helps to establish causal parameter estimates.

We are still in the process of data collection and can therefore not fully test all of our hypotheses. In this draft we only show some preliminary non-causal results. Those results suggest that advertising services indeed reduce their activities on piracy websites, however only those that are most likely to be affected the most. Perhaps as a strategic reaction, services that have not appeared on piracy websites before do so after the regulation comes into effect.

If confirmed in a more careful empirical set up, this seems to suggest that there is further scope for public policy to improve the effectiveness of “follow the money” enforcement.

2 Background

2.1 Online copyright enforcement

2.1.1 Evidence on effects and effectiveness of previous enforcement efforts

In the last decade, firms and governmental authorities have used a variety of different approaches to enforce copyright online. The aim has been to counter the potentially large negative revenue effects of online piracy, which may reduce long-run incentives to invest in the creation of content. A growing body of literature evaluates these policies and arrives at sobering conclusions.

A first set of actions was targeted at the demand-side, involving stricter anti-piracy legislation, and blocking access to certain websites. According to [McKenzie \(2017\)](#), the announcement and introduction of stricter laws in France, New Zealand, South Korea, Taiwan, UK, and the US did not affect movie revenues at the box office. For digital music sales, however, [Danaher et al. \(2014\)](#) report an increase of 22-25% as a result of a new law in France. [Adermon and Liang \(2014\)](#) find a similar effect on music sales caused by stricter legislation in Sweden, but they also show that the effect diminishes after 6 months. Evaluating the effects of DNS-blocking in the Netherlands, [Poort et al. \(2014\)](#) find a small decrease in self-reported piracy-levels of music, video, games and books. But also in their data, the effect vanishes after 6 months for most content types.

A second line of research investigates the effects of supply-side measures, such as the takedown of specific content, entire hosting services and link directories. [Reimers \(2016\)](#) finds that private efforts to remove infringing content increases sales of electronic books by 14%. [Danaher and Smith \(2014\)](#) find that the number of downloads of licensed movie content increases by 7-9% after the shutdown of Megaupload. In the same context, [Peukert et al. \(2017\)](#) find a similarly sized effect on the theatrical revenues of widely released blockbuster movies, however, revenues of the average movie even decline. The authors argue that piracy can have positive externalities through word-of-mouth, which are more prevalent for movies with smaller advertising budgets. Evaluating the shutdown of a link directory, [Aguiar et al. \(2017\)](#) use individual-level clickstream data to show that a reduction in consumption of unlicensed content is relatively small, mainly because consumers

quickly switch to other piracy websites. They further don't find much evidence for an increase in the demand for licensed services.

As a result, it is not surprising that online piracy is still prevalent despite the plethora of enforcement efforts: According to a report published by the UK Intellectual Property Office, a consumer survey from the first half of 2016 suggests that of all consumers who consumed online entertainment content in past three months, 25% have done so at least once via unlicensed channels.¹

2.1.2 “Follow the money” enforcement

A relatively new idea is to take the so-called “follow the money approach”, i.e. indirect enforcement targeted towards the financial incentives of piracy websites. According to a joint report of the music and advertising industry in 2012, more than two thirds of infringing websites are predominately financed by advertising.² This has led the European Commission to put forward a self-regulatory initiative of the online advertising industry to minimize the funding of copyright infringing websites:

“On 14 March 2016, the Commission held a stakeholders’ general meeting on online advertising and IPR, bringing all the interested parties together (the advertising industry, intermediaries, content protection sector, online media, right owners, civil society, consumer organisations, brands and advertisers). The stakeholders discussed the possibility of establishing a voluntary agreement at EU level in order to avoid the misplacement of advertising on IP-infringing websites, thereby restricting the flow of revenue to such sites while safeguarding the reputation of the advertisers and the integrity of the advertising industry. At the next general meeting on 21 October, stakeholders agreed on the basic principles of an agreement to help to restrict the income that IP-infringing websites get from online advertising.”³

It seems that the effectiveness of a “follow the money” policy, although intuitively appeal-

¹Intellectual Property Office (2016), “Online Copyright Infringement Tracker, Latest wave of research Mar 16–May 16, Overview and key findings”, p. 19.

²Google and PRS (2012), “The six business models for copyright infringement”, https://docs.google.com/file/d/0Bw8Krj_Q8UaENDhEOG1LVFRhVku/view.

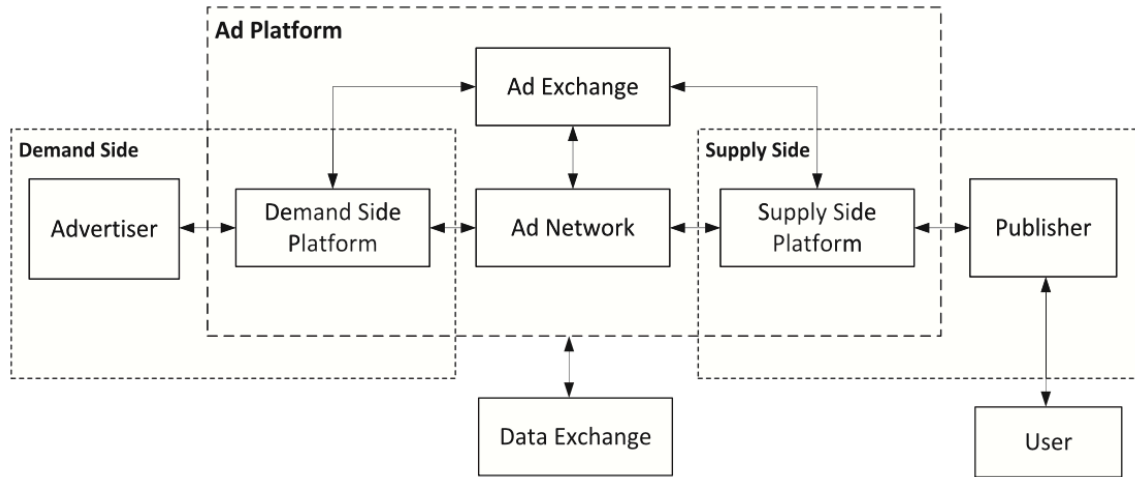
³https://ec.europa.eu/growth/industry/intellectual-property/enforcement_en

ing, hinges critically on the appropriate incentives of the involved parties in the online advertising market. The trade-offs faced by individual entities in the market are functions of the complex relationships between advertisers, publishers, advertising networks and auxiliary advertising services. It is therefore crucial to develop an empirical understanding of the advertising market as a whole when trying to evaluate recent self-regulatory initiatives and eventually provide guidance for the further development of market mechanisms and policy instruments.

2.2 The economics of online advertising

2.2.1 The market for online advertising and web tracking

Figure 1: The online advertising ecosystem



Source: Estrada-Jiménez et al. (2017), figure 3.

Figure 1 shows the complex organization of the online advertising industry. Publishers connect advertisers and consumers/users, e.g. via ads shown in the context of an unlicensed streaming website.

A complex system of entities operates the advertising platform that connects publishers and advertisers: *Advertising networks* take a central position in this system in that they connect a range of auxiliary advertising services that intermediate and facilitate a market between advertisers, publishers and data providers. *Advertising exchanges* and *data exchanges* are auction-based marketplaces where advertisers or advertising networks bid to

place ads in the space of publishers or for access to data about consumers, which they can match to their information about a particular user (group). Advertisers and publishers sometimes buy services from third-parties (*demand-side* and *supply-side platforms*) that enable more effective targeting, most prominently from *tracking services* which collect and consolidate behavioral data of users. For example, a tracking-plugin installed on a website places a cookie on a consumer’s machine, which allows to follow the same consumer across other websites that operate the same plugin. Prominent examples of tracking services are the “like” and “share” buttons operated by social media websites ([Chaabane et al., 2012](#); [Roosendaal, 2012](#)). These data are sold to advertisers or advertising networks, which use it to predict the “value” of a consumer, which affects their willingness to pay for the consumer’s eyeballs.

It is important to note that firms are often vertically integrated and assume multiple roles at the same time. For example, an advertising network can collect data from the tracking of consumers, use this data to enhance its prediction models, but also sell it to advertisers or other advertising networks.

Much of what is going on in real-time as a user visits a publisher’s website is not directly visible for the consumer, nor are the identities of the involved parties that collect and trade personal information completely transparent. As a result, a growing body of literature in economics and related fields has focused on privacy issues (see [Acquisti et al., 2016](#) for a summary).

An empirical literature in computer science has established some stylized facts about the industrial organization of the online advertising industry (e.g. [Gill et al., 2013](#); [Libert, 2015](#); [Schelter and Kunegis, 2016](#); [Lerner et al., 2016](#)): Over the last 20 years, the industry has grown by a factor of 50 in terms of number of distinct services on the market, and by a factor of 4 in terms of the median number of distinct services used per website. Further, the industry has become considerably concentrated with the top 20 services covering about a third of the market. As of May 2014, about 88% of the top 1 million most popular websites operate some kind of advertising technology, contacting an average of 9.47 distinct third-party domains when visited ([Libert, 2015](#)). Beyond that, we don’t have much evidence on the competitive forces in the online advertising industry.

2.2.2 Previous theoretical literature

The starting point for an evaluation of the effects of “follow the money” enforcement is to develop an understanding of the industrial organization of online advertising and markets for data (see [Goldfarb, 2014](#) and [Peitz and Reisinger, 2015](#) for summaries).

[de Cornière and de Nijs \(2016\)](#) look at the problem of a monopoly publisher that can decide on whether to install a tracking service that allows advertisers to access information about consumers. In their model, the tracking service enables advertisers to charge higher prices on the product market. However, only with sufficient competition across advertisers, the publisher can fully capture this rent, which drives incentives to install the tracking services.

[Athey et al. \(forthcoming\)](#) have a model with two competing publishers that sell ad space to heterogeneous advertisers. Both publishers can perfectly track consumers on their own platform, but not beyond. With multi-homing consumers, advertising is inefficient in that multi-homing advertisers may not reach some consumers and may advertise too often to others. As a result, low-type (low valuation per consumer) advertisers single home, high-type advertisers multi home (because of opportunity cost of not informing consumers).

The set-up of [Ghosh et al. \(2015\)](#) is similar, but publishers can decide to share their information about consumers with other publishers. In this model, when advertisers value consumers homogenously, all publishers agree, i.e. either all will decide to share information, or none. However, with heterogeneous advertisers, in the sense that some advertisers value high-type consumers more than other advertisers, sharing information makes some publishers better off, some worse off. This is because advertisers can better distinguish between high-type and low-type consumers visiting the publisher, which increases their average willingness to pay, reducing demand for ads at competitive publishers. Put differently, because of the leakage effect of information sharing, the advertiser will be able to target the same consumer elsewhere at a cheaper price.

[Xiang and Sarvary \(2013\)](#) look at the effects of upstream and downstream competition in model with two information sellers and two information buyers (e.g. advertisers). The sellers compete on accuracy of information, and their products can be complimentary to buyers, leading to a high competitive price. However, with downstream competition, a buyer gains more through information when the competitor is less informed. This

makes the information of the two sellers strategic substitutes, leading to harsh competition between the sellers.

The existing theoretical work has been very helpful to develop an understanding of how (not) sharing of information with advertisers affects publishers' competition for advertisers, and how competition in the market for data affects competition across advertisers. However, when compared to the complex empirical structure of the online advertising industry, it becomes evident that these models can only be used to look at distinct segments of the market. An empirical evaluation of the effects of "follow the money" enforcement will help to understand how competition in the market for advertising services affects the competition of publishers for competing advertisers. Beyond the specific application of online piracy, this may have significant insights for the further development of theory.

3 Hypotheses

We can combine the empirical insights about the online advertising ecosystem and its market structure with the lessons from theoretical work to arrive at testable hypotheses regarding the effects of (self-) regulation in the industry.

3.1 Direct effects

Exiting the market of rogue publishers has direct negative effects on revenues for many firms operating within the advertising ecosystem. The decision to exit will depend on the relative importance of this market for the firm's overall profits. Revenue effects seem important when compared against the expected sanctions for non-compliance, which are likely to be only weakly enforceable in an industry where revenues of different firms are correlated. Therefore some firms would ex ante not take part in a self-regulatory effort that would put their existence at risk. Even if some firms find it beneficial to exit the market, all else equal, some other firms will find it profitable to enter.

Hypothesis 1: Advertising services for which the self-regulation is more binding will decrease their presence on unlicensed content websites to a higher degree.

Hypothesis 2: Advertising services that have previously not operated on unlicensed content websites will start to do so.

3.2 Indirect effects

Depending on the business model, there can be important indirect effects. Firms that specialize or are vertically integrated in consumer tracking will have low incentives to stop serving rogue publishers, because the information that a user visited a (rogue) website can be valuable to advertisers seeking to advertise elsewhere. For example, knowing that a user has navigated to an unlicensed live sports streaming website can be indicative of that user’s interest in soccer. In other cases, observing visits to rogue websites can be of value for negative targeting, i.e. to make sure that this user is not served the ad. Therefore, if an integrated ad network drops out of the market for rogue websites, it indirectly also reduces the value of its service to legitimate sites. Again, this effect should be stronger for firms with a higher market share in rogue websites.

Hypothesis 3: The number of tracking services operating on unlicensed content websites will decrease to a lesser extent than the number of pure advertising services operating on unlicensed content websites.

4 Identification strategy and data

4.1 Identification strategy

Reduced form causal evidence on the effects of the recent implementation of the “follow the money” enforcement approach can be generated in difference-in-differences analysis. The econometric model is based on a before and after comparison, i.e.

$$NumberOfWebsites_{it}^k = \alpha + \delta (After_t \times Treat_i) + \nu_t + \mu_i + \varepsilon_{it}, \quad (1)$$

where $NumberOfWebsites_{it}^k$ is the number of websites of type k (piracy, placebo) served by third-party service i at time t . $After_t$ indicates the time period after the first step of

the self-regulation came into effect. The preferred specification further includes time fixed effects ν_t and third-party service fixed effects μ_i . Therefore we cannot separately identify coefficients for $Treat_i$ and $After_t$ as they are absorbed by third-party service and time fixed effects, respectively. The error term ε_{it} has the standard assumptions, we report White-robust estimates.

We propose three different ways of defining the treatment group. First, we define $Treat_i = EU_i$, i.e. an indicator whether the third-party service operates within the European Union. The idea is that this should be the group of services that is affected (the most) by the self-regulation. Second, we define $Treat_i = Served_i$, i.e. an indicator whether the third-party service has served piracy websites before the self-regulation came into force. Third, we combine both to define $Treat_i = Served_i \times EU_i$. This is our preferred definition as it should pin down the (most) affected group of third-party services most accurately.

To be able to observe variation in $Served_i$, we also collect data about the third-party services used by “placebo” websites, i.e. websites that do not provide unlicensed content, but are otherwise similar to (e.g. in terms of audience size) to piracy websites in our data. As a robustness check, we compare the changes in the number of “placebo” websites served by treated and untreated third-party services before and after the date of the first step of the self-regulation.

4.2 Data collection

Although we are mostly interested in changes in the strategy of third-party services, our data collection begins at the publisher level.

We start with a list of domains of piracy websites obtained from the IT security firm *Shalla* (shallalist.de, “warez” category). To construct our sample of placebo websites, we match this with the list of the top 1 million most frequently visited websites from market research firm *Alexa*. The idea is to find similarly popular legitimate websites for each piracy website. We use domains that are placed one position above and below in the Alexa top 1m ranking list and are not a piracy website.

Adapting a software tool developed by [Libert \(2015\)](#), we visit each domain in the piracy and placebo group to collect data on the third-party HTTP requests the landing page

makes. We did this at three points in time, on August, 10th 2016, November, 17th 2016, and February, 2nd 2017, such that we observe data from 2.5 months before and 2.5 months after the self-regulatory agreement was signed on October, 21, 2016. Obviously, three (random) snapshots of before and after are not the optimal sample to work with. For the next draft of this paper, we plan to access historical data from *CommonCrawl.org* to have a longer time dimension in the data.

To be able to distinguish between advertising-related and non-advertising-related third party requests, we match the domains of third-party requests to a list of known advertising services provided by *Ghostery*, a market research firm that provides a browser extension that allows consumers to make HTTP requests visible and blockable.

We infer the geographical location of an advertising service from the top level domain (TLD). That is, we assume that a domain that ends with *.co.uk* is operated by a firm legally based in the United Kingdom, *.de* in Germany, *.it* in Italy, etc. Obviously, this is not the best measure, because some country-specific TLDs can be registered by bodies outside the country, and there are international TLDs such as *.com*. We plan to refine this measure in future versions of the paper using manually collected external data.

Our final sample consists of 247 piracy websites and 492 placebo websites. We observe 496 distinct domains of third-party services, out of which 359 (72.4%) are pure advertising services. The rest are website tools (e.g. audio/video players, customer interaction, comments, site analytics, social media).

5 Preliminary Evidence

In what follows we present some preliminary evidence regarding the hypotheses laid out above. Note that we are not making any causal statements at this point.

5.1 Direct effects

We begin with a test of the direct effects described in hypotheses 1 and 2.

The unit of observation in table 1 is an advertising service that has served at least one piracy or placebo website before and after October 2016. The dependent variable is the number of piracy websites that are served by this service. On average, the advertising

services in our sample serve 4.4 piracy websites (std dev. 11.68), the minimum is 0 and the maximum is 130. 53.6% (223) advertising services in our sample serve piracy websites, 4.2% (23) use TLDs that are specific to countries in the European Union. Only 3.88% (15) services are in both categories.

Table 1: Direct effects, only advertising services

	(1) Piracy	(2) Piracy	(3) Piracy	(4) Placebo	(5) Placebo	(6) Placebo
After \times EU	-2.809*** (0.761)		0.287** (0.118)	-0.424 (0.322)		-0.740 (0.648)
After \times Served		2.110*** (0.569)	2.352*** (0.598)		-0.088 (0.313)	-0.077 (0.327)
After \times Served \times EU			-4.405*** (0.974)			0.392 (0.764)
Observations	933	933	933	933	933	933
$\overline{R^2}$	0.895	0.897	0.898	0.977	0.977	0.977

Dependent variables: Number of piracy/control websites served per third-party service.

After indicates period after October 21, 2016, *EU* indicates a TLD of a EU-country.

Served indicates whether the service has serviced piracy websites before October 21, 2016.

All models include third-party service, top-level-domain and time period fixed effects.

White robust standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

In column 1, we define the treatment group as advertising services operating within the EU. Obviously, this is likely to be endogenous, because TLD choice, or firm location, is not random. However, the firm-fixed effect should capture some unobserved heterogeneity. As expected, the results suggest that “European” advertising services reduce their activities on piracy websites. The point estimate is quite sizable, equivalent to almost two-thirds of the mean, or almost a quarter of the standard deviation.

In column 2, we estimate a difference-in-differences model where we define firms that have served at least one piracy website before the self-regulation as the treatment group. Here we find, perhaps quite surprisingly, a sizable positive effect. Again, note that serving piracy websites is perhaps not random, which is why cannot interpret the parameter estimates as causal. However, again, the firm-fixed effect should help to reduce the bias.

In column 3 we look at the triple difference, i.e. define the treatment group as EU-firms that served at least one piracy website before the regulation came into effect. Looking

at the coefficient of $After \times Served \times EU$, we see a reduction in the number of piracy websites served regarding those advertising services for which the regulation is probably most binding (EU-TLD and history of serving piracy websites). Together with the results in column 1, this suggests some evidence in favor of hypothesis 1.

Comparing columns 2 and 3, the results suggests some support of hypothesis 2 as well. First, the positive coefficient of $After \times EU$ in column 3 may directly imply that firms that did not serve the piracy industry before start to do so. Second, the fact that the positive coefficient of $After \times Served$ in column 2 continues to hold in column 3, together with the negative coefficient of $After \times Served \times EU$ suggests that firms with a history of serving piracy websites react differently to the regulation. Those, who are supposedly directly affected (EU-based), reduce their services for this type of publisher, those who are less (not) affected (non-EU based) step up. Maybe this is evidence for a strategic response to the regulation.

To put those results into perspective, it is important to keep in mind that the group of advertising firms, for which we observe a negative trend over time represents only tiny fraction of the overall market.

In columns 4–6 of table 1, we repeat the same exercises, but instead use the number of placebo (non-piracy) websites served per advertising service as the dependent variable. The fact that we do not find any statistical significant coefficients here makes use somehow confident that the results regarding piracy websites in columns 1–3 are not entirely driven by some kind of spurious correlation.

5.2 Indirect effects

We now proceed with some preliminary evidence regarding indirect effects, depending on the type of third-party service, as described in hypothesis 3.

The results reported in table 2 are based on the full sample of third-party services; advertising and all other services. On this sample we estimate a version of the model in columns 3 and 6 of table 1. The difference is that we introduce a fourth interaction with $Ad - Service$, an indicator whether the third-party service mainly serves advertising.

The results in column 1 provide support for hypothesis 3. $After \times Served \times EU$ is not

Table 2: Indirect effects, contrasting advertising and other third-party services

	(1) Piracy	(2) Placebo
After \times Served	-0.469 (0.882)	-0.349 (0.466)
After \times EU	-0.041 (0.240)	-0.037 (0.351)
After \times Served \times EU	0.219 (1.178)	1.849 (1.163)
After \times Ad-Service	0.177** (0.079)	0.208 (0.203)
After \times Served \times Ad-Service	2.818*** (1.065)	0.269 (0.569)
After \times EU \times Ad-Service	0.323 (0.292)	-0.708 (0.681)
After \times Served \times EU \times Ad-Service	-4.621*** (1.530)	-1.454 (1.362)
Observations	1313	1313
$\overline{R^2}$	0.949	0.986

Dependent variables: Number of piracy/placebo websites served per third-party service.

After indicates period after October 21, 2016.

After indicates period after October 21, 2016, *EU* indicates a TLD of a EU-country.

Served indicates whether the service has serviced piracy websites before October 21, 2016.

All models include third-party service, top-level-domain and time period fixed effects.

White robust standard errors in parentheses.

* $p < 0.10$, ** $p < 0.05$ *** $p < 0.01$

significant, implying that non-advertising services do not change their strategy of serving piracy websites as a response to the regulation. The estimate of *After* \times *Served* \times *EU* \times *Ad-Service* provides a direct test of hypothesis 3. Advertising services react significantly more strongly than other third-party services to the regulation. Again, when estimating the same model, but looking at the number of placebo websites as the dependent variable in column 2, we do not find any significant coefficients.

6 Conclusions

In this paper, we aim to evaluate the effectiveness of “follow the money” copyright enforcement. We use a unique dataset that allows us to observe the interconnections between

advertising services and websites between piracy websites and legitimate websites. We compare the dynamics before and after representatives of the advertising industry has agreed on basic principles of an self-regulatory agreement to reduce the financial incentives for piracy websites.

The hypothesized mechanisms lead to the conclusion that the self-regulatory efforts are not likely to be very effective in reducing the incentives to create rogue content online. Preliminary non-causal evidence suggests that advertising services indeed reduce their activities on piracy websites, however only those that operate within the European Union and are affiliated with industry associations.

There is much work yet to be done in this project. A major next step will be the collection of more historical data points before and after the first steps of the self-regulation came into effect. This will be helpful to understand whether the underlying assumptions that allow causal interpretation of coefficients in difference-in-differences models, e.g. common pre-trends make sense in our case.

Most obviously, we plan to extend our analysis to evaluate any further steps in the process of the regulatory effort.

We believe that we will ultimately be able to draw conclusions about the dynamic effects such a regulatory intervention could have. First, concentrated segments of the advertising market could become more concentrated because of the exit of some players. Second, tracking services could grow in importance relative to services that combine advertising and tracking. Third, social network services may increase their presence in the tracking market, which potentially translates into a more dominant position in the advertising market. Because the market for online piracy is large, those effects may be important in the aggregate.

We believe that our study could not only contribute to the literature on online copyright enforcement, but may also provide implications for different settings where the availability of online advertising creates incentives to produce content that has negative externalities. A recent example are “fake news”, that seemed to be play an important role in the 2016 US presidential elections, or least in the corresponding public debate.

References

- Acquisti, A., Taylor, C., and Wagman, L. (2016). “The Economics of Privacy.” *Journal of Economic Literature*, 54(2), 442–492.
- Adermon, A., and Liang, C.-Y. (2014). “Piracy and Music Sales: The Effects of an Anti-Piracy Law.” *Journal of Economic Behavior & Organization*, 105, 90–106.
- Aguilar, L., Claussen, J., and Peukert, C. (2017). “Catch Me If You Can: Effectiveness and Consequences of Online Copyright Enforcement.” *Working paper*.
- Athey, S., Calvano, E., and Gans, J. (forthcoming). “The Impact of Consumer Multi-homing on Advertising Markets and Media Competition.” *Management Science*.
- Chaabane, A., Kaafar, M., and Borelli, R. (2012). “Big Friend is Watching You: Analyzing Online Social Network Tracking Capabilities.” *WOSN 12 Proceedings*.
- Danaher, B., and Smith, M. (2014). “Gone in 60 Seconds: The Impact of the Megaupload Shutdown on Movie Sales.” *International Journal of Industrial Organization*, 33, 1–8.
- Danaher, B., Smith, M. D., Telang, R., and Chen, S. (2014). “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.
- de Cornière, A., and de Nijs, R. (2016). “Online Advertising and Privacy.” *RAND Journal of Economics*, 15(3), 311–327.
- Estrada-Jiménez, J., Parra-Arnau, J., Rodríguez-Hoyos, A., and Forné, J. (2017). “Online Advertising: Analysis of Privacy Threats and Protection Approaches.” *Computer Communications*, 100, 32–51.
- Ghosh, A., Mahdian, M., McAfee, R. P., and Vassilvitskii, S. (2015). “To Match or Not to Match: Economics of Cookie Matching in Online Advertising.” *ACM Transactions on Economics and Computation*, 3(2), article 12.
- Gill, P., Erramili, V., Chaintreau, A., Krishnamurthy, B., and Papagiannaki, D. (2013). “Follow the Money ? Understanding Economics of Online Aggregation and Advertising.” *Proceedings of the 2013 ACM Internet Measurement Conference*.
- Goldfarb, A. (2014). “What is Different About Online Advertising?” *Review of Industrial Organization*, 44, 115–129.
- Lerner, A., Kornfeld Simpson, A., Kohno, T., and Roesner, F. (2016). “Internet Jones and the Raiders of the Lost Tracker: An Archaeological Study of Web Tracking from 1996 to 2016.” *Proceedings of the 25th USENIX Security Symposium*.
- Libert, T. (2015). “Exposing the Hidden Web: An Analysis of Third-Party HTTP Requests on One Million Websites.” *International Journal of Communication*, 9, 3544–3561.
- McKenzie, J. (2017). “Graduated Response Policies to Digital Piracy: Do They Increase Box Office Revenues of Movies?” *Information Economics and Policy*, 38(1), 1–11.

- Peitz, M., and Reisinger, M. (2015). “The Economics of Internet Media.” In S. Anderson, J. Waldfogel, and D. Stromberg (Eds.), *Handbook of Media Economics*, vol. 1B, 445–530, Elsevier.
- Peukert, C., Claussen, J., and Kretschmer, T. (2017). “Piracy and Box Office Movie Revenues: Evidence from Megaupload.” *International Journal of Industrial Organization*, 52, 188–215.
- Poort, J., Leenheer, J., van der Ham, J., and Dumitru, C. (2014). “Baywatch: Two approaches to measure the effects of blocking access to The Pirate Bay.” *Telecommunications Policy*, 38, 383–392.
- Reimers, I. (2016). “Can Private Copyright Protection be Effective? Evidence from Book Publishing.” *Journal of Law and Economics*, 59(2), 411–440.
- Roosendaal, A. (2012). “We Are All Connected to Facebook... by Facebook!” In S. Gutwirth, R. Leenes, P. DeHert, and Y. Pouillet (Eds.), *European Data Protection: In Good Health?*, 3–19, Springer, New York.
- Schelter, S., and Kunegis, J. (2016). “On the Ubiquity of Web Tracking: Insights from a Billion-Page Web Crawl.” *Working paper*.
- Xiang, Y., and Sarvary, M. (2013). “Buying and Selling Information Under Competition.” *Quantitative Marketing and Economics*, 11, 321–351.