■第 31 回東京国際映画祭　共催企画　第 8 回 MPA セミナー

日　時：2018 年 10 月 26 日（金）

主　催：モーション・ピクチャー・アソシエーション（MPA）

　　日本国際映画著作権協会（JIMCA）

共　催：公益財団法人ユニジャパン

後　援：米国大使館

　　一般社団法人映画産業団体連合会

■講演 1

Mr. Brett Danaher,

Assistant Professor of Economics and Management Science at Chapman University

Good afternoon. My name is Brett Danaher. I'm a professor of economics and management science at Chapman University in the United States. I'm honored to be here today and have a chance a little bit of my research to you and talk to you a little bit about measuring the effectiveness of piracy website blocking using data and evidence-based approaches.

Before I do that, there's one question I think I should answer, which is as an academic, we have the freedom to research almost anything that we want and that's is sort of part of our job. That's why we go into academia. Why have I spent so much of my life researching antipiracy enforcement and what's effective and what's not?

The answer to that is that looking around the world and looking at the debate over copyright enforcement, I found that it to be a very visceral debate. I found that tones of people's arguments tend to be almost

theological in nature.　People have very strong beliefs about copyright, about piracy, whether it should be enforced or should not be enforced.

My training is as a scientist.　As a scientist, I feel that a little data and a little analysis can go a long way in trying to help us all come to the same side or at least respect each other's opinions on topics like this.

One of the things I want to say first is that I'm not here today as an advocate for copyright or for the entertainment industries, nor am I here as an advocate for pirates or for the internet or anything like that.　The only thing I am is an advocate for data and an advocate for evidence-based policy making.

As a little background, my colleagues—I have several coauthors for the research that I've been doing.　My colleagues are Carnegie Melon University in the United States, and we've spent close to 10 years studying various types of antipiracy enforcement in a very scientific manner where it's not so much that we care about whether they work or not.　We're just interested in knowing whether they work or not.　Which ones are effective and which ones are not effective?

We like to think of this as two different types of – we like to say that we see antipiracy enforcement in two different flavors.　One is demand-site antipiracy enforcement.　This is actions or policies that target consumers of pirated content.　We've studied a number of these.　We've studied the HADOPI law in France, which had penalties – which required that internet service providers monitor users' activity and impose penalties upon individuals that were caught pirating content.

We've looked at as well other more softly paternalistic demand-site

measures such as search rank manipulation. What would happen if search engines like Google were forced to promote legal links and demote legal links and what effects can that have on consumers' choice between pirated content and legitimate content?

Also, we have the flavor that we call supply-site antipiracy policies. These are antipiracy enforcement actions that target sites or protocols that facilitate piracy – kind of the other side of the coin. We've looked a number of these. We've looked at the effectiveness of shutting down megaupload.com in 2012. Before that it was the world's largest piracy cyber locker, and we've looked at whether it was effective in shutting it down, whether that caused a change in legitimate consumption or legal purchases.

We've looked at several instances of court- or government-ordered piracy website blocking, and we've also looked at county-specific shut downs of piracy sites that largely served up content specific to a particular culture.

We've done a lot of research in this area. This is just a small sampling of what we've done. We're also very well aware of the research of other independent colleagues of ours at other universities. I would say this is my bread and butter. I eat, sleep and breathe antipiracy research.

One of the things that I've noticed in having the chance to study so many different antipiracy enforcements actions, I've been able to get sort of an overarching viewpoint on what seems to be effective, what's not effective, but also what seems to be politically viable. One of the things that we've noticed is that demand-side antipiracy policies, the policies that target consumers, they are effective at increasing legitimate sales, but they don't appear to be politically very viable.

In most of the cases that have been studied, while we see an increase in legitimate consumption and in legal purchases, following these policies, we often find that enforcement grows lax over time or that the laws are actually overturned over time because it appears as if demand-side antipiracy enforcement simply there is not a strong political (00:05:00) taste for it among many countries.

On the other hand, when we have studies supply-side antipiracy enforcement, we've seen something different. We've seen that it's much more widely used around the world. Targeting pirate sites or protocols that support internet piracy, we found that to be relatively politically viable around the world, that a number of countries have implemented these sorts of policies.

One interesting thing that I'm going to try to show you today is that piracy websites blocking appears to be much more agile enforcement action than actually piracy website shutdown. Actually, entirely shutting down a piracy site seems to be a lot more difficult and less agile of a policy than just asking ISPs or requiring ISPs to block access to those sites.

One of the things I'm going to try to show you in the data today is that in antipiracy enforcement, agility is actually a very important asset. Let's see if I can show you what I mean by that.

First, in terms of piracy website blocking, there are number of countries around the world where this has been implemented, which I think there's a lot of things that can be said about that but one is that to a scientist that means there's a lot of experiments for me to potentially study. At this point, I've only studied the United Kingdom, but we're actually

working on a study on piracy website blocking in Australia as well.

The question that a lot of policymakers have is is piracy website blocking effective. I should step back for a moment and say that I understand there are a lot of questions regarding piracy website blocking. I know that there are legal questions about it, and those may be specific to each country. I know that there are technical questions about it, and I know that some people have raised ethical questions about it. I think that those are all very good discussions.

I know a few things about the answers to those questions, but the truth is I'm not an expert on any of those things. But one of the questions that seem to come up everywhere that this is implemented is will it be effective? That's something that I think data and some empirical approaches from economics and from data science can help us to figure out.

If we want to ask if piracy website blocking is effective, I would argue there are three questions that we need to ask. The first is does piracy website blocking actually cause a decline in visits to the blocked sites? That's a very basic question. Does it work? Can we actually get it to work? Can we keep people from getting to the sites that we've blocked?

The second question is does it cause a decline in total pirated downloads? If I block access to one site and everybody who was pirating there simply goes to another site, I would argue that that's not a very meaningful change in behavior. The question is when we block access to some piracy sites and there are certainly other piracy sites out there, do we actually cause a decrease in piracy?

Finally the third, and I would argue the most important question is, does it cause an increase in legitimate consumption? Does it cause an increase in visits to legal sites, purchases from legal sites? Again I would argue if we don't see this, then I would argue piracy website blocking isn't really effective. If I can stop a lot of piracy but I can't increase legitimate consumption, then I would argue that content is getting to less people and nobody is benefitting as a result of this policy. I think it's really important to ask whether we get an increase in legitimate consumption.

The problem here of course is that measuring the impact of any policy intervention, and I don't care what the policy intervention is, but measuring the impact of any policy intervention is difficult. Many of us have probably heard the maxim correlation is not causation. In other words, if I just find correlations in the data, that doesn't necessarily mean that I can prove that the thing that I'm looking at caused the changes that I see in the data.

For example, if I'm looking at entertainment and if I'm looking at legitimate consumption and legal sales of movies or books or music, all of these outcome variables are changing over time for lots of reasons. Right? There's declining interest in older content as it ages. There are spikes of interest in new content. You can imagine that when Game of Thrones releases new episodes, legal sales across the world go up as a result of a single show.

And then, of course, there is this natural seasonality around holidays and periods of time when students are in school or not in school. So, there are a lot of reasons that the outcome variables that we care about when thinking about piracy website blocking are changing that have nothing to do with website blocking.

That’s where data science and the tools of econometrics and empirical methods come in really handy, and the question that I'm going to pose to you and I'm going to try to help you answer here is how do we actually measure the causal effect of piracy website blocking on things like total piracy and on legitimate consumption of media.

(00:10:03)

Well, we have three experiments that we studied in our research. First in May 2012, courts in the United Kingdom ordered internet service providers to block access to the Pirate Bay, which was the most popular piracy site in the UK at the time. That was the first experiment that we were able to study. Then about a year and a half later in November 2013, the courts ordered internet service providers to block access to 19 major video piracy sites.

Finally, one year later, because they had sort of figured out how to do piracy website blocking efficiently, the courts ordered ISPs to block access to 53 more meaningful video piracy sites. I obtained a data set surrounding each one of these experiments, and I'm going to show you a little bit of how I used it to figure out the effect of these actions on legitimate consumption.

The data that I obtained is from an internet-panel tracking company. Essentially, thousands of internet users were tracked in terms of their visits to various sites month over month in the United Kingdom. The trick that I was able to exploit is the fact that some of these users in this panel didn’t access the block, they didn’t use the blocked sites at all. So, for example, if I'm studying the blocking of the Pirate Bay, in my data,

some of the users weren't using the Pirate Bay at all anyway. When these people were blocked from the Pirate Bay, you could think of them as almost like a control group. Any changes in their legitimate consumption aren't being caused by the blocks. That's my measure of what the seasonal trend is or what the market trends in consumption of legitimate content at the time are.

Then we had some consumers in our data that were actually making a few visits to the Pirate Bay. Perhaps they visited the Pirate Bay once or twice in the months before it was blocked. We consider these users to be lightly treated. When the Pirate Bay was blocked, they received a small shock to their piracy levels.

Then we had really heavy users of the Pirate Bay in our data. For example, there are users in our data that were accessing the Pirate Bay 200 or 300 times per month in the months before it was blocked. When that site was blocked, these users were very heavily treated. That was a very meaningful action targeting their piracy levels.

If website blocking is effective, then we would think that the heavy users of the blocked sites would make larger changes in their behavior than the lighter users of the blocked sites and the lighter users would make larger changes in their behavior than the people who weren't using the blocked sites.

That's how I am going to take the correlations in the data and try to tell you a story of causality, the actual effect of website blocking on the users that were blocked from these sites.

I observed in my data month by month by consumer visits to the blocked

piracy sites, visits to remaining unblocked piracy sites, visits to paid legal streaming sites and as a small bonus, I also observed visits to virtual private network sites. One of the arguments about piracy website blocking is that it might not be effective because people will subscribe to virtual private networks to circumvent the blocks on websites. One of the things I can do is look in the data and see whether that story is true as well.

You can't trust anybody's empirical argument unless you understand what they are trying to do with the data. I want to really make sure that you understand what I'm trying to do here. Imagine I just have three users. I don't. I have thousands of users, but imagine I just have three users. If one is a nonuser of the Pirate Bay, that's like my control group in a medical trial. When the Pirate Bay is blocked, they were given a sugar pill, a placebo pill.

Then I've got users who were light users. When the Pirate Bay was blocked, it's like they were given a pill that has 20% medicine but 80% sugar in it. Then I have these really heavy users of the Pirate Bay. When the pirate Bay was blocked, it's like they were given a pill and it was all medicine.

If you were to do a medical trial like this, you would ask, does the recovery of patients after being given the pills correlate with the amount of medicine that was in the pill that were given. That's exactly what I'm doing with the data that I have.

If you were to see a pattern like this, if you were to see the nonusers make only small changes in legitimate or even no change in legitimate consumption, but that light users make larger changes and heavier users

make the largest changes, then you would conclude the piracy website blocking is increasing legal behavior. On the other hand, if you see any other pattern other than this, you would argue that there's not much evidence of an increase in legitimate consumption as a result of the blocks.

(00:15:03)

I'm going to share a little bit of my results with you. Obviously, there's a lot of sophisticated, as an academic, you have to use a lot of Greek symbols and a lot of sophisticated methodology in your papers. This is not my statistics class, so I'm going to try to just give you a small insight, a few graphs that show you a little bit about what I saw in the data.

First thing I want to say is that the blocks were effective in reducing traffic to the blocked sites. At the very least, it does appear as if ISPs can generally keep people from getting to blocked sites. After each wave of blocks, traffic to blocked sites decreased by about 90%.

Now, the other 10%, where is that coming from? There are a couple of reasons. Some ISPs didn’t comply right away. Some users did use virtual private networks to get around the blocks but for the most part there was a 90% drop in visits to the blocked sites. But as I argued to you, that's not really the most important question. The real question is do I cause decreases in total piracy and do I cause increases in legal site visits and legal consumption.

Here are a couple scatter plots and let me talk you through these. In each of these scatter plots, each dot actually represents a group of hundreds of consumers. Even though, you only see ten dots on each

plot, this is actually representing in total thousands of consumers that I observed in my data.

In each of these scatter plots, the X-axis is pre-blocked visits to the blocked sites. In other words, before the blocked were enacted, how often, how frequently were the users in this group visiting those blocked sites? If this is zero, they're like my control group and if it's really high, these are the people that were really accessing the Pirate Bay or whatever sites were being blocked a lot. The Y-axis is the change in total piracy levels. What was the actual change in total piracy levels of those individuals after the blocks were enacted?

In both cases here you can see negative relationships on the scatter plots. In both cases, it is true that the blocks caused decreases in total piracy but the slopes of the best fit lines for each of these graphs are very different and they tell very different stories.

When the Pirate Bay alone was blocked in May 2012, the slope of this line tells me that out of every 12 blocked visits to the Pirate Bay, total piracy only went down by one visit. In other words, out of every 12 blocked visits to the Pirate Bay, 11 of those people found other piracy sites to go to. Blocking the Pirate Bay had only a very small effect on total piracy levels.

However, in November 2013, when 19 of the major piracy sites were blocked at once, the slope of the line is actually very different. When I look at the slope of that line, it tells me that for every three blocked visits to the blocked sites, there were 2.5 fewer visits to all piracy sites after the blocks. In other words, out of every six blocked visits, five of them were an actual decrease in total piracy. Only one of those people ended up

finding some other piracy site to go to.

Now, the question you might ask yourself is that's fine if I decreased total piracy or don't decrease total piracy but what happened to legitimate consumption. And again I don't want to give a full lecture on statistics here but there is an interesting pattern in the graphs. In these scatter plots, the X-axis is the total piracy change before and after the blocks and the Y-axis is the change in legitimate consumption, visits to legal sites like Netflix or iTunes. Here you see very different stories.

When the Pirate Bay was blocked in 2012, there is no correlation between the change in total piracy and the change in legal visits. What that tells me is that the blocking of the Pirate Bay didn't cause any uplift in legal consumption that I can measure and that's not surprising because 11 out of 12 piracy site visitors found another site to go to anyway.

In November 2013, when 19 sites were blocked, you see a very different pattern in the data here. You can almost see these points lining up along a negatively-sloped line.

What this says is that after the blocks happened, the more that you decreased your total visits to piracy sites, the more that you increased your visits to legal sites. People who only decreased their piracy by a little bit didn't increase their legal very much and people who decreased their visits to piracy sites by a lot had large changes in visits to sites like Netflix and iTunes. This is where we actually see the causal impact come into play.

(00:20:02)

Now, we do a lot of thing in the paper to make sure that this is really, truly an effect of the blocks that were happening.　I don't want to delve into all of it, but to give you an example, the relationships that I just showed you don't exist in months before the blocks happened.　If you look, I've got a line here; this is essentially measuring the slope of those lines in that scatter plot.　In all the months before the blocks happened, you don't see any real relationship between those two things.　Only when the blocks happen do you see a relationship between the decrease in total piracy and the change in visits to legal consumption sites like Netflix.

If you're interested in more of the evidence, the paper is available on the internet.　I encourage you to take a look.　But if you just want to result summarized for you, here's what we find.　Here are our conclusions.

In May 2012, when the Pirate Bay was blocked, it caused a 90% decrease in visits to the Pirate Bay, but it only caused a very small decrease in total piracy and it caused no increase in legitimate consumption.　Netflix didn't appear to benefit from this at all.

In November 2013, the story is quite different.　In November 2013, the blocking of 19 sites actually caused a decrease of 30% in visits to all piracy sites, and it caused a 12% increase in visits to legal sites like Netflix and iTunes.

Finally, in November 2014, when 53 sites were blocked at once, I see a relatively similar story.　Again, I see a 10% increase in usage of paid legal streaming sites and a 12% increase in usage of legal ad-supported sites like YouTube Movies or since this is the UK, the BBC iPlayer.

How do I explain this?　How do I explain that some actions work and

some actions do not?　I think the answer here makes a lot of sense.　It's almost common sense.　When you block access to one piracy site, even if it's incredibly popular, it's not that hard for people to find another piracy site.　In fact, if I ask my students in my class what's your favorite piracy site, they all have an answer.　If I ask them what's your second favorite piracy site, most of them have an answer.　But then I always do one thing.　I say what's your 25th favorite piracy site and usually people don't have much to say at that point.

I would argue that that's exactly what we're seeing here is that if you block one site, no matter how popular it is, that people can find their way around that block.　But when you block a number of the most reliable sites at once, people find it difficult to find another site.　It's not that it's hard to find a piracy site, but to find one that's reliable, one that you trust, one that you're not worried about getting a virus from, one that you know how to use properly, that is rather difficult.

I would argue, while some of those people continue to pirate, a number of them actually make the decision to come over to legal sites.　That's where you see the effect of piracy website blocking.

So, the conclusions from this research, blocking just one popular site had little impact on total piracy and no impact on legal consumption. Simultaneous blocking of a number of sites significantly reduced piracy in the UK and increased legal consumption.　And it does appear to be the case that site blocking can be an effective tool in converting consumption from illegal channels into legal ones.

But it comes with caveats.　Clearly antipiracy enforcement works when you can make piracy difficult enough for the user.　Weak actions, weak

enforcement actions may not make piracy terribly difficult. It's stronger enforcement actions that block access to a number of sites that where you actually see a result.

I put out just one other paper here because I want to tell you, as academics, we trust our own research but it's also important that this research is validated. There are other papers and peer-reviewed journals that find very similar results to what we've been finding.

There is a paper here that found a very similar result. When one site was blocked in Germany, they actually found people just went to other piracy sites. And that's the most recently published work on this. Our research shows that that is true but that the full story is more nuanced, that while one site can't have much of an impact, blocking multiple sites does.

Now, I told you that I was going to speak to the idea that agility is important in antipiracy enforcement. What I would argue here is that the reason that piracy website blocking is more agile than actually shutting down sites is that if I want to shut down (00:25:00) I raid the servers and sue or shut down 20 different sites, that's probably almost 20 times as expensive as shutting down a single site. But ordering ISPs to block access to 20 different sites is nearly the same cost as ordering them to block access to one site, very little difference in cost.

I would argue that since our research seems to show that it's important to go after a number of sites at once and probably that it's important to keep doing it over time as people find new piracy sites and start to get comfortable with them, I would argue that site blocking is the more agile tool relative to actually trying to shut down piracy sites and remove their

content from the internet.

The real conclusion from this is when implemented correctly, piracy website blocking appears to be more effective than site shutdowns, and since we've studied site shutdowns in our research, I can tell you that the percentage impact does appear to be larger here with piracy website blocking.

An important component of website blocking appears to be that it must be maintained over time. As you can imagine, while we found these impacts on legal consumption, if we never went back to site blocking, eventually another dominant piracy site would emerge and over time people would get comfortable with it. So, it is continued action that appears to cause the uplift that we saw in legal consumption.

Thank you very much for the chance to talk here today.