The Short- and Long-Term Effectiveness of Anti-Piracy Laws and Enforcement Actions

Abstract Film studios have spent the past two decades lobbying extensively to establish new legislation restricting access to copyrighted materials online. While there is growing evidence of the effect film piracy has on studio profits, the evidence on the impact of anti-piracy legislation is limited. If anti-piracy legislation is having the film industry's desired impact we would expect film revenues to be consistently higher following the passage of major laws that restrict access to pirated content, or major enforcement actions, such as the shutdown of websites that provide illegal content for download. This paper applies an intervention analysis approach to weekly data on movie box-office revenues in the US to determine whether the passage of new anti-piracy policy has generated significant changes in box-office revenues during the period from 1997 to the present. These effects are evaluated in both the short and long term, which allows an assessment of the duration of effectiveness of government actions. The results show that four of the six included policies are ineffective in the long-term and those policies that do impact revenues in the short term often harm film studios, rather than help them.

Keywords Online piracy, Motion-pictures industry, Intervention analysis, Intellectual property rights

JEL Classification C22, K11, L82, Z18

1 Introduction

Film industry leaders have long held that strong legislation restricting access to pirated content is necessary for the industry to maintain its profitability. The industry spends millions of dollars each year lobbying for increased enforcement of existing policies and the creation of new legislation. This legislation has largely been viewed as a necessary solution to a difficult problem, but little academic research has evaluated the efficacy of such programs. This paper seeks to evaluate the historical record of major anti-piracy legislation in the US.

Early investigations of film piracy evolved from discussions of unenforceable copyrights. Varian (2005) provides a concise summary of the study of copyright enforcement by economic scholars. Research has since expanded to more directly address the particular issues facing the film industry. McKenzie's (2012) literature review of the economics of the film industry provides a succinct explanation of these challenges. Work specifically focusing on film piracy initially produced mixed results. Bounie *et al.* (2006) found that the piracy of films generated increased demand for purchasing films, either in theaters or through home media. Their study, which focused on students at a French university, found that students who illegally pirated films are actually more likely to rent or buy copies of films. They found that films that are downloaded benefit from increased demand. However, DeVany and Walls (2007) found evidence that contradicted this claim. They found that an increased availability of downloads produces a marked decrease in fim-industry revenue as measured by box-office revenue, DVDs and legal downloads. This work is particularly pertinent to the analysis of the effectiveness of anti-piracy policy because it identifies a direct causal link between revenues and the ease of downloading a film (not the number of downloads). Their study confirms the film industry's conventional wisdom: more available downloads translate into lower revenue for production studios.

McKenzie and Walls (2014) find similar results. Their paper evaluates the impact of a film

being downloaded on its Australian box-office revenue. In addition, the paper utilizes a pair of Australian court cases on online piracy as a part of their identification strategy. They find that downloads have a small but statistically significant negative impact on the revenues in the film industry in Australia. In addition, while their use of the Australian court cases is primarily used as a tool in identification, they do find that these cases have a statistically significant impact on the number of downloads of films. This finding indicates that government policies do have a noticeable impact on the flow of illegal content online.

The missing link in the economics literature thus far has been evaluating whether copyright law is able to remain effective over the long term. The law literature on major anti-piracy laws has generally agreed on two major conclusions: First, the ability of most anti-piracy policies to have a dramatic impact on actual downloads is limited by the unenforceable nature of much of online copyright law (Heneghan 2002). Second, anti-piracy policies tend to have a large immediate impact, but their effects diminish over time as the anti-piracy actions are rendered less effective by evolving technologies for sharing files online (Sudler 2013). This fits with the conventional wisdom on online piracy: innovations by those who infringe on copyright eventually circumvent existing laws and technology. For every government preventive action there is a method of getting around it, and it is only a matter of time before that method is discovered and employed on a large scale.

This paper tests this conventional wisdom empirically using data on film box-office revenues. The paper uses an intervention analysis approach to evaluate six major anti-piracy policies. The paper proceeds as follows: Section 2 discusses intervention analysis and its application in this paper; Section 3 discusses the data; Section 4 fits a number of intervention function models to the data and interprets the results of these econometric models; and section 5 concludes.

2 Data

In order to assess the effectiveness of anti-piracy actions by the government in the film industry, this paper examines weekly data on U.S. box-office revenues. Revenue data is obtained from Box-Office Mojo (2013) which collects publicly released revenues of individual films and releases, aggregated daily, weekly and monthly. Weekly data is ideal for evaluating movie industry performance because it is the time frame that studios typically use in evaluating the success of films.

Box-office data ranges from Week 1 of 1997, which allows us nearly a year of data before the enactment of the first policy, to Week 39 of 2013, for a total of 874 weeks of data. Revenues are adjusted for inflation, with 2012 as the base year. The series of log revenues, measured in millions of real dollars, is shown in Figure 1.

Figure 1: Weekly US Box-Office Revenue, 1997 to 2013

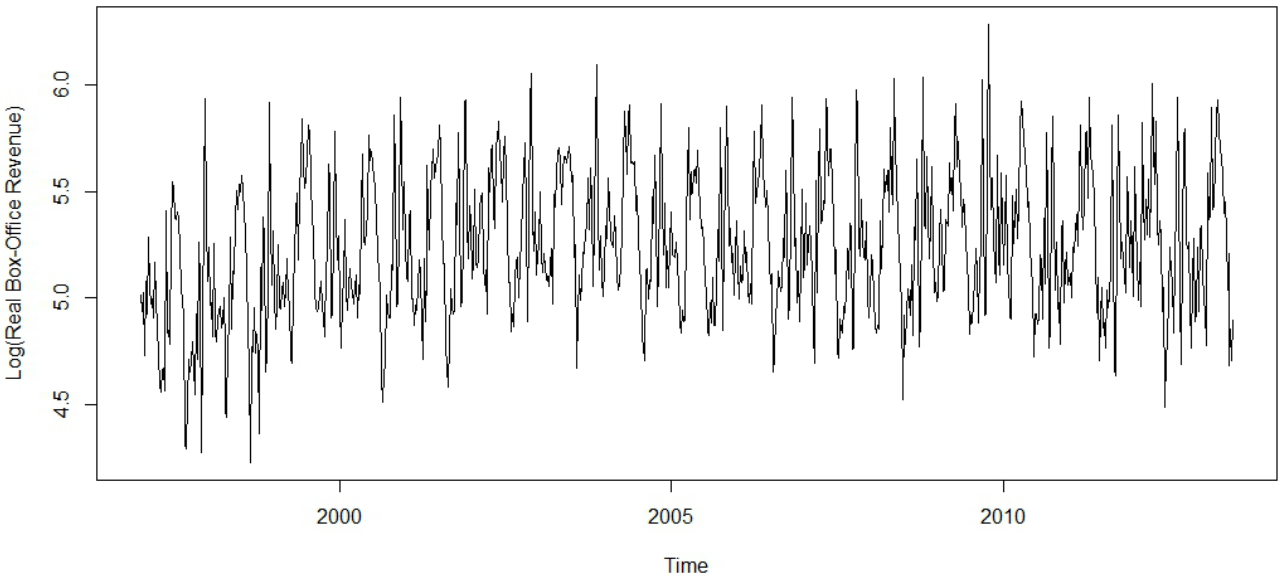

There is a high degree of seasonality in the revenues of the film industry, with summer and holiday releases receiving much higher returns on average. The data is seasonally adjusted using 53 weekly dummy variables, as well as holiday variables for those holidays that may fall on different

weeks in different years (e.g. Christmas falls in week 52 of 1999, but in week 51 of 2000). Dummy variables are also used for films produced before 1998 and films produced before 2001. Box-Office Mojo's data collection method changed in these years, resulting in more total films being accounted for in the box-office revenue figures. The dummy variables account for this change in data collection[1]. The inflation and seasonality adjusted data is shown in Figure 2. There no longer appears to be any regular pattern in the data, and the mean does not change over time. The data shown in Figure 2 appears, based on visual inspection, to be stable and stationary, a requirement for most time-series applications.

Figure 2: Seasonally Adjusted US Box-Office Revenue, 1997 to 2013

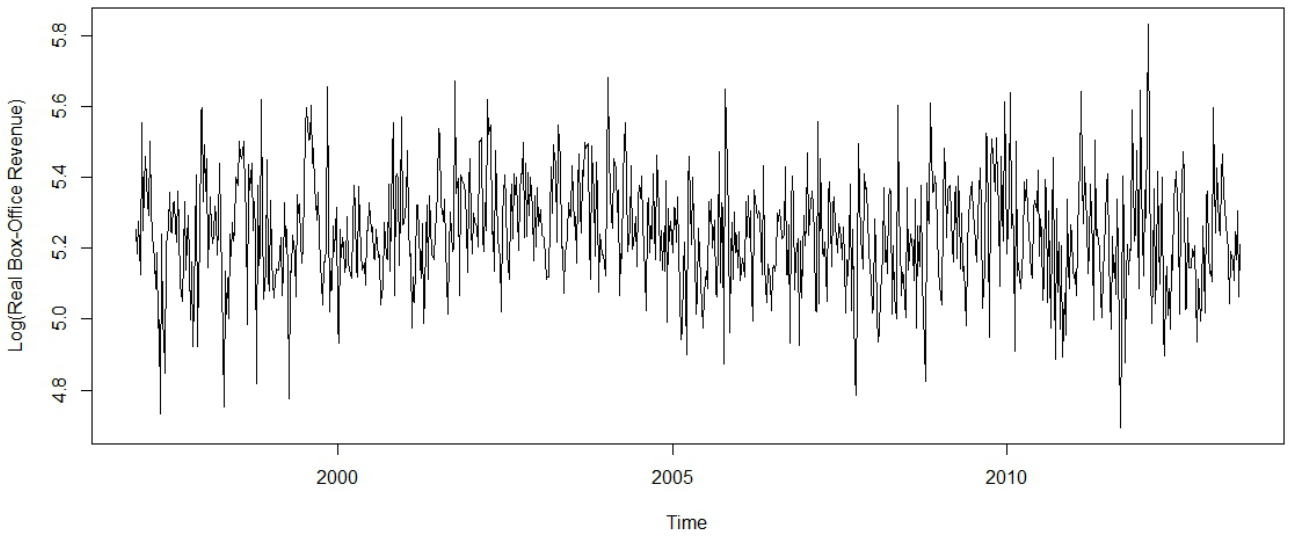

This data is used to evaluate the long-term effectiveness of six government actions. Each action falls into one of two categories: laws, which establish rules for how copyrighted material is regulated by the government; and enforcement actions, which are direct interventions by government officials to remove websites that violate copyright law. While laws and enforcement actions clearly affect content

1 Prior to 1998 only the highest grossing films are recorded, usually 8-15 films per week. From 1998 to 2001 more films, but not all, are counted, resulting in data on 40-70 films per week. After 2001 all films, even those showing in a limited number of theaters, with limited revenues are included. Over this time period revenues are routinely reported for 90-130 films per week. This problem would produce selection bias (it excludes the lowest grossing films from early data) if it is not accounted for with dummy variables.

differently, the goal is the same: to reduce the availability of pirated content online. All six policies have the primary goal of limiting the supply of pirated material and so this paper treats their impacts similarly. The six policies are considered to be the widest-reaching and most important anti-piracy policies in recent years. They are: the No Electronic Theft (NET) Act; the Digital Millenium Copyright Act (DMCA); the shutdown of the website Napster; the Family Entertainment and Copyright Act; the Prioritizing Resources and Organization for Intellectual Property (PRO-IP) Act; and the shutdown of the website Megaupload.

The NET Act of 1997 was the first policy to facilitate the prosecution of only copyright violations in the same way as copyright infringement of traditional media (NET Act, 1997). The policy re-defined the use of the terms "commercial advantage" and "monetary gain" in previous copyright law. This closed a previous loophole whereby actions that disseminated copyrighted material online, without any direct monetary gain, could not be prosecuted. The result was that uploading and downloading copyrighted material became justification for intellectual property based lawsuits. This law was thus the genesis of intellectual property protection online because it extended the applicability of laws protecting copyright holders so they could be reasonably applied to online activities.

With the DMCA, passed in 1998, the US adopted two major international treaties on copyright protection, with some additional policy actions (Digital Millenium Copyright Act, 1998). The law included a multitude of policy changes, but a small number of major changes particularly impacted the availability of online content. The law strengthened penalties for violating copyright law, made it illegal to circumvent anti-piracy controls (even if the copyright is not violated after this circumvention), and added liability to those who assist in the distribution of copyrighted materials. Perhaps its largest impact was the creation of the system of notices and subsequent lawsuits against those hosting copyrighted material, which is still in place today. The DMCA was a major change in copyright law that dramatically strengthened the enforceability of copyright law online.

The website Napster was, in the late 1990's and early 2000's, the largest source of illegally downloaded content online (MPAA, 2013). The website allowed individuals to share files between one another. While sharing MP3 music files was by far the most popular type of transfer on the site, movies were also shared. In early 2000, a number of recording studios filed a lawsuit against Napster and after a protracted legal battle the website was shut down in July 2001 as a result of an injunction imposed in federal court. This shutdown eliminated the primary source of copyrighted material online, and while the impact was likely not as large on the film industry as it was on the recording industry, it probably had a significant effect on the supply of pirated materials.

The Family Entertainment and Copyright Act, passed in 2005, had two major provisions (Family Entertainment and Copyright Act, 2005). First, the policy imposed stricter punishments for copyright violations. Second, and more important, the law made it a federal crime to record copyrighted material at a theater. This law banned the use of camcorders in theaters and provided that individuals could be punished without proof of intent to distribute. Some states had already passed similar restrictions on camcorders, but this policy had both harsher penalties and wider reach than these state policies[2]. Video recordings of films are a major source of pirated content online (MPAA, 2013) and this policy sought to restrict the number of recordings available.

The PRO-IP Act, passed in 2008, was an omnibus bill that took a large number of actions all designed to protect copyrighted materials (PRO-IP, 2008). The major tenets of the law were to increase penalties for violations of copyright, allow for the seizure of property and financial assets related to copyright violation, create an executive-level position for the coordination of copyright enforcement, allocate funds to anti-piracy projects and expand the types of content distribution that are classified as

2 While this policy is largely an expansion of existing state legislation, rather than an all new policy, it is treated the same as the other policies in the study. This is done for three reasons: 1. The state policies were not nearly as well enforced as the new national policy. 2. The national press associated with a major law like this likely had a major impact on film revenues that a state policy would not have. 3. Given the ease of copying a filmed movie it is largely irrelevant whether a movie is recorded in one state or in 20. The federal policy established for the first time that it was difficult to record films in theaters across the nation.

copyright violating. The law specifically targeted large-scale efforts to distribute copyrighted material and made it much easier for the government to both investigate and prosecute those who had violated copyrights.

Megaupload was another website which, like Napster before it, was a major source of copyright-infringing downloads. Similar to Napster, in that it hosted content that was uploaded by users, the major difference between it and peer to peer (p2p) file sharing sites was that it allowed for streaming of content, rather than downloads. US federal prosecutors filed charges against Megaupload on January 19, 2012 resulting in an FBI seizure of servers, taking the website offline (Kravets, 2012). The shutdown of the website likely reduced the number of available downloads significantly.

The goal of this paper is to evaluate the effectiveness of these government actions.

3 Methodology

In addition to providing a test of the immediate effectiveness of these anti-piracy government and legislative actions, this paper seeks to evaluate whether these effects are sustained over time. The theoretical argument for why the effectiveness of laws is expected to change over time is a simple one. Laws are only able to effectively legislate technologies that are currently used and understood. Laws that attempt to create rules that apply, regardless of the technology, will almost inevitably be unenforceable (Heneghan 2002). Further, anti-piracy laws cannot predict future technologies for creating and distributing illegal content, and thus will be unable to produce enforceable rules for these technologies. It follows that as new technologies emerge that facilitate the distribution of illegal content, past laws will become ineffective in regulating these new technologies, and new laws will be required to limit their use. The difficulty of creating enforceable and meaningful policies is even more pronounced in the case of website shutdowns. Shutting down a website may be expected to reduce the availability of downloads in the short term; however, due to the legal and physical enforcement costs of

shutting down a website, it is impossible for authorities to shut down all sites hosting illegal content. Thus after a fairly short time, we expect illegal content to migrate to active file sharing sites, and the website shutdown's effect will be reduced.

To evaluate the long-term effectiveness of anti-piracy laws and enforcement actions, this paper utilizes intervention analysis. Intervention analysis, first developed by Box and Tiao (1975) allows for the dependent variable to be influenced by one or more exogenous variables (interventions). Further, it allows for a number of functional forms to model the interventions. Enders (2010) provides a succinct summary of intervention analysis and the various intervention functions that are commonly used to model the effects of an exogenous shock.

This paper utilizes weekly data on inflation-adjusted US box-office revenues and determines whether government policies have a significant impact on these revenues. This is done by including a variable z for each policy. In intervention analysis, these exogenous z variables are called intervention functions. Rather than restricting z to be a dummy variable which is either zero or one, we allow a number of specifications for z.

This paper models box-office revenue with the equation of the form,

$$y_t = a_0 + A(L)y_{t-1} + B(L)z_{it} + C(L)\varepsilon_t \quad (1)$$

where y_t is the dependent variable (film revenues), z_{it} are the six intervention functions for our exogenous variables, ε_t is a normally distributed error term and A(L), B(L) and C(L) are vectors of coefficients multiplied by lag functions. This general form of the equation can be re-written in more traditional notation if the lag-orders of y, z and ε are known. For example: if we plan to model y_t using an ARMA(2,2) form with one lag of a single exogenous variable then we can re-write equation (1) as:

$$y_t = a_0 + a_1 y_{t-1} + a_2 y_{t-2} + b_0 z_t + b_1 z_{t-1} + c_1 \varepsilon_{t-1} + c_2 \varepsilon_{t-2} + \varepsilon_t \quad (2)$$

A key consideration in investigating the long-term impact of exogenous shocks is determining the shape of the intervention function. This paper considers multiple possible forms for intervention

functions for each *z* included in the final regression equation. The analysis considers six possible intervention functions for each of our policies. In all six, the intervention variable *z* is zero prior to the passage of the policy in question. In the first intervention function *z* is one for all following weeks. This is the "dummy variable specification." The second intervention function has *z* equal to one on the date of passage and zero for all following weeks, this is the "pure pulse specification." The third and fourth intervention functions assign *z* a value of one on the date of passage, then the value stays at one for one and two years, respectively. If *z* is one for a single year we will refer to this as the "one-year pulse specification." If *z* is one for two years following the passage of the policy it is the "two-year pulse specification." The fifth intervention function assigns *z* a value of one on the date of passage, but the value of *z* declines steadily over the following two years. At all points over two years from the date of passage the value of *z* is zero. This is the "declining pulse specification." The final intervention function has *z* steadily ascend from zero to one over the two years following policy passage. The value of *z* stays at one for the remaining weeks. This is the "increasing dummy specification." A visual depiction of the value of *z* following policy passage for each intervention function is included in Figure 3.

This method of allowing a number of specifications for *z* is not uncommon, but the included specifications for the *z* variable is different than usually used. It is quite common in empirical research to allow the *z* variable to be a dummy variable, or to use the "increasing dummy specification" (Enders, 2010). This standard method assumes that since the policy has not been repealed it must stay in force for the duration of the data-set. Information criteria are used to determine how long it takes for the policy to become fully effective, but there is no consideration of the policy becoming ineffective without some sort of subsequent legislative action or policy change. This paper breaks from the standard method by allowing a number of *z* specifications described earlier, which have the *z* variable

Figure 3: Value of z_i in Different Intervention Function Specifications

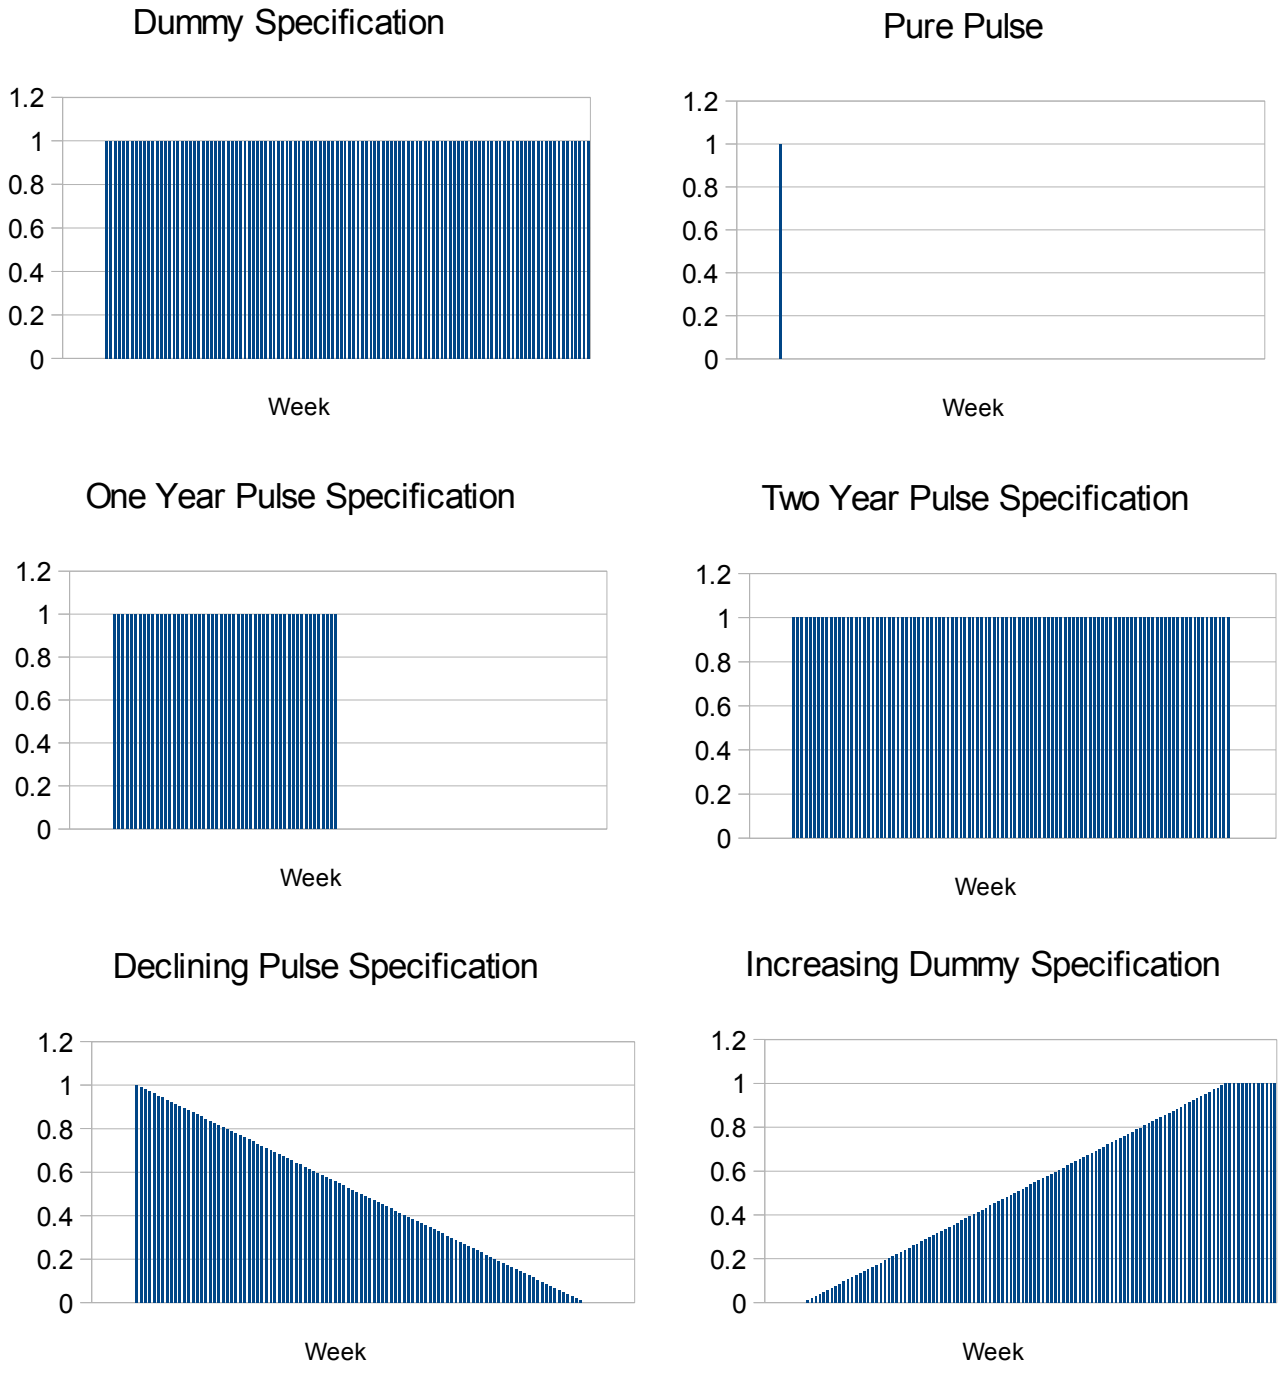

decrease from 1 to 0 either gradually or suddenly. In the standard intervention analysis these *z* variables would not be considered unless the policy were repealed.

The situation of anti-piracy policy is unique, however. As discussed earlier, there are a number of theoretical reasons that a policy could become ineffective or completely irrelevant without the passage of a new law or the repeal of the policy. Because there is the possibility of a de-facto repeal of the policy, which would not be evident in the legal record, this paper considers a number of intervention function specifications that would ordinarily only be available if a law were taken off the books. Other papers have used information criteria to select among intervention functions (e.g. Worthington & Valadkhani, 2004); this paper simply allows for a broader range of possible *z* variables due to the unique situation created by the evolving technology of online piracy and law enforcement.

Each of these different intervention function specifications implies a different duration of effectiveness of the policy that it models. The dummy specification has the policy effective immediately and its effectiveness is permanent. The pure pulse specification implies that the policy is truly temporary, and only has an effect in the time period in which it is passed, though the autoregressive nature of the model allows this short term impact to gradually fade (as it does in all of the intervention function specifications). The one-year pulse specification implies that the policy is only effective for one year, then becomes immediately ineffective. The two-year pulse specification has an identical interpretation to the one-year pulse, but with the time frame of two years rather than one. The declining pulse specification has the policy only fully effective in the week that it is passed, but its effectiveness does not disappear overnight. Instead the policy steadily declines in effectiveness over two years. Finally, the increasing dummy specification implies that the policy is ineffective when it is first passed; over the next two years the policy steadily becomes effective, then stays effective.

Given these different implications about the longevity of effectiveness of policies it is important that we choose the correct intervention function for each policy; however, given the computational

intensity (there are 6^6 possible specifications of the full model) it is not feasible (nor is it econometrically advisable) to test all possible models. Thus, this paper individually evaluates models using an initial estimation that has an equation of the form,

$$yt = a_0 + A(L)y_{t-1} + B(L)z_t + C(L)\varepsilon_t \quad (3)$$

which is identical to the more general model specified in Equation 1, except that it only includes the intervention function for a single policy rather than for all six.

The model specified in Equation 3 is estimated six times for each policy, once for each specification of the intervention function *z*. For example: to determine the best intervention function to model the effect of the NET Act, a regression is run using the form,

$$y_t = a_0 + A(L)y_{t-1} + B(L)z_{NET,t} + C(L)\varepsilon_t \quad (4)$$

where *z* is specified by the dummy variable specification. That is, z takes a value of zero before the passage of the NET Act and one after the passage of the NET Act. A separate regression is then estimated using the form of Equation 4 but with a pure pulse specification. Similar regressions are estimated for the four other intervention function specifications. These six equations are thus specified identically, other than the values of the *z* variable. The results of these regressions are compared according to information criteria. The best performing model, according to these information criteria, is assumed to be the best fit for the effectiveness of the NET Act. Once *z* for the NET Act is chosen, an identical procedure is used to choose *z* for each of the five other policies in our dataset. This gives us six *z* variables that can be used in our complete model.

This complete model has the form specified in Equation 1, with the specific intervention functions *z* being those chosen in our initial estimations. Once intervention functions are specified, and the coefficients are estimated, the coefficients on the intervention functions tell us the immediate effect of the interventions, and long-term impacts can be investigated by constructing impulse response functions.

4 Results

The application of this methodology to the impact of government anti-piracy actions begins with an evaluation of the appropriate ARMA order for the revenue data. Different information criteria choose different ARMA orders in this case. The Akaike Information Criterion (AIC) indicates an ARMA(1,2) process is best to model the data; the Bayesian Information Criterion (BIC) chooses AR(1). Visual inspection of the Autocorrelation Function (ACF) and Partial Autocorrelation Function (PACF) of the residuals from these two models indicates that an ARMA(1,2) process provides a better fit, so this paper proceeds with an ARMA(1,2) process to model box-office revenues. The paper does utilize results from an AR(1) process as a robustness check[3].

Two major assumptions of the intervention analysis model must be considered before intervention functions can be fitted. First, the intervention variables must be exogenous. That is, the passage of anti-piracy legislation cannot be in response to changes in film revenues. While there is some concern that the film industry would increase lobbying for anti-piracy government actions during times when their revenues have fallen, the time lag between increased desire for action and the passage of legislation is long and variable. It is unlikely that the government is able to respond to declines in film revenue quickly enough for there to be a meaningful effect of our dependent variable on the timing of policy actions.

Second, the method for estimating the coefficients of the intervention functions assumes that the only exogenous shocks are those included in the estimation equation. If there are other exogenous shocks affecting film revenues that are not controlled for, it is possible that these shocks would cause our coefficients to be biased. This paper controls for demand-side shocks by including both the previously discussed seasonal adjustment and also yearly dummy variables. Given the steady level of

3 The results do not change significantly with an AR(1) rather than ARMA(1,2) process. Concerns about overfitting by using the AIC should be dramatically reduced due to the similarity of results. It is also worth noting that since the goal of this paper is explanatory, not predictive, the possibility of overfitting is less of a concern.

inflation-adjusted revenues, it seems unlikely that there are other major demand shocks that would influence our estimates. Similarly, on the supply side, the film industry has been remarkably stable in the time frame considered in this paper. The six major studios have not changed during this time period, no major technological changes seem to have shifted the cost of producing films and there appear not to have been major changes in the release schedules or strategies of major film studios. With any intervention analysis there are concerns about other exogenous shocks, but the film industry in the late 1990's and early twenty-first century seems to have been especially stable other than the emergence of copyright-infringing activity online.

With an ARMA(1,2) process we can identify the proper forms of the intervention functions. For each policy we fit equations of the form[4]:

$$y_t = a_0 + a_1 y_{t-1} + b_0 z_t + b_1 z_{t-1} + b_2 z_{t-2} + b_3 z_{t-3} + c_1 \varepsilon_{t-1} + c_2 \varepsilon_{t-2} + \varepsilon_t \quad (5)$$

For each policy, six regressions of the form given in Equation 5 are run, but with z values given by a different intervention function specification. These regressions are compared with one another, using AIC and BIC, to choose which specification for z is best to model the impact of the policy considered. This procedure is undertaken for each of our six policies so that there is an intervention function specification for the z variable used to model each policy.

In all specifications of this equation, other than the pure pulse specifications, all coefficients on the lags of the intervention function (b_2 and b_3) are statistically insignificant. This result holds throughout for all of our initial and combined regressions, so lags of the intervention function will not be included in the results shown in this paper. These lags are included in the regressions, but always drop out when a preferred version of the model is specified.

These initial regressions identify which intervention function is the best fit for each policy. The AIC used to choose these intervention functions (generated from our initial regressions) are included in

4 This version of the equation includes three lags of the z variable, but this choice is arbitrary. Many versions of this equation were tested with many different lag orders of z.

Table 1. Since the number of variables in each equation is identical, BIC and AIC indicate the same models in all cases.

The dummy specification is chosen for the NET Act and the Family Entertainment and Copyright Act. This indicates that these policies had an immediate effect that remained permanent over time. This result seems to be consistent with the basic understanding of what these laws do. The NET Act did not apply to specific actions or technologies, but rather extended a set of rules that had already been in place (and that are constantly evolving) into the online community, regardless of the monetary benefits that website hosts and uploaders of content stand to gain. It makes sense that this sort of policy would be less susceptible to a decline in effectiveness than the other policies in the dataset. The result is slightly more surprising in the case of the Family Entertainment and Copyright Act. This law did directly target a method of uploading content (camcorders) and thus we might expect that the law would decline in effectiveness as new methods of uploading content replaced camcorder content theft. It may be that methods of obtaining illegal copies of a film are less likely to evolve over time than are other technologies surrounding piracy, or there may be another explanation. Further research into the language and effectiveness of the Family Entertainment and Copyright Act would be useful to provide context for this result.

Two policies are best modeled by the two-year pulse specification. Both the DMCA and the Napster shutdown are fitted best with this intervention function. This indicates that while these policies did have an immediate effect, the laws became ineffective after a period of two years. This result suggest that, while these policies have some impact on the time trend of box-office revenues, they do not have a permanent effect. This result is consistent with the hypothesis that policies become obsolete as technology evolves around them.

Table 1: AIC Values in Initial Regressions

	Policy Evaluated					
	NET	DMCA	Napster	Family	PRO-IP	Megaupload
Dummy	**-859.89**	-858.06	-858.37	**-862.59**	-858.40	-858.25
Pure Pulse	-858.01	-858.24	-858.07	-858.04	-858.02	-858.02
1 Year Pulse	-859.75	-858.01	-861.96	-858.42	-858.03	-858.02
2 Year Pulse	-859.18	**-859.09**	**-865.29**	-859.13	-858.63	-858.25
Declining Pulse	-859.32	-858.76	-865.02	-858.84	-858.35	-858.02
Increasing Dummy	-858.05	-858.10	-858.29	-861.40	**-859.00**	**-858.68**

Finally, both the PRO-IP Act and the Megaupload shutdown are modeled by the increasing dummy specification of intervention function. This result indicates that the policies take some time to become effective before maintaining a permanent effect. It is important to note, however, that these two policies are the most recent in the dataset. It may be that this specification is driven not by the long-term staying power of the policies, but rather by the fact that we do not have enough data to fully evaluate the form of the intervention function of these policies. The best fitting intervention functions are included in the combined regression, the regression with the same form as Equation 1, to serve as controls, but it is unlikely that we will be able to evaluate the full impact of these two policies for a number of years.

The combined regression can then be constructed using the intervention functions identified in the initial regressions. The estimating equation is of the form given in Equation 1. Results are presented in Table 2. The combined model produces a significantly stronger fit than with our simple ARMA(1,2) equation without intervention functions. It is preferred by both AIC and BIC. Figure 4 provides the residuals from the reduced regression, as well as the ACF and PACF functions of these residuals. Visual inspection of the residuals shows that they are white noise. The lack of statistically significant values in both the ACF and PACF support this conclusion. The residuals indicate a strong fit of the model, and do not appear to indicate misspecification or that our results are spurrious.

Table 2: Combined Regression Results

	Combined Regression		Reduced Specification	
Variable	Coefficient	S.E.	Coefficient	S.E.
a_1	0.7284***	0.1019	0.7391***	0.0975
c_1	-0.4479***	0.1078	-0.457***	0.1038
c_2	-0.1127**	0.0474	-0.115**	0.0469
a_0	5.1955***	0.0329	5.1956***	0.0333
z_{NET}	0.0819**	0.037	0.0928***	0.036
z_{DMCA}	-0.0586**	0.0283	-0.0693***	0.0268
$z_{Napster}$	0.0297	0.0283		
z_{Family}	-0.0606**	0.0237	-0.0669***	0.0181
z_{PRO-IP}	0.0155	0.0281		
$z_{Megaupload}$	-0.0286	0.0597		
AIC	-864.50		AIC	-869.06
BIC	-812.00		BIC	-830.87

N = 874

*** = 1% significance, ** = 5% significance, * = 10% significance

Three policies, the Napster Shutdown, the PRO-IP Act, and the Megaupload shutdown, are all found to have statistically insignificant effects and so are removed from our reduced specification. This is not to say that these policies had no impact on the number of downloads, but they did not make a noticeable impact on the box-office revenues of the film industry. Two of these policies are the most recent, and thus our results may be unable to fully evaluate their impact, but further research should be dedicated to understanding why these policies, which include both of the website shutdowns in our dataset, did not impact revenues in motion-pictures exhibition.

The NET Act, the DMCA, and the Family Entertainment and Copyright Act are all found to have statistically significant impacts on box-office revenues. Their coefficients, given in Table 1, are measures of the immediate impact of these policies. For example, when the NET Act was passed, revenues in the film industry increased by 9% in the following week. This short-term effect (impact effect) is a measure of the impact the policy had immediately, but due to the autoregressive nature of

Figure 4: Residuals from Reduced Specification

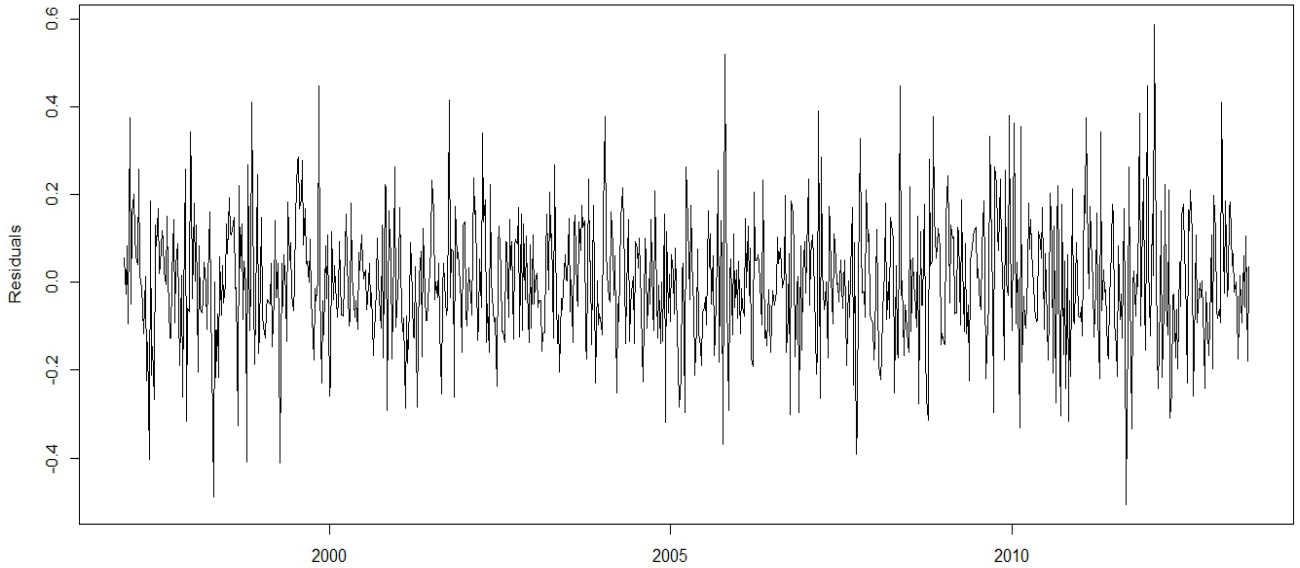

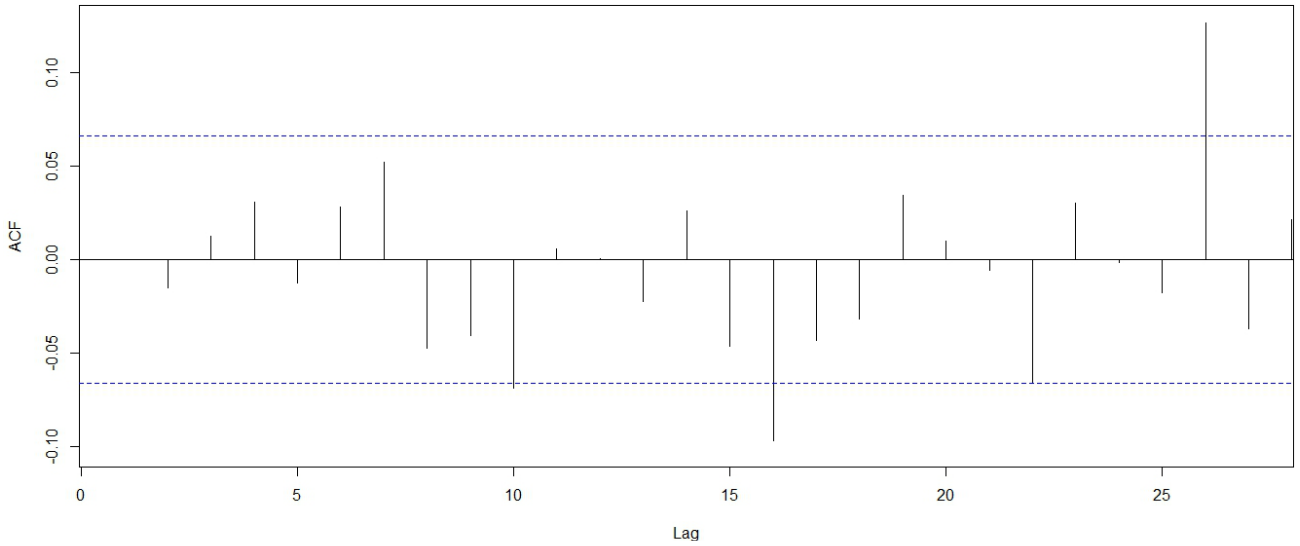

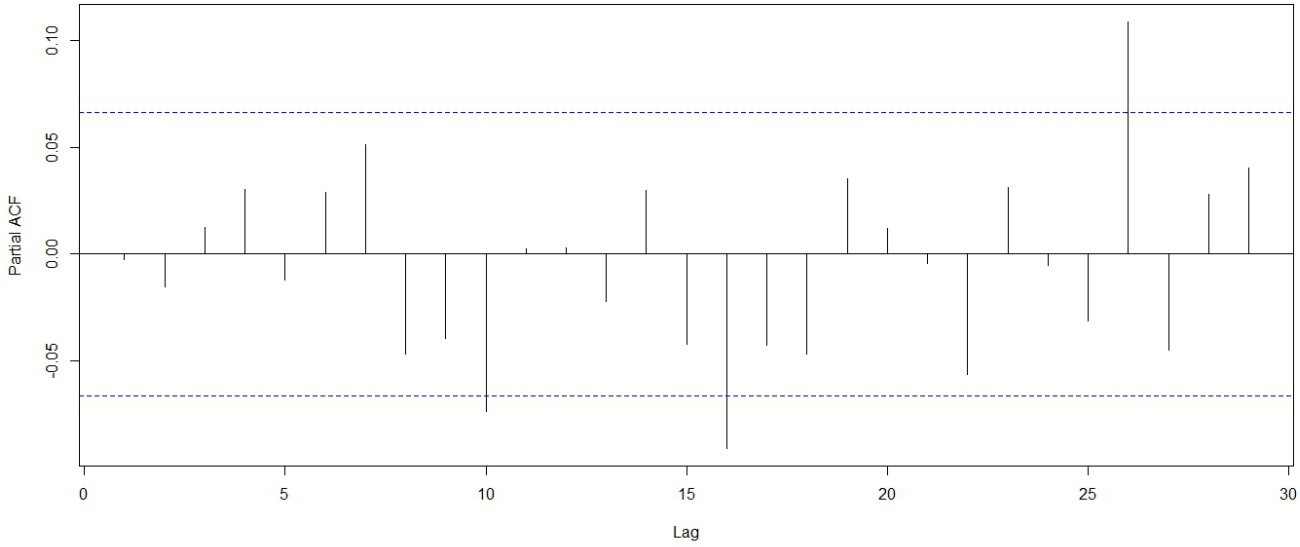

our model, the long-term impact is quite different. The evolution of the impacts of these policies can be evaluated by constructing impulse response functions to demonstrate how their impact evolves over time. Impulse response functions are shown in Figure 5. These graphs illustrate the effect that these three policies had in the weeks following their passage.

In addition, we can calculate the long-run effect of these policies using the equation,

$$\text{Long-Run Effect} = b_i/(1-a_1) \tag{6}$$

were b_i is the coefficient of the policy we are evaluating and a_1 is the coefficient of the autoregressive term. This long-run effect gives the largest impact that the policy has. For the NET Act and the Family Entertainment and Copyright Act, this is a measure of the impact that the law has permanently once the law becomes fully effective. For the DMCA, the long-run effect is a measure of the weekly impact of the law being in place before its effectiveness declines after two years. Short- and long-run effects of the statistically significant policies are given in Table 3.

Since our dependent variable is logged, but our intervention variable is not, the interpretations of these effects are that a change of one in the independent variable will result in a percent change in revenues equal to the short- and long-term effects in Table 3. The NET Act, upon passage, raised revenues of the film industry by 9.3%, and once the policy had reached its peak effectiveness, the law was able to increase weekly revenue over 35%. Since the policy is modeled by the dummy variable intervention function, this increase was permanent. The NET Act is thus a policy-maker's ideal scenario. It had a major positive impact on film revenues that has stood the test of time.

The DMCA, on the other hand, immediately reduced box-office revenues by 6.9%, and at the peak of its effectiveness, film revenues had fallen by 26.6% as a result of its passage. This effect was not permanent, however, and after two years the policy's impact on film revenues quickly declined. The impermanence of this effect is seen in the oddly shaped impulse response function for the DMCA. Because the *z* variable goes from 1 to 0 abruptly at a certain point, we have an impulse response

Figure 5: The Evolution of the Impacts of Policies Over Time

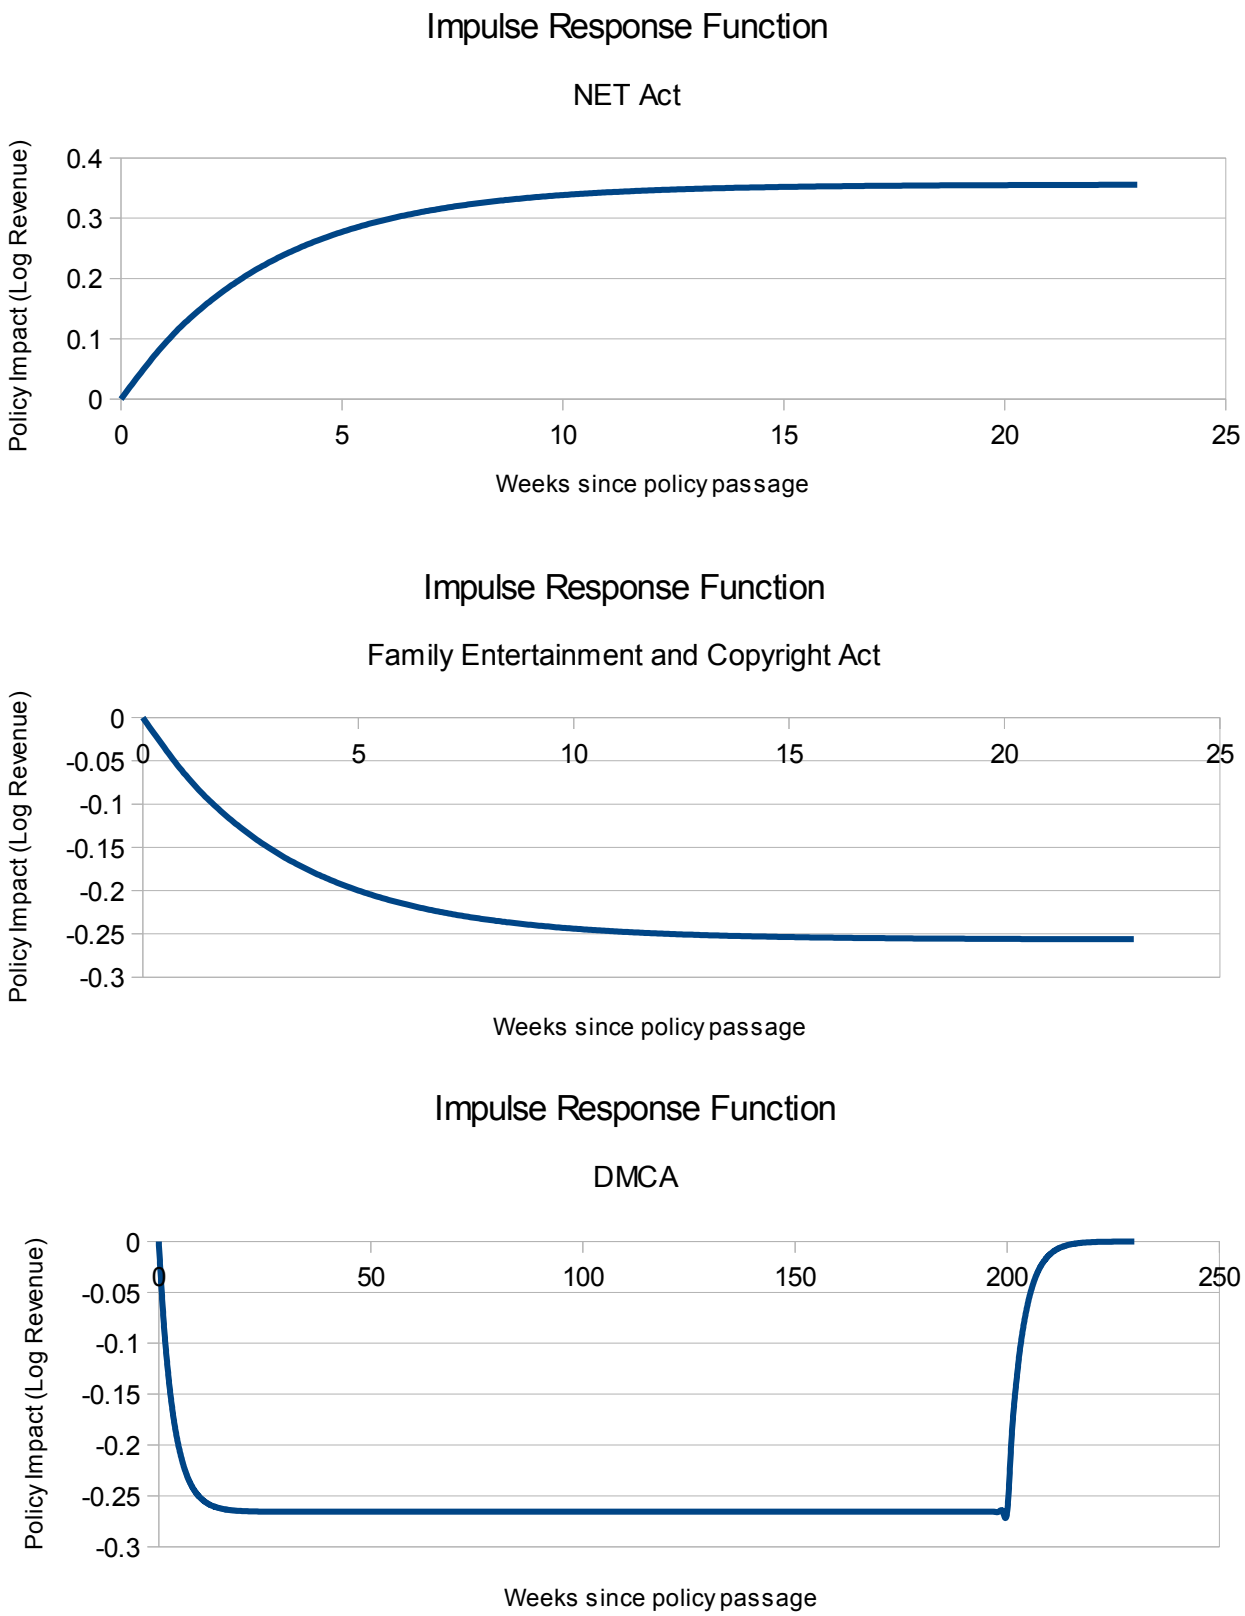

Table 3: Short and Long-Run Effects of Anti-Piracy Policies

Percent Change in Revenues

	Short-run Effect	Long-Run Effect
NET	0.093	0.356
DMCA	-0.069	-0.266
Family	-0.067	-0.256

function that would ordinarily only be seen if a policy were repealed. In this case, however, it indicates that the policy became ineffective due to logistical or technological issues preventing the DMCA from having a meaningful impact on film revenues.

Finally, the Family Entertainment and Copyright Act had a permanent depressing effect on box-office revenues. Immediately following its passage, revenue fell by 6.7%, and the long-term impact of the policy is an over 25% decline in box-office revenues.

5 Discussion and Conclusion

The results show that in the long run, four of the six anti-piracy policies considered have no impact on box-office revenues, and that even in the short run, where the NET, DMCA and Family Entertainment and Copyright Acts all have some impact, two of the policies actually harmed the revenues of the film industry, rather than improving them. These results are perhaps counter-intuitive in that they indicate that policies designed to limit piracy actually decreased box-office revenue. It is possible that these policies are having the reverse of their intended effect and in some way are increasing the number of downloads of films. More likely, however, is that some interaction between the substitution effect observed by DeVany and Walls (2007) and the word-of-mouth effect noted by Bounie *et al.* (2006) is taking place. There are likely some consumers who will only watch a film for free online (their marginal benefit from viewing another film is well below the market price). Anti-

piracy laws stop these consumers from sharing their impressions of films (which may increase demand), but do not have any substitution effect on revenues, because these consumers would not have bought a movie ticket regardless of online availability. Other consumers may be willing to purchase a ticket to see a movie, but will embrace an illegal zero-cost option if it is presented to them. Restricting access to copyright-infringing material online does cause these individuals to substitute into buying theater tickets, and since the person is seeing the film regardless of the policy, there is at best a marginal word-of-mouth effect on the demand for films. The relative sizes of the substitution and word-of-mouth effects, as well as quantity of affected consumers who fall into these categories, is what determines whether the policies have a positive or negative effect on revenues. Future research should investigate these factors to determine how future policies will impact box-office revenue.

Despite a lack of specific theoretical research modeling the behavior of film goers in response to anti-piracy legislation, it is reasonable to hypothesize that different policies have different effects along these margins. If different policies impact certain consumer groups more than others, this would imply that some policies (e.g. NET) are able to reduce substitution out of theater attendance without harming the ability of some pirated copies of films to be discussed and generate word-of-mouth demand. Other policies (e.g. DMCA) may limit the ability of individuals to access pirated material, but they do so in a way that limits the ability of pirated films from generating good buzz, which can drive individuals to theaters. Further research should investigate the interaction between piracy, word of mouth, demand, and government policy to determine why some laws are more or less likely to raise revenue when they limit illegal access to copyrighted material.

In addition to investigating the trade-off between preventing substitution away from theater-going and preventing demand-generating information flows, this paper raises the question of how anti-piracy laws are formulated. The policies in this paper are considered to be the most wide-reaching and most effective US laws to address online piracy, and yet most of them either have no impact or have an

effect that is unintended. A more careful analysis of the texts of these laws may be able to answer why the impacts of these laws have been, on the whole, ineffective. This research would also be useful to policymakers seeking to design future laws with a permanent impact of raising revenues, like the NET Act, while avoiding the pitfalls of other policies that fall short of their goals.

Another key finding of this paper is that it appears that shutting down websites, seemingly the most direct method of reducing illegal access to copyrighted materials, does not have any noticeable impact on the revenues received by the US film industry. These results are likely to hold true for other developed countries with similar internet utilization, but policies may be more or less effective in developing countries. Further research could investigate why these policies have been ineffective and should evaluate how readily individuals change download sites following legal action against their preferred site. In addition, international investigation of online piracy and its impact on film revenues would provide a more clear picture of how film studios could best respond to growing piracy.

The only policy that had the desired long-term effect is the first policy in the dataset. It may be that the NET Act has already done a good job of taking the low-hanging fruit and limiting the most harmful online piracy. Laws subsequent to the NET Act may have gone too far and limited the ability of online piracy to help spread the word about quality films. These laws may also be targeting piracy that is hard to identify and prosecute. Future research should evaluate the NET Act, in particular, to determine whether its introduction has changed the evolution of piracy technology in such a way that future laws are either unnecessary or are more difficult to craft effectively.

Similar studies in other industries could be crucial to understanding whether or not it is simply the structure of the motion-pictures industry that is driving these results. Analysis of related industries such as television, music, live entertainment and others may find that these and other policies are effective in other arenas, or it may confirm that the problem of constructing effective anti-piracy policies fundamentally lies in the development of technology that can circumvent the goals of such

policies.

This paper applies an intervention analysis approach to determine the short- and long-term effectiveness of six major US anti-piracy government actions on the box-office revenues of the motion pictures industry. The results show that, with one notable exception, these laws have either been ineffective or counter-productive, from the perspective of the film studios. These results have widespread implications for both policymakers and those in the industry who are interested in limiting piracy and raising revenue. Future research into the policies surrounding online piracy should be able to expand upon these results and provide better explanations as to which policies are effective and how policy can be better crafted to benefit the owners of intellectual property and holders of copyrights.

References

Bounie, David, Bourreau, Marc, and Waelbroeck, Patrick. (2006). Piracy and the Demand for Films: Analysis of Piracy Behavior in French Universities. *Review of Economic Research on Copyright Issues* 3(2), 15-27.

Box, G.E.P., Tiao, G.C. (1975). Intervention analysis with applications to economic and environmental problems. *Journal of the American Statistical Association* 70, 70–74 .

Box-Office Mojo. http://boxofficemojo.com/weekly/ Accessed: November 1, 2013.

DeVany, Arthur S., Walls, W. David. (2007). Estimating the Effects of Movie Piracy on Box-Office Revenue. *Review of Industrial Organization* 30(4), 291-301.

Digital Millenium Copyright Act. Public Law 105–304. (1998).

Enders, Walter. (2010*). Applied Econometric Time Series*. Danvers: Wiley.

Family Entertainment and Copyright Act. Public Law 109-9. (2005).

Heneghan, Brian P. (2002). The NET Act, Fair Use, and Willfulness- Is Congress Making a Scarecrow of the Law? *Journal of High Technology Law* 1(1)*,* 27-46.

Kravets, David. (2012). Feds Shutter Megaupload, Arrest Executives. *Wired.*

McKenzie, Jordi. (2012). The Economics of Movies: A Literature Survey. *Journal Of Economic Surveys* 26(1), 42-70.

McKenzie, Jordi, Walls, W.D. (2013). File-Sharing and Film Revenues: An Empirical Analysis, Working Paper.

MPAA. http://www.mpaa.org/news/pr Accessed November 20, 2013.

NET Act. Public Law 105–147. (1997).

Prioritizing Resources and Organization for Intellectual Property Act of 2008. Public Law 110-403. (2008).

Sudler, Hasshi. (2013). Effectiveness of Anti-Piracy Technology: Finding Appropriate Solutions for Evolving Online Piracy, *Business Horizons* 56(2), 149-157.

Varian, Hal. (2005). Copying and Copyright, *Journal of Economic Perspectives* 19(2), 121-38.

Worthington, A., & Valadkhani, A. (2004). Measuring the Impact of Natural Disasters on Capital Markets: An Empirical Application Using Intervention Analysis. *Applied Economics,* 36(19), 2177-2186.