P E R S P E C T I V E S

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

Piracy and Movie Revenues: A Critical Review of *A Tale of the Long Tail*

George S. Ford, PhD*

September 16, 2013

Introduction

Theft isn't good for the victim. Yet, a recent study on piracy entitled *Piracy and Movie Revenues: Evidence from Megaupload: A Tale of the Long Tail?* (hereinafter "*Tale of the Long Tail*") conducted by researchers at the University of Munich and the Copenhagen Business School serves up the counter-intuitive claim that piracy may increase box office sales for *some* films.[1] To reach this conclusion, *Tale of the Long Tail* uses as an experiment the January 2012 shutdown of what was likely the single largest online distributor of pirated content—the file sharing website Megaupload.com.[2] Using data on *weekend* box office revenues for thousands of films shown in about fifty countries, the authors split this data into pre- and post-shutdown periods, and then apply econometric methods to test for a difference between the two samples, attributed any observed change to the shutdown.[3]

Naturally, as other studies have shown, the expectation is that piracy, if anything, reduces box office revenues, but the authors of the study offer a surprising, counter-intuitive conclusion.[4] Specifically, while the revenues of "small" and "blockbuster" movies (defined in reference to the number of theaters in which the film opened using a country-specific penetration rate) were mostly unaffected by the Megaupload shutdown, *Tale of the Long Tail* claims that the shutdown of the pirate site actually *reduced* box office revenues for "mid-size" films. The authors contend that this result suggests that pirated copies are a form of advertising (thus increasing sales), but the promotional effect applies only to films released in a middling number of theaters.

> ***Upon close inspection of Tale of the Long Tail, … the "piracy helps box office sales" result is revealed to be an artifact of a poorly-designed statistical model….***

While the unexpected conclusion drew significant media attention,[5] the result is sufficiently bizarre to make an audience of professional economists suspicious. Indeed, although theft may offer benefits to the thief, economists generally do not normally view theft as beneficial to the victim or to society as a whole.[6] Accordingly, a claim of such an effect obviously faces a high evidentiary standard.

Upon close inspection of *Tale of the Long Tail*, it is clear the authors have failed to meet this burden. To the contrary, the "piracy helps box office sales" result is an artifact of a poorly-designed statistical model, which is, in part, a consequence of the study's authors ignoring the basic economics of the box office. The defect in the *Tale of the Long Tail*'s model is revealed in the dubious implication of that model that a

"blockbuster" movie (as the authors define it) will earn no more revenue than a "small movie" and far less revenue than a "mid size" movie. In fact, according to the analysis in *Tale of the Long Tail,* what the authors define to be a "blockbuster" movie makes only one-fourth (1/4) of the revenues of mid-sized movies.[7] The bizarre result about the benefits of piracy is based upon this dubious implication of the statistical model.

In this PERSPECTIVE, I detail this problem with *Tale of the Long Tail*, and outline a few other defects. There are far too many problems with the study's empirical approach to cover them all in detail, including the fundamental question regarding whether the general "before" and "after" approach can identify the effect of piracy generally, and whether the authors have estimated a model capable of capturing the effect.[8] A thorough analysis of the study is unnecessary since the study's central (and most controversial) finding is easily dismissed.

Box Office Economics

To begin, let's establish the basic principle of profit maximization at the box office. Movie distributors and exhibitors face the problem of assigning available films to available screens, on an ongoing basis. Although there are various complications and constraints involved in these processes, from the economic point-of-view this problem has a very basic structure, and that structure, in the limit, yields a very simple characterization of profit-maximizing behavior. This characterization, in turn, implies a series of simple propositions about the relationships between theater (or screen) exhibitions and film revenues.[9] It is precisely this economic characterization which should be examined statistically, and should form the basis of any econometric analysis of this problem.

For simplicity, suppose an exhibitor has k potential films, which he may show on $n > 1$ potential screen venues. Further, suppose the exhibition costs, royalties, and so on are proportional to the revenues earned from ticket sales (this is an unnecessary but simplifying assumption). How should he allocate films among venues?

In this case, the basic principle is obvious—the exhibitor wishes to maximize the total of revenues earned on the n screens. In the limit, he will assign films to screens so that the revenues earned per screen are similar: if one screen showing film x earned far less than another showing film y, then he should consider replacing x with y when he is free to do so. Although it is reasonable to assume that showing the same film on many screens will result in continually diluted demand, and a fall in per screen revenues, the basic principle can be simplified to read as follows: show each film with sufficiently high demand (single screen revenue) until, on the margin, each film being exhibited has equal marginal revenue. This problem is very familiar to economists, and resembles the well-known problem of allocating a total output among a group of factories with varying cost structures.

The defect in the Tale of the Long Tail's model is revealed in the dubious implication of that model that a "blockbuster" movie (as the authors define it) will earn no more revenue than a "small movie" and far less revenue than a "mid size" movie. In fact, according to the analysis in Tale of the Long Tail, what the authors define to be a "blockbuster" movie makes only one-fourth (1/4) of the revenues of mid-sized movies.

Although the actual problem of scheduling movies and screens is more complex than this,

the basic principle is a very powerful one. It can also be tested, if only approximately. If this conceptualization is correct, then one expects the following to roughly hold true:

First, it should be possible to predict a film's box office revenues per weekend (or week) with good accuracy using only the numbers of screens on which the film is shown *that week*. Although films "age" and become less popular over time, rational exhibitors react to this "aging" effect by reducing screens to keep the income per screen roughly equal to that available from other films.

Second, one should be able to explain a film's *gross* box office revenue fairly well using *total* "screen weeks" the film is exhibited.

Third, deviations from a good fit for these predictions should be related to, and explained by, variables plausibly related to the other constraints under which exhibitors and theaters operate (producer contracts, holidays, cast popularity, and so forth).

Figure 1. Revenues (*R*), Theaters (*N*)

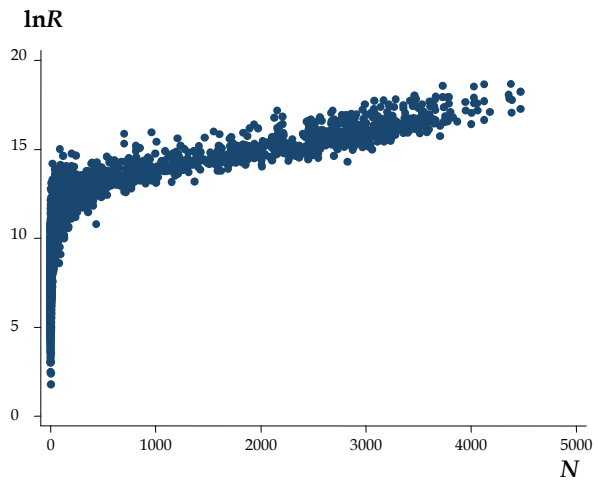

This process has a predictable effect on observed revenues: we should expect that the number of theaters in which a movie is displayed will be a very good predictor of a movie's (weekend) box office revenue. In Figure 1, I show this to be true. In the figure, the (natural log of) weekend gross revenues (ln*R*) for a sample of 634 films shown in the U.S. in the year 2010 are plotted against the number of theaters (*N*) in which the film appeared each weekend.[10] Clearly, theater count is a potent determinant of weekend revenues.

To illustrate the point, let's do some tests of the observations listed above regarding the expectations of the economics of the box office by using some simple regression analysis. First, using data for U.S. films over the 2003-2013 time period, theater counts alone explain 87% of the variability in weekend revenues.[11] Second, the total number of theater weeks (computed as theaters multiplied by weeks played) explains 80% of gross box office revenue for the same set of films. Thus, predications of the theory are supported—box office revenues are determined, in large part, by the number of theaters in which a film is shown. This is the "first-order" prediction of economic theory: when exhibitors behave optimally, they strive to assign films to screens so that the marginal profitability of each exhibited film is equal.

Given the economics of the box office ..., it would seem a natural starting point for a model explaining box office revenues to start with theater counts. Theater counts explain almost all the variation in weekend revenues (87% in the simplest of models). Yet, this is not what Tale of the Long Tail does.

Implausible Implications

Given the economics of the box office just described, and the relationship shown so clearly

by Figure 1, it would seem a natural starting point for a model explaining box office revenues to start with theater counts. Theater counts explain almost all the variation in weekend revenues (87% in the simplest of models). Yet, this is not what *Tale of the Long Tail* does. Instead, the study replaces theater counts with a variable measuring the number of theaters in which a film is shown in its opening weekend. The authors then divide this first-week theater count by the largest first-week theater count in the sample (for each country individually) so that the variable is expressed in terms a country-specific penetration rate.[12]

Tale of the Long Tail's statistical model implies that the revenue of a "blockbuster" will be equal, and in some cases below, the revenues of a "small" film. *** This result is obviously senseless.

This penetration variable plays a key role in the study; it is used to assign "sizes" to the movies. Size is important, since the controversial result from *Tale of the Long Tail* regards the differential impacts of piracy across the "size" of movies (that is, mid-size movies are alleged to actually benefit at the box office from piracy). If a movie appears in most theaters (a penetration closer to 1.0), then it is labeled a "blockbuster"; whereas if it appears in few theaters (a penetration rate closer to 0.0), then it is labeled "small."[13] A mid-size movie, then, lies in the middle between 0.0 and 1.0. Thus, this measure of first-week "penetration" is the study's measure of "size," a variable labeled "*S*" in the study. In effect, the authors intend the variable to be a measure of the exhibitors' estimates of the market potential of the film, made prior to any actual box office feedback.

Tale of the Long Tail's statistical model attempts to explain weekend box office revenues using this *size* variable (and a few other factors not material to the counter-intuitive result, so I ignore them). Recall that the sample is divided into pre- and post-shutdown groups, and the revenue-to-size relationship is estimated for each group. In Figure 2, I illustrate the estimated relationship for the pre-shutdown period (labeled "Pre"). The natural log of revenue is on the vertical axis, and the size variable *S* is on the horizontal.

Figure 2. Revenues and Movie Size (Munich Study)

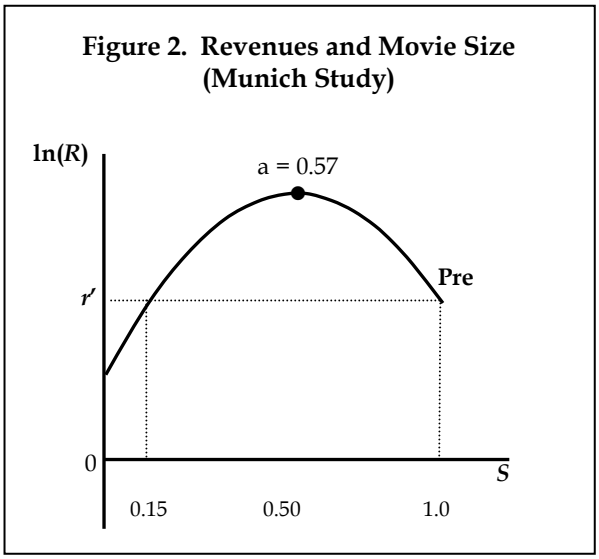

Your alarm bells might be ringing. The "Pre" curve in Figure 2 is concave (a hump), and this hump reaches its maximum for films opening in 57% of theaters.[14] Let's consider the implications of this result. *Tale of the Long Tail's* statistical model implies that the revenue of a "blockbuster" will be equal, and in some cases below, the revenues of a "small" film.[15] This can be seen in the figure that shows equal revenues (at *r'*) for a film shown in 15% of theaters or in 100% of theaters.[16] This result is obviously senseless. Blockbuster films appear in many theaters and earn significantly higher revenues than other films—that's the definition of a "blockbuster." But in *Tale of the Long Tail,* the authors claim that a movie which has limited consumer interests (small size) potentially earns just as much if not much more revenue as does a

movie with great consumer interest appearing in the vast majority of theaters.[17]

Furthermore, the study's model implies that a "mid size" film earns substantially greater revenues than a "blockbuster" film (about four-times as much revenue). This result is simply implausible. A statistical model which claims that a "blockbuster" movie earns revenues smaller than "mid-size" and many "small size" movies point plainly to a bad statistical model and the need to conduct diagnostics in an effort to reformulate. The humped relationship in film "size" is sufficient to dismiss the entire study.

I believe that one primary source of the dubious result in the treatment of size in the model. Even in the U.S., the study's measure of size isn't measuring size at all. (The problems are more profound in the international context.) A few examples demonstrate the problem. Consider the 2012 film *Obama's America*. This movie opened in only *one* theater, earning $31,610. Since this was opening weekend, *Tale of the Long Tail* would take this movie to be a "small film," holding its theater count to one theater over the course of its run. However, in the film's eighth week, it grossed $7.5 million from viewings in 2,017 theaters. It grossed $33 million over its box office run. This is no small film. A similar case is the 2009 film *The Princess and the Frog*, which opened in only *two* theaters so again it would be considered a "small film" by the study's standard. Yet, the film went on to gross $104 million and was shown in 3,475 theaters in a single weekend.

There are many other examples of this sort, like the film *Juno*, which grossed $143 million but opened in only seven theaters. On average, about 70% of movies grossing over $100 million will eventually appear in more theaters than they open in. At the other end of the spectrum, some movies earning less than $1 million opened in thousands of theaters. While perhaps these may be labeled exceptions, such exceptions are caused by a poor modeling choice and, thus, avoidable. Plainly, the S variable is a poor measure of size, and a poor measure of what actually determines weekend box office revenues (mainly, theater or screen counts).

Box office data also belies the statistical model's results. In the U.S., the average film over the 2003-2013 period grossed $1.5 million if in about 700 theaters ($S \approx 0.15$), $2.9 million if in about 2,500 theaters ($S \approx 0.57$), and $24 million if in almost every theater (> 4,250 theaters, $S \approx 1.0$).[18] There's no hump in the relationship. Films opening big do not earn one-quarter of the revenues of films in 57% of theaters, or earn revenues equal to films in 15% of theaters. Something is obviously amuck in the study's model.[19]

...the study's model implies that a "mid size" film earns substantially greater revenues than a "blockbuster" film (about four-times as much revenue). This result is simply implausible. A statistical model which claims that a "blockbuster" movie earns revenues smaller than "mid-size" and many "small size" movies point plainly to a bad statistical model and the need to conduct diagnostics in an effort to reformulate.

It is this mysterious result shown in Figure 2 that drives the counter-intuitive claim regarding the benefits of piracy. In Figure 3, I add to Figure 2 the "Post" revenue-to-size curve and the difference between the "Pre" and "Post" curves (the Δ curve, which is the marginal effect). (See Figure 5 in *Tale of the Long Tail* for comparison.) The Post curve exhibits the same implausible property, with a maximum at 67% of film "size." As shown in the figure, the Post-

curve is flatter than the Pre curve, and its maximum is shifted over.[20]

Figure 3 also illustrates the controversial effect attributed by *Tale of the Long Tail* to the Megaupload shutdown by the curve labeled Δ, which is simply the subtraction of the Pre from the Post curve. The U-shape of Δ implies that the effect of the shutdown is positive for firms with either a small *S* or a very large *S*, whereas firms with a middling *S* see a revenue reduction from the shutdown. This is the study's proverbial "money shot." (If fact, it is the only shot, since the other results are quite weak.) As stated in *Tale of the Long Tail,* "[t]he marginal shutdown effect follows a u-shaped form in movie size that is only significant for medium-sized movies."[21]

Figure 3. Revenue, Size and the Marginal Effect

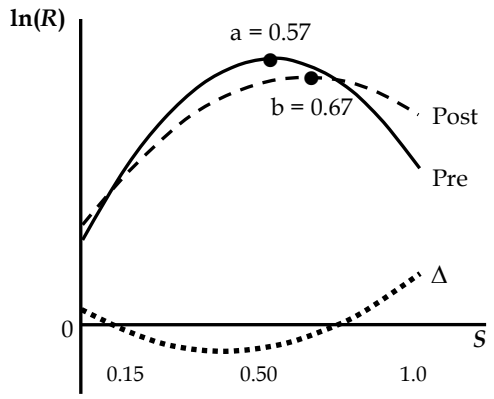

As made obvious by the figure, the counter-intuitive U-shaped relationship is based exclusively on the nonsensical hump-shaped revenue-to-size relationships. As a consequence, the U-shape relationship is as implausible as the humps—a statistical artifact of a mis-specified model.

What I looked for and didn't find in the *Tale of the Long Tail* was some explanation as to what effect piracy should have on the relationship between penetration (or size) and weekend revenues in the chosen model. Is piracy, or its reduction, expected to have a particular influence on the shape of the curve? On its maximum? On small size film versus large films? No explanation or expectation is provided, and this is problematic. (The authors find something, and then make an attribution unique to the result.) For example, given the model's specification (i.e., the quadratic functional form), if the reduction in piracy has the primary effect of increasing the revenues of "blockbuster" films (films with a large *S*), but no effect on small and mid-size films, then the curve will shift in the same manner as shown in the Figure 3 and the "mid-size" benefit argument could appear.[22] Yet, by definition, there is no effect; the mid-size effect is purely a statistical artifact of a mis-specified model. The lack of a theoretical basis for this work prohibits any meaningful interpretation of the results, even if the statistical model wasn't as defective as it is.

It is difficult to accept that exhibitors, acting in an economically rational manner, would ever choose levels of penetration that reduce film revenues. Results of this kind should prompt a careful re-examination of the authors' model. The proper place to start is almost always with the basic economic principles involved.

The bizarre implications of the statistical model regarding the revenue-to-size relationship strongly suggest an underlying problem with the specification. It is difficult to accept that exhibitors, acting in an economically rational manner, would ever choose levels of penetration that reduce film revenues. Results of this kind

should prompt a careful re-examination of the authors' model. The proper place to start is almost always with the basic economic principles involved.

Empirics and Economics

As discussed above, under profit maximization, the revenues of a movie will be related to the number of theaters the movie plays in, not the number of theaters it opens in. Theater count alone explains 87% of the variation in weekend revenues for a sample of movies appearing in theaters over the period 2003-2013, and total theater weeks over the box office live of a film explain 80% of the variability in *gross* revenues. Here, I compare the explanatory power of a model explaining revenues in terms of theater counts to the model used in the study. I limit the sample to the U.S. films over the 2003-2011 period (to avoid the Megaupload shutdown).

...the benefit of the revenue-to-theater/screen model is that it is consistent with basic economic theory; the Tale of the Long Tail's model is not. Revenues are not determined by an opening week penetration rate: they are determined by the number of theaters (or screens) a movie is in.

From the discussion above we know that a simple regression of log-revenues on log theater counts explains 87% of the variation of the dependent variable over this period, and the model is plainly consistent with the pattern shown in Figure 1.[23] In contrast, the *Tale of the Long Tail*'s model (including S, S^2, and the natural log of weeks in the theater) explains only 47% of the variability of log revenues of the U.S. films over the 2003-2011 period. This evidence strongly suggests that the study's model could be better specified and likely suffers from a functional form mis-specification.[24] More importantly, the benefit of the revenue-to-theater/screen model is that it is consistent with basic economic theory; the *Tale of the Long Tail's* model is not. Revenues are not determined by an opening week penetration rate: they are determined by the number of theaters (or screens) a movie is in.[25] Nor are revenues from blockbusters lower than the revenues from small or mid-sized films.

Why Not Theaters?

Why didn't the authors of *Tale of the Long Tail* just use the number of theaters for each weekend instead of the size variable? It is obvious that weekend revenues are determined mostly (about 87%) by the number of the theaters in which the movie plays. Economic theory certainly points to the use of theaters or screens.[26]

Theft is not normally viewed as beneficial to the victim. Tale of the Long Tail's take on piracy, however, suggests otherwise, and this result should be a "red flag" to perform careful diagnostics on the model before touting its findings. Having done so, it seems readily apparent that the counter-intuitive result found in Tale of the Long Tail is not "fact" but "artifact."

The argument provided in *Tale of the Long Tail* is that the "absolute number of screens per country per week [is] potentially endogenous to the shutdown because theater owners can quickly adjust the number of screens as a response to changes in demand."[27] The argument is unexplained and, in my view, makes no sense. Indeed, the fact that theater

owners adjust the number of theaters to changes in demand is the reason why revenues are so successfully modeled as a function of the number of theaters. Such responses are the data generating process because they reflect the profit-maximizing behavior of exhibitors.

If it is true that the number of theaters is "endogenous to the shutdown," then why doesn't the study model that process, including the multiple equations implied by the argument, or at least provide a rigorous demonstration as to why the revenue equation can be estimated alone? If the variable is in fact endogenous and relevant, then the proper method is to include it and apply proper statistical methods to deal with endogeneity. Moreover, replacing the number of theaters with opening week theaters is not a solution at all. The movie industry selects the number of theaters in which to play a movie based on expectations, so the opening number of screens is just as endogenous as the actual number of screens per weekend. Furthermore, the endogeneity argument would be equally valid for the number of weeks a movie is screened as it would for the number of theaters in which the movie appears, yet the authors include that variable in the analysis.[28]

Conclusion

Theft is not normally viewed as beneficial to the victim. *Tale of the Long Tail's* take on piracy, however, suggests otherwise, and this result should be a "red flag" to perform careful diagnostics on the model before touting its findings. Having done so, it seems readily apparent that the counter-intuitive result found in *Tale of the Long Tail* is not "fact" but "artifact." In this PERSPECTIVE, I have pointed out one major defect in the study, but there are many more problems. Indeed, the improbable result is the consequence of the compounding of poor modeling choices, and a neglect of the basic economics of the movie industry. Given the unreasonable implications of the model, *Tale of the Long Tail* adds nothing constructive to the debate—save a little excitement.

NOTES:

[*] **Dr. George S. Ford** ***is the Chief Economist of the Phoenix Center for Advanced Legal and Economic Public Policy Studies. The views expressed in this*** **PERSPECTIVE** ***do not represent the views of the Phoenix Center or its staff.***

[1] C. Peukert, J. Claussen, and T. Kretschmer, *Piracy and Movie Revenues: Evidence from Megaupload: A Tale of the Long Tail?*, Unpublished Working Paper (August 20, 2013) (hereinafter "*Tale of the Long Tail*") (available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2176246).

[2] B. Sisario, *7 Charged as FBI Closes a Top File-Sharing Site*, NEW YORK TIMES (January 19, 2012) (available at: http://www.nytimes.com/2012/01/20/technology/indictment-charges-megaupload-site-with-piracy.html).

[3] In order to estimate the fixed-effects model of the study (that is, to identify the shutdown variable), some movies must be in theaters both before and after the shutdown date (and this crossover could occur across countries).

[4] B. Danaher and M D. Smith, *Gone in 60 Seconds: The Impact of the Megaupload Shutdown on Movie Sales*, Working Paper (March 6, 2013)(available at: http://ssrn.com/abstract=2229349).

[5] *See, e.g.*, C. Dewey, Study: *Piracy Actually Helps Small Films Make Money,* WASHINGTON POST (August 26, 2013) (available at: http://www.washingtonpost.com/blogs/the-switch/wp/2013/08/26/study-piracy-actually-helps-small-films-make-money); D. Cohen, *Closing Piracy Powerhouse Actually Hurt Movie Revenues,* VARIETY (August 27, 2013) (available at: http://variety.com/2013/digital/news/megaupload-closing-piracy-hurt-film-revenues-1200590076); L. Bahr, *Study Suggest Piracy Might Be Good For Indies, Bad For Blockbusters,* ENTERTAINMENT WEEKLY (August 27, 2013) (available at: http://insidemovies.ew.com/2013/08/27/piracy-study).

[6] *See, e.g.*, G. Tullock, *The Welfare Costs of Tariffs, Monopolies, and Theft*, 5 WESTERN ECONOMIC JOURNAL 224-232 (1967).

[7] Computed based on estimates from the "Pre" period at $S = 1$ and $S = 0.57$.

[8] These problems are both statistical and practical. Some issues include the fact that the variable measuring Megaupload's popularity across countries is based on data from only after the shutdown, not before when the popularity of the site matters most to piracy. Also, it appears that many of the movies in the sample were not available on pirating sites during their box office run.

[9] Since a movie can show on multiple screens (e.g., standard and 3D versions), screens is probably the preferred measure.

[10] The sample is from boxofficemojo.com.

[11] I suspect the explanatory power of the model may be even higher with data on screens.

[12] *Tale of the Long Tail*, *supra* n. 1 at p. 12 ("A value of 1 indicates that a movie has the largest all-time first-week audience in a given country."). There are potentially serious problems with this approach, including serving as the source for the odd result discussed here, but a detailed analysis of that issue is not possible without all the data and it's not the main issue of discussion now.

[13] *Id.* at pp. 8, 12.

[14] *Id.* at p. 12 ("We find the expected non-linear relationship with a maximum of 0.57.")

[15] This interpretation of the result is consistent with the description provided in the study. *Tale of the Long Tail*, *id.* at p. 14-5, discussing Figure 5. Another interpretation of Figure 2 is that is represents the relationship between revenues and the penetration rate of opening theaters (S) for a representative movie, not for different movies. In this view, the relationship suggests a movie will earn more money if its opens in fewer rather than more theaters; an unreasonable and counter-factual finding. Thus, there is no solace in this alternative interpretation. Notably, it appears that the relationship between revenues and S arises solely from the cross-country variation in S, which could be the source of the peculiar relationship between revenues and size. With movie fixed effects, the S variable can't be estimated for an individual country, since the movie fixed effect and S are collinear.

[16] The value of the "Pre" curve is 1.07 (ignoring the constants) for S values of 1 and 0.144.

[17] For example, the 2009 film *Avatar* could have increased its revenues by nearly 30% by opening in 25% fewer theaters.

NOTES CONTINUED:

[18] The bands are 600-800 theaters, 2400-2600 theaters, and greater than 4,250 theaters.

[19] The U.S. data suggests the presence of multiple countries in the sample is a potential cause of the problem.

[20] The shape of the Post curve is determined by: $0.22 + (8.56\text{-}2.61)S + (-7.49+3.07)S^2$.

[21] *Tale of the Long Tail, supra* n. 1 at p. 15. Statistically, the only relationship shown to be non-zero is the effect on middle-sized S values. While both small-and-large S films have positive values, the authors are unable to conclude that the effect is not zero due to the wide confidence bands, which are not shown in the figure. Confidence intervals are tighter closer to the mean of the data, so the wider intervals at the extremes of S are expected.

[22] This scenario is easily demonstrated using a Monte Carlo study.

[23] For consistency with *Tale of the Long Tail*, I include calendar week and year fixed effects.

[24] Goodness of fit is not the only determinant of a well-specified model. The revenue-to-theater model is consistent with theory and fits the data; the *Tale of the Long Tail*'s specification does not.

[25] The model also includes a measure of weeks the movie has played, which is correlated with theater count (one is falling over time, the other rising). Still, this decay rate is not uniform across films, and the opening week is not always consistent with the true "size" of a film.

[26] Forcing common slope coefficients across all countries is suspect. Even so, the model could have estimated separate coefficients for countries with theater data and for countries with screen data, creating only two sets of coefficients.

[27] *Tale of the Long Tail, supra* n. 1 at p. 15.

[28] To see this, consider a slight modification of the authors' claim about theaters: "the absolute number of [weeks] [is] potentially endogenous to the shutdown because theater owners can quickly adjust the [number of week] as a response to changes in demand." *Id.*