PERSPECTIVES

PHOENIX CENTER FOR ADVANCED LEGAL & ECONOMIC PUBLIC POLICY STUDIES

The Economic Impact of Expanding Fair Use in Singapore: More Junk Science for Copyright Reform

George S. Ford, PhD*

February 16, 2016

Introduction

As a consequence of digitization and the diffusion of the Internet, nations across the globe are taking a fresh look at their respective intellectual property laws. The ease with which such property can be obtained, used, modified, distributed and simply pirated has fundamentally altered the businesses that depend on intellectual property protections. Modification of existing laws to better fit modern society seems like a sensible strategy, but simply changing the laws that protect intellectual property is not the same as improving them. Better laws require a better understanding of the shortcomings of existing laws and enforcement mechanisms as well as empirical evidence to quantify theoretical parameters and the tradeoffs inherent to legal changes.

To be useful, empirical evidence must focus on the right questions and be conducted with competence, aiming to estimate, to the extent practical, a causal relationship. Sadly, these essential features of good empirical work are often assigned little weight, allowing poorly constructed and executed "empirical analysis"—often cloaked as scholarly work—to influence the policy process. In this PERSPECTIVE, I review an example of a frequently cited yet particularly poor empirical study published in the on-line journal LAWS entitled *A Counterfactual Impact Analysis of Fair Use Policy on Copyright Related Industries in Singapore* on exceptions and limitations to copyright's grant of exclusive rights, known as "fair use."[1]

The empirical analysis by Ghafele and Gibert in their Singapore Study is of stunningly poor quality.

In the study, authors Roya Ghafele, a consultant at the firm Oxfirst (at the time), and Benjamin Gibert, a fellow at the Lisbon Council, claim to estimate the economic effects of amending Singapore's copyright law in 2005 to include United States-style fair use exceptions. They do so for certain industry groups before and after 2005 using growth rates based on a simple linear trend computed using only six annual data points. The authors conclude this "difference-in-differences" methodology "seems to support" that the legal change was responsible for substantial economic gains in Singapore for those industry segments facilitating private copying of copyrighted content (e.g., disk drives and media).[2] Any causal interpretation is explicitly rejected by the authors, thereby (and perhaps inadvertently) rendering the findings in their *Singapore Study* worthless for policymakers.[3]

The empirical analysis by Ghafele and Gibert in their *Singapore Study* is of stunningly poor quality.[4] First, there were many changes in Singapore's copyright law (including an extension of copyright term and the addition of statutory damages), as well as the nation's trademark and patent laws, coincident with the modifications to fair dealing. Toss a bursting tech bubble in 2001 and a global recession in 2008 into the mix and it is clear to see that a simple "before and after" calculation based on almost no data cannot possibly quantify the effect of any single event. The error is fatal.

Second, while Ghafele and Gibert claim to "test this hypothesis using a differences-in-differences methodology," they neither test a hypothesis nor employ a proper difference-in-differences methodology. The control group is improperly used and ineptly chosen. Moreover, by Ghafele and Gibert's own admission, the most basic and necessary assumption of the difference-in-differences method is violated (the parallel paths assumption). For this reason alone, the results are meaningless. Also, Ghafele and Gibert make a glaring and basic error in their treatment of first-differenced data.

… by Ghafele and Gibert's own admission, the most basic and necessary assumption of the difference-in-differences method is violated …

Third, while there are 12 years of data available for 23 distinct industry segments (276 total observations), Ghafele and Gibert use aggregate data to such a level that only six data points (1999-2004) are used to forecast a single future value five years out (2010) based on a linear trend, thereby ignoring the non-linear patterns and tossing out the vast majority of the available data. In so doing, the authors foreclose the ability to use common statistical testing techniques due to very small samples. Ghafele and Gibert likewise ignore the most basic fundamentals of difference-in-differences analysis, statistical forecasting, and time series data.[5] Additionally, the treated and control groups have sizeable differences in their mean values, so it is essentially impossible for the control group to serve any meaningful purpose in the calculations used.

In an effort to remedy some of the many (but not all) shortcomings of Ghafele and Gibert's *Singapore Study*, I apply the standard regression-based difference-in-differences test so that statistical testing can be performed within a more proper difference-in-differences regression model.[6] Using the same data as the *Singapore Study* (which is not good data) but using it along with the standard difference-in-differences regression model, I am unable to reject the hypothesis that the changes in Singapore's laws had "no effect" on the industry groups.

Under Ghafele and Gibert's logic, … the massive rise in the production of copying technologies would suggest that Singapore's new copyright law was taken to be a license to pirate copyrighted content with impunity.

Finally, ignoring Ghafele and Gibert's defective statistical analysis for the moment, one could easily argue that the results of their *Singapore Study* show why relaxed fair use policies are a bad idea. In effect, the results of Ghafele and Gibert's *Singapore Study* (if they were legitimately estimated) suggest that Singapore's expanded fair use policies coincided with a massive expansion in the production of technologies that permit widespread copying of intellectual property but had no effect on the creation of new, transformative works. Fair use is *not* about the wholesale copying of intellectual property; fair use is mostly about the use of very *limited* pieces of content to create something new (i.e., transformative), or perhaps limited use in educational or research settings.[7] Under Ghafele and Gibert's logic, however, the massive rise in the production of copying technologies would suggest that Singapore's new copyright law was taken to be a license to pirate copyrighted content with impunity. Notably, less than a decade after the new fair use policy was codified, Singapore's Ministry of Law initiated a consultation regarding widespread digital piracy in that country; the law was subsequently amended to address these concerns.[8]

A Summary of the *Singapore Study*

In testing the efficacy of a new medication, researchers gather a sample of persons, split them randomly into two groups, and then give half the group the "treatment" of a real drug and the other half a placebo (or nothing at all). Changes in health outcomes—both those intended and any side effects—between the treated and control groups are then measured to determine whether the treatment is effective and safe. The approach is scientifically valid because of the random assignment of the groups, implying the two groups are, for all practical purposes, identical.

Randomization experiments are typically not possible in legal and economic environments; treatments are not randomly assigned. On occasion, however, some quasi-experimental situations do arise in the course of economic activity.[9] Say, for example, a regulation is imposed in one set of cities but not another.[10] If a sample of otherwise identical cities can be obtained to serve as a control group for the newly regulated cities, then the potentially "causal" effect of the regulatory change can be quantified by looking for changes in economic outcomes over time between the two types of cities—the treated group and the control group.

As an example, say there are two sets of cities, both having a 10% infant mortality rate.[11] In one set of cities—the treatment group—a policy in implemented (say, privatization) to improve the quality of water filtration. No change in policy occurs in the other group, which serves as the control group. Five years after the policy change, mortality is measured and the treated cities are found to have a mortality rate of only 5%, but the mortality rate of the control group is 9%. The improvement in mortality over the five-years for the treated group is 0.05 (= 0.10 – 0.05), but for the control group is only 0.01 (= 0.10 – 0.09). The difference between the outcomes of the two groups of cities is 0.04 (= 0.05 – 0.01). This difference—often referred to as the Average Treatment Effect ("ATE")—may be causally assigned to the change in water treatment policy if the two sets of cities are identical in all respects save the treatment.

This procedure is called a difference-in-differences approach, and the treatment effect is computed using

$$\delta = (Y_1^T - Y_0^T) - (Y_1^C - Y_0^C) \qquad (1)$$

where δ is the difference-in-differences estimator, the Y^T are the outcomes of the treated group and the Y^C the control group.[12] The subscripts 0 and 1 indicate, respectively, the outcomes before and after the treatment. Equation (1) shows clearly why the method is referred to as a difference-in-difference estimator; it is literally the difference between two differences.

[By] tossing out the vast majority of the available data … [Ghafele and Gibert] foreclose the ability to use common statistical testing techniques due to very small samples.

The key component of the difference-in-difference approach is the quality of the control group. One critical assumption regarding the control group is the *parallel paths assumption*. This assumption holds that the pattern of outcomes for the control group after the treatment is an unbiased estimate of what would have happened to the treated group had the treated group not received the treatment. In these types of statistical models, the point is not simply to compare the average outcomes between a treated group and a control group; rather, the control group is intended to be a stand-in for the treated group had the treated group not received the treatment. Ideally, the only difference between the control and treated groups is that one receives the treatment and the other does not. Otherwise, the "two groups should be similar in every respect,"[13] so that the control group serves as a valid *counterfactual*.[14] Say, for example, the control cities from the water-treatment example were all from a country that experienced a famine following the change in policy for the treated group, altering child mortality. Or, cities in the treated group were all in developing countries while the control group was made up of cities from advanced economies. If so, then the control group is invalid and the measured treatment effect of privatization is biased (i.e., inaccurate).

This parallel paths assumption cannot be formally tested. In addition to the apparent comparability of the treated and control groups across a range of characteristics, the commonality of pre-treatment trends and any odd activity post-treatment are often used as a check on the assumption. If pre-treatment trends are not equal across the two groups, then there is no reason to expect the control group to serve its purpose. If the control group's outcomes change substantially after the treatment, then there is good reason to believe something else was influencing the outcomes. Of course, if the treated and control groups are very different along any other observed dimensions, then the validity of the parallel paths assumption is dubious. Plainly, the control group must be very carefully selected and the statistical methods must be properly applied and suited to the nature of the data.[15] A poorly-selected control group renders the analysis incapable of measuring anything meaningful.

Ghafele and Gibert believe that such a quasi-experiment occurred in Singapore in 2005, when the country altered the language of the pre-existing multi-factor test of fair use in §35 of the Singapore Copyright Act in a manner that made the multi-factor test more broadly applicable. Ghafele and Gibert lay out their empirical approach as follows,

> … we posit that flexible fair use in copyright law has two additional effects in the economy beyond those posited by traditional fair use analyses. Flexible fair use exemptions may: (1) increase the growth rate of private copying technology industries; and (2) increase the growth rate of copyright markets. We test this hypothesis using a differences-in-differences methodology that is applied to the 2005 fair use amendments to the Singapore Copyright Act and test its implications on private copying technology and copyright sectors in Singapore.[16]

The time series data used in the report spans 1999 to 2010, and the *Singapore Study* marks the year 2005 as the treatment year. Value-added as a percentage of Gross Domestic Product ("GDP") is the outcome of interest. The three industry groups represented include: (1) Private Copying Technology Industry Group; (2) Copyright Industry Group; and (3) Control Industry Group. The *Private Copying* group includes sectors of the economy that produce and sell technologies—e.g., disk drives and media—that may benefit from fair use in that it permits more copying of copyrighted content without permission. The *Copyright* group includes industry segments—e.g., recordings, books, magazines, and movies—that may be harmed by fair use since the fair use exception may impact sales. Finally, a single *Control* group is constructed from industry segments—e.g., office equipment—believed to be unaffected by fair use policies (though these industries received the treatment).

A number of problems with the data and analysis in the *Singapore Study* are covered by Barker (2013) and Barker and Png (2012).[17] For example, about 95% of Singapore's electronics production is exported, leaving only a small share of production for domestic copying. Since Singapore's laws apply only to Singapore, the high export rate of the products analyzed in the *Singapore Study* makes the data of questionable value. Barker (2013) and Barker and Png (2012) discuss many other problems with the data as well as with the analysis. I will focus on some other problems not covered in these earlier works and will take the data as a given; the validity of the analysis is evaluated from that point on.

The empirical approach used by Ghafele and Gibert in their *Singapore Study* is as follows. Using six years of data prior to 2005, the following linear regression is estimated:

$$y_t = \beta_0 + \beta_1 t + e_t, \qquad (2)$$

where y_t is the value-added figure and t is a time indicator ($t = 1, 2, \ldots, 6$). Using the estimated β_0 and β_1 coefficients from Equation (2), the authors compute the value of y in the final period (y_{12}, six years later in year 2010, ignoring the data from 2005 through 2009), and this prediction serves as the pre-treatment value of y (or Y_0^T from Eq. 1). For example, data from the *Private Copying* group produces the "prediction" equation:

$$y_t = 0.06338 - 0.00388t. \qquad (3)$$

Inserting the value of 12 for t, the prediction for 2010 (five years into the future) is 0.0168. The actual value for 2010, however, was 0.0502 (or Y_1^T), so the authors conclude the "difference" of the change in fair use law for the *Private Copying* group was 0.0334, or 3.34%.

Ghafele and Gibert use the linear prediction for 2010 as the Y_0^T (from Eq. 1) in the difference-in-differences calculation (the *Private Copying* outcome absent the treatment). In fact, this difference between the forecast and the actual value is better described as a forecast error based on out-of-sample information (2010 in not in the estimation sample). Forecasters often test the validity of forecast models by seeing how well the models predict known future values that are purposely left out of the sample so that such calculations can be made.[18] The huge miss of Equation (3) suggests, most plausibly, that the linear model is a very poor one for this data, especially for the *Private Copying* group.

These same computations are made for the *Control* and *Copyright* groups. For the *Control* group, the difference is very small [0.0044 – 0.0043 = 0.0001]. The difference-in-difference estimator for the *Private Copying* group is, then, 0.0333 [= 0.0334 - 0.0001]. Ghafele and Gibert conclude that the addition of fair use to Singapore's copyright law was "correlated with a 3.33% increase in value–added (as a % of GDP) for private copying technology industries."[19] The calculations are summarized in Table 1.

Table 1. Estimated Effects from the *Singapore Study*
(Value added as Percent of GDP)

Industry Group	2010 Actual	2010 Predicted	Impact	Impact (Adj. for Control)
Private Copying	5.02%	1.68%	3.34%	3.33%
Copyright	0.55%	0.78%	-0.23%	-0.25%
Control	0.44%	0.43%	0.01%	…

Ghafele and Gibert conduct the same calculations for the *Copyright* group, using the same *Control* group, and their *Singapore Study* reports "no significant change in growth rates for the copyright group in the pre- and post-intervention time-periods in terms of real economic growth."[20] Based on these findings, the conclusion of their *Singapore Study* is that expanded fair use is good for the *Private Copying* group and does not harm rights holders. An alternative interpretation of the results is that the analysis merely reveals that the linear forecast is an especially poor forecast model for the *Private Copying* group relative to *Copyright* and *Control* groups.

While evidence on the impacts of changes to fair use policies is desirable, the question remains as to whether Ghafele and Gibert's empirical analysis is properly done so as to render any meaningful evidence. For the reasons discussed already and detailed more below (among many others), it was not. The analysis in the *Singapore Study* is poorly conceived and incompetently performed; it's junk science. Attentive readers may notice immediately a couple of problems revealed in Table 1. First, the mean of the *Private Copying* group is 5.02% and the *Control* group is 0.55%. In order to offset the change in the *Private Copying* group, the *Control* group would have to change many multiplies of its mean. Also, the same control group is used for two very different industries. Unless the *Copyright* and *Private Copying* groups respond identically to economic stimulus, then a single control group cannot serve as a valid counterfactual for both. I discuss these obvious problems and some others below.

Multiple Treatments

Returning to the drug testing example for the moment, imagine a drug company wishing to test the efficacy of a new treatment. Once more, following standard protocol, they gather a group of test subjects and then randomly divide them into treatment and control groups. So far, so good, but now say that the company administers four different drugs to the treatment group—the drug of interest and three others. After the allotted time, the outcomes are measured and it is found that the treated group has some improvement in the condition of interest. This evidence is presented to the Federal Drug Administration ("FDA") for approval of the treatment. Of course, the FDA would instantly reject the study. The observed outcomes and side effects are the result of taking four medications, and it is impossible to determine which drug is responsible for what outcomes.

The problem of multiple treatments is obvious enough, though apparently undetected by the authors of the *Singapore Study*. There were many significant changes in the 2005 revision of Singapore's copyright law, including an extension of the copyright term from 50 to 70 years after death (§28) and the introduction of statutory damages (§119), both thought to be a strengthening of copyright.[21] In addition, Singapore's patent and trademark laws were updated at this same time, and at least some of the industry groups (e.g., electronics) analyzed in the *Singapore Study* depend on these forms of intellectual property. On top of these changes in intellectual property law, there was a bursting technology bubble in 2001 and a global recession in 2008, neither of which are addressed in the *Study's* calculations.[22] So, while fair use policy did change in 2005, it was only one of many changes that occurred at that time and over the period analyzed in the *Singapore Study*. That is, multiple drugs were administered. There is simply no way to determine the impact of the change in fair use policy independent of all these other material changes.

[W]hile fair use policy did change in 2005, it was only one of many changes that occurred at that time and over the period analyzed in the Singapore Study. *** There is simply no way to determine the impact of the change in fair use policy independent of all these other material changes.

Scale Effects and Sample Sizes

Though the *Singapore Study* considers 23 different industry classifications, it aggregates these industries into three groups. No explanation is provided as to why tossing out data in this way is helpful, and there are many reasons to believe it is harmful. For example, aggregating the data reduces the sample sizes to a puny count (i.e., six observations are used to forecast a single observation, so only seven observations are employed per group). If using all the data, these seven observations could be increased to 96 observations for the *Private Copying* group.

Another problem with aggregation is that it combines large industry segments with small ones. In effect, many of the segments are all but lost in aggregation with their relatively large companions. For example, two of the industry segments in the *Private Copying* group account for 70% of the total. Four segments account for less than 4% of the total, so the variation in these smaller segments is essentially washed out in the aggregation. Yet, if the impact of fair use is as Ghafele and Gibert presume, then these impacts should be detectable in all (or most) of the individual industry segments (i.e., each SSIC). Aggregation prohibits such measurements, masking the effects on smaller components with the effects on bigger ones.

Another important scale effect, discussed briefly above, renders the comparisons made in the *Singapore Study* utterly meaningless. In computing the treatment effect, Ghafele and Gibert subtract the changes in the *Control* group from that of the *Private Copying* group using the formula from Equation (1). But, notice from Table 1 that the mean value of the *Private Copying* group is 0.05 while the mean value of the *Control* group is only 0.005—a ten-fold difference. By Equation (1), in order to offset the 0.0334 change in the *Private Copying* group, the *Control* group's outcome would also need to change by 0.0334. Given its relatively small mean, however, the Control group's outcome must increase by a whopping 668% (= 0.0334/0.005 = 6.68) to offset the 67% increase in the *Private Copying* group's outcome (= 0.0334/0.05 = 0.67). Thus, it is practically impossible for the *Control* group to offset a change in the *Private Copying* group.

If this issue is not clear yet, consider an analogy. Say a pharmaceutical company wishes to test a new weight-loss drug for dogs. For the treated group, the company gathers a sample of fifty Great Danes with an average weight of 150 pounds. As a control, the company chooses a group of fifty Chihuahuas with an average weight of 5 pounds. Three months after the Great Danes are treated (the Chihuahuas get a placebo), we find that the larger breed has lost 20 pounds. The Chihuahuas, however, have lost only 0.2 pounds, on average. Looking at these results, the company concludes the drug is effective, reducing weight by 19.8 pounds (or 13%). Plainly, this experiment is pointless. Yet, this error is exactly the one made by Ghafele and Gibert in their *Singapore Study*.

This type of scale effect is well covered in the literature on difference-in-difference estimators. Data transformations and econometric methods are available to correct for such scale differences, if necessary. Ideally, the means between the control and treated groups would not be very different in scale as such disparities lead to questions about the legitimacy of the control group to produce a valid counterfactual.

The parallel trends assumption is a necessary condition for the validity of the difference-in-difference estimator, but Ghafele and Gibert reject the validity of this assumption yet are apparently unaware this rejection renders their entire analysis meaningless.

Poor Choice of Control Group

Ghafele and Gibert also make a completely inappropriate choice of control group. As stated in the study, the control group is selected because these industry segments are purportedly "non-beneficiaries of the fair use policy intervention."[23] That is not to say that the control group did not receive the treatment—copyright law in Singapore applies to all industries in Singapore. The control group is treated, it is just that Ghafele and Gibert *believe* that these industry segments are unaffected because they did not benefit from the policy intervention despite the treatment. A valid control, however, is given a placebo, not the treatment. For this group to be a valid control, we must believe that those industry segments in the control group are identical to the treated group in all respects but one –the response to fair use. Such a scenario is implausible.[24]

Figure 1. Pre-Treatment Trends

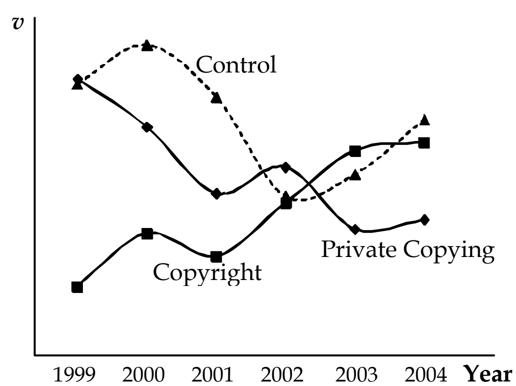

Even when researchers go to great lengths to find a proper control (say, by carefully choosing a set of cities, schools, persons, and so forth), it is common practice to evaluate the trends in the data prior to the treatment to see if they are the same. Visual inspection of pre-treatment trends is made possible by Figure 1.[25] From the figure, it is apparent that the *Control* group has a very different trend than either of the treated groups. Thus, based on visual inspection, the parallel paths assumption is not satisfied for either the *Private Copying* or *Copyright* groups.

Figure 1 also shows that the *Control* group could not possibly serve as a valid control for both the *Private Copying* and *Copyright* groups. The pre-treatment trends of the *Private Copying* and *Copyright* groups are about as inconsistent as they could be, one rising and one falling over much of the pre-treatment period. The *Control*, in contrast, is rising, then falling, then rising, over the pre-treatment period. It is fairly obvious that the *Control* group is a poor one for either of the "treatment" groups.[26]

In fact, Ghafele and Gibert acknowledge that the parallel paths assumption is violated, stating the "[a]nalysis of growth rates shows that the parallelism assumption does not hold. [T]he private copying group was shrinking at a faster rate before 2005 than the control group."[27] *For a difference-in-difference analysis, this statement demands a full stop; nothing about the analysis has any validity past this point.* The parallel trends assumption is a *necessary* condition for the validity of the difference-in-difference estimator, but Ghafele and Gibert reject the validity of this assumption yet are apparently unaware this rejection renders their entire analysis meaningless.

Researchers also look at post-treatment behaviors to see if there is something amiss. In Figure 2 the data for all years on the three groups is illustrated. Over the entire period, the *Copyright* group has a generally positive trend. The trend for the *Private Copying* group initially is falling but turns positive prior to 2005, a pre-treatment directional change that revealing the treatment effect is biased. For both the *Private Copying* and *Copyright* groups, the upward trends in the pre-treatment period are generally continued on into the treatment period. If a change in activity is detected prior to the treatment and carried over into the treatment period, then we have fairly strong evidence that the policy change did not cause the change in activity.[28]

Figure 2. Full Data Trends

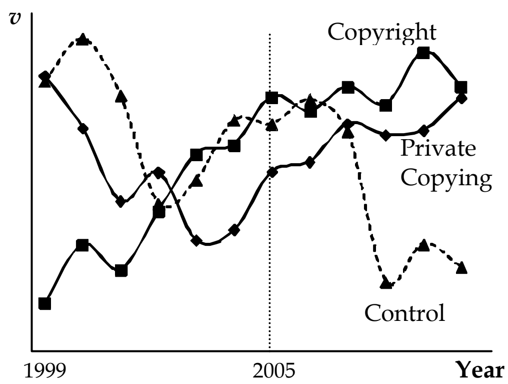

There appears to be no visible change after 2005 for either group, other than perhaps a slowed growth for the *Copyright* group. In contrast, three years after the treatment period begins value added for the *Control* plummets. Since only the 2010 data are used from the post-treatment period, this apparent shock to the control group greatly influences the reported outcomes in the *Singapore Study*. Overall, these data suggest neither a strong treatment effect nor a meaningful control group.

Despite their own rejection of the parallel paths assumption, Ghafele and Gibert repeatedly qualify their findings with the statement: "assuming parallel growth rates between beneficiaries and non-beneficiaries…."[29] *They do so despite their own explicit rejection of the of parallel paths assumption*. It seems quite clear that Ghafele and Gibert either do not fully comprehend the empirical techniques they are using or are purposefully misleading the reader and policymakers as to the strength of their findings.

A Review of the *Singapore Study's* Empirics

Close (if not casual) inspection and replication of the *Singapore Study's* empirical analysis reveals a number of severe defects. While my review is not exhaustive, the defects already mentioned are fatal, so additional errors, of which there are many, need not be covered in any detail. However, something might be learned by trying to remedy some of the defects. Here, I use the raw data from the *Singapore Study* to estimate a proper difference-in-differences regression model so that statistical tests can be performed. In doing so, I can determine whether there is any meaningful and quantifiable difference like that reported in the *Singapore Study* (ignoring the other problems) within the general confines of the study's empirical approach.

The standard format for the difference-in-difference regression is

$$y_{it} = \delta D_{it} + \beta X_{it} + \lambda_i + \mu_i + \varepsilon_{it}\,, \qquad (4)$$

where y_{it} is the outcome for observation i at time t, D_{it} is a dummy variable that equals 1 if the observation is treated in year t (0 otherwise), X_{it} is a vector of control variables that vary by observation and time, μ_i is fixed effect for each observation i, λ_t is a time effect common to all observations in time t, and ε_{it} is the econometric disturbance term that is assumed to be distributed independently of all μ and λ.[30] The δ and β_1 are estimated parameters. This model is a two-way fixed effects model (that is, there are dummy variables for both i and t). Assuming the assumptions of the difference-in-differences method are satisfied, the coefficient δ is the difference-in-differences estimate.

In the *Singapore Study*, all the data is national data from Singapore and the analysis includes no other exogenous factors, so the X_{it} falls out of Equation (4).[31] The data is measured as "value added" of various sectors in the economy, which I label v_{it}. Thus, Equation (4) can be rewritten as,

$$v_{it} = \delta D_{it} + \lambda_i + \mu_i + \varepsilon_{it}\,, \qquad (5)$$

where i is a SSIC-level industry group and D_{it} is a dummy variable equal to 1 for the *Private Copying* or *Copyright* groups beginning in year 2005 (0 otherwise). The regression is a two-way (SSIC, time) fixed effects regression with a dummy variable that equals 1 for the treatment groups during the treatment period (2005 to

2010).[32] Standard errors are clustered on the SSICs, as is recommended in the literature.[33]

There are 23 SSICs in the sample: (a) eight for *Private Copying*; (b) ten for *Copyright*; and (c) five for the *Control* group. With twelve years of data (with 2 observations missing for 1999), the estimation sample includes 274 observations. The R^2 of the regression is 0.05 and the F-statistic is 4.34 (prob < 0.01). The estimated treatment effect (δ) for the *Private Copying* group is 0.00074 with a t-statistic of 0.30 (prob = 0.77); the effect is not statistically different from zero. The null hypothesis of "no change" in the *Private Copying* group's outcomes before and after 2005 cannot be rejected.

For the *Copyright* group, the estimated treatment effect (δ) is equal 0.00034 with a t-statistic of 0.81; the effect is not statistically different from zero (prob = 0.43). The null hypothesis of "no change" in the *Private Copying* group's outcomes before and after 2005 cannot be rejected.

Equation (5) is also estimated replacing the dependent variable with its natural log transformation. The results are not much changed, though the regression is not statistically significant (F = 1.16, prob = 0.37). The ATE for the *Private Copying* group is 0.174 with a t-statistic of 0.48 (prob = 0.637); for the Copyright group, the ATE is 0.171 with a t-statistic of 0.66 (prob = 0.515).

The null hypotheses of the test is that the change in Singapore's copyright law—including fair use policies, term, statutory damages, and so forth—had no effect on the economic outcomes after 2005 for the *Private Copying* and *Copyright* groups relative to the *Control* group. The hypothesis for neither the *Private Copying* nor the *Copyright* group is rejected. Thus, the host of modifications to Singapore's copyright law (and trademark and patent laws), whatever those changes may be, *had no effect* on the economic outcomes that Ghafele and Gibert deemed of interest, at least when using the data employed in that study and more properly applying the difference-in-differences methodology.

[T]he host of modifications to Singapore's copyright law, whatever those changes may be, had no effect on the economic outcomes that Ghafele and Gibert deemed of interest [when] more properly applying the difference-in-differences methodology.

Particularly Puzzling Parts

In Appendix Tables A4 and A5 of the *Singapore Study*, Ghafele and Gibert present econometric results that play no role in their analysis. These reported regression results in Table A4, and the data in Table A5, do not serve as the basis for the estimated impacts. Nor are the results discussed in the paper. What purpose the tables serve is unclear.

What is most telling about this analysis is the information contained in Table A5 ("Data Used in the Regression Analysis"). The study has data for the period 1999 through 2010. In the regression analysis from Table A4, the authors estimate a first-difference model, which simply means that the data used in the estimation is calculated by subtracting each year's data from the prior year's data ($y_t - y_{t-1}$). Without data from 1998, this produces a missing value for year 1999, which reduces the sample size from 12 years to 11 years.

Looking at Table A5, however, rather than excluding 1999 from the data as a missing value, the authors have inserted "0" as a data point for that year. The regression results from Table A4 indicate there are 12 observations in the regression, which implies that the "0" observation for 1999 was, in fact, included in the estimation. (I was able to replicate the results to confirm this error). While the results from this regression analysis serve no apparent purpose in the *Singapore Study*, this error in the handling of the data suggests the authors are unskilled in statistical analysis.

Also, in multiple instances the *Singapore Study* stresses that its analysis has no causal interpretation, noting that the analysis "can only establish a correlation between a policy and a given outcome. Our findings must be read in this light."[34] While it is true that causal effects are very difficult to estimate, the advantage of a properly implemented difference-in-differences analysis is that it lends much credibility to the causal nature of the estimates. Causality is critical for policy work: only if a relationship is *causal* can we conclude that implementing the same policy elsewhere will produce the same results. As observed by Angrist and Pischke in their book *Mostly Harmless Econometrics*:

> Although purely descriptive research has an important role to play, … [a] causal relationship is useful for making predictions about the consequences of changing circumstances or policies; it tells us what would happen in alternative (or "counterfactual) worlds.

Correlation (a descriptive statistic) may exist between entirely unrelated things. For example, *U.S. Spending on Science, Space and Technology* has a near perfect correlation (r = 0.998) with *Suicides by Hanging, Strangulation and Suffocation*, but we would never reduce technology spending to curb suicides because we know the correlation is spurious.[35] Ghafele and Gibert's aggressive rejection of a causal interpretation of their results makes the *Singapore Study* irrelevant for reforming laws.

What if the Results Were Right?

Assume, for the moment, that the results reported in the *Singapore Study* are legitimate in the sense that Singapore's fair use reform led to massive growth in manufacturing of copying technologies yet had no effect on copyright industries. What do the findings of the *Singapore Study* imply about greater flexibility of fair use policies?

Ghafele and Gibert's aggressive rejection of a causal interpretation of their results makes the Singapore Study irrelevant for reforming laws.

Fair use is *not* about the wholesale copying of intellectual property; instead, fair use grants rights to the *limited use* of content to create something new (i.e., transformative) out of existing content, or perhaps limited use in educational settings without having to compensate the rights' holder. Singapore's laws are no different in that regard. For example, one of the factors applied in Singapore's fair dealing law is the "amount and substantiality of the part copied taken in relation to the whole work" and the "effect of the dealing upon the potential market for, or value of, the work."[36] When using copyrighted content for research or study, Singapore's fair dealing law limits the copying to 10% of the published work. Fair use is a somewhat *limited* exemption.

What a well-designed fair use policy *should do* is lead to the creation of new, transformative works, thereby growing the copyright sector. Fair use policies *should not*, if properly designed, lead to a massive increase in the sale of pirating technology. Yet, according to Ghafele and Gibert, Singapore's new law led to exactly that perverse outcome; their results imply that subsequent to the increase in the flexibility of fair use in Singapore, while consumers rushed out and purchased vast quantities of technology used to copy intellectual property (the *Private Copying* group) such increased flexibility had no effect on the creation of new works. I can think of no plausible explanation for these results for a *legitimate* application of fair use.

Thus, there is an alternative and more troubling take away from the results of the *Singapore Study*: consumers took Singapore's new law as a license to copy and distribute copyrighted content (i.e., piracy) without consequence. One could read Ghafele and Gibert's study as evidence that expanded fair use laws *encourages* digital piracy of intellectual property without expanding the creation of new, transformative works. The evidence presented in the study would confirm such a hypothesis (though the results are meaningless and should be ignored). Also, Singapore amended its copyright law less than a decade subsequent to modifying its fair use policy to address rampant digital piracy. Thus, if the *Singapore Study* tells us anything, then it is that Singapore's 2005 fair use amendment was bad policy.

While evidence on fair use policies is welcome and critical to informed policy reform, Ghafele and Gibert's empirical analysis is so poorly done that it fails to shed any light on copyright laws. Governments reviewing their copyright laws should dismiss the Singapore Study as junk science.

Conclusion

The 2014 *Singapore Study* by Ghafele and Gibert on the economic impacts of Singapore's change in its fair use policies claims to show a large effect on industries that manufacture goods useful for private copying of copyrighted works and no effect on the copyright industries. While evidence on fair use policies is welcome and critical to informed policy reform, Ghafele and Gibert's empirical analysis is so poorly done that it fails to shed any light on copyright laws. Governments reviewing their copyright laws should dismiss the *Singapore Study* as junk science.

The defects in the analysis are numerous, but I have focused on a few fatal defects.

First, there were many changes to Singapore's copyright law in 2005, not simply its fair use policies. In the presence of multiple coincident treatments, it is impossible to assign differences in outcomes to any one particular change.

Second, the chosen treatment of the data results in puny samples, foreclosing standard statistical testing (no effort is made in the study to test hypotheses). Data aggregation also leads to two types of scale effects, one of which makes the comparison of treated and control groups meaningless.

Third, a fundamental requirement of the difference-in-difference estimator is the "parallel paths" assumption, since if the assumption fails, any measured treatment effect is biased. The authors acknowledge that this assumption fails, but they fail to grasp the significance of this finding.

Fourth, using the *Singapore Study*'s data and applying the standard regression model for the difference-in-differences estimator, I am unable to reject the hypothesis that the changes in the copyright law had no effect on economic outcomes.

Finally, the most plausible interpretation of Ghafele and Gibert's findings, ignoring the defects in the data and the analysis which render invalid the study's reported results, is that the expanded fair use policy was incorrectly interpreted by consumers as a license to pirate and distribute intellectual property without consequence. Less than a decade after the new fair use policy was implemented, Singapore amended its copyright law to address widespread digital piracy.

Protections offered by intellectual property laws support substantial economic activity and employment.[37] Such laws do much to encourage the creative activity of humankind. But laws are man-made, imperfect, and in need of an occasional update. The evidence used to inform legal reform must be relevant and carefully conducted. Since policymakers are rarely skilled in statistical analysis, there is a trust that must be built between the research and the policymaking communities. Careless and unskilled studies, like the *Singapore Study*, threaten that trust. One of the reasons empirical work plays such a small role in policymaking is the prevalence of this type of junk science. The formulation of public policy deserves better.

NOTES:

* **Dr. George Ford** ***is 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. Dr. Ford has testified as an expert before the copyright royalty authorities in the U.S. and in Canada. Dr. Ford may be contacted at*** ford@phoenix-center.org.

1 R. Ghafele and B. Gibert, *A Counterfactual Impact Analysis of Fair Use Policy on Copyright Related Industries in Singapore*, 3 LAWS 327-352 (2014) (available at: http://www.mdpi.com/2075-471X/3/2/327) (hereinafter "*Singapore Study*"). This published study was a revised draft of an earlier 2012 study by the same authors entitled *The Economic Value of Fair Use in Copyright Law: Counterfactual Impact Analysis of Fair Use Policy On Private Copying Technology and Copyright Markets in Singapore,* OXFIRST LIMITED (October 2012) (available at: http://infojustice.org/download/copyright-flexibilities/articles/Roya%20Ghafele%20and%20Benjamin%20Gibert%20-%20The%20Economic%20Value%20of%20Fair%20Use%20in%20Copyright%20Law.pdf).

2 *Singapore Study*, *id*. at p. 327.

3 *Id*. ("We caution that a counterfactual analysis cannot be used to attribute a causal relationship. It can only establish a correlation between a policy and a given outcome. Our findings must be read in this light (at p. 327)"; "Note: correlation is not causality! (at. ft. 2)"). In fact, a properly-implemented difference-in-differences (counterfactual) analysis is explicitly intended to lend a causal interpretation to statistical results. *See, e.g.*, G.W. Imbens and J.M. Wooldridge, *Recent Developments in the Econometrics of Program Evaluation*, 47 JOURNAL OF ECONOMIC LITERATURE 5-86 (2009).

4 Mr. Gibert is apparently building a well-earned reputation for low-quality empirical work. G. S. Ford, *The Lisbon Council's 2015 Intellectual Property and Economic Growth Index: A Showcase of Methodological Blunder*, PHOENIX CENTER POLICY PERSPECTIVE No. 15-03 (June 29, 2015) (available at: http://www.phoenix-center.org/perspectives/Perspective15-03Final.pdf).

5 Note that I am ignoring here a huge number of problems with the *Singapore Study*'s treatment of time series data, which can only be described as profoundly incompetent. From a statistical perspective, using a linear trend to forecast five-years into the future using only six data points is simply absurd. For example, the large miss of the forecast for the *Private Copying* group would be used by forecasters as a signal of a bad prediction model. W. Enders, APPLIED ECONOMETRIC TIME SERIES (2004) at p. 79 ("models with poor out-of-sample forecasts should be eliminated"). At the risk of lending any credibility to the *Singapore Study*'s approach, *see also* P. Saffro, *Six Rules for Effective Forecasting*, HARVARD BUSINESS REVIEW (July-August 2007) (available at: https://hbr.org/2007/07/six-rules-for-effective-forecasting) ("So when you look back for parallels, always look back at least twice as far as you are looking forward.").

6 *See, e.g.*, B.D. Meyer, *Natural and Quasi-Experiments in Economics*, 13 JOURNAL OF BUSINESS & ECONOMIC STATISTICS 151-161 (1995); J.D. Angrist and J.S. Pischke, MOSTLY HARMLESS ECONOMETRICS: AN EMPIRICIST'S COMPANION (2008); J.D. Angrist and A.B. Krueger, *Empirical Strategies in Labor Economics*, in HANDBOOK OF LABOR ECONOMICS (Volume 3A)(1999)(O. Ashenfelter and D. Card, eds.) at Ch. 23.

7 *Asking Permission and Using Copyright*, Intellectual Office of Singapore (Viewed on January 22, 2016) (available at: http://www.ipos.gov.sg/AboutIP/TypesofIPWhatisIntellectualProperty/Whatiscopyright/Askingpermissionandusingcopyright.aspx). *See also*, P.N. Leval, *Toward a Fair Use Standard*, 103 HARVARD LAW REVIEW 1105-1136 (1990) (available at: http://docs.law.gwu.edu/facweb/claw/levalfrustd.htm); P. Aufderheide and P. Jaszi, RECLAIMING FAIR USE (2011) at Ch. 6.; *What is Fair Use?*, Stanford University Library (Viewed January 22, 2016) ("In its most general sense, a fair use is any copying of copyrighted material done for a limited and "transformative" purpose, such as to comment upon, criticize, or parody a copyrighted work") (available at: http://fairuse.stanford.edu/overview/fair-use/what-is-fair-use).

8 C. Mann, *Singapore Proposes Copyright Amendments*, ADVANCED TELEVISION (April 9, 2014) (available at: http://advanced-television.com/2014/04/09/singapore-proposes-copyright-amendments/); *Singapore Passes Copyright Amendment*, TRADEMARKS & BRANDS ONLINE (August 7, 2014) ("The prevalence of online piracy in Singapore turns customers away from legitimate content and adversely affects Singapore's creative sector") (available at: http://www.trademarksandbrandsonline.com/news/singapore-passes-copyright-amendment-4011).

9 *Supra* n. 7.

NOTES CONTINUED:

10 Or, the effect on the labor market of a boatlift of Cubans to Miami in 1980. *See* D. Card, *The Impact of the Mariel Boatlift on the Miami Labor Market*, 43 INDUSTRIAL AND LABOR RELATIONS REVIEW 245-257 (1990). *See also* S. Galiani, P. Gertler, and E. Schargrodsky, *Water for Life: The Impact of the Privatization of Water Services on Child Mortality*, 113 JOURNAL OF POLITICAL ECONOMY 83-123 (2005) (available at: http://sekhon.berkeley.edu/causalinf/papers/GalianiWater.pdf).

11 For a study on this very scenario, *see* Galiani, *supra* n. 10.

12 Meyer, *supra* n. 6, at p. 155.

13 *A Guide to Social Experimentation*, European Commission (September 2011) (available at: http://ec.europa.eu/social/BlobServlet?docId=7112&langId=en).

14 Theoretically, the comparison here is not simply between the outcomes of the two groups. Rather, the control group is being used as a proxy for the treated group in the case where the treated group does not receive the treatment. This proxy (or counterfactual) is necessary because we only observe the outcomes of the treated group when it is treated. Obviously, the selection of the control group is really important. The strength of the evidence from a difference-in-difference analysis hangs on the legitimacy of the control group.

15 The importance of the selection of control group is detailed in citations (among many others) listed in *supra* n. 6.

16 *Singapore Study*, *supra* n. 1 at p. 328.

17 G.R. Barker, *Agreed Use and Fair Use: The Economic Effects of Fair Use and Other Copyright Exceptions in the Digital Age*, Unpublished Working Paper (July 9, 2013) (available at: SSRN:http://ssrn.com/abstract=2298618 or http://dx.doi.org/10.2139/ssrn.2298618); G.R. Barker and I. Png, *ALRC Inquiry into Copyright and the Digital Economy Submission in response to Discussion Paper 79: Unreliable Evidence on Fair Use*, Submission Received by the Australian Law Reform Commission, Copyright and the Digital Economy (DP 79) (2013) (available at: https://www.alrc.gov.au/sites/default/files/subs/507._g_r_barker__prof_i__png.pdf).

18 Enders, *supra* n. 5.

19 *Singapore Study*, *supra* n. 1 at p. 340.

20 *Id.* at p. 341.

21 Copyright Act (Singapore), Chapter 63, Revised Edition 2006 (available at: http://statutes.agc.gov.sg/aol/search/display/view.w3p;page=0;query=DocId%3A%22e20124e1-6616-4dc5-865f-c83553293ed3%22%20Status%3Ainforce%20Depth%3A0;rec=0).

22 The study acknowledges "a shock induced by the global financial crisis," but makes no attempt to adjust the analysis for it. *Singapore Study*, *supra* n. 1 at p. 343.

23 *Id.* at p. 338.

24 A possible case where both the treated and control groups receive the treatment but the control group remains valid (possibly) is if the control group is genetically immune from the disease being treated (or the treatment itself). Even so, the control group would be suspect.

25 Each series is mean-centered (on unity) for illustration purposes.

26 I put treatment in parenthesis because the *Control* group is also treated.

27 *Singapore Study*, *supra* n. 1 at p. 343.

28 My own statistical analysis—consistent with Figure 2—indicates that the most probable upward change in the linear trend for the *Private Copying* group occurred in 2004.

29 *See, e.g.*, *Singapore Study*, *supra* n. 1 at p. 38 (and similar statements elsewhere).

NOTES CONTINUED:

30 Angrist and Krueger, *supra* n. 6 at 1294. Angrist and Krueger provide an excellent and intuitive derivation of the difference-in-differences estimator. This is also the model estimated in Galiani *et al.*, *supra* n. 10, and many other studies using the difference-in-differences estimator.

31 With a valid control group, the *X* vector might include GDP across countries, states, or cities, or the incomes of individuals.

32 This approach avoids one criticism of Barker (2012), *supra* n. 17, related to the relative sizes of value-added by SSIC by estimating a fixed-effect for each SSIC.

33 M. Bertrand, E. Duflo and S. Mullainathan, *How Much Should We Trust Differences-In-Differences Estimates?* 119 QUARTERLY JOURNAL OF ECONOMICS 249-275 (2004). The models are estimated using the 'xtreg' command in STATA 12.

34 *Singapore Study*, *supra* n. 1, at p. 327.

35 Other interesting cases are provided at: http://www.tylervigen.com/spurious-correlations.

36 *Asking Permission*, *supra* n. 7.

37 *Intellectual Property and the U.S. Economy: Industries in Focus*, U.S. Patent & Trademark Office (March 2012) (available at: http://www.uspto.gov/sites/default/files/news/publications/IP_Report_March_2012.pdf).