

A replication plan for “Does social media reduce corruption?” (Information Economics and Policy, 2017)

Randall J. Hannum

Abstract

The importance of replicating economic research to improve the validity of findings has been the topic of an ongoing discussion, but there is not a consensus about what that means in practice. This article discusses a rationale for replicating a study and offers a plan of how one might go about approaching a replication of an actual study.

(Published in Special Issue [The practice of replication](#))

JEL B40 D73

Keywords Replications; corruption

Authors

Randall J. Hannum, New York City College of Technology, CUNY, Brooklyn, NY, USA, rhannum@citytech.cuny.edu

Citation Randall J. Hannum (2018). A replication plan for “Does social media reduce corruption?” (Information Economics and Policy, 2017). *Economics: The Open-Access, Open-Assessment E-Journal*, 12 (2018-49): 1–7. <http://dx.doi.org/10.5018/economics-ejournal.ja.2018-49>

Received August 24, 2017 Published as Economics Discussion Paper September 28, 2017

Revised July 4, 2018 Accepted July 16, 2018 Published August 7, 2018

© Author(s) 2018. Licensed under the [Creative Commons License - Attribution 4.0 International \(CC BY 4.0\)](#)

1 Introduction

The importance of replicating empirical research in economics has been advanced by numerous studies. The difficulties associated with replicating studies have also been discussed at length (Duvendack et al., 2017; Lindsay and Ehrenberg, 1993 among others). The topic this paper will address is an approach to actually doing a meaningful replication.

Before looking at a specific approach to performing a replication, however, it is important to consider a rationale for the replication. For this, assume a model with three actors: an original researcher, a replicator, and a consumer of the research output, who is dependent on the researcher and the replicator to determine a course of action (i.e. if this is done, we can expect that to happen). Replications are of value to the consumer if the consumer gains more certainty about the reliability of the results obtained by the researcher. In an ideal world, value is obtained as a result of an independent verification of the results and possible clarification of the context in which the results were obtained. In this scenario, the replication may support the original research or may identify data errors or methodological errors that affect the results obtained and the conclusions drawn.

One factor that may undermine the replication effort is that research has found that replications that refute findings of the original are more likely to be published than those that support the original (Galiani et al., 2017). This bias in publishing success tends to raise questions about the credibility of replications given the incentive to find a basis on which to refute part or all of the original research findings. The challenge, therefore, is how to perform a credible replication that creates value for the consumer.

2 Replication Principles

Since there is no agreed upon definition of what a replication is, the replicator first needs to define the scope of the replication. A number of papers dealing with replications have presented frameworks for classifying types of replication studies (Reed, 2017; Clemens, 2017; Morrison et al., 2010). Replications are usually differentiated along two dimensions in terms of the data used and the methods employed to analyze the data. This paper will use a slightly different approach by starting with the question that replication studies should be dealing with, which is, are the results of the original study reliable? Or, put in a more operational form, do the results hold up to constructive scrutiny? Looked at from this perspective, a replication can range from a quality control exercise to an expanded robustness check based on an evaluation of the data considered and the methods employed, to further extensions of the original study. Types of replications can differ based on how far along that continuum the replicator journeys.

In deciding the scope of the replication, one should consider the validity of the approach to be taken, which is why the above framework is suggested. If the replication doesn't start with the original data and the methods employed in the original study to determine if the results can be truly replicated, what Reed (2017) refers to as *Reproduction*, then the results from the replication can easily be dismissed if they differ from the results of the original study. By confirming that the original results can be duplicated using the original data and, optimally, the

original code, a replicator can ensure that any criticism of the original study is based on full knowledge regarding how the original results were obtained. If the original results cannot be replicated as a result of honest attempts to do so (a discussion with the authors of the original study should be part of this process), then it is safe to say that the results of the original study do not hold up to constructive scrutiny and, therefore, are unreliable.

Since replications are not universally performed on all published empirical research in economics, an additional issue that the replicator should consider is the process used to choose a particular study to replicate. One method has been to replicate all papers in a particular journal or dealing with a particular topic. By applying the same approach to a number of papers, the replicator avoids the criticism of singling out a particular paper for excessive scrutiny. While this approach has proven useful in providing evidence that replications should more routinely be performed, it may be impractical if the objective is to go beyond seeing whether the researcher's data and code can be used to obtain identical results.

Since one motivation for selecting a particular study to replicate is that something seems “wrong” with the results, it may seem that the replicator is merely trying to undermine the results of the original study rather than constructively scrutinizing how they were obtained. Given this potential bias in selecting a particular study, the replicator should clearly indicate where the results of the replication of the original study end and where anything new, in terms of data or methodology, is being introduced. Subjecting the original results to constructive scrutiny can be seen as equivalent to cross-examination as employed by the court system. Just as in courts there are rules to make the system fairer, there should be rules that replicators follow to preserve the integrity of the process. The rest of this paper expands on these guidelines by showing how they can be used to develop a plan to replicate an actual paper.

3 Replication Plan

The paper chosen for this purpose is “Does Social Media Reduce Corruption?” by Jha and Sarangi (2017). As the title suggests, Jha and Sarangi set out to answer the question: Does social media reduce corruption? The author chose this particular paper for the replication discussion because it seems to make a significant claim, that social media is socially beneficial, and to highlight a number of issues in developing a replication plan. As will be seen, these issues relate to the data, the methods, and a consideration of the conclusions drawn. Before discussing a replication plan for this paper, a brief summary of the main result is provided.

The primary finding of the Jha and Sarangi paper is that Facebook penetration in a country, used as a proxy for social media penetration, is negatively correlated with corruption levels. To analyze this relationship, they start by estimating the following equation:

$$\text{Corruption}_i = \alpha + \beta \text{Facebook}_i + \delta \text{Internet}_i + \gamma_1 \log(\text{GDPPC}_i) + \gamma_2 \text{Political Rights}_i + \gamma_3 \log(\text{Population}_i) + \gamma_4 \text{Christian}_i + \gamma_5 \text{Muslim}_i + \varepsilon_i$$

The Corruption measure used is the negative transformation of the 2012 Control of Corruption Index, which is published by the World Bank and, as published, ranges from –2.5 to

2.5 with 2.5 representing the lowest level of corruption. The Facebook variable represents Facebook penetration and comes from an organization called Quintly. The Internet variable represents internet penetration, and GDPPC is per capita GDP. The data for both of these come from the World Bank, along with the data for the Population variable. The Political Rights variable is published by Freedom House and ranges from 1 to 7, with 1 representing the best level of political rights. The Christian and Muslim variables represent the proportion of the population identified as such and are published by the Association of Religion Data Archive. The year for the data is 2012, except for religion variables, which is 2005. Data for all of these variables were available for 170 countries.

The OLS estimates of this model found the coefficient of Facebook to be highly significant ($p < 0.01$) and of the expected (negative) sign to indicate a negative relationship between Facebook penetration and level of corruption. Alternative specifications of the model presented also show Facebook penetration to be significant and of the correct sign. The coefficients of Facebook presented in the paper range from -0.0114 to -0.0163 . As the authors point out, this means that by the most conservative estimate of the coefficient (-0.0114), a one-standard deviation (18.2%) increase in Facebook penetration (mean = 20.62%) increases the Control of Corruption Index by 0.2.

The replication plan being presented for this study would best be described as replication with constructive scrutiny to determine the extent to which the results could be described as reliable. The replication plan assesses only the evidence presented and, where appropriate, scrutinizes choices made regarding data considered or methods used that differ from the choices made in papers the authors reference. In terms of the replication classification scheme presented by Reed (2017), the plan presented here would best be described as a (1) reproduction and (2) robustness analysis using the same data.

The starting point in the replication would be to assess the argument made in light of the evidence presented. As mentioned in the summary of the Jha and Sarangi paper, the authors present evidence, based on data from more than 150 countries, that Facebook penetration, used as a proxy for social media, is negatively correlated with corruption. While the authors point out that they are not claiming that their evidence supports the hypothesis that Facebook penetration causes less corruption, their interest is in the significance of this social media variable.

The next step would be to obtain the data and code used to generate the results presented in the paper. Some journals require these to be submitted but are sometimes rather lax about actually obtaining them. Often, the author needs to be contacted to obtain the data and the code. The data and code for the chosen paper, for instance, are not available from the journal, but the authors did make them available to me upon request. Contacting the author, even if the data are publicly available, should be done anyway to let the author know what your plans are and to raise any preliminary questions you may have. An alternative approach is to try to recreate the data based on the descriptions presented in the paper. While this is a recommended step further along in the replication plan, it is more efficient to start with the original data and code given that the descriptions provided in papers may not be sufficient to recreate identical data given possible revisions, corrections, transformations, etc.

Once the original data are obtained, the next step would be to determine whether the results can be duplicated or reproduced. Does the identical data and code yield the published results? If they cannot, any discrepancies should be clarified with the author of the original rather than

assuming that the replicator has discovered flaws in the original. If honest attempts to reconcile the discrepancies still yield different results, it should then be determined whether the differences alter the conclusion of the original study. In the case of the paper being considered, this would be whether the social media variable is significant or not. If the variable is not statistically significant, the original has failed the reproduction stage.

Following the reproduction stage of the replication, the next step is to look at whether there are any issues with the data. The chosen paper, for instance, relies on data from a number of sources that have been merged together. Has that been done correctly? This step is to determine whether reconstructing the data from the original sources, based on the descriptions provided in the original paper, yields an identical data set. Errors can be introduced as a result of trying to combine data with corrupted, missing or duplicate unique data identifiers. Again, if there are any differences or discrepancies, the source should be identified along with the effect, if any, the differences have on the results that are obtained. At this stage of the replication, the social media variable would be expected to still be significant for the replication to be judged a success.

The next step would be to compare the methods discussed in the paper to the code used to analyze the data. This step would involve checking if there are any differences in the variables used or statistical methods discussed. Any discrepancies should be investigated, and if they remain, it should be determined if they have any effect on the results. In the case of the paper being evaluated, the important consideration would be whether the differences have any effect on the significance of the social media variable.

Once the original research has been examined to determine if there are any errors or discrepancies between the data and methods described and the data and methods used in the analysis, the next phase of the replication plan is to examine both the data and methods used in the analysis to see how sensitive the results are to the choices made regarding the data included and the methods employed. This phase of the replication would correspond to the Reed (2017) classification of robustness analysis – same data set. For instance, in the paper chosen, it is noted that Argentina was not included in the analysis. It is also noted that the regression model does not include the Protestant percentage of the population by country, which is a variable included in many of the papers cited. Both would be checked to determine the effect on the results. If the social media variable remains significant, the replication would still be judged a success.

Any analysis beyond what has been described would depend on the objectives that the replicator has beyond simply verifying the results of the original research. For instance, with the paper chosen, it might seem reasonable to wonder whether changes in corruption are correlated with changes in Facebook penetration. Since this question goes beyond the scope of what the authors are arguing, it might be an unreasonable basis on which to claim that the replication rejects the original if it is found that corruption is not significantly correlated with changes in Facebook penetration. This would also be the case if the Facebook penetration variable is replaced with another social media variable, and the new variable is found to not be significant. Results that do not support the findings of the original would raise questions about limitations of the original results but would not be sufficient to reject them.

Judging Replication Success

In this paper it is being argued that a replication of the chosen paper could “fail” in several different ways. The most obvious way is if the results obtained using exactly the same data and code failed to yield a statistically significant coefficient on the Facebook variable. It could also be said to have failed if it is found that there are errors in the data, either different from those that are published or assembled incorrectly, meaning data from one country was incorrectly matched up with another country, which, again, yields a result at odds with the main results of the paper. A third way is if the actual methods used in the analysis differ from what was discussed in the paper in a way which, when changed to the described methods, yield results that are no longer significant. If the replication yields identical results or results that do not undermine the results of the original, then the replication has “successfully” replicated the original study.

4 Conclusion

In the end, the result of a replication is whether it confirms or rejects all or parts of the original study. In the strictest sense, the confirmation or rejection should apply only to whether or not the original results can be reproduced using the original data and code. In a broader sense, a rejection can also apply to whether there are any issues with the original data and/or code that undermine the validity of the original results. Any issues beyond that should be seen as clarification or an extension of the original study. The value of the replication lies in its persuasiveness from the perspective of the consumer. If the replication results show no significant discrepancies from the original, it may seem as though effort was wasted needlessly, but the replication does provide evidence that the results are reliable. If the replication casts doubts on the reliability or validity of the original results, it should clearly point out what the specific issues are, which, hopefully, has the end result of clarifying the value of the original effort.

Acknowledgements I would like to thank Benjamin Wood, Brian D. Haig, an anonymous referee and W. Robert Reed for their many helpful comments. Any mistakes, though, are my own.

References

- Clemens, M. (2017). The meaning of failed replications: a review and proposal. *Journal of Economic Surveys* 31 (1): 326–342. <https://doi.org/10.1111/joes.12139>
- Duwendack, M., Palmer-Jones, R., and Reed, R. (2017). What is meant by ‘replication’ and why does it encounter such resistance in economics? *American Economic Review* 107 (5), 46–51. <https://doi.org/10.1257/aer.p20171031>
- Galiani, S., Gertler, P.J., and Romero, M. (2017). Incentives for replication in economics. Available at SSRN: <https://ssrn.com/abstract=2999062> or <http://dx.doi.org/10.2139/ssrn.2999062>
- Jha, C.K., and Sarangi, S. (2017). Does social media reduce corruption? *Information Economics and Policy* 39: 60–71. <https://doi.org/10.1016/j.infoecopol.2017.04.001>
- Lindsay, R.M. and Ehrenberg, A.S.C. (1993). The design of replicated studies. *The American Statistician* 47 (3): 217–228. <https://doi.org/10.2307/2684982>
- Morrison, R., Matuszek, T., and Self, D. (2010). Preparing a replication or update study in the business disciplines. *European Journal of Scientific Research* 47 (2): 278–287. http://www.academia.edu/download/31222294/ejsr_47_2_12.pdf
- Reed, W.R. (2017). Replication in labor economics. *IZA World of Labor* 2018: 413. <https://doi.org/10.15185/izawol.413>

Please note:

You are most sincerely encouraged to participate in the open assessment of this article. You can do so by either recommending the article or by posting your comments.

Please go to:

<http://dx.doi.org/10.5018/economics-ejournal.ja.2018-49>

The Editor