

# Replications in Development Economics

By SANDIP SUKHTANKAR\*

Replications in the social sciences have garnered considerable attention lately. In psychology, the “Many Labs” and Open Science Collaboration projects involved 270 researches attempting to replicate 100 studies published in the top three journals (Open Science Collaboration, 2015). In economics, Camerer et al. (2016) attempted to replicate 18 papers in experimental economics. To aid efforts in replication, all of the “Top 5” journals in economics have now adopted a Data Access Policy (DAP), requiring authors to make available data and code to all researchers for the purposes of replication.<sup>1</sup> Yet the extent to which replications occur within economics is uncertain.

In this paper, I calculate the replication rate for papers in the field of development economics published since the year 2000 in one of the top 5 or next 5<sup>2</sup> general interest economics journals, categorize the types of replications, and examine the correlates of replicated papers. The issue of replications in development economics is particularly interesting for a number of reasons. First, there has been a marked increase in the number of empirical papers published in development, driven by advances in empirical methods, data availability, as well as the proliferation of Randomized Control Trials (RCTs) (Figures 1 and A1). Second, development economists work in a variety of contexts, with different institutional and market structures; whether results generalize across these various contexts is thus important to determine. Finally, the increase in the ease of conducting RCTs offers the hope that direct replication of interventions and tests of theories across contexts is possible.

In order to determine the replication rate, one has to define what is meant by “replication.” Social psychologists have proposed a conceptual definition: “direct replication is the attempt to duplicate the conditions and procedure that existing theory and evidence anticipate as necessary for obtaining the effect” (Nosek and Lakens, 2014). Economists have come up with various working definitions: for example, Clemens (2015) separates replications from robustness tests, classifying

\* University of Virginia, Monroe Hall, Charlottesville, VA 22904; sandip.sukhtankar@virginia.edu. I thank Lucas Coffman and Brian Nosek for helpful discussions, and Hanna Charankevich and Mia Giurato for excellent research assistance.

<sup>1</sup>The *American Economic Review* (AER) was the first to adopt this policy in 2005. It states that “It is the policy of the American Economic Association to publish papers only if the data used in the analysis are clearly and precisely documented and are readily available to any researcher for purposes of replication.” The *Quarterly Journal of Economics* (QJE) and the *Journal of Political Economy* (JPE) have recently adopted this wording verbatim; *Econometrica* uses similar wording. The *Review of Economic Studies* (ReStud) has a DAP but it does not make clear whether data need to be made available only to the journal or to other researchers too.

<sup>2</sup>These are the *American Economic Journal: Applied Economics* (AEJAE), the *American Economic Journal: Economic Policy* (AEJEP), the *Economic Journal* (EJ), the *Journal of the European Economic Association* (JEEA), and the *Review of Economics and Statistics* (ReStat).

“verification” (replication) as those using the same sample, population, and empirical specification; “reproduction” (replication) as those using different samples from the same population but using the same specification; “reanalysis” (robustness) using the same sample and population but different specifications; and “extension” (robustness) using different samples and populations but the same specification. Hamermesh (2007) separates “pure replications” (same methods, same sample and population) from “statistical replications” (different sample, same methods and population) and “scientific” replications (different sample and population, similar but not identical methods). While seemingly incongruent, both seem to agree on what is “pure replication” or “verification”; the “statistical replication” of Hamermesh (2007) corresponds to the “reproduction” of Clemens (2015); while the “scientific replication” of Hamermesh (2007) incorporates the “robustness” categories defined by Clemens (2015), although it appears to be broader.

For the purposes of this paper, I consider all papers conforming to any of the Clemens (2015) classifications - including those he classifies as robustness tests - as replications. These categories define replication more narrowly than, for example, Berry et al. (2017) in this issue, who consider a replication to be “any project that reports results that speak directly to the veracity of the original paper’s main hypothesis.” In addition, my method for finding replications relies to a large extent on replicating papers self-identifying as replications; I discuss the implications for the replication rate below. Finally, when calculating the replication rate, I only consider whether a paper was replicated or not; I do not consider whether the paper *successfully* replicated the analysis.

## I. Data

Using Econlit, I searched for every paper published in the above-referenced ten journals that were published between the years 2000 and 2015 (inclusive) and had an “O” Journal of Economic Literature classification. I dropped papers that were themselves replications or comments. I classified the remaining papers into those that were purely theoretical and those that contained some empirical analysis, and the empirical papers into those that involved an RCT or not. This yielded a total of 1,390 papers, with 252 pure theory papers, 120 RCTs, and 1,018 empirical papers not involving RCTs (Table A1).

Figures 1 and A1 show the evolution of numbers of papers in each of the categories over the period I study. Contrary to perceptions, theoretical papers have not been crowded out by empirical papers and RCTs; the number of theory papers is more or less constant over the years.<sup>3</sup> The number of empirical papers does seem to have increased significantly, driven largely by the increase in papers published based on RCTs. The introduction of the *American Economic Journals* in 2009 also seems

<sup>3</sup>Note that my sample does not include the *American Economic Journal: Microeconomics*, which mainly publishes theoretical papers; hence it is possible that the number of theoretical papers published has also increased.

TABLE 1—REPLICATIONS BY RCT STATUS

	No		RCT Yes		Total	
	#	%	#	%	#	%
	Not Replicated	962	94.5	105	87.5	1,067
Replicated	56	5.5	15	12.5	71	6.2
Total	1,018	100.0	120	100.0	1,138	100.0

to have driven much of this change.

For the empirical papers in my sample, I searched for published or working papers that attempted to replicate the empirical analysis or intervention. The main method involved doing a reverse citation search on Google Scholar, and searching within the ensuing list for “replication” or alternative cognates. I supplemented this method by examining websites of various projects dedicated to replications in economics. The Online Appendix describes data and methods in detail.

## II. Results

Of the 1,138 empirical papers, I found that 71 (6.2%) were replicated in a published or working paper. The number drops to 37 (3.3%) if considering only already published papers. RCTs seem to be replicated at a higher rate, with 15 of the 120 RCTs (12.5%) being replicated.<sup>4</sup>

Table 2 classifies the 71 replicating papers into the Clemens (2015) categories. The majority of replications involve reanalyses using different econometric specifications or reconfigurations of the data. A few are extensions in the sense that a particular intervention is tried in a different context. A common theme is a reanalysis using a new econometric technique. No paper solely does a verification or pure replication; this is usually done as a precursor to the reanalysis. None of the papers involves a reproduction, i.e. the same methods applied to a different sample from the same population. Finally, the replicating papers confirm the widely held view that a positive replication - one that simply confirms the findings of the original study - is basically unpublishable in economics journals: only one of the published replications wholeheartedly confirms the original findings.<sup>5</sup>

At face value the overall rate of replication seems low, for example when compared to the nearly 30% rate found by Berry et al. (2017) for one volume of the AER. There are at least two reasons why the rate here is lower; the first is the narrower definition of “replication”, and the second

<sup>4</sup>Five of the replicated RCTs are replicated by a single paper that uses an alternative statistical model for analysis; even without including these five, the rate of replication for RCTs is still higher, and statistically different from that of non-RCTs.

<sup>5</sup>This is the confirmation by Chu, Henderson and Wang (2016) of the Nunn and Qian (2014) results on US food aid and conflict, published in the *Journal of Applied Econometrics*, which is the journal that publishes the most replications in economics (Duvendack, Palmer-Jones and Reed, 2017).

TABLE 2—REPLICATIONS BY TYPE

	%
ReplicationVerification	32.4
ReplicationReproduction	0
RobustnessReanalysis	77.5
RobustnessExtension	36.6

involves the search method which relies chiefly on papers self-identifying as replications. Given publication incentives, it makes sense that authors do not wish their work to appear derivative, and only choose to self-identify a paper as a replication when their paper attempts to replicate analysis in the original paper in a narrow statistical sense rather than a broader conceptual sense.

Whether scarce researcher resources should be devoted to increase this rate of replication is of course an entirely different although important question. In the rest of this section I restrict myself to analyzing the correlates of the current set of replicated papers.

A reasonable basis for determining whether a paper should be replicated or not might consider the impact of the paper as well as the uncertainty of the empirical analysis; i.e., influential papers with large standard errors should be replicated. While it is a task of considerable magnitude to determine the “uncertainty” of the empirical analysis in over a thousand papers, measuring influence via citations is a much more attainable task. I find that papers that are replicated are far more highly cited than papers that are not; on average, a replicated paper has nearly 4 times as many Google Scholar citations as a non-replicated paper (833 vs 232). On this basis, the “right” papers are being replicated.

I run some simple regressions in order to assess the statistical relevance of the above data, and examine the correlates of replications in more detail. Table 3 confirms that having a higher citation count as well as being an RCT are both strongly correlated with being replicated. It also reveals that the linear term for year of publication is significant, suggesting that more recent papers are more likely to be replicated. Being published in a “Top 5” journal by itself does not seem to be predictive after taking into account citation counts. The RCT and citation count results survive the inclusion of year and journal dummies, as well as a logistic rather than a linear probability specification.

### III. Discussion

While the reason for replications of papers with high citation counts is clear, why RCTs as well as more recent papers tend to be replicated more is less clear. One possible explanation is data availability. While the AEA adopted its DAP in 2005, other top journals have only adopted these policies as recently as 2015. For example, one paper that replicates multiple RCTs notes: “Due to

TABLE 3—SIMPLE REPLICATION OLS REGRESSION

	(1) Replicated
year	0.00399* (0.00170)
Citation Count	0.000145*** (0.0000150)
RCT	0.0658** (0.0228)
Paper is from top five journal	0.0262 (0.0143)
Constant	-8.002* (3.420)
Observations	1138

Standard errors in parentheses

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

the policies of the two journals that published these papers - the AEJ:Applied and Science - all the microdata from these RCTs is freely available online” (Maeger, 2016). In general, the donors that fund RCTs require authors to make data available for others to use, not simply for replication but to maximize their bang for the buck.

Other reasons for why newer papers are replicated more may simply have to do with recency bias. For example, the econometric theory papers that replicate and reanalyze empirical papers in order to make a methodological point seem to just pick a recent paper published in the AER that satisfies the criterion they need.<sup>6</sup>

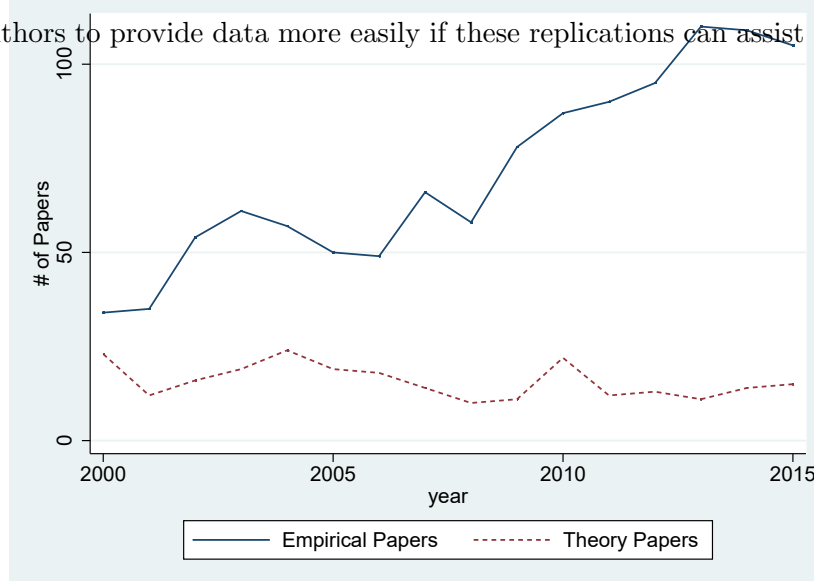
As surmised in the introduction, RCTs do seem to allow for direct replication of interventions in different contexts. For example, the *AEJAE* dedicated an issue (January 2015) to six RCTs (five in developing countries) that each tested the impact of expansion of microfinance in various contexts. While the interventions were not identical, the basic theory being tested was the same. An example in which the identical intervention was tested in six countries is Banerjee et al. (2015), who examined an intervention that assisted the extreme poor to “graduate” out of poverty via sustainable self-employment opportunities.

What might be done to encourage more replications? Other papers in this session and elsewhere directly examine incentives to replicate. However, one source of data that is currently available but

<sup>6</sup>For example, Sloczynski (2014) replicates Martinez-Bravo (2014) in a paper titled “New Evidence on Linear Regression and Treatment Effect Heterogeneity,” where the criteria for selecting this paper appear to be as simple as the ones outlined here.

untapped is the set of replications undertaken by PhD students for coursework everywhere. Many top PhD program in economics have classes in which graduate students replicate papers and run robustness tests or extensions.<sup>7</sup> A recent blog post describes one such exercise in a macro course at the New Economic School, in which the professor asked students to replicate and test robustness for papers published in the AER.<sup>8</sup>

Indeed, Fecher, Frassdorf and Wagner (2016) argue that replication exercises should be a mandatory part of PhD coursework in economics. Such replications can at the very least serve as verification tests on the data and code used, but currently there is no way to access them. A repository of these replications - for example through the Center for Open Science (<https://cos.io/>) - might help increase our confidence in currently published papers. Such a repository could provide incentives to graduate students to complete and write up replication efforts as papers, as well as incentives for authors to provide data more easily if these replications can assist in validating their results.



<sup>7</sup>For example, a fellow graduate student and I were able to precisely replicate the results in Abadie and Gardeazabal (2003) in a graduate course at Harvard in 2006.

<sup>8</sup>See <https://replicationnetwork.com/2016/12/27/campbell-is-the-aer-replicable-and-is-it-robust-evidence-from-a-class-project/> for details.

FIGURE 1. BREAKDOWN OF PAPER TYPES BY YEAR

## REFERENCES

- Abadie, Alberto, and Javier Gardeazabal.** 2003. "The Economic Costs of Conflict: A Case Study of the Basque Country." *American Economic Review*, 93(1): 113–132.
- Banerjee, Abhijit, Esther Duflo, Nathanael Goldberg, Dean Karlan, Robert Osei, William Parienté, Jeremy Shapiro, Bram Thuysbaert, and Christopher Udry.** 2015. "A multifaceted program causes lasting progress for the very poor: Evidence from six countries." *Science*, 348(6236).
- Berry, James, Lucas Coffman, Douglas Hanley, Rania Gihleb, and Alistair Wilson.** 2017. "Assessing the Rate of Replication in Economics." *American Economic Review*, 107.
- Camerer, Colin F., Anna Dreber, Eskil Forsell, Teck-Hua Ho, Jürgen Huber, Magnus Johannesson, Michael Kirchler, Johan Almenberg, Adam Altmeld, Taizan Chan, Emma Heikensten, Felix Holzmeister, Taisuke Imai, Siri Isaksson, Gideon Nave, Thomas Pfeiffer, Michael Raza, and Hang Wu.** 2016. "Evaluating replicability of laboratory experiments in economics." *Science*.
- Chu, Chi-Yang, Daniel J. Henderson, and Le Wang.** 2016. "The Robust Relationship Between US Food Aid and Civil Conflict." *Journal of Applied Econometrics*.
- Clemens, Michael A.** 2015. "The Meaning of Failed Replications: A Review and Proposal." *Journal of Economic Surveys*.
- Duvendack, Maren, Richard Palmer-Jones, and W. Robert Reed.** 2017. "What is Meant by 'Replication' and Why Does It Encounter Resistance in Economics?" *American Economic Review*, 107.
- Fecher, Benedikt, Mathis Frassdorf, and Gert Wagner.** 2016. "Perceptions and Practices of Replication by Social and Behavioral Scientists: Making Replications a Mandatory Element of Curricula Would Be Useful." IZA Discussion Paper 9896.
- Hamermesh, Daniel S.** 2007. "Viewpoint: Replication in economics." *Canadian Journal of Economics*, 40(3): 715–733.
- Maeger, Rachel.** 2016. "Understanding the Impact of Microcredit Expansions: A Bayesian Hierarchical Analysis of 7 Randomised Experiments." MIT.
- Martinez-Bravo, Monica.** 2014. "The Role of Local Officials in New Democracies: Evidence from Indonesia." *American Economic Review*, 104(4): 1244–87.

- Nosek, Brian, and Daniel Lakens.** 2014. "Registered reports: A method to increase the credibility of published results." *Social Psychology*, 45(3): 137–141.
- Nunn, Nathan, and Nancy Qian.** 2014. "US Food Aid and Civil Conflict." *American Economic Review*, 104(6): 1630–66.
- Open Science Collaboration.** 2015. "Estimating the reproducibility of psychological science." *Science*, 349(6251).
- Sloczynski, T.** 2014. "New Evidence on Linear Regression and Treatment Effect Heterogeneity." IZA Discussion Paper 9491.



**Online Appendix: Not for Publication**

## TABLES AND FIGURES

TABLE A1—BREAKDOWN OF PAPER TYPES

	RCT			Total
	-	No	Yes	
Empirical	0	1018	120	1138
Theory	252	0	0	252
Total	252	1018	120	1390

TABLE A2—“CITED BY” COUNTS

	RCT		
	No	Yes	Total
Not Replicated	236.1	190.2	231.6
Replicated	954.7	372.1	831.6
Total	275.6	213.0	269.0

TABLE A3—REPLICATION OLS REGRESSION WITH YEAR AND JOURNAL DUMMIES

	(1) Replicated
Citation Count	0.000146*** (0.0000151)
RCT	0.0566* (0.0232)
Observations	1138

Standard errors in parentheses

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

TABLE A4—SIMPLE REPLICATION LOGIT REGRESSION

	(1) Replicated
year	0.0679 (0.0368)
Citation Count	0.00141*** (0.000251)
RCT	0.946** (0.330)
Paper is from top five journal	0.655* (0.280)
Constant	-140.2 (74.02)
Observations	1138

Standard errors in parentheses

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$ 

TABLE A5—REPLICATION LOGIT REGRESSION WITH YEAR AND JOURNAL DUMMIES

	(1) Replicated
Citation Count	0.00190*** (0.000341)
RCT	0.792* (0.359)
Observations	1037

Standard errors in parentheses

\*  $p < 0.05$ , \*\*  $p < 0.01$ , \*\*\*  $p < 0.001$

FIGURE A1. QUANTITY AND SHARE OF EMPIRICAL PAPERS WHICH ARE RCTs BY YEAR




FIGURE A2. SHARE OF EMPIRICAL PAPERS WHICH ARE REPLICATED BY YEAR




FIGURE A3. AVERAGE NUMBER OF CITATIONS PER YEAR FOR ALL EMPIRICAL PAPERS



## DATA

A list of all papers satisfying the following conditions was compiled using EconLit:

- 1) The paper was published between 2000 and 2015 (inclusive)
- 2) The paper was published in one of the following journals
  - a) “Top five” journals
    - the American Economic Review (AER)<sup>9</sup>
    - Econometrica
    - the Journal of Political Economy (JPE)
    - the Quarterly Journal of Economics (QJE)
    - the Review of Economic Studies (ReStud)
  - b) Five other highly regarded general interest journals
    - the American Economic Journal: Applied Economics (AEJAE)
    - the American Economic Journal: Economic Policy (AEJEP)
    - the Economic Journal (EJ)
    - the Journal of the European Economic Association (JEEA)
    - the Review of Economics and Statistics (ReStat)
- 3) The paper had an “O” EconLit classification code

Papers that were themselves comments or replications were dropped. Papers were then manually labeled as an empirical paper or a theory paper, and the empirical papers as RCT or non-RCT.

## METHOD FOR DETERMINING WHETHER PAPERS WERE REPLICATED

I used the following methods to determine whether a paper classified as “empirical” in the list above was replicated, and to find the corresponding replication paper:

- 1) I searched for each paper separately in Google Scholar, and then searched for the terms *replicate OR replicates OR replicated OR replication OR replicating* within the set of papers that cited this paper. The abstracts and introductions of papers satisfying both criteria were then examined to determine whether the citing paper did indeed attempt a replication of the original paper. The majority of replications were found using this method.
- 2) I supplemented the above method using the following websites to search for replications:

<sup>9</sup>All “Papers and Proceedings” papers are excluded.

- a) [http://replication.uni-goettingen.de/wiki/index.php/Main\\_Page](http://replication.uni-goettingen.de/wiki/index.php/Main_Page). The website allows you to search by JEL category, and I examined search results for every single “O” classification.
- b) <https://replicationnetwork.com/>
- c) <http://www.3ieimpact.org/en/publications/3ie-replication-paper-series/>

Both published papers as well as unpublished working papers were counted as replication papers. If multiple versions of the same paper were found, the published version was included in the data.