

Science Deserves Better: The Imperative to Share Complete Replication Files

Allan Dafoe, *Yale University*

In April 2013, a controversy arose when a working paper (Herndon, Ash, and Pollin 2013) claimed to show serious errors in a highly cited and influential economics paper by Carmen Reinhart and Kenneth Rogoff (2010). The Reinhart and Rogoff paper had come to serve as authoritative evidence in elite conversations (Krugman 2013) that high levels of debt, especially above the “90 percent [debt/GDP] threshold” (Reinhart and Rogoff 2010, 577), posed a risk to economic growth. Much of the coverage of this controversy focused on an error that was a “perfect made-for-TV mistake” (Stevenson and Wolfers 2013) involving a simple error in the formula used in their Excel calculations. The real story here, however, is that it took three years for this error and other issues to be discovered because replication files were not publicly available, nor were they provided to scholars when asked. If professional norms or the *American Economic Review* had required that authors publish replication files, this debate would be advanced by three years and discussions about austerity policies would have been based on a more clear-sighted appraisal of the evidence.

An essential characteristic of science is the commitment to transparency. Assumptions should be clearly stated, evidence should be publicly verifiable, and the basis for inferences should be explicit. Independent researchers should be able to reproduce, at least in principle, the structure of inferences linking assumptions, prior theory, other findings, data collection, data processing, and data analysis, to an alleged scientific finding. Much of the institution of science exists to promote transparency, such as the strong norms around citations, the requirement to describe methods, the esteem for formal methods of inference (statistics and formal theory), expectations about maintaining (laboratory or field) notebooks, the expectation to publish proofs of theorems, and the condemnation reserved for nonreproducible results.

Transparency is a foundation for a number of core features of science: *refutability*, *openness*, *cumulation*, and *minimal barriers to entry*. Transparency makes scientific work more *refutable*—more subject to detailed criticism—which is the basis for scientific progress. Transparency makes a scientific enterprise more *open* to exploration by others, facilitating divergent interpretations of results and alternative uses of data and tools. The *cumulation* of data, tools, and findings is essential for the progress of science. The sharing of findings is incentivized through rewards to publications. There are not, however, adequate individual incentives for the sharing of data and tools. The sharing of data and tools is an essential

public good for science; a commitment to transparency would promote this public good. Transparency, by making more steps of the scientific process publicly observable, *reduces the barriers to entry* for students and novices. Keeping scientific discussion accessible improves the scientific enterprise by reducing the costs to training new scientists, by bringing in new perspectives, by permitting more cross-disciplinary conversation, and by keeping fields open to external criticism.

This article makes a simple argument. Political science should take its commitment to transparency more seriously by insisting that researchers publish complete replication files, making every step of research as explicit and reproducible as is practical. In return, political science will become more refutable, open, cumulative, and accessible. Science deserves this commitment from us.

This article proceeds as follows. The first section reviews some evidence about the current state of replication practices in political science. The second section elaborates on the benefits of greater transparency through the sharing of complete replication files. Specific recommendations for authors, journals, and universities are provided in the third section. Online Appendix A discusses exceptions to the prior recommendations for confidential, costly, or proprietary data or code. Online Appendix B discusses a proposal for Replication Audits.

The recommendations presented here apply to any domain of science in which some feature of inference could be practically made more explicit and reproducible. These recommendations apply especially to modes of inference that use computers, because any processing involving a computer can be codified and made reproducible. For this reason this article focuses primarily on replication practices in statistical studies, although the recommendations apply equally to computational theory (theoretical models using computer simulations or solutions). Noncomputational modes of inference can also be made more transparent. For example, Moravcsik (2010) offers valuable recommendations to qualitative researchers.

EVIDENCE FROM POLITICAL SCIENCE

What is the current state of replication practices in political science? Gherghina and Katsanidou (2013) found that only 18 of 120 political science journals have a replication policy posted on their websites, to say nothing about enforcing those policies. To provide additional data on the state of replication practices in political science, I collected data on the availability of

replication files for publications in two leading journals—*American Political Science Review* (APSR) and *American Journal of Political Science* (AJPS)—and scholars’ attempts to replicate publications.

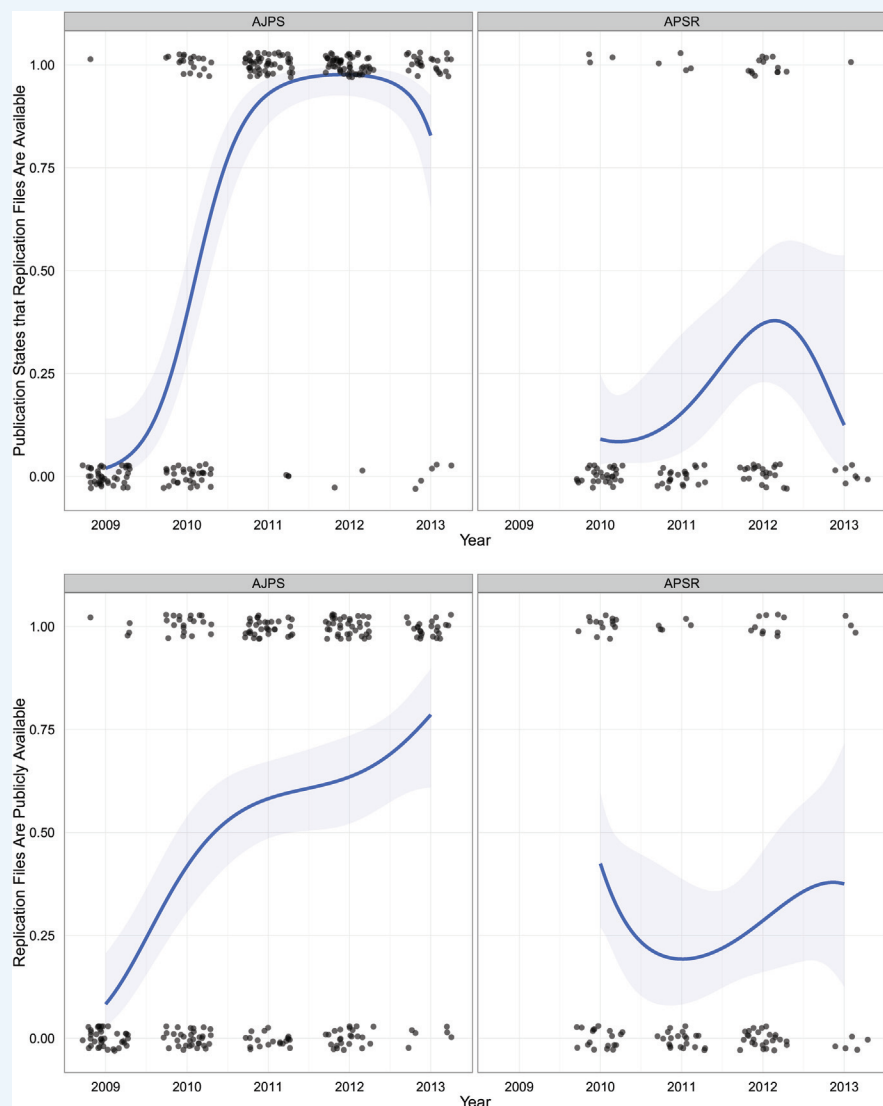
Replication Practices at APSR and AJPS

APSR does not have a policy of requiring replication files, although it encourages them. For example, in the submission guidelines it states that authors “are expected to address the issue of data availability. You must normally indicate both where (online) you will deposit the information that is necessary to reproduce the numerical results and when that information will be posted.” By contrast, the policy at AJPS as of 2010 states that articles “will not be published unless the first footnote explicitly states where the data used in the study can be obtained,” the acceptance letter provides instructions for posting files to AJPS’s Dataverse site, and the editor, Rick Wilson, frequently double checks that files are posted and has held up publications that have not posted replication files. The beneficial effects of this policy and editorial involvement are evident in figure 1.

Data was collected¹ on the availability of replication files for recent publications in the two top political science journals, the APSR since 2010 and the AJPS since 2009. We found that 48% of publications using statistical analysis stated on their first page that replication files were available; we were able to find replication files for 68% of these. We were also able to find replication files for 18% of the publications that did not state that replication files were available.

As figure 1 shows, publications at APSR are much less likely than AJPS to state that replication files are available, and somewhat less likely to provide replication files. Since 2011, nearly 100% of publications at AJPS state that replication files are available, increasing dramatically from 2009 before the new replication policy. This policy seems to have substantially increased the actual availability of

Figure 1
Replication Practices at APSR and AJPS



Proportion of published articles employing statistical analysis in APSR and AJPS that (top row) state that replication files are publicly available, and (bottom row) for which replication files are publicly available. Confidence intervals provide an assessment of whether observed proportions could have come from the same underlying distribution. Results: 30% of the publications at APSR provided replication files, 65% of the publications at AJPS 2011–2013 provided replication files; AJPS’s 2010 policy substantially increased provision of replication files.

replication files, although we were still unable to find replication files for about 35% of the publications in AJPS 2011–2013.

Robustness of Published Results to Replication

The data in figure 1 does not speak to the quality of the replication files that are provided, nor the actual robustness of results. Addressing this question would require a more systematic evaluation of the literature, such as from a Replication Audit (see Online Appendix B).

To offer some preliminary data on this question I surveyed three groups of scholars about their experiences attempting to replicate statistical studies; these groups were students from

my PhD methods class, students from Gary King's PhD methods class, and subscribers to the Political Methodology listserve. These numbers should be interpreted with caution because it is not from a representative sample: respondents selected into the survey, and respondents selected the work they wanted to replicate. See Online Appendix D for more details about this survey.

This data suggests a mixed conclusion. Of those who responded to the reproducibility of the result, about 52% reported that they were "able to precisely reproduce the main results" and only 13% reported that they were "not able to approximately reproduce the main results." This suggests that many results in political science can be, at least superficially, reproduced, but also that many seem only approximately reproducible. Of those who responded to the robustness of the results,

The scientific paper is currently a snapshot of a data landscape. Instead, the scientific paper should become an open safari of the data landscape, from which readers are encouraged to depart at any time and explore the landscape on their own. Providing that level of freedom would allow scholars to explore and build on the data to answer their own questions, and otherwise to be much more active participants in the scientific journey of the authors.

36% reported that "most or all of the key results were robust", 20% that there were "major technical errors though these didn't change the main results", and 56% that results were not robust (responses 5–7). This is encouraging in how many results were found to be robust, while also reinforcing the value of strong transparency norms so that the many fragile results can be more easily uncovered and examined. Also, given that more reproducible and robust work is more likely to share replication files (Piwowar, Day, and Fridsma 2007), these numbers are probably an optimistic appraisal of the reproducibility and robustness of statistical work in political science.

BENEFITS OF SHARING COMPLETE REPLICATION FILES

This section discusses some of the specific benefits of the sharing of complete replication files, which are defined as files that make as explicit and objective as practical every step of research from initial data collection to final statistical output.² These benefits include greater refutability, openness, cumulation, and reducing barriers to entry.

A primary benefit of sharing replication files is that it makes research more refutable, and therefore makes the body of non-refuted findings more informative. Fragile, misleading, and nonreplicable³ statistical analyses can be largely eliminated by the simple requirement that authors be required to submit complete replication files before publication.⁴ Doing so will deter many scholars from publishing unreliable analyses, and the scientific community can be relied on to expose many of those remaining.

Consider the cases of serious fraud that have been uncovered in psychology. Uri Simonsohn used data analysis techniques (Shea 2012; Simonsohn 2012) to detect suspicious data

patterns in psychology and has uncovered incidents of fraudulent data which has led to multiple retractions of articles and the resignation of prominent tenured professors. One psychologist was found guilty of fabricating data for more than 50 publications; this professor did not just "massage" the data, or report only convenient analyses, he literally made the data up and then gave it to his students to analyze for their dissertations (Bhattacharjee 2013). Simonsohn (2012) argues that "requiring authors to post the raw data" will "make fraud much less likely to go undetected."

Fraudulent science probably only makes up a tiny proportion of the scientific output. However, it threatens to dramatically reduce the public credibility of science. Of potentially greater concern (Stevenson and Wolfers 2013) is the unknown proportion of fragile, misleading, or nonreproducible results.

The data reported earlier in this article suggests this proportion is not trivially small in political science.

The sharing of complete replication files has the additional benefit that it opens up scientific research to the questions, insights, and exploration of others. Rather than confine reviewers and readers to a snapshot of the data that has been carefully prepared by the authors, a reader of an empirical analysis with a question or insight could immediately go to the data to evaluate it. Readers might want to double check that primary data has been merged correctly, evaluate how known coding issues have been addressed, identify influential observations, examine particular subsets of the data, implement an alternative conditioning strategy or estimator, or plot the data in potentially insightful ways. Of course, as with any analysis, any findings from such a reanalysis should follow from a principled and well-motivated empirical strategy, and scholars should be wary of the biases from multiple-comparisons (data dredging).

The full set of questions that a reader might have is vast, and cannot, even in principle, be answered in the text of a paper or supplementary materials. However, authors need not anticipate every possible question. Instead, they need only make a sufficient effort to warrant publication, and provide clear and complete replication files so that the scientific community is able to evaluate and build on their work. The scientific paper is currently a snapshot of a data landscape. Instead, the scientific paper should become an open safari of the data landscape, from which readers are encouraged to depart at any time and explore the landscape on their own. Providing that level of freedom would allow scholars to explore and build on the data to answer their own questions, and otherwise to be much more active participants in the scientific journey of the authors.

Scientific productivity exhibits network externalities. The cumulation of data and tools provides substantial benefits beyond the intentions of the creators. A scholar will collect some data for some specific research purpose; however, often those data can be used to answer other questions. This is especially the case for those kinds of observational data, such as cross-national data, that are relevant to multiple research programs. For example, the statistical study of international relations has benefited greatly from the creation of large standardized datasets based on the cumulated work of hundreds of scholars. These externalities are also present in experimental research. Experimental manipulations can be “reused” for down-stream experiments in which other causal effects are investigated.

Sharing complete replication files is also likely to incentivize scholars to be more careful (e.g., see Andrew Gelman here): we face trade-offs in how we invest our time and we are likely to invest more effort in those stages of our research that are

to broadly promote transparency norms such as is articulated in the most recent *Guide to Professional Ethics in Political Science* (APSA 2012): “openness is an indispensable element of credible research and rigorous analysis, and hence essential to both making and demonstrating scientific progress.” Specifically, I recommend the following transparency maxim for statistical and computational work.

Transparency Maxim: Good research involves publishing complete replication files, making every step of research as explicit and reproducible as is practical.

The transparency maxim is both a descriptive statement that good research tends to publish replication files, and a normative statement that the publishing of replication files is a necessary component of good research.

The transparency maxim is likely to be partly self-enforcing. Researchers are more likely to publish complete replication files as they are more technically proficient, more concerned

Transparency Maxim: Good research involves publishing complete replication files, making every step of research as explicit and reproducible as is practical.

most subject to scrutiny. By making more of the research process subject to scrutiny, scholars will have greater incentives to be cautious with those parts of the research process. This incentivizing effect may be one of the primary benefits of stronger replication norms.

Sharing code for analysis and presentation lowers the barrier to entry for students and others, and promotes the dissemination of useful techniques. Students especially benefit from having access to replication code because it allows them to see precisely how prominent scholars execute their empirical analyses and provides an opportunity for junior scholars to contribute to the research frontier (Rich 2013). Instead of needing to “study under” leading scholars to learn their statistical methods, scholars will be able to learn by working through replication files.

Papers that share their replication data and code have greater visibility and more citations (Gleditsch, Metelits, and Strand 2003; Piwowar, Day, and Fridsma 2007). This is probably partly a selection effect, but also probably partly a causal effect. It is much harder to build off of a study for which replication files are not available. In addition, the sharing of replication files provides a public signal about the quality, confidence, and professionalism of a scholar.

Transparent replication practices are a scientific public good: the benefits are large and shared by many, the costs are small but born largely by the authors. While the benefits vastly outweigh the costs, transparency will likely be underprovided unless individual’s incentives are aligned with the group’s. Strengthening of formal incentives could help, such as if journals, universities, and funders insist that replication files be publicly posted as part of the publication, promotion, and funding process. Ultimately, however, scientific practice follows scientific norms. To incentivize adequate transparency we need

about the quality of their work, more confident in their work, more concerned with the scientific enterprise, and more concerned with being perceived as producing good research. These motives generate a correlation between replication files and good research. For example, Wicherts, Bakker, and Molenaar (2011) report that willingness to share data is positively associated with the strength of the evidence and the quality of the reporting of statistical results. Similarly, some scholars adopt judgmental heuristics based on the availability of replication files (e.g., [here](#)).

As this descriptive association becomes stronger, publishing replication files will send a positive signal about the quality of one’s research (or the failure to publish replication files will send a negative signal). Low quality research cannot easily “fake” this signal because the very act of publishing replication files makes it much easier to evaluate the quality of the research. This signal will then encourage scholars, journals, and universities that produce good research and wish to be perceived as such to publish replication files. Publishing replication files is thus an informative signal of the quality of one’s research.

RECOMMENDATIONS

The following are some specific recommendations about how to produce good replication files for researchers engaged in statistical analysis. This advice is similarly applicable to scholars engaged in computational theory. Some advice is also offered here for journals, universities, and funders about how best to promote these practices. The American Political Science Association has also recently revised its *Guide to Professional Ethics, Rights and Freedoms* to emphasize and clarify researchers’ “ethical obligation to facilitate the evaluation of their evidence-based knowledge claims through data access,

production transparency, and analytic transparency so that their work can be tested or replicated” and offers additional guidance on these topics for quantitative and qualitative researchers. For advice for qualitative researchers, see Moravcsik (2010). Davenport and Moore (2013), Hook et al. (2010), and ICPSR offer advice on data preparation and archiving.

Recommendations for Statistical Studies

1. Do all data preparation and analysis in code. Even if analysis is done by “clicking and pointing,” most statistics programs (such as Stata and SPSS) produce the code required to replicate each step.
2. Adopt best practice for coding. Some recommendations to keep in mind are:
 - Use comments and functions to make your code clear. Keep your code clean and clear. Comment liberally to remind yourself and communicate to others what your code is (supposed to be) doing. Use functions to execute specific commands, especially when these commands are repeatedly used.
 - Test your code. Build in routine tests to make sure that your code is doing what you think it is doing. Execute the same procedure in multiple redundant ways to reduce

many files to the final reported output. In our case, this chain of custody should be fully documented and easily reproducible.

4. Fully describe your variables. Somewhere—in variable labels, a codebook, paper, or comments in the code—the meaning of variables needs to be clearly communicated. The original sources or coding rules for variables should be provided. A reader should be able to trace a variable back to its original creation, and the author who first created a variable should clearly document the rules by which the variable was constructed. It is unacceptable to share data files for which it is unclear precisely where a variable came from, let alone one in which the variables names are an indecipherable character string.
5. Document every empirical claim. Every empirical claim in a paper based on the data should be explicitly produced somewhere in the replication code. This includes all graphics and tables, but also any in-text reference to some feature of the data. It should be easy for others to link empirical claims in the paper to the relevant portion of the code. One way to do this is to include a quote in the replication code of the sentence in the paper in which the empirical claim is made; this way a reader can search for the text of the empir-

Journals are the key site for improving replication practices.

the risks of a mistake. It is not uncommon for results to be driven by a misimplemented routine such as the mishandling of missing data.

- Run your final code all the way through from scratch. Before finalizing the paper, the entire replication code should be rerun from beginning to end. Make sure you set a seed, and make a log file (in Stata) or use something like knitr (in R), so that this final run is recorded and fully reproducible. Also make sure that all relevant data files are included in the replication files. One way to ensure this is to only call data files from within the folder where the code is stored, and then to upload the entire replication folder for archiving.
 - For a helpful discussion of strategies to improve replication practices and code, see Bowers (2011), Appendix A of Shalizi (2013), and Gandrud (2013).
3. Build all analysis from primary data files. Download data files from the original source, and include a precise reference in your code or paper to this original source. Lock these primary files to prevent accidental changes to them. This way any errors that occur will take place in your code, which can be diagnosed and corrected. Share these original data files along with all other files in your replication files. What good is sharing a final data file and replication code, as many scholars currently do, if the crucial decisions and errors were made earlier in the merging and cleaning of the data? We should take the attitude toward a statistical estimate as legal courts take toward evidence: there should be a clearly documented “chain of custody” from trusted pri-

tical claim in the replication code. Another strategy is to have one’s replication code follow the structure of the paper.

6. Archive your files. Upload this finalized set of analysis code and data files to a reliable third-party site such as ICPSR or Dataverse.
7. Encourage coauthors to adopt these standards. Maintaining good replication practices is more difficult when one’s collaborators do not have the same replication practices. Scholars may be understandably reluctant to impose strict requirements on their colleagues. I recommend sharing with them papers such as this one that outline good replication practices and the reasons behind them, and otherwise leading by example.

Recommendations for Journals

Journals are the key site for improving replication practices. I describe here a set of policies that journals could readily adopt. This package of policies was crafted to have minimal cost to the journal, to maintain the status quo with respect to when data and code should be shared with the public (at time of publication), and to maximally improve replication practices. Adopting these recommendations will improve the quality of work being published, will signal that the journal has higher standards, and will likely increase the prominence of the journal (Gherghina and Katsanidou 2013).

1. Require complete replication files before acceptance. The simplest policy is a nominal requirement that the authors make replication files available. For example, *AJPS* requires

that “the first footnote explicitly states where the data used in the study can be obtained for purposes of replication.” As is suggested by the data in figure 1, this change made a substantial difference for *AJPS* (compare years before and after 2010). However, despite that nearly 100% of *AJPS* articles now explicitly state the location for replication files, less than 75% actually provide replication files. Others have similarly noted the limits of requiring authors to agree to or sign statements of intent to share data (Wicherts et al. 2006; Savage and Vickers 2009). As such, to achieve sufficiently high compliance, journals need to actually ensure that replication files are posted or to post it themselves. For example, the *Quarterly Journal of Political Science* ensures that replication files are available and that it is possible to replicate the results before an article is published. This process of checking or posting replication files can be partly automated in the workflow programs used by journals. Journals may want to follow the journal *Biostatistics* by indicating on the first page of a publication (with a “D”, “C”, and “R”) whether data and/or code is available, and whether the publication has passed a “reproducibility review” (Peng 2011). Journals may also want to adopt the policy at *Nature* of requiring an “accession number” or URL for the replication files at the time of first submission; the files can then be released at the time of acceptance.

2. Encourage high standards for replication files. The journal should articulate its expectations about the quality of replication files to authors. Ideally the journal will encourage high standards, such as those articulated above. Replication files could be made available to reviewers after a revise and resubmit decision, allowing reviewers the option to include the quality of replication files in their assessment of the publication.⁵
3. Implement a replication audit. A replication audit involves assembling a replication team of trusted researchers to evaluate the reproducibility and robustness of a random subset of publications from the journal. By guaranteeing regular space in the journal for the replication audit the journal (1) helps reward the act of evaluating the reproducibility and robustness of published work, (2) incentivizes authors to invest additional effort to make sure that their results are reproducible, robust, and that their inferences are not misleading, and (3) provides a diagnostic of the (hopefully improving) quality of empirical work in the journal. The replication audit is described in Online Appendix B. A replication audit is preferred to the exclusive publishing of replication articles on a case-by-case basis because the latter process is more susceptible to publication bias that will over-represent “interesting” replications that claim to overturn earlier studies.
4. Retract publications with nonreproducible analyses.⁶ If an analysis cannot be reproduced, even by the original authors when given ample opportunity to do so, the results from the study cannot be trusted and the study should no longer be a part of the public scientific record. Publications based on nonreproducible analyses should, therefore, be retracted. David Laitin, in a personal communication, recommended the adoption of a standard retraction procedure for politi-

cal science. Doing so would make the retraction policy more transparent, remove editorial discretion, and insulate editors from legal retaliation. The standards for such a retraction policy could be set so that only the most egregious cases of nonreproducible analyses are retracted: where the finding cannot be approximately reproduced by the original authors. Even if rarely activated, a retraction policy would promote replication practices by establishing the norm that authors are responsible for providing adequate replication files and by providing strong sanctions against the worst kinds of nonreproducibility.

Recommendations for Universities

1. Universities can provide institutional support for producing complete replication files. For example, the Institution for Social and Policy Studies (ISPS) at Yale University provides a service in which they help produce and publish complete replication files in both R and Stata for ISPS funded research (see here). Harvard’s Institute for Quantitative Social Science has built the Dataverse Network Project which provides long-term archiving of replication files and other services.
2. Universities could encourage and expect high replication standards from their students and scholars. Students should be encouraged to submit replication files for course papers. Departments could have a policy recommending publication of complete replication files for all published work.
3. Norms of scholarly evaluation could place more emphasis on transparency and specifically the provision of replication files.

Recommendations for Funders

Require recipients of funding to commit to transparency and specifically to publish replication files. The National Science Foundation, for example, now requires a data management plan as part of any proposal, although of course this is not sufficient (for a satire of one, see [here](#)).

CONCLUSION

The study of politics rightly aspires to be scientific: it aims to establish generalizable causal insights from the nonsubjective, replicable, and transparent empirical evaluation of precise and logical theories (Gerring 2011, 11). Relative to many natural sciences, however, political science faces daunting methodological challenges. We are less able to rely exclusively on experiments to resolve many of our questions. We can rarely isolate mechanisms and processes in a controlled setting. Our subjects of study rarely follow simple mathematical patterns, are highly context dependent, and adapt to our interventions and theories. However, political science has better replication practices than many of the natural and social sciences. We should be proud of this and continue to lead the way.

ACKNOWLEDGMENTS

Replication Files available at <http://hdl.handle.net/1902.1/22160>. Online appendix available at: <http://ssrn.com/abstract=2318223>. For input I thank Natalia Bueno, Mats

Hammerström, Nicole Janz, Gary King, David Laitin, Daniel Masterson, Gwyneth McClendon, Guadalupe Tuñón, Rick Wilson, Baobao Zhang, Magnus Öberg, and especially the editorial committee at the *Journal of Peace Research*, John Bullock, Don Green, and Arthur Lupia. ■

NOTES

1. Data was collected by Guadalupe Tuñón, Peter Repucci, and myself.
2. See text in the section Recommendations for Statistical Studies for more discussion of what constitutes “complete replication files.”
3. We might distinguish between the “replication of a study” in which the research design is replicated on new data (also called “broad replication”), and the much less informative but nonetheless important “replication of an analysis” (or “narrow replication”) in which the analysis is replicated on the same data. Other fields such as computational science distinguish between “reproducibility” (replication of analysis) vs “replication” (replication of a study); I use the term “reproducible” to refer specifically to research for which the analysis is replicable. Sharing of replication files foremost promotes reproducible research (replications of analysis), though it might also promote replications of studies if the greater transparency facilitates the execution of the study on a new sample.
4. This article focuses on replication practices for statistical empirical work. Qualitative scholarship would also be much improved by the adoption of stronger practices of data transparency (see Moravcsik 2010).
5. Another option is to require that reviewers have access to replication files from first submission. However, a number of scholars have expressed concern over this proposal because they are not comfortable with others having access to replication materials so much in advance of publication.
6. I thank David Laitin for raising this idea.

REFERENCES

- APSA. 2012. “A Guide to Professional Ethics in Political Science.” URL: <http://www.apsanet.org/content86135.cfm>.
- Bhattacharjee, Yudhijit. 2013. “The Mind of a Con Man.” *New York Times Magazine*, April 26. URL: <http://www.nytimes.com/2013/04/28/magazine/diederik-stapels-audacious-academic-fraud.html?pagewanted=all&r=1&>.
- Bowers, Jake. 2011. “Six Steps to a Better Relationship with Your Future Self.” *Political Methodologist* 18 (2): 9.
- Bueno de Mesquita, Bruce, Nils Petter Gleditsch, Patrick James, Gary King, Claire Metelits, James Lee Ray, Bruce Russett, Håvard Strand, and Brandon Valeriano. 2003. “Symposium on Replication in International Studies Research.” *International Studies Perspective* 4 (1): 72–107. URL: <http://onlinelibrary.wiley.com/doi/10.1111/1528-3577.04105/full>.
- Davenport, Christian, and Will H. Moore. 2013. “Conflict Consortium Data Standards & Practices.” Manuscript.
- Gandrud, C. 2013. *Reproducible Research with R and Rstudio*. Chapman & Hall/CRC.
- Gerring, John. 2011. *Social Science Methodology: A Unified Framework*. New York: Cambridge University Press.
- Gherghina, Sergiu, and Alexia Katsanidou. 2013. “Data Availability in Political Science Journals.” *European Political Science* 12: 333–49.
- Gleditsch, Nils Petter, Claire Metelits, and Håvard Strand. 2003. “Posting Your Data: Will You Be Scooped or Will You Be Famous?” *International Studies Perspectives* 4 (1): 9–95. URL: <http://gking.harvard.edu/files/replvdc.pdf>.
- Herndon, Thomas, Michael Ash, and Robert Pollin. 2013. “Does High Public Debt Consistently Stifle Economic Growth? A Critique of Reinhart and Rogoff.” Political Economy Research Institute Working Paper Series (322).
- Hook, Les A., Suresh K. Santhana Vanna, Tammy W. Beaty, Robert B. Cook, and Bruce E. Wilson. 2010. “Best Practices for Preparing Environmental Data Sets to Share and Archive.” URL: daac.ornl.gov/PI/BestPractices-2010.pdf. An earlier version of this was published in: 2001. *Bulletin of the Ecological Society of America* 82 (2): 138–41.
- King, Gary. 1995. “Replication, Replication.” *PS: Political Science and Politics* 28 (3): 444–52.
- Krugman, Paul. 2013. “The Excel Depression.” *New York Times*. URL: <http://www.nytimes.com/2013/04/19/opinion/krugman-the-excel-depression.html>.
- Moravcsik, Andrew. 2010. “Active Citation: A Precondition for Replicable Qualitative Research.” *PS: Political Science and Politics* 43 (1): 29–35.
- Peng, Roger D. 2011. “Reproducible Research in Computational Science.” *Science* 334 (6060): 1226–27.
- Piowar, Heather A, Roger S. Day, and Douglas B Fridsma. 2007. “Sharing Detailed Research Data Is Associated with Increased Citation Rate.” *PLoS One* 2 (3): e308.
- Reinhart, Carmen M., and Kenneth S. Rogoff. 2010. “Growth in a Time of Debt.” *American Economic Review: Papers & Proceedings* 100 (2): 573–78.
- Rich, Timothy S. 2013. “Publishing as a Graduate Student: A Quick and (Hopefully) Painless Guide to Establishing Yourself as a Scholar.” *PS: Political Science & Politics* 46 (2): 376–79.
- Savage, C. J., and A. J. Vickers. 2009. “Empirical study of data sharing by authors publishing in PLoS journals.” *PLoS One* 4 (9): e7078.
- Shalizi, Cosma Rohilla. 2013. “Advanced Data Analysis from an Elementary Point of View.” URL: <http://www.stat.cmu.edu/~cshalizi/ADAFaEPoV/>.
- Shea, Christopher. 2012. “The Data Vigilante.” *The Atlantic*, December. URL: <http://www.theatlantic.com/magazine/archive/2012/12/the-data-vigilante/309172/>.
- Simonsohn, Uri. 2012. “Just Post It: The Lesson from Two Cases of Fabricated Data Detected by Statistics Alone.” *Psychological Science* 24: 1875–88.
- Stevenson, Betsey, and Justin Wolfers. 2013. “Six Ways to Separate Lies from Statistics.” Bloomberg. URL: <http://www.bloomberg.com/news/2013-05-10/six-ways-to-separate-lies-from-statistics.html>.
- Wicherts, Jelte M., Denny Borsboom, Judith Katsand, and Dylan Molenaar. 2006. “The Poor Availability of Psychological Research Data for Reanalysis.” *American Psychologist* 61 (7): 726–28.
- Wicherts, Jelte M., Marjan Bakker, and Dylan Molenaar. 2011. “Willingness to Share Research Data Is Related to the Strength of the Evidence and the Quality of Statistical Results.” *PLoS One* 6 (11): e26828.