

Active Maintenance: A Proposal for the Long-Term Computational Reproducibility of Scientific Results

Limor Peer, *Yale University*

Lilla V. Orr, *Yale University*

Alexander Coppock, *Yale University*

ABSTRACT

Computational reproducibility, or the ability to reproduce analytic results of a scientific study on the basis of publicly available code and data, is a shared goal of many researchers, journals, and scientific communities. Researchers in many disciplines including political science have made strides toward realizing that goal. A new challenge, however, has arisen. Code too often becomes obsolete within only a few years. We document this problem with a random sample of studies posted to the Institution for Social and Policy Studies (ISPS) Data Archive; we encountered nontrivial errors in seven of 20 studies. In line with similar proposals for the long-term maintenance of data and commercial software, we propose that researchers dedicated to computational reproducibility should have a plan in place for “active maintenance” of their analysis code. We offer concrete suggestions for how data archives, journals, and research communities could encourage and reward the active maintenance of scientific code and data.

A study is computationally reproducible if a sufficiently savvy third party could download the code and data from a public repository and successfully execute the analysis with minimal difficulty. Initiatives in political science such as DA-RT (Elman, Kapiszewski, and Lupia 2018; Lupia and Elman 2014) reflect a growing consensus that computational reproducibility is a shared goal—at least for research that relies on the formal analysis of quantitative or qualitative data that can be safely made public (Alvarez, Key, and Núñez 2018; Christensen, Freese, and Miguel 2019; Dafoe 2014; Rohlfing et al. 2020).

Although “the scientific enterprise depends on the ability of the scientific community to scrutinize scientific claims” (National Academies of Sciences, Engineering, and Medicine 2019, 6), evaluating published scientific results can be difficult because readers must infer from written descriptions of statistical procedures what

the authors must have actually done. Even when those descriptions are accurate, they are necessarily an approximation of what occurred. The move toward computational reproducibility as a precondition of publication has been successful in revealing analysis errors and ensuring that shared materials are given persistent links (Gertler and Bullock 2017).

We consider what happens *after* materials have been deposited in public archives. A study that was computationally reproducible on the day of deposit has already cleared a high bar. In addition to rigorous scrutiny of the manuscript, the code and data—or “replication archive,” as it is commonly referred to in political science—were demonstrated by a trusted party to reproduce the numerical results reported in the paper. However, computational reproducibility is dynamic and sometimes elusive. Code that successfully executes on the day of deposit may break a year, a decade, or a generation later. Operating systems change and software packages are updated or not maintained at all. Data used for analysis may become unintelligible if the software used to create the data cannot be run. Workflows that depend on particular online resources may fail if the resources are moved or taken down. The cost of these failures is the potential for loss of knowledge and a diminished ability to use these materials in future settings.

Limor Peer  is associate director for research at the Institution for Social and Policy Studies at Yale University. She can be reached at limor.peer@yale.edu.

Lilla Orr  is a PhD candidate in political science at Yale University. She can be reached at lilla.orr@yale.edu.

Alexander Coppock  is assistant professor of political science at Yale University. He can be reached at alex.coppock@yale.edu.

This article proposes “active maintenance” as a solution to this problem. We suggest that replication archives should be inspected periodically to ensure they run in contemporary computing environments according to a maintenance plan. Active maintenance builds on common practices in the data-preservation (Conway

typically requires that researchers maintain organized data, use version control, and test code according to best practices of reproducible research (Christensen, Freese, and Miguel 2019; Dafoe 2014). Reproducibility is challenged *after* deposit if future users cannot make use or sense of the files.

Code that successfully executes on the day of deposit may break a year, a decade, or a generation later.

2000) and code-development communities (Fowler and Foemmel 2006). Both recognize that supporting materials is an ongoing effort that requires resources, infrastructure, and standards.

Whether active maintenance is appropriate in a given case depends on the relative benefits and costs. The benefits of active maintenance are data and code that continue to provide scholars with tools to understand, critique, and reuse existing scientific knowledge. The costs associated with maintenance, of course, could be substantial. Whether the benefits exceed the costs will ultimately depend on the value of the scientific knowledge contained in the replication archive and on the difficulty of the required maintenance. We argue that studies deemed as important justify an investment in long-term reproducibility. We discuss who among the authors, archives, publishers, scholarly communities, funders, universities, and other stakeholders might take responsibility for the active maintenance of digital materials.

MOTIVATING EXAMPLE: “SHY TRUMP VOTERS”

We begin with a case study that illustrates over-time degradation in computational reproducibility and how active maintenance could address it. Coppock (2017) reported the results of a study that attempted to discern whether some portion of the polling misses in the 2016 US presidential election could be attributed to survey respondents who misreported their support for Donald Trump for fear of being perceived as racist or sexist.

The ISPS Data Archive reviewed the replication materials for Coppock (2017) in July 2017 as part of its routine process (Peer and Green 2012). Archive staff confirmed that all files necessary to replicate the reported results were available. The analysis consisted primarily of descriptive summaries and a comparison of direct and list experimental estimates with bootstrapped standard errors. Two years later, archive staff revisited these materials as part of a training exercise. Newly hired staff members encountered an error that prevented the code from running in full. Archive staff contacted the author, who traced the error to the removal of the bootstrap function from the `broom` package for R (Robinson and Hayes 2019) and then rewrote the bootstrap code using the `rsample` package instead (Kuhn, Chow, and Wickham 2020). ISPS staff then confirmed that results reported in the paper could be computationally reproduced and posted the updated materials.

This case is an example of unplanned active maintenance. Close working relationships between archive staff and ISPS researchers facilitated a simple troubleshooting and update process. Extending the computational reproducibility of this study cost perhaps 10 emails and an hour of debugging.

COMPUTATIONAL REPRODUCIBILITY IS DYNAMIC

Achieving computational reproducibility at the time of deposit is itself a feat. Even before deposit, computational reproducibility

Reproducibility breaks down if datasets are damaged at the bit level or contained in proprietary file formats incompatible with current systems. After a period of time, the original software might be totally unavailable (Peng 2011) or restricted (e.g., closed source), or the ability to use it may be constrained by numerous components and dependencies (Hinsen 2019; Ivie and Thain 2018). Changes to statistical programming languages can ripple through the many add-on modules and packages on which analysts rely. For example, R updated how it generates random samples in 2019 to allow for greater uniformity of uniform random samples (Smith 2019). As a result, analyses that depend on random-number generation will yield slightly different results with a more recent version of R, even if the replication archive included a random-number seed. Furthermore, the interpretability of study documentation and metadata is itself dynamic. Even if the documentation makes sense to a contemporary user, it may be less interpretable to users 50 years in the future—or even to the original researcher a few years after deposit (Bowers 2011). These breakdowns can prevent future researchers from computationally reproducing the original results and from using the materials in a meaningful way.

EVALUATION OF THE LONG-TERM COMPUTATIONAL REPRODUCIBILITY OF 20 STUDIES

To assess long-term computational reproducibility in the ISPS Data Archive, we attempted to execute the analysis files associated with a sample of studies, all of which were computationally reproducible upon deposit. We compiled a list of all 97 archived studies, excluded three books, and then randomly selected 20 studies (Peer, Orr, and Coppock 2021). The sample included studies archived between 2009 and 2019, with from four to 88 files associated with each study. These studies were archived using versions of Stata (18), R (one), and SPSS (one) that were available at the time of deposit.

Between December 2019 and March 2020, archive staff reviewed the program files using Stata 15, R 3.6.1 and R 3.6.2, and SPSS 26. For each study, they reviewed documentation and downloaded data and analysis files. After adjusting working directories and file names, they attempted to run each program file and recorded any errors. Errors were classified based on whether a moderately trained replicator would be able to overcome them without substantially updating the script for the statistical analysis or after receiving help from the study authors (table 1).

Of these 20 studies, 13 could be reproduced with present-day hardware and software. Various challenges to computational reproducibility emerged in the remaining seven studies. In six studies, program files produced errors that could be resolved by seasoned analysts. Necessary adjustments, such as updating

Table 1
Computational Reproducibility in 20 Studies

	Author Provided Code	Archive Added Code
	N = 20	N = 11
No errors	13 (65%)	3 (27%)
Resolvable errors	6 (30%)	8 (74%)
Errors requiring further information	3 (15%)	0
Data-preparation scripts not archived	3 (15%)	0

Stata's `outreg` command to `outreg2`, are trivial for advanced users. In three of the studies, program files produced errors that could not be resolved by present-day users without further information or documentation. For example, variables that now appear to be missing would be supplied or explained most easily by study authors. In three of the 20 studies, raw data and associated cleaning files were not deposited. Because it often is necessary to exclude raw data from replication files, we did not classify the resulting errors as indications of failed computational replication.

For 11 of the studies, archive staff added R files at the time of deposit as part of an effort to convert statistical code to open-source formats. Of the eight instances in which the R programs did not run, several included simple typos. However, emerging challenges to computational reproducibility were generally more difficult to overcome. For example, three studies relied on packages that had been removed from CRAN (e.g., `Design`) or updated in ways that made it impossible to run the original syntax (e.g., `Zelig`).

In summary, we found that even in the relatively short 10-year period, replication materials that originally allowed for computational reproducibility broke down. Most of the challenges were

In summary, we found that even in the relatively short 10-year period, replication materials that originally allowed for computational reproducibility broke down.

encountered in the older studies. Sophisticated users would have little difficulty resolving many of the errors that appeared; however, even some of the user-resolvable errors would have required substantial troubleshooting and recoding.

TECHNOLOGY-BASED APPROACHES

A valuable if not wholly satisfying solution to the degradation over time of computational reproducibility is containerization. The replication archive and the computing environment as it existed at the moment of deposit (i.e., a “snapshot”) are placed in a virtual container guaranteeing that even if the code is run on different machines, the same results will be reproduced. Emulation—that is, running legacy software on modern computers—is a procedurally distinct approach that similarly allows users to execute code on a host system. Both approaches have the principal virtue of dramatically increasing the probability that old code will run in the future.

Nevertheless, these approaches face inherent limitations when the goal is for future users to make use or sense of the files. As Katz (2017) observed, containers “provide bitwise reproducibility but aren’t scientifically useful because, as black boxes, you can’t really remix the contents.” For example, consider Brady and Ansolabehere’s (1989) analysis of voter preferences, which was conducted in FORTRAN. Containerization was not available at the time, but if it had been, in principle we would be able to reproduce the analysis within the confines of the 1989 time capsule. However, extending the analysis to modern data or incorporating later methodological advances would not be possible.

Technology-based approaches also face a coordination problem. For example, containerization and packaging tools (e.g., Research Compendium, Reprozip, and WholeTale) are proliferating but, importantly, any such tool will be most reliable when it is popular. If a technological solution falls out of favor in the scientific community or fails to gain traction in the first place, it is not likely to be maintained. With no training or support networks for users in the somewhat distant future, reopening containers or accessing any technological tool may become increasingly difficult over time.

ACTIVE MAINTENANCE

Our proposed solution to the problem of degradation over time in computational reproducibility is active maintenance. Unlike static images, these strategies afford more robust options for future reuse. Active maintenance requires monitoring materials and deploying a combination of strategies drawn from two fields: data preservation and code development. Recent recommendations on how to sustain content over time parallel our active-maintenance suggestion (Daigle et al. 2018).

Our thinking about active maintenance is influenced by digital archivists’ approaches to preserving the usability of digital materials. “Active preservation” is a curatorial responsibility and relies on strategies including copying and migration. Copying entails saving a copy of the data in a system-independent, nonproprietary

format that is likely to be machine readable in future decades. This procedure has advantages but some information may be lost in translation. Migration goes further and “includes refreshing the media but also addresses the internal structure of the files so that the information within can be read on subsequent computer platforms, operating systems, and software” (Green, Dionne, and Dennis 1999). Because code is instrumental for computational reproducibility, more archives and repositories are applying these strategies to software (Chassanoff et al. 2018).

Software developers are keenly aware of how changes to the computational environment affect the performance of source code. Software-testing techniques such as continuous integration involve regularly recompiling software packages to verify that the software runs without errors. Key principles include frequent automatic testing, documenting specifications and dependencies, and saving all components in one location. Adopting code-testing principles makes science more robust (Kraczyk et al. 2019).

Committing to automatically rerunning computational analysis whenever relevant changes are made can increase the probability that the code remains available and usable (Beaulieu-Jones and Greene 2017).

DISCUSSION

Computational reproducibility is a special type of knowledge that offers future analysts a way to scrutinize and understand scientific studies more intimately and tangibly than a journal article can provide. However, computational reproducibility is time varying: we showed in the case of the ISPS Data Archive how reproducibility degraded over time. As political science and other fields develop shared values and norms around computational reproducibility, we must recognize that our commitment to this knowledge should not end on the day we deposit the study in the archive.

As political science and other fields develop shared values and norms around computational reproducibility, we must recognize that our commitment to this knowledge should not end on the day we deposit the study in the archive.

We propose active maintenance as a way to enhance long-term computational reproducibility. Active maintenance requires effort, which raises two important questions: Is the effort worth it and who should be making the effort?

We think that the labor required to open up replication archives every so often to determine whether they still run will differ from study to study. As a general rule, “clean,” well-documented code with proper metadata will be easier to maintain and should be a matter of routine. Whether the substantial effort required to update a complicated replication archive is worth it will depend on the importance of the knowledge contained in it as well as various communities’ ideas about how the materials should be used in the future. In our view, important replication archives should be actively maintained, but we refrain from passing judgment on which studies those are.

Who should be making these efforts depends on collective values, priorities, and incentives. Research teams could take the responsibility for enhancing reproducibility. Under this model, the onus is on authors to adhere to “best practice” guidelines and to stay current with advances in contemporary statistical computing. Currently, authors do not face incentives that encourage this type of sustained investment, although their incentives are—at least in part—a function of disciplinary norms that could change. Under an alternative model, active maintenance could be crowdsourced to the wider academic community. Much like user contributions to software made on GitHub, we can imagine the intended audience of the replication archives (i.e., future scholars) contributing up-to-date analysis scripts. Many courses have replication-project assignments. Updating analysis code and uploading it to the data repository could be an apt capstone to such an assignment. This model faces the obvious difficulty that it relies on the goodwill of the unspecified “crowd” to contribute to knowledge, not to mention the cooperation of authors and archives to review their revisions. Under a third model, independent third parties such as the Odum Institute for Research in Social Science or CASCAD could provide active-maintenance services on behalf of journals, universities, and scholarly associations.

Currently, these services certify computational reproducibility at the time of deposit, although we could imagine expanding their range of responsibilities. Finally, the responsibility for active maintenance could fall to archive or repository staff. Data archives are well positioned to provide services verifying computational reproducibility as part of standard curation practice (Peer, Green, and Stephenson 2014). Archives staff can conduct an initial review, incorporate tools, and enforce policies conducive to active maintenance—for example, requiring particular file formats or specific information in a *readme* file at the time of deposit. Many archives already conduct automated bit-level checks and potentially could expand to perform routine reproducibility checks.

In light of the dynamic nature of computational reproducibility, we urge the political science community to recognize that supporting data and code is an ongoing effort that requires

resources, infrastructure, standards, and policies. Studies that were computationally reproducible on the day they were archived often cannot be reproduced on modern software and hardware only a few years later. For some of these studies, we think computational reproducibility is worth actively maintaining.

ACKNOWLEDGMENTS

We thank Jamie Druckman, Vicky Steeves, and Ann Green for helpful feedback on previous drafts and Yuntian Xia and Brian Shih for excellent research assistance.

DATA AVAILABILITY STATEMENT

Replication materials are available on Harvard Dataverse at DOI:10.7910/DVN/JLLFGK. ■

REFERENCES

- Alvarez, R. Michael, Ellen M. Key, and Lucas Núñez. 2018. “Research Replication: Practical Considerations.” *PS: Political Science & Politics* 51 (2): 422–26.
- Beaulieu-Jones, Brett K., and Casey S. Greene. 2017. “Reproducibility of Computational Workflows Is Automated Using Continuous Analysis.” *Nature Biotechnology* 35 (4): 342–46.
- Bowers, Jake. 2011. “Six Steps to a Better Relationship with Your Future Self.” https://cpb-us-e1.wpmucdn.com/blogs.rice.edu/dist/d/2418/files/2013/09/tpm{_}v18{_}n2.pdf.
- Brady, Henry E., and Stephen Ansolabehere. 1989. “The Nature of Utility Functions in Mass Publics.” *American Political Science Review* 83 (1): 143–63.
- Chassanoff, Alexandra, John Borghi, Yasmin AlNoamany, and Katherine Thornton. 2018. “Software Curation in Research Libraries: Practice and Promise.” *Journal of Librarianship and Scholarly Communication* 1 (eP2239). <https://doi.org/10.7710/2162-3309.2239>.
- Christensen, Garret, Jeremy Freese, and Edward Miguel. 2019. *Transparent and Reproducible Social Science Research: How to Do Open Science*. Oakland: University of California Press.
- Conway, Paul. 2000. “Overview: Rationale for Digitization and Preservation.” In *Handbook for Digital Projects*, ed. Maxine K. Sitts, 5–20. Andover, MA: Northeast Document Conservation Center.
- Coppock, Alexander. 2017. “Did Shy Trump Supporters Bias the 2016 Polls? Evidence from a Nationally Representative List Experiment.” *Statistics, Politics and Policy* 8 (1): 29–40.
- Dafoe, Allan. 2014. “Science Deserves Better: The Imperative to Share Complete Replication Files.” *PS: Political Science & Politics* 47 (1): 60–66.

- Daigle, Bradley, Karen Cariani, Carol Kussmann, Nathan Tallman, and Lauren Work. 2018. "Levels of Digital Preservation." <https://ndsa.org/publications/levels-of-digital-preservation>.
- Elman, Colin, Diana Kapiszewski, and Arthur Lupia. 2018. "Transparent Social Inquiry: Implications for Political Science." *Annual Review of Political Science* 21 (1): 29–47.
- Fowler, Martin, and Matthew Foemmel. 2006. "Continuous Integration." <https://martinfowler.com/articles/continuousintegration.html>.
- Gertler, Aaron L., and John G. Bullock. 2017. "Reference Rot: An Emerging Threat to Transparency in Political Science." *PS: Political Science & Politics* 50 (1): 166.
- Green, Ann, JoAnn Dionne, and Martin Dennis. 1999. *Preserving the Whole: A Two-Track Approach to Rescuing Social Science Data and Metadata*. Washington, DC: Digital Library Federation.
- Hinsen, Konrad. 2019. "Dealing with Software Collapse." *Computing in Science & Engineering* 21 (3): 104–108.
- Ivie, Peter, and Douglas Thain. 2018. "Reproducibility in Scientific Computing." *ACM Computing Surveys* 51 (3): 1–36. DOI:10.1145/3186266.
- Katz, Daniel. 2017. "Is Software Reproducibility Possible and Practical?" <https://danielskatzblog.wordpress.com/2017/02/07/is-software-reproducibility-possible-and-practical>.
- Krafczyk, Matthew, August Shi, Adhithya Bhaskar, Darko Marinov, and Victoria Stodden. 2019. "Scientific Tests and Continuous Integration Strategies to Enhance Reproducibility in the Scientific Software Context." In *Proceedings of the 2nd International Workshop on Practical Reproducible Evaluation of Computer Systems*, 23–28. P-Recs '19. New York: Association for Computing Machinery.
- Kuhn, Max, Fanny Chow, and Hadley Wickham. 2020. *rsample: General Resampling Infrastructure*. <https://cran.r-project.org/package=rsample>.
- Lupia, Arthur, and Colin Elman. 2014. "Openness in Political Science: Data Access and Research Transparency: Introduction." *PS: Political Science & Politics* 47 (1): 19–42.
- National Academies of Sciences, Engineering, and Medicine. 2019. *Reproducibility and Replicability in Science*. Washington, DC: The National Academies Press.
- Peer, Limor, and Ann Green. 2012. "Building an Open Data Repository for a Specialized Research Community: Process, Challenges and Lessons." *International Journal of Digital Curation* 7 (1): 151–62.
- Peer, Limor, Ann Green, and Elizabeth Stephenson. 2014. "Committing to Data Quality Review." *International Journal of Digital Curation* 9 (1): 263–91.
- Peer, Limor, Lilla Orr, and Alexander Coppock. 2021. "Replication Data for: Active Maintenance: A Proposal for the Long-Term Computational Reproducibility of Scientific Results." Harvard Dataverse. DOI:10.7910/DVN/JLLFGK.
- Peng, Roger D. 2011. "Reproducible Research in Computational Science." *Science* 334 (6060): 1226–27.
- Robinson, David, and Alex Hayes. 2019. *broom: Convert Statistical Analysis Objects into Tidy Tibbles*. <https://cran.r-project.org/package=broom>.
- Rohlfing, Ingo, Lea Königshofen, Susanne Krenzer, Jan Schwalbach, and Ayjeren Bekmuratovna R. 2020. "A Reproduction Analysis of 106 Articles Using Qualitative Comparative Analysis, 2016–2018." *PS: Political Science & Politics. First View* 1–5. <https://doi.org/10.1017/S1049096520001717>.
- Smith, David. 2019. "What's New in R 3.6.0." <https://blog.revolutionanalytics.com/2019/05/whats-new-in-r-360.html>.