

Harvard Data Science Review • Issue 6.3, Summer 2024

In-Class Data Analysis Replications: Teaching Students While Testing Science

Kristina Gligorić¹ Tiziano Piccardi¹ Jake M. Hofman² Robert West³

¹Stanford University, Stanford, California, United States of America,

²Microsoft Research, New York, New York, United States of America, ³EPFL, Lausanne, Switzerland

The MIT Press

Published on: Jun 06, 2024

DOI: <https://doi.org/10.1162/99608f92.f9720d1f>

License: [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

ABSTRACT

Science is facing a reproducibility crisis. Overcoming it would require concerted efforts to replicate prior studies, but the incentives for researchers are currently weak, as replicating prior studies requires considerable time and effort without providing the same level of recognition as *de novo* research. Previous work has proposed incorporating data analysis replications into classrooms as a potential solution. However, despite the potential benefits, it is unclear what the involved stakeholders—students, educators, and scientists—should expect from it. What are the costs and benefits? And how can this solution help benchmark and improve the state of science?

In the present study, we incorporated data analysis replications in the project component of the CS-401 Applied Data Analysis course (ADA) taught at EPFL (École Polytechnique Fédérale de Lausanne) enrolling $N = 354$ students. First, we report preregistered findings based on surveys administered throughout the course. We find discrepancies between what students expect of data analysis replications and what they experience by doing them, along with changes in expectations about reproducibility. Second, we provide information for educators about how much overhead is needed to incorporate replications into the classroom and identify concerns that replications bring, as compared to more traditional assignments. Finally, we discuss potential implications of in-class data analysis replications for scientific communities, such as insights about replication barriers in scientific work that should be avoided going forward.

Overall, we demonstrate that incorporating replication tasks into a large data science class can increase the reproducibility of scientific work as a by-product of data science instruction.

Keywords: reproducibility, data analysis, education, open science, citizen science

1. Introduction

The low reproducibility rates of scientific publications have raised concerns across a number of fields (Baker, 2016; Open Science Collaboration, 2012, 2015). Although scientific publishing plays a key role in advancing science, the publication process has multiple weaknesses that may influence the validity of conclusions. The focus on novel, confirmatory, and statistically significant results leads to substantial bias in the scientific literature (Thornton & Lee, 2000), in fields ranging from basic (Begley & Ioannidis, 2015) and biomedical (Goodman et al., 2016), to management and organizational sciences (Bergh et al., 2017). This inclination may lead to bad research practices (Bishop, 2019), such as *p*-hacking (Head et al., 2015; Loken & Gelman, 2017), or developing post hoc hypotheses to fit known results (Kerr, 1998).

Recently, Patil et al. (2019) introduced a framework to consider the key components of a scientific study pipeline that tend to vary across studies and disciplines: the intent of a study (including research question, experimental design, and analysis plan) and what was actually performed when conducting the study (when data is collected, analyses are conducted, estimates are made, and conclusions are asserted). Replication challenges exist throughout the entire pipeline, all the way to data analysis, given previously collected and publicly available data. Data analysis replication, in particular, entails different analysts using their independently written data analysis code to reproduce the original estimates and claims, using the same data and the same analysis plan (Hofman et al., 2021). Such data analysis replication corresponds to a computational reproduction based on the original data, but without the original code (National Academies of Sciences, Engineering, and Medicine, 2019).

Significant variation in the results of data analysis replication has been proven difficult to avoid, even when the incentives are well-aligned (Silberzahn et al., 2018). Researchers are increasingly encouraged to share code and materials (Nosek et al., 2018) for other researchers to perform direct data analysis replication, as a way to improve the credibility of the corresponding research findings. However, replicating the data analysis reported in the publications of others requires considerable time and effort, without providing a particularly rewarding outcome, that is a publication, because of a presumed lack of originality (Janz, 2016) and novelty (Open Science Collaboration, 2015). Researchers are thus usually not incentivized to perform data analysis replications. Ultimately, published replications are rare across fields (King, 1995; Lemons et al., 2016; Makel & Plucker, 2014; Perry et al., 2022; Plucker & Makel, 2021) and the incentives are not yet in place to address this issue.

A recent body of work (Hofman et al., 2021; Quintana, 2021) has proposed one step toward a solution: educating undergraduate and graduate students to perform data analysis replications. Universities are well positioned to introduce replications as class assignments in methods training in order to establish a culture of replication (Ball, 2023; Mendez-Carbajo & Dellachiesa, 2023), reproducibility, and critical thinking (Chopik et al., 2018; Janz, 2016; Smith et al., 2021; Stojmenovska et al., 2019). In-class replications have previously been proposed for college-level education (Meng, 2020) and for psychology education (Frank et al., 2024; Hawkins et al., 2018), to understand correlates of replicability (Boyce et al., 2023; Frank & Saxe, 2012). Furthermore, data analysis replication efforts have previously been used for comprehensive meta-analyses (Wagge, Baciú, et al., 2019; Wagge, Brandt, et al., 2019), based on multiple studies rather than on a single replication attempt (Boyce et al., 2023; Perry & See, 2022; Shrout & Rodgers, 2018).

However, despite the postulated advantages of this solution, it is unclear what the involved stakeholders—students, educators, and scientists—should expect from it. First, in terms of *students*, it is unclear, what type of effort does this require on their end? What do students learn from the process, that is, how do their beliefs differ before vs. after engaging in data analysis replication exercises? What outcomes do students expect before the activity, and how do actual outcomes differ from prior expectations?

Second, in terms of *educators*, there are open questions regarding required investments vs. potential advantages over traditional exercises. For instance, what is required on the educator's end to run successful data analysis replications? How can data analysis replications be incorporated into existing large university classes? What should educators expect their students to learn and take away from data analysis replications? How much of the educator's time and effort is in-class replication expected to take, and what challenges might the educator face (Stojmenovska et al., 2019)?

Lastly, in terms of *scientists*, it remains to be determined how this solution can help benchmark and improve the state of science. What are the main sources of error or confusion that students identify? How can these replication barriers in scientific work be avoided going forward?

To provide new insights about the in-class data analysis replication approach, we incorporated data analysis replications in the project component of the *Applied Data Analysis* course (CS-401) taught at EPFL, the Swiss Federal Institute of Technology in Lausanne.

CS-401 class: Background. This course taught the basic techniques, methodologies, and practical skills required to draw meaningful insights from data. The course had the following prerequisites: an introduction to databases course, a course in probability and statistics, or two separate courses that include programming projects. Also, programming skills were required (in class, we mostly used Python). Most students were first-semester students in computer science or data science master's programs (although registration was open to students from other programs who meet the requirements). At the start of the class, a typical student had strong programming skills and was familiar with fundamental concepts related to algorithms, computer systems (e.g., databases), and the fundamentals of probability and statistics.

During the semester, the students learned the methods during lectures and were introduced, in the lab sessions, to the data analysis software tools. In parallel, the students worked on an applied data analysis project. In a regular iteration, for the project component, students proposed and executed meaningful analyses of a real-world data set. These required creativity and the application of the methods and tools encountered in the course. The outcome of this team effort was a project report and a publicly available code repository.

Lastly, at the end of the semester, students took a 3-hour final exam where they completed a data analysis pipeline on a data set they have never worked with before. By the end of the class, a student is typically able to construct a coherent understanding of the techniques and software tools required to design a data science pipeline.

Our approach in integrating data analysis replications into CS-401. As part of the project component of the class, instead of the standard unconstrained data analysis project leveraging a real-world data set, students individually performed data analysis replications.¹ Class setup was otherwise unchanged, except for adjustments necessary to run the replication exercises (cf. Section 5).

Based on a set of surveys conducted over the course of the semester, our main goal was to understand students' expectations about the difficulty of the exercise before performing the replication vs. their impressions of how hard the task actually was, once completed. Through preregistered analyses of survey responses, we pose the following specific research questions about the impact that data analysis replication tasks have on the students. Our guiding research question is: How large are the discrepancies between students' expectations and the reality of data analysis replication, in terms of time investment, perceived difficulty, tasks, and outcomes (RQ1)? Additionally, we explore the following questions: Do the discrepancies (if any) persist in subsequent replication tasks, after the first one is solved (RQ2)? Can students anticipate in what ways peer-reviewed data science papers might be hard to replicate (RQ3)? Finally, are the effects stronger for the same type of data analysis as performed in the replication exercise, or is there an attitude shift for expectations regarding different data analysis types (RQ4)?

Any discrepancies between expectations and reality (RQ1–2) and any changes in expectations about reproducibility in general (RQ3–4) serve as evidence of shifts in students' attitude. Identifying such indicators of behavioral changes is essential to understanding students' experiences of performing data analysis replications.

Overview of study design. The replication activity was performed as part of the graded class project. Replication exercises were conducted individually. The study design is outlined in Figure 1. The study started with a bidding process where students expressed preferences for papers (Step 1). Afterwards, each student focused on one scientific paper, assigned to them by the class instructors. After reading the assigned paper (Step 2), presurveys recorded the individual students' expectations about the time required, the difficulty of replicating findings from data science papers, and about the perceived reproducibility of papers in the field.

Then, students performed the replications (Steps 3 and 4). Replications were performed and reported individually by each student. We specified two figures or tables to replicate, a basic one (replicated in Step 3) and an advanced one (replicated in Step 4). Students then individually recorded their results and working times in postsurveys, which we compared with students' expectations from before they started as expressed in the presurveys. Lastly, students proposed and conducted creative extension projects, which students built on top of the replicated analyses (Step 5) and presented at the end of the class.

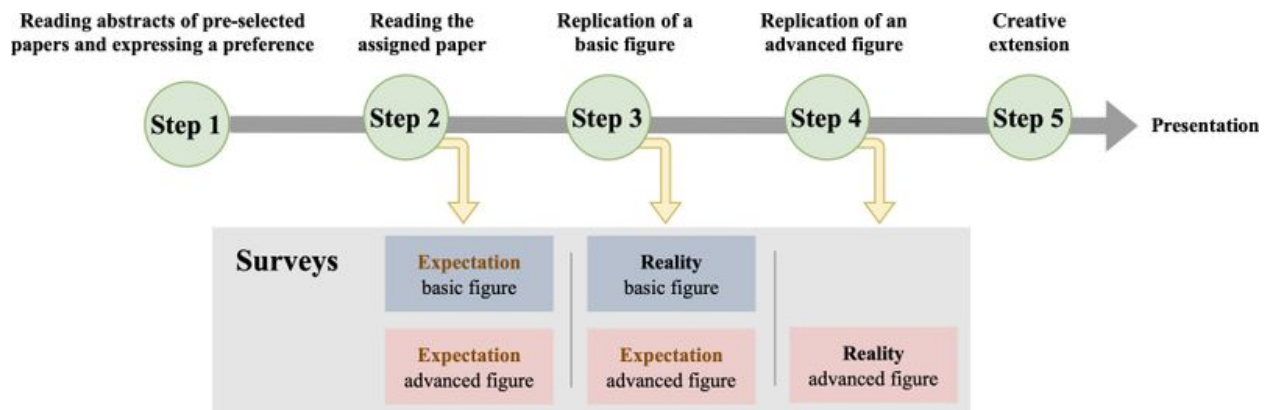


Figure 1. Study design summary. The timeline is visualized from the students' perspective. The semester progresses from the left to the right. The surveys were administered upon submission of the respective assignment step.

Contributions. Concretely, we describe in-class data analysis replication and report ‘lessons learned’ as relevant for students, educators, and scientists. Our findings are based on the work and responses of 354 consenting students who produced data analysis replications of 10 peer-reviewed publications.² Moreover, creative replication extensions performed at the end of the class are contrasted with standard, unconstrained projects, conducted the following year.

Students. In total, 98% of students reported having replicated exactly or qualitatively the basic figure, and 87% the advanced figure. A small fraction of replications failed, and in cases where there were known issues with papers, students correctly identified them. We found that it takes students on average about 10.5 hours to replicate a main result (cf. Section 3.1), and further 8.5 hours to replicate the second result (19 hours in total). Discrepancies between expectations and reality, and changes in expectations about reproducibility arose among students.³ On average, students underestimated the time they would take to reproduce, overestimated how long data wrangling would take, and underestimated how long it would take to iteratively analyze and interpret results (Section 3.1).

The identified attitude shifts signal students’ enhanced appreciation for the challenges involved in the scientific process. Exploratory analyses of open-text responses (cf. Section 3.2) then let us identify how the students perceived this activity and understand the specific challenges that the students faced, including resource, expertise, and time constraints. Further exploratory analyses of creative extensions on top of data analysis replication show that replication extensions might be both more methodologically advanced and scientifically meaningful than unconstrained projects conducted the following year.

Educators. On the educator side, we provide realistic information about how much overhead is needed in teacher-to-student ratios for overseeing replications, how much effort is required to select papers for replications, and some concerns that replications bring over more traditional assignments. We offer further ‘lessons learned’ that can be useful to other educators, putting particular emphasis on reflections regarding cost–

benefit tradeoffs. The insights about the discrepancies between expectations and true outcomes, as well as the associated attitude shifts, will be informative for future efforts aiming to incorporate data analysis replications into existing educational practices. For example, since the replication activity took students longer than expected, instructors should carefully plan the course timeline and clearly communicate the expected workload to students, to avoid stress and frustration (cf. Section 5).

Science. Lastly, we identified potential ways how the scientific communities could benefit from this and similar efforts. Overall, we demonstrated that incorporating replication tasks into a large data science class has the potential to increase the reproducibility of scientific work as a by-product of data science instruction.

2. Methods

2.1. Study Design

In preparation for the study, we identified 10 data science publications suitable for the course, in terms of the difficulty of data analysis tasks required, a variety of covered topics, and data availability. The publications were split into five tracks, with two publications each:

1. Natural language processing and machine learning (Muchlinski et al., 2016; Niculae et al., 2015)
2. Computational social science (Choi & Varian, 2012; Pierson et al., 2020)
3. Networks (Cho et al., 2011; Leskovec et al., 2010)
4. Social media and Web (Liang & Fu, 2015; Penney, 2016)
5. Health (Aiello et al., 2020; Cattaneo et al., 2009)

We identified two key figures or tables from each of the publications that are important for the overall message of the publication. Teaching assistants (master's students who took the course the previous year) aimed to replicate (exactly or qualitatively) the selected figures before the class started, which ensured that the selected figures and tables were qualitatively reproducible. We developed pre- and postsurveys by conducting a pilot with student assistants.

The data analysis replication activity was composed of six steps. We asked students to fill out repeated short surveys, each part of a project milestone deadline. Each student was assigned one paper to replicate (around 36 students per replicated paper). In each paper we selected a primary and secondary figure or table. The primary figure or table requires basic skills taught in the lectures and exercises before the replication task was performed (limited to counting, hypothesis testing, visualizing, and fitting regressions). The secondary figure or table requires potentially more advanced data analysis such as nonstandard resampling and error estimation techniques, examination of feature importance, and network analysis. Note that henceforth we refer to the basic figure/table and the advanced figure/table as simply *basic* and *advanced figures* (although the result might be presented in a table).

Additionally, the 10 assigned papers were divided into two conditions based on the paper type (referred to as type A and type B). Paper type refers to the type of analysis necessary to perform the replication. For basic figures in type A papers, to reproduce a result, students were required to count items, perform hypothesis testing, and make a visualization and interpretation of the result (papers: Niculae et al., 2015, Liang and Fu, 2015, Cho et al., 2011, Aiello et al., 2020, and Leskovec et al., 2010). For type B papers, to reproduce a result, students had to fit a regression model and make a visualization and interpretation of the result (papers: Cattaneo et al., 2009; Choi & Varian, 2012; Muchlinski et al., 2016; Penney, 2016; and Pierson et al., 2020). In addition to the main assigned paper (referred to as ‘Paper 1’), each student was assigned two control papers (referred to as Papers 2 and 3) that the student does not replicate. One control paper was of the same type as the replicated paper, and one of the other type.

The study consisted of five steps, outlined in Figure 1:

Step 1: Reading abstracts of preselected papers and expressing a preference. Students were instructed to read abstracts of all the 10 preselected papers to get an idea of what the papers are about. Students then ranked the 10 papers by their preference of working on them for the project. After this, students were assigned a main paper (‘Paper 1’). We assigned the same number of students per paper. We calculated the average rank of preference for each paper across the students, and assigned papers to students in a balanced way, to minimize the total average rank since smaller rank implies higher preference. We also assigned to each publication two assistants who were in charge of mentoring students working on the respective data analysis replication.

Step 2: Reading the assigned paper. Students were instructed to read the assigned paper. Students were pointed to the freely accessible PDF and the data set repository. Students wrote a short summary (at most 500 characters). Upon submission of the summary, the students completed the presurvey measuring expectations of the replication of the assigned figure and general attitudes toward reproducibility.

Step 3: Replication of a basic figure. Students individually performed a replication of the assigned basic figure from the assigned paper (‘Paper 1’). Students prepared a replication report in the form of a Jupyter Notebook containing independently written code and text. Students were instructed to log their hours spent doing the replication, on a piece of paper, in a digital sheet, or using time-tracking software. To elicit truthful time log reports, it was clarified to the students that the number of reported hours would in no way affect the grading of the work. Upon submitting the replication report, the students completed the first postsurvey, which measured outcomes of the replication of the basic figure and expectations for the advanced figure. The main analyses (RQ1) contrast the postsurvey responses after replication of the basic figure with the presurvey responses given before the replication exercise.

Step 4: Replication of an advanced figure. In order to work on other graded assignments in the course, students formed groups of four students (some groups exceptionally comprised three students) working together throughout the semester. In their group, the students then proposed a creative extension of the analysis

performed in the paper, placing their data science skills into practice (Kolaczyk et al., 2021).⁴ When submitting the short project proposal, the students also completed the second postsurvey, a repeated measurement of the expectations for the advanced figure. Analyses in RQ2 contrast the second postsurvey with the expectation for the advanced figure.

Step 5: Creative extension. Students conducted the proposed creative extension in their group. Additionally, individually and following identical instructions as in step 3, the students replicated the advanced figure from the assigned paper. Students were again instructed to log their hours spent doing the data analysis replication. The students completed the third postsurvey, measuring outcomes of the replication of the advanced figure, and general attitudes toward reproducibility. General outcomes toward reproducibility are studied to address RQ3.

After each step, students were additionally asked about their expectations about the two control papers that they did not replicate (Paper 2 and Paper 3). We explore answers related to these control papers in order to address RQ4. In preparation for the study, we tested this pipeline with five student assistants.

2.2. Inclusion and Exclusion Criteria

The study took place at EPFL in the fall semester of 2020, between September 2020 and January 2021. In total, 384 students took the class. Out of 384, 30 students (7.81%) opted out from the study (resulting in 354 consenting students). Data from all the enrolled students were analyzed in the study, except from those who chose to opt out. We also excluded students who did not submit all four surveys or whose responses did not pass validation checks. With these restrictions, we analyzed responses from $N = 329$ consenting students.

2.3. Consent Statement and Information Sheet

Students were provided with the following information about the study and its purpose: “As part of ADA 2020, we introduced data analysis replications as a way of making you interact with real data science research. In order to understand the effectiveness of this new learning paradigm, we will analyze your solutions and survey responses, and we aim to publish a research paper about our findings. No personal data will be made public; we will only release aggregate, anonymized information. Every data point is valuable for us, but if you would nonetheless like to retract your data from the analysis, you can indicate this by checking the following box. Checkbox: I would like to be excluded from the analysis of the ADA data analysis replications.” An information sheet about the study was provided to students.⁵

3. Results: Data Analysis Replications

3.1. Preregistered Findings: Discrepancies Between Expectations and the Reality of Data Analysis Replication

Before analyzing the data collected via surveys, we formed and preregistered a set of primary and secondary hypotheses, each relating to one of the four research questions (RQ1–RQ4).⁶ We then executed the analyses

following the plan. Our primary confirmatory analysis tests the hypothesis that there are discrepancies between students' expectations and the reality of data analysis replication (H1 [RQ1]). An overview of the results is presented in Table 1. Replication package including code and data is publicly available (Gligorić, 2024).

Testing the preregistered hypotheses, we first asked: *Is there a significant difference between the time students take to perform the data analysis replication and the time they expect to take (H1a)?* We found that there is a significant difference ($p = .0309$; full distributions in Figure A1a and A1b). On average, students expected to take 9.01 hr, but actually took 10.53 hr. The median expected time is 5 hr and median time taken is 8 hr. In total, 62% of students took longer than expected, 7.30% the same, and 30.70% less than expected. So overall, students on average underestimate the time it would take to reproduce the basic figure.

Second, we compared how challenging students thought that it would be to reproduce the basic figure from the assigned paper with the reported true level of challenge. Specifically, we asked: *Is there a significant difference between how challenging performing data analysis replication tasks is and how challenging students expect it to be (H1b)?* We found that there is a significant difference (illustrated in Figure A1c and A1d)—interestingly, performing data analysis replication tasks was less challenging than expected ($p = 3.70 \times 10^{-5}$). The average expected score on the 1–5 scale is 3.39 (median 4), whereas the average score after performing the task is 3.11 (median 3).

Third, we conceptualized the data analysis replication task as being composed of three core activities: data wrangling (understanding the data structure, preprocessing steps, feature engineering), data analysis (exploratory analysis, statistical tests, developing and training models, evaluating model performance), and interpretation (evaluating results and comparing them with the results in the paper, interpreting findings, and redoing the analysis if necessary). Then, we then asked: *Are there discrepancies between the predicted and the true distribution of time spent on the three core activities: data wrangling, data analysis, and interpretation (H1c)?* We found discrepancies ($p < 10^{-307}$)—in relative terms, students overestimated how much time data wrangling would take, and underestimated how much time data analysis and interpreting results would take (Table 2). This finding shines light on why replication took more time than expected, but was less challenging than expected. Students took more time iteratively redoing the data analyses, interpreting their results, which was perceived as time-consuming, although not technically challenging.

Finally, we asked: *Are there discrepancies between predicted and true outcomes of the replication (H1d)?* First, we found that 98% of students reported having replicated exactly or qualitatively the basic figure, and 87% the advanced figure. We did not find significant discrepancies between predicted and true outcomes of the replication ($p = .0747$; illustrated in Figure A1e and A1f). A possible explanation is that the papers were preselected to be (with enough effort) at least partially qualitatively replicable. Students were not exposed to randomly sampled papers from the field. Rather, the selected papers were already found to be qualitatively

reproducible in our paper selection process. Further statistics are available in the Appendices, Appendix A: Primary Hypotheses—statistics and data distribution visualization.

Table 1. Results overview: Modified expectations (H1). Summary of the results comparing pre- and postreplication expectations, across the four hypotheses (H1a–H1c).

| Outcome | Summary Statistics |
|---------------------------|---|
| H1a: Time spent | Students underestimate the time it would take to reproduce: 1.52 hour increase ($p = .0309$). Pre test: $M = 9.01hr$, post test: $M = 10.53hr$. |
| H1b: Level of challenge | Data analysis replication was less challenging than expected: 0.28 point decrease ($p = 3.70 \times 10^{-5}$). Pre test: $M = 3.39$, post test: $M = 3.11$, on 1–5 scale. |
| H1c: Time distribution | Students overestimated time for data wrangling, and underestimated time for data analysis and interpreting results: Significant disturbance in the ranking ($p < 10^{-307}$). Pre test ranking in decreasing order: Wrangling, Analysis, Interpretation, post test ranking in decreasing order: Analysis, Interpretation, Wrangling |
| H1d: Replication outcomes | Difference not significant ($p = .0747$). Pre test: $M = 1.81$, post test: $M = 1.75$, on a 1–3 scale. |

We also considered a set of secondary hypotheses (H2–4). First, we hypothesized (H2 [RQ2]) that discrepancies between predictions and true outcomes persist as students solve replication tasks (complete statistics available in Appendices, Appendix B: Secondary hypotheses). Overall, when reproducing the advanced figure after the basic one, discrepancies between expectations and outcomes persisted (although some in the opposite direction). Most notably, there were discrepancies between the predicted and the true distribution of time spent on the core activities and between predicted and true outcomes of the replication, since there were replication failures that the students did not expect. Second, we found no evidence supporting the hypothesis that the replication task affects the students' expectations on the fraction of peer-reviewed data science papers that are reproducible (H3 [RQ3]; Appendices, Appendix B: Secondary Hypotheses).

Third, we hypothesized that there is a spillover effect as expectations are modified across the board, to papers that students did not replicate (H4 [RQ4]). Overall, we indeed found that there is a spillover effect as expectations regarding time spent and time distribution across the activities are modified for the papers that students did not replicate (summarized in Table B1).

Summary. Overall, we found that data analysis replication tasks take longer, but are less challenging than expected. Compared to the expectations, students spent more time analyzing and modeling the data and interpreting the results, and less time in data wrangling activities. We did not find significant discrepancies between predicted and true outcomes of the replication. The considerable amount of time spent modeling and interpreting the results may explain why replication took more time than expected, while simultaneously being less challenging than expected. We found that students took time iteratively redoing the data analyses, and interpreting their results, which was perceived as time-consuming, although not necessarily technically challenging. The identified discrepancies and attitude shifts signal students' enhanced appreciation for the challenges involved in the scientific process.

3.2. Exploratory Findings: Understanding the Students' Experience

Next, we complement the previous findings with an exploratory study identifying the challenges students experienced during the replication activity, to understand the gaps between the expectations and the reality of data analysis replication. In this analysis, we qualitatively investigate the open-text responses to two questions we included in the postsurvey: (1) "What was challenging?" (2) "What may explain the differences?" Students replied to these questions after replicating the second figure and completing the replication assignment.

To understand what topics the students mentioned, two of the authors of this study qualitatively coded the students' answers using a grounded-theory approach. For both questions, we independently repeated the following process. The researchers autonomously read a random sample of 100 answers and produced a list of topics mentioned in the students' descriptions. These topics were then compared and discussed until an agreement on their representativeness was reached. This process led to merging similar topics and refining the names describing them. Then, each researcher assigned the obtained topics to the answers. Multiple topics (or none) could be assigned to an answer. Finally, the label assignments were compared and, in case of discrepancies, discussed until a final agreement was reached. At the end of the process, the answers not assigned to any previously agreed topics were examined to extract new possible labels. If new topics were identified, the process was repeated; otherwise, the process terminated by leaving these answers unlabeled. The outlined topic coding approach was applied to two open-ended questions included in the postsurvey: (1) "What was challenging?" (2) "What may explain the differences?" We report topics assigned to at least 5% of the answers.

What was challenging about the data analysis replication? In this question, students were asked to describe in two or three sentences what they found challenging during the replication task. Most students (77%)

described challenges assigned to at least one of the topics. Inspecting the unassigned responses (23%) did not lead to introducing additional themes. Rather, the unassigned responses were short and vague (e.g., “probability issue”) or uninformative (e.g., “It did not replicate at all”).

We identified four frequent topics: **Poor Description**, **Expertise Requirements**, **Time Requirements**, and **Limited Resources**. In the following paragraphs, we report more details about the four themes and the relative commonness in the questions assigned to at least one topic. Since each answer can be assigned to multiple topics, the percentages of assignments do not sum to 100%.

Poor Description (60%): Students pointed out that the main challenge in replicating the authors’ results was a poor description of the process. This issue includes missing details about the parameters used in the modeling (e.g., size of the random forest model), little information on the data preprocessing steps, inconsistency between the data released and the description in the article, and explicit mistakes of the authors in reporting the method details (e.g., wrong start date in a time series analysis). This issue was summarized by one student as: “It’s almost a guessing game as to what method or inclusion I might be doing differently. This lack of hints was fairly difficult to navigate.”

Expertise Requirements (37%): Many students mentioned their lack of expertise as one challenge they encountered during the replication. Their descriptions varied from specific issues, such as the need to be confident in manipulating and plotting the data (e.g., how to plot timestamps on the x -axis), to more complex problems, such as the use of some advanced techniques (e.g., domain-specific hypothesis testing).

Time Requirements (17%): Students frequently mentioned the amount of time they spent working on the replication as a challenge. This problem is often associated with a poor description and is often described as many trial and error attempts.

Limited Resources (11%): Finally, some students found working with the data provided challenging because of its scale. The computation time required to process large data sets represented a limitation for students working with personal laptops.

What may explain the differences between the original and the reproduced result? In this question, we investigate what the students believed could explain the differences between the figure in the paper and the one they obtained in the replication task. First, we asked as a multiple-choice question if they were able to replicate the results exactly (a), qualitatively (b), or not at all (c). Then, students were asked to describe in two or three sentences what may explain the differences. The most common outcome is that figure “replicated qualitatively but not exactly” (b, 73.2%), followed by “did not replicate at all” (c, 13.9%), and “replicated exactly” (a, 12.7%).

In this second exploratory analysis, we focus on students who obtained similar results (b) or failed to reproduce the figure assigned (c). We identified five recurrent topics mentioned by the students who could not replicate

the figure exactly: **Poor Description, Data Issues, Authors' Mistakes, Tool Differences, and Students' Skills**. As for the previous analysis, each answer can mention multiple problems. We found that 83% of the answers are assigned to at least one theme, while the remaining 17% were not informative and could not be assigned to new topics.

Poor Description (55%): Similarly to what we observed in the answers to the previous question, students blame the limited description for the mismatch between their results and the article's figure. Answers in this category frequently mention a lack of details on the models' parameters used by the authors. Students who managed to reproduce the results only qualitatively pointed out that it was impossible to reproduce the figure exactly when the code and seeds used for 'random initialization' are unavailable. Another common observation was the limited description of all the steps and choices involved in the preprocessing pipeline. These aspects include how authors sampled data, handled missing values, what qualifies as outliers, and what numeric rounding steps are involved.

Data Issues (30%): Many students attribute their impossibility of reproducing the results to problems associated with the data. These problems come from issues with the data release that does not entirely match the description in the paper or from an incomplete release of the data necessary to reproduce all the results. Students encountering this last limitation went as far as trying to collect their own data set with all the associated challenges—especially if depending on an outdated automated programming interface (API).

Authors' Mistakes (24%): A significant portion of students assigned the blame for the impossibility of reproducing the results to the authors of the research. Answers assigned to this category mentioned possible embellishment of the results by the researchers and both genuine mistakes in reporting or plotting (e.g., "The authors interchanged a row at some point which messed up their analysis") and bad-faith adjustments (e.g., "The authors did some shady-ish things, for example hard coding the plot").

Tool Differences (11%): Some students suspected that the discrepancy between the tool used for the replication and the originally used tool may play a role in obtaining different results. They speculated on potential differences in the model and optimizer implementations available in Python, R, and Stata.

Students' Skills (7%): Lastly, some students believe mistakes on their side can be a possible reason for the differences. Some of them mention general mistakes in their code, whereas others describe their inexperience in doing effective data preprocessing and using libraries or methods that are not explicitly covered in the course material (e.g., "Researchers used a very advanced algorithm from another paper and I would be surprised if any student fully implemented it."). This topic is relatively infrequent, likely due to the fact that, in preparation for the study, we identified publications suitable for the course in terms of the difficulty of tasks required. We additionally ensured that the course lectures ahead of the replication covered the crucial skills necessary to perform the replication.

Paper-specific common feedback. Lastly, we aim to understand whether there were blocking factors that made it impossible for students to replicate the result that cannot be addressed simply by taking more time.⁷ We reexamined the students' explanations separately per paper in order to identify issues that students consistently mention when the result is not replicated. Such consistent issues that are reported many times might be authors' own mistakes or a true lack of information. We list recurring issues for five papers where more than 10% of the students self-reported that they did not manage to replicate at all any of the assigned figures (Table 2). Explanations for the remaining five papers did not contain any repeatedly occurring explanations.

Examining the recurring explanations, we identified two recurring issues—a cross-validation mis- take (Muchlinski et al., 2016) and counting error (Leskovec et al., 2010), which were known to the instructors in advance and were correctly identified by students, while the other recurring issues mainly reflect a lack of information or other preprocessing discrepancies.

Table 2. Challenges students encountered, separately by paper (exploratory analysis). The frequency of self-reported outcomes across students. For each paper, specific shared challenges that the students identified in their open-form responses when asked to explain the discrepancies between the original result and their result, and to speculate as to why the differences arose. Note that qualitative replication is the most frequent self-reported outcome for each of the figures.

| Paper | Figure | Replicated exactly | Replicated quantitatively | Did not replicate |
|--|---------|--------------------|---------------------------|-------------------|
| Muchlinski et al., 2016 | Fig. 2 | 3.7% | 96.3% | 0% |
| | Fig. 4 | 0% | 55.56% | 44.44% |
| Random forest parameters and random seed are not stated in the paper. Specific feature importance metric is not explained in detail. Programming language or library incompatibility meant that students could not reproduce the order of features. Authors' mistakes in the cross-validation procedure. | | | | |
| Cho et al., 2011 | Fig. 2A | 7.69% | 80.77% | 11.54% |
| | Fig. 3B | 0% | 84.62% | 15.38% |
| Outlier removal and the preprocessing steps are not explained in sufficient detail. Null baseline is not explained in detail. | | | | |
| Leskovec et al., 2010 | Tab. 1 | 6.9% | 93.1% | 0% |
| | Tab. 3 | 6.9% | 55.17% | 37.93% |
| Authors' mistakes in data processing and counting. | | | | |

| | | | | |
|--|---------|--------|--------|--------|
| Aiello et al., 2020 | Fig. 5 | 81.82% | 18.18% | 0% |
| | Fig. 4 | 4.55% | 59.09% | 36.36% |
| Not sufficient information provided in order to reproduce the figure. Scaling of the distributions not explained in detail. | | | | |
| Penney, 2016 | Fig. 3 | 0% | 96.43% | 3.57% |
| | Fig. 4A | 0% | 89.29% | 10.71% |
| The original data set is not available. The data set that the students used contained slight discrepancies. | | | | |

4. Results: Creative Replication Extensions

As part of Step 5, in their groups, students conducted creative extensions of the analysis performed in the paper. According to the instructors' anecdotal experience, this creative component of the project—which students built on top of the replicated papers—was in many cases more technically advanced and meaningful compared to the unconstrained project in regular iterations.

To confirm these observations, we conducted several follow-up exploratory analyses. We analyzed structured project descriptions provided by the students in a consistent format. Across class iterations, these were submitted at the start, and updated at the end of the project. The descriptions were provided in a structured `README.md` document and contained a title, abstract, and a description of the research question(s), data set(s), and methods.

First, we developed a systematic method to automatically code structured project descriptions for the type of approaches each one uses and their scientific contributions (using GPT-4 model; Appendices, Appendix C: Annotation Details), an approach evaluated for similar text classification tasks (Gilardi et al., 2023). Specifically, we developed two custom prompts for GPT-4 and applied each prompt to the structured project descriptions that students are required to write. We confirmed that the GPT-4-generated annotations had high agreement with independent human annotations on a subset of descriptions. In particular, two authors annotated a random sample of project descriptions. The sample was also annotated leveraging the GPT-4 model, using the same instructions (see Appendices, Appendix C: Annotation Details for specific prompts and model parameters). We measured a substantial agreement between the manual and automated annotations ($\kappa = .70$ and $\kappa = .77$). Complete instructions and details about the agreement metrics are listed in Appendices, Appendix C: Annotation.

Then, we applied this method to the structured project descriptions. Each project description was annotated to indicate what type of data analysis method the project leveraged, among those covered in the class. The methods ranged from simple descriptive approaches, over less simple approaches (inference and prediction), to more technically advanced causal inference techniques. Following the same approach, each project description

was also annotated to indicate whether it is scientifically relevant, by considering whether the project potentially pushes the boundaries of current scientific knowledge, as adapted from National Science Foundation definition of transformative research (U.S. National Science Foundation, 2024).

We annotated descriptions of replication extensions conducted after data analysis replications performed as part of the project component of the class (2020, $N_1 = 115$), and descriptions of standard, open-ended projects, which were conducted in the following year (2021, $N_2 = 114$), when data analysis replications were not integrated into the class, but students instead independently proposed and conducted a project topic of choice. Students had the freedom to select their own project topic such that it did not rely nor build on data analysis replication.⁸ We then compared the results between the two years, contrasting replication extension projects with open-ended projects as conducted in other iterations of the class.

Creative replication extensions are more technically advanced than unconstrained projects. In line with our anecdotal experience, we found that creative replication extensions are significantly less likely to be focused on descriptive statistics and data visualization (e.g., simple statistical tests and correlations) compared to unconstrained projects (5.22% vs. 20.35%; $\chi^2 = 11.58$, $p = 6.67 \times 10^{-4}$). Simultaneously, replication extensions are more likely to focus on causal inference and counterfactual techniques (e.g., effect estimation and matching) compared to unconstrained projects (9.57% vs. 0.88%; $\chi^2 = 8.70$, $p = 3.18 \times 10^{-3}$). No significant differences were observed for statistical modeling and inference (38.26% vs. 41.59%; $\chi^2 = 0.21$, $p = 0.64$) and machine learning and prediction (46.96% vs. 37.17%; $\chi^2 = 2.41$, $p = 0.12$). Complete histogram across the two years is visualized in Figure 2. These findings were robust to the inclusion of specific methods as examples, and to an alternative format where all data analysis types that apply are selected, as opposed to one that applies the most. This confirmed that the more advanced data analysis type—causal inference—is the least frequent analysis type, but more frequent among replication extensions than unconstrained projects (Table C1).

In summary, this exploratory analysis is aligned with the insight that creative replication extensions tend to be more methodologically advanced compared to unconstrained projects. Replication extension projects are associated with an increased use of more advanced causal data analysis methods, and a decreased use of less advanced descriptive methods.

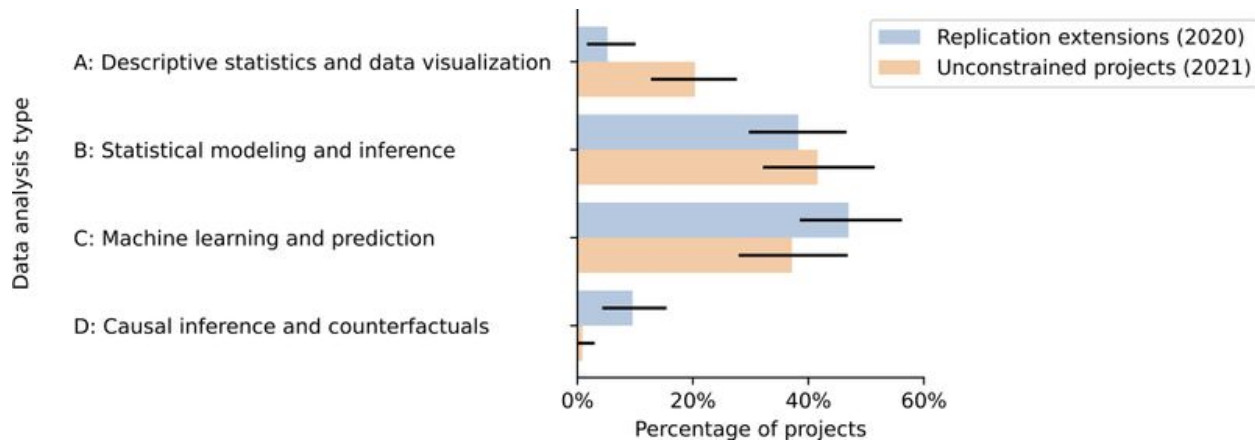


Figure 2. Data analysis types, between years. Histogram of the data analysis type across projects, in 2020, the year of creative replication extensions (blue), and 2021, the year on unconstrained projects (orange). Error-bars mark bootstrapped 95% CI. Creative replication extensions are more technically advanced than unconstrained projects, as captured by a decreased use of less advanced descriptive methods (A), and an increased use of more advanced causal data analysis methods (D).

Creative replication extensions are more meaningful than unconstrained projects. Compared to unconstrained projects, replication extensions were significantly more likely to be classified as scientifically relevant (15.65% vs. 5.26%; $\chi^2 = 6.59$, $p = 1.03 \times 10^{-2}$), again confirming the insight that replication extensions are more meaningful than unconstrained projects.

Lastly, to explore whether projects differ in further ways, beyond those tested, we annotated project descriptions with adjectives that best capture the strengths of the project (Appendices, Appendix C: Annotation). We identified four adjectives occurring with a significantly different frequency between the two years ($p < .05$; full results in Table C2): “practical” and “methodical” (more frequent among replication extensions), and “insightful” and “comprehensive” (more frequent among unconstrained projects). This analysis again points to unconstrained projects being more descriptive (comprehensive and providing insights), while replication extensions of replicated work focus on systematically executing advanced methods, and are of practical value and relevant (methodical and practical), as hypothesized.

Summary. Exploratory analyses point in the same direction as the instructors’ experience—replication extensions are both more technically advanced and scientifically meaningful than unconstrained projects. Perhaps counterintuitively, creative extensions built on top of replications might be more meaningful than unconstrained projects (Rosso, 2014). Unconstrained, students test many potential paths since they have not yet performed a viable data analysis. In contrast, extensions of data analysis replications allow going further beyond numerous shallow analyses, and are therefore more meaningful, allowing students to start from a strong foundation.

We note that this analysis is exploratory. Many other factors could contribute to differences between creative replication extensions and the more traditional, open-ended data analysis projects. These could include fundamental differences in the student body, instruction, and broader factors related to the class and the external environment.

For instance, students performed projects using different data sets, and data set type could impact both the methodologies and the scientific relevance. However, in principle, different data sets allow performing all the data analysis types. When considering the data set type and limiting the projects only to those that primarily analyze the most common data set type (textual data, 113 in total), we still see consistent patterns such that the leveraging causal inference methods (5.88% vs. 1.04%) and scientific relevance (35.29% vs. 5.21%) are more common among replication extensions. Similarly, grading guidelines, instruction, and student prerequisites were otherwise unchanged. Nonetheless, other factors could impact these patterns and future work is needed to truly identify advantages of replication extensions compared to traditional assignments.

5. Considerations for Educators

Having described how students experience data analysis replications, we now report insights and further ‘lessons learned’ that can be useful to educators, with a particular emphasis on the necessary considerations to integrate data analysis replications into a class. The outlined points are based on the instructor and assistants’ experiences and discussions, students’ anonymous feedback, and the results of the surveys administered throughout the class. Although data analysis replications might have their advantages (cf. Section 4), integrating them into an existing course is challenging. Based on our study, we highlight five major considerations.

5.1. Logistics

When designing and conducting in-class data analysis replications, it is necessary to carefully reevaluate and implement changes in the order in which the concepts are taught throughout the semester, since replicating data analyses requires specific skills (such as statistical tests, regression modeling, or counting items). One has to ensure that at the time when students start working on it, they have the required knowledge, which can lead to tradeoffs. In the study, in addition to modifying the class schedule, we carefully reconsidered other logistical aspects of the class, including group size and student assignment to projects and advisors.

5.2. Human resources

In-class data analysis replication activities may require additional human resources. In our class (with $N = 354$ consenting students), two teaching assistants dedicated half-time of their teaching assistantship to coordinating the project component of the course, as part of which the replication analysis was conducted. This amounted to around 8 hours per week. Additionally, around 30 students were assigned to each teaching assistant. The teaching assistants provided ongoing support specifically to the replicated paper throughout the semester, as

well as performed grading, troubleshooting, technical support, and data analysis replications in preparation for the class.

5.3. Added Constraints

We note that, if implemented as part of a standard component of a class (e.g., project or homework), data analysis replications may constrain the topics, as students cannot perform a project of choice, but have to build on top of the data analysis replication. Additionally, student level needs to be considered, and the activity designed to be appropriate.

5.4. Ethical Challenges

Data analysis replication activities call for ethical consideration. First, we had doubts about assigning students to papers that we knew were likely not to replicate at all, because we did not want to give students tasks we knew were unlikely to succeed. The ethical issue of potentially knowingly exposing students to stress and frustration limits the pool of paper candidates. Second, since the replication activity takes students more time than expected, instructors should carefully plan the course timeline and communicate the expected workload clearly to students, to avoid any stress and frustration.

5.5. Grading

Grading guidelines were adapted to the replication exercise. Each student submitted a computational notebook containing well-commented code to create the figure or the table that was replicated, textual descriptions guiding through the process, and the figure/table that is the result of the replication.

Grading was independent of the replication outcome. Students were instructed that they would be graded based on the overall quality of the replication, textual descriptions, and code. It was noted that it would not be graded whether or not students managed to replicate the results from the paper, but only whether they had made an honest and diligent attempt at replicating, given the information available in the paper. We developed grading guidelines that specified the mapping between grades and the quality of textual descriptions and code, and provided graders with examples, which helped reduce subjectivity.

Moreover, we advise caution in grading when assigning multiple papers within the same class. The selection of papers such that they are of comparable difficulty with regard to reproducibility is challenging, given that there are many paths one could take during a data analysis, and students are bound to face challenges that were not anticipated (Merrill et al., 2021). Students are sensitive to a perceived uneven workload across groups, might prioritize performance in class, and in other ways feel that it is unfair that there is variance across groups in the amount of time they had to spend. In our study, this was the only aspect of the data analysis replication activity that the students reflected on negatively in the anonymous feedback. Alternatively, a single class-wide project would address the issue of an uneven workload, but might not fit specific students' interests.

6. Discussion

Our study characterizes students' experiences performing data analysis replications and derives insights and necessary considerations for educators aiming to incorporate them into classes.

Data analysis replications from the students' perspective. First, testing our primary preregistered hypothesis, we found a significant difference between the expectations and reality of data analysis replications (Figure A1). The activity was more time-consuming and less challenging than anticipated, likely because the tasks were laborious and iterative (Section 3.2). It is noteworthy that the attitude shifts extended beyond the specific papers that the students replicated, into control papers, where, following the initial replications, students' expected time to perform the replication increased by about two hours (Figure B1). The identified discrepancies between expectations and reality, and the observed changes in expectations about reproducibility, serve as evidence of students' attitude shifts that have the potential to promote students' appreciation for the challenges involved in the scientific process.

Second, the creative component of the project, which students built on top of the replicated papers, was more technically advanced and meaningful than what students do in a fully unconstrained project in regular iterations, according to the instructors' experience and exploratory analyses of produced artifacts. This implies that data analysis replications might serve as one way to prepare students for addressing methodologically advanced and scientifically relevant problems.

Data analysis replications from the educators' perspective. Integrating data analysis replications into an existing course requires thoughtfulness and can run into challenges. We outline essential considerations for educators. Overall, we emphasize the need for careful logistics planning, allocating sufficient human resources, addressing ethical challenges, and devising appropriate grading strategies. We advise grading based on effort and methodology rather than replication outcomes. Moreover, we highlight and discuss necessary adjustments in course design, including the sequence in which concepts are taught and group sizes. We strongly emphasize the need for appropriate teaching assistants to support students and manage workloads, alongside carefully considering and selecting in advance publications that match both students' skill levels and individual interests.

Data analysis replications from the scientists' perspective. Moving forward, the scientific communities could potentially benefit from this and similar efforts. Teaching students to do data analysis replications can increase the overall number of conducted replications. Further advantages include a potential shift of norms and incentives if the auditing paradigm becomes more prevalent. If researchers are aware of large data analysis replication attempts and more replications are done, more attention may be paid to reproducibility in the future. Lastly, students' own experiences with replications may have an impact on their understanding and appreciation of reproducibility problems, and lead them to take measures to ensure that their own work is reproducible.

6.1. Limitations

We note that we are not measuring how replication exercises prepare students to practice computational science. Are replication exercises effective in teaching coding skills, deepening understanding, or gaining confidence in conducting independent research? While our study does not address these questions, we paint an initial picture of how students experience data analysis replications, and how that experience enhances students' understanding of what a data analysis replication entails.

Similarly, our study does not disentangle the educational impact of a data analysis replication task from the educational impact of another comparable data analysis task. We contrast measurements before and after the activity, without randomly assigning students to the experimental conditions. Randomized assignment to the replication activity vs. another type of data analysis activity was considered but ruled out due to ethical challenges and to avoid student frustration. Conversely, self-selection into a condition (replication vs. standard data analysis) would introduce biases and was hence also ruled out. Nonetheless, carefully designed cross-sectional longitudinal comparisons (Section 2) surfaced insights about the impact of replications within a set of students who all performed the activity.

6.2. Future Work

Our study opens the door for several future directions aimed at understanding how to conduct in-class replication activities. First, exploratory analyses of students' perceptions of their ability to reproduce results revealed a tension between attributing inconsistencies to either the authors' mistakes or the students' perceived lack of skill. This raises an important follow-up question: when can a replication attempt be considered complete, allowing a student to stop, rather than assuming the inconsistencies are due to the students' (lack of) skills or mistakes? It remains unclear what constitutes a sufficient and satisfactory time investment. How can we prevent students from committing unlimited time to unproductive replication attempts?

One proposed solution may involve providing students with a limited number of submissions to a platform for corrected checks. This can involve allocating a 'budget' with the number of attempts submitted to a platform to evaluate the data analysis results (similar to leaderboards where participants submit the predictions on a test set for evaluation to a platform). This approach would, however, require the instructors to *a priori* know the correct results of the data analysis to be performed, which would in turn defeat the power of data analysis replications to serve as a way of scaling up reproducibility checks.

Second, future efforts should consider building a crowdsourced cohort of university students to standardize and unify similar efforts (Berkeley Initiative for Transparency in the Social Sciences, 2024; Höffler, 2017; Schooler, 2014). Such efforts to redesign undergraduate courses for reproducibility and collaboration across institutes can result in fostering open science (Button, 2018).

Third, our study was based on 10 preselected publications tested in advance. In the future, we envision development of an auditing paradigm where classrooms are fundamentally integrated into the scientific process to evaluate comprehensive samples of published scientific findings, beyond the carefully selected pool used here.

Finally, future research integrating tools to support replication attempts is called for, including the usage of software containers, cloud computing, and checkpoints. These tools make it possible to standardize the computing environment around each submission (Hofman et al., 2021; Liu & Salganik, 2019). Standardizing the computing environment becomes particularly relevant in the age of closed-access large language models increasingly used as part of data analysis and modeling pipelines.

6.3. Conclusion

Our study explores the paradigm of in-class data analysis replications with a double purpose: to teach students while improving the scientific process. We show that incorporating replications tasks into the project component of a large data science class has the potential to establish and increase the reproducibility of scientific work as a natural by-product of data science instruction. We hope this article will inspire other instructors to consider including data analysis replications in their classes.

6.4. Ethics

This study was approved by the EPFL Human Research Ethics Committee. We obtained consent for using the produced materials and survey responses for conducting research. Students were able to opt out of their data being analyzed. Students were provided with an information sheet (Methods, “Information sheet for students”). The analyzed data is anonymized. Furthermore, the students were informed that any survey analyses would be conducted only after the class had already finished and the grades had been formed.

Disclosure Statement

Kristina Gligorić, Tiziano Piccardi, Jake M. Hofman, and Robert West have no financial or non-financial disclosures to share for this article.

References

- Aiello, L. M., Quercia, D., Schifanella, R., & Del Prete, L. (2020). Tesco Grocery 1.0, a large-scale dataset of grocery purchases in London. *Scientific Data*, 7(1), Article 57. <https://doi.org/10.1038/s41597-020-0397-7>
- Baker, M. (2016). 1,500 scientists lift the lid on reproducibility. *Nature*, 533(26), 452–454. <https://doi.org/10.1038/533452a>

Ball, R. (2023). “Yes We Can!”: A practical approach to teaching reproducibility to undergraduates. *Harvard Data Science Review*, 5(3). <https://doi.org/10.1162/99608f92.9e002f7b>

Begley, C. G., & Ioannidis, J. P. (2015). Reproducibility in science: Improving the standard for basic and preclinical research. *Circulation Research*, 116(1), 116–126. <https://doi.org/10.1161/circresaha.114.303819>

Bergh, D. D., Sharp, B. M., Aguinis, H., & Li, M. (2017). Is there a credibility crisis in strategic management research? Evidence on the reproducibility of study findings. *Strategic Organization*, 15(3), 423–436. <https://doi.org/10.1177/1476127017701076>

Berkeley Initiative for Transparency in the Social Sciences. (2024). Social Science Reproduction Platform. Retrieved February 18, 2024, from <https://www.socialsciencereproduction.org/about>

Bishop, D. (2019). Rein in the four horsemen of irreproducibility. *Nature*, 568(7753), 435. <https://doi.org/10.1038/d41586-019-01307-2>

Boyce, V., Mathur, M., & Frank, M. C. (2023). Eleven years of student replication projects provide evidence on the correlates of replicability in psychology. *Royal Society Open Science*, 10(11), Article 231240. <https://doi.org/10.1098/rsos.231240>

Button, K. (2018). Reboot undergraduate courses for reproducibility. *Nature*, 561(7723), 287–288. <https://doi.org/10.1038/d41586-018-06692-8>

Cattaneo, M. D., Galiani, S., Gertler, P. J., Martinez, S., & Titiunik, R. (2009). Housing, health, and happiness. *American Economic Journal: Economic Policy*, 1(1), 75–105. <https://doi.org/10.1257/pol.1.1.75>

Cho, E., Myers, S. A., & Leskovec, J. (2011). Friendship and mobility: User movement in location-based social networks. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1082–1090). Association for Computing Machinery. <https://doi.org/10.1145/2020408.2020579>

Choi, H., & Varian, H. (2012). Predicting the present with Google Trends. *Economic Record*, 88(1), 2–9. <https://doi.org/10.1111/j.1475-4932.2012.00809.x>

Chopik, W. J., Bremner, R. H., Defever, A. M., & Keller, V. N. (2018). How (and whether) to teach undergraduates about the replication crisis in psychological science. *Teaching of Psychology*, 45(2), 158–163. <https://doi.org/10.1177/0098628318762900>

Frank, M. C., Braginsky, M., Cachia, J., Coles, N., Hardwicke, T., Hawkins, R., Mathur, M. B., & Williams, R. (2024). *Experimentology: An open science approach to experimental psychology methods*. MIT Press. <https://doi.org/10.7551/mitpress/14810.001.0001>

Frank, M. C., & Saxe, R. (2012). Teaching replication. *Perspectives on Psychological Science*, 7(6), 600–604. <https://doi.org/10.1177/1745691612460686>

Gilardi, F., Alizadeh, M., & Kubli, M. (2023). ChatGPT outperforms crowd workers for text-annotation tasks. *Proceedings of the National Academy of Sciences*, 120(30), Article e2305016120. <https://doi.org/10.1073/pnas.2305016120>

Glgorić, K. (2024). *Replication data for: In-class data analysis replications: Teaching students while testing science* (Version V1) [Data set]. Harvard Dataverse. <https://doi.org/10.7910/DVN/A6VMD9>, V1.

Goodman, S. N., Fanelli, D., & Ioannidis, J. P. (2016). What does research reproducibility mean? *Science Translational Medicine*, 8(341), 341ps12. <https://doi.org/10.1126/scitranslmed.aaf5027>

Hawkins, R. X. D., Smith, E. N., Au, C., Arias, J. M., Catapano, R., Hermann, E., Keil, M., Lampinen, A., Raposo, S., Reynolds, J., Salehi, S., Salloum, J., Tan, J., & Frank, M. C. (2018). Improving the replicability of psychological science through pedagogy. *Advances in Methods and Practices in Psychological Science*, 1(1), 7–18. <https://doi.org/10.1177/2515245917740427>

Head, M. L., Holman, L., Lanfear, R., Kahn, A. T., & Jennions, M. D. (2015). The extent and consequences of p-hacking in science. *PLoS Biology*, 13(3), Article e1002106. <https://doi.org/10.1371/journal.pbio.1002106>

Höfler, J. H. (2017). ReplicationWiki: Improving transparency in social sciences research. *D-Lib Magazine*, 23(3). <https://doi.org/10.1045/march2017-hoeffler/4>

Hofman, J. M., Goldstein, D. G., Sen, S., Poursabzi-Sangdeh, F., Allen, J., Dong, L. L., Fried, B., Gaur, H., Hoq, A., Mbazor, E., Moreira, N., Muso, C., Rapp, E., & Terrero, R. (2021). Expanding the scope of reproducibility research through data analysis replications. *Organizational Behavior and Human Decision Processes*, 164, 192–202. <https://doi.org/10.1016/j.obhdp.2020.11.003>

Janz, N. (2016). Bringing the gold standard into the classroom: Replication in university teaching. *International Studies Perspectives*, 17(4), 392–407. <https://doi.org/10.1111/insp.12104>

Kerr, N. L. (1998). HARKing: Hypothesizing after the results are known. *Personality and Social Psychology Review*, 2(3), 196–217. https://doi.org/10.1207/s15327957pspr0203_4

King, G. (1995). Replication, replication. *PS: Political Science & Politics*, 28(3), 444–452. <https://doi.org/10.2307/420301>

Kolaczyk, E. D., Wright, H., & Yajima, M. (2021). Statistics practicum: Placing ‘practice’ at the center of data science education. *Harvard Data Science Review*, 3(1). <https://doi.org/10.1162/99608f92.2d65fc70>

- Lemons, C. J., King, S. A., Davidson, K. A., Berryessa, T. L., Gajjar, S. A., & Sacks, L. H. (2016). An inadvertent concurrent replication: Same roadmap, different journey. *Remedial and Special Education, 37*(4), 213–222. <https://doi.org/10.1177/0741932516631116>
- Leskovec, J., Huttenlocher, D., & Kleinberg, J. (2010). Signed networks in social media. In *CHI '10: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (pp. 1361–1370). Association for Computing Machinery. <https://doi.org/10.1145/1753326.1753532>
- Liang, H., & Fu, K.-w. (2015). Testing propositions derived from Twitter studies: Generalization and replication in computational social science. *PLoS One, 10*(8), Article e0134270. <https://doi.org/10.1371/journal.pone.0134270>
- Liu, D. M., & Salganik, M. J. (2019). Successes and struggles with computational reproducibility: Lessons from the fragile families challenge. *Socius, 5*. <https://doi.org/10.1177/2378023119849803>
- Loken, E., & Gelman, A. (2017). Measurement error and the replication crisis. *Science, 355*(6325), 584–585. <https://doi.org/10.1126/science.aal3618>
- Makel, M. C., & Plucker, J. A. (2014). Facts are more important than novelty: Replication in the education sciences. *Educational Researcher, 43*(6), 304–316. <https://doi.org/10.3102/0013189X14545513>
- Mendez-Carbajo, D., & Dellachiesa, A. (2023). Data citations and reproducibility in the undergraduate curriculum. *Harvard Data Science Review, 5*(3). <https://doi.org/10.1162/99608f92.c2835391>
- Meng, X.-L. (2020). Reproducibility, replicability, and reliability. *Harvard Data Science Review, 2*(4). <https://doi.org/10.1162/99608f92.dbfce7f9>
- Merrill, M. A., Zhang, G., & Althoff, T. (2021). MULTIVERSE: Mining collective data science knowledge from code on the web to suggest alternative analysis approaches. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining* (pp. 1212–1222). Association for Computing Machinery. <https://doi.org/10.1145/3447548.3467455>
- Muchlinski, D., Siroky, D., He, J., & Kocher, M. (2016). Comparing random forest with logistic regression for predicting class-imbalanced civil war onset data. *Political Analysis, 24*(1), 87–103. <https://doi.org/10.1093/pan/mpv024>
- National Academies of Sciences, Engineering, and Medicine. (2019). *Reproducibility and replicability in science*. National Academies Press. <https://nap.nationalacademies.org/catalog/25303/reproducibility-and-replicability-in-science>

- Niculae, V., Kumar, S., Boyd-Graber, J., & Danescu-Niculescu-Mizil, C. (2015). Linguistic harbingers of betrayal: A case study on an online strategy game. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing* (pp. 1650–1659). Association for Computational Linguistics. <https://doi.org/10.3115/v1/P15-1159>
- Nosek, B. A., Ebersole, C. R., DeHaven, A. C., & Mellor, D. T. (2018). The preregistration revolution. *Proceedings of the National Academy of Sciences*, 115 (11), 2600–2606. <https://doi.org/10.1073/pnas.1708274114>
- Open Science Collaboration. (2012). An open, large-scale, collaborative effort to estimate the reproducibility of psychological science. *Perspectives on Psychological Science*, 7(6), 657–660. <https://doi.org/10.1177/1745691612462588>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), Article aac4716. <https://doi.org/10.1126/science.aac4716>
- Patil, P., Peng, R. D., & Leek, J. T. (2019). A visual tool for defining reproducibility and replicability. *Nature Human Behaviour*, 3(7), 650–652. <https://doi.org/10.1038/s41562-019-0629-z>
- Penney, J. W. (2016). Chilling effects: Online surveillance and Wikipedia use. *Berkeley Technology Law Journal*, 31(1), 117. https://btlj.org/data/articles2016/vol31/31_1/0117_0182_Penney_ChillingEffects_WEB.pdf
- Perry, T., Morris, R., & Lea, R. (2022). A decade of replication study in education? A mapping review (2011–2020). *Educational Research and Evaluation*, 27(1-2), 12–34. <https://doi.org/10.1080/13803611.2021.2022315>
- Perry, T., & See, B. H. (2022). Replication study in education. *Educational Research and Evaluation*, 27(1-2), 1–7. <https://doi.org/10.1080/13803611.2021.2022307>
- Pierson, E., Simoiu, C., Overgoor, J., Corbett-Davies, S., Jenson, D., Shoemaker, A., Ramachandran, V., Barghouty, P., Phillips, C., & Shroff, R. (2020). A large-scale analysis of racial disparities in police stops across the United States. *Nature Human Behaviour*, 4(7), 736–745. <https://doi.org/10.1038/s41562-020-0858-1>
- Plucker, J. A., & Makel, M. C. (2021). Replication is important for educational psychology: Recent developments and key issues. *Educational Psychologist*, 56(2), 90–100. <https://doi.org/10.1080/00461520.2021.1895796>
- Quintana, D. S. (2021). Replication studies for undergraduate theses to improve science and education. *Nature Human Behaviour*, 5(9), 1117–1118. <https://doi.org/10.1038/s41562-021-01192-8>

Rosso, B. D. (2014). Creativity and constraints: Exploring the role of constraints in the creative processes of research and development teams. *Organization Studies*, 35(4), 551–585.

<https://doi.org/10.1177/0170840613517600>

Schooler, J. W. (2014). Metascience could rescue the ‘replication crisis’. *Nature*, 515(7525), 9.

<https://doi.org/10.1038/515009a>

Shrout, P. E., & Rodgers, J. L. (2018). Psychology, science, and knowledge construction: Broadening perspectives from the replication crisis. *Annual Review of Psychology*, 69, 487–510.

<https://doi.org/10.1146/annurev-psych-122216-011845>

Silberzahn, R., Uhlmann, E. L., Martin, D. P., Anselmi, P., Aust, F., Awtrey, E., Bahník, Š., Bai, F., Bannard, C., & Bonnier, E. (2018). Many analysts, one data set: Making transparent how variations in analytic choices affect results. *Advances in Methods and Practices in Psychological Science*, 1(3), 337–356.

<https://doi.org/10.1177/2515245917747646>

Smith, L. M., Yu, F., & Schmid, K. K. (2021). Role of replication research in biostatistics graduate education. *Journal of Statistics and Data Science Education*, 29(1), 95–104.

<https://doi.org/10.1080/10691898.2020.1844105>

Stojmenovska, D., Bol, T., & Leopold, T. (2019). Teaching replication to graduate students. *Teaching Sociology*, 47(4), 303–313. <https://doi.org/10.1177/0092055X19867996>

Thornton, A., & Lee, P. (2000). Publication bias in meta-analysis: Its causes and consequences. *Journal of Clinical Epidemiology*, 53(2), 207–216. [https://doi.org/10.1016/S0895-4356\(99\)00161-4](https://doi.org/10.1016/S0895-4356(99)00161-4)

U.S. National Science Foundation. (2024). *Learn about transformative research*. Retrieved June 18, 2024, from <https://new.nsf.gov/funding/learn/research-types/transformative-research>

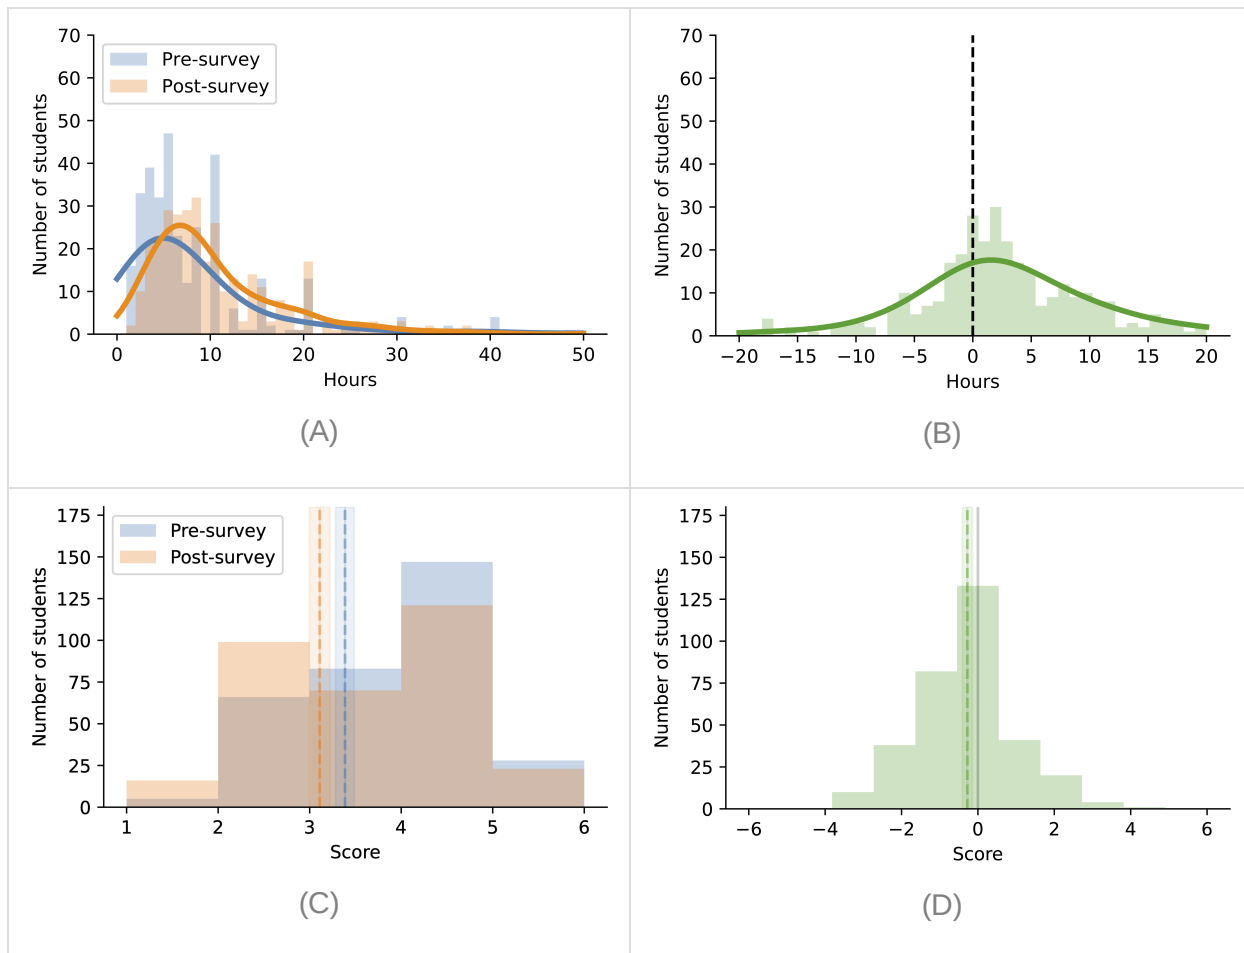
Wagge, J. R., Baciú, C., Banas, K., Nadler, J. T., Schwarz, S., Weisberg, Y., IJzerman, H., Legate, N., & Grahe, J. (2019). A demonstration of the Collaborative Replication and Education Project: Replication attempts of the red-romance effect. *Collabra: Psychology*, 5(1), Article 5. <https://doi.org/10.1525/collabra.177>

Wagge, J. R., Brandt, M. J., Lazarevic, L. B., Legate, N., Christopherson, C., Wiggins, B., & Grahe, J. E. (2019). Publishing research with undergraduate students via replication work: The collaborative replications and education project. *Frontiers in Psychology*, 10(1), 247. <https://doi.org/10.3389/fpsyg.2019.00247>

Appendices

Appendix A. Primary Hypotheses–Statistics and Data Distribution Visualization

Below, we list details about statistical analysis of our collected variables. All statistical tests were run with preregistered significance level $p = 0.05$. The unit of analysis is a student.



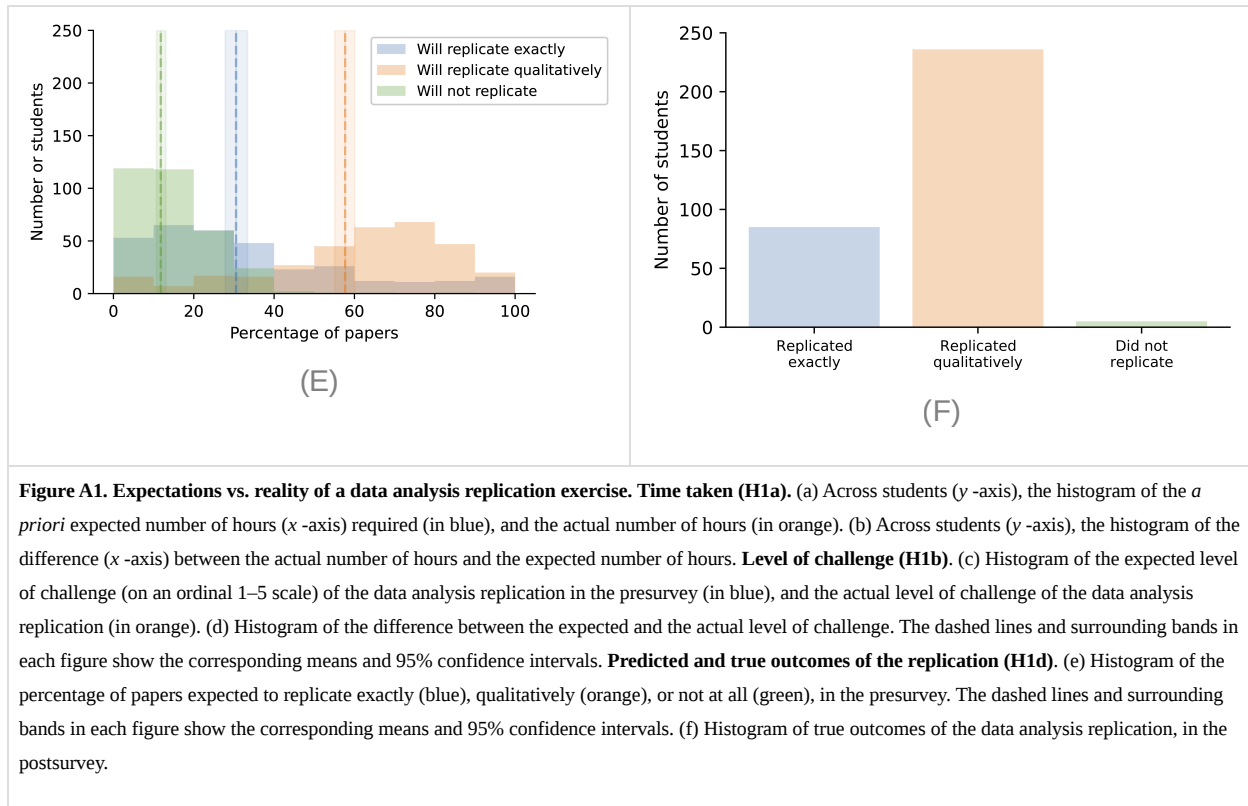


Figure A1. Expectations vs. reality of a data analysis replication exercise. Time taken (H1a). (a) Across students (y -axis), the histogram of the *a priori* expected number of hours (x -axis) required (in blue), and the actual number of hours (in orange). (b) Across students (y -axis), the histogram of the difference (x -axis) between the actual number of hours and the expected number of hours. **Level of challenge (H1b).** (c) Histogram of the expected level of challenge (on an ordinal 1–5 scale) of the data analysis replication in the presurvey (in blue), and the actual level of challenge of the data analysis replication (in orange). (d) Histogram of the difference between the expected and the actual level of challenge. The dashed lines and surrounding bands in each figure show the corresponding means and 95% confidence intervals. **Predicted and true outcomes of the replication (H1d).** (e) Histogram of the percentage of papers expected to replicate exactly (blue), qualitatively (orange), or not at all (green), in the presurvey. The dashed lines and surrounding bands in each figure show the corresponding means and 95% confidence intervals. (f) Histogram of true outcomes of the data analysis replication, in the postsurvey.

Preregistered analysis plan: Time spent: Students reported the expected number of hours in the pre-survey and the actual number of hours in the post-survey. Across students, we compared the expected number of hours to reproduce the basic figure from the assigned paper with the actual number of hours it took to reproduce. Specifically, we conducted a paired, two-sided t test on the difference between actual and anticipated number of hours, with a null hypothesis of no mean difference.

Preregistered analysis plan: Level of challenge: Students reported the perceived level of challenge on an ordinal scale (1: very straightforward, 2: somewhat straightforward, 3: neither straightforward nor challenging, 4: somewhat challenging, 5: very challenging). Specifically, we conducted a paired, two-sided t test on the difference between the actual and anticipated level of challenge, with a null hypothesis of no mean difference.

Preregistered analysis plan: Distribution of time across core activities: Students were asked to sort three core activities with respect to the amount of time they expected to spend on them (before the analysis), and with respect to the amount of time they actually spent on them (after the analysis). The three activities can be ranked in six possible ways. We treated each of the six ranking configurations as a categorical variable. Our main hypothesis here relates to a disturbance in the rank of the three core activities. The ranking configurations in the pretest and the posttest were paired across students in a 6×6 contingency table. We then performed the Stuart–Maxwell test for marginal homogeneity in the contingency table. The null hypothesis is that the activity rank configuration frequencies for pretest and posttest are the same.

Preregistered analysis plan: Replication outcomes: There are three possible self-reported outcomes of the data analysis replication: the analysis replicated exactly (the replication attempt produced results that agreed exactly with the paper, up to the decimals printed in the paper or shown in the figures), the analysis replicated qualitatively (the replication attempt produced results that had small differences with the paper, but these still agreed with the abstract-level findings of the paper), and the analysis did not replicate at all (the replication attempt produced results that were in conflict with the abstract-level findings of the paper).

We considered these outcomes as ordinal variables (1: the analysis replicated exactly, 2: the analysis replicated qualitatively, 3: the analysis did not replicate at all). In the pre-survey, students attributed a probability to each of the possible outcomes. We calculated the outcome expectation on the ordinal scale for each student by multiplying each possible outcome (1, 2, and 3) with the probability the student attributed to it and summing up. In the post-survey, students selected one of the outcomes. We compared the anticipated and the true value across students, for the basic figure from the assigned paper. We performed a paired two-sided t test.

Appendix B. Secondary Hypotheses

Here we provide detailed statistics and analyses addressing a set of secondary hypotheses (H2–4) summarized in the main text.

H2 (RQ2): *Discrepancies between predictions and true outcomes persist as students solve replication tasks.*

In the first replication task, the students replicated the basic figure, and in the second replication task, they replicated the advanced figure. We compared the predictions and true outcomes for the advanced figure in the assigned paper by repeating the same analyses and statistical tests described in H1a–d, but now for the advanced rather than the basic figure. We then explored the ways how the second replication task differs from the first replication task. In other words, we explored how the discrepancies between expectations and outcomes vary as students gain experience in conducting data analysis replication tasks.

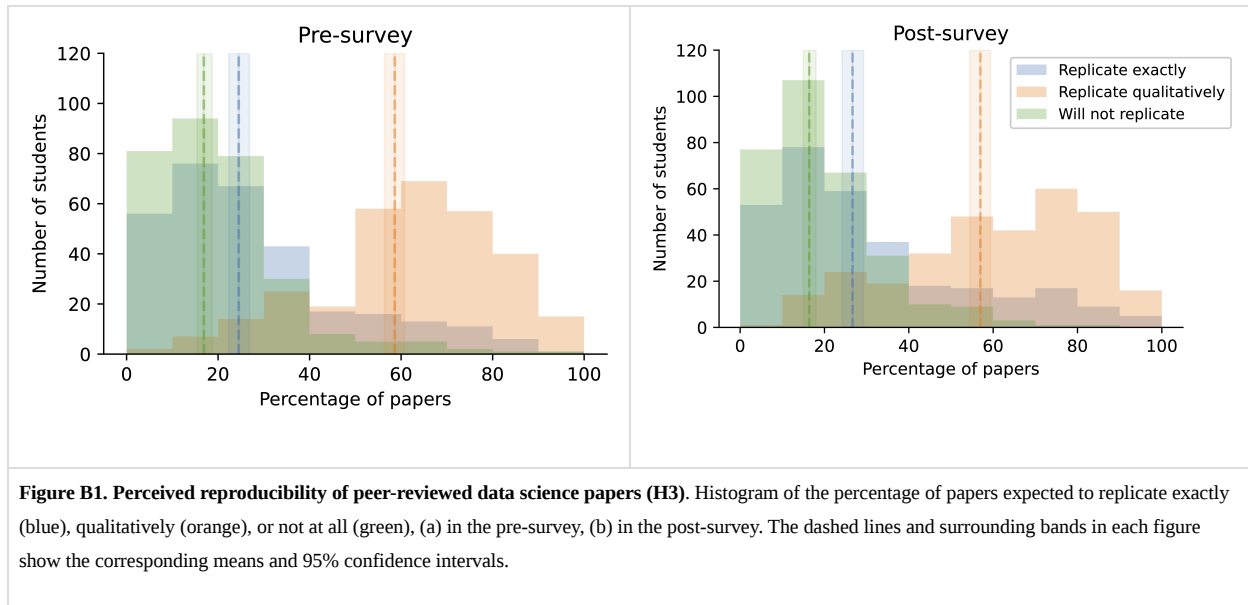
1. We found that a significant difference between the time students take to perform the advanced data analysis replication and the time they expect to take ($p = .0311$). On average, students expected to take 9.43 hr and took 8.54 hr. That is, after underestimating the time it takes to reproduce the basic figure, students overestimated the time it would take to reproduce the advanced figure (i.e., students overshoot after they initially underestimated).
2. We found that, after performing the replication of the basic figure, there was a significant difference between how challenging performing data analysis replication of the advanced figure is, and how challenging students expect it to be. Performing data analysis replication tasks was again less challenging than expected ($p = .00766$) as students overestimated how challenging it would be. The average expected score on a 1–5 scale is 3.10, whereas the average score after performing the task is 2.91. For comparison, in the case of the basic figure, the average expected score on the 1–5 scale was 3.39 and the average score after performing the task was 3.11.

3. For the advanced figure, we again found discrepancies between the predicted and the true distribution of time spent on the three core activities: data wrangling, data analysis, and interpretation ($p = 9.54 \times 10^{-7}$). In particular, on average, data wrangling and data analysis took less time than expected, while interpreting results took more time than expected. For the advanced figure, students again overestimated how much time data wrangling would take, and underestimated how much time interpreting the results would take. We find no significant difference for the data analysis component.
4. For the advanced figure, we found discrepancies between predicted and true outcomes of the replication ($p = 1.17 \times 10^{-13}$). As a reminder, we considered these outcomes as ordinal variables (1: the analysis replicated exactly, 2: the analysis replicated qualitatively, 3: the analysis did not replicate at all). The pretest average score is on average 1.76, whereas the posttest average score is 2.01. Overall, the outcomes were less successful than expected. That is, with the advanced figure, students faced more reproducibility issues than with the basic figure, as we expected.

H3 (RQ3): *The replication task affects the students' expectations on the fraction of peer-reviewed data science papers that are reproducible.*

At the beginning and at the end of the study, we asked the following question: “Out of 100 peer-reviewed data science papers published in 2020, in how many of these papers do you think the analysis would replicate exactly, the analysis would replicate qualitatively, and the analysis would not replicate at all?”

As before (H1d), we considered the outcomes as ordinal variables (1: the analysis replicated exactly, 2: the analysis replicated qualitatively, 3: the analysis does not replicate at all). We calculated the outcome expectation on the ordinal scale for each student, by multiplying each possible outcome (1, 2, and 3) with the probability the student attributed to it and summing up. Students were instructed to carefully verify that the three numbers add up to 100, and we excluded students whose responses do not pass this validation check. We then performed a paired two-sided t test on the outcome expectation at the beginning and at the end of the study. We did not find evidence that the replication task affects the students' expectations of the fraction of peer-reviewed data science papers that are reproducible ($p = .143$; illustrated in Figure B1).



H4 (RQ4): *There is a spillover effect as expectations are modified across the board, to papers that students did not replicate.*

Upon performing the replication tasks, we monitored any simultaneous changes in the expectations (as described in H1a–d) for the two control papers that students did not reproduce. We explored whether there were any changes in expectations regarding quantities described in H1a–d by repeating the same tests as outlined above, for the two control figures. One of the two control papers (‘Paper 2’) entails data analysis of the same type as the replicated paper (‘Paper 1’) and the other (‘Paper 3’) one of a different type (counting items and hypothesis testing vs. regression modeling). By contrasting the two control papers, we explored the presence of any spillover effects to different types of data analysis replication, beyond the specific type of analysis that the student worked on.

Overall, we found that there is a spillover effect as expectations regarding time spent and time distribution across the activities are modified across the board, for the papers that students did not replicate (summarized in Table B1). It is noteworthy that, even though the figures students were asked about were not replicated, the expectations changed after vs. before the replication activity. The expectations were modified in the same direction as for the replicated papers (by about 2 hours’ increase in expectation, and more time expected to spend in analysis and interpretation). The effects are not stronger for the same type of data analysis as performed in the replication exercise. We found that, overall, there was an attitude shift across data analysis types

Table B1. Spillover effects: Modified expectations regarding papers that students did not replicate (H4). Summary of the results comparing pre- and post-replication expectations, across the four hypotheses (H4a–H4c), for the two types of control papers that were not replicated.

| Control paper of the same type as replicated | |
|--|--|
| H4a: Expected time | 1.73 hour increase ($p = .0129$). Pre test: $M = 8.81$, post test: $M = 10.54$. |
| H4b: Expected level of challenge | Difference not significant ($p = .836$). |
| H4c: Expected distribution | Significant disturbance in the ranking ($p < 10^{-307}$). Wrangling: +0.36, Analysis: -0.04, Interpretation: -0.32. |
| H4d: Expected outcomes | Difference not significant ($p = .0804$). |
| Control paper of a different type than replicated | |
| H4a: Expected time | 2.07 hour increase ($p = .000434$). Pre test: $M = 8.75$, post test: $M = 10.82$. |
| H4b: Expected level of challenge | Difference not significant ($p = .161$). |
| H4c: Expected distribution | Significant disturbance in the ranking ($p < 10^{-307}$). Wrangling: +0.37, Analysis: -0.07, Interpretation: -0.31. |
| H4d: Expected outcomes | Difference not significant ($p = .0841$). |

Appendix C. Annotation Details

Two authors independently annotated a set of 20 project descriptions (10 each). We calculated Cohen’s kappa coefficient to probe interrater reliability between the authors and annotations conducted with GPT-4. For both questions (data analysis type and novelty of the scientific question), we found substantial interrater reliability between GPT-4 annotations and the authors’ annotations ($\kappa = .70$ and $\kappa = .77$, respectively). Following this small-scale evaluation, we adopted automated annotation for exploratory analyses of project descriptions. We note that further evaluation is necessary to robustly validate this approach and extend it beyond the exploratory analyses described here.

Prompting parameters. All the annotations were collected using OpenAI’s ChatCompletion API endpoint. Model GPT-4 (‘gpt-4-0613’) was used, with default parameters (default temperature of 1). For reproducibility, we list the complete prompt texts below.

Prompt text: How technically advanced is a project? Which of the following types of data analysis applies to the described activities?

Select the one that applies the most.

- A) Descriptive statistics and data visualization (e.g., Statistical tests, Correlation)
- B) Statistical modeling and inference (e.g., Regression analysis, Logistic regression)
- C) Machine learning and prediction (e.g., Predictive modeling, Clustering)
- D) Causal inference and counterfactuals (e.g., Effect estimation, Matching) Description: <README.md text>

Answer:

Prompt text: How scientifically meaningful is a project? Is the proposed project pushing the boundaries of current scientific knowledge?

Answer YES or NO.

Description: <README.md text>

Answer:

Prompt text: Open-ended adjective generation. List between one and five adjectives that best capture the strengths of this project. Focus on the questions, methods, results, and possible impact. Output a comma-separated list of adjectives. Description: <README.md text>

Answer:

Prompt text: Dataset type. Which data type applies to the described activities? Output a number corresponding to the most relevant data type.

- 1) Tabular data
- 2) Networks
- 3) Textual data
- 4) Other data types

Description: <README.md text>

Answer:

Table C1. Alternative annotation scheme. Data analysis types, their frequencies across the two years, χ^2 statistic, and the corresponding p value.

| Does the data analysis type apply? | Frequency, 2020 | Frequency, 2021 | χ^2 | p value |
|--|-----------------|-----------------|----------|-----------------------|
| A) Descriptive statistics and data visualization | 6.96% | 22.12% | 10.41 | 1.26×10^{-3} |
| B) Statistical modeling and inference | 42.61% | 46.90% | 0.35 | .55 |

| | | | | |
|---|--------|--------|------|-----------------------|
| C) Machine learning and prediction | 42.61% | 30.09% | 4.05 | 4.42×10^{-2} |
| D) Causal inference and counterfactuals | 7.83% | 0.88% | 6.62 | 1.01×10^{-2} |

Adjectives. Complete statistics for all the tested adjectives are listed in Table C2.

Table C2. Adjectives describing the projects. Adjectives, their frequencies across the two years, χ^2 statistic, and the corresponding p value. Terms with statistically significant difference in frequency ($p < .05$) are marked in bold. For completeness, all the terms are listed.

| Adjective | Frequency, 2020 | Frequency, 2021 | χ^2 | p value |
|----------------------|-----------------|-----------------|----------|-----------|
| insightful | 6.09% | 12.81% | 1513.19 | .000100 |
| practical | 1.57% | 0.18% | 638.62 | .011501 |
| methodical | 8.87% | 5.44% | 506.80 | .024372 |
| comprehensive | 13.39% | 18.07% | 472.97 | .029647 |
| detailed | 2.43% | 3.86% | 190.84 | .167144 |
| informative | 0.87% | 1.58% | 119.27 | .274793 |
| inquisitive | 0.17% | 0.53% | 102.11 | .312251 |
| methodological | 0.52% | 0.18% | 98.61 | .320704 |
| quantitative | 0.52% | 0.18% | 98.61 | .320704 |
| robust | 0.52% | 0.18% | 98.61 | .320704 |
| collaborative | 3.13% | 2.28% | 78.46 | .375728 |
| impactful | 9.91% | 11.40% | 66.80 | .413761 |
| innovative | 14.78% | 13.16% | 62.86 | .427885 |
| relevant | 1.57% | 2.11% | 46.37 | .495906 |

| | | | | |
|-------------------|--------|--------|-------|---------|
| strategic | 0.17% | 0.35% | 34.30 | .558086 |
| systematic | 0.17% | 0.35% | 34.30 | .558086 |
| multidisciplinary | 0.35% | 0.18% | 32.55 | .568311 |
| resourceful | 0.35% | 0.18% | 32.55 | .568311 |
| inclusive | 0.35% | 0.18% | 32.55 | .568311 |
| in-depth | 0.52% | 0.35% | 19.22 | .661088 |
| rigorous | 1.04% | 0.88% | 8.32 | .773025 |
| thorough | 1.91% | 2.11% | 5.37 | .816694 |
| analytical | 14.09% | 14.39% | 2.10 | .884885 |
| data-driven | 1.57% | 1.58% | .03 | .985101 |
| forward-thinking | 0.17% | 0.18% | .00 | .995068 |
| detail-oriented | 0.17% | 0.18% | .00 | .995068 |
| timely | 0.17% | 0.18% | .00 | .995068 |
| meaningful | 0.17% | 0.18% | .00 | .995068 |

Data Repository/Code

Anonymized responses and the analysis code necessary to reproduce the results are deposited, and publicly available (Gligoric, 2024).

©2024 Kristina Gligorić, Tiziano Piccardi, Jake M. Hofman, and Robert West. This article is licensed under a [Creative Commons Attribution \(CC BY 4.0\) International license](#), except where otherwise indicated with respect to particular material included in the article.

Footnotes

1. The replication-based project contributed to 25% of the final course grade (the remaining 75% was split between the homework, quizzes, and the final exam). ↵

2. <https://dlab.epfl.ch/teaching/fall2020/cs401/reports/> ↵
3. Preregistration: <https://osf.io/usm4k/> ↵
4. Recall that the group members individually replicated the same paper in the previous steps. ↵
5. <https://go.epfl.ch/ada2021-replic-info-sheet> ↵
6. Pre-registered data analysis plan and survey materials are publicly available: <https://osf.io/usm4k/> ↵
7. Note that our goal is *not* to discover any specific replication failures, but rather to understand if the activity allows discriminating between potential true issues in the original publication from students' own issues. ↵
8. In the year preceding the data analysis replication exercise (2019) the project description had a slightly different structure. We, therefore, opted to analyze 2020 and 2021, two iterations when the structure was consistent ↵