

Meta-research: Replication studies of original “null results” – Absence of evidence or evidence of absence?

Samuel Pawel*, Rachel Heyard*, Charlotte Micheloud, Leonhard Held

* contributed equally

Epidemiology, Biostatistics and Prevention Institute
Center for Reproducible Science
University of Zurich

March 24, 2023

This is a preprint which has not yet been peer reviewed.

Abstract

Absence of evidence is not evidence of absence – the title of the 1995 paper by Douglas Altman and Martin Bland has since become a mantra in the statistical and medical literature. Yet the misinterpretation of statistically non-significant results as evidence for the absence of an effect is still common and further complicated in the context of replication studies. In several large-scale replication projects, non-significant results in both the original and the replication study have been interpreted as a “replication success”. Here we discuss the logical problems with this approach. It does not ensure that the studies provide evidence for the absence of an effect and “replication success” can virtually always be achieved if the sample sizes of the studies are small enough. In addition, the relevant error rates are not controlled. We show how methods, such as equivalence testing and Bayes factors, can be used to adequately quantify the evidence for the absence of an effect and how they can be applied in the replication setting. Using data from the Reproducibility Project: Cancer Biology we illustrate that most original and replication studies with “null results” are in fact inconclusive. We conclude that it is important to also replicate studies with statistically non-significant results, but that they should be designed, analyzed, and interpreted appropriately.

Keywords: Bayesian hypothesis testing, equivalence testing, null hypothesis, replication success

1 Introduction

The misconception that a statistically non-significant result indicates evidence for the absence of an effect is unfortunately widespread (Altman and Bland, 1995). Such a “null result” – typically characterized by a p -value of $p > 5\%$ for the null hypothesis of an absent effect – may also occur if an effect is actually present. For example, if the sample size of a study is chosen to detect an assumed effect with a power of 80%, null results will incorrectly occur 20% of the time when the assumed effect is actually present. Conversely, if the power of the study is lower, null results

will occur more often. In general, the lower the power of a study, the greater the ambiguity of a null result. To put a null result in context, it is therefore critical to know whether the study was adequately powered and under what assumed effect the power was calculated (Hoenig and Heisey, 2001; Greenland, 2012). However, if the goal of a study is to explicitly quantify the evidence for the absence of an effect, more appropriate methods designed for this task, such as equivalence testing or Bayes factors, should ideally be used from the outset.

The contextualization of null results becomes even more complicated in the setting of replication studies. In a replication study, researchers attempt to repeat an original study as closely as possible in order to assess whether similar results can be obtained with new data (National Academies of Sciences, Engineering, and Medicine, 2019). There have been various large-scale replication projects in the biomedical and social sciences in the last decade (Prinz et al., 2011; Begley and Ellis, 2012; Klein et al., 2014; Open Science Collaboration, 2015; Camerer et al., 2016, 2018; Klein et al., 2018; Cova et al., 2018; Errington et al., 2021, among others). Most of these projects suggested alarmingly low replicability rates across a broad spectrum of criteria for quantifying replicability. While most of these projects restricted their focus on original studies with statistically significant results (“positive results”), the *Reproducibility Project: Psychology* (RPP, Open Science Collaboration, 2015), the *Reproducibility Project: Experimental Philosophy* (RPEP, Cova et al., 2018), and the *Reproducibility Project: Cancer Biology* (RPCB, Errington et al., 2021) also attempted to replicate some original studies with null results.

The RPP excluded the original null results from its overall assessment of replication success, but the RPCB and the RPEP explicitly defined null results in both the original and the replication study as a criterion for “replication success”. There are several logical problems with this “non-significance” criterion. First, if the original study had low statistical power, a non-significant result is highly inconclusive and does not provide evidence for the absence of an effect. It is then unclear what exactly the goal of the replication should be – to replicate the inconclusiveness of the original result? On the other hand, if the original study was adequately powered, a non-significant result may indeed provide some evidence for the absence of an effect when analyzed with appropriate methods, so that the goal of the replication is clearer. However, the criterion does not distinguish between these two cases. Second, with this criterion researchers can virtually always achieve replication success by conducting two studies with very small sample sizes, such that the p -values are non-significant and the results are inconclusive. This is because the null hypothesis under which the p -values are computed is misaligned with the goal of inference, which is to quantify the evidence for the absence of an effect. We will discuss methods that are better aligned with this inferential goal in Section 3. Third, the criterion does not control the error of falsely claiming the absence of an effect at some predetermined rate. This is in contrast to the standard replication success criterion of requiring significance from both studies (also known as the two-trials rule, see chapter 12.2.8 in Senn, 2008), which ensures that the error of falsely claiming the presence of an effect is controlled at a rate equal to the squared significance level (for example, $5\% \times 5\% = 0.25\%$ for a 5% significance level). The non-significance criterion may be intended to complement the two-trials rule for null results, but it fails to do so in this respect, which may be important to regulators, funders, and researchers. We will now demonstrate these issues and potential solutions using the null results from the RPCB.

2 Null results from the Reproducibility Project: Cancer Biology

Figure 1 shows standardized mean difference effect estimates with confidence intervals from two RPCB study pairs. Both are “null results” and meet the non-significance criterion for replication success (the two-sided p -values are greater than 5% in both the original and the replication study), but intuition would suggest that these two pairs are very much different.

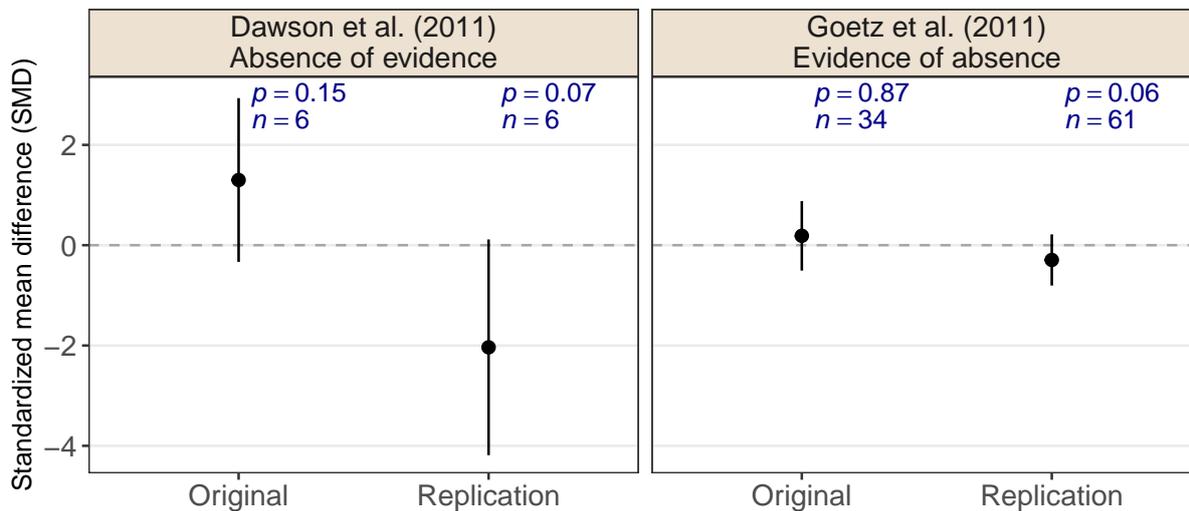


Figure 1: Two examples of original and replication study pairs which meet the non-significance replication success criterion from the Reproducibility Project: Cancer Biology (Errington et al., 2021). Shown are standardized mean difference effect estimates with 95% confidence intervals, total sample size, and p -values for the null hypothesis that the standardized mean difference is zero.

The original study from Dawson et al. (2011) and its replication both show large effect estimates in magnitude, but due to the small sample sizes, the uncertainty of these estimates is very large, too. If the sample sizes of the studies were larger and the point estimates remained the same, intuitively both studies would provide evidence for a non-zero effect. However, with the sample sizes that were actually used, the results seem inconclusive. In contrast, the effect estimates from Goetz et al. (2011) and its replication are much smaller in magnitude and their uncertainty is also smaller because the studies used larger sample sizes. Intuitively, these studies seem to provide some evidence for a zero (or negligibly small) effect. While these two examples show the qualitative difference between absence of evidence and evidence of absence, we will now discuss how the two can be quantitatively distinguished.

3 Methods for assessing replicability of null results

There are both frequentist and Bayesian methods that can be used for assessing evidence for the absence of an effect. Anderson and Maxwell (2016) provide an excellent summary of both approaches in the context of replication studies in psychology. We now briefly discuss two possible approaches – frequentist equivalence testing and Bayesian hypothesis testing – and their application to the RPCB data.

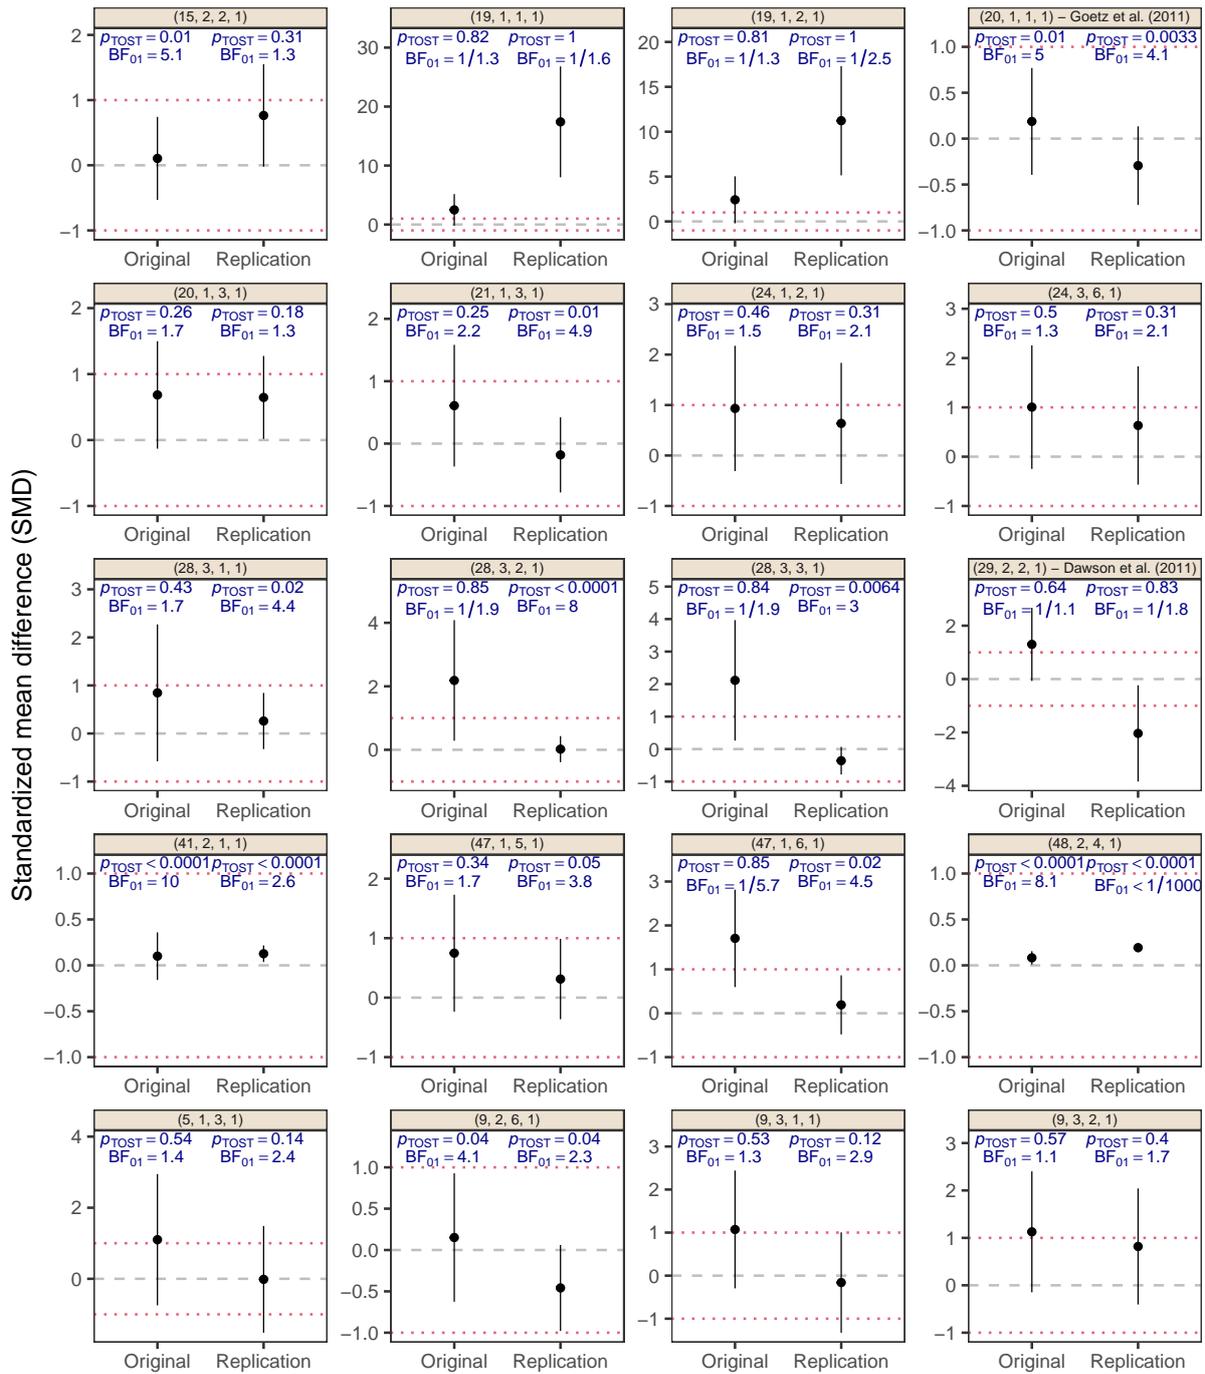


Figure 2: Standardized mean difference (SMD) effect estimates with 90% confidence interval for the “null results” (those with original two-sided p -value $p > 5\%$) and their replication studies from the Reproducibility Project: Cancer Biology (Errington et al., 2021). The identifier above each plot indicates (Original paper number, Experiment number, Effect number, Internal replication number). The two examples from Figure 1 are indicated in the plot titles. The dashed grey line depicts the value of no effect (SMD = 0) whereas the dotted red lines depict the equivalence range with margin $\Delta = 1$. The p -values p_{TOST} are the maximum of the two one-sided p -values for the effect being smaller or greater than $+\Delta$ or $-\Delta$, respectively. The Bayes factors BF_{01} quantify evidence for the null hypothesis $H_0: SMD = 0$ against the alternative $H_1: SMD \neq 0$ with normal unit-information prior assigned to the SMD under H_1 .

3.1 Equivalence testing

Equivalence testing was developed in the context of clinical trials to assess whether a new treatment – typically cheaper or with fewer side effects than the established treatment – is practically equivalent to the established treatment (Westlake, 1972; Schuirmann, 1987). The method can also be used to assess whether an effect is practically equivalent to the value of an absent effect, usually zero. Using equivalence testing as a remedy for non-significant results has been suggested by several authors (Hauck and Anderson, 1986; Campbell and Gustafson, 2018). The main challenge is to specify the margin $\Delta > 0$ that defines an equivalence range $[-\Delta, +\Delta]$ in which an effect is considered as absent for practical purposes. The goal is then to reject the composite null hypothesis that the true effect is outside the equivalence range. To ensure that the null hypothesis is falsely rejected at most $\alpha \times 100\%$ of the time, one either rejects it if the $(1 - 2\alpha) \times 100\%$ confidence interval for the effect is contained within the equivalence range (for example, a 90% confidence interval for $\alpha = 5\%$), or if two one-sided tests (TOST) for the effect being smaller/greater than $+\Delta$ and $-\Delta$ are significant at level α , respectively. A quantitative measure of evidence for the absence of an effect is then given by the maximum of the two one-sided p -values (the TOST p -value).

Returning to the RPCB data, Figure 2 shows the standardized mean difference effect estimates with 90% confidence intervals for the 20 study pairs with quantitative null results in the original study ($p > 5\%$). The dotted red lines represent an equivalence range for the margin $\Delta = 1$, for which the shown TOST p -values are computed. This margin is rather lax compared to the margins typically used in clinical research; we chose it primarily for illustrative purposes and because effect sizes in preclinical research are typically much larger than in clinical research. In practice, the margin should be determined on a case-by-case basis by researchers who are familiar with the subject matter. However, even with this generous margin, only four of the twenty study pairs – one of them being the previously discussed example from Goetz et al. (2011) – are able to establish equivalence at the 5% level in the sense that both the original and the replication 90% confidence interval fall within the equivalence range (or equivalently that their TOST p -values are smaller than 5%). For the remaining 16 studies – for instance, the previously discussed example from Dawson et al. (2011) – the situation remains inconclusive and there is neither evidence for the absence nor the presence of the effect.

3.2 Bayesian hypothesis testing

The distinction between absence of evidence and evidence of absence is naturally built into the Bayesian approach to hypothesis testing. A central measure of evidence is the Bayes factor (Kass and Raftery, 1995), which is the updating factor of the prior odds to the posterior odds of the null hypothesis H_0 versus the alternative hypothesis H_1

$$\underbrace{\frac{\Pr(H_0 | \text{data})}{\Pr(H_1 | \text{data})}}_{\text{Posterior odds}} = \underbrace{\frac{\Pr(H_0)}{\Pr(H_1)}}_{\text{Prior odds}} \times \underbrace{\frac{p(\text{data} | H_0)}{p(\text{data} | H_1)}}_{\text{Bayes factor BF}_{01}}.$$

The Bayes factor quantifies how much the observed data have increased or decreased the probability of the null hypothesis H_0 relative to the alternative H_1 . If the null hypothesis states the absence of an effect, a Bayes factor greater than one ($\text{BF}_{01} > 1$) indicates evidence for the absence of the effect and a Bayes factor smaller than one indicates evidence for the presence of the effect ($\text{BF}_{01} < 1$), whereas a Bayes factor not much different from one indicates absence of evidence for either hypothesis ($\text{BF}_{01} \approx 1$).

When the observed data are dichotomized into positive ($p < 5\%$) or null results ($p > 5\%$), the Bayes factor based on a null result is the probability of observing $p > 5\%$ when the effect is indeed absent (which is 95%) divided by the probability of observing $p > 5\%$ when the effect is indeed present (which is one minus the power of the study). For example, if the power is 90%, we have $\text{BF}_{01} = 95\%/10\% = 9.5$ indicating almost ten times more evidence for the absence of the effect than for its presence. On the other hand, if the power is only 50%, we have $\text{BF}_{01} = 95\%/50\% = 1.9$ indicating only slightly more evidence for the absence of the effect. This example also highlights the main challenge with Bayes factors – the specification of the alternative hypothesis H_1 . The assumed effect under H_1 is directly related to the power of the study, and researchers who assume different effects under H_1 will end up with different Bayes factors. Instead of specifying a single effect, one therefore typically specifies a “prior distribution” of plausible effects. Importantly, the prior distribution, like the equivalence margin, should be determined by researchers with subject knowledge and before the data are observed.

In practice, the observed data should not be dichotomized into positive or null results, as this leads to a loss of information. Therefore, to compute the Bayes factors for the RPCB null results, we used the observed effect estimates as the data and assumed a normal sampling distribution for them, as in a meta-analysis. The Bayes factors BF_{01} shown in Figure 2 then quantify the evidence for the null hypothesis of no effect ($H_0: \text{SMD} = 0$) against the alternative hypothesis that there is an effect ($H_1: \text{SMD} \neq 0$) using a normal “unit-information” prior distribution (Kass and Wasserman, 1995) for the effect size under the alternative H_1 . There are several more advanced prior distributions that could be used here, and they should ideally be specified for each effect individually based on domain knowledge. The normal unit-information prior (with a standard deviation of 2 for SMDs) is only a reasonable default choice, as it implies that small to large effects are plausible under the alternative. We see that in most cases there is no substantial evidence for either the absence or the presence of an effect, as with the equivalence tests. The Bayes factors for the two previously discussed examples from Goetz et al. (2011) and Dawson et al. (2011) are consistent with our intuitions – there is indeed some evidence for the absence of an effect in Goetz et al. (2011), while there is even slightly more evidence for the presence of an effect in Dawson et al. (2011), though the Bayes factor is very close to one due to the small sample sizes. If we use a lenient Bayes factor threshold of $\text{BF}_{01} > 3$ to define evidence for the absence of the effect, only one of the twenty study pairs meets this criterion in both the original and replication study.

Among the twenty RPCB null results, there is one interesting case (the rightmost plot in the fourth row (48, 2, 4, 1)) where the Bayes factor is qualitatively different from the equivalence test, revealing a fundamental difference between the two approaches. The Bayes factor is concerned with testing whether the effect is *exactly zero*, whereas the equivalence test is concerned with

whether the effect is within an *interval around zero*. Due to the very large sample size in this replication study, the data are incompatible with an exactly zero effect, but compatible with effects within the equivalence range. Apart from this example, however, the approaches lead to the same qualitative conclusion – most RPCB null results are highly ambiguous.

4 Conclusions

We showed that in most of the RPCB studies with “null results” (those with $p > 5\%$), neither the original nor the replication study provided conclusive evidence for the presence or absence of an effect. It seems logically questionable to declare an inconclusive replication of an inconclusive original study as a replication success. While it is important to replicate original studies with null results, our analysis highlights that they should be analyzed and interpreted appropriately.

For both the equivalence testing and the Bayes factor approach, it is critical that the parameters of the procedure (the equivalence margin and the prior distribution) are specified independently of the data, ideally before the studies are conducted. Typically, however, the original studies were designed to find evidence for the presence of an effect, and the goal of replicating the “null result” was formulated only after failure to do so. [Campbell and Gustafson \(2021\)](#) discuss various approaches to post-hoc specification of equivalence margins, such as motivating it using data from previous studies or using field conventions. [Hauck and Anderson \(1986\)](#) propose a sensitivity analysis approach in the form of plotting the TOST p -value against a range of possible margins (“equivalence curves”). Post-hoc specification of a prior distribution for a Bayes factor may likewise be based on historical data, field conventions, or assessed visually with sensitivity analyses.

While the equivalence test and the Bayes factor are two principled methods for analyzing original and replication studies with null results, they are not the only possible methods for doing so. For instance, the reverse-Bayes approach from [Micheloud and Held \(2022b\)](#) specifically tailored to equivalence testing in the replication setting may lead to more appropriate inferences as it also takes into account the compatibility of the effect estimates from original and replication studies. In addition, there are various other Bayesian methods which could potentially improve upon the considered Bayes factor approach. For example, Bayes factors based on non-local priors ([Johnson and Rossell, 2010](#)) or based on interval null hypotheses ([Morey and Rouder, 2011](#); [Liao et al., 2020](#)), methods for equivalence testing based on effect size posterior distributions ([Kruschke, 2018](#)), or Bayesian procedures that involve utilities of decisions ([Lindley, 1998](#)). Finally, the design of replication studies should align with the planned analysis ([Anderson and Maxwell, 2017](#); [Anderson and Kelley, 2022](#); [Micheloud and Held, 2022a](#); [Pawel et al., 2022](#)). If the goal of the study is to find evidence for the absence of an effect, the replication sample size should also be determined so that the study has adequate power to make conclusive inferences regarding the absence of the effect.

Acknowledgements

We thank the contributors of the RPCB for their tremendous efforts and for making their data publicly available. We thank Maya Mathur for helpful advice with the data preparation. This work was supported by the Swiss National Science Foundation (grant #189295).

Conflict of interest

We declare no conflict of interest.

Data and software

The data from the RPCB were obtained by downloading the files from <https://github.com/mayamathur/rpcb> (commit a1e0c63) and executing the R script `Code/data_prep.R` with the line 632 commented out so that also original studies with null results are included. This then produced the file `prepped_outcome_level_data.csv` which was used for the subsequent analyses. The effect estimates and standard errors on SMD scale provided in this data set differ in some cases from those in the data set available at <https://doi.org/10.17605/osf.io/e5nvr>, which is cited in [Errington et al. \(2021\)](#). We used this particular version of the data set because it was recommended to us by the RPCB statistician (Maya Mathur) upon request. For the [Dawson et al. \(2011\)](#) example study and its replication ([Shan et al., 2017](#)), the sample sizes $n = 3$ in the data set seem to correspond to the group sample sizes, see Figure 5A in the replication study (<https://doi.org/10.7554/eLife.25306.012>), which is why we report the total sample sizes of $n = 6$ in Figure 1.

The code and data to reproduce our analyses is openly available at <https://gitlab.uzh.ch/samuel.pawel/rsAbsence>. A snapshot of the repository at the time of writing is available at <https://doi.org/10.5281/zenodo.XXXXXX>. We used the statistical programming language R version 4.2.0 ([R Core Team, 2022](#)) for analyses. The R packages `ggplot2` ([Wickham, 2016](#)), `dplyr` ([Wickham et al., 2022](#)), `knitr` ([Xie, 2022](#)), and `reporttools` ([Rufibach, 2009](#)) were used for plotting, data preparation, dynamic reporting, and formatting, respectively.

References

- Altman, D. G. and Bland, J. M. (1995). Statistics notes: Absence of evidence is not evidence of absence. *BMJ*, 311(7003):485–485. doi:10.1136/bmj.311.7003.485.
- Anderson, S. F. and Kelley, K. (2022). Sample size planning for replication studies: The devil is in the design. *Psychological Methods*. doi:10.1037/met0000520.
- Anderson, S. F. and Maxwell, S. E. (2016). There’s more than one way to conduct a replication study: Beyond statistical significance. *Psychological Methods*, 21(1):1–12. doi:10.1037/met0000051.

- Anderson, S. F. and Maxwell, S. E. (2017). Addressing the “replication crisis”: Using original studies to design replication studies with appropriate statistical power. *Multivariate Behavioral Research*, 52(3):305–324. doi:10.1080/00273171.2017.1289361.
- Begley, C. G. and Ellis, L. M. (2012). Raise standards for preclinical cancer research. *Nature*, 483(7391):531–533. doi:10.1038/483531a.
- Camerer, C. F., Dreber, A., Forsell, E., Ho, T., Huber, J., Johannesson, M., Kirchler, M., Almenberg, J., Altmejd, A., et al. (2016). Evaluating replicability of laboratory experiments in economics. *Science*, 351:1433–1436. doi:10.1126/science.aaf0918.
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T., Huber, J., Johannesson, M., Kirchler, M., Nave, G., Nosek, B., et al. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behavior*, 2:637–644. doi:10.1038/s41562-018-0399-z.
- Campbell, H. and Gustafson, P. (2018). Conditional equivalence testing: An alternative remedy for publication bias. *PLOS ONE*, 13(4):e0195145. doi:10.1371/journal.pone.0195145.
- Campbell, H. and Gustafson, P. (2021). What to make of equivalence testing with a post-specified margin? *Meta-Psychology*, 5. doi:10.15626/mp.2020.2506.
- Cova, F., Strickland, B., Abatista, A., Allard, A., Andow, J., Attie, M., Beebe, J., Berniūnas, R., Boudesseul, J., Colombo, M., et al. (2018). Estimating the reproducibility of experimental philosophy. *Review of Philosophy and Psychology*. doi:10.1007/s13164-018-0400-9.
- Dawson, M. A., Prinjha, R. K., Dittmann, A., Giotopoulos, G., Bantscheff, M., Chan, W.-I., Robson, S. C., wa Chung, C., Hopf, C., Savitski, M. M., Huthmacher, C., Gudgin, E., Lugo, D., Beinke, S., Chapman, T. D., Roberts, E. J., Soden, P. E., Auger, K. R., Mirguet, O., Doehner, K., Delwel, R., Burnett, A. K., Jeffrey, P., Drewes, G., Lee, K., Huntly, B. J. P., and Kouzarides, T. (2011). Inhibition of BET recruitment to chromatin as an effective treatment for MLL-fusion leukaemia. *Nature*, 478(7370):529–533. doi:10.1038/nature10509.
- Errington, T. M., Mathur, M., Soderberg, C. K., Denis, A., Perfito, N., Iorns, E., and Nosek, B. A. (2021). Investigating the replicability of preclinical cancer biology. *eLife*, 10. doi:10.7554/elife.71601.
- Goetz, J. G., Minguet, S., Navarro-Lérida, I., Lazcano, J. J., Samaniego, R., Calvo, E., Tello, M., Osteso-Ibáñez, T., Pellinen, T., Echarri, A., Cerezo, A., Klein-Szanto, A. J., Garcia, R., Keely, P. J., Sánchez-Mateos, P., Cukierman, E., and Pozo, M. A. D. (2011). Biomechanical remodeling of the microenvironment by stromal caveolin-1 favors tumor invasion and metastasis. *Cell*, 146(1):148–163. doi:10.1016/j.cell.2011.05.040.
- Greenland, S. (2012). Nonsignificance plus high power does not imply support for the null over the alternative. *Annals of Epidemiology*, 22(5):364–368. doi:10.1016/j.annepidem.2012.02.007.

- Hauck, W. W. and Anderson, S. (1986). A proposal for interpreting and reporting negative studies. *Statistics in Medicine*, 5(3):203–209. doi:10.1002/sim.4780050302.
- Hoenig, J. M. and Heisey, D. M. (2001). The abuse of power. *The American Statistician*, 55(1):19–24. doi:10.1198/000313001300339897.
- Johnson, V. E. and Rossell, D. (2010). On the use of non-local prior densities in Bayesian hypothesis tests. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 72(2):143–170. doi:10.1111/j.1467-9868.2009.00730.x.
- Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795. doi:10.1080/01621459.1995.10476572.
- Kass, R. E. and Wasserman, L. (1995). A reference Bayesian test for nested hypotheses and its relationship to the Schwarz criterion. *Journal of the American Statistical Association*, 90(431):928–934. doi:10.1080/01621459.1995.10476592.
- Klein, R. A., Ratliff, K. A., Vianello, M., Adams, R. B., Bahník, v., Bernstein, M. J., Bocian, K., Brandt, M. J., Brooks, B., et al. (2014). Investigating variation in replicability: A “many labs” replication project. *Social Psychology*, 45:142–152. doi:10.1027/1864-9335/a000178.
- Klein, R. A., Vianello, M., Hasselman, F., Adams, B. G., Reginald B. Adams, J., Alper, S., Aveyard, M., Axt, J. R., Babalola, M. T., et al. (2018). Many labs 2: Investigating variation in replicability across samples and settings. *Advances in Methods and Practices in Psychological Science*, 1(4):443–490. doi:10.1177/2515245918810225.
- Kruschke, J. K. (2018). Rejecting or accepting parameter values in Bayesian estimation. *Advances in Methods and Practices in Psychological Science*, 1(2):270–280. doi:10.1177/2515245918771304.
- Liao, J. G., Midya, V., and Berg, A. (2020). Connecting and contrasting the Bayes factor and a modified ROPE procedure for testing interval null hypotheses. *The American Statistician*, 75(3):256–264. doi:10.1080/00031305.2019.1701550.
- Lindley, D. V. (1998). Decision analysis and bioequivalence trials. *Statistical Science*, 13(2). doi:10.1214/ss/1028905932.
- Micheloud, C. and Held, L. (2022a). Power calculations for replication studies. *Statistical Science*, 37(3):369–379. doi:10.1214/21-sts828.
- Micheloud, C. and Held, L. (2022b). The replication of non-inferiority and equivalence studies. doi:10.48550/ARXIV.2204.06960. arXiv preprint.
- Morey, R. D. and Rouder, J. N. (2011). Bayes factor approaches for testing interval null hypotheses. *Psychological Methods*, 16(4):406–419. doi:10.1037/a0024377.
- National Academies of Sciences, Engineering, and Medicine (2019). *Reproducibility and Replicability in Science*. National Academies Press. doi:10.17226/25303. URL <https://doi.org/10.17226/25303>.

- Open Science Collaboration (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251):aac4716. doi:10.1126/science.aac4716.
- Pawel, S., Consonni, G., and Held, L. (2022). Bayesian approaches to designing replication studies. doi:10.48550/ARXIV.2211.02552. arXiv preprint.
- Prinz, F., Schlange, T., and Asadullah, K. (2011). Believe it or not: how much can we rely on published data on potential drug targets? *Nature Reviews Drug Discovery*, 10(9):712–712. doi:10.1038/nrd3439-c1.
- R Core Team (2022). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Rufibach, K. (2009). reporttools: R functions to generate L^AT_EX tables of descriptive statistics. *Journal of Statistical Software, Code Snippets*, 31(1). doi:10.18637/jss.v031.c01.
- Schuirman, D. J. (1987). A comparison of the two one-sided tests procedure and the power approach for assessing the equivalence of average bioavailability. *Journal of Pharmacokinetics and Biopharmaceutics*, 15(6):657–680. doi:10.1007/bf01068419.
- Senn, S. (2008). *Statistical issues in drug development*, volume 69. John Wiley & Sons.
- Shan, X., Fung, J. J., Kosaka, A., and and, G. D.-D. (2017). Replication study: Inhibition of BET recruitment to chromatin as an effective treatment for MLL-fusion leukaemia. *eLife*, 6. doi:10.7554/elife.25306.
- Westlake, W. J. (1972). Use of confidence intervals in analysis of comparative bioavailability trials. *Journal of Pharmaceutical Sciences*, 61(8):1340–1341. doi:10.1002/jps.2600610845.
- Wickham, H. (2016). *ggplot2: Elegant Graphics for Data Analysis*. Springer International Publishing. doi:10.1007/978-3-319-24277-4.
- Wickham, H., François, R., Henry, L., and Müller, K. (2022). *dplyr: A Grammar of Data Manipulation*. URL <https://CRAN.R-project.org/package=dplyr>. R package version 1.0.10.
- Xie, Y. (2022). *knitr: A General-Purpose Package for Dynamic Report Generation in R*. URL <https://yihui.org/knitr/>. R package version 1.40.

Computational details

```
cat(paste(Sys.time(), Sys.timezone(), "\n"))

## 2023-03-24 09:59:39 Etc/UTC

sessionInfo()

## R version 4.2.0 (2022-04-22)
## Platform: x86_64-pc-linux-gnu (64-bit)
## Running under: Ubuntu 20.04.4 LTS
##
## Matrix products: default
## BLAS: /usr/lib/x86_64-linux-gnu/openblas-pthread/libblas.so.3
## LAPACK: /usr/lib/x86_64-linux-gnu/openblas-pthread/liblapack.so.3
##
## locale:
## [1] LC_CTYPE=en_US.UTF-8 LC_NUMERIC=C
## [3] LC_TIME=en_US.UTF-8 LC_COLLATE=en_US.UTF-8
## [5] LC_MONETARY=en_US.UTF-8 LC_MESSAGES=en_US.UTF-8
## [7] LC_PAPER=en_US.UTF-8 LC_NAME=C
## [9] LC_ADDRESS=C LC_TELEPHONE=C
## [11] LC_MEASUREMENT=en_US.UTF-8 LC_IDENTIFICATION=C
##
## attached base packages:
## [1] stats graphics grDevices utils datasets methods base
##
## other attached packages:
## [1] reporttools_1.1.3 xtable_1.8-4 dplyr_1.0.9 ggplot2_3.3.6
## [5] knitr_1.39
##
## loaded via a namespace (and not attached):
## [1] magrittr_2.0.3 tidyselect_1.1.2 munsell_0.5.0 colorspace_2.0-3
## [5] R6_2.5.1 rlang_1.0.2 fansi_1.0.3 highr_0.9
## [9] stringr_1.4.0 tools_4.2.0 grid_4.2.0 gtable_0.3.0
## [13] xfun_0.31 utf8_1.2.2 DBI_1.1.2 cli_3.3.0
## [17] withr_2.5.0 ellipsis_0.3.2 digest_0.6.29 assertthat_0.2.1
## [21] tibble_3.1.7 lifecycle_1.0.1 crayon_1.5.1 farver_2.1.0
## [25] purrr_0.3.4 vctrs_0.4.1 glue_1.6.2 evaluate_0.15
## [29] labeling_0.4.2 stringi_1.7.6 compiler_4.2.0 pillar_1.7.0
## [33] generics_0.1.2 scales_1.2.0 pkgconfig_2.0.3
```