# LETTER



# Decoupling Visualization and Testing when Presenting Confidence Intervals

David A. Armstrong II<sup>1</sup> and William Poirier<sup>2</sup>

<sup>1</sup>Professor, Canada Research Chair in Political Methodology, Department of Political Science, Western University, London, Ontario, Canada; <sup>2</sup>Ph.D. Student, Department of Political Science, Western University, London, Ontario, Canada

Corresponding author: David A. Armstrong II; Email: dave.armstrong@uwo.ca

(Received 15 May 2024; revised 17 August 2024; accepted 25 September 2024)

### Abstract

Confidence intervals are ubiquitous in the presentation of social science models, data, and effects. When several intervals are plotted together, one natural inclination is to ask whether the estimates represented by those intervals are significantly different from each other. Unfortunately, there is no general rule or procedure that would allow us to answer this question from the confidence intervals alone. It is well known that using the overlaps in 95% confidence intervals to perform significance tests at the 0.05 level does not work. Recent scholarship has developed and refined a set of tools for *inferential confidence intervals* that permit inference on confidence intervals with the appropriate type I error rate in many different bivariate contexts. These are all based on the same underlying idea of identifying the multiple of the standard error (i.e., a new confidence level) such that the overlap in confidence intervals matches the desired type I error rate. These procedures remain stymied by multiple simultaneous comparisons. We propose an entirely new procedure for developing inferential confidence intervals that decouples the testing and visualization that can overcome many of these problems in any visual testing scenario. We provide software in R and Stata to accomplish this goal.

Keywords: confidence intervals; visual testing; statistical inference

Edited by: Jeff Gill

# 1. The Problem

Confidence intervals are ubiquitous in the presentation of data, models, and effects in the social sciences.<sup>1</sup> Consider, for example, Gibson's Gibson's (2024) Figure 1(b) where he shows the proportion of people agreeing with the proposition that it might be better to do away with the US Supreme Court when it starts making decisions that most people disagree with. We reproduce this figure below in Figure 1.

While we can see how the proportion agreeing decreases from mid to late 2020 and then increases into July 2022, we might wonder which of these estimates are different from the others. For example, is the change from July 2020 to March 2021 significant? Ideally, we could look at whether the confidence intervals for the two estimates overlap—if they do, the difference between the two estimates is not significant, if they do not, the difference is significant. Unfortunately, this will not always lead us to the right conclusion. Often, the difference between two confidence intervals that greatly overlap

<sup>&</sup>lt;sup>1</sup>In 2024, volume 2 of *The American Political Science Review*, 57% of their research articles (72% of quantitative articles) used graph presenting a point estimate surrounded by confidence bounds to present treatment effects, predicted probabilities, and regression coefficients.

<sup>©</sup> The Author(s), 2025. Published by Cambridge University Press on behalf of The Society for Political Methodology.

This is an Open Access article, distributed under the terms of the Creative Commons Attribution-ShareAlike licence (https://creativecommons.org/ licenses/by-sa/4.0), which permits re-use, distribution, and reproduction in any medium, provided the same Creative Commons licence is used to distribute the re-used or adapted article and the original article is properly cited.



Figure 1. Proportion agreeing -do away with the Supreme Court for unpopular decisions.

is insignificant. Conversely, the difference between two confidence intervals that do not overlap is significant. The problem lies in between those extremes. Appendix 1 of the Supplementary Material describes the problem in greater detail. The 95% confidence intervals for July 2020 and March 2021 do overlap and the difference between those two proportions is not significant. However, the confidence intervals for July 2020 and July 2022 also overlap, but those two estimates are statistically different from each other.<sup>2</sup> What we know is that when two 95% confidence intervals do not overlap, the difference between the two estimates is significant, but when the intervals overlap somewhat, we cannot necessarily conclude that the two estimates represented by the two confidence intervals are statistically indistinguishable from each other (Browne 1979; Radean 2023; Schenker and Gentleman 2001). How then can we allow readers to visually compare confidence intervals while retaining their statistical properties as much as possible?

# 1.1. Previous Attempts at Visual Testing

We refer to the procedure of judging the statistical significance of the difference between two estimates by whether their confidence intervals overlap as *visual testing*. Below, we use the term *inferential confidence interval*, coined by Tyron (2001), to refer to  $(1-\gamma) \times 100\%$  confidence intervals whose (non-) overlaps correspond with tests at the level  $\alpha$ , where generally  $\gamma \neq \alpha$ . For nearly a century, scholars have grappled with the idea that readers will try to make inferences about differences in estimates using confidence intervals (Dice and Laraas 1936; Simpson and Roe 1939).<sup>3</sup> The first systematic scholarship here identified 84% confidence intervals as useful to represent tests at the 5% level in a narrow set of

<sup>&</sup>lt;sup>2</sup>We are calculating statistical significance with a pairwise *t*-test:  $t = \frac{\bar{x}_2 - \bar{x}_1}{\sqrt{\operatorname{var}(\bar{x}_2) + \operatorname{var}(\bar{x}_1) - 2\operatorname{cov}(\bar{x}_2, \bar{x}_1)}}$ 

<sup>&</sup>lt;sup>3</sup>This bears some resemblance to the so-called reference category problem. We discuss the similarities and differences in Appendix 2 of the Supplementary Material.

circumstances (Browne 1979; Tukey 1991).<sup>4</sup> The ensuing decades saw several attempts to generalize this procedure to take account of differences in sample size, uncertainty, covariance and the functional form of the distribution of the difference (Afshartous and Preston 2010; Goldstein and Healy 1995; Payton, Greenstone, and Schenker 2003; Radean 2023).

The main takeaway from the line of research discussed above is that for approximately normally distributed estimates, we can identify an appropriate confidence level  $(1 - \gamma)$  such that two confidence intervals do not overlap with a given probability,  $\alpha$ , under the null hypothesis. Following Afshartous and Preston (2010) and Radean (2023), we compute the inferential confidence level required for any pair of estimates.

$$Z_{\gamma} = \left[\frac{F^{-1}\left(\frac{\alpha}{2}\right)}{\frac{\theta}{\sqrt{\theta^2 + 1 - 2\rho\theta}} + \frac{\frac{1}{\theta}}{\sqrt{1 + \theta^{-2} - 2\rho\theta^{-1}}}}\right]$$
(1.1)

$$\Pr(\text{Overlap}) = 2\left(1 - F\left(Z_{\gamma}\frac{\theta}{\sqrt{\theta^2 + 1 - 2\rho\theta}} + \frac{\frac{1}{\theta}}{\sqrt{1 + \theta^{-2} - 2\rho\theta^{-1}}}\right)\right)$$
(1.2)

where  $\theta$  is the ratio of standard errors for the two estimates,  $\rho$  is the correlation between the two estimates,  $Z_{\gamma}$  is the multiplier for the standard error for the inferential confidence level  $(1 - \gamma)$ , F() is the CDF of the appropriate *t* or normal distribution and  $F^{-1}()$  is its quantile function.

For any pair of estimates, the appropriate value of  $Z_{\gamma}$  will differ depending on the ratio of their variances and their covariance. To the extent that we are trying to find a single value of  $Z_{\gamma}$  that appropriately represents all pairwise tests, this variation is problematic. If multiple pairs are present, Afshartous and Preston (2010) would have us average over the values of  $Z_{\gamma}$ . In all but the most optimistic cases, the tests produced will not all have the same type I error rate. This leaves the user in essentially the same situation in which she started—not knowing whether estimates are different at a particular level of  $\alpha$ . Below, we develop a procedure that 1) works on an arbitrarily large collection of intervals, 2) directly identifies tests that are not appropriately captured by the inferential confidence intervals, and 3) is agnostic to inferential paradigm.<sup>5</sup>

### 2. Inferential Confidence Intervals

The main goal of previous research in this field is to identify an inferential confidence level  $(1-\gamma)$  such that intervals overlap with probability  $\alpha$  under the null hypothesis. This couples the testing procedure to the visualization. Our innovation is to decouple the testing from the visualization. For most quantities of interest in social science, it is easy to compute pairwise tests of difference. We suggest using the appropriate tests to make pairwise inferences and then attempt to identify the inferential confidence level (or levels) such that overlapping intervals correspond with insignificant differences and non-overlapping intervals correspond with significant differences. Since we are not using the confidence intervals to do the test, the probability with which any pair of intervals overlaps under the null hypothesis is of no real concern.

Importantly, the main beneficiary of this kind of display is not the researcher. We imagine that researchers will have a good sense, through investigation of their models, which of their intended inferences are significant and which are not. Instead, this tool empowers readers to use published results to make valid inferences about comparisons that may not have been anticipated by or important to the

<sup>&</sup>lt;sup>4</sup>Appendix 3 of the Supplementary Material demonstrates why 84% confidence intervals may not always produces the desired result.

<sup>&</sup>lt;sup>5</sup>This article is not a defense of the NHST, it only acknowledges that what may be a logically flawed practice still dominates statistical decision-making in our field and others—see Gill (1999) for a comprehensive critique. We demonstrate how this procedure works in the Bayesian context in Appendix 5 of the Supplementary Material.

researcher. If we acknowledge that readers are already engaging in this kind of visual analysis as a matter of course, our procedure will allow them to do so in a constructive and inferentially valid fashion.

In any situation where visual testing is desirable, the following steps may be implemented.

- 1. **Conduct all pairwise tests between estimates.** Use whatever method you like to identify significant/interesting from insignificant/uninteresting differences.<sup>6</sup> The tests could include a reference estimate of zero, with sampling variability equal to zero and with zero covariance with all other estimates. This would ensure that all univariate tests against zero are also respected by the procedure. Once the appropriate *p*-values ( $p_{ij}$ 's) are calculated for  $b_j b_i \forall i < j$ , we define  $s_{ij}$  as 1 if  $p_{ij} < \alpha$  (the desired type I error rate for the test) and 0 otherwise; s is a vector of  $s_{ij}$  values.<sup>7</sup>
- 2. Find the inferential confidence level(s). Once the baseline results of pairwise tests are computed, we find  $(1-\gamma)$  (the inferential confidence level) as the solution(s) to the following optimization:

$$\underset{(1-\gamma)}{\operatorname{arg\,max}} \sum_{j=2}^{J} \sum_{i=1}^{j-1} I(s_{ij} = s_{ij}^{*})$$
(2.1)

where  $s_{ij}^*$  is 0 if the  $(1 - \gamma) \times 100\%$  confidence intervals for  $b_i$  and  $b_j$  overlap and 1 if they do not (where **b**, the vector of estimates, is ordered from largest to smallest). That is, we find the value(s) of  $\gamma$  that maximizes the agreement between  $s_{ij}$  the "correct" indicator of significance for the difference between  $b_i$  and  $b_j$  based on a pairwise test and  $s_{ij}^*$ , the indicator of (non-)overlapping of the inferential confidence intervals for  $b_i$  and  $b_j$ .

3. Pick the inferential confidence level that is most useful. If there are multiple levels identified by step 2, we should try to identify which one is most useful. In Appendix 4 of the Supplementary Material, we discuss several different options, but any of the identified values would work. As a good place to start, we suggest using the value halfway between smallest and largest acceptable values.

If we can find a level such that *all* the pairwise tests are appropriately represented by whether or not the inferential confidence intervals overlap, then using that interval in a coefficient plot or similar display would produce the desired result—readers could easily identify whether pairs of estimates are different from each other based on whether or not the intervals overlap. We produced software in R and Stata to perform these calculations easily after most models. The vignettes and help pages for the software provide examples and guides for use and interpretation. A brief software demonstration for both R and Stata can be found in Appendix 7 of the Supplementary Material.

#### 3. Case Study: Iyengar and Westwood (2014)

Below, we describe a case where the inferential confidence intervals provide much more clarity in testing.<sup>8</sup> Iyengar and Westwood (2015) provide implicit, explicit, and behavioral indicators of affective polarization in the US. Of interest here is their second experiment where they explicitly asked respondents to choose a high school senior to receive a scholarship. The GPA of the two students were randomly varied (either 3.5 or 4.0). Each student was either identified as being president of the Young Republicans or Young Democrats as a partisan identity marker. The two students could either be equally qualified (with a 3.5 or 4.0 GPA) or one could be more qualified than the other. The authors construct a binary

<sup>&</sup>lt;sup>6</sup>These tests could employ clustered/robust standard errors, multiplicity adjustments etc... Explaining these various inferential tweaks is beyond the scope of this article. We mention them to ensure readers that our procedure is compatible with any kind of pairwise testing. See Bretz, Hothorn, and Westfall (2010) for a discussion of multiplicity adjustments.

<sup>&</sup>lt;sup>7</sup>In the notation above  $b_i$  and  $b_j$  as any estimates for which pairwise differences can be calculated.

<sup>&</sup>lt;sup>8</sup>In the interest of space, we present one case study here. Appendix 6 of the Supplementary Material demonstrates another example where we apply the methodology we describe to a Bayesian analysis.



**Figure 2.** Iyengar and Westwood (2015)'s predicted probabilities for Partisan Winner selection. Note:Inferential confidence intervals at 84% level visually representing results of all 95% level pairwise tests between same treatment party ID. Red x's mark comparisons where the overlaps and test results diverge.

dependent variable for whether the respondent chose the Democrat candidate (0) or the Republican candidate (1), a treatment condition variable with three levels (Democrat more qualified, both equally qualified, Republican more qualified), and the respondent's partisan identification (Democrat, Lean Democrat, Independent, Lean Republican, Republican). Figure 6 of Iyengar and Westwood (2015) presents the predicted probabilities computed from a logistic regression model where the respondent's choice is regressed on an interaction between the respondent's partisan identification and her received treatment condition.

The top panels of Figure 2 reproduces Iyengar and Westwood (2015)'s results with their original 95% confidence intervals. There are five estimates and 10 possible pairwise tests in each panel resulting in 30 pairwise tests of potential interest. Using 95% confidence intervals to do these tests, we would get seven of them wrong. Using our method, we find that there are a range of confidence intervals that perfectly account for all 30 tests (Equally Qualified: [0.590, 0.863], Republican More Qualified: [0.817, 0.878], Democrat Most Qualified: [0.744, 0.869]). All tests across all three panels can be accommodated by any value in the range [.817, .863]. This range includes (and is centered on) the 84% confidence interval that is often used.<sup>9</sup> Our procedure ensures the appropriate type I error rate then optimizes the display to correspond with the tests.

For example, we can see from the inferential confidence intervals that the estimates for Republicans and Independents are statistically different when the Republican was most qualified, but the 95% intervals overlap.

<sup>&</sup>lt;sup>9</sup>Using the methods described in previous research, the 84% confidence intervals would have type I error rates between 4.7% and 7.9% depending on the comparison. While the 84% intervals work, they will not always be appropriate.

## 4. Conclusion

There is a long literature in statistics and more recently in political science that identifies inferential confidence intervals—intervals meant to permit visual testing with the desired type I error rate. Despite continued development and refinement, all attempts continue to suffer from an important flaw—the inferential intervals are only defined for a single pair and vary (perhaps interestingly) across pairs of estimates from the same analysis. Rather than grafting an inferential framework onto the overlaps of confidence intervals, our approach is to focus on the full set of tests to identify the inferential confidence level that maximally corresponds with properly done pairwise tests. Our approach is sufficiently flexible that "properly done" could take on any number of meanings across inferential paradigms, multiplicity corrections and operationalizations of the variance–covariance matrix of the estimates. When tests are inappropriately characterized by our procedure, those tests are directly identified and can be flagged by the analyst; see Appendix 6 of the Supplementary Material for an example. Ultimately, this benefits readers who will be able to make valid inferences about comparisons that may not have been anticipated by the researcher.

Acknowledgments. We would like to thank Matt Lebo, Ryan Bakker, and Arthur Spirling for helpful comments and suggestions.

Funding. This project was funded by the Social Sciences and Humanities Research Council of Canada, grant CRC-2022-00299.

Data Availability Statement. Replication code and data for this article have been published in the Political Analysis Dataverse at https://doi.org/10.7910/DVN/GFLSLH (dataverse).

Supplementary Material. For supplementary material accompanying this paper, please visit https://doi.org/10.1017/pan.2024.24.

### References

- Afshartous, D., and R. A. Preston. 2010. "Confidence Intervals for Dependent Data: Equating Non-Overlap with Statistical Significance." *Computational Statistics and Data Analysis* 54: 2296–2305.
- Armstrong II, D. A., and W. Poirier. 2024. "Replication Data for: Decoupling Visualization and Testing when Presenting Confidence Intervals." Version DRAFT VERSION. https://doi.org/10.7910/DVN/GFLSLH.
- Bretz, F., T. Hothorn, and P. Westfall. 2010. Multiple Comparisons Using R. Boca Raton, FL: Chapman & Hall.
- Browne, R. H. 1979. "On Visual Assessment of the Significance of a Mean Difference." Biometrics 35 (3): 657-665.
- Dice, L., and H. Laraas. 1936. "A Graphic Method for Comparing Several Sets of Measurements." *Contributions from the Lab* of Vertebrate Genetics 3: 1–3.
- Gibson, J. L. 2024. "Losing Legitimacy: The Challenges of the Dobbs Ruling to Conventional Legitimacy Theory." *American Journal of Political Science*. 68 (3): 1041–1056.
- Gill, J. 1999. "The Insignificance of Null Hypothesis Significance Testing." Political Research Quarterly 52 (3): 647-674.
- Goldstein, H., and M. J. Healy. 1995. "The Graphical Presentation of a Collection of Means." *Journal of the Royal Statistical Society, Series A* 158 (1): 175–177.
- Iyengar, S., and S. J. Westwood. 2015. "Fear and Loathing across Party Lines: New Evidence on Group Polarization." American Journal of Political Science 59 (3): 690–707.

Payton, M. E., M. H. Greenstone, and N. Schenker. 2003. "Overlapping Confidence Intervals or Standard Error Intervals: What do they Mean in Terms of Statistical Significance?" *Journal of Insect Science* 3 (1): 34.

- Radean, M. 2023. "The Significance of Differences Interval: Assessing the Statistical and Substantive Difference between Two Quantities of Interest." *Journal of Politics* 85 (3): 969–983.
- Schenker, N., and J. F. Gentleman. 2001. "On judging the Significance of Differences by Examining the Overlap between Confidence Intervals." *The American Statistician* 55 (3): 182–186.
- Simpson, G. G., and A. Roe. 1939. Quantitative Zoology. Revised edition. New York: McGraw-Hill.
- Tukey, J. 1991. "The Philosophy of Multiple Comparisons." Statistical Science 6 (1): 100-116.

Tyron, W. W. 2001. "Evaluating Statistical Difference, Equivalend and Indeterminacy Using Inferential Confidence Intervals: An Integrated Alternative Method of Conducting Null Hypothesis Statistical Tests. "*Psychological Methods* 6 (4): 371–386.

Cite this article: Armstrong II, D. A. and Poirier, W. (2025). Decoupling Visualization and Testing when Presenting Confidence Intervals. *Political Analysis*, 1–6. https://doi.org/10.1017/pan.2024.24