Sensitivity analysis is a useful tool for detecting influential observations

Samuel A. Mehr

Department of Psychology, Harvard University

Data Science Initiative, Harvard University

Simmons, Nelson, and Simonsohn (2011) argue convincingly that false positive results are unacceptably easy to publish given current reporting standards in psychological science. They conducted a real experiment with an impossible result: that listening to the song “When I’m Sixty-Four” can *decrease* the age of a listener. On this basis of this and other compelling evidence, the authors propose ten guidelines for authors and reviewers, including the reporting of analyses both with and without covariates, the reporting of all manipulations and variables, the use of sufficiently large samples, and so on. When these guidelines are applied to their experiment, it becomes evident that the initial finding was spurious.

The authors' conclusions are correct but they omit a key analysis: a search for atypical observations along with a systematic analysis of those observations' influence on the result. At present, psychologists have no standard method for the detection of influential observations (i.e., outliers). For instance, Cumming's (2013) “new statistics” make no mention of appropriate treatment of influential observations, and Simmons et al. (2011) found alarmingly haphazard decisions concerning outlier deletion in response time data reported in *Psychological Science*: “what constituted ‘too slow’ varied enormously: the slowest 2.5% or 10%, or 2 or 2.5 or 3 standard deviations slower than the mean, or 1.5 standard deviations slower from that condition's mean, or slower than 1,000 or 1,200 or 1,500 or 2,000 or 3,000 or 5,000 ms” (p. 1360).

There exists a systematic process for influential observation detection, already widely used in education research and econometrics: sensitivity analysis (reviews: Berk, 2004; Chatterjee & Hadi, 2009; Cohen, Cohen, West, & Aiken, 2003). To demonstrate this method's utility, I use it here to show that Simmons et al's (2011) false positive result was attributable not only to the questionable research practices they described, but also to the presence of a highly influential observation in their data.

U. Simonsohn (personal communication, October 24, 2013) shared the data to enable this analysis. First, I attempted to replicate the main analysis of Study 2 in Simmons et al (2011). The replication was successful: a multiple linear regression (i.e., ANCOVA) comparing age between subjects randomly assigned to listen to “When I’m Sixty-Four” to a control group revealed a statistically significant age difference between the two groups, adjusting for father’s age ($F(1,17) = 4.92, p = .04$; Wald test), identical to the original finding.

I then conducted a sensitivity analysis¹. First, I produced scatterplots for both of the independent variables, which I inspected visually to find any atypical observations (e.g., Anscombe, 1973). One such observation was evident (see Figure 1). To determine the degree of influence of this observation — i.e., whether or not one should be concerned that the main effect was substantially influenced by its presence — I obtained Cook’s D statistics (Cook, 1977) for each observation in the dataset. An observation’s D provides an estimate of the degree of change in model fit, were that single observation deleted from the dataset. For the observation in question, D was nearly an order of magnitude higher than the next-most influential observation, indicating that the observation exerted an extremely high degree of influence on the result.

¹For brevity, here I only show a simple sensitivity analysis, leaving aside topics of influence statistics other than Cook’s D and influence of groups of atypical points, for instance, which are discussed at length in the reviews cited above.

Figure 1. A sensitivity analysis comparing the unaltered, Winsorized, and trimmed models.

The sensitivity analysis proceeded as follows. I first removed the observation, to determine whether the main effect held without its influence; it did not ($F(1,16) = 0.96, p = .34$; Wald test). Removing an influential point presents two problems, however, in small samples: (a) it reduces statistical power, and (b) some portion of the influence of a particular point may have a valid effect on the dependent variable, such that removing *all* of its influence could *reduce* precision (Dixon, 1960). A useful alternative to trimming is Winsorization (Yale & Forsythe, 1976), in which atypical observations are replaced with the next-highest (or next-lowest) observation in the dataset, maintaining rank but reducing influence: I continued the sensitivity analysis with Winsorization of the influential observation; here again, the main effect did not

hold ($F(1,17) = 1.64, p = .22$; Wald test). To present the complete sensitivity analysis, I plotted the three models together (i.e., unaltered, Winsorized, or trimmed; Figure 1). The stark difference in effect sizes across the models, demonstrates the high degree of influence of the single atypical observation.

Thus, while Simmons et al.'s impossible result may indeed be attributable to lax methodology and reporting standards, it might also have been avoided by using sensitivity analysis. This method could be introduced as a standard post-hoc analysis in psychological science, to help resolve the ambiguity present in current methods for the treatment of outliers. Crucially, sensitivity analysis can be applied to *all linear models* (the vast majority of analyses used in psychology), is straightforward to conduct, and is supported by all common statistical analysis packages.

Using sensitivity analysis is consistent with Simmons et al.'s (2011) requirement that improved reporting imposes “a minimal cost to authors, readers, and reviewers” (p. 1365). In cases where a dataset has no detectable influential observations, reporting of a sensitivity analysis adds integrity to the analyses without much increase in word count: in six words an author might state, “Sensitivity analyses revealed no influential observations.” In cases where influential observations *are* detected, authors can avoid making biased decisions on outlier exclusions by reporting the main analysis in three ways (Figure 1): an unaltered model, a Winsorized model, and a trimmed model. This strategy would allow readers and reviewers to make better-informed judgments as to the statistical and practical significance of results.

Acknowledgements

Thanks to Uri Simonsohn for sharing the data and to Alia Martin for feedback on the manuscript.

References

- Anscombe, F. J. (1973). Graphs in statistical analysis. *The American Statistician*, 27(1), 17–21.
doi:10.1080/00031305.1973.10478966
- Berk, R. A. (2004). *Regression analysis: a constructive critique*. Thousand Oaks, Calif.: Sage Publications.
- Cohen, J. (2003). *Applied multiple regression/correlation analysis for the behavioral sciences*. Mahwah, N.J.; London: Lawrence Erlbaum Associates.
- Cook, R. D. (1977). Detection of influential observation in linear regression. *Technometrics*, 19(1), 15. doi:10.2307/1268249
- Cumming, G. (2013). The new statistics: Why and how. *Psychological Science*, 0956797613504966. doi:10.1177/0956797613504966
- Dixon, W. J. (1960). Simplified estimation from censored normal samples. *The Annals of Mathematical Statistics*, 31(2), 385–391. doi:10.1214/aoms/1177705900
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22(11), 1359–1366. doi:10.1177/0956797611417632
- Yale, C., & Forsythe, A. B. (1976). Winsorized regression. *Technometrics*, 18(3), 291–300.
doi:10.2307/1268738