Commentary on „Heterogeneity in direct replications in psychology and its association with effect size” (Olsson-Collentine, Wicherts & van Assen, 2020, Psychological Bulletin, online first publication, July 23, 2020.)

¹Erich H. Witte & ²Frank Zenker

¹Institute for Psychology, University of Hamburg, Von-Melle-Park 5, 20146 Hamburg, Germany; ORCID: 0000-0002-9446-3211;

²Dpt. of Philosophy, Boğaziçi University, 34342 Istanbul, Turkey; ORCID: 0000-0001-7173-7964; * corresponding author: frank.zenker@boun.edu.tr

Abstract: We argue that the observed correlation between the effect size and the heterogeneity of replicated studies is artificial. Specifically, that the observed heterogeneity for larger effects exceeds the observed heterogeneity for zero or small effects fails to provide a reason to invest comparatively less trust in parameter-estimates for large(r) effects.

Keywords: effect size; heterogeneity; parameter-estimation

1. Status quo

The public significance-statement of Olsson-Collentine, Wicherts & van Assen’s (2020, 1) states that “[d]ifferences in [empirical experimental] findings are unlikely if the overall effect is (close to) 0, whereas these differences are more likely if the overall effect is large.” Based on studying the heterogeneity of independently replicated effects, Schauer and Hedges (2020, 15) make the related observation that “[...] the experiments for which heterogeneity was often small tended to have effect estimates near zero.”

Both author-teams thus report an observed correlation between a zero effect and a homogeneous effect, on one hand, as well as between a non-zero effect and a non-homogenous (i.e., a heterogeneous) effect. If this correlation held generally in human behavioral experimental studies, then original effects would be *robust* only if the replication effects are either zero or small, whereas if the effect size increases, then such effects would be comparatively more heterogeneous. How one might interpret this observation? Is there anything newsworthy?

2. Prima facie considerations

For the two extremal cases, the focal correlation suggests that one can obtain a *stable* (homogeneous) replication effect, only if the data from the replication studies show a zero or

a small effect; whereas all larger effects tend to be more or less heterogeneous. This would suggest that larger effects generally provide comparatively less trustworthy parameter-estimates than zero or small effects. If so, then one's a priori expectation should be that future replication-studies can successfully replicate a zero or small effect; whereas a replication-studies that exceeds a small effect provides a less trustworthy parameter-estimate, because such effects display large(r) heterogeneity.

Before deciding to therefore altogether *ignore* the effect size-“chase” that appears to rule most of our present-day literature, one should probably ask whether the focal correlation is best conceived of as a characteristic of the effect size? The empirical fact is that most replications found a zero or a small effect and these replications are robust (homogeneous). Where replications found a larger effect, by contrast, these replicated effects were heterogeneous, and seemingly non-trustworthy. We argue that the correlation is artificial, and thus without further meaning for the effect size's stability, because the observed heterogeneity presents a statistical artefact.

3. An artificial correlation

Cohen's d -measure $d=(m_1-m_2)/s$ is widely used to quantify an effect's size. Here ignoring the variability of s in the denominator, it counts that the variance of the nominator is clearly discernible as the variance of a mean-difference. Generally, if an effect is small, then so is the numerator; otherwise it is null. In case the observed effect is null ($d=0$), the variance associated to (m_1-m_2) collapses into the variance of a single distribution, because one statistically treats only the distribution centered on zero. One can therefore trace the heterogeneity of a zero-effect *only* to this distribution's variance.¹

As an effect increases beyond zero, the two distributions around the means m_1 and m_2 come to overlap ever less. Test-results obtained from these distributions thus become gradually more independent (see the U-measures in Cohen, 1977, 22; Table 2.2.1). As independence increases, so does the effect size's heterogeneity, because the mutual separation of all individual values in both distribution increases. Put differently, increasing the separation of normally distributed values formally entails that the distribution of the differences displays greater variability. Thus, the effect's observed heterogeneity increases.

¹ The main argument is the same if, instead of mean differences, one uses the differences of Fisher-z-transformed correlations.

This is precisely what the data bear out (Olsson-Collentine et al., 2020, 14; Fig. 3). Irrespective of the different measures of heterogeneity, the product-moment correlation here is always positive and large. Specifically, the correlation between the absolute effect size and the standard deviation across reported studies is $r=0.77$. The observed qualitative regularity between the observed effect size and observed heterogeneity can therefore be explained by a measurement regularity, which is in turn reflected by a statistical regularity.

This regularity, however, here leads to a correlation between the effect size and heterogeneity that is *artificial*. After all, it is well-known from the classical analysis of variance (ANOVA)² (Winer, 1962) that decreasing the overlap between distributions increases the variance of sample results for different groups:

$$\begin{aligned} s^2_{\text{total}} &= s^2_{\text{intern}} + s^2_{\text{m}} & (1) \\ s^2_{\text{total}} &= \sum [(X_{ij} - M)^2 / (N-1)] \\ s^2_{\text{m}} &= \sum [(m_p - M)^2 / (p)] \end{aligned}$$

s^2_{total} : joint variance of all measurements in control and experimental group

s^2_{intern} : variance around the means of both groups

s^2_{m} : variance of the means

m_p : mean of the groups

X_{ij} : single measurements

M : grand mean

p : total number of means

To see why the variance increases, it should suffice to read the ANOVA-design backwards. Given a classification into groups with larger mean differences, the total variance (s^2_{total}) increases because the mean differences increase, rather than because the internal variance (s^2_{intern}) increases. This leads to the F-test which compares the variance between means to the internal variance.

² Using the ANOVA as a basic model, the underlying design is best understood as a between-subjects design of independent groups. If the design is a within-subjects design, by contrast, then the total variance of the differences must be corrected: $s^2_{\text{diff}} = s^2_1 + s^2_0 - 2 r_{10} \cdot s^2_1 \cdot s^2_0$.

4. Non-artificial heterogeneity

As effects increase in size, the total variance thus automatically increases, too. But this occurs for artificial reasons, rather than for substantial reason that might motivate finding parameter-estimates for large(r) effects comparatively less trustworthy. Indeed, that the variance increases is a purely quantitative fact that is related to the effect's size. By contrast, one may interpret the observed heterogeneity as a qualitative characteristic of an effect that is less stable than a zero or a small effect. In this case, the variance around the means, s^2_m , will increase independently of how the means themselves vary.

The solution, then, is to focus on the *internal variance* (s^2_{intern}) around the means of both groups. This variable should be correlated with the absolute effect size, i.e., the difference d between both means:

$$\begin{aligned} s^2_{\text{total}} - s^2_m &= s^2_{\text{intern}} \\ r_{(s^2_{\text{intern}};d)} &= 0 \text{ (expected)} \end{aligned} \tag{2}$$

Hence, (2) presents a way of eliminating the artificial correlation between the observed heterogeneity and the observed effect size. Where (2) is used in practice, however, and if a non-zero correlation between the effect size (measured as d) and heterogeneity (measured as s^2_{intern}) is observed, then a comparatively increased effect no longer explains this correlation. In this case, a reasonable interpretation might be that random influences *generally* disturb large(r) effects more than they disturb zero or small effects, and therefore provide less trustworthy parameter-estimates.

References

- Cohen, J. (1977). *Statistical power analysis for the behavioral sciences*. (rev. ed.). New York. NY: Academic Press.
- Olsson-Collentine, A., Wicherts, J.M. & van Assen, M.A.L.M. (2020, July 23). Heterogeneity in direct replications in psychology and its association with effect size. *Psychological Bulletin*. Advance online publication. <http://dx.doi.org/10.1037/bul0000294>.
- Schauer, J.M. & Hedges, L.V. (2020, April, 9). Assessing heterogeneity and power in replications of psychological experiments. *Psychological Bulletin*. Advance online publication. <http://dx.doi.org/10.1037/bul0000232>.
- Winer, B.J. (1962). *Statistical Principles in experimental design*. New York. NY: McGraw-Hill.