Commentary on BBS target article: Zwann et al. *Making Replication Mainstream*

A Bayesian decision-making framework for replication

Tom E. Hardwicke, Michael Henry Tessler, Benjamin N. Peloquin, Michael C. Frank
Stanford University

Correspondence addresses:

Tom E. Hardwicke, Meta-Research Innovation Center at Stanford (METRICS), Stanford School of Medicine, Stanford University, CA, USA, 94305. tom.hardwicke@stanford.edu

Michael Henry Tessler, Benjamin N. Peloquin, & Michael C. Frank, Department of Psychology, Stanford University, CA, USA, 94305. mhtessler@stanford.edu, bpeloqui@stanford.edu, mcfrank@stanford.edu

Abstract

Replication is the cornerstone of science – but when and why? Not all studies need replication, especially when resources are limited. We propose that a decision-making framework based on Bayesian philosophy of science provides a basis for choosing which studies to replicate.

Main text

Direct replications have an important place in our scientific toolkit. Given limited resources, however, scientists must decide when to replicate vs. conduct novel research. A Bayesian viewpoint can help clarify this issue. On the Bayesian view, scientific knowledge is represented by a probability distribution over theoretical hypotheses, given the available evidence (Strevens, 2006). This distribution – called the *posterior* – can be decomposed into the product of the prior probability of each hypothesis and the likelihood of the data given the hypothesis. Evidence from a new study can then be integrated into the posterior, making hypotheses more or less probable. The amount of change to the posterior can be quantified as a study's *information gain*. Using this formalism, one can design “optimal experiments” that maximize information gain relative to available resources (e.g., McKay, 1992). One attraction of this quantitative framework is that it captures the dictum that researchers should design studies that can adjudicate between competing hypotheses (Platt, 1964).

Some good intuitions fall out of the Bayesian formulation. A study designed to select between *a priori* likely hypotheses (e.g., those well-supported by existing data) can lead to high information gain. By contrast, a study whose data provide strong support for a hypothesis that already has a high prior probability, or weak support for a hypothesis that has a low prior probability, provides much less information gain. Larger samples and more precise measurements will result in greater information gain, but only in the context of a design that can distinguish high prior probability hypotheses.

Even well-designed studies can be undermined by errors in execution, reporting, or analysis. If a study is known to be erroneous, then it clearly leads to no information gain, but the more common situation is some uncertainty about the possibility of error. A Bayesian framework can capture this uncertainty by weighting information gain by a study's *credibility*. Concerns about questionable research practices, analytic error, or fraud thus all decrease the overall information gain from a study.

Direct replications are special in this framework only in that they follow a pre-existing set of design decisions. Thus, the main reason to replicate is simply to gather more data in a

promising paradigm. In cases where the original study had low credibility or a small sample size, a replication can lead to substantial information gain (Klein et al., 2014). Replicating the same design again and again will offer diminishing returns, however, as estimates of relevant quantities become more precise (Mullen, Muellerleile, & Bryant, 2001). If a study design is not potentially informative, for example because it cannot in principle differentiate between hypotheses, then replicating that design will not lead to information gain. Finally, when a particular finding has substantial *applied* value, replicators might want to consider an *expected value analysis* wherein a replication's information gain is weighted by the expected utility of a particular outcome.

Replications have one unique feature, though: They can change our interpretation of an original study by affecting our estimates of the original study's credibility. Imagine a very large effect is observed in a small study and an identical but larger replication study then observes a much smaller effect. If both studies are assumed to be completely credible, the best estimate of the quantity of interest is the variance-weighted average of the two (Borenstein, Hedges, Higgins, & Rothstein, 2009). But if the replication has high credibility – e.g., because of preregistration, open data, etc. – then the mismatch between the two may result from the earlier study lacking credibility due to error, analytic flexibility, or another cause. Such explanations would be taken into account by downweighting the information gain of the original study by that study's potential lack of credibility. Of course, substantial scientific judgment is required when sample, stimulus, or procedural details differ between replication and original (cf. Gilbert, King, Pettigrew, & Wilson, 2016; Anderson et al., 2016). Often, multiple studies that investigate reasons for the failure of a replication are needed in order to understand disparities in results (see e.g., Baribault et al., in press; Lewis & Frank, 2016; Phillips, Ong, et al., 2015).

The prior probability of hypotheses will not be universally agreed upon and can lead to disagreements about whether a particular result should be replicated. One researcher may believe that a study with low information gain – perhaps due to a small sample size – deserves to be “rescued” by replication because it addresses a plausible hypothesis. By contrast, a more skeptical researcher who assigned the original hypothesis a lower prior probability might see no reason to replicate. Or that researcher might replicate simply to convince others that the original study lacks credibility, especially in the case that it is influential within academia or the general public. Overall, as long as studies are appropriately conducted and reported, and all studies are considered, then the Bayesian framework will accumulate evidence and converge to an estimate of the true posterior.

Replication is an expensive option for assessing credibility, however. Assessing analytic reproducibility and robustness may be a more efficient means of ensuring that errors or specific

analytic choices are not responsible for a particular result (Stodden et al., 2016; Steegen, Tuerlinckx, Gelman, & Vanpaemel, 2016). Forensic tools like p-curve or the test for excess significance (Simonsohn, Nelson, & Simmons, 2014; Ioannidis & Trikalinos, 2007) can also help in assessing credibility.

How should an individual researcher make use of this Bayesian framework? When thinking about replication, researchers should ask the same questions they do when planning a new study: Does my planned study differentiate between plausible theoretical hypotheses, and do I have sufficient resources to carry it out? For a replication, this judgment can then be qualified by whether a re-evaluation of the credibility of the original study would be a net positive, since downweighting the credibility of an incorrect or spurious study also leads to overall information gain. Adopting such a framework to guide a rough assessment of information value (even in the absence of precise numerical assignments) can help researchers decide when to replicate.

References

- Anderson, C. J., Bahník, S., Barnett-Cowan, M., Bosco, F. A., Chandler, J., Chartier, C. R...Zuni, K. (2016). Response to comment on "Estimating the reproducibility of psychological science." *Science*, 351, 1037–1037.
- Baribault, B., Donkin, C., Little, D. R., Trueblood, J., Oravecz, Z., van Ravenzwaaij, D., White, C. N., de Boeck, P., & Vandekerckhove, J. (in press). Metastudies for robust tests of theory. *Proceedings of the National Academy of Sciences*.
- Borenstein, M., Hedges, L. V., Higgins, J. P. T., & Rothstein, H. R. (2009). *Introduction to meta-analysis*. Chichester, England: John Wiley & Sons.
- Gilbert, D. T., King, G., Pettigrew, S., & Wilson, T. D. (2016). Comment on "Estimating the reproducibility of psychological science". *Science*, 351, 1037–1037.
- Ioannidis, J. P., & Trikalinos, T. A. (2007). An exploratory test for an excess of significant findings. *Clinical Trials*, 4, 245–253.
- Klein, R. A., Ratliff, K. A., Vianello, M., Adams, R. B., Jr, Bahník, Š., Bernstein, M. J...Nosek, B. A. (2014). Theory building through replication: Response to commentaries on the "Many Labs" replication project. *Social Psychology*, 45(4), 299–311.

Lewis, M. L., & Frank, M. C. (2016). Understanding the effect of social context on learning: A replication of Xu and Tenenbaum (2007b). *Journal of Experimental Psychology: General*, 145, e72– e80.

MacKay, D. J. (1992). Information-based objective functions for active data selection. *Neural Computation*, 4(4), 590-604.

Mullen, B., Muellerleile, P., & Bryant, B. (2001). Cumulative meta-analysis: a consideration of indicators of sufficiency and stability. *Personality and Social Psychology Bulletin*, 27(11), 1450-1462.

Platt, J. R. (1964). Strong inference. *Science*, 146(3642), 347-353.

Phillips, J., Ong, D. C., Surtees, A. D. R., Xin, Y., Williams, S., Saxe, R., & Frank, M. C. (2015). A second look at automatic theory of mind. *Psychological Science*, 26(9), 1353–1367.

Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). P-curve: A key to the file-drawer. *Journal of Experimental Psychology: General*, 143(2), 534–547.

Steegen, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing transparency through a multiverse analysis. *Perspectives on Psychological Science*, 11(5), 702–712.

Stodden, V., McNutt, M., Bailey, D. H., Deelman, E., Gil, Y., Hanson, B...Taufer, M.(2016). Enhancing reproducibility for computational methods. *Science*, 354(6317), 1240–1241.

Strevens, M. (2006). The bayesian approach to the philosophy of science. D. M. Borchert (Ed.) *Encyclopedia of Philosophy*. Detroit, MI: Macmillan Reference USA.