The relevance and the pitfalls of replications: a commentary on “The convergent validity and nomological net of two methods to measure retroactive influences” by Matt C. Howard, 2018

Patrizio E. Tressoldi¹

¹Science of Consciousness Research Group

Dipartimento di Psicologia Generale, Università di Padova, Italy

Email: patrizio.tressoldi@unipd.it

Abstract

This commentary aims to demonstrate that Howard's (2018) studies are neither a "close" nor a "conceptual" replications of Bem's test of retroactive influence (RI) as far as it concerns task 1 and that all tasks in the three studies can be classified as "very far conceptual replication" of a general class of phenomena classified as forced-choice precognition (conscious cognitive awareness) of a future event that could not otherwise be anticipated through any known inferential or perceptual process.

As to the results related to the convergent validity (correlation) between the tasks used to test the RI, we posit that without a measure of intra-task reliability any measure of inter-task reliability is uninterpretable.

Keywords: replication; retroactive influence; reliability; precognition.

Benefits of replication

Replication of previous findings is a cornerstone of the scientific process. No one would dispute this statement, particularly these days where different disciplines, from psychology to medicine, and from experimental economy to ecology, along with many others are struggling to face up to and deal with issues of replicability (e.g. Munafò et al., 2017; Pashler & Wagenmakers, 2012). Hence, the first impressions of Howard's (2018) paper is that it may provide a welcomed contribution to the debate. However, a close (exact) replication is not a trivial endeavour and much has been written and continues to be published about how to correctly replicate previous findings (e.g. LeBel, McCarthy, Earp, Elson, & Vanpaemel, 2018; Zwaan, Etz, Lucas, & Donnellan, 2018).

With this commentary, we will demonstrate that Howard's (2018) studies are neither a "close" nor a "conceptual"¹ replications of Bem's (2011) test of retroactive influence (RI). In fact, according to the criteria of LeBel et al., (2018), all tasks in the three studies by Howard (2018), can be classified as "very far conceptual replications". Research focusing on such retroactive effects has, since the seminal work of Bem (2011), produced some robust evidence of a real effect though small in size (Bem, Tressoldi, Rabeyron, & Duggan, 2016) though not all find such effects (Galak et al., 2012; Ritchie et al., 2012) which in turn has led to calls for more research to be conducted in an effort to help clarify the process (Franklin, Baumgart, Schooler, & Broadway, 2014).

Bem's versus Howard's task differences

The focus of the replication was Bem's retroactive facilitation of recall task (experiments 8 and 9). This task involved giving participants a set of words to remember which was immediately followed by a free recall test. Following the test, participants were then given a randomly selected sub-set of

¹ "Close (direct) direct replication" is a study that attempts to recreate the critical elements (e.g., samples, procedures, and measures) of an original study where those elements are understood according to "a theoretical commitment based on the current understanding of the phenomenon under study, reflecting current beliefs about what is needed to produce a finding. A "conceptual replication" is a study where there are changes to the original procedures that might make a difference with regard to the observed effect size. Conceptual replications span a range from having one theoretically meaningful change with regard to the original experiment (e.g., a different dependent measure) to having multiple changes. From Zwaan et al., (2018), pag. 3.

words to practice with. The prediction was that the post-test practice exercises would retroactively facilitate the recall of those words, and, hence, participants would recall more of the *to-be-practiced* words than those not practiced. For a good example of a close replication of Bem's experiments 8 and 9 see Galak, LeBoeuf, Nelson, & Simmons (2012).

Table 1: Differences between Bem's and Howard's Study 1 and 2 experimental designs.

Bem (2011)	Howard (2018) Study 1 & 2 task 1	Howard (2018) Study 1 & 2 task 2
1- Participants and experimenter present in the lab	Mechanical Turk website	Mechanical Turk website
2- Instructions: This experiment tests for ESP by administering several tasks involving common everyday words. At the end of the session, I will explain to you how this procedure tests for ESP	Instructions: Tests your ability to predict the future	Instructions: Tests your ability to predict the future
3- Three-minutes relaxation before the task	No relaxation	No relaxation
4- Presentation of 48 nouns one at a time for 3 sec each divided in four categories.	Presentation of a list of 20 words;	Presentation of 12 series of 4 pictures.
5- Unexpected free recall of all remembered words	Select the five words that will appear on the next page;	Select one of four photos that will appear on the page following the question.
6- Four repeated reading of 24 randomly selected words, with the request to identify and retype those belonging to a specific category	Type once the five randomly selected words	

Characteristics 1 to 3, were never explicitly considered by Bem's as core aspects for testing RI, even if the difference in the Howard's task may influence the quality of participants' performance given the lack of quality checks. In fact, in the (Galak et al., 2012) close replication study, it emerged a clear difference between the effect sizes of the experiments executed online: effect size = -.02; 95% CI: -.09 - .05 with respect those executed in the lab, effect size .09; 95% CI: .03 - .14.

Also characteristic 4, number of items and their division in specific conceptual categories, was never declared at the core of the RI effect. However, differences 5 and 6, clearly show that both for study 1 and 2 of Howard's tasks are completely different and cannot be considered neither a close nor a

conceptual replication of that used by Bem, lacking the unexpected free recall of all presented words and, more importantly, the repeated practice of a subgroup of words, considered as the key component of the supposedly RI effect.

As to task 2 of both Study 1 and 2 and to task 1 and 2 of Study 3 (see Table 2), Howard correctly specified that they referred to a separate category of experiments which require the participants to guess the future targets choosing among different items by using forced-choice methods in a normal state of consciousness (see Storm, Tressoldi, & Di Risio, 2012,).

Table 2: Main characteristics of Howard’s Study 3

Howard (2018)
Study 3
Mechanical Turk website
Instructions: Tests your ability to predict the future
No relaxation
Presentation of a list of 20 words (Task 1);
Presentation of 12 series of 4 pictures (Task 2)
Select the 5 words that you like the best. (Task 1); Select the most liked picture (Task 2).
Presentation of the randomly selected picture

In conclusion, Howard’s claim that the results of his work draws “prior investigations into question” (p.10) is a somewhat surprising over interpretation of his null results falling in the “Absence of evidence is not evidence of absence” fallacy, in particular if we consider that a statistical power of at least .80 for detecting the expected small effect size of this category of phenomena which is approximately .04, will require a minimum of 3800 participants.²

Convergent validity

² Estimated with G*Power software v.3.1.9.2, for a one-tailed, one-sample t-test

As for testing the convergent validity of the two methods (task 1 and task 2 in the three Studies) in order to provide a more robust test of RI effects, this is a good idea and Howard should be commended for proposing this. However, he espouses a strong assumption, namely that participant's performance on the two tasks remains stable. Apart from the differences in the two tasks in all three studies (see Tables 1 and 2), the hypothesis that every, or most of the participants perform at the same level, in order to find a positive correlation between the two tasks is hardly tenable. First of all, it would have been necessary to measure the intra task reliability before checking the inter-task one. If the intra-task reliability is low, any inter-task reliability makes no sense. Furthermore, the shape of the data distribution should be checked in order to decide whether parametric or nonparametric statistics should be applied. Moreover, given that neither of the tasks were able to elicit/show evidence of RI, it should therefore not be surprising that they do not correlate with such variables.

Running cross correlations tells us nothing about the potential suitability of such tasks for measuring RI, but simply that performance in this instance was not related. Moreover, these correlations are very similar: .09: -.09 in the text (-.10 in Table 2) and .09 respectively for the three studies and this rare coincidence should accurately be checked.

Concluding comments

Independent close (exact/direct) replications of published studies must be valued important as much as the original studies (Zwaan et al., 2018), and editors must encourage them in particular when the phenomena under investigations are controversial or have important practical and societal impact. However, this type of replications requires specific methodological and statistical tools, starting from the preparation of the experimental design possibly with the collaboration or the agreement of the original authors, up to the statistical comparison of the outcomes, possibly preregistering the study in order to reduce the degrees of freedom inherent in every experimental study that enhance the risk of p-hacking and HARKing (Hypothesizing After the Results are Known; Nosek, Ebersole, DeHaven,

& Mellor, 2018) and making all data available open access. A very good example to be imitate of a close replication of another Bem's study is the Transparent PSI project (Kekecs et al., 2016).

We think to have demonstrate that Howard's paper doesn't satisfy these requirements and can be considered as only a statistical underpowered contribution to the study of a generic class of phenomena defined as conscious forced-choice precognition.

Acknowledgments: I thank you David Vernon and an anonymous reviewer for their comments and suggestions to a previous version of this paper.

References

- Bem, D. J. (2011). Feeling the future: Experimental evidence for anomalous retroactive influences on cognition and affect. *Journal of Personality and Social Psychology*, 100(3), 407–425.
<https://doi.org/10.1037/a0021524>
- Bem, D., Tressoldi, P., Rabeyron, T., & Duggan, M. (2016). Feeling the future: A meta-analysis of 90 experiments on the anomalous anticipation of random future events [version 2; referees: 2 approved]. <https://doi.org/10.12688/f1000research.7177.2>
- Franklin, M. S., Baumgart, S. L., Schooler, J. W., & Broadway, J. M. (2014). Future directions in precognition research: more research can bridge the gap between skeptics and proponents. *Frontiers in Psychology*, 5 :907. <https://doi.org/10.3389/fpsyg.2014.00907>
- Galak, J., LeBoeuf, R. A., Nelson, L. D., & Simmons, J. P. (2012). Correcting the past: Failures to replicate psi. *Journal of Personality and Social Psychology*, 103(6), 933–948.
<https://doi.org/10.1037/a0029709>
- Kekecs, Z., Aczel, B., Palfi, B., Barnabas, S., Zrubka, M., Szecsi, P., & Kovacs, M. (2016). Transparent Psi Project. OSF. <https://doi.org/none>
- LeBel, E. P., McCarthy, R. J., Earp, B. D., Elson, M., & Vanpaemel, W. (2018). A Unified Framework to Quantify the Credibility of Scientific Findings. *Advances in Methods and Practices in Psychological Science*, 1(3), 389–402. <https://doi.org/10.1177/2515245918787489>
- Munafò, M. R., Nosek, B. A., Bishop, D. V. M., Button, K. S., Chambers, C. D., Percie du Sert, N., ... Ioannidis, J. P. A. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, 1(1), 0021. <https://doi.org/10.1038/s41562-016-0021>
- Nosek, B. A., Ebersole, C. R., DeHaven, A. C., & Mellor, D. T. (2018). The preregistration revolution. *Proceedings of the National Academy of Sciences of the United States of America*,

115(11), 2600–2606. <https://doi.org/10.1073/pnas.1708274114>

Pashler, H., & Wagenmakers, E. (2012). Editors' Introduction to the Special Section on Replicability in Psychological Science. *Perspectives on Psychological Science*, 7(6), 528–530. <https://doi.org/10.1177/1745691612465253>

Storm, L., Tressoldi, P. E., & Di Risio, L. (2012). Meta-analysis of ESP studies, 1987–2010: Assessing the success of the forced choice design in parapsychology. *Journal of Parapsychology*.

Zwaan, R. A., Etz, A., Lucas, R. E., & Donnellan, M. B. (2018). Making replication mainstream. *Behavioral and Brain Sciences*, 41, e120. <https://doi.org/10.1017/S0140525X17001972>