Mischaracterizing replication studies leads to erroneous conclusions

Brian Nosek & Elizabeth Gilbert

Published in Fall 2015, [the Reproducibility Project: Psychology](#) reported the first systematic effort to generate an estimate of reproducibility by replicating 100 published studies from 2008 issues of three prominent journals in psychology. We are two of the co-authors of this 270 author project. Overall, the reproducibility rate was ~40% across 5 distinct criteria. In their [critique](#) published last week, Dan Gilbert (no relation), Gary King, Stephen Pettigrew, and Tim Wilson (2016) suggested that effective reproducibility rate was not distinguishable from 100% reproducibility because of flaws in the methodology.

Many co-authors of the Reproducibility Project: Psychology [published a response](#) countering the critique. Within days, independent commentaries have also emerged challenging Gilbert and colleagues' methodology and conclusion by [Sanjay Srivastava](#), [Uri Simonsohn](#), [Daniel Lakens](#), [Simine Vazire](#), [Andrew Gelman](#), [David Funder](#), [Rolf Zwaan](#), and [Dorothy Bishop](#).

We will address a simple point not covered by these commentaries. Gilbert and colleagues suggested that dramatic differences between the original and replication studies caused lower reproducibility. To support their point, they briefly described differences between six original and replication studies including this one:

"An original study that asked Israelis to imagine the consequences of military service was replicated by asking Americans to imagine the consequences of a honeymoon"

And, then elaborated in their [on-line rebuttal](#) to [our response](#):

"In one original study, researchers asked Israelis to imagine the consequences of taking a leave from mandatory military service (Shnabel & Nadler, 2008). The replication study asked Americans to imagine the consequences of a taking a leave to get married and go on a honeymoon. ... And not surprisingly, the replication study failed."

Based on Gilbert and colleagues' description of such a dramatic difference, one can't help but conclude that the replication team was incompetent. And, if one replication could be fouled up this badly, it would seem likely that there are more problems in the other 99. With this description, it is easy to believe that the Gilbert and colleagues' press release headline "[Researchers overturn landmark study on the replicability of psychological science](#)" is not just bloviating.

One of us (Elizabeth) led this replication attempt, and the other of us (Brian) is the only person on the planet that thinks that military service and honeymoons are the same thing. Together, we disagree with their characterization of this study. It is certainly possible that we are just feeling defensive about our work. We will describe this original study and replication. You can assess whether Gilbert and colleagues describe the study accurately, and came to a reasonable conclusion that the replication was "dramatically different" from the original.

The Original Study

The original study was about what psychological resources increase the chance that perpetrators and victims will reconcile. Across four vignette studies, the original authors found evidence that restoring victims' sense of status and power and restoring perpetrators' sense of public moral image increased the likelihood that they were willing to reconcile. These results supported their hypothesis that victims' and perpetrators' have different psychological needs. This is an important topic and a creative study, and our single failure to replicate it does not mean that the original finding is wrong.

How did the original authors test this question? All the original and replication materials [are available on the Open Science Framework](#). Here is a summary of the key parts. In the original study, 75 female and 19 male Israeli students (average age 23.5) came to the laboratory, read a short story, and answered questions about the story including what they would do in one character's position. The story was about a coworker (the perpetrator) who took credit for another coworker's work (the victim). The two coworkers had worked together for several years and recently had been collaborating on a joint project, but shortly before the project was due, the victim had to leave work for a while. The perpetrator then turned in the joint work, receiving all the credit for it and a promotion, whereas the victim got demoted when s/he returned. The key part of the experiment was whether the participant imagined being the victim or the perpetrator in the story.

Here is the original and replication "victim" stories. The "perpetrator" story is identical except for switching Assaf/Anat and the participant's role ("you") in the story. The materials were translated by the original author from Hebrew to English. Male participants read about Assaf and reserve duty, female participants read about Anat and maternity leave. Highlighted in bold is the feature of that Gilbert and colleagues used to explain what the study was about:

Original Study "Victim" condition

You and Assaf/Anat are working together for several years in a successful advertising company. Recently you have been collaborating on a joint project that included a campaign to a thriving fashion brand. Just prior to the completion of the project you had to **leave for [reserve duty]/[maternity leave]** earlier than expected and you asked Assaf/Anat to take care of two tasks that you haven't managed to complete. Assaf/Anat has completed the tasks with distinction, and thus, one month after you left for the **[reserve duty]/[maternity leave]** the boss offered him/her to take over your role and responsibilities and demote you to an underling role in a different department. Although Assaf/Anat knows you might be hurt by it, s/he accepts the promotion including everything that goes with it. When you learn about it you feel betrayed, hurt, and very angry at him/her because s/he agreed to accept your role. You know that now, despite your proven successes, you will be far less confident in job interviews.

After reading this story or the "perpetrator" version, all participants answered survey questions such as how much "I was hurt by Assaf/Anat" and whether "Assaf/Anat perceived me as

completely moral.” Then, participants read another short vignette that Assaf/Anat later provided public feedback that either praised the participants’ professional skills (the condition intended to increase sense of power) or their interpersonal skills (the condition intended to increase public moral image). Participants then reported their willingness to reconcile with Assaf/Anat.

The Replication Study

We conducted the replication study with 144 students at the University of Virginia (82 men, 62 women). Based on existing theory and evidence, we had no reason to expect cultural differences in whether restoring status and power to victims and restoring moral acceptance to perpetrators increases chances of reconciliation. However, the scenario text was not perfect for our sample. In the U.S., reserve duty isn’t common for men, and demotions for maternity leave are illegal. So, we needed to alter the scenarios to have a reason for being away from work that was relatable. We chose one that would work for both men and women, a honeymoon. We retained gender matching between the participant and the co-worker, but also altered the names to be more relatable to Americans. We made a couple of additional phrasing edits, here is the full text, again just the “victim” condition with the same revisions made to the “perpetrator” condition:

Replication “Victim” condition

You have graduated from college and have been working at a successful advertising company for several years. During the last year, you and your colleague, Amy/Andy, have been collaborating on a joint project that included a campaign for a thriving fashion brand. You also are **recently engaged to be married**, and unfortunately your **wedding and honeymoon** are scheduled to occur right before the completion of the project. You need to take a leave of absence for this, so you ask Amy/Andy to take care of two tasks that you haven't managed to complete. Amy/Andy has completed the tasks with distinction. Before you return two weeks later, your boss offered her/him to take over your role and responsibilities and demote you to an underling role in a different department. Although Amy/Andy knows you might be hurt, she/he accepts the promotion, along with all benefits that come with it. When you return you feel betrayed, hurt, and very angry at her/him because she/he agreed to accept your position. You know that now, despite your proven successes, you will be far less confident in job interviews.

The follow-up survey, vignette about public praising, and final questions about possible reconciliation were the same as the original study. We shared this revised design with lead author of the original research for feedback. She asked us to pilot test the materials to make sure our U.S. participants responded to it the same way her Israeli participants responded to the original. That is, did U.S. participants find it easy to imagine themselves in the situation, how angry and hurt would they would feel in this situation, etc.? We pilot tested the revised scenario and confirmed that it met those criteria. The original author approved this revision to the design for conducting the replication. We [pre-registered the protocol](#) noting the changes, conducted the replication, and wrote the [final report](#).

Gilbert and colleagues’ (mis-)characterization

Gilbert and colleagues asserted that this study illustrated “dramatic differences” between the original and the replication. They had described the study as being about imagining “the consequences of military service” versus “the consequences of a honeymoon.” In our opinion, this is quite misleading. First, the study was about how victims and perpetrators respond when someone else takes credit for their work. Second, the reason for being away from the office (reserve duty, maternity leave, or honeymoon) was an incidental feature of the scenario to allow someone to take credit for another person’s work. Third, that incidental feature was held constant across the experimental conditions (victim or perpetrator). Fourth, 80% of the original participants were women, so their reason for being away was maternity leave, not military service. Fifth, we conducted manipulation checks to make sure that the scenario had similar effects as the originals before conducting the study. And, sixth, the lead original author was very generous with time and advice, and ultimately approved the changes - something that Gilbert and colleagues argued was essential for reproducibility.

The Conclusion

We believe that Gilbert, King, Pettigrew, and Wilson’s characterization of this replication and others is unfair. And, as we pointed out in [our published response](#), among the other five studies that they call out as flawed, some were endorsed by the original authors, and another replicated successfully. So who is right about the “fidelity” of these replications? You don’t need to take our word for it, or theirs. You can review [all the protocols, materials, data, and scripts yourself](#) on the Open Science Framework.

We suggest that this was a fair replication. Simultaneously, our assertion does not mean that the replication overturned the original result. It is perfectly reasonable to wonder if the effect is, for example, culturally constrained to Israelis or other more collectivistic societies. It is even reasonable to wonder if it is dependent on the reason the victim is out of the office. The fact that new interpretations are generated after the fact is not a threat to science; it is the engine of discovery.

When a replication produces a different result than an original study it may challenge our existing understanding of the phenomenon. As we said in our original article:

It is also too easy to conclude that a failure to replicate a result means that the original evidence was a false positive. Replications can fail if the replication methodology differs from the original in ways that interfere with observing the effect. We conducted replications designed to minimize a priori reasons to expect a different result by using original materials, engaging original authors for review of the designs, and conducting internal reviews. Nonetheless, unanticipated factors in the sample, setting, or procedure could still have altered the observed effect magnitudes.

We can use the different result to update our beliefs or try to preserve our pre-existing beliefs by identifying methodological differences that could explain the difference. As long as we recognize that such exploration is a means of *generating* hypotheses, then we don’t risk overconfidence and premature conclusions.

The Reproducibility Project: Psychology offered descriptive evidence of challenges for reproducibility. We and our 268 co-authors recognized that our evidence was not sufficient to draw strong conclusions about why that reproducibility rate was observed. To emphasize those limitations [our original article](#)'s conclusion opened thusly:

After this intensive effort to reproduce a sample of published psychological findings, how many of the effects have we established are true? Zero. And how many of the effects have we established are false? Zero. Is this a limitation of the project design? No. It is the reality of doing science, even if it is not appreciated in daily practice. Humans desire certainty, and science infrequently provides it. As much as we might wish it to be otherwise, a single study almost never provides definitive resolution for or against an effect and its explanation. The original studies examined here offered tentative evidence; the replications we conducted offered additional, confirmatory evidence. In some cases, the replications increase confidence in the reliability of the original results; in other cases, the replications suggest that more investigation is needed to establish the validity of the original findings. Scientific progress is a cumulative process of uncertainty reduction that can only succeed if science itself remains the greatest skeptic of its explanatory claims.

Gilbert, King, Pettigrew, and Wilson's error was not that they offered an optimistic interpretation of reproducibility from the Reproducibility Project: Psychology data. Their error was mistaking their *hypothesis that was generated from the data* for a *conclusion that was demanded by the data*. Accurate calibration of claims with evidence is what makes science special, and credible, as we all continue stumbling toward understanding.