Opinions on the value of direct replication: A survey of 2,000 psychologists

Brett Buttlere

Jelte Wicherts

Tilburg University, Netherlands

Words: 8,338

Abstract:

A major debate has emerged in Psychology on the merits of direct replication, with much discussion but little empirical data. Here we conducted a survey among 2,056 psychologists to gather opinions on several important themes in the debate such as: whether direct replications are valuable, whether obtained results can be the outcome of context or expertise, the intentions and reputational interests of the researchers involved, the severity of methodological problems, and what should be done to help move the field forward.

Descriptively, the results of our survey suggest that most psychologists trust the seminal findings in the field and are sensitive to the contextual variables which may lead to a failed replication of such a finding; while also agreeing that more direct replications should be done. Also of interest was what opinions were related to valuing direct replications. The application of structural equation models, developed on a subsample then confirmed across subfields, suggests that the respondents' perceptions of the severity of the methodological problems positively predicts support for direct replication, while respondents' view on replicator's intentions (e.g., replication bullying) negatively predicts support for direct replication.

Keywords: Replication, Scientific communication, Survey, Meta-science

Opinions on the value of direct replication: A survey of 2,000 psychologists

Several fraud cases (Levitt committee, 2012; Simonsohn, 2013), papers highlighting Questionable Research Practices (QRPs; John, Loewenstein, & Prelec, 2012; Simmons, Nelson, & Simonsohn, 2011; Bakker & Wicherts, 2011), and several failures to replicate seminal findings (e.g., Doyen, Klein, Pichon, & Cleeremans, 2012; Harris, Coburn, Rohrer, & Pashler, 2013; Johnson, Cheung, & Donnellan, 2014; Klein et al., 2014) have led to a major debate in Psychology on the need for -and value of- (direct) replications. What has been called a “Crisis of confidence” has occurred and several special issues in major journals have been released on how to improve the field (Pashler & Wagenmakers, 2012; Nosek & Lakens, 2014; Ledgerwood, 2014). These debates were renewed when a recent attempt to directly replicate important findings in the field (Nosek & Lakens, 2014) led to at least one researcher feeling bullied and her work unfairly defamed after a failure to replicate her earlier findings (Schnall, 2014).

These most recent debates have taken place largely outside peer reviewed journals, with quicker exchanges on blogs and in other social media environments (e.g., Twitter, blogs, and the Facebook group of the International Social Cognition Network). In fact, most of what has led to the increased tensions has happened in these alternative communication mediums. For instance, some of what led the original allegation of bullying was a blog post (titled “Go big or go home,” from an author of the replication) suggesting that the effort to replicate Schnall, Benton, & Harvey (2008) was an “epic failure” (Donnellan, 2014). On the other side of things, those doing the replications have been called, (to quote the first line of one blog post about the affair) “‘Replication police.’ ‘P-squashers.’ ‘Hand-wringers.’ ‘Hostile replicators.’ And of course, who can ever forget, ‘shameless little bullies.’” (pg. 1, Funder, 2014). This last was phrased by a well-known social psychologist (Gilbert, 2014a) who indicated receiving numerous other emails

expressing worry about “a grim McCarthyite nightmare” (Pg. 1, Gilbert, 2014b) amidst what has been called an, “atmosphere of wary and suspicious disbelief” (Pg. 1, Coan, 2014) caused by (and resulting in) attempts to replicate findings published in the psychological literature.

Despite this large amount of discussion, both in journals and via social media, there is almost no empirical data on (how the community generally feels about) these topics. Our goal here, then, is to gather data on what active researchers feel about some of the topics that have been debated. Specifically, we wanted to better understand the community’s opinions about: the value of direct replication, the importance of context effects and sample differences in psychological results, the role that expertise plays in psychological experimentation, the potential positive or negative outcomes for those researchers whose work is being replicated, the intentions of those doing the replications, the extent of the methodological problems in the field, and what should be done to help move the field forward. We studied these views by developing a brief but comprehensive 30-item survey that was then administered to psychologists who had recently published in the psychological literature.

Survey: Methods and Questions

In late November 2013 the authors of this project also ran another survey of psychologists examining problems in the publication system of psychology (Buttliere & Wicherts, in preparation). The sample for that survey was a randomly selected set of 10,000 respondents drawn from a list, based on Web of Science, of 23,295 corresponding authors of articles published in the psychological literature between January and November 2013. For the current survey, the emails of those authors we did not contact in that earlier survey (13,395) were again randomly ordered and the first 10,000 were contacted.

The survey was distributed using the Qualtrics mailer and emails were sent out between July 16th and July 30th, 2014. A reminder email was sent approximately one week after the initial email and data collection ceased on August 20th 2014. Thus, the majority of the data were collected within two months of the special issue of *Social Psychology* and the original Schnall (2014) blog post. At this time the discussion had settled somewhat from its start but was still at the forefront of minds (for instance, a lengthy Slate article published July 31, 2014 was shared over 2,000 times on Facebook as of November 2014; Meyer & Chabris, 2014). Below we focus more closely on each one of these contested questions, first defining the questions and outlining the various positions in the debate, then presenting the abstracted statements that were presented to participants to *agree* or *disagree* with (on a 100 point scale with breaks every 25 points). Specifically, the first question we wanted to help answer was:

Are direct replications important?

One of hottest debated topics of the months before the survey was sent out was the value of direct replications. The appearance of a special issue of *Social Psychology* that consisted exclusively of pre-registered replication projects is evidence of the sizable portion of psychologists who believe replications are important (Nosek & Lakens, 2014; Klein et al., 2014). These projects are the work of researchers concerned about the integrity of the field and its important findings after several papers have suggested that research practices that artificially inflate effect sizes were commonplace (John et al., 2012; Simmons et al., 2012) along with several demonstrations of outright fraudulent papers (Levelt committee, 2012; Simonsohn, 2013). Additional work has suggested that only about 1% of all Psychology papers are direct replications (Makel, Plucker, & Hegarty, 2012) and many authors have discussed the need for more replications (Ioannidis, 2005; Asendorpf et al., 2013; Nosek & Lakens, 2014).

Data about a lack of replications (Makel, Plucker, & Hegarty, 2012), though, can be interpreted as either a need for more replications, or evidence that the psychological community, for whatever reason, does not highly value direct replication. Instead, conceptual and partial replication seems to be the norm. For instance, Neuliep and Crandall (1993) found that almost 75% of articles in *JPSP* were some form of replication, along with suggesting that relatively similar studies must be conducted for methods like meta-analysis to be applied. Similarly, Jones, Derby, and Schmidlin (2010) found that more than half of the papers in *Human Factors* had been replicated at least once, though they were generally not reported as such. Mitchell (2014) has more recently suggested that given the many reasons a replication might fail, a direct replication failing actually says very little about the reality of the effect.

In order to assess whether the psychological community values direct replication, we asked psychologists directly how much they *agreed* or *disagreed* (on a 100 point scale with breaks every 25 points) with the following statements: “Direct replication is the best method of determining which effects are real”, “Psychological effects reported in ‘flagship’ journals should be widely replicable across labs, stimuli, cultures, and time”, and, “In order to help move the field forward, we should: conduct more direct replications”.

Do context effects and cultural differences matter in experimentation?

Of particular importance in the recent discussions has been how differences in context or sample might lead to diverging results. In fact, both Bargh (Yong, 2012) and Schnall (2014) have pointed toward differences in context as potential reasons for their works failing to replicate. For instance, the original Bargh, Chen, & Burrows (1996) social priming studies were done in the United States, while the failed replication was done in Belgium (Doyen et al., 2012). Similarly, the Schnall, et al. (2008) experiments were run in the United Kingdom, while the failed

replication was run in the United States (and online rather than in lab; Johnson, et al., 2014). Several authors have pointed out the difficulties in making these cross-cultural and contextual comparisons about stereotypes of the elderly or the morality of an act, for instance (Cesario, 2014; Stroebe & Strack, 2014). Schnall (2014) has further pointed out substantial differences in the distributions between the original study and the failed replication, which suggest that the samples might actually be from different populations. Here we investigate how strongly psychologists think type of sample, context, and cultural differences can make in experimental results by asking them to indicate their agreement or disagreement with the following statements: “The same exact materials can mean different things to participants across samples, contexts, and time,” and “Failed direct replications are often due to unknown substantive moderators (e.g., sample/ experimental setting differences).”

Does expertise play an important role in psychological experimentation?

Another important topic in the discussion has been the role that expertise on the part of the researcher can have in the experimentation process. For instance, when Doyen et al., (2012) failed to replicate Bargh et al. (1996), Bargh (in a blog post which has since been removed; Yong, 2012) suggested that the replicators had not primed or maintained the effect as should be done (also related to context differences; Yong, 2012). Others have reiterated this point saying, “I also believe in a point that John Bargh made in his response to Cleeremans, that priming effects are subtle and that their design requires high-level skills” (p. 1, Kahneman, 2012). There are some psychological protocols that require a special knowledge (or expertise) to implement (e.g., fMRI, LIWC, IAT, Stroop), but it is an open question whether psychologists believe that a supposed lack of expertise on the part of the replicator can be a major reason for a replication to fail. In order to learn more about psychologists’ perceptions of the role expertise plays in

psychological science, we asked them to *agree* or *disagree* with the following statements: “Some effects can only be achieved by an expert in the field” and “Psychological effects sometimes require a craftsmanship to achieve that some replicators might lack”.

What are the potential effects of a replication for the original author?

Another important aspect are the outcomes of (successful or unsuccessful) replication for the original researcher. While most agree that a successful replication is only good for the original author, several commentators have pointed out that a failed replication can cast serious doubt not only on a particular researcher or research line (Schnall, 2014; Mitchell, 2014), but an entire field (Garfield, & Welljams-Dorof, 1990; Kahneman, 2012; Yong, 2012). For instance, when Doyen, Klein, Pichon, and Cleeremans (2012) reported failing to replicate Bargh et al.’s (1996) elderly priming study, they (unintentionally) called into question not only the original researcher’s integrity and abilities as researchers, but the integrity of the entire field (Yong, 2012; Kahneman, 2012; Pashler, & Wagenmakers, 2012; Chin, 2013). Similarly, when Shanks et al. (2013) failed to replicate Dijksterhuis & van Knippenberg (1998), one journalist at Nature went so far as to suggest that the results, “call into question the underpinnings of unconscious-thought theory” (Abbott, 2013). While there are few empirical data demonstrating that a failure to replicate leads to lowered outcomes, there are several anecdotal examples (e.g., Donnellan’s ‘epic failure’ comment; ‘Controversial ESP study fails yet again’, Radford, 2012) and authors have expressed concern about themselves or their graduate students suffering negatively from these experiences (Schnall, 2014). Psychologists’ perceptions about the repercussions of a replication for the original researcher’s reputation were investigated by asking psychologists to *agree* or *disagree* with the following statements: “A successful independent replication is likely to increase other researcher’s trust in the original researcher’s integrity”, “An independent

replication that fails is likely to lower the career outcomes of the original researchers”, “A failed replication may lead to accusations of fraud”, and finally, “Some opponents of the replication movement are only worried about the replicability of their own effects.”

What are the perceived intentions of those doing the replications?

Science has shown that in judging the value of an act, not only the outcomes, but also the intentions are important (Darley, Klossen, & Zanna, 1978). There are many potential reasons for engaging in an independent replication, for instance to replicate an effect before building on it in one’s own experimental work, to test an effect one doubts, or to intentionally harm another researcher (Bohannon, 2014). If, as has been suggested, researchers who doubt an effect are less likely to find it (as is suggested by the confirmation bias; Mitchell, 2014; Mynatt, Doherty, Tweney, 1977), it can raise serious questions about the perceived ‘independence’ of independent replication attempts. There have also been suggestions that some replications are undertaken because they are an easy way to increase one’s career outcomes. Specifically, undertaking a replication allows one the opportunity to work closely with prominent researchers who are pro-replication and even affords the possibility of failing to replicate, which almost assures the work will get a lot of attention (e.g., Doyen et al., 2012 been viewed more than 37,000 times when writing this). The prevalence of concerns about replicator’s intentions were investigated by asking psychologists whether they *agreed* or *disagreed* with the suggestions that: “Some replications are conducted by researchers who have bad intentions”, “Replications are usually conducted based on doubts about the original finding”, and “Researchers who report failed replications are mostly interested in career advancement.”

How serious are the methodological problems perceived to be?

Potentially being the starting point for the whole debate are views on the severity of the methodological problems in the field. While many problems in the field have been discussed (Simmons et al., 2011) and several cases of fraud are certain (Funder, et al., 2013; Levelt committee, 2012), notably lacking are any hard data about how replicable the findings in the field *actually* are, and there are little data demonstrating that psychologists believe the field contains many non-replicable effects (but see Fuchs, Jenny, & Fieldler 2012 and Buttlere & Wicherts, in prep). John et al. (2012) presented survey data suggesting that ‘Questionable Research Practices’ were widespread in Psychology, but these findings have never been replicated and, given the high leave before completion rate and somewhat low response rate in this survey, there is some discussion about the actual extent of the problems.

In order to help gauge the field’s trust in its literature and perceptions about the methodological problems, we asked psychologists directly to *agree* or *disagree* (with the same 100 point scale with breaks every 25 as above) with the suggestions that: “The seminal (textbook) effects in my subfield are probably replicable”, “There are many non-replicable results in the psychological literature”, “Most psychologists have engaged in data collection or analysis practices that inflate effect sizes and the rate of false positives (Type 1 errors)”, and “Many psychologists in my subfield engage in Questionable Research Practices”.

What is the future of Psychology?

Aside from examining questions directly related to the most recent debate, we were also interested in gathering opinions on what psychologists think should be done to help move the field forward more generally. Thus, here we asked respondents about 10 of the most discussed proposals for improving the reproducibility and replicability of findings (e.g., Nosek & Lakens,

2014; Asendorpf et al., 2013; Simmons, Nelson, & Simonshon, 2011; Giner-Sorolla, 2012; Levelt committee, 2012; Buttlere, 2014; Kriegskorte, 2012).

All of the suggestions asked to the participant were framed with the stem, “In order to help move the field forward, we should:”, again using a 100 point *strongly disagree* to *strongly agree* scale with breaks every 25 points. The 10 suggestions the participant rated were to: “Conduct more direct replications” (always displayed first; Nosek & Lakens, 2014), “Register studies, hypotheses and analyses plans before studies are conducted.” (Asendorpf et al., 2013), “Demand that authors report how they determined sample size in their studies” (Simmons, et al., 2011), “Demand that authors disclose all data exclusions, all manipulations, and all measures in their studies (if any)” (Simmons, et al., 2011), “Consider sufficient statistical power of studies as a publication criterion” (Simmons, et al., 2011), “Publish the data and stimuli of our papers online” (Asendorpf et al., 2013), “Ensure that the academic incentive system only rewards publishing replicable work” (Giner-Sorolla, 2012), “Establish an official system to investigate and arbitrate accusations of scientific misconduct (e.g., fraud, unprofessional personal attacks, plagiarism)” (Levelt committee, 2012; Funder et al., 2013), “Determine what moderates/mediates replicability of results (e.g., culture, time).” (Klein et al., 2014), and, “Focus on the future rather than the past” (Buttlere, 2014; Kriegskorte, 2012).

Procedure

The introductory email (available at osf.io/4y7gn/) to respondents first mentioned the purpose of the survey, “to assess psychological researchers' views on the replication of psychological studies,” the average time to complete the survey (about 4 minutes) and how many participants had taken part in the survey thus far (with an urge to help us reach the next milestone; 500, 1,000, 1,500). The respondent was ensured their data would, “be confidential and

aggregated at the presentation of results,” and that their “participation is completely voluntary.” After clicking through the email, the participant saw an information and consent page which largely repeated the above information (again available at osf.io/92c4y/). Before the psychologists were asked to indicate their *agreement* or *disagreement* with the statements, they were asked the number of papers they had published to date and their main subfield of research in Psychology.

The first block of questions presented to the participant consisted of the 10 statements focusing on what should be done in order to help move the field forward. The second and third blocks consisted of the 17 statements we designed to help answer some of the open questions that have been debated. Each item was first randomly assigned to a block, and then randomly presented within the block, except that the first question in the second block (ten questions) was always, “Direct replication is the best method of determining which effects are real” and the first question of the third block (seven questions) was always, “There are many non-replicable results in the psychological literature.” Again, participants were asked to indicate on a 100 point *strongly disagree* to *strongly agree* scale how they felt about each of the abstracted statements. After the questions as outlined above, participants were asked if they had any final comments on the subject of replication or the field more generally. The final page of the survey asked the participants if they had any comments about our survey specifically, thanked them again for their participation, and provided a link by which they could provide their email for the results.

Results

Response rate and sample characteristics

Of the 10,000 surveys we sent out, 2,307 were started, and 2,056 were completed entirely, yielding a 20% complete response rate. Especially as the survey was conducted in late

July and early August (when many researchers are on vacation), this suggests a substantial interest in the survey and the topics it covered. An 11% leave rate, while smaller than other surveys (e.g., John et al., 2012; Fuchs, et al., 2012; Mulligan et al., 2012; Ware & Monkman, 2008) still raises questions about how these participants might affect the results. In order to examine this impact, the means and correlations were examined depending on whether all available information was used or only those complete participants. None of the descriptive means changed more than .5 point out of 100, and none of the correlations change more than .02. Given the small impact these individuals have on the data, we move forward using list wise deletion, also in order to limit researcher degrees of freedom (Simmons et al., 2011).

The average psychologist in the survey reported having between 11 and 25 articles. The smallest group was those with over 100 publications, with 214 respondents in this category. There were 356 cognitive psychologists and 354 social psychologists, representing the two largest two subfields in our sample. Only 29 respondents reported being forensic psychologists and only 59 participants indicated being quantitative psychologists. The ‘Other’ subfield category was utilized 274 times by psychologists from such fields as: Counseling, Behavior Analysis, Sport Psychology, Geography, and Psychiatry.

Examining the univariate distributions revealed kurtotic spikes in the data around the breaks in the scale (every 25 points; Figure 1). In order to examine the extent of this problem, we counted the number of times each psychologist indicated a score within three points of one of these breaks (e.g., 22-28, 42-48, 72-78) and found that the average participant answered within these areas 19 times (of 27 questions total). This indicates that the scale was, for many psychologists, utilized more like a 4 button scale, and then going just beyond or below the break to nuance their answer. Though a correction might be possible (to move those who use these

break point areas more up or down based upon whether they are one or two points above or below the break point), we forego the possibility in order to minimize researcher degrees of freedom.

Results in brief

The results of our survey suggest that the community thinks positively of direct replication (and those who do them), but that they also trust the literature, generally speaking, and recognize that there are many intervening factors which might lead to a failed replication (Tables 1 & 2). The overall most supported suggestions/ statements were that, “The same exact materials can mean different things to participants across samples, contexts, and time” ($M = 75.97$; $SD = 21.19$; all SE ’s about 1 point), that, “A successful independent replication is likely to increase other researcher’s trust in the original researcher’s integrity” ($M = 75.84$; $SD = 21.70$), and that, “The seminal (textbook) effects in my subfield are probably replicable” ($M = 65.09$; $SD = 21.16$).

More evidence that the movement is generally accepted comes from the fact that the least endorsed statement was that, “Researchers who report failed replications are mostly interested in career advancement” ($M = 29.62$; $SD = 23.21$). Nearly as unpopular were the ideas that “Some effects can only be achieved by an expert in the field” ($M = 36.43$; $SD = 28.52$), and that “Many psychologists in my subfield engage in Questionable Research Practices” ($M = 37.92$; $SD = 27.56$). With this general picture in mind, we more directly examine the answers we found to the questions we posited above.

Do psychologists value direct replication? Generally speaking, yes.

The results of the survey suggest that the majority of psychologists agree that direct replications are valuable and that replicability is important. Both the suggestions that, “Direct

replication is the best method of determining which effects are real” ($M = 63.51$; $SD = 24.14$; Table 1) and that “Psychological effects reported in ‘flagship’ journals should be widely replicable across labs, stimuli, cultures, and time” ($M = 62.20$; $SD = 25.28$) were among the top five most endorsed statements in the set (though a large portion of psychologists indicated not thinking that flagship results need to be widely replicable). The suggestion that psychologists should conduct more direct replications to help the field move forward was also among the most endorsed suggestions for moving forward ($M = 69.73$; $SD = 23.11$). Especially the notion that psychologists should conduct more replications in the future was one of the more agreed upon statements (low SD).

Are context and cultural differences important in obtaining results? Yes.

The results suggest that psychologists believe context and sample differences can play an important role in obtaining certain results. The most agreed to statement in the entire survey was that, “The same exact materials can mean different things to participants across samples, contexts, and time,” ($M = 75.97$; $SD = 21.19$). Nearly as strongly endorsed was the notion that, “Failed direct replications are often due to unknown substantive moderators (e.g., sample/ experimental setting differences)” ($M = 60.02$; $SD = 20.75$). These two items were again in the top 5 most endorsed statements and suggest that psychologists are sensitive to the contextual and cultural components of the science. These conclusions were also bolstered by the fact that the second most endorsed way to move forward was to, “Determine what moderates/ mediates replicability of results (e.g., culture, time)” ($M = 60.02$; $SD = 20.75$).

Is expertise or craftsmanship important for obtaining results? Most say no.

While those questions about context were among the most agreed to, the ones about expertise were among the least agreed to. For instance, the second least endorsed statement in

the entire survey was that “Some effects can only be achieved by an expert in the field” ($M = 36.43$; $SD = 28.43$). Opinions were more divided as to whether, “Psychological effects sometimes require a craftsmanship to achieve that some replicators might lack” ($M = 47.70$; $SD = 25.45$). Note that while these are among the least agreed to overall, the standard deviations for these questions are notably larger than those for context or thinking that direct replications are valuable, suggesting that this is one of the more contested ideas in the survey (J. Buder, personal communication, 19 September, 2014).

What are the outcomes for original researchers? Good, or at least not bad.

Most psychologists feel a successful replication is good for a researcher’s reputation, and that a failed replication does not (or should not) hurt them. For instance, the notion that, “A successful independent replication is likely to increase other researcher’s trust in the original researcher’s integrity” ($M = 75.84$; $SD = 21.70$) was the second most endorsed statement of the set and the most leptokurtic in the dataset. Conversely, the community hesitantly disagreed that, “A failed replication may lead to accusations of fraud” ($M = 42.36$, $SD = 26.03$) or that “An independent replication that fails is likely to lower the career outcomes of the original researchers” ($M = 42.16$, $SD = 23.96$). Note that these results indicate that many people believe a failed replication leads to fraud accusations and worse outcomes (which might help explain some of the defensiveness). More psychologists agreed that, “Some opponents of the replication movement are only worried about the replicability of their own effects” ($M = 48.20$, $SD = 23.85$) than that failed replications are actually harmful, which potentially is revealing of their attitudes concerning those working against the movement. These results match well with the results concerning context, where the community recognized there are many potential reasons for a failure to replicate.

Are replications undertaken from bad motives? Most respondents say no.

Suggestions that those doing replications have poor intentions were relatively disagreed with, indicating relatively positive attitudes overall. For instance, the least endorsed statement in the survey was that, “Researchers who report failed replications are mostly interested in career advancement” ($M = 29.62$; $SD = 23.21$). Also among the least endorsed suggestions was that “Some replications are conducted by researchers who have bad intentions” ($M = 41.48$; $SD = 26.01$). Opinions were generally split concerning the notion that, “Replications are usually conducted based on doubts about the original finding” ($M = 47.11$; $SD = 26.37$).

How serious are the methodological problems in the field? Opinions diverge.

It appears that most psychologists believe the seminal effects are replicable, while also recognizing that there are probably many non-replicable results in the literature. For instance, the third most endorsed statement in the set overall was that, “The seminal (textbook) effects in my subfield are probably replicable” ($M = 65.09$, $SD = 21.16$). Still agreed to, though slightly less so ($M = 61.78$, $SD = 23.10$), was the suggestion that, “There are many non-replicable results in the psychological literature.” Similar to the results of Buttlere and Wicherts (in prep), those statements suggesting psychologists’ behaviors are the problem (i.e., that they use or used QRPs) were less endorsed; psychologists (naturally?) avoid placing the blame on themselves. The results are sensible though in that the community more strongly agreed QRPs *were utilized* ($M = 52.64$, $SD = 25.38$) in the past than that QRPs *are utilized* currently ($M = 37.92$, $SD = 27.56$). Notably, these questions have large standard deviations, indicating more disagreement in the community.

What *should* Psychology do next? Enforce rules.

What is first notable is that the means are on the positive side, suggesting that psychologists are optimistic about potential changes to the system, as even the least popular suggestion, which was to, “Register studies, hypotheses and analyses plans before studies are conducted” was only 5 points below the middle of the scale (Table 2). Generally, the results suggest that Psychologists are ready to conduct more replications and that they want someone to enforce the rules. These conclusions are evidenced by the fact that the three most endorsed suggestions were to: “Demand that authors disclose all data exclusions (if any), all manipulations, and all measures in their studies” ($M = 72.39$, $SD = 25.15$; All SE ’s about 1 point), “Determine what moderates/ mediates replicability of results (e.g., culture, time)” ($M = 69.85$, $SD = 23.47$), and “Conduct more direct replications” ($M = 69.73$, $SD = 23.11$). Note that there is more disagreement in the community concerning this explicit governing step (larger SD). Standard errors for all statements are about half a point ($N = 2,056$).

While the most endorsed statement (demand disclosure) left it unclear as to who or what exactly would enforce and demand this standard, the fourth most popular suggestion was to, “Establish an official system to investigate and arbitrate accusations of scientific misconduct (e.g. fraud, unprofessional personal attacks, plagiarism)” ($M = 64.13$, $SD = 27.04$). Note that again this question has an even higher SD (and thus more disagreement) than the one suggesting that psychologists need to disclose.

The majority of psychologists also agreed that data and stimuli of our papers should be published and sufficient power should be seen as a publication criterion (Table 2). “Register(ing) studies, hypotheses and analyses plans before studies are conducted” received less support, being the only suggestion with an average definitively below the midpoint of the scale ($M = 44.16$, SD

= 30.24). Notably, this is the question that has the largest standard deviation, which suggests that opinions diverge strongly on this notion. Also regarded with general ambivalence were the suggestions that we should, “Focus on the future rather than the past” ($M = 49.35$, $SD = 25.39$) and, “Ensure that the academic incentive system only rewards publishing replicable work” ($M = 51.69$, $SD = 29.15$). Overall, the results are similar to our last survey (Buttlere and Wicherts, in preparation) in that there appears to be high enthusiasm, but no clear consensus of what to do next.

Structural equation modelling

Besides examining the results descriptively, we also examined the relationships between the questions and hypothesized concepts. For this analysis we used structural equation modelling as implemented in the R package Lavaan (Rosseel, 2012). On the suggestion of Anderson and Gerbing (1988), the model was first developed on a 600 person random subsample, and then confirmed across subfields in the remaining 1,456 person sample. Maximum likelihood estimation was used for the parameter estimates, but the Satorra-Bentler correction was utilized for the standard errors because the assumption of multivariate normality was not upheld (as indicated by the Mardia and Henze-Zirkler multivariate tests of normality; Steiger 1990).

Assessment of model fit was done according to the suggestions of McDonald and Ho (2002), Hooper, Coughlan, and Mullen (2008), and Boomsma (2012) who suggest always reporting the chi square, the degrees of freedom, the RMSEA, and then one or two discretionary indices. Here we chose to use the CFI, as it is relatively unaffected by sample size, and the SRMR in order to more directly understand the residuals of the model (also suggested by Steiger, 1990 and Boomsma, 2012). When developing the model, the latent variances are set to 1 in order

to help interpret the latent factor covariance matrix (which is then essentially a correlation matrix).

Developing a good fitting model

The original model we hypothesized included one factor for each of the questions about direct replication examined above, meaning six overall factors which correspond to the individual's perceptions about: the Value of direct replication, the importance of Context effects, the importance of Expertise, the potential Reputation damage for original author, the perceived Intentions of replicators, and the researcher's concern about the Methodological problems. Implementing this model yielded a mediocre to poor fit (Table 3, Row 1). On the suggestion of McDonald and Ho (2002) and Hooper, Coughlan, and Mullen (2008), the first step we took toward a better fitting model was to remove those items that did not load well on their factor or were not unidimensional.

The summary, modification indices, and residuals of this model suggested several modifications which were implemented in a stepwise manner as described here: First, it was found that the statement about seminal findings being mostly replicable (Item no. 3; ordered by popularity) did not load strongly on the Methods factor ($b = -.15$, $se = .053$, $Z = -2.835$; standardized coefficients shown) and cross loaded on the factors representing the researcher's valuation of direct replication, how much they think Context matters, and the potential Reputation damage that could be done. Removing this statement (Item no. 3) led to an improvement in fit in terms of CFI, RMSEA, and SRMR (Table 3, Row 2) and we thus retained the modification. Both in the original model and after removing Item 3, the statement suggesting that a successful replication is good for the original author's reputation (Item no. 2) did not load well on the potential Reputation damage factor ($b = .25$, $se = .057$, $Z = 4.33$), and it loaded on the

Valuation of direct replications and with the statement suggesting that many psychologists use QRPs. Removing Item 2 also resulted in a better fitting model overall (Table 3, Row 3) in terms of the fit indices. Finally, it was also apparent in all three models that the statement about opponents of replication only being worried about their own effects (Item no. 9) loaded relatively poorly on its own factor ($b = .43$, $se = .053$, $Z = 7.22$), while it also loaded on the Value Direct Replications and the Methods factors. Removing this question again led to a better fitting model in terms of RMSEA, SRMR, and CFI (Table 3, Row 4).

Removing these three ill-fitting questions from the model increased the overall fit of the model to the point that it is considered acceptable by most standards; $\chi^2 (75, n = 600) = 124.94$, RMSEA = .040 [.030, .050], SRMR = .042, CFI = .949 (Table 3, Row 4). From here the goal was to simply identify any small and theoretically sound changes that can be made in order to further improve model fit. The modification indices of this model proposed allowing a link between the residuals of the item suggesting that the field should conduct more replications (Item no. 4s) and the item suggesting that there are many non-replicable results in the literature (Item no. 6). There was also a relatively large modification index which recommended a link between the statements suggesting that flagship results should be replicable (Item no. 5) and that replications are often conducted from doubts (Item no. 11).

The theoretical link between thinking there are many non-replicable results in the literature (Item no. 6) and desiring to conduct more replications (Item no. 4s) makes good theoretical sense and is the larger of the two modification indices; it was thus implemented, leading to a significant decrease in model misfit (Table 3; Row 5). Still the residuals recommended a covariance between the residuals of Item no. 5 (flagship results should be replicable) and Item no. 11 (replications conducted from doubts), but allowing this covariance

did not make as much sense theoretically, so we decided not to implement this change in the final model (Table 3; Row 6). The final measurement model decided upon in this exploratory phase retained the six hypothesized factors, had excellent fit to the data in our exploratory subsample (Table 3; Row 5), and is displayed in Figure 2 (loadings are for the final invariant model across subfields, as described next).

Confirming SEM model across subfields

After having slightly tweaked our originally hypothesized model to fit well in the 600 person cross-field subsample, we sought to confirm this model across subfields using the rest of the sample ($n = 1,456$). This multi-group analysis also allowed us to ensure that the measurement and structural parameters are invariant across subfields, as well as allowing us to explore where subfield differences are. In order to maintain sample size in this analysis, smaller subgroups were combined in the following manner: Clinical and Health Psychology ($n = 328$), Cognitive and Neuroscience Psychology ($n = 364$), Industrial Organizational, Personality and Social Psychology ($n = 351$), and the rest of the subfields including Developmental, Forensic, Quantitative, and Other Psychology ($n = 413$).

The model was first tested for configural invariance across subfields by simultaneously testing the model across the four groups. The overall model fitted well (Table 5, Line 1), suggesting that the factor structure across subfields is largely invariant. The model fitted worst in the Social, Personality, and Industrial Organizational group χ^2 (DF = 74) = 132.3 and best in the Neuro/Cognitive Psychology group χ^2 (DF = 74) = 90.30. Metric equivalence across subgroups was next examined by constraining the factor loadings to be equal across subgroups. The results

suggested that metric equivalence was tenable, meaning that the indicators were largely invariant in how they loaded on the factors across groups (Table 5, Line 2).

Next, strong factorial invariance was tested by restricting the measurement intercepts to be equal across groups. This resulted in a significant deterioration in fit (Table 5, Line 3), suggesting that there is a group based uniform difference (bias) in certain items. To investigate these group based biases in response we inspected the modification indices related to the constrained intercepts and the summary of the model with unconstrained intercepts. These investigations showed that, relative to respondents in the other subfields, cognitive and neuroscience psychologists more strongly endorsed the notion that replications are conducted based upon doubts about the original finding (Item no. 11). Similarly, cognitive/neuro psychologists more strongly endorsed the notion that results published in flagship journals should be highly replicable (Item no. 5) compared to the other groups. Proceeding in a stepwise manner, we released only the intercept in the group that was variant, leaving the other three groups constrained to be equal, first for Item no. 11 (Table 5, Line 3.1), then for Item no. 5 (Table 5, Line 3.2). Both of these modifications significantly reduce the amount of misfit in the model. While there are other small variations (e.g., social/personality/IO group worried more that failed replications would lead to fraud accusations, social/personality/IO group endorsed more the notion that many psychologists in their subfield currently use QRPs, cognitive/neuro group endorsed less the notion that the same materials can mean different things across contexts) none of these allowances modify the model enough to ensure that we are not capitalizing on chance ($MIs < 12$, Bonferroni correction for 45 parameters suggests MI cutoff of 10.63; apriori selected $MI > 15$ and theoretically sound as conditions for freeing the group intercept). After freeing the two intercepts in the cognitive/neuro group, we constrained the residual variances of the

indicators to be equal across the four subfields; which did not lead to any appreciable increase in misfit, suggesting that strict factorial invariance holds (apart from those two intercepts in the single group; Table 5, Line 4).

The next step in the analysis was to examine whether the relationships between the latent constructs in the analysis were invariant across subfields. First, the five regression weights of the structural model (using the factors to predict Valuation of direct replication) were constrained to be equal across groups, which did not lead to a significant increase in misfit (Table 5, Line 5). The latent (residual) covariances were next constrained, which were also found to be invariant across the groups (Table 5, Line 6). Finally, the means of the latent variables were constrained to be equal across groups and these were found to vary. Hence, we explored the differences in subfield on the latent traits (standardized means in Table 7). An examination of these group latent means revealed the following: First, clinical and health psychologists endorsed less strongly those statements suggesting that failed replications will lead to potential Reputation damage. Second, social, personality, and I/O psychologists along with neuro/ cognitive psychologists endorsed more strongly the notion that those doing replications have bad Intentions. Third, the social, personality, and I/O group indicated a greater belief that Expertise can be important in obtaining particular results than the other groups. Finally, the miscellaneous subgroup considered the methodological problems in their field to be significantly less threatening than the other fields. Notably, these latent means were not calculated utilizing the two statements whose intercepts were variant across groups (4 & 11), though this is not a problem as intercept variance was only found on those factors that had three indicators.

With this mostly invariant model (Table 6), we are able to begin answering our final research question, namely, “what leads a researcher to value direct replications?” The results of

this model (Figure 2: Table 6) suggest that the best predictor of psychologists Valuing direct replication is their perceptions of the severity of the Methodological problems in the field. The larger the person perceived the Methodological problems to be, the more they indicated Valuing direct replications ($b = .40$, $se = .06$, $Z = 7.31$). Psychologists perceptions of the Intentions of replicators' intentions ($b = -.35$, $se = .14$, $Z = -2.59$) and their intuitions about how much Expertise is required for experiments affect the results achieved ($b = .10$, $se = .06$, $Z = -1.73$; marginally significant) were also of some importance in predicting psychologist's valuation of direct replication across subfields. It is noteworthy that despite these significant regression weights, the model still only explained about 15% of the variance in Valuing direct replications.

Discussion

Here we gathered psychologists' opinions about: the value of replication, how much context and expertise can influence the results of research efforts, the effects of a replication for the original and replicating researchers, the extent of the methodological problems in the field, and what the field should do next. As far as we know, this is the first published survey of psychologists' opinions on the value of direct replication and what the field should do next, despite the large amount of discussion these topics have elicited, both within the peer-reviewed journals and alternative learning platforms (e.g., Wagenmakers et al., 2011; Pashler, & Wagenmakers, 2012; Roediger, 2012; Simmons, Nelson, & Simonshon, 2011; Buttlere, 2014).

Descriptively, the results indicate that the majority of psychologists agree that direct replications are valuable and should continue being done in the future. At the same time, most psychologists appreciate the contextually sensitive nature of psychological stimuli and experimentation. Concerning the actors involved, the results suggest that the community does not

view those doing replications as having negative intentions, with a corresponding belief that a failed replication does not (and perhaps should not) harm the original author's reputation. Finally, results suggested a general trust in the seminal effects in the field, with an understanding that there are probably also many non-replicable results in the literature. Concerning potential change, psychologists appear to be generally enthusiastic, with even the least popular potential change (pre-registering hypotheses and studies) being only 5 points below the middle of the scale.

The relationships between the indicators and latent variables were found to be largely invariant across subfields. Thinking that the Methodological problems in the field are larger was found to be positively related to valuing direct replications while the perceived negative Intentions of replicator were found to be negatively related to valuing direct replications. Unfortunately, the model could only explain about 15% of the variance in an individual's belief about whether more direct replications should be undertaken or not, indicating that more research is needed to better understand researcher's opinions on these issues. Concerns about potential Reputation damage or Context effects were not found to relate to respondents' valuing direct replications.

This project has been designed as an attempt to apply the scientific method to some of the many open questions the field faces in this time of strife and reformation. It is the opinion of the authors that more work of this nature is needed, not only to better understand the questions the field faces (e.g., What are the largest problems? Should we do more direct replications?), but also to help design the best potential solutions to them (e.g., Which solutions are most likely to work? How can we best implement those solutions?). High enthusiasm for change, as evidenced

by psychologists endorsing almost all potential changes, means little if it cannot be channeled into useful behavioral outcomes.

Limitations

There are several limitations associated with researches of this type. The first limitation we will mention is the 20% overall response rate, which leaves open questions about what the other 80% who did not respond think and whether they may be systematically different than those who participated. This question is difficult to answer, though it is possible to say that a 20% response rate from active researchers in the middle of the summer is relatively high (John et al. had only 23% full response and is one of the best response rates for this type of survey). Similarly, there is a possibility that the 11% of participants who left the survey before completing it were different from those who completed the survey entirely; though again it is important to say that this leave rate is quite good, as even John et al. (2012) had over a third of their participants leave before completing their survey, and that some examination suggested they did not have a large effect on the data. Another limitation is the odd distributions resulting from the breaks in the scale (Figure 1). Certainly if the project were to be run again, it would be better to run it either with more (e.g., 9, 10) or less (e.g., 1 or 0) breaks in the scale.

Probably the largest limitation in this sort of survey (with only 30 questions) is the amount of detail that can be obtained. This limitation is particularly felt in the large number of questions we were unable to answer. Some of the notable questions we could not answer in this research, but which would probably be useful for the field to investigate, are:

- Why are direct replications not undertaken? There have been the suggestion that reviewers and editors are biased against replication (Greenwald, 1975; Neuliep &

Crandall, 1990); but, if they were read and cited, then editors would (probably) publish them. Knowing more about why they are not done can help encourage them.

- What are the relative strengths of direct vs. conceptual replication? There is behavioral evidence that psychologists engage in many more conceptual replications than direct replications (Hubbard & Armstrong, 1994; Neuliep & Crandall, 1993). Why this is the case can tell us more about how to best encourage direct replications (if we want).
- Do those who do replications differ from those who they are replicating? A comparison of average citations between those doing replications and those they are replicating might provide evidence for or against the notion that some replicators lack the expertise necessary to achieve some results (as has been suggested; Yong, 2012).
- Why do researchers in psychology often fail to publish “null” findings? The traditional wisdom is that editors and reviewers are biased against the null (Greenwald, 1975), but several journals have been founded specifically to alleviate the file drawer problem and these efforts have not been particularly fruitful (Simons, 2012). Another potential explanation is that researchers do not want to log their failures or that they believe there was some fundamental flaw in the design or execution.
- Finally, what factors make a study replicable, and how can we maximize this in our experimentation generally? One of the most endorsed suggestions in our survey was to examine what moderates and mediates the replicability of a finding. One way researchers may begin is by running priming and stereotype studies across cultures (where the stimuli are clearly inappropriate or unclear), to examine the extent of context dependent effects.

A final limitation is that when speaking about the future of the field, none of the proposed solutions asked about here had an implementation plan outlined. Research can be done to develop several potential plans and then to ask the community what they think of them, eventually coming up with a 5 year and 10 year plan for the field. In this way the enthusiasm can be channeled into useful outcomes, rather than dispersing in many directions.

Conclusion

Here we have gathered data about the replication debate by asking psychological researchers their opinions about some of the most hotly contested issues in the debate on replication. We found a generally positive attitude toward (direct) replication, with a recognition that there are many reasons a direct replication may fail and the belief that most of the seminal

effects in each subfield are replicable. We hope that this survey will spur more researchers into gathering data surrounding proposed change plans, in order to ensure that the field is making the best decisions possible.

References

- Abbott, A. (2013). Disputed results a fresh blow for social psychology. *Scientific American*, retrieved from <http://www.scientificamerican.com/article/disputed-results-a-fresh-blow-for-social-psychology/>
- Amir, Y., & Sharon, I. (1990). Replication research: A "must" for the scientific advancement of psychology. *Journal of Social Behavior & Personality*, 5(4), 51-69.
- Anderson, J. C., & Gerbing, D. W. (1988). Structural equation modeling in practice: A review and recommended two-step approach. *Psychological Bulletin*, 103(3), 411.
- Asendorpf, J. B., Conner, M., De Fruyt, F., De Houwer, J., Denissen, J. J., Fiedler, K., ... & Wicherts, J. M. (2013). Recommendations for increasing replicability in psychology. *European Journal of Personality*, 27(2), 108-119.
- Bakker, M., & Wicherts, J. M. (2011). The (mis) reporting of statistical results in psychology journals. *Behavior Research Methods*, 43(3), 666-678.
- Bargh, J. A., Chen, M., & Burrows, L. (1996). Automaticity of social behavior: Direct effects of trait construct and stereotype activation on action. *Journal of Personality and Social Psychology*, 71(2), 230.
- Boomsma, A. (2000). Reporting analyses of covariance structures. *Structural Equation Modeling*, 7, 461 - 483.
- Buttlere, B. (2014). How Science and Psychology can be used to improve the dissemination and evaluation of scientific work. *Frontiers in Computational Neuroscience*.
- Buttlere, B., & Wicherts, J. (in prep). Psychological researchers' views on peer review: A forward focused survey.

- Cesario, J. (2014). Priming, replication, and the hardest science. *Perspectives on psychological science*, 9(1), 40-48.
- Chin, J., (2014). Psychological science's replicability crisis and what it means for science in the courtroom. *Journal of Psychology, Public Policy, and Law*, *Forthcoming*.
- Coan, J., (2014, May 30). Negative Psychology: The atmosphere of wary and suspicious disbelief. Retrieved from <https://medium.com/@jimcoan/negative-psychology-f66795952859>
- Darley, J. M., Klosson, E. C., & Zanna, M. P. (1978). Intentions and their contexts in the moral judgments of children and adults. *Child Development*, 66-74.
- Dijksterhuis, A., & van Knippenberg, A., (1998) The relation between perception and behavior, or how to win a game of Trivial Pursuit. *Journal of Personality and Social Psychology* 74, 865–877. doi: 10.1037//0022-3514.74.4.865
- Donnellan, B. (2013, December, 11). Go big or go home – a recent replication attempt. Retrieved from <http://traitstate.wordpress.com/2013/12/11/go-big-or-go-home-a-recent-replication-attempt/>
- Doyen, S., Klein, O., Pichon, C. L., & Cleeremans, A. (2012). Behavioral priming: it's all in the mind, but whose mind? *PloS ONE*, 7(1), e29081.
- Fuchs, H. M., Jenny, M., & Fiedler, S. (2012). Psychologists are open to change, yet wary of rules. *Perspectives on Psychological Science*, 7(6), 639-642.
- Funder, D. (2014, July 11). The real source of the replication crisis. Retrieved from <http://funderstorms.wordpress.com/2014/07/12/the-real-source-of-the-replication-crisis/>

Funder, D., Levine, J., Mackie, D., Morf, C., Sansone, C., West, S., Vazire, S., (2013, June 13).

Task force on publication and research practices. Retrieved 21-10-14 from

<http://www.spsp.org/?page=PubResPracTaskForce>

Garfield, E., & Welljams-Dorof, A. (1990). The impact of fraudulent research on the scientific literature: The Stephen E. Breuning case. *Journal of the American Medical Association*, 263(10), 1424-1426.

Gilbert, D., [DanTGilber] (2014, May 24). sychology's replication police prove to be shameless little bullies: <http://www.psychol.cam.ac.uk/cece/blog> (corrected link). Retrieved from <https://twitter.com/DanTGilbert/status/470199929626193921>

Gilbert, D. (2014, May 28). Some thoughts on shameless bullies. Retrieved from <http://www.wjh.harvard.edu/~dtg/Bullies.pdf>

Giner-Sorolla R. (2012). Will we march to utopia, or be dragged there? Past failures and future hopes for publishing our science. *Psychological Inquiry* 23, 263–266
10.1080/1047840x.2012.706506

Greenwald, A. G. (1975). Consequences of prejudice against the null hypothesis. *Psychological Bulletin*, 82(1), 1.

Hooper, D., Coughlan, J., & Mullen, M. (2008). Structural equation modelling: guidelines for determining model fit. *Articles*, 2.

Hubbard, R., & Armstrong, J. S. (1994). Replications and extensions in marketing: Rarely published but quite contrary. *International Journal of Research in Marketing*, 11(3), 233-248.

Ioannidis, J. P. (2005). Why most published research findings are false. *PLoS Medicine*, 2(8), 124.

- Ioannidis, J.P. (2014) How to Make More Published Research True. *PLoS Medicine*, 11(10): e1001747. doi:10.1371/journal.pmed.1001747
- John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, 23(5), 524-532.
- Johnson, Cheung, & Donnellan (2014). Cleanliness primes do not influence moral judgment. *Social Psychology*, 45(3), 209-215.
- Jones, K. S., Derby, P. L., & Schmidlin, E. A. (2010). An investigation of the prevalence of replication research in human factors. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 52(5), 586.
- Kahneman, D. (2012, September 26). A proposal to deal with questions about priming effects. Retrieved from:
http://www.nature.com/polopoly_fs/7.6716.1349271308!/suppinfoFile/Kahneman%20Letter.pdf
- Klein, R. A., Ratliff, K. A., Vianello, M., Adams Jr, R. B., Bahník, Š., Bernstein, M. J., ... & Woodzicka, J. A. (2014). Investigating variation in replicability. *Social Psychology*, 45(3), 142-152.
- Kriegeskorte N. (2012). Open evaluation: a vision for entirely transparent post-publication peer review and rating for science. *Front. Comput. Neurosci.* 6:79 10.3389/fncom.2012.00079
- Lee, S. W., & Schwarz, N. (2011). Wiping the slate clean psychological consequences of physical cleansing. *Current Directions in Psychological Science*, 20(5), 307-311.

Lee, S. W.S., & Schwarz, N. (2010a). Dirty hands and dirty mouths: Embodiment of the moral–purity metaphor is specific to the motor modality involved in moral transgression.

Psychological Science, 21, 1423–1425.

Lee, S. W.S., & Schwarz, N. (2010b). Washing away post decisional dissonance. *Science*, 328, 709 - 712.

Levelt Committee, (2012). *Flawed science: The fraudulent research practices of social psychologist Diederik Stapel*. <https://www.commissielevelt.nl/>

Makel, M. C., Plucker, J. A., & Hegarty, B. (2012). Replications in psychology research: How often do they really occur? *Perspectives on Psychological Science*, 7(6), 537-542.

McDonald, R. P., & Ho, M-H. R. (2002). Principles and practice in reporting structural equation analyses. *Psychological Methods*, 7, 64 – 82.

Meyer, N. M., Chabris, C., (2014, July 31). Why psychologists’ food fight matters: “Important findings” haven’t been replicated, and science may have to change its ways. Retrieved from http://www.slate.com/articles/health_and_science/science/2014/07/replication_controversy_in_psychology_bullying_file_drawer_effect_blog_posts.html

Mitchell, J., (2014). On the emptiness of failed replications. Retrieved from http://wjh.harvard.edu/~jmitchel/writing/failed_science.htm

Mulligan, A., Hall, L., & Raphael, E. (2013). Peer review in a changing world: An international study measuring the attitudes of researchers. *Journal of the American Society for Information Science and Technology*, 64(1), 132-161.

- Mynatt, C. R., Doherty, M. E., & Tweney, R. D. (1978). Consequences of confirmation and disconfirmation in a simulated research environment. *The Quarterly Journal of Experimental Psychology*, 30(3), 395-406.
- Neuliep, J. W., & Crandall, R. (1990). Editorial bias against replication research. *Journal of Social Behavior & Personality*, 5(4), 85-90
- Neuliep, J. W., & Crandall, R. (1993). Everyone was wrong: There are lots of replications out there. *Journal of Social Behavior & Personality*, 8(6), 1-8.
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology*, 2(2), 175.
- Nieuwenhuis, S., Forstmann, B. U., & Wagenmakers, E. J. (2011). Erroneous analyses of interactions in neuroscience: a problem of significance. *Nature Neuroscience*, 14(9), 1105-1107.
- Nosek, B. A., & Lakens, D. (2014). Registered reports. *Social Psychology*, 45 (3), 137-141.
- Pashler, H., & Wagenmakers, E. J. (2012). Editors' introduction to the special section on replicability in psychological science a crisis of confidence? *Perspectives on Psychological Science*, 7(6), 528-530.
- Radford, B., (2012; September 12). Controversial ESP study fails yet again. Discovery News. Retrieved from <http://news.discovery.com/human/psychology/controversial-esp-study-fails-yet-again-120912.htm>
- Ritter, R. S., & Preston, J. L. (2011). Gross gods and icky atheism: Disgust responses to rejected religious beliefs. *Journal of Experimental Social Psychology*, 47(6), 1225-1230.
- Roediger, H. L. (2012). Psychology's woes and a partial cure: The value of replication. *APS Observer*, 25(2), 9.

- Schnall, S., Benton, J., & Harvey, S. (2008). With a clean conscience? Cleanliness reduces the severity of moral judgments. *Psychological Science*, 19 (12), 1219-1222.
- Schnall, S. (2014, May 22). *An experience with a registered replication project*. Retrieved from <http://www.psychol.cam.ac.uk/cece/blog>
- Shanks, D. R., Newell, B. R., Lee, E. H., Balakrishnan, D., Ekelund, L., Cenac, Z., ... & Moore, C. (2013). Priming intelligent behavior: An elusive phenomenon. *PloS one*, 8(4), e56515.
- 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.
- Simonsohn, U. (2013). Just Post It The Lesson From Two Cases of Fabricated Data Detected by Statistics Alone. *Psychological science*, 24(10), 1875-1888.
- Smith, N. C. (1970). Replication studies: A neglected aspect of psychological research. *American Psychologist*, 25(10), 970.
- Steiger, J. H. (1990). Structural model evaluation and modification: An interval estimation approach. *Multivariate Behavioral Research*, 25(2), 173-180.
- Stroebe, W., & Strack, F. (2014). The alleged crisis and the illusion of exact replication. *Perspectives on Psychological Science*, 9(1), 59-71.
- Ware, M., and M. Monkman, (2008). Peer review in scholarly journals: Perspective of the scholarly community, an international study.
- Yong, E., (2012). A failed replication draws a scathing personal attack from a psychology professor. *Discover blog*, accessed one 13-8-14 at: <http://blogs.discovermagazine.com/notrocketscience/2012/03/10/failed-replication-bargh-psychology-study-doyen/#.U-uohPmSx8E>

Zhong, C. B., Strejcek, B., & Sivanathan, N. (2010). A clean self can render harsh moral judgment. *Journal of Experimental Social Psychology*, 46(5), 859-862.

Tables and Figures

Table 1: Descriptive statistics for those statements particularly aimed at answering our six main questions. Full questions can be seen in the survey section. Factor names: C = Context, V = Value direct replications, E = Expertise, R = Reputation Damage, I = Intentions of replicators, M = Methodological problems. $N = 2,056$

Variable	M (SD)	Med.	Skew	Kurt.	F
1. The same materials can mean different things across contexts.	75.97 (21.19)	76	-0.98	1.05	C
2. Successful replication likely increases trust in orig. res. integrity.	75.84 (21.70)	76	-1.10	1.28	R
3. The seminal effects in my subfield are probably replicable.	65.09 (21.16)	66	-0.54	0.28	M
4. Direct replication is the best method of determining which real.	63.51 (24.14)	64	-0.53	-0.08	V
5. Effects in ‘flagship’ journals should be widely replicable.	62.20 (25.28)	64	-0.55	-0.13	V
6. There are many non-replicable results in the literature.	61.78 (23.10)	63	-0.35	-0.13	M
7. Failed replications often due to unknown moderators.	60.02 (20.75)	60	-0.42	0.39	C
8. Most psychologists engaged in practices that inflate effect sizes.	52.64 (25.38)	50	-0.19	-0.40	M
9. Some opponents are worried about replicability of their effects.	48.20 (23.85)	50	-0.14	-0.18	R
10. Effects sometimes require a craftsmanship to achieve.	47.70 (25.45)	50	-0.19	-0.64	E
11. Replications are usually based on doubts about the original.	47.11 (26.37)	50	-0.08	-0.80	I
12. A failed replication may lead to accusations of fraud.	42.36 (26.03)	50	0.10	-0.78	R
13. A failed replication lowers outcomes of original researcher.	42.16 (23.96)	49	0.06	-0.58	R
14. Some replications are conducted by researchers with bad intent.	41.84 (26.01)	50	0.08	-0.66	I
15. Many psychologists in my subfield engage in QRPs.	37.92 (27.56)	36	0.36	-0.80	M
16. Some effects can only be achieved by an expert in the field.	36.43 (28.52)	34	0.33	-0.99	E
17. Reporters of failed replications want career advancement.	29.62 (23.21)	25	0.51	-0.37	I

Table 2: Descriptive statistics for potential future actions. All statements were stemmed with, “In order to help move the field forward, we should...” and randomly presented except that “Conduct more direct replications” was always first. Here questions are presented in order of popularity. The scale was 0 (*Strongly disagree*) to 100 (*Strongly agree*). $N = 2,056$

Variable name	M (SD)	Med.	Skew	Kurt.
1. Publications				
2. Demand disclose all data exclusions, measures	72.39 (25.15)	76	-0.92	0.36
3. Determine what moderates replicability of results.	69.85 (23.47)	75	-0.86	0.57
4. Conduct more direct reps.	69.73 (23.11)	75	-0.72	0.21
5. Est. system to investigate accts. of misconduct.	64.13 (27.04)	67	-0.57	-0.31
6. Publish the data and stimuli of our papers online.	62.50 (27.92)	66	-0.45	-0.64
7. Consider sufficient power a publication criterion.	60.68 (27.39)	64	-0.47	-0.55
8. Demand report how determined sample size.	59.44 (27.89)	62	-0.37	-0.67
9. Ensure incentive sys only rewards replicable work.	51.69 (29.15)	50	-0.15	-0.91
10. Focus on the future rather than the past.	49.35 (25.39)	50	-0.12	-0.38
11. Register plans before studies are conducted	44.16 (30.24)	50	0.16	-1.03

Table 3: Fit measures of steps toward a well-fitting model in the 600 person cross-field exploratory subsample. Model 5 is chosen as best and confirmed across subfields.

Model	df	χ^2	Δ df	$\Delta \chi^2$	CFI	AIC	rmsea	SRMR	GFI
1. Orig.	120	371.73			.818	28881	.066	.070	.920
2. Drop m25	104	290.32	16		.858	27142	.061	.063	.933
3. Drop r18	89	203.16	15		.903	25491	.053	.054	.952
4. Drop r19	75	124.94	14		.949	23831	.040	.042	.968
<u>5. Cov d4 m26</u>	<u>74</u>	<u>115.76</u>	<u>1</u>	<u>08.65**</u>	<u>.956</u>	<u>23822</u>	<u>.037</u>	<u>.041</u>	<u>.971</u>
6. Cov d15 r22	73	111.06	1	05.17*	.959	23819	.036	.039	.972

Table 4: Pearson product moment correlations of the final model in the exploratory subsample $n = 600$. All $SEs < .09$.

	1.	2.	3.	4.	5.	6.
1. Valuing direct replication	1					
2. Context effects matter for results	-.09	1				
3. Expertise matters for results	-.11	.50	1			
4. Reputational damage, potential	-.02	.30	.33	1		
5. Intentions of replicators	-.13	.43	.45	.80	1	
6. Methodological problems	.30	.07	.20	.36	.32	1

Table 5: Model fit and change parameters for those models testing invariance across subgroups. All model χ^2 s are significant at the .001 level (because of the sample size). Overall, all models fit quite well. $n = 1,456$

Model	df	χ^2	Δdf	$\Delta\chi^2$	cfi	aic	rmsea	srmr
1. Model across fields	296	465.9			.931	58982	.049	.046
2. Equal loadings	323	488.2	27	21.53	.932	58952	.047	.048
3. Loadings, Intercepts	350	582.66	27	114.45***	.913	58995	.051	.052
3.1 Ldngs, Ints, Loosen group 4 r22	349	562.25	-01	-20.41***	.919	58975	.049	.051
3.2 Ldngs, Ints, Loosen group 4 d15	348	546.45	-01	-15.80***	.923	58960	.048	.050
4. Ldngs, Ints, Residuals	393	603.06	45	55.33	.919	58929	.046	.053
5. Ldngs, Ints, Rsdls, Regressions	408	615.89	15	13.53	.919	58916	.045	.054
6. Ls, Is, Rs, Rgns, Latent Covariances.	438	654.21	30	38.61	.914	58902	.045	.059
7. Ls, Is, Rs, Rgns, L Cov., Latent Means	456	715.07	18	72.07***	.903	58928	.047	.064

Table 6: Full summary of final model across subfields. Only two differences allowed are the intercepts of 5 and 11. Displayed graphically in Figure 2, fit statistics can be seen in Table 5.

Measurement Model	Est.	SE	Z	p
Value direct replications				
4. Direct replication is the best method of determining which real.	1			
4s. Conduct more direct reps.	.95	.07	14.21	0
5. Effects in ‘flagship’ journals should be widely replicable.	.61	.05	12.50	0
Context matters				
7. Failed replications often due to unknown moderators.	1			
1. Same exact materials can mean different things across contexts.	.55	.08	6.69	0
Expertise matters				
10. Effects sometimes require a craftsmanship to achieve.	1			
16. Some effects can only be achieved by an expert in the field.	.81	.07	12.42	0
Reputation damage authors				
12. A failed replication may lead to accusations of fraud.	1			
13. A failed replication lowers outcomes of original researcher.	.93	.06	14.85	0
Bullying				
14. Some replications are conducted by researchers with bad intent.	1			
17. Reporters of failed replications want career advancement.	.88	.06	14.62	0
11. Replications are usually based on doubts about the original.	.71	.06	12.34	0
Methods				
15. Many psychologists in my subfield engage in QRPs.	1			
8. Most psychologists engaged in practices that inflate effect sizes.	.95	.06	15.36	0
6. There are many non-replicable results in the literature.	.76	.05	13.98	0
Regression equation predicting valuing direct replication	Est.	SE	Z	p
Context matters	.03	.08	.44	.66
Expertise matters	-.10	.06	-1.75	.08
Reputation damage authors	.14	.13	1.08	.23
Intentions of replicators	-.35	.14	-2.52	.01
Methodological problems	.40	.06	7.31	0
(Latent) Covariances	Est.	SE	Z	p
4s. (Do more Reps) ~~ 6. (many non-replicable results)	.09	.05	.392	.70
Context matters				
Expertise matters	.28	.03	9.31	0
Reputation damage original authors	.15	.03	5.24	0
Intentions of replicators	.20	.03	7.40	0
Methodological problems	.03	.03	1.32	0.19
Expertise matters in effect creation				
Reputation damage original authors	.17	.025	6.94	0

PSYCHOLOGISTS ON REPLICATION

Intentions of replicators	.23	.025	9.32	0
Methodological problems	.11	.023	4.84	0
Reputation damage original authors				
Intentions of replicators	.355	.026	13.59	0
Methodological problems	.171	.022	7.70	0
Intentions of replicators				
Methodological problems	.15	.022	6.98	0
<hr/>				
Intercepts:	Est.	SE	Z	p
4. Direct replication is the best method of determining which real.	.08	.05	1.6	.09
5. Effects in 'flagship' journals should be widely replicable.	-.01	.04	-.21	.83
Group 4	.23	.06	3.96	0
4s. Conduct more direct reps.	.07	.04	1.69	.09
7. Failed replications often due to unknown moderators.	-.02	.05	-.45	.66
1. Same exact materials can mean different things across contexts.	-.03	.04	-.76	.45
10. Effects sometimes require a craftsmanship to achieve.	-.05	.05	-1.03	.31
16. Some effects can only be achieved by an expert in the field.	-.04	.05	-.80	.43
12. A failed replication may lead to accusations of fraud.	-.22	.05	-4.48	0
13. A failed replication lowers outcomes of original researcher.	-.19	.05	-4.13	0
14. Some replications are conducted by researchers with bad intent.	-.12	.05	-2.46	.01
17. Reporters of failed replications want career advancement.	-.11	.05	-2.43	.02
11. Replications are usually based on doubts about the original.	-.15	.04	-3.52	0
Group 4	.14	.07	2.15	.03
15. Many psychologists in my subfield engage in QRPs.	.02	.05	.43	.67
8. Most psychologists engaged in practices that inflate effect sizes.	.01	.05	.27	.79
6. There are many non-replicable results in the literature.	.02	.04	.54	.59
Valuing direct replications	0			
Context matters	0			
Expertise matters	0			
Reputation damage authors	0			
Intentions of replicators	0			
Methodological problems	0			
<hr/>				
Variances:	Est.	SE	Z	p
4. Direct replication is the best method of determining which real.	.43	.04		
4s. Conduct more direct reps.	.50	.04		
5. Effects in 'flagship' journals should be widely replicable.	.79	.04		
7. Failed replications often due to unknown moderators.	.55	.08		
1. Same exact materials can mean different things across contexts.	.89	.04		
10. Effects sometimes require a craftsmanship to achieve.	.37	.05		
16. Some effects can only be achieved by an expert in the field.	.63	.04		
12. A failed replication may lead to accusations of fraud.	.55	.04		
13. A failed replication lowers outcomes of original researcher.	.63	.04		

PSYCHOLOGISTS ON REPLICATION

14. Some replications are conducted by researchers with bad intent.	.56	.04
17. Reporters of failed replications want career advancement.	.67	.04
11. Replications are usually based on doubts about the original.	.79	.03
15. Many psychologists in my subfield engage in QRPs.	.51	.04
8. Most psychologists engaged in practices that inflate effect sizes.	.58	.04
6. There are many non-replicable results in the literature.	.73	.03
Valuing direct replications	.46	.05
Context matters	.51	.08
Expertise matters	.63	.06
Reputation damage authors	.45	.04
Intentions of replicators	.43	.04
Methodological problems	.49	.04

PSYCHOLOGISTS ON REPLICATION

Table 7: Latent means per subfield. All SEs < .073. $n = 1,456$.

	Clinical	Cognitive	Social	Other
Value direct replication	0	-.06	-.17*	.02
Context matters	0	.02	.08	-.03
Expertise matters	0	.06	.17*	.05
Reputation damage authors	0	.27***	.34***	.19**
Intentions of replicators	0	.072	.25***	.08
Methodological problems	0	.05	.02	-.14*

PSYCHOLOGISTS ON REPLICATION

Figure 1: Distribution of responses for the question, ‘Some opponents of the replication movement are only worried about the replicability of their own effects.’ The mean of this question is 48.20 with a standard deviation of 23.85.

Figure 2: Depiction of the final model tested across subfields. An * indicates that the intercept for this question in the cognitive/neuro group was released, which can be investigated in Table 6, where the full model output can be seen. All loadings except the residual covariance are significant at the .001 level (***).