Cleansing effects still lack empirical support: Rejoinder to Lee and Schwarz' rejoinder

Ivan Ropovik & Hans IJzerman

Rejoinder to Lee, S., & Schwarz, N. (2021). Grounded procedures in mind and society. *Behavioral and Brain Sciences*, 44(E29). doi:10.1017/S0140525X20000643

In Spike W. S. Lee and Norbert Schwarz' recently published BBS target article "Grounded procedures: A proximate mechanism for the psychology of cleansing and other physical actions" (2020), the authors outline proximal mechanisms underlying so-called cleansing effects. In this rejoinder, we present a rejoinder to their rejoinder, first briefly discussing the portion of the target article we commented on, then discussing the disagreement we voiced in our commentary on their target article, then we discuss Lee and Schwarz' rejoinder, and we finish with our rejoinder to their rejoinder. Before anything else, we want to voice our appreciation for Lee and Schwarz' meta-analytic assessment of the literature and their rejoinder to our voiced criticism. These disagreements are vital to identify stronger versus weaker theories in our science. In this rejoinder to their rejoinder, we make explicit the differences between our appraisal of evidence and theirs and briefly discuss why authors cannot make their set of inferences a "moving target". All in all, it is clear to us that the empirical foundations for cleansing effects, as Lee and Schwarz present them in their BBS target article, are extremely shaky.

Context of the target article we commented on

In their target article and elsewhere, Lee and Schwarz acknowledge that there are more than 200 experiments on cleansing effects yielding more than 500 effects (Lee et al., 2020). Although Lee and Schwarz acknowledge numerous replications of cleansing effects that failed to find an effect, they argued that several successful replications make it difficult to dismiss cleansing effects off hand. In rebuffing our critique, Lee and Schwarz made it appear that our selection criteria were unclear or that we cherry-picked evidence (see their RA.3); because of this, we repeat here the effects we included, which were *entirely based on their own presentation*.

Because they did not give us access to the data underlying their in-progress meta-analysis (see also Explanation 1 below), we identified all *publicly available empirical evidence* that Lee and Schwarz verbally presented to rebut replicability concerns. As they were unwilling to share their data that formed the basis of their claims in their target article, it

left us in the dark whether their claims were backed up by robust evidence. As a result, we simply took all the studies that *they identified* as successful replications.¹

This body of evidence on the replicability of publicly available cleansing effects, selected by Lee and Schwarz, was therefore the focus of our inference as we also stated in our commentary.

Disagreement voiced in our commentary

In our commentary, we thus examined the empirical evidence behind the replication studies that Lee and Schwarz cite as evidence for their claims. Based on the assessment of the evidential value using the p -curve technique (Simonsohn et al., 2014), as well as a data simulation, we concluded the following: based on the evidence Lee and Schwarz lay out in the target article there is a lack of robust evidence for the replicability of cleansing effects and the pattern of data underlying the successful replications of cleansing effects is improbable and most consistent with selective reporting.

The p -curve that we generated based on their own focus on rebutting replicability concerns looks like this:

¹ Unfortunately, their conceptual definition of “replication” was vague. For example, when they try to address replicability concerns, they write: “For example, regarding Schnall et al. (2008), one paper (Johnson et al., 2014b) reported direct replications using American samples (as opposed to the original British samples) and found effect sizes (Cohen’s d s) of .009 and -.016 (as opposed to the original .606 and .852). Another paper (J. L. Huang, 2014) reported extended replications of Schnall et al.’s Experiment 1 by moving the setting from lab to online and by adding a measure (Experiment 1) or a manipulation (Experiments 2 & 2a) of participants’ response effort” and “Yet another paper reported a conceptual replication of the original Experiment 3 by having participants first complete 183 ratings of their own conscientiousness and nominating others to rate their personality (Fayard et al., 2009). This paper also reported a conceptual replication of the original Experiment 4 by changing the design from one-factor (wipe vs. no wipe) to 2 (wipe vs. no wipe) x 2 (scent vs. no scent) x 2 (rubbing vs. no rubbing). Relevant effect sizes were .112 and .230, as opposed to the original .887 and .777.” Thus, to rebut replicability concerns they included both conceptual and extended replications. We thus followed their lead by including these in the p -curve. In one specific instance, this vague conceptual definition meant that our selection led to the inclusion of a different effect than theirs. Specifically, for the following sentences, where they claim a replication: “This finding was replicated with a German sample (Marotta & Bohner, 2013). A conceptual replication with an American sample showed the same pattern and also found that it was moderated by individual differences (De Los Reyes et al., 2012)”, we picked the effect that reported a replication with a US sample, but they include the p -value for a different, moderation effect. We are unsure why they would prefer the moderation by individual differences over the closer replication.

Lee and Schwarz' Rejoinder

Lee and Schwarz wrote a rejoinder to our commentary as well as to the other commentaries. We invite you to read their well-crafted response in full ([rejoinder](#), [supplementary material](#)). They had a couple of points of criticisms on our approach²:

1. “[Ropovik et al.] draw [their] conclusion [that there is lack of robust evidence for the replicability of cleansing effects] on the basis of a *p*-curve analysis of a small subset of the entire body of experimental research on the psychological consequences and antecedents of physical cleansing (namely, seven out of several hundred effects)
2. “..., which included only some of the replication studies and excluded all of the original studies.”
3. “The procedures they applied to the selected studies did not follow core steps of best practice recommendations (Simonsohn, Nelson, & Simmons, 2014b, 2015).”
4. “[They] included *p*-values that should be excluded”
5. “[They] excluded *p*-values that should be included.”

As they suggest we made mistakes that ostensibly completely invalidates our conclusion, they conducted a new curve analysis, which looks like this:

² Lee and Schwarz identified a mistake on our part concerning Camerer et al.'s (2018) failed replication and mentioned it in Footnote 2 in their Response Supplement. We included them as independent, which they should not have been. However, as the results were non-significant, they did not end up in any of our analyses anyways. Our dataset also incorrectly contained a note saying that Arbesfeld et al. (2014) and Besman et al. (2013) did not disclose the use of a one-tailed test. We are sorry about both of those slips and thank Lee and Schwarz for pointing them out. Both of these mistakes leave the *p*-curve identical.

These two *p*-curves demonstrate the completely opposite pattern. While our *p*-curve shows evidence of selective reporting, theirs shows evidence of evidential value. How can this be?

Rejoinder to Lee and Schwarz' Rejoinder

Their critique 3 is easily rebuffed, as we simply took what they saw as replications of cleansing effects (and we did provide a disclosure table, see also Explanation 2). Beyond that, we thought it might be helpful to make the assumptions and interpretational consequences of both of our approaches explicit and list the resulting changes to the *p*-curve data needed to get from our *p*-curve to theirs *p*-curve two to address their concerns 1, 2, 4, and 5. We will also articulate why we see our approach as a more adequate way to appraise the merits of a set of published claims – one that is also much more in concord with the substantive inferences drawn in the original studies and the target article by Lee and Schwarz. Here is a summary of the changes to our data what Lee and Schwarz did to arrive at their *p*-curve from the target article to their rejoinder, which led to the more favorable *p*-curve:

1. In their original article, they describe three effects as successful replications: “[this effect was] successfully replicated in two other direct replications (Arbesfeld et al., 2014; Besman et al., 2013)” and “This finding was replicated with a German sample (Marotta & Bohner, 2013)”. Marotta and Bohner (2013) reported a significant effect (at $p = .05$), which was treated as a successful replication by both the original authors and by Lee and Schwarz in their target article. Yet, for the rejoinder Lee and Schwarz recomputed the *p*-value and treated it as no longer significant. This is problematic for two reasons. First, they changed the interpretation from target article to rejoinder. Second, based on the available information, it is unclear whether the *p*-value was truly above or below .05 (as these studies are not published as full papers, there is very little information about the design and analysis). Similarly so, Arbesfeld et al. (2014) and Besman et al. (2013) each formulated one-tailed hypotheses themselves, which Lee and Schwarz inappropriately transformed into two-tailed hypotheses.
2. Another mistake they made is that they selected a different effect than what was appropriate for the target of inference – evidence on the replicability of cleansing effects. In their *p*-curve, they selected a three-way interaction instead of what was the replication effect from De Los Reyes et al. (2012). For the three-way interaction (which included an addition of individual differences, which was not part of the original concept of the cleansing effect) a *p*-value of .021 was reported; for what they

cite in their target article as a replication, the replicated two-way interaction had a p -value of .048.

3. They included three p -values from a single, publicly unavailable conference poster (it was linked in the reference list, but produced an error when visiting the link). As it was not publicly available and Lee and Schwarz were unwilling to share their data, it did not form part of our inference as we clearly stated in our commentary.

Nevertheless, after examining their dataset, it shows a $N = 10$ per cell, all yielding significant and small p -values with extraordinarily – and unbelievably – large effect sizes, equivalent to $d = 1.55, 1.49$, and 1.84 . Furthermore, for their rejoinder, they decided to add Experiment 1 to their p -curve, while in the target article they only considered Experiment 2 as a conceptual replication. For the full list of L&S' changes to our p -curve set, see Table 1.

Table 1. L&S' changes to our p -curve data

Paper	Reported p -value	Mean N per cell	Note	Changes by L&S
Huang, 2014	.04	54		None
Huang, 2014	.04	109	Effect found only in low-effort condition	None
Huang, 2014	.006	111	Effect found only in low-effort condition	None
Arbesfeld et al., 2014	.03	28	One-tailed hypothesis	Re-computed as a different, two-tailed hypothesis, $p = .071$. Considered a successful replication by L&S in the target article and elsewhere.
Besman et al., 2013	.039	28	One-tailed hypothesis	Re-computed as a different, two-tailed hypothesis, $p = .077$. Considered a successful replication by L&S in the target article and elsewhere.
Marotta & Bohner, 2013	.05	38		Re-computed as $p = .057$. Considered a successful replication by L&S in the target article and elsewhere, while re-computed differently as .054 in Lee & Schwarz, 2018)
De Los Reyes, 2012	.048	12		L&S took the wrong effect ($F(1, 44) = 5.77, p = .021$; page 5). The results of the replication study are reported on p. 4, part "Replicating Lee and Schwarz's (2010) Clean Slate Effects". Should be as in our original analysis $F(1,46) = 4.14, p .048$.
Moscatiello & Nagel, 2014	.0062 (re-computed)	10		Added data from publicly unavailable conference poster.
Moscatiello & Nagel, 2014	.0036 (re-computed)	10		Added data from publicly unavailable conference poster.
Moscatiello & Nagel, 2014	.0006 (re-computed)	10		Added data from publicly unavailable conference poster.

Note. Grey color = effects presented by Lee and Schwarz' target article as successfully replicated and used in our p -curve analysis. Orange color = changes to our p -curve set by Lee and Schwarz. Green color = effects added by Lee and Schwarz. P -values in **bold** represent the effect set used by Lee and Schwarz. *Italicized* p -values were common to both analyses.

But let's try to accept their transformations from target article to rejoinder. If, per Lee and Schwarz, there are two published articles and one conference poster, yielding *a mere 7*

effects evidencing a successful replication with a median $N = 12$ per cell³, we honestly don't see why L&S denote our argument ("lack of robust evidence for the replicability of cleansing effects") as being "strong" or even controversial. In fact, even *if* their p -curve demonstrates an effect, such extremely modest sample sizes with incredibly large effect sizes in only a few studies that they describe in a target article should prompt anyone to investigate more carefully and question the efficacy of the p -curve under such conditions. The meta-analytic techniques we suggested in our commentary allow them to do just that.⁴

Conclusion

All in all, we have clearly shown here again why there is a lack of evidence for the replicability of cleansing effects based on the evidence Lee and Schwarz present. We again show why the successful replications are in fact consistent with the non-successful replications. The answer to our challenge of the data is not to apply an analysis approach that changes the inference criteria they themselves set forth. Instead, close, pre-registered replications will provide a better answer.

Finally, we wanted to comment on these constantly changing inference criteria. Lee and Schwarz likely did not consider the discrepancies between their article and their re-computation, as well as the addition of a study, important enough to warrant a mention in their response. So, the interpretation from the target article that these effects are evidence in favor of the replicability of cleansing effects remains unchallenged and may continue to misguide readers. What this process illustrates is one of the symptoms of an all-too-common problem – a loose derivation chain from theoretical premises, to statistical instantiation of these premises, to substantive inferences. In such instances, the same evidence can be used as a rhetorical device to support exactly the opposite stances. Such moving targets create weak theories, and rebuff solid critiques of one's work. In the end, it all comes down to what one considers adequate empirical evidence for a scientific claim.

Further explanations

Explanation 1

As an additional response to their first critique ("[Ropovik et al.] draw [their] conclusion [that there is lack of robust evidence for the replicability of cleansing effects] on

³ The median N for the non-significant replication effects happens to be 12 times higher, $N = 144$.

⁴ In our commentary, we critiqued the analytical methods they describe in their target articles, as "both their bias tackling workhorses, fail-safe N and trim-and-fill, are known to rest on untenable assumptions and are long considered outdated". We reasoned that the authors should instead apply state-of-the-art correction methods like the regression-based (Stanley & Doucouliagos, 2014) and especially the multiple-parameter selection models (e.g., McShane et al., 2016) by default to examine their claims. Such methods can help detect extremely shaky evidence, such as the case for a mere 7 effects with a median $N = 12$ per cell.

the basis of a *p*-curve analysis of a small subset of the entire body of experimental research on the psychological consequences and antecedents of physical cleansing (namely, seven out of several hundred effects))), we wanted to complement our critique by discussing the history of our conversation with Lee and Schwarz.

After reading their target article and before writing our commentary, we asked Lee and Schwarz to share the data underlying their recent meta-analysis of which the conclusions they incorporated in the target article. We strongly believed that their evidence, based on what they described, was not as strong as they claimed it to be. As the meta-analysis was one of the core components of their target article, we deemed independent verification to be of crucial importance. They declined our invitation for independent verification as “the meta-analytic review was still being written up and any quantitative presentation of its results would prevent them from submitting the manuscript to *Psychological Bulletin*”.

We accepted their refusal to share the data. As we believed their bias-tackling workhorses to rest on untenable assumptions and to be outdated⁵, we chose instead to assess the evidential value of the *replication evidence as Lee and Schwarz present it*. There are very few replications of cleansing effects (with only a minority showing success). Being the leading experts in their field, Lee and Schwarz either reported all the successful ones or chose to present a subset that we reasonably thought would be the best ones.

Maybe there are other replication studies with feeble evidence or problematic designs. Maybe there are far more failed replications. We don’t know as we did not receive the data from the authors and we simply analyzed their “qualitative insights” (Communication with original authors, 2020). On the one hand, that makes it a non-standard way of synthesizing the evidence. But on the other hand, we regard it as a steel-man way to appraise only the merits of the evidence behind the studies that Lee and Schwarz themselves hand-picked as prominent examples of the literature to support the vital auxiliary assumption of their theory – replicability.⁶

⁵ In their rejoinder to our commentary, they indicated that our observation that trim-and-fill and fail-safe *N* are long considered outdated more reflects our sentiments than the standards of the field, because recent meta-analyses published in *Psychological Bulletin* still employ these methods. We thought this was a pretty funny way of arguing a point. Perhaps the authors missed this part of our commentary, but Becker (2005), Ferguson and Heene (2012), and Stanley and Doucouliagos (2014) have clearly shown these methods to be outdated and we prefer to rely on science over engaging in *argumentum ad populum*. It is however true, as they state, that psychology sometimes uses outdated methods. For example, while McDonald’s Omega should be used instead of Cronbach’s Alpha in most instances (Dunn et al., 2014; Revelle & Zinbarg, 2009; Sijtsma, 2009), some researchers stubbornly resist from updating their methodology (e.g., Hauser & Schwarz, 2020). Or consider the fact that it has been known for years that sufficiently powering one’s studies is necessary to reduce the chance of obtaining a false positive. Researchers still stubbornly persist in underpowering their research, even years even after the Bem (2011) and Simmons et al. (2011) articles (e.g., Lee & Schwarz, 2014).

⁶ That said, we fully agree that we drew the conclusion about the evidence for the replicability of cleansing effects based on a small (rather tiny) subset of the relevant literature. We regard it self-evident that if the target of inference is evidence of replicability, original studies are to be excluded. Why didn’t we search for all conducted replications? Because the target of our inference was replication evidence that Lee and Schwarz presented as such and because a sizeable proportion of studies were not part of public record. Apparently, there are only a handful of studies that set out to replicate an experiment on cleansing effects, and the only ones that seemed successful were severely underpowered.

Explanation 2

Lee and Schwarz claimed that we didn't follow best practice because we (1) haven't put together a *p*-curve disclosure table and because (2) we did not re-compute the *p*-values that represent the input for the *p*-curve. The first claim is simply false. Still, this point detracts from the main point of disagreement. The point of a disclosure table is to identify the target effect in a study to ensure that the synthesized effect was the focal effect of the study. In this case, Lee and Schwarz, not us, were the ones who identified the focal effects in their target article. We just followed their lead. For every individual effect, our table clearly identifies the paper and study it comes from, quotes the text string where the effect is reported in the text, effect size, reported *p*-value, *N* for the given test, and the author's inference whether the effect was found or not. We also coded numerous other data about the measurement properties of the dependent measure.

Regarding the second objection, this is a more interesting disagreement. Of course, we understand the importance of re-computing effect sizes or test statistics for any other ordinary evidence synthesis as we have done elsewhere. Refraining from re-computing the focal results of the studies listed by Lee and Schwarz, and taking the reported evidence at face value was, however, a conscious choice. There were two reasons for that. As we explicitly stated, our goal in this very specific analysis was to appraise the merits of a *finite set of empirical evidence, as used by Lee and Schwarz to support their proposed theory*. There was no goal to infer beyond that finite set or estimating some true underlying effect size. In such a case, it makes most sense to take the relevant evidence as it stands. First and foremost, the biasing selection process is not guided by re-computed *p*-values. Second, few practitioners or members of the public re-compute *p*-values when they read the conclusions of a study and adjust their reading accordingly. So do few colleagues making decisions about what hypotheses to pursue next or creating theories (just like the one concerning "Grounded Procedures"). In their target article, Lee and Schwarz seemed to form no exception (but now reading their rejoinder, we sometimes wonder whether the target article and the rejoinder were written by a different set of persons).

What's more, a sizable proportion of significant effects that both – Lee and Schwarz (in the target article) as well as replication authors – presented as successful replications, turn non-significant after their re-computation. Lee and Schwarz likely did not consider the discrepancy between their article and their re-computation important enough to warrant a mention in their response. So, the interpretation from the target article that these effects are evidence in favor of the replicability of cleansing effects remains unchallenged and may continue to misguide readers. What this process illustrates is one of the symptoms of an all-

too-common problem – loose derivation chain from theoretical premises, to statistical instantiation of these premises, to substantive inferences. In such instances, the same evidence can be used as a rhetorical device to support exactly the opposite stances.

Further, any evidence synthesis requires that there is at least basic information regarding the study design and analytical approach. In this case, half of the successfully replicated effects came from unpublished studies with no full-fledged empirical paper available. Re-computation of p -values would require taking a leap of faith because critical pieces of information were frequently missing. For instance, the replication authors may have not assumed equal group variances in a t -test (like the re-computation assumes) and instead of reporting the df for the Welsch's test (not a whole number), they just reported $N - 2$ as df . The analytic sample size may not have been equal to, e.g., $df + 2$ in a two-sample t -test. The replication authors might have excluded some participants for legitimate reasons.

Lee and Schwarz also assumed zero effect of rounding the test statistics on the p -value by their re-computation. They further presume that a two-tailed value was always the proper statistical translation of the substantive hypothesis. In some instances, it was not clear what exact statistical model they used and whether it was parametric at all. Lastly, as is obvious from the p -curve, our conclusion was not contingent upon the decision to re-compute the p -values. Namely, we would have arrived at the *exact* same conclusion even if we had re-computed the p -values – lack of robust evidence for the replicability of cleansing effects. We think that it is just fair that we gave the authors of the replication studies (and Lee and Schwarz) the benefit of the doubt and take the reported results of inferential tests at face value. So, if an exact p -value was available, in agreement with the authors' inferences at the given alpha level, and in agreement with the substantive inference made by Lee and Schwarz (as leading experts in the field) in their target article, we took it at face value. Just like the integrity of that replication evidence itself.

Explanation 3

All the above-discussed culminates in our final explanation – the critique that we included p -values that should be excluded, and excluded p -values that should be included. Despite the apparent eloquence of Lee and Schwarz' critique, we thought it might be helpful for a reader to see a transparent and more detailed presentation of changes to the set of p -values by L&S.

1. Arbesfeld et al. (2014) and Besman et al (2013) both tested a directional hypothesis for which they found support ($p = .030$ and $.039$). They claim that the effect

replicated. So do Lee and Schwarz. However, by re-computing the p -value, Lee and Schwarz effectively ignored the fact that the replication authors regarded a one-tailed test as a proper instantiation of the substantive hypothesis. As the biasing selection process functions at a different alpha level for directional hypotheses, the application of the selection model should not force an irrelevant publication threshold. In this case, by forcing a two-tailed test, these effects dropped out of the p -curve set, as this method only includes significant effects. Of course, there are sometimes issues with the use of one-tailed tests in general and in p -curve in particular⁷, and one can discuss how to deal with them. But more importantly, Lee and Schwarz did not find it sufficiently noteworthy to notify the reader about the disconnect between what is claimed in the target article (“successfully replicated”) and the implication of their reanalysis (“these two effects ceased to be successful replications”).

2. Marotta and Bohner (2013) is not a part of the public record. The result is publicly reported only in several of the lead author’s (Spike Lee) papers. In Lee and Schwarz (2018) NHB paper, they report this effect as being associated with $p = .054$. In the present table, the re-computed p -value equaled .0575. However, in their 2018 mini meta-analysis as well as in some other papers (Dong & Lee, 2017; Schwarz & Lee, 2018), it was explicitly stated that the result replicated the original finding. Because it was unclear and the fact that .054 and, e.g., .04999999 is statistically the same effect (Gelman & Stern, 2006), we once again consistently applied the benefit-of-the-doubt principle and regarded it as a significant effect. Namely, it is the substantive inference that practically matters way more than minuscule differences at the 3rd decimal place. Regardless of whether the reader sees this decision as substantiated or not, it is unfortunate that Lee and Schwarz claim successful replication when it suits them.
3. For the De los Reyes study (2012), they synthesized the wrong effect ($F(1, 44) = 5.77, p = .021$; page 5) when in fact, the results of the replication study are reported on p. 4, section “Replicating Lee and Schwarz’s (2010) Clean Slate Effects” where the following (attenuation) interaction effect ($F(1,46) = 4.14$) should have been selected. The former was the moderation by an individual difference variable, the latter the ostensible replication. This focal replication effect is, however, associated with a much higher p -value of .048.
4. The ultimate game-changer was, however, the inclusion of publicly unavailable data from another conference poster (Moscatiello & Nagel, 2014). Again, the target of our

⁷ These include the general bias towards evidential value and different density in the upper part of the p -value distribution under the alternative hypothesis which is, however, irrelevant in this context.

inference was publicly available information, and we thus did not include this. Nevertheless, let's look at these experiments in more detail. First, in their target article, they only considered Experiment 2 as a conceptual replication. Thus, Experiment 1 should not have been included. Nevertheless, they included both Experiment 1 and 2 (where it is not even clear whether the samples were independent), which yielded 4 p -values (.0061, .3613, .0036, and .0006).

Given that all of these p -values were based on an $N = 10$ per cell design, the effect sizes had to be relatively very large for the three significant effects, with an equivalent of d equal to 1.55, 1.49, and 1.84 (we assume between-subjects design). As an additional note, the latter two are the main effect for a focal reversal 2×2 interaction with an effect size that is so large as to be incredulous, $d = 1.66$ ($n_p^2 = .434$). We leave it up to the reader how to judge the merits of such study and the probability of observing 3 such uncommonly large effect sizes using $N = 10$ per cell in this research domain.

Before publishing our rejoinder, we gave Lee and Schwarz 1.5 weeks to address our concerns. After posting, they published a response on PsyArxiv (available [here](#) and in the comments below). We think that at this point, the reader has sufficient information to judge the replicability of cleansing effects and we will write no further reply. We only note that Lee and Schwarz have now concluded twice that they don't consider Arbesfeld et al. (2014), Besman et al. (2013), and Marotta and Bohnner (2013) as significant, while they considered them successful replications in their BBS target article. We think that at the very least this warrants a correction of their BBS article, as Lee and Schwarz no longer consider them successful replications.

Conflict of interests

The second author, IJzerman, is friends with Lee.