Critical review of “Regression to the mean does not explain away non-conscious processing”
by Sklar, Goldstein & Hassin (<https://psyarxiv.com/evcgy/>)

Guido Hesselmann^{1*} & Marcus Rothkirch²

¹ Department of General & Biological Psychology, Psychologische Hochschule Berlin (PHB),
10179 Berlin, Germany

² Visual Perception Laboratory, Department of Psychiatry and Psychotherapy, Charité –
Universitätsmedizin Berlin, corporate member of Freie Universität Berlin, Humboldt-
Universität zu Berlin, and Berlin Institute of Health , 10117 Berlin, Germany

* To whom correspondence should be addressed: g.hesselmann@phb.de

Keywords

Regression to the mean; non-conscious processing; subliminal priming; consciousness

Abstract

In our review, we show that Sklar, Goldstein, and Hassin’s (SGH) critical discussion of Shanks (2017) is lacking strong arguments that would allow to conclude that “regression to the mean does not explain away non-conscious processing” (title). In its current form, the manuscript by SGH (<https://psyarxiv.com/evcgy/>; DOI: 10.31234/osf.io/evcgy) is not a useful contribution to the field.

Major points

1) SGH make the point that data exclusions based on awareness are common practice in research on unconscious processing: “A common solution to this problem is to analyze the results of the group of participants for whom the stimuli remained subliminal” (abstract). And: “experimenters usually exclude from analyses participants who are aware of the stimuli, and/or trials in which participants were aware of the stimuli” (p.3). The fact that the exclusion of data is “common practice” says very little about its correctness (another example: the commonly neglected problem of multiple comparisons in multifactorial ANOVAs in psychology; Cramer et al., 2016). In fact, it has been pointed out previously why invisibility judgments cannot be interpreted in isolation (Schmidt, 2015). In a nutshell, a stimulus in a given experimental setup (e.g., in a CFS experiment) does not become “more invisible” by excluding more and more “aware” participants from the data set. As Schmidt (2015) puts it: “at the heart of this flawed procedure is a sampling fallacy, the erroneous assumption that a restriction of a sample on the basis of measured characteristics is valid on the population level”.

One could argue that Exp.6 by Sklar et al. (2012) represents an extreme case: here, 25 out of 42 participants were excluded, nearly 60%. It seems reasonable to say that this is a suboptimal procedure (also because it is a waste of resources), even though SGH present their strategy to obtain an “unaware group” as a feature, not a bug (footnote 7). However, we might indeed all be guilty of data exclusions (including us, GH and MR). We too have excluded participants (or, trials) when there were signs of awareness of a suppressed stimulus. But the fact that we are all in this mess together should not make us want to shoot the messenger (e.g., Shanks, Schmidt, or Wagenmakers), but to carefully rethink and revise our research designs and data analysis strategies. We have the impression that SGH chose to shoot the messenger in an aim to rehabilitate their “unconscious arithmetic effect” (2012), which has been criticized before on various other grounds (e.g., Moors & Hesselmann, 2017; also see Moors & Hesselmann, 2019).

2) Quoting Cohen et al. (2003), Shanks (2017) provides a short summary of regression to the mean (Rttm): “it is a mathematical necessity that whenever two variables correlate less than perfectly, cases that are at one extreme on one of the variables will, on the average, be less extreme on the other” (p.757). In the case of Exp.6 by Sklar et al., the two ingredients are there: the correlation between awareness and performance is less than 1 (it is 0.20), and a subset of “unaware” participants is analyzed in isolation (17 out of 42 participants). Hence, there will be a Rttm-induced bias in the data, and SGH do not dispute this. What then are SGH arguing against when they claim that Shanks’ conclusion is mistaken? Shanks (2017) provides a “test”

(in the following, “Rttm test”) to check whether the observed priming effect presented by Sklar et al. (2012) exceeds the Rttm-bias or not. Based on the Rttm formula, a predicted standardized performance is calculated. Shanks (2017) suggest that one can check whether this value falls inside or outside of the 95% confidence interval (CI) around the observed standardized performance. Please note that this procedure is similar to conducting a t-test (observed data versus predicted value). In the case of Exp.6, the value is -0.155 (or -0.1574 according to our calculation, see table below; our Matlab script is on OSF) and falls within the CI (-0.163; 0.525). The associated t-test is not significant ($p > .05$, see table below). Based on this result, Shanks (2017) concludes: “Thus the observed level of priming in the “unaware” subgroup is not sufficiently greater than the level predicted purely by regression to the mean to rule out the latter as the sole process in operation” (p.764).

SGH conclude (also based on their data simulations, see our next point) that the Rttm test itself is severely biased, in the sense that “datasets in which the awareness-performance relationship is linear (or non-existent) will almost always to lead to the conclusion that there is no evidence for non-conscious effect, regardless of the actual evidence” (p.8). Two points are important here. First, we can easily show that SGH’s conclusion is incorrect, at least as far as the dataset from their Exp.6 is concerned. One can see that the “Rttm test” is only marginally non-significant ($p = 0.0535$), because the predicted value almost falls outside the CI. To make the CI narrower, we simply added one more participant to the “unaware” group (so that $N = 17+1$). For the sake of simplicity, we “mirrored” one participant (i.e., we included his/her data twice, performance = 10.35 ms, awareness = 0.41; our Matlab script is on OSF). As can be seen in the table below, the Rttm test now gets significant ($p = 0.0496$), and thus indicates that the observed effect exceeds the Rttm-induced bias. In our opinion, however, there is one important caveat. Scientific arguments should not solely be based on whether the p-value – measured on a continuous scale - passes a threshold (e.g., 0.05). This thought leads us to our second point. As Shanks (2017) points out: “The regression account can be thought of as a “null” model, which assumes that awareness and performance depend on a common source. In focusing on this null model, it is important to avoid inferring that all data that are consistent with regression to the mean are uniquely supportive of this theoretical position. Put differently, finding that performance in unaware participants is approximately in line with what one would expect from regression to the mean does not automatically support the regression model over an alternative unconscious processing hypothesis” (p.771). In other words, the “Rttm test” (whatever its outcome) is not the whole story. If one can come up with a more plausible explanation for one’s subliminal priming effect than the Rttm-induced bias, then things are fine.

For example, the data from Exp.6 seem to contain one “oddity”. The total sample is – on average, as a sample – aware of the priming stimulus (N=42, awareness = 0.65), and the mean priming effect for subtraction is 11 ms. The sample of “unaware” participants (N=17, awareness = 0.50) shows a slightly larger priming effect (15 ms), while the sample of “aware” participants (N=25, awareness = 0.75) shows a slightly smaller priming effect (8 ms). How can this observation be explained? One could argue that conscious reflection interferes with priming effects, so that priming by subliminal information becomes larger than priming by supraliminal information. Alternatively (and arguably more parsimoniously), one can argue that such data pattern is exactly what is predicted by Rttm. At the end of the day, the two models compete, as mentioned by Shanks (2017). One solution to bring the debate forward – also mentioned by Shanks (2017, p.773) – could be to gather more data with a research design that avoids or minimizes Rttm, e.g., by adjusting the contrast of the priming stimulus individually for each participant. Due to the large interindividual variability in interocular suppression (Yamashiro et al., 2014), we have already followed this approach in a number of CFS studies (e.g., Rothkirch & Hesselmann, 2018).

	Mean standardized awareness	A-P-correlation	Predicted standardized performance	Observed standardized performance	CI of observed standardized performance	Significance test (observed versus predic)
Exp.6 (N=17 unaware)	-1.05	0.2019	-0.1574	0.1811	[-0.1631; 0.5253]	t(16)=2.08, p = 0.0535
Exp.6 (N=17+1 unaware)	-0.7724	0.1985	-0.1559	0.1725	[-0.1552; 0.5003]	t(17)=2.11, p = 0.0496

3) SGH present two new data simulations (data sets 2 and 3) in addition to Shanks’ original data simulation (data set 1) of Exp.6. Shanks (2017) used this simulation to show that even in the absence of a true subliminal priming effect, Rttm can lead to a false positive. Based on their simulation, SGH write: “Our simulations show that the procedure leads to the conclusion that there is no evidence for non-conscious processing even when there is very strong evidence for non-conscious processing in the data” (p.4). However, this conclusion is not justified. SGH simply added 10 ms or 1000 ms to the performance measure (P) in the data simulation (p.15, appendix). When 1000 ms are added, there is a clear priming effect, both for “aware” and

“unaware” participants (i.e., every participant has a priming effect of at least 950 ms). In this case, no one would dispute that there is subliminal priming. However, the Rttm test is not significant, how can that be, will it never be significant? Do we have to conclude that the Rttm test is severely biased? In our opinion, there are at least two potential misunderstandings. First, as mentioned previously, the Rttm test is not an ultimate tool to detect true or discard false “subliminal effects”. In data sets 2 and 3 (where there is “true” simulated subliminal priming effect), the Rttm test simply indicates that there is a considerable overestimation of the effect due to an Rttm-induced bias. Importantly, it does not indicate an absence of subliminal priming. Second, when looking at the Rttm formula, it is clear that adding 10, 1000, or even 1000000 ms to the performance scores will not make any difference, as the formula is using z-scores in a correlation. Thus, one can conclude that data sets 2 and 3 are uninformative.

To our understanding, Shanks (2017) used data set 1 to show that subliminal priming effects can seemingly appear “out of nothing” (i.e., when a “null model” is used). Rttm induces a bias and eventually pushes the “effect” above 0 ms. Shanks’ simulation uses a “null model” that mimics an existent data set (Exp.6 by Sklar et al.). As outlined above, we do not think that we can learn anything from turning this “null model” into an “alternative model” by adding a constant. The data simulation will show a significant Rttm effect when a non-linear term is added (e.g., “10.^base” to the performance score; our Matlab script is on OSF). But again, the “null model” presented by Shanks (2017) is useful because it is parsimonious, and because it shares important features of the original data set (Exp.6). Shanks (2017) does not claim that this “null model” is the only possible model underlying the observed data. If SGH can come up with an alternative model that a) is parsimonious, and b) shares important features of the original data set, and c) on which the Rttm test does not get significant, then this would be more informative.

4) In our opinion, Shanks’ article (2017) deserves as much attention as the influential “False Positive Psychology” paper by Simmons, Nelson & Simonsohn (2011). We are confident that it will stimulate better research designs and analyses. However, it should not be embraced blindly but critically be examined, just like any other methodological work (e.g., “excessive significance”, Francis et al., 2014). Important questions about the impact of Rttm in subliminal perception research remain to be addressed. For example, what is the outcome when other published data sets are analyzed using Shanks’ “Rttm test”? How does the Rttm-induced bias depend on the percentage of excluded data points? How does Rttm distort the effects when

trials, not participants, are excluded? To what degree does the “Rttm test” depend on statistical assumptions (e.g., about the distribution of the data points)?

Minor points

1) SGH write that there is a “crucial difference” between “post-hoc data selection” and “post-hoc analyses” (footnote 1). They argue that “post-hoc data selection” is determined a-priori, while “post-hoc analyses” are determined after seeing the data, and thus are data-contingent. In an attempt to rehabilitate post-hoc data selection, the authors introduce the term “a-priori selection rules”. We tend to disagree. Without preregistration, it will be impossible to say how and when an author decided to define “unawareness”. The decision to exclude participants could be based on particular performance levels (such as >55%, >60%, and so on) or on statistical tests like binomial or z-tests. Therefore, we think that replacing “post-hoc data selection” with “a-priori selection rules” is unnecessary and potentially misleading.

2) SGH write: “Thus, [Shanks’] simulation fails to replicate a crucial aspect of Sklar et al.’s design – obtaining an unaware group. In fact, if these were true results, no reasonable research would claim they suggest non-conscious processing, because the unaware group was simply aware” (footnote 7). How exactly did Sklar et al. obtain their “unaware group”? In the 2012 paper, they write: “Participants’ lack of awareness was confirmed by two strict criteria. First, they completed an objective test block, in which they repeatedly made two-alternative forced-choice judgments regarding the primes. Second, participants completed a subjective measure that consisted of direct questions about the nature of the stimuli and the tasks. We used the binomial distribution to determine whether each participant performed better than chance on the objective block and excluded from analyses all those participants who did. Also excluded were participants who reported any form of subjective awareness of the stimuli” (p.3). Using a binomial test with a fixed number of trials amounts to a fixed cutoff, e.g., 60%. There is no guarantee that this procedure results in an “unaware group”, if unawareness is defined as a non-significant difference from 50% performance. Without data-contingent decisions, it might as well have ended up 55.19% (as in the Shanks simulation), so it is clearly not a “crucial aspect of Sklar et al.’s design” (footnote 7).

3) SGH write: “In recent decades the modal view suggests that non-conscious mental processes are an inherent part of human psychology” (p.3), and then provide a long list of 13 references from very different approaches and paradigms. We believe that it would be helpful to differentiate between non-conscious mental processes as seen by social psychology, and as seen by cognitive psychology. The important differences have been pointed out elsewhere (Doyen et al., 2014).

4) SGH claim that the overestimation of unconscious processing is based on a positive correlation between awareness and performance. By pointing out previous studies in which this correlation was non-significant and apparently around zero (Sklar et al., 2012; Clark & Squire, 1998), the authors intend to show that a basic assumption of Rttm is often violated and that this aspect is neglected by Shanks (2017). However, Shanks (2017, p.753) does in fact mention that non-significant correlations between awareness and performance do not provide sufficient evidence for the absence of such correlations, especially when statistical power is low and confidence intervals are large.

5) SGH conclude that Rttm should not be a major concern if the reliability of the awareness measure is high and the correlation between awareness and performance is negligible. However, Shanks (2017) points out that Rttm is always a statistical inevitability whenever two variables are imperfectly correlated and a cutoff is applied on one dimension, “regardless of the underlying relationship between the latent variables they measure” (Shanks, p. 757). The example of a highly unreliable measure, for instance when awareness is defined on the basis of a coin toss, is only an extreme case to further illustrate the problem of post-hoc data selection, as low correlations between awareness and performance can for instance occur due to low reliability of the awareness measure. The inverse conclusion that Rttm is unproblematic if the reliability of the awareness measure is established, cannot be drawn, however. Thus, Rttm will have an impact when there is an imperfect correlation between awareness and performance, independent of the reason for this imperfect correlation.

6) SGH write: “We conclude by briefly outlining a new description of the influence of the artifact, and how it should be statistically addressed” (abstract). We are not sure where this part of the manuscript can be found.

Final note

The manuscript by SGH reads like an attempt to rehabilitate the work by Sklar et al. (2012), which was criticized in Shanks' (2017) article. This is understandable, since Shanks used Sklar et al.'s data to demonstrate the consequences of Rttm. Nevertheless, even if the data by Sklar et al. (2012) were uncontaminated by Rttm-related artifacts, this would not absolve their work from other methodological weaknesses (e.g., Moors & Hesselmann, 2017). Their manuscript should thus not come across as if it could remove all doubts related to the study (Sklar et al., 2012).

Data

Shanks' (2017) data simulation can be found at OSF (<https://osf.io/6c9wt/>). We had access to the original data of Exp.6 because one of us (GH) already published a reanalysis of the Sklar data set (Moors & Hesselmann, 2017). Our Matlab script can be found at <https://osf.io/akrdj/> (private, login required).

References

Cramer, A. O. J., van Ravenzwaaij, D., Matzke, D., Steingroever, H., Wetzels, R., Grasman, R. P. P. P., Waldorp, L. J., & Wagenmakers, E.-J. (2016). Hidden multiplicity in exploratory multiway ANOVA: Prevalence and remedies. *Psychonomic Bulletin & Review*, 23(2), 640–647. <https://doi.org/10.3758/s13423-015-0913-5>

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*. <https://doi.org/10.1177/0956797611417632>

Francis, G., Tanzman, J., & Matthews, W. J. (2014). Excess Success for Psychology Articles in the Journal Science. *PLOS ONE*, 9(12), e114255. <https://doi.org/10.1371/journal.pone.0114255>

Moors, P., & Hesselmann, G. (2017). A critical reexamination of doing arithmetic nonconsciously. *Psychonomic Bulletin & Review*, Epub ahead of print. <https://doi.org/doi:10.3758/s13423-017-1292-x>

Moors, P., & Hesselmann, G. (2019). Unconscious arithmetic: Assessing the robustness of the results reported by Karpinski, Briggs, and Yale (2018). *Consciousness and Cognition: An International Journal*, 68, 97–106. <https://doi.org/10.1016/j.concog.2019.01.003>

Hesselmann, G., & Rothkirch, M. (2018). No evidence for dorsal-stream-based priming under continuous flash suppression. *Conscious Cogn.* <https://doi.org/10.1016/j.concog.2018.05.011>

Doyen, S., Klein, O., Simons, D. J., & Cleeremans, A. (2014). On the Other Side of the Mirror: Priming in Cognitive and Social Psychology. *Social Cognition*, 32, 12–32.

Yamashiro, H., Yamamoto, H., Mano, H., Umeda, M., Higuchi, T., & Saiki, J. (2014). Activity in early visual areas predicts interindividual differences in binocular rivalry dynamics. *Journal of Neurophysiology*, 111(6), 1190–1202. <https://doi.org/10.1152/jn.00509.2013>