COMMENTS ON “Differences in common psychoacoustical tasks by sex, menstrual cycle, and race” by D. McFadden *et al.*, 2018, and methodological pitfalls in human population research

Etienne Gaudrain^{a,b}, Jaime Undurraga^c, Nicolas Grimault^a, Deniz Başkent^b

- a. Lyon Neuroscience Research Center, CNRS UMR5292, Inserm U1028, Université Lyon 1, Université de Lyon, Lyon, France
- b. University of Groningen, University Medical Center Groningen, Groningen, Netherlands
- c. Department of Linguistics, Macquarie University, Sydney, Australia

INTRODUCTION

In psychology, *replicability* of studies (Pashler and Wagenmakers, 2012) has become as important as *personalised medicine* in the clinical field (Hamburg and Collins, 2010). Replicability in psychology is thought to be hindered by low-powered statistics which is due to high inter-individual variability. Similarly, the efficiency of medical treatments varies greatly across patients, which can lead to the rejection of potential treatments that are seen as no-better than placebo at the group level, or vice-versa, which can lead to prescribing treatments that will turn out to be inadequate for a particular patient. There is thus a growing

consensus that, to address these issues, more needs to be known about the specifics of the individuals taking part in empirical studies, so that the extent of the conclusions can be appropriately limited to a more specific population. In other words, the purpose is to avoid overgeneralization of findings from a population that may not be representative of, or perhaps even misleading for, all individuals.

Researchers willing to embark on that endeavour are confronted with a difficult dilemma: how to choose which factors to include in describing or tailoring a study population? This is a problem that geneticists commonly face: there are so many phenotypic combinations that it is not possible to recruit enough individuals to just randomly probe for genes that may affect a particular outcome. If one is to also consider, in addition to genetic factors, acquired, cultural and social factors, it becomes clear that there are just too many potential factors that can influence the findings of a population study. Simply randomly picking a subset to include in a study poses a great methodological risk that has been long documented by scientists, philosophers and statisticians: the false positive rate increases as the number of tests increases. Nothing has illustrated this risk better than Tyler Vigen's website (<http://tylervigen.com/spurious-correlations>), which shows that if you correlate enough random variables with each other, you may find very strong correlation between obviously unrelated datasets, just by chance.

To reduce this risk, a number of steps need to be taken. (1) Hypotheses have to be formulated before data is collected. In other words, current knowledge must create a reasonable expectation that a specific factor may influence the studied phenomenon, and this expectation, rather than a random selection, guides our choice of factor. (2) The experiment — and especially the recruitment of participants — must be performed in a way that ensures no (or minimal) bias beyond the studied factor. In psychoacoustical experiments, for example, one needs to carefully manipulate auditory stimuli so that only a single parameter can be responsible for the observed results. This rigorous approach must also be used when considering group factors or individual factors amongst human participants. (3) When interpreting the results, researchers must be immune to their biases and expectations for the findings. They need to exert caution when evaluating whether a result is statistically significant (which is related to how likely it is for the result to be replicable; Colquhoun, 2017), and whether it is meaningful.

With this framework in mind, the study reported by McFadden, Pasanen, Maloney, Leshikar and Pho [2018, *JASA* 143(4):2338-54] fails on all of these points, at least where it concerns “race” as a factor. This study investigated the effect of “sex, menstrual cycle and race” on performance in

seven psychoacoustical tasks. Out of 35 comparisons related to the “race” factor, 2 were deemed significant by the authors. Based on these observations, the authors concluded that:

[...] racial background needs to be considered when making generalizations about human auditory performance [...]

Surprisingly, though, the same authors also admitted that the study had serious flaws in various sections: “*all outcomes relevant to race,*” they warn the reader, “*must be interpreted with caution,*” and should be seen as preliminary data that requires further investigation. Given this contradiction in the paper and for the reasons that will be further elaborated below on methodological flaws, we wonder if the data was published as such because, as the authors reported, “*this lab now is closed due to retirements, we must leave any future investigations to others.*” Yet, we have observed that “others” have at times taken the conclusions for granted without scrutinising the data, which prompted us to write this commentary.

Although we have no reason to doubt the good faith of the authors, we are concerned that these conclusions, supported by questionable evidence, may be used to justify the exclusion of ethnic groups from perceptual studies. Admittedly, as we argued above, that could be an acceptable risk if this study raised the prospect of a potential benefit, especially for the said ethnic groups. But again, the authors admit that they could not foresee any:

There exist no obvious everyday issues or problems that would be solved or better addressed by additional detailed knowledge about the mechanisms underlying the small race, sex, or menstrual-cycle differences reported here.

We are thus left with conclusions that the authors admit are not entirely supported by the data, and whose further investigation, left to others, may not raise the prospect of any particular benefit for any population. Even presuming the authors had only the best intentions, we think that there is a tangible risk that these results may be used for discrimination against population subgroups. Given that stake, the lack of benefit, and the weakness of the evidence, we recommend that this paper should be retracted.

In the sections that follow, we are reporting a careful review of the study’s flaws, most of which have already been reported by the authors themselves, who must be commended for their transparency.

PRIOR HYPOTHESES

To reduce the risk of false-positives, clear hypotheses must be formulated during study design. The factors of interest that are considered need to be justified using either prior literature, or by detailing potential mechanisms that may render the factor relevant in the task at hand.

The title of the study, “Differences in common psychoacoustical tasks by sex, menstrual cycle, and race”, mentions race as a factor, yet, in the abstract, “race” only appears in the second half, amid the results, but not in the rationale, suggesting that this factor was not initially included in the design of the study. Similarly, in the Introduction of the paper, a long literature review provides arguments about the influence of sex and menstrual cycles and concludes that there was a need to design a new study examining “psychoacoustical and physiological measures during different phases of the cycle,” but there is not mention of “race” in the initial justification of the study.

In fact, from the authors own admission, the effect of “race/ethnicity” only came as an afterthought, through casual examination of the data. This is where the risk of spurious correlation is at the highest: when factors are included retrospectively because they seem to contribute to the result. This practice creates a report bias that prevents the scientific community from evaluating how likely it is for the conclusions to be actually true.

The authors then provide a number of references to justify why race should be considered as a factor. However, none of the cited references makes a particularly strong point about race. Royster *et al.* (1980) and Agrawal *et al.* (2008) both report auditory thresholds statistics over large populations without investigating whether any other factors than race, for example, a sound exposure bias, could explain their results. Such population studies are useful from an epidemiological point of view but may not be really sufficient to justify looking at race as a factor on its own. DaCosta *et al.* (2008) focused on the possible association between eye pigmentation and susceptibility to noise-induced hearing-loss. These authors explain themselves that they did not consider race or ethnic origin as a factor, so it does not help McFadden *et al.* make a case for choosing race as a factor. The study by Lin *et al.* (2012) seems the only one with some merit in that regard: Based on epidemiological reports, they made the hypothesis that melanocytic functioning may play a role in hearing preservation through noise exposure, and stratified their data both in terms of skin pigmentation and race. The authors also report non-behavioural studies that may support the rationale for investigating the role of skin pigmentation in psychoacoustical tasks.

To their credit, the authors end their Introduction with these words of caution:

Even with this knowledge of the auditory literature, the current study was not initially designed to explore race differences; we did not recruit with race in mind. Consequently, we are not able now to make some obvious comparisons that otherwise might have been possible, and all outcomes relevant to race must be interpreted with caution.

This cautionary message is repeated in the Discussion, but somehow failed to bring some nuance to the conclusions.

STUDY DESIGN AND STUDY POPULATION

Formulating clear hypotheses allows researchers to design their experiment in a way that ensures focused interpretation of their data, excluding confounding factors as much as possible.

The “race” factor

The after-thought hypothesis about race presented in the introduction concerns skin pigmentation as a proxy for melanocytic functioning (as detailed by Lin et al., 2012). However, unlike Lin and colleagues, McFadden *et al.* did not measure skin pigmentation, which would be the main factor of interest. Instead, they used standard labels “White” and “Non-White” as a proxy for skin pigmentation. In their Introduction, they also admit:

In Sec. IV we make the important point that race categories are a poor proxy for the likely actual basis for these apparent “race” differences: individual differences in pigmentation.

Proxies can already be problematic in their own rights, but proxies of proxies add yet another level of uncertainty. The second level of proxy here, which consists in associating self-declared categorical labels with skin pigmentation, is a particularly icky one: the “Non-White” label included (amongst others) “Asian”, “American Indian”, “Black”, “More than one race”, “Other or Unknown”. It can be expected that there would be enormous variability in skin pigmentation within the “Non-White” group.

The authors give a bit more details about the makeup of the groups:

For the females in this study, approximately 45% self-identified as Asian, 8% as Black, 36% as White, and 11% as “More Than One.” For the males, approximately 27% self-identified as Asian, 9% as Black, 2% as Pacific Islander, 46% as White, and 16% as “More Than One.”

It appears, first, that the proportions of White vs. Non-White are fairly different across sexes, which makes it problematic in interpreting any interaction between race and sex. This is already an important issue because, according to the authors, this interaction is what motivated them to look to race in the first place.

Second, because the majority of participants categorized as Non-White are self-identified as Asian, and as the authors explained, because Asian is a very broad description, the link between that categorisation and skin pigmentation can be expected to be extremely weak. In other words, the White/Non-White classification is already a poor descriptor of race, and it is an even poorer descriptor of skin pigmentation.

At this point of the study, the link between the (*a posteriori*) hypothesis and the selection criteria for the population is weak to nonexistent, and results will not be able to be interpreted in relationship to the research question.

Confounds

When selecting a subgroup of the population, researchers have to be particularly vigilant to recruitment biases that may lead to confounded variables. Even when the factors defining groups are very well defined, such confounds will occur if they are not controlled for.

This is for instance the case in studies comparing Musicians to Non-Musicians. While there are robust definitions that allows researchers to clearly define the two contrasted populations, of factors such as race, education level (also that of the parents), and socio-economic status often correlate with musical proficiency (e.g. Elpus and Abril, 2011). This means that Musicians often do not differ from Non-Musicians only by their musicalness, but also by a number of other factors that could contribute to the observed differences between the groups as well. To avoid this problem, an effort must be made to match the groups across all factors that are likely to affect the results so that only musicality separates them.

The same bias exists in the current study. Because White and Non-White participants were not recruited to be matched, there is no guarantee that they do not systematically differ in non-reported criteria. For instance, education level tends to be higher amongst people reporting as Asian compared to White (see Table 1, in Williams et al., 2016). Because the test

population considered by McFadden *et al.* was composed of university students, it can reasonably be assumed that education level could translate into higher motivation in participating in research studies at their university. It could also be, for instance, that a proportion of the participants reporting as Asian may have been exposed to a tonal language during childhood, because tonal languages are more prominent in Asia, and amongst Americans of asian descent. All these factors have nothing to do with “race” *per se*, let alone with melanocytic functioning, and yet could potentially explain the observed results.

Sample size

It is also questionable whether the sample size of the study was sufficiently large. Estimating effect size for skin pigmentation from the data provided by Lin *et al.* (2012), the largest effect, found in their Table 3 as the effect of skin pigmentation on high-frequency average pure tone threshold (PTA) within the Hispanic group, we arrived at a value of Cohen’s $d=0.22^1$. If McFadden *et al.* had measured skin pigmentation, aiming for a power of 0.95, and ignoring the fact that they performed many experiments and thus need to correct for multiple comparisons, they would have had to recruit 538 participants, in *each* group. A slightly less conservative power of 0.80 would yield 326 participants (in each group). They recruited a total of 140 participants. This is a commendably large number compared to most psychoacoustical studies, but even in the best case scenario, if we assumed that all participants were included (which was not the case), they could only have expected to achieve a power of 0.25. In other words, without examining the results, *a priori* evaluation indicates that the sample size was most likely insufficient.

Furthermore, some of the confounds mentioned above could be accounted for in the statistical analyses to exclude their potential influence, provided the data was collected, and provided the sample was large enough to allow enough variability within groups. The sample size used here is already insufficient to consider a single effect, and is thus absolutely insufficient to consider potential effects from possible confounds.

These statistical weaknesses undermine how strong the conclusions are, i.e. how likely their replication is, or how likely these results reflect a *true* effect in the population.

¹ This is a best case scenario. Only the mean (-3.1 dB) and 95% confidence interval ([-5.3, -0.8]) are given. Assuming a normal distribution was used to calculate confidence interval yields $d=0.22$. Using a Student’s *t* distribution with 143 degrees of freedom (the harmonic mean of the *Ns* of each group) yields a $d=0.11$.

Voltmeter malfunction

In addition to the problems described above, the collection of experimental data was plagued by a hardware malfunction that affected calibration. The authors, again commendable for their transparency, openly disclosed that after the malfunction, the stimuli had a level 8.1 dB lower than before the malfunction. They estimated when the malfunction happened using “*numerous lines of evidence*” and separated their data on this basis. Most supervisors would probably not attempt to publish such a compromised dataset, but we cannot fault the authors for not disclosing the incident. Where full disclosure stops, though, is in the determination of *when* the malfunction occurred which was used to separate the data in two groups. This is particularly important because the effects that lead to some of the claims made in the conclusions of the paper only show up in the data *preceding* the malfunction, but not in the data following the malfunction.

Even more troubling, the (non-significant) difference in auditory thresholds between White and Non-White in male participants only appeared before the malfunction, despite the fact that the malfunction is supposed to have only resulted in a constant shift in levels.

There is also a discrepancy between the data pooled over race and the data segregated across race categories. The authors explain that this discrepancy comes from the lack of response to the “race/ethnicity” question in the recruitment form. It is however rather unfortunate that *all* these non-answers that happened in this experiment happened for the male subjects *before* the malfunction.

Despite the effort of the authors to be open about all these problems, these raise serious suspicions that the dataset may not be reliable and hampers the potential validity of the conclusions.

DATA ANALYSIS AND INTERPRETATION

The first set of tools available to researchers to interpret their data, beyond visual examination, are statistical tools. Once statistical analysis results are presented, they need to be critically examined to understand what they actually mean, and, critically, if alternative explanations may be equally likely.

McFadden *et al.* resorted to statistical methods that do not warrant any reasonable level of certainty that the reported effects are not false positives.

Statistical analyses

The experimental design used by the authors implies that there are a large number of comparisons performed with the same subjects. As the authors acknowledge:

When a large number of measurements are made on the same subjects, and multiple pairwise comparisons are made on those measurements, it is difficult to know how to protect against false-positive errors.

To solve this potential issue, they report using a *resampling* method that is described step by step. There is some possible ambiguity as to how the method was exactly implemented. Namely, it is not entirely clear whether all the data samples were pooled, across groups *and* conditions. That seems rather unlikely because the different results of the different tests are scaled differently and expressed in different units. Instead, it seems more likely that the data was resampled within the subset of interest (e.g. pooled across groups for a within-condition comparison). This interpretation also fits the resampled data set partition they described:

(2) a sample the size of one of the two groups was drawn at random, and without replacement, from the pool; (3) the remaining values were taken to represent the second group;

We used this method with simulated data to estimate how much of a protection against false-positives it offered, and found that it offered none. With 7 repetitions, i.e. 7 measurements, the chance of spurious significance was close to 35%.

Therefore, we conclude that either the statistical method is not described accurately, or the results that are deemed “significant” are probably not so.

Interpretation of the results

As reported above, assuming that the statistical analysis method they used were not as flawed as it seems, the authors are fully transparent about the weakness of the evidence underlying the “race” effect. In the Discussion, the authors repeat the word of caution of the Introduction:

Because this study was not designed initially to study race/ethnicity, we were not able to make some obvious comparisons that might have been possible had we recruited subjects with race in mind. So, many uncertainties remain.

Or further:

We remind the reader that because subjects were not recruited with race in mind, all conclusions about possible race effects must be viewed as tentative.

And even:

When it comes to the auditory system, the standard categories of race are at best poor proxies for the relevant factor(s).

Yet, the *race* factor made it to the title, and to the abstract. In the final section of the paper, despite all the warnings provided by the authors, they still insist to conclude:

The decibel differences admittedly were small, but the existence of race differences does deserve note.

CONCLUSION

While the transparency the authors demonstrated in the reporting of their results is laudable, it actually makes the best case of why the conclusions they reached are not adequately supported by evidence.

We believe that publishing preliminary or null results is extremely valuable and should be supported by journals and editors. However, whatever the data that is presented, the conclusions should always be commensurate to the effects they describe. This is a widely agreed upon standard for the quality of empirical scientific investigations.

Moreover, when the conclusions reached have the potential of being used in a discriminatory way against a group of people, we believe that it is desirable to use a stricter threshold of certainty in the interpretation of the presented evidence. Scrutiny of the data presented in McFadden *et al.*, largely provided by the authors themselves, reveals that there is not a level of certainty that is sufficient to conclude that race should be taken into account when performing psychoacoustical experiments. As far as the data indicates, at this point, it does not seem necessary to include, or exclude, specific ethnic groups in order to reduce inter-individual variability in psychoacoustical experiments.

REFERENCES

- Colquhoun, D. (2017). "The reproducibility of research and the misinterpretation of *p* -values," R. Soc. Open Sci., 4, 171085. doi:10.1098/rsos.171085
- Elpus, K., and Abril, C. R. (2011). "High School Music Ensemble Students in the United States:

- A Demographic Profile,” *J. Res. Music Educ.*, **59**, 128–145.
doi:10.1177/0022429411405207
- Hamburg, M. A., and Collins, F. S. (2010). “The Path to Personalized Medicine,” *N. Engl. J. Med.*, **363**, 301–304. doi:10.1056/NEJMp1006304
- Keita, S. O. Y., Kittles, R. A., Royal, C. D. M., Bonney, G. E., Furbert-Harris, P., Dunston, G. M., and Rotimi, C. N. (2004). “Conceptualizing human variation,” *Nat. Genet.*, **36**, S17–S20. doi:10.1038/ng1455
- Lin, F. R., Maas, P., Chien, W., Carey, J. P., Ferrucci, L., and Thorpe, R. (2012). “Association of Skin Color, Race/Ethnicity, and Hearing Loss Among Adults in the USA,” *J. Assoc. Res. Otolaryngol.*, **13**, 109–117. doi:10.1007/s10162-011-0298-8
- Pashler, H., and Wagenmakers, E. (2012). “Editors’ Introduction to the Special Section on Replicability in Psychological Science: A Crisis of Confidence?,” *Perspect. Psychol. Sci.*, **7**, 528–530. doi:10.1177/1745691612465253
- Williams, D. R., Priest, N., and Anderson, N. B. (2016). “Understanding associations among race, socioeconomic status, and health: Patterns and prospects,” *Health Psychol.*, **35**, 407–411. doi:10.1037/hea0000242