## Behavior Research Methods

## You don't need t- or F-tests to show that your groups are matched

| Journal: | Behavior Research Methods |
| ---: | :--- |
| Manuscript ID | Draft |
| Manuscript Type: | Original Manuscript |
| Date Submitted by the Author: | n/a |
| Complete List of Authors: | Politzer-Ahles, Stephen; The Hong Kong Polytechnic University, Chinese <br> and Bilingual Studies <br> Chen, Si; The Hong Kong Polytechnic University, Chinese and Bilingual <br> Studies |
|  |  |

SCHOLARONE ${ }^{m}$
Manuscripts

# You don't need $\boldsymbol{t}$ - or F-tests to show that your groups are matched 

Address correspondence to:

Stephen Politzer-Ahles

Department of Chinese and Bilingual Studies

The Hong Kong Polytechnic University

Hong Kong

E-mail: stephen.politzerahles@polyu.edu.hk

Phone: +852 27667891


#### Abstract

Researchers who want to show that two groups are matched on some control variable often attempt to do so by using statistical tests (e.g. a $t$-test or analysis of variance) in which the difference between groups is not statistically significant. This approach is useless for the purpose it is usually intended for (i.e., proving that differences in some outcome variable are attributable to theoretical interesting differences between the groups and not attributable to confounds) because inferential statistics are about population effects whereas the question of interest in this situation is about sample differences. Instead of using irrelevant inferential tests to assume that samples are matched, it is more useful to directly include confounding variables in the statistical model.


Keywords: matching, inferential statistics, between-groups designs

## Introduction

When psychological research involves comparisons of outcomes between groups (e.g., between participants or between stimuli), researchers often want to be able to show that differences between those groups are due to the manipulation of interest and not to other confounds. For instance, if a researcher measures language ability in children with and without Autism Spectrum Disorder and wants to argue that the disorder is associated with lower language ability, they generally will want to make sure that any differences they observe in language ability cannot
be attributed to other relevant variables such as age or nonverbal intelligence, which might also differ between the groups. To take another example, imagine an experiment in which participants are shown a real words and made-up words (one at a time, in random order) and instructed to indicate as quickly and accurately as possible if each word is real or made-up. If the researcher wants to demonstrate that the speed of word comprehension (as measured by response speed in the task) is different between nouns and verbs, they will want to make sure any differences in response time between nouns and verbs cannot instead be attributed to other, possibly confounding, factors, like how long the words are.

Researchers often do this by producing "matched" samples, e.g., finding children with and without Autism Spectrum disorder such that the two groups have similar age and nonverbal intelligence, or selecting nouns and verbs such that the two sets of words are similar in length. Furthermore, researchers often attempt to demonstrate that these samples are well matched by showing non-significant inferential statistical tests comparing the groups on relevant variables. For example, if we were carrying out the abovementioned word recognition study using 36 nouns and 36 verbs, we might make a statement such as the following:

The mean length in letters was $8.3(S D=2.1)$ for nouns and $7.7(S D=2.3)$ for verbs; this difference was not significant (Welch $\mathrm{t}(68.6)=0.90$, $\mathrm{p}=.372$ ).
and then go on to be confident that any reaction time differences we observe between nouns and verbs cannot be due to the words' length.

This use of inferential statistics is common in much research in experimental psychology. For instance, we reviewed the latest volume (2017) of Language, Cognition and Neuroscience, a journal in psycholinguistics and neurolinguistics, and found that out of 63 regular articles presenting empirical experiments, $48 \%$ had at least one instance of using an inferential statistical test to conclude that samples of materials or participants were matched or were different (Supplementary File 1). This is likely an underestimate of how often this is done in between-groups designs (since some of the papers that did not include such statistics were reporting experiments with fully within-stimuli and within-participants designs, which would not necessitate matching anyway). While this analysis only involved one journal, the pattern is consistent with our experience with papers many other journals in psycholinguistics, phonetics, and experimental psychology. In short, inferential statistical tests are widely used by researchers wanting to show that their samples are or are not matched on some critical dimension.

Such inferential tests, however, are useless for demonstrating what the researcher is trying to demonstrate. Below we explain why these tests are useless, and offer some suggestions for better approaches to this problem.

## Why inferential statistics don't show that samples are matched

These tests provide inferences about population effects, when the researcher is interested in samples. Inferential statistics are meant to evaluate hypotheses
about populations (Hoekstra et al., 2006). ${ }^{1}$ For example, the $t$-test comparing noun and verb samples tests the null hypothesis that the 36 nouns and 36 verbs came from a population in which nouns and verbs have the same length; the $p$-value of .372 indicates that if they were indeed randomly sampled from such a population, there is a $37.2 \%$ chance that the observed length difference between nouns and verbs could have been obtained in a random sample. The properties of the population, however, are irrelevant to the research question at hand. ${ }^{2}$ In this example, the researcher wants to know if the difference in the sample is enough to cause meaningful difference in reaction times; a statistical test providing inferences about the population does not answer this question.

This issue alone is enough to reject the use of inferential tests for demonstrating that groups are matched. Nevertheless, there is another problem with this approach as well: the even if a difference between groups is not

[^0]statistically significant, it may still be enough to substantially influence the results of interest.

Statistical non-significance does not entail practical non-importance. Most researchers are familiar with the admonition that "statistical significance does not entail practical importance" (Wasserstein \& Lazar, 2016). Likewise, statistical nonsignificance does not entail practical unimportance. Even if the difference in length between a sample of nouns and a sample of verbs is non-significant, it may still be enough to trigger significant differences in some dependent variable that it is associated with, like reaction times. Supplementary File 2 includes several simulated datasets which all have this property: each dataset has two samples (verbs and nouns) with a non-significant difference in length and a significant or marginal difference in reaction speed. The common approach described aboveshowing that the difference in length is non-significant and then attributing the difference in reaction time to the manipulation of interest-would lead the researcher to conclude that the difference in reaction time must be due to something other than length (e.g., to an underlying difference between verbs and nouns). In fact, the data were simulated such that the differences in reaction time would be attributable to length. Specifically, data with an $r=.35$ correlation between length and reaction time were simulated, and then grouped arbitrarily into two samples such that nouns had a slightly, but not significantly, higher mean length and a significantly higher mean reaction latency. Regressing reaction time on both part of speech (noun vs. verb) and length, as shown in the right-hand side of Figure 1 , reveals that the reaction time difference between nouns and verbs is not significant once length is taken into account. In this case, rather than concluding
that nouns and verbs are intrinsically associated with different reaction speeds, the researcher might instead conclude that apparent differences between nouns and verbs in terms of reaction speed might just be an epiphenomenon of differences in length (i.e., that nouns and verbs may both fall along the same length - reaction time regression line).


Figure 1. Simulated noun/verb reaction time data. Each point represents one word (red for nouns, blue for verbs). At the top of each panel is shown a t-test comparing the $y$-axis variable between nouns and verbs (i.e., the t-test on the left panel tests whether nouns and verbs have the same length, and the tests on the center and right panels test whether nouns and verbs have the same reaction time).

The left panel shows that the two samples do not significantly differ in length; the long black lines show the mean, and the error bars show difference-adjusted 95\% confidence intervals. The center panel shows that the two samples do marginally differ in reaction time. The rightmost panel shows a regression of reaction time on length and part of speech; shaded regions represent the confidence interval of each line. In this panel, the difference between nouns and verbs is no longer significant. While the first two panels suggest that the two samples differ in reaction time, the
regression suggests that the apparent differences between the samples in reaction time are probably due to the confounding difference between samples in length.

## Better ways to deal with control variables

If providing a non-significant $t$-test or analysis of variance does not help prove that the experiment results are not due to confounds, what else can a researcher do? Surely researchers cannot simply ignore confounds entirely. Nonetheless, a non-significant inferential test does nothing to prove that the confound has been taken care of, and researchers should not mistake such a test as a license to stop worrying about potential confound. Simply describing the data distribution for each sample provides more relevant information than the inferential statistics without descriptions do. E.g., in the example above, if the $t$-test were removed and the means and standard deviation left, the reader would not lose anything that helps them understand the results. Thus, simply providing the descriptive statistics (or, better yet, the full list of stimuli/participants and relevant variables for each; Morey et al., 2016; Vazire, 2017) and letting the reader judge for themselves whether or not there is a problem, is more useful than simply providing an inferential test. Furthermore, it avoids giving anybody the false impression that a confound has been eliminated.

If descriptive statistics are not enough to satisfy a researcher or their readers that the results of a study are not due to some confounding variable, what more can be done? In such a case, instead of resorting to an irrelevant inferential
statistic, researchers can instead include the confounding variables in their statistical model, or researchers with directional hypotheses can select samples where confounds cannot account for results of interest.

Including control variables in the statistical model. The best option is to explicitly account for relevant variables in the statistical model. A regression model can take into account many variables at once and try to identify whether one variable of interest has a unique association with the outcome variable when other variables are accounted for (Keith, 2006; Westfall \& Yarkoni, 2016). For example, the researcher in the word recognition experiment could construct one model predicting reaction times as a function of length and word category (as well as relevant random effects; Baayen et al., 2008, Judd \& Westfall, 2012), and another model predicting reaction times as a function of length only. The researcher would then compare the two models' goodness of fit to test whether the inclusion of word category significantly improves how well the model fits the data (Keith, 2006). If it does, the researcher can infer that noun/verb status has a significant association with reaction times over and above the confounded effect of length. On the other hand, if the inclusion of this factor does not significantly improve the model fit, the researcher does not have evidence to make such a conclusion; this is the case for all datasets in Supplementary File 2. Indeed, taking multiple and potentially confounded variables into account at once is one of the main purposes of regression models (Christenfield et al., 2004; Keith, 2006). An additional benefit of this practice is that it obviates the need to form matched samples, which can be quite difficult when the number of relevant control variables is high (Balota et al., 2004; Pourhoseingholi et al., 2012).

This solution is not perfect, because "regressing out" some control variable like is not the same as regressing out the actual confounding construct that is really influencing reaction times (Westfall \& Yarkoni, 2016). ${ }^{3}$ For example, length in letters is only a rough measure of how "big" a word actually is on the page or computer screen; if words are written in a font that is not monospace, words with the same length in letters might take up more space when typed out on the page. While this example is, by design, very simple, in actual research many constructs we are interested in, like personality, working memory capacity, language proficiency, how frequently a word is used, etc., are unobservable and can only be estimated using rough measures. While this poses a challenge for researchers trying to use regression models to control for confounding variables (Westfall \& Yarkoni, 2016), the exact same shortcoming applies to studies trying to create samples of stimuli that are matched on length or some other confounding variable: the samples can only be matched on a rough measure, not on the true construct, which is un-observable.

Another limitation of this approach is that some models may not be able to account well for the effect of the confounding variable. For example, if our sample of words has very little variation in length, it may be difficult to model this variable well; likewise, it may be difficult to fit a model to the length effect if a researcher originally dichotomized the variable by choosing samples of long and short words, thus ending up with a sample in which length is bimodally distributed. Furthermore, some such models may be difficult to implement. For example, if a manipulation is both between participants and between stimuli, and both of these random factors
has potentially confounding variables, then a mixed model (Baayen et al., 2008) would be needed to account for all of these variables together, and this may be difficult to apply to situations which require some sort of aggregation over participants or stimuli first (such as cluster-based permutation tests used on eventrelated potential data; Maris \& Oostenveld, 2007). In cases like this, putting the confounding variable into the model may not solve the problems, and the researcher might have no choice but to only offer descriptive statistics, as mentioned above. This would still be better than offering a non-significant inferential statistic that gives the researcher false confidence that the groups are well matched.

Taking advantage of directional hypotheses. Another option researchers can use if they are unable to use a regression model is to simply make sure that any difference between groups is not confounded with their predicted result. In the word recognition example, imagine that the researcher knows that less-frequentlyused words are usually recognized more slowly, and the researcher's hypothesis is that nouns will be recognized more slowly than verbs. If the researcher intentionally samples stimuli such that the nouns are more frequently used than verbs, ${ }^{4}$ and observes slower reaction times for nouns in the experiment, then the reaction time difference cannot be attributed to length. This approach, however, requires having a directional research hypothesis, and requires great confidence in the direction of the possibly confounding effect (e.g., that common words are responded to more

[^1]quickly, not more slowly). Therefore, it is not applicable in situations where either the outcome of interest in the experiment, or the impact of the confounding variable, might be expected to go in either direction.

## Conclusion

While researchers are right to worry about confounding variables, the common approach of using non-significant $t$ - or F-tests to try to show that groups are "matched" on these variables does not resolve this problem, as it provides no information germane to the question of how these confounds are influencing the effects of interest. Instead of trying to use such tests to claim that groups are matched, it is more productive to directly include these confounding variables in the statistical model.

## References

Austin, P. (2011). An introduction to propensity score methods for reducing the effects of confounding in observational studies. Multivariate Behavioral Research, 46, 399-424.

Baayen, H., Davidson, D., \& Bates, D. (2008). Mixed-effects modeling with crossed random effects for subjects and items. Journal of Memory and Language, 59, 390412.

Balota, D., Cortese, M., Sergent-Marshall, S., Spieler, D., \& Yap, M. (2004). Visual word recognition for single-syllable words. Journal of Experimental Psychology: General, 133, 283-316.

Christenfeld, N., Sloan, R., Carroll, D., \& Greenland, S. (2004). Risk factors, confounding, and the illusion of statistical control. Psychosomatic Medicine, 66, 868-875.

Hoekstra, R., Kiers, H., Johnson, A., Groenier, M. (2006) Problems when interpreting research results using only p-value and sample size. In: ICOTS-7 International Conference on Teaching Statistics: Working Cooperatively in Statistics Education.

Judd, C., Westfall, J., \& Kenny, D. (2012). Treating stimuli as a random factor in social psychology: a new and comprehensive solution to a pervasive but largely ignored problem. Journal of Personality and Social Psychology, 103, 54-69.

Keith, T. (2006). Multiple regression and beyond. Boston: Pearson.

Maris, E., \& Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. Journal of Neuroscience Methods, 164, 177-190.

Morey, R., Chambers, C., Etchells, P, Harris, C., Hoekstra, R., Lakens, D., ..., \& Zwaan, R. (2016). The Peer Reviewers' Openness Initiative: incentivising open research practices through peer review. Royal Society Open Science, 3, 150547.

Pourhoseingholi, M., Baghestani, A., \& Vahedi, M. (2012). How to control confounding effects by statistical analysis. Gastroenterology and Hepatology: From Bed to Bench, 5, 79-83.

Vazire, S. (2017). Quality uncertainty erodes trust in science. Collabra, 3, 1. Wasserstein, R., \& Lazar, N. (2016). The ASA's statement on p-values: context, process, and purpose. The American Statistician, 70, 129-133.

Westfall, J., \& Yarkoni, T. (2016). Statistically controlling for confounding constructs is harder than you think. PLoS ONE, 11, e0152719.

Wiener, S., \& Ito, K. (2015). Do syllable-specific tonal probabilities guide lexical access? Evidence from Mandarin, Shanghai and Cantonese speakers. Language, Cognition, and Neuroscience, 30, 1048-1060.

| 2 | Paper | Journal | JournalVolume | Year |
| :---: | :---: | :---: | :---: | :---: |
| 3 | Brown-Schmidt and Toscano | LCN | LCN32_10 | 2017 |
| 4 | Henry et al. | LCN | LCN32_10 | 2017 |
| 5 | Middleton et al. | LCN | LCN32_10 | 2017 |
| 7 | Martín-Loeches et al. | LCN | LCN32_10 | 2017 |
| 8 | Dufour et al. | LCN | LCN32_10 | 2017 |
| 9 | Kapnoula et al. | LCN | LCN32_10 | 2017 |
| 10 | Wong et al. | LCN | LCN32_10 | 2017 |
| 11 | Arantzeta et al. | LCN | LCN32_10 | 2017 |
| 12 | Grieco-Calub | LCN | LCN32_10 | 2017 |
| 14 | Caffarra et al. | LCN | LCN32_9 | 2017 |
| 15 | Scaltritti et al. | LCN | LCN32_9 | 2017 |
| 16 | Shahin et al. | LCN | LCN32_9 | 2017 |
| 17 | Croot et al. | LCN | LCN32_9 | 2017 |
| 18 19 | Mitterer and Reinisch | LCN | LCN32_9 | 2017 |
| 20 | Yue et al. | LCN | LCN32_9 | 2017 |
| 21 | Yu et al. | LCN | LCN32_9 | 2017 |
| 22 | Lawyer and Corina | LCN | LCN32_9 | 2017 |
| 3 | Frank and Willems | LCN | LCN32_9 | 2017 |
| 24 | Bellanger et al. | LCN | LCN32_9 | 2017 |
| 26 | Ito et al. | LCN | LCN32_8 | 2017 |
| 27 | van Rijswijk et al. | LCN | LCN32_8 | 2017 |
| 28 | Calloway and Perfetti | LCN | LCN32_8 | 2017 |
| 29 | Turnbull et al. | LCN | LCN32_8 | 2017 |
| 30 | Fieder et al. | LCN | LCN32_8 | 2017 |
| 32 | Gu et al. | LCN | LCN32_8 | 2017 |
| 33 | Liu et al. | LCN | LCN32_6 | 2017 |
| 34 | Bürki | LCN | LCN32_6 | 2017 |
| 35 | Gould et al. | LCN | LCN32_6 | 2017 |
| 36 37 | Vivas et al. | LCN | LCN32_6 | 2017 |
| 38 | Bühler et al. | LCN | LCN32_6 | 2017 |
| 39 | Zellou et al. | LCN | LCN32_6 | 2017 |
| 40 | Akhavan et al. | LCN | LCN32_6 | 2017 |
| 41 | Ghitza | LCN | LCN32_5 | 2017 |
| 42 | Rommers et al. | LCN | LCN32_5 | 2017 |
| 44 | Lam et al. | LCN | LCN32_5 | 2017 |
| 45 | Lewis et al. | LCN | LCN32_5 | 2017 |
| 46 | Tsang et al. | LCN | LCN32_5 | 2017 |
| 47 | Fang et al. | LCN | LCN32_5 | 2017 |
| 49 | Monaco et al. | LCN | LCN32_5 | 2017 |
| 50 | Babineau et al. | LCN | LCN32_4 | 2017 |
| 51 | Reifegerste et al. | LCN | LCN32_4 | 2017 |
| 52 53 | Declerck and Philipp | LCN | LCN32_4 | 2017 |
| 54 | Olson | LCN | LCN32_4 | 2017 |
| 55 | Smolka and Libben | LCN | LCN32_4 | 2017 |
| 56 | Riesenhuber and Glezer | LCN | LCN32_3 | 2017 |
| 57 | Gubian et al. | LCN | LCN32_3 | 2017 |
| 5 | Vankov and Bowers | LCN | LCN32_3 | 2017 |

Brehm and Bock
Schiller et al.
Weisberg et al.
Ivanova et al.
Borovsky
Pycha
Zhuang and Devereux
Bishop
Miwa et al.
Achim et al.
Chen and Yeh
Rose and Rahman
Tsang et al.
Shantz and Tanner
Wu et al.
Hwang

| LCN | LCN32_2 | 2017 |
| :--- | :--- | :--- |
| LCN | LCN32_2 | 2017 |
| LCN | LCN32_2 | 2017 |
| LCN | LCN32_2 | 2017 |
| LCN | LCN32_2 | 2017 |
| LCN | LCN32_2 | 2017 |
| LCN | LCN32_2 | 2017 |
| LCN | LCN32_2 | 2017 |
| LCN | LCN32_1 | 2017 |
| LCN | LCN32_1 | 2017 |
| LCN | LCN32_1 | 2017 |
| LCN | LCN32_1 | 2017 |
| LCN | LCN32_1 | 2017 |
| LCN | LCN32_1 | 2017 |
| LCN | LCN32_1 | 2017 |
| LCN | LCN32_1 | 2017 |

## has inferential statistic for making conclusions about samples?



| 2 | Y | x | cond |
| :---: | :---: | :---: | :---: |
| 3 | 0.504044 | 0.115988 | A |
| 4 | 0.597747 | 0.183254 | A |
| 6 | 0.605813 | -0.03392 | A |
| 7 | 0.721443 | -0.1191 | A |
| 8 | 0.738369 | -0.14843 | B |
| 9 | 0.749678 | 0.225042 | A |
| 11 | 0.770495 | -0.03131 | B |
| 12 | 0.781449 | -0.02786 | A |
| 13 | 0.795879 | -0.13236 | B |
| 14 | 0.801702 | 0.027689 | A |
| 15 | 0.802044 | 0.066873 | B |
| 17 | 0.905733 | -0.07725 | A |
| 18 | 0.926484 | 0.029461 | B |
| 19 | 0.929668 | 0.10807 | A |
| 21 | 0.957128 | -0.07128 | B |
| 22 | 0.963321 | -0.07461 | A |
| 23 | 0.99234 | -0.09293 | B |
| 24 | 0.999539 | 0.204705 | A |
| 25 | 1.018425 | -0.01128 | B |
|  |  |  |  |
| 27 | 1.022739 | -0.03652 | A |
| 28 | 1.023101 | -0.01877 | B |
| 29 | 1.026082 | -0.11643 | A |
| 30 | 1.039974 | -0.15389 | B |
| 32 | 1.070918 | -0.19812 | A |
| 33 | 1.090003 | -0.00791 | B |
| 34 | 1.13825 | 0.071231 | A |
| 35 | 1.143477 | -0.09521 | B |
| 37 | 1.185487 | 0.006076 | A |
| 38 | 1.186985 | -0.00045 | B |
| 39 | 1.256148 | -0.0502 | A |
| 40 | 1.258514 | -0.02055 | B |
| 42 | 1.28407 | 0.003 | A |
| 43 | 1.286365 | 0.090335 | B |
| 44 | 1.30271 | 0.011283 | A |
| 45 | 1.309562 | 0.096223 | B |
| 47 | 1.330447 | -0.14638 | A |
| 48 | 1.417858 | 0.017078 | B |
| 49 | 1.472239 | -0.26579 | A |
| 50 | 1.525241 | -0.02916 | B |
| 52 | 1.574016 | -0.14771 | A |
| 53 | 1.651307 | 0.160608 | B |
| 54 | 1.669936 | -0.08774 | A |
| 55 | 1.727708 | 0.033348 | B |
| 57 | 1.76846 | 0.078238 | A |
| 58 | 1.855449 | -0.0568 |  |
| 59 | 1.899281 | -0.05975 | A |

```
1.933937 -0.00858 B
1.945686 0.087497 A
1.977582 0.04809 B
2.036047 0.159971 B
2.153764 0.076605 В
2.170337 0.075684 B
```

| 2 | Y | $x$ | cond |
| :---: | :---: | :---: | :---: |
| 3 | 0.423114 | 0.032769 | A |
| 4 | 0.540773 | -0.1018 |  |
| 6 | 0.603295 | 0.101789 | A |
| 7 | 0.66598 | 0.111585 |  |
| 8 | 0.669595 | -0.16077 | B |
| 9 | 0.673531 | 0.031569 | A |
| 11 | 0.698546 | -0.12798 | B |
| 12 | 0.778257 | -0.13825 | A |
| 13 | 0.790263 | 0.176789 | B |
| 14 | 0.79949 | 0.152493 | A |
| 16 | 0.814503 | 0.007965 | B |
| 17 | 0.829434 | 0.007874 | A |
| 18 | 0.838239 | 0.121257 | B |
| 19 | 0.841395 | -0.1127 | A |
| 21 | 0.866606 | 0.006526 | B |
| 22 | 0.938719 | -0.12751 | A |
| 23 | 0.953075 | 0.037643 | B |
| 24 | 0.960099 | -0.17012 | A |
| 25 | 0.985214 | 0.023912 | B |
| 27 | 1.02132 | 0.026639 | A |
| 28 | 1.048531 | -0.05399 | B |
| 29 | 1.081667 | 0.186066 | A |
| 30 | 1.101311 | -0.13764 | B |
| 32 | 1.112555 | -0.02012 | A |
| 33 | 1.148615 | 0.080326 | B |
| 34 | 1.153951 | -0.12871 | A |
| 35 36 | 1.171355 | 0.08466 B | B |
| 37 | 1.197304 | -0.00218 | A |
| 38 | 1.239122 | -0.14545 | B |
| 39 | 1.241997 | 0.154367 | A |
| 40 | 1.248006 | -0.00173 | B |
| 42 | 1.253114 | -0.08527 | A |
| 43 | 1.264995 | -0.01879 | B |
| 44 | 1.27206 | -0.0637 | A |
| 45 | 1.303732 | 0.258468 | B |
| 47 | 1.314489 | -0.11428 | A |
| 48 | 1.351257 | 0.137476 | B |
| 49 | 1.369353 | 0.112902 | A |
| 50 | 1.376125 | -0.09039 | B |
| 52 | 1.395906 | -0.0368 | A |
| 53 | 1.416968 | 0.126259 | B |
| 54 | 1.435761 | 0.068511 | A |
| 55 | 1.460717 | -0.02423 | B |
| 57 | 1.483648 | -0.02606 | A |
| 58 | 1.500215 | 0.182193 | B |
| 59 | 1.503537 | 0.02982 | A |


| 1.58792 | 0.031239 B |
| ---: | ---: |
| 1.59541 | -0.15495 A |
| 1.601028 | 0.027073 B |
| 1.618038 | 0.284162 B |
| 1.706993 | -0.01348 B |
| 1.732961 | 0.042319 B |


| 1 |  |  |
| :--- | ---: | ---: |
| 2 | Y Cond |  |
| 3 | 0.164284 | -0.18472 A |
| 4 | 0.286197 | 0.006731 A |
| 5 | 0.325619 | -0.13995 A |
| 6 | 0.552318 | 0.03579 A |
| 7 | 0.554854 | 0.060489 B |
| 8 | 0.566297 | -0.01083 A |
| 9 | 0.598027 | 0.257287 B |
| 10 | 0.61398 | -0.19421 A |
| 11 | 0.65203 | -0.04182 B |
| 12 | 0.719211 | -0.00285 A |
| 13 | 0.736169 | 0.070882 B |
| 14 | 0.750403 | 0.120045 A |
| 15 | 0.769202 | 0.100609 B |
| 16 | 0.798556 | -0.03513 A |
| 17 | 0.809214 | -0.19671 B |
| 18 | 0.849928 | -0.04712 A |
| 19 | 0.887225 | 0.007683 B |
| 20 | 0.899017 | -0.05598 A |
| 21 | 0.981483 | -0.07275 B |
| 22 | 1.000652 | -0.00683 A |
| 23 | 1.04609 | 0.143729 B |
| 24 | 1.108824 | -0.14558 A |
| 25 | 1.125854 | -0.15454 B |
| 26 | 1.618619 | 0.13547 A |
| 27 |  |  |

Page 25 of 25


[^0]:    ${ }^{1}$ Inferential statistics are most often based on $p$-values, or on confidence intervals (which are directly mathematically related to $p$-values). A $p$-value tests the compatibility between a sample statistic and a statistical model (Wasserstein \& Lazar, 2016). In practice, this model is usually a hypothetical population in which the relevant statistic is zero.
    ${ }^{2}$ A potential exception to this point is that properties of a sample can be relative to a population. For example, imagine that a researcher finds a sample of nouns and a sample of verbs such that the two samples have the exact same average length, but in the language at large verbs tend to be shorter. In this case, the nouns and verbs have the same absolute length but different relative lengths: the verbs are long for verbs and the nouns are short for nouns. It necessarily follows that, if two groups of items or subjects are matched on some absolute measure, they are not matched on the relative measure. There are indeed contexts where the relative measures may be more important (e.g., frequencies of lexical tones relative to a given syllable: Wiener \& Ito, 2015). In any case, however, the kinds of statistical tests described here do not solve this problem, and indeed attempts to match samples mask the problem.

    Another case in which inferential tests may be relevant is for judging the effectiveness of a matching procedure, as opposed to the matched-ness of a sample itself. For example, propensity score matching is often used in observational studies to generate matched samples out of a larger sample (Austin, 2011); inferential statistics can be used to test (and hopefully fail to reject) the hypothesis that the procedure itself generates samples that are the same on some characteristic. The issue remains, however, that this would not be an inference about the sample itself.

[^1]:    ${ }^{4}$ Or used with the exact same frequency as verbs, but this is usually impossible. For control variables that are continuous, no two observations can ever be exactly the same. For instance, if two people were born at the "same" time down to the second, they still do not have the exact same age; someone using a more fine-grained clock (e.g. with nanosecondlevel resolution) would eventually be able to detect a small difference in time of birth when measuring with sufficient detail.

