Why most experiments in psychology failed: sample sizes required for randomization to generate equivalent groups as a partial solution to the replication crisis

Gjalt-Jorn Ygram Peters^{1, 2*}

¹ Department of Methodology & Statistics, Faculty of Psychology & Education Science, Open University of the Netherlands, Heerlen, the Netherlands

² Department of Work & Social Psychology, Faculty of Psychology & Neuroscience, Maastricht University, Maastricht, the Netherlands

* Corresponding author: gjalt-jorn@behaviorchange.eu

Acknowledgements: I would like to thank Stefan Gruijters for his invaluable contributions to the conception and development of the ideas in this paper.

Abstract

The experimental method is one of the staple methodological tools of the scientific method, and is as such prevalent in the psychological literature. It relies on the process of randomization to create equivalent groups. However, this procedure requires sufficiently large samples to succeed. In the current paper, we introduce tools that are based on the sampling distribution of Cohen's d and that enable computing the likelihood that randomization succeeded in creating equivalent groups and the required sample size to achieve a desired likelihood of randomization success. The required sample sizes are considerable, and to illustrate this, we compute the likelihood of randomization failure using data from the Reproducibility Project: Psychology. It is shown that it is likely that many original studies but also many replications failed to successfully create equivalent groups. For the replications, the mean likelihood of randomization failure was 44.54% (with a 95% confidence interval of [35.03%; 54.05%]) in the most liberal scenario, and 100% in the most conservative scenario. This means that many studies were in fact not experiments: the observed effects were at best conditional upon the value of unknown confounders, and at worst biased. In any case replication is unlikely when the randomization procedure failed to generate equivalent groups in either the original study or the replication. The consequence is that researchers in psychology, but also the funders of research in psychology, will have to get used to conducting considerably larger studies if they are to build a strong evidence base.

Health psychology, specifically behaviour change science, is the study of human behaviour and its causal antecedents with the goal of predicting and changing behaviour. As bachelor-level methodology and statistics courses dutifully teach psychology students, ‘correlation does not imply causation’, and studying causality requires experimental designs¹. Consequently, experimental designs are central to a considerable number of studies ranging from randomized controlled trials evaluating behaviour change interventions (de Bruin et al., 2010) and fundamental studies into defensive processing of threatening stimuli (Kessels, Ruiter, Brug, & Jansma, 2011) to classic studies into processes underlying behaviour change methods such as modelling (Bandura, Ross, & Ross, 1963). Experimental designs enable conclusions regarding causality because the procedure of randomization enables creating equivalent groups of participants, which can then be differentially exposed to stimuli (‘manipulated’). If subsequent measurement of psychological variables reveals differences between the groups, logic dictates that this difference must be caused by the manipulation. This allows the conclusion that the manipulated variable is a causal antecedent of the variable(s) measured after that manipulation.

This logic holds only under the assumption that the randomization procedure rendered the groups equivalent before manipulation took place. In this contribution, we will show that randomization frequently fails to produce equivalent groups, so that what researchers think are experiments are in fact quasi-experiments. We will also show how this phenomenon may partly explain the so-called ‘replication crisis’, and finally, we will provide guidelines to try to prevent this in future research.

¹ But note that when practical considerations render such designs unattainable, such as in the case of studying the causal relationship between smoking and lung cancer, causality can still be established through synthesis of a vast amount of multidisciplinary theoretical and empirical evidence. This route, however, is substantially more expensive both in terms of funds and time.

Conclusions as to a causal relationship between two variables require establishing three empirical conditions. First, the causal antecedent must temporally precede the causal consequent. Second, the antecedent and consequent must covary: in other words, when the antecedent occurs, the consequent must occur. Not necessarily all of the time, but with some degree of regularity (after all, the association may be moderated). Third, alternative explanations for this covariance must be excluded. For example, the observation that days where many ice-creams are sold are often followed by days where many people drown satisfied the first two conditions, but fails to satisfy the last one (in this case because nice weather causes both phenomena). The first two things could be established by measuring both variables, in the correct order, a sufficient number of times to establish their association: a longitudinal observational design would suffice. However, excluding alternative explanations is impossible in an observational design: even if all known potential confounders are measured, it is logically impossible to know for certain that no unknown confounders exist.

The experimental method affords the right to draw conclusions about causality because randomizing people into groups rules out most alternative explanations for covariance between the independent variable, the levels of which manifest as groups, and the dependent variable. Because an adequate randomization procedure produces equivalent groups, any covariance that is observed between the manipulated independent variable and the observed dependent variable must necessarily be a consequence of a causal effect of the independent variable on the dependent variable, assuming the described conditions required for causality such as temporal order of cause and effect². Therefore, a sufficiently powered

² Assuming the researcher assumes that time is linear, which is a widely accepted assumption.

study yields an effect size and confidence interval informative regarding the strength of the causal association between the independent and dependent variable. Such effect size estimates can then be meta-analysed to arrive at sufficiently reliable estimates to inform practice and policy.

However, if the groups are *not* equivalent before the manipulation of the independent variable takes place, the experiment no longer allows conclusions as to causality. In such a case, the observed association between the independent variable and the dependent variable need not be caused by the independent variable. If the groups are not equivalent, the design is no longer an experimental design, but rather a quasi-experiment, where one or more unknown variables confound the independent variable, rendering its manipulation an invalid operationalisation. In other words, the different groups no longer solely reflect variation in the independent variable: they also reflect variation in the confounders. This implies that those confounders can therefore be the causes of any observed variation in the dependent variable(s), which makes it impossible to draw conclusions as to the causal role of the manipulated variable.

The most obvious example of why this is a problem occurs when the groups already differ on the dependent variable at baseline measurement: in this scenario, even if the manipulation of the independent variable has no effect whatsoever on the dependent variable, an association between the independent and the dependent variables is still observed. Of course, because the dependent variable is known in advance, it can be measured at baseline, and ANCOVA can be used to correct for this difference (Van Breukelen, 2006; but note that this is more complicated than it might seem, see e.g. Elwert & Winship, 2014, and Westfall & Yarkoni, 2016).

In situations where the groups are equivalent with respect to the dependent variable, or where ANCOVA is used to correct non-equivalence, the supposed experiment can still in fact be a quasi-experiment if the groups differ on one or more so-called 'nuisance variables': unknown confounder(s). If inadequate randomization results in groups that differ on nuisance variables, there is no remedy, because being unknown to the researchers, those confounders cannot be measured and corrected for. This is exactly why only experimental designs with adequate randomization procedures (i.e. procedures resulting in equivalent groups) permit conclusions as to causality (but see footnote 1).

If the groups differ on one or more nuisance variables (unknown confounders), an association between the independent and dependent variable(s) can have several causes other than the straightforward explanation that the independent variable causally predicts the dependent variable. Even if that causal association exists, the observed effect size can over- or underestimate the true effect because of an interaction with one or more confounders. It is also possible that the confounders that differ between the groups cause the dependent variable to change differentially over time, resulting in a difference at post-manipulation measurement that is independent of the manipulation itself. There are more alternative explanations; and exactly because unknown confounders and their effects are unknown, it seems futile to try and compile such a list. The bottom line is that in experiments, group equivalence is important; so important that non-equivalent groups are a serious problem (Krause & Howard, 2003).

From this perspective, it makes sense that it is common to examine baseline differences to verify this equivalence. Common, but wrong. An earlier draft of a recent paper that aimed to explain this in an accessible fashion (Gruijters, 2016) coined the term

'randomization paranoia' to describe researchers' preoccupation with verifying baseline group equivalence. The term 'paranoia' implies somewhat of an over-reaction, but given the crucial role of group equivalence, is some preoccupation not justified? Or, a more useful question: when should researchers worry that their groups may not be equivalent? Of course, when twenty people are randomized into two groups, the probability that these groups differ on at least one variable is considerable, probably even close to a hundred percent if enough variables are considered. But which sample size is required to be reasonably sure that both groups *are* equivalent?

To our knowledge, there have been few attempts to provide guidance to researchers who are planning their studies and desire a high likelihood of equivalent groups. Most work regarding randomization has been theoretical, and characterized by quite extreme positions. It has been argued that there is no cause to randomize (Worrall, 2007) and even that randomization can do more harm than good (Saint-Mont, 2015), although the value of randomization as methodological instrument has also been vigorously defended (Papineau, 1994; Senn, 2013). To our knowledge, only two papers examined the sample sizes required for randomization to successfully produce equivalent groups (Hsu, 1989; Nguyen et al., 2017). Both were from the medical literature, but their conclusions were dramatically different. Hsu (1989) optimistically concluded that "when samples are large (40 per group), randomization and random sampling appear to be effective methods of creating groups that are equivalent on the maximum number of nuisance variables examined in this article. Even when samples are of moderate size (20-40 per group), randomization and random sampling appear to work well provided the number of nuisance variables is small" (Hsu, 1989, pp. 136-137). However, Hsu (1989) only considered dichotomous nuisance variables (arguing

that continuous nuisance variables can be dichotomised), only considered the situation where all nuisance variables were orthogonal, and did not consider different effect sizes that a researcher could deem acceptable. Instead they defined “groups as non-equivalent [...] if the proportion of subjects who fall in one category of the dichotomous nuisance variable in one group is at least twice that of the other group” (Hsu, 1989, p. 132). This is a very large effect size indeed: much smaller degrees of confounding can already be problematic depending on the robustness of the effect that is studied.

Because processing power was rare and expensive in the eighties, the approach chosen by Hsu (1989) made sense. However, today’s processing power enables looking at continuous nuisance variables (and median splits have even become widely considered to be unacceptable; MacCallum, Zhang, Preacher, & Rucker, 2002; Maxwell & Delaney, 1993). Nguyen et al. (2017) ran such simulations. First, they sampled repeatedly from datasets from large trials, comparing the distribution of nuisance variables and examining bias resulting from failure of the randomization procedure to create equivalent groups (henceforth: randomization failure). Second, they conducted a series of Monte Carlo simulations, which confirmed their results. On the basis of these analyses, Nguyen et al. (2017) recommend a much more pessimistic minimum required sample size of 1000 participants. They proceeded to examine the literature, concluding that over half of the phase 3 trials did not meet this standard.

In the current paper, we introduce a different approach that does not require simulations. Although we started out with simulations, we realised that in fact, the sampling distribution of Cohen’s d can be used to estimate the probability of group equivalence through randomization in simple cases. In the present paper we will introduce this

approach. We present a table and figures and provide a free and easy-to-use tool enabling researchers to compute the sample size required to achieve the desired likelihood of generating equivalent groups, for a given definition of equivalence.

Methods

Initially, we programmed a series of simulations in R (R Development Core Team, 2017) using the graphical user interface (GUI) provided by the open source software RStudio (RStudio Team, 2016). We simulated multivariately normal datasets using the MASS package (Venables & Ripley, 2002), supported by the plyr (Wickham, 2011), reshape2 (Wickham, 2007), and userfriendlyscience (G.-J. Y. Peters, 2017) packages, and using the ggplot2 package (Wickham, 2009) to produce the graphics. The script, results, and a supplementary file describing the simulations are available at the Open Science Framework at <https://osf.io/a8km5>. On the basis of these simulations, we concluded that although the differences were relatively small, the likelihood of equivalent groups was lowest when the nuisance variables were not associated. Roughly when we finalized these simulations, we realised that in fact, the likelihood of equivalent groups could be computed exactly using the Cohen's d distribution.

Our simplest simulation simply generated a random, normally distributed vector, which was then randomly divided into two. We then computed the difference between the means and expressed this as Cohen's d , replicating this 100000 times for each sample size in our simulation, and counting the proportion of samples where randomization yielded non-equivalent groups, where non-equivalence was defined as a difference between group means of $d > .1$ to $d > 1$. This proportion is in fact equal to the proportion of the central Cohen's d distribution for a given sample size that exceeds the specified d value. Functions to work

with the Cohen's d distribution had just been added to the `userfriendlyscience` package (G.-J. Y. Peters, 2017), so we extended these with two functions. One, `prob.randomizationSuccess`, computes the likelihood of group equivalence based on sample size, number of nuisance variables, and definition of group equivalence (i.e. maximum acceptable difference between the groups expressed as Cohen's d). The other, `pwr.randomizationSuccess`, computes the required sample size to achieve the desired probability of successful randomization (i.e. equivalent groups), given a number of nuisance variables and a definition of group equivalence (i.e. maximum acceptable difference between the groups expressed as Cohen's d). To use the functions, the following commands can be used in an R analysis script or entered in the R console:

```
install.packages('userfriendlyscience');  
require('userfriendlyscience');
```

The first of these commands downloads the package and installs it. This command only needs to be run once: the package will remain installed. The second command loads the package in the current session: this command has to be repeated in every session where the user wishes to use this package. After loading the package, the following command can be used to request sample sizes:

```
pwr.randomizationSuccess(dNonequivalence = 0.2,  
                          pRandomizationSuccess = 0.95,  
                          nNuisanceVars = 100);
```

The command above requests the sample size required to obtain a probability of 95% that two groups are equivalent, assuming that 100 nuisance variables exist and considering groups non-equivalent if they differ at $d > 0.2$ (typically considered a small effect size, which,

given the relatively small effects typically obtained by psychological interventions, seems reasonable). For studies that have already been conducted, the likelihood that two groups are equivalent can be obtained with the following command:

```
prob.randomizationSuccess(n = 1000,  
                           dNonequivalence = .2,  
                           nNuisanceVars = 100);
```

The command above requests the likelihood of equivalent groups (defined as groups differing at most with $d = 0.2$) in a situation with 100 nuisance variables with a total sample size of 1000 participants (i.e. 500 per group).

Results

We then proceeded to use these functions to generate Figures 1, 2, and 3 and Table 1. As can be seen from Figure 1, when just one nuisance variable exists, likelihood of non-equivalence drops quickly already at low sample sizes of less than a hundred participants. If a researcher works with effects that are very large and robust, and is willing to accept group differences of up to half a standard deviation on that single nuisance variable, already 63 participants (slightly over 30 per condition) suffice to achieve 95% likelihood of randomization success. However, most psychological manipulations and interventions have effects that are neither very large nor robust, so perhaps requiring groups to differ with at most a fifth of a standard deviation (Cohen's $d < 0.2$) seems more realistic. In that case, even if only a single nuisance variable is assumed to exist, already 385 participants are required, or almost 200 per condition.

Figure 1: Probability of one nuisance variable differing between groups with Cohen's d values of at most 0.1 (light green area), 0.2 (green area), 0.3 (yellow area), 0.4 (orange area), at most 0.5 (light red area) and more than 0.5 (red area) for sample sizes from 20-1000.

Of course, in most situations, assuming that only one nuisance variable exists is almost as naive as assuming that no nuisance variables exist. Figure 2 shows what happens when 10 nuisance variables exist, and Figure 3 shows what to expect if 100 nuisance variables exist. Researchers who study effects that are very subtle or sensitive to moderation, only considering groups equivalent if they differ at most with Cohen's $d = 0.2$, and assuming that 10 such potential moderators (i.e. nuisance variables) exist, require a total sample size of 787 participants (almost 400 per condition). Researchers who also assume 10 nuisance variables exist but study effects that are very robust, accepting group differences of up to half a standard deviation ($d = 0.5$), only require 128 participants, or 64 per condition. If 100 nuisance variables exist, the first group of researchers would require over a thousand participants (Table 1 shows that 1212 participants would be required, over 600 per condition), and the second group 198 participants (100 per condition).

Figure 2: Probability of one or more of ten nuisance variables differing between groups with Cohen's d values of at most 0.1 (light green area), 0.2 (green area), 0.3 (yellow area), 0.4 (orange area), at most 0.5 (light red area) and more than 0.5 (red area) for sample sizes from 20-1000.

Figure 3: Probability of one or more of one hundred nuisance variables differing between groups with Cohen's d values of at most 0.1 (light green area), 0.2 (green area), 0.3 (yellow area), 0.4 (orange area), at most 0.5 (light red area) and more than 0.5 (red area) for sample sizes from 20-1000.

How many nuisance variables exist can, by definition, never be known: if all nuisance variables were known, they could simply all be measured (although correcting for them would not be straightforward; Elwert & Winship, 2014; Westfall & Yarkoni, 2016). While clearly more than 1 nuisance variable exists for psychological experiments, whether the number is more likely to 10 or 100 can depend on the situation. Table 1 allows researchers to quickly look up the required sample size for scenarios ranging from situations where non-equivalent groups are very undesirable (requiring 99% likelihood of equivalence) to scenarios where non-equivalence is quite acceptable (accepting non-equivalent groups in one in five studies), for studies into subject matter that is very subtle or sensitive to moderators (considering group differences of $d = 0.1$ problematic) to studies into subject matter that is very robust (accepting group differences of $d = 0.5$), and assuming one, ten, or a hundred nuisance variables exist. In addition, the function `pwr.randomizationSuccess` in the `userfriendlyscience` R package can be used for situations not in this Table, for example when 42 or 67 nuisance variables exist. Which raises the question: how can one know how many nuisance variables exist?

Table 1: Required sample size to achieve 99%, 95%, 90%, and 80% likelihood of randomization success ($P(\text{equivalence})$ or $P(\text{eq})$), with non-equivalence defined as group differences exceeding Cohen's $d = 0.1, 0.2, 0.3, 0.4$ and 0.5 , and assuming one, ten, or a hundred nuisance variables exist.

	$d = 0.1$	$d = 0.2$	$d = 0.3$	$d = 0.4$	$d = 0.5$
1 nuisance variable					
$P(\text{eq}) = 0.99$	2656	666	297	168	109
$P(\text{eq}) = 0.95$	1538	385	172	97	63
$P(\text{eq}) = 0.9$	1083	271	121	68	44
$P(\text{eq}) = 0.8$	657	164	73	41	26
10 nuisance variables					
$P(\text{eq}) = 0.99$	4332	1086	485	275	178
$P(\text{eq}) = 0.95$	3138	787	351	199	128

P(eq) = 0.9	2623	657	293	166	107
P(eq) = 0.8	2098	526	235	133	86
100 nuisance variables					
P(eq) = 0.99	6057	1519	679	385	249
P(eq) = 0.95	4832	1212	541	307	198
P(eq) = 0.9	4297	1078	481	273	176
P(eq): 0.8	3744	939	419	237	153

Given that the very nature of psychological variables remains a debated topic (De Vries, 2017; Fried, 2017; Gruijters, 2017; G.-J. Y. Peters & Crutzen, 2017a, 2017c; Trafimow, 2017), and the question specifically concerns as yet unknown variables, we think this question is at present impossible to answer. Assuming only ten exist seems too liberal, whereas assuming a hundred exist seems too conservative. When considering the lower two thirds of Table 1, then, the more pertinent question becomes how large a difference between groups is deemed acceptable. As indicated before, this depends on the robustness of the effect under study. Very robust effects may barely be biased even when randomization failure turns a nuisance variable into a confounder that differs between the groups with half a standard deviation (Cohen's $d = 0.5$). Very subtle effects, on the other hand, might require groups to be much more equivalent, causing researchers to accept at most a difference of a tenth of a standard deviation ($d = 0.1$) on a nuisance variable-turned-confounder. In such a situation, desiring a large likelihood that randomization succeeds (e.g. accepting only 1% probability of failure) would mean that at least 2000 participants would be required per condition. Accepting a larger likelihood of randomization failure, for example when a field is willing to risk an invalidated manipulation in one in five studies, would cut the number of required participants per study in half, but would also mean that more studies would be required. Studying subtle psychological effects clearly has a price.

A randomization-based perspective on the replication crisis

These large sample sizes required for replication to ‘work’ mean that even when researchers power their studies quite highly for the purpose of null hypothesis significance testing (NHST), their manipulations have a high likelihood of being invalid (i.e. confounded by a nuisance variable), or at least conditional upon an unknown confounder. Because most studies are in fact underpowered (Bakker, Dijk, & Wicherts, 2012), randomization failure is likely common. From this point of view, the somewhat depressing results of large-scale replication studies (e.g. Open Science Collaboration, 2015) seem to make sense. Most replicated experiments (the original studies, that is) were extremely underpowered, even from an NHST point of view. This means that regardless of whether those original studies were correct regarding whether the hypothesized associations exist, the likelihood of widespread randomization failure means that the manipulations were often invalid or at least conditional upon unknown confounders.

To explore this, we selected all original studies with two-cell designs included in the Reproducibility Project: Psychology (Open Science Collaboration, 2015), available at <https://osf.io/fgjvw>. We extracted the sample sizes from the studies with two-cell designs and used these to estimate the likelihood of randomization success assuming one, ten, or a hundred nuisance variables exist, and for non-equivalence tolerances of $d = 0.2$ and $d = 0.5$. The results are shown in Figure 4, which shows that when assuming that for any given experiment in psychology, between ten and a hundred nuisance variables exist, in both the original studies and the replications randomization failure is likely to be widespread unless groups are allowed to differ with half a standard deviation or more. For subtle effects, the mean likelihood of randomization failure is 93.45% (95% confidence interval = [90.65%;

95.25%], median = 99.98%) assuming 10 nuisance variables, and 100% assuming 100 nuisance variables (this likelihood was 100% for all included studies). For robust effects, the mean likelihood of randomization failure is 44.54% (95% confidence interval = [35.03%; 54.05%], median = 38.54%) assuming 10 nuisance variables, and 69.07% (95% confidence interval = [59.25%; 78.88%], median = 99.16%) assuming 100 nuisance variables.

Figure 4: Probability of randomization failure for the original studies and the replications from the Reproducibility Project: Psychology.

Discussion

Thus, randomization requires large sample sizes. This means that studies that were designed as experiments but have low sample sizes are likely to, in fact, be quasi-experiments, where groups differ on one or more unknown confounders. This has two important implications. The first can be considered the good news. Recently a number of until recently considered well-established psychological phenomena were subjected to series of replications (Klein et al., 2014; Open Science Collaboration, 2015). Few of those replications

did in fact replicate the original findings. The low sample sizes imply that at least for some of these studies, randomization failed to produce equivalent groups in either the original study or the replication. One or more unknown confounders varied between conditions, causing supposedly equivalent groups to differ, sometimes perhaps only with a tenth of a standard deviation, but sometimes an entire standard deviation or more. In other words, only some of the replications represented real replications. If indeed arbitrary unknown variables confounded some of these replications, one would expect quite a heterogeneous pattern of effect sizes to emerge, which is indeed what was observed (Klein et al., 2014; McShane & Böckenholt, 2014). Note that although we used data from the Reproducibility Project: Psychology as an example, non-equivalent groups can help explain results of other replication projects (e.g. Hagger et al., 2016) as well.

The implication is that if a series of replications with low sample sizes fail to confirm a study's findings, that need not mean the original findings can be considered the result of a Type-1 error (or worse, scientific misconduct). However, if the original study, too, had a low sample size, those original results may be conditional upon the unknown confounders' distributions as they were in the original study, which of course is by definition unknown and therefore unreplicable. The solution is to no longer conduct experiments with low sample sizes, which brings us to the second implication, which can be considered the bad news: experiments in psychology require substantially larger samples than power analyses indicate (McShane & Böckenholt, 2014). An independent samples t-test with two groups achieves 80% power to detect a moderate effect size of Cohen's $d = 0.5$ with a total of 128 participants; if the dependent variable is also measured at baseline, only 54 participants are required to achieve 80% power to detect a moderate effect size. With 54 participants, the

probability that one arbitrary confounder differs between the groups with a Cohen's d of 0.2 or more is 45%, and 5% for a Cohen's d of 0.5. If ten potential confounders exist, the probability that at least one of them differs with a Cohen's d of 0.2 is 99.8%, and with 0.5, 50%.

In experiments with such small sample sizes, it is likely that the design's validity is flawed. Drawing conclusions that fail to acknowledge that, therefore, could be considered ethically circumspect. The only way to prevent this is recruiting more participants. This means that adequate experiments require considerably more resources than are currently typically spent. This is likely an unwelcome message to researchers (more time and effort for each study means fewer publications), funders (more funding per study means fewer studies), and editors (larger studies mean a lower likelihood of sensational findings). On the other hand, even when larger sample sizes would become commonplace, psychology would not suddenly become an exceptionally expensive field. Creating the conditions necessary for studying the subject matter of a field has costs. In some fields, researchers require clean rooms (e.g. T.-J. Peters & Tichem, 2016) or magnetic resonance imaging equipment (and also many participants; Szucs & Ioannidis, 2017). In psychology, many participants are required for randomization to succeed. The 'law of large numbers' (Nguyen et al., 2017), then, really is a law of 'large' numbers; with 'large' having no lesser meaning than 'at least a three-digit and sometimes a four-digit numerical value'.

The current paper did not explore the severity of the bias introduced by randomization failure. Nguyen et al. (Nguyen et al., 2017) concluded that bias minimized at 1000 participants, but simulations (e.g. using the R functions introduced in this paper) can be designed to estimate bias in different scenarios, enabling more precise guidelines for

different subfields of psychology. This will enable establishing, for example, how strongly a nuisance variable needs to moderate the manipulation's effect before bias can be expected, which will enable more accurate estimations of the required sample sizes. It might even turn out that group non-equivalence of $d = 0.5$ or even $d = 1.0$ can commonly be tolerated – but the results and conclusions of Nguyen et al. (2017) do not bode well.

Guidelines and randomization power

So how many participants are required? Partly because of the high sample sizes that are required to obtain replicable results in experimental designs (G.-J. Y. Peters & Crutzen, 2017b), researchers who want to be on the same side might be best off following the recommendation given by Nguyen et al. (2017). With a thousand participants, the likelihood of randomization success is 99% when assuming ten nuisance variables exist and when only tolerating group differences of up to a fifth of a standard deviation ($d < 0.2$), and still 90% when assuming a hundred nuisance variables exist. When studying phenomena that are known to be very robust, already 250 participants suffice to achieve randomization success in 99% of the studies even when assuming that 100 nuisance variables exist. Even the resulting 125 participants per design cell will likely take some getting used to for most researchers and funders, though. And where fewer participants can often be compensated by measuring each participant more frequently when planning studies for power or precision, this is not possible when planning for randomization success. In between subject-designs, increasing the number of data points by adding measurement moments does not increase the likelihood of randomization success. In within subject-designs, participants are not randomized in the first place, and so nuisance variables existing within participants are not a problem literally by design, but such designs are often not feasible.

It seems unlikely that experiments with samples of the sizes common in the literature can be relied on to randomize successfully. For experimental designs, power analyses based on null hypothesis significance testing (NHST) will almost always yield unrealistically low sample sizes (i.e. underestimates of what is required for a valid design). When researchers are interested in how effective a manipulation is, sample size should be planned based on the desired accuracy in parameter estimation (G.-J. Y. Peters & Crutzen, 2017b), not NHST; and when researchers want to conduct an experiment, the sample size should be sufficient to make it likely that the randomization procedure succeeds in producing equivalent groups.

A helpful way to express the ‘randomization power’ of a design is to compute the likelihood of equivalent groups both for a situation where ten nuisance variables exist, and a situation where a hundred nuisance variables exist. This recommendation is based on the fact that the number of nuisance variables will always, by definition, be unknown. However, it *will* often be possible to determine how robust the effect under study is. Therefore, researchers will often be able to decide how equivalent exactly they want the groups in their experiment to be. Because of the two parameters that are required to compute the likelihood that randomization succeeded for a given sample size, only the number of nuisance variables is impossible to estimate, it seems sensible to compute both a conservative (100) and liberal (10) estimate and report both. As an example, 128 participants yield 80% power for detecting a moderate effect size. However, if an experiment is conducted with 128 participants, if the phenomenon under study is robust (accepting group differences up to $d = 0.5$), the randomization power would be [60%; 95%], and if the phenomenon under study is subtle (accepting group difference of at most $d = 0.2$), the randomization power would be [0%; 5%]. This means that even for robust phenomena, randomization may fail to produce equivalent

groups in two out of five studies, and for subtle phenomena, randomization will yield equivalent groups in at most one out of every twenty studies.

On the bright side, once researchers resign to conducting larger studies, accurate conclusions as to effect size estimates also become possible (G.-J. Y. Peters & Crutzen, 2017b). In fact, in two-cell experiments with only 128 participants, effect size estimates will vary wildly between studies (G.-J. Y. Peters & Crutzen, 2017b), so even when studying robust effects, it seems advisable to use at least 750 participants, even if randomization success were guaranteed. Note that although in this paper, we use the simple example of a design with a dichotomous independent variable (which manifests as two groups or conditions), the logic holds for any experimental design: other designs will just require more participants.

Conclusion

Table 1 and the functions we introduced in this paper can assist researchers with their study planning. These same tools can assist editors and reviewers when scrutinizing submitted papers and determining which conclusions and claims are permitted by the collected dataset. The tools can be used by funders when calls for proposals are drafted to ensure that the conducted studies can yield replicable conclusions. Finally, ethical committees and institutional review boards can use these tools in their roles as gatekeepers, to make sure that the scarce (often public) resources that are invested in research are not wasted on studies with designs that are clearly invalid (e.g. with sample sizes so low that randomization success is unlikely).

Combined with full disclosure of materials and data (Crutzen, Peters, & Abraham, 2012; G.-J. Y. Peters, Abraham, & Crutzen, 2012) and complete transparency regarding the research proceedings (G.-J. Y. Peters, Kok, Crutzen, & Sanderman, 2017), conducting studies

with sufficiently large sample sizes to make randomization failure unlikely enables us to build a solid basis of empirical evidence. If this is combined with careful testing, development (Earp & Trafimow, 2015) and application (G.-J. Y. Peters & Crutzen, 2017c) of theory, this can yield a theory- and evidence base that can then confidently be used in the development of behaviour change interventions, eventually contributing to improvements in health and well-being. Conducting studies with sufficiently large sample sizes is as close to a guarantee of replication one is likely to come. This is an important message to funders as well: if the goal is to build a strong, replicable evidence base in psychology, it is necessary to fund studies with sample sizes that are considerably larger than what was funded in the past. It is also an important message for ethical committees: if a study has a flawed design (e.g. an experiment with randomization that is unlikely to fail because the sample size is too low), approving of that study as ethical is not straightforward. However, although the price is high (literally), the promised rewards are plentiful.

References

- Bakker, M., Dijk, A. Van, & Wicherts, J. M. (2012). The Rules of the Game Called Psychological Science. *Perspectives on Psychological Science*, 7, 543–554.
<https://doi.org/10.1177/1745691612459060>
- Bandura, A., Ross, D., & Ross, S. A. (1963). Imitation of film-mediated aggressive models. *Journal of Abnormal and Social Psychology*, 66(1), 3–11.
- Crutzen, R., Peters, G.-J. Y., & Abraham, C. (2012). What about trialists sharing other study materials? *BMJ (Online)*, 345(7887). <https://doi.org/10.1136/bmj.e8352>
- de Bruin, M., Hospers, H. J., van Breukelen, G. J. P., Kok, G., Koevoets, W. M., & Prins, J. M.

- (2010). Electronic monitoring-based counseling to enhance adherence among HIV-infected patients: a randomized controlled trial. *Health Psychology*, 29(4), 421.
- De Vries, H. (2017). Thinking is the best way to travel: towards an ecological interactionist approach; a comment on Peters and Crutzen. *Health Psychology Review*, 0(0), 1–6.
<https://doi.org/10.1080/17437199.2017.1306719>
- Earp, B. D., & Trafimow, D. (2015). Replication, falsification, and the crisis of confidence in social psychology. *Frontiers in Psychology*, 6(May), 621.
<https://doi.org/10.3389/fpsyg.2015.00621>
- Elwert, F., & Winship, C. (2014). Endogenous Selection Bias: The Problem of Conditioning on a Collider Variable. *Annual Review of Sociology*, 40(1), 31–53.
<https://doi.org/10.1146/annurev-soc-071913-043455>
- Fried, E. I. (2017). What are psychological constructs? On the nature and statistical modeling of emotions, intelligence, personality traits and mental disorders. *Health Psychology Review*, 0(0), 1–11. <https://doi.org/10.1080/17437199.2017.1306718>
- Gruijters, S. L. K. (2016). Baseline comparisons and covariate fishing: Bad statistical habits we should have broken yesterday. *European Health Psychologist*, 18(5), 205–209.
- Gruijters, S. L. K. (2017). The reasoned actions of an espresso machine: A comment on Peters & Crutzen (2017). *Health Psychology Review*, *In press*.(0), 1–15.
<https://doi.org/10.1080/17437199.2017.1306716>
- Hagger, M. S., Chatzisarantis, N. L. D., Alberts, H., Anggono, C. O., Birt, A., Brand, R., ... Cannon, T. (2016). A multi-lab pre-registered replication of the ego-depletion effect. *Perspectives on Psychological Science*, 25(4), 1227–1234.

<https://doi.org/10.1177/0956797614526415>.Data

Hsu, L. (1989). Random sampling, randomization, and equivalence of contrasted groups in psychotherapy outcome research. *Journal of Consulting and Clinical Psychology*, 57(1), 131–137.

Kessels, L. T. E., Ruiter, R. A. C., Brug, J., & Jansma, B. M. (2011). The effects of tailored and threatening nutrition information on message attention. Evidence from an event-related potential study. *Appetite*, 56(1), 32–38. <https://doi.org/10.1016/j.appet.2010.11.139>

Klein, R. A., Ratliff, K. A., Vianello, M., Adams, R. B., Bahník, Š., Bernstein, M. J., ... Nosek, B. A. (2014). Investigating variation in replicability: A “many labs” replication project. *Social Psychology*, 45(3), 142–152. <https://doi.org/10.1027/1864-9335/a000178>

Krause, M. S., & Howard, K. I. (2003). What random assignment does and does not do. *Journal of Clinical Psychology*, 59(7), 751–766. <https://doi.org/10.1002/jclp.10170>

MacCallum, R. C., Zhang, S., Preacher, K. J., & Rucker, D. D. (2002). On the practice of dichotomization of quantitative variables. *Psychological Methods*, 7(1), 19–40. <https://doi.org/10.1037//1082-989X.7.1.19>

Maxwell, S. E., & Delaney, H. D. (1993). Bivariate median splits and spurious statistical significance. *Psychological Bulletin*, 113(1), 181–190. <https://doi.org/10.1037//0033-2909.113.1.181>

McShane, B. B., & Böckenholt, U. (2014). You Cannot Step Into the Same River Twice: When Power Analyses Are Optimistic. *Perspectives on Psychological Science*, 9(6), 612–625. <https://doi.org/10.1177/1745691614548513>

- Nguyen, T.-L., Collins, G. S., Lamy, A., Devereaux, P. J., Daurès, J.-P., Landais, P., & Le Manach, Y. (2017). Simple Randomization Did not Protect Against Bias in Smaller Trials. *Journal of Clinical Epidemiology*. <https://doi.org/10.1016/j.jclinepi.2017.02.010>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251), aac4716. <https://doi.org/10.1126/science.aac4716>
- Papineau, D. (1994). The Virtues of Randomization. *The British Journal for the Philosophy of Science*, 45(2), 437–450. <https://doi.org/10.1093/bjps/45.2.437>
- Peters, G.-J. Y. (2017). userfriendlyscience: Quantitative analysis made accessible. Retrieved from <http://cran.r-project.org/package=userfriendlyscience>
- Peters, G.-J. Y., Abraham, C. S., & Crutzen, R. (2012). Full disclosure: doing behavioural science necessitates sharing. *The European Health Psychologist*, 14(4), 77–84.
- Peters, G.-J. Y., & Crutzen, R. (2017a). Confidence in constant progress: or how pragmatic nihilism encourages optimism through modesty. *Health Psychology Review*.
- Peters, G.-J. Y., & Crutzen, R. (2017b). Knowing exactly how effective an intervention, treatment, or manipulation is and ensuring that a study replicates: accuracy in parameter estimation as a partial solution to the replication crisis. Retrieved from <http://osf.io/preprints/psyarxiv/cjsk2>
- Peters, G.-J. Y., & Crutzen, R. (2017c). Pragmatic Nihilism: How a Theory of Nothing can help health psychology progress. *Health Psychology Review*.
<https://doi.org/10.1080/17437199.2017.1284015>
- Peters, G.-J. Y., Kok, G., Crutzen, R., & Sanderman, R. (2017). Health Psychology Bulletin:

- improving publication practices to accelerate scientific progress. *Health Psychology Bulletin*, 1(1), 1–6. <https://doi.org/10.5334/hpb.2>
- Peters, T.-J., & Tichem, M. (2016). Electrothermal Actuators for SiO₂ Photonic MEMS. *Micromachines*, 7(11), 200. <https://doi.org/10.3390/mi7110200>
- R Development Core Team. (2017). R: A Language and Environment for Statistical Computing. Vienna, Austria. Retrieved from <http://www.r-project.org/>
- RStudio Team. (2016). RStudio: Integrated Development Environment for R. Boston, MA. Retrieved from <http://www.rstudio.com/>
- Saint-Mont, U. (2015). Randomization does not help much, comparability does. *PLoS ONE*, 10(7), 1–24. <https://doi.org/10.1371/journal.pone.0132102>
- Senn, S. (2013). Seven myths of randomisation in clinical trials. *Statistics in Medicine*, 32(9), 1439–50. <https://doi.org/10.1002/sim.5713>
- Szucs, D., & Ioannidis, J. P. (2017). Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLoS Biology*, 15(3), e2000797. <https://doi.org/10.1371/journal.pbio.2000797>
- Trafimow, D. (2017). Why I Am Not a Fan of Pragmatic Nihilism. *Health Psychology Review*, 7199(March), 1–5. <https://doi.org/10.1080/17437199.2017.1306717>
- Van Breukelen, G. J. P. (2006). ANCOVA versus change from baseline: more power in randomized studies, more bias in nonrandomized studies [corrected]. *Journal of Clinical Epidemiology*, 59(9), 920–5. <https://doi.org/10.1016/j.jclinepi.2006.02.007>
- Venables, W. N., & Ripley, B. D. (2002). *Modern Applied Statistics with S* (Fourth). New York:

Springer.

Westfall, J., & Yarkoni, T. (2016). Statistically controlling for confounding constructs is harder than you think. *PLoS ONE*, 11(3), e0152719.

<https://doi.org/10.1017/CBO9781107415324.004>

Wickham, H. (2007). Reshaping Data with the {reshape} Package. *Journal of Statistical Software*, 21(12), 1–20. Retrieved from <http://www.jstatsoft.org/v21/i12/>

Wickham, H. (2009). *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York. Retrieved from <http://ggplot2.org>

Wickham, H. (2011). The Split-Apply-Combine Strategy for Data Analysis. *Journal of Statistical Software*. *Journal Of Statistical Software*, 40(1), 1–29.

Worrall, J. (2007). Why there's no cause to randomize. *The British Journal for the Philosophy of Science*, 58(3), 451–488. <https://doi.org/10.1093/bjps/axm024>