

Replicability Across Multiple Studies

Marina Bogomolov¹ and Ruth Heller²

Abstract. Meta-analysis is routinely performed in many scientific disciplines. This analysis is attractive since discoveries are possible even when all the individual studies are underpowered. However, the meta-analytic discoveries may be entirely driven by signal in a single study, and thus non-replicable. Although the great majority of meta-analyses carried out to date do not infer on the replicability of their findings, it is possible to do so. We provide a selective overview of analyses that can be carried out towards establishing replicability of the scientific findings. We describe methods for the setting where a single outcome is examined in multiple studies (as is common in systematic reviews of medical interventions), as well as for the setting where multiple studies each examine multiple features (as in genomics applications). We also discuss some of the current shortcomings and future directions.

Key words and phrases: Composite null, false discovery rate, meta-analysis, multiple hypothesis testing, replicability analysis.

1. INTRODUCTION

Meta-analysis is routinely performed in many scientific disciplines. In the health sciences, it is common to rely on systematic reviews and meta-analyses in order to inform practitioners of the best standard of care. For example, the Cochrane Collaboration carries out systematic reviews of health care interventions (Higgins et al., 2022). In large-scale analyses it is common to combine several independent studies that examine the same outcome in order to increase the discovery power. For example, in genome wide association studies (GWAS), every study examines many genotypes, and by combining several studies that examine the same phenotype it is typically possible to discover more associations between genotype and phenotype than in an individual study (see, e.g., Franke et al., 2010). The great majority of meta-analyses carried out to date do not infer on the replicability of their findings. Realizing that inference towards replicability of findings is possible from the studies meta-analyzed (Benjamini, Heller and Yekutieli, 2009), there are more and more works that target replicability from the available set of studies for meta-analysis.

Lack of replicability of scientific findings has been of great concern in the last two decades. In an influential

paper, Ioannidis (2005) claimed that the fraction of false positive findings in many domains in science is too high. The false positives in the published research may occur due to publication bias, or selective reporting, or other questionable research practices that inflate reported effect sizes, or mere bad luck. A common practice for corroborating findings is to bring forth prior findings published that support the novel findings. This practice assumes that there are no false positives in the published research and may therefore be problematic, potentially causing proliferation of false findings as ‘scientific’ facts. In order to reduce the lack of fieldwise replicability, there was a suggestion to lower the standard level for publication as a discovery, α , from 0.05 to 0.005. This suggestion was supported by 72 researchers that are authors of the paper by Benjamin et al. (2018), who believe that a leading cause for published false discoveries is that the statistical standards of evidence for claiming new discoveries in many fields of science are too low. A clear consequence of application of their suggestion is great power loss to make true discoveries. See Remark 4.1 for a further discussion of this suggestion.

An alternative way of raising the standard of evidence is to aim the analysis at replicable, rather than single, discoveries. For example, in GWAS, the demand for establishing replicability was often enforced by funders and journals, and GWAS findings are highly replicable (Marigorta et al., 2018). Suppose we have at our disposal two or more studies that examine the same (or a similar) problem. These may be a primary and then a follow-up

Marina Bogomolov is Associate Professor, Faculty of Data and Decision Sciences, Technion—Israel Institute of Technology, Haifa, Israel (e-mail: marinabo@technion.ac.il). Ruth Heller is Professor, Department of Statistics and Operations Research, Tel-Aviv University, Tel-Aviv, Israel (e-mail: ruheller@gmail.com).

study, or a group of studies that are meta-analyzed together. *Replicability analysis* refers to an analysis carried out towards identification of the replicable scientific findings. The following two observations regarding *replicability analysis* are important:

- It is more challenging than the typical analysis carried out in order to identify that a signal is present in at least one study, since replicability analysis takes care to rule out the possibility that the finding is entirely driven by a single study.
- It can always be carried out along with a meta-analysis, adding to the meta-analytic findings a quantitative measure of strength of evidence towards replicability of the findings.

When repeating a study, each repeated study necessarily differs from the original study (Nosek et al., 2022). Therefore, the efforts to corroborate true scientific findings may fail because these findings are very specific to the study, that is, they cannot be generalized to the conditions of a repeated study (which differs from the original study in many ways). Even if the innumerable differences from the original study are believed to be irrelevant for obtaining the evidence about these true scientific findings, the power may not be high enough in order to corroborate the findings. Recognizing the statistical challenge of potentially low power, the statistical community has developed methods specifically aimed at discovering the replicable scientific findings (i.e., the findings that are corroborated by multiple studies). We shall provide our selective review of these methods. They can be useful for practitioners, and if indeed used often then we expect that they will lead to a reduction in the fieldwise false positive findings.

After discussing the definition of replicability and introducing our notation, we review separately the case of a single feature (in Section 4), and of multiple features (in Sections 5 and 6). Can replicability be established if only a single study is available? Controlling for inflation of false positives in the single study is the first important step towards replicable research, and there is an abundance of multiple comparisons methods for this purpose. However, even with a severe multiple testing correction, one can obtain findings that will not be replicable in other studies. This may happen, for example, because of hidden biases. Of course, it is not possible to establish replicability by splitting the sample at random and thus creating multiple ‘studies’ examining the same problem, because the same biases will be present in all these ‘studies.’ However, if it is possible to split the data into subgroups based on certain important covariates (e.g., age, gender, ethnicity), establishing replicability is possible even within a single study. This setting is addressed in Section 7.

2. THE DEFINITION OF REPLICABILITY OF SCIENTIFIC FINDINGS

Since studies necessarily differ, we should not expect the same results across studies. We should not even expect the underlying (unknown) effects to be the same, since we cannot rule out the possibility that the sampling and measurement variations across studies had an effect on the underlying parameters of interest.

There is no consensus on how to quantify or measure replicability. Various definitions of replicability, and no-replicability, were considered in Nosek et al. (2022), Hung and Fithian (2020), Mathur and VanderWeele (2019), Patil, Peng and Leek (2019), Hedges and Schauer (2019b), Goodman, Fanelli and Ioannidis (2016), Simonsohn (2015) and references within. It is necessary to define what replicability actually is in order to be able to assess the extent of the evidence towards replicability of findings. The answer should be in terms of relevant parameters, which are the underlying quantities that relevant summary statistics aim to estimate. For example, for a single feature (or outcome) in two studies: a reasonable definition is that the null hypothesis is false in both studies (with both parameters pointing in the same direction); a relevant summary statistic is the maximum p -value; and a natural approach is to declare a finding as replicable if the maximum of the two p -values is at most α . Definitions of no-replicability include: nonnegligible statistical heterogeneity of effects; mixed effect directions across studies; an effect decline from an original study to a follow-up study.

2.1 r Out of n Replicability

We consider the setting where m features are examined in $n \geq 2$ independent studies. Let θ_{ij} denote the true parameter for feature i in study j . In GWAS, the parameter θ_{ij} may be a measure of association of the i th genotype with the phenotype in the j th study. In a systematic review that examines the effects of a health care intervention, θ_{ij} is the average treatment effect on the i th outcome in the j th study.

The research question regarding feature i in study j is addressed by testing the null hypothesis $H_{ij} : \theta_{ij} \in \Theta_{ij}^0$, for $i = 1, \dots, m, j = 1, \dots, n$, where Θ_{ij}^0 is a certain set of possible values of θ_{ij} . If the alternative is right sided, then $H_{ij} : \theta_{ij} \leq \theta_{ij}^0$, for a certain constant θ_{ij}^0 , and H_{ij} is false if $\theta_{ij} > \theta_{ij}^0$. Let $n_i^+ \in \{0, \dots, n\}$ be the number of parameters for feature i for which H_{ij} is false in this case, that is, $n_i^+ = |\{j \in (1, \dots, n) : \theta_{ij} > \theta_{ij}^0\}|$. If the alternative is left sided, then the inequalities above are reversed, and the number of false null hypotheses is defined as $n_i^- = |\{j \in (1, \dots, n) : \theta_{ij} < \theta_{ij}^0\}|$. If the alternative is two sided, then $H_{ij} : \theta_{ij} = \theta_{ij}^0$, and H_{ij} is false if $\theta_{ij} \neq \theta_{ij}^0$. In this case, the number of false null hypotheses for feature i is $n_i^+ + n_i^-$.

DEFINITION 2.1 (The r out of n replicability, $2 \leq r \leq n$). If the alternative is right sided, then we have r/n replicability for feature i if $n_i^+ \geq r$. The r/n no-replicability null hypothesis is

$$(1) \quad H_i^{r/n,+} : n_i^+ \leq r - 1,$$

so we have r/n replicability if $H_i^{r/n,+}$ is false. If the alternative is two sided, then we have r/n replicability for feature i if $n_i^+ \geq r$ or $n_i^- \geq r$, that is, at least r of the parameters in the n studies addressing feature i are non-null in the same direction. The r/n no-replicability null hypothesis is

$$(2) \quad H_i^{r/n} : \{n_i^+ \leq r - 1\} \cap \{n_i^- \leq r - 1\},$$

so we have r/n replicability if $H_i^{r/n}$ is false.

The *minimal no-replicability null hypothesis* for feature $i \in \{1, \dots, m\}$ is given by (1) or (2) with $r = 2$. The case $r = 1$ reduces to the global null hypothesis which is typically tested in meta-analysis. However, rejection of the global null hypothesis does not provide evidence towards replicability since the rejection may be due to signal in a single study.

For $n = 2$, the only possible replicability claims are minimal replicability claims. However, for $n > 2$, one can address the question whether the r/n replicability claim is true for different values of r , specifically, $r \in \{2, \dots, n\}$. Obviously, the higher is the value of r , the stronger is the corresponding replicability claim. For example, if three studies address the question whether a certain relationship holds, a claim that this relationship holds in all the three studies is stronger than a claim that it holds in at least two of the three studies. Thus, it may be attractive to test only the n/n no replicability null hypotheses, rejection of which gives the strongest possible claims regarding replicability. However, the power to reject the r/n no replicability null hypothesis decreases as r increases (since the composite null is larger, $H_i^{r/n} \subset H_i^{(r+1)/n}$). Therefore, even if the n/n no replicability null hypothesis is false, the power may be too low to detect it in practice. It is worth to take these power considerations into account when making the choice of r for $n > 2$. In Section 4 and Section 5 we show methods that do not require fixing r in advance. These methods provide a lower bound on the number of studies with signal for each selected feature, with a confidence guarantee.

Searching for replicated effects across studies has been considered in several application fields, where different terms have been used for what we call r/n replicability. This is referred to as generalizability in Nosek et al. (2022), Sofer et al. (2017), who consider generalization of effects to different ethnic groups. In the GWAS literature, it is referred to as pleiotropy when the phenotypes studied are distinct. For example, the single nucleotide polymorphisms (SNPs) that are associated with more than one

psychiatric disorder are pleiotropic SNPs (e.g., the SNPs associated with both Schizophrenia and bipolar disorder, Andreassen et al., 2013). It is referred to as consistency or external validation when the different studies correspond to different environments (Li et al., 2022).

Assessing replicability can be useful for increasing the evidence for causality of effects. For example, higher rates of leukemia among radiologists and among survivors of Hiroshima and Nagasaki provide replicable evidence of the increased risk of leukemia caused by radiation (Rosenbaum (2001), see more examples in Rosenbaum (2022)). For assessing causality in etiological epidemiology, Lawlor, Tilling and Davey Smith (2017) have shown the usefulness of identifying that an effect is replicated in studies with different biases, some of which are, desirably, in different directions. Lawlor, Tilling and Davey Smith (2017) use the term *triangulation* to describe obtaining more reliable or accurate answers to research questions by comparing results in studies from two or more different epidemiological approaches. Since triangulation will mostly provide a qualitative assessment of the strength of evidence regarding causality (Lawlor, Tilling and Davey Smith, 2017), the methods reviewed in this paper could be used in order to complement this assessment with quantitative measures regarding replicability strength across the different studies.

Definition 2.1 is useful when methods of measurements differ and may not be on the same scale from one study to the next, since θ_{i1}^0 need not equal θ_{i2}^0 . Moreover, it is useful when only the qualitative conclusion matters, so it is enough to know that the association/phenomenon exists in at least r of the studies. Simonsohn (2015) argues that this is the case for many studies in psychology. As an example, he discusses the (hypothetical) phenomenon of levitation: if one experiment concludes that people can levitate on average 9 inches above ground, and another concludes that the average is 0 inches, then the meta-analytic average is 4.5 inches but there is no replicated evidence for the phenomenon. There would have been $2/2$ replicability only if both experiments concluded that people can levitate on average a positive amount above ground (and this positive amount may differ across the studies).

REMARK 2.1. According to our definition, a true r/n replicability claim is made if at least $r \geq 2$ of the parameters in the n studies are nonnull (in the same direction), even if the nonnull parameters differ across studies. We realize that this definition of replicability may seem unsatisfying for some purposes. Hedges and Schauer (2019a) argue that if effects are in the same direction but vary in several orders of magnitude, the result may not be considered replicated. If the primary study has a positive effect, and a follow-up study has a positive yet much smaller effect, they would like to define it as nonreplicated. This is in contrast with our definition which, with enough data,

will consider it replicated. One way to alleviate their concern, when the hypothesis testing problems are one sided, is to define a feature as replicated if *the effect is at least of a meaningful (predefined) magnitude in at least r studies* (Mathur and VanderWeele, 2019, Jaljuli et al., 2022).

REMARK 2.2. We concentrate on identification of the replicated signals. For assessing fieldwise replicability measures we refer the reader to Hung and Fithian (2020) and references within. These measures are important when focusing on the reliability of a field of science. This is not the focus of this review, though it is related in the sense that by routinely applying the procedures discussed in this review, which aim to detect replicated signals while controlling for false replicability claims, the reliability of the field may increase. Specifically, the fraction of published replicable relationships that are false is expected to be smaller than the fraction of published relationships that are false.

REMARK 2.3. One of the settings addressed in Section 5.3 considers the null hypothesis H_{ij} to be conditional independence between an outcome and a predictor i , given all other predictors, in study j (Li et al., 2022). With some abuse of notation, the r out of n replicability in this case is that $n_i^+ \geq r$, where n_i^+ represents the number of studies in which predictor i and the outcome are conditionally dependent.

2.2 Reasons for Lack of Replicability

Since there is unavoidable heterogeneity across studies, it can happen that the effect is present in exactly one of the studies, and therefore it is clearly not a replicated effect. Formally, $H_i^{r/n,+}$ may be true for $r \geq 2$, even though $n_i^+ = 1$. This explains why rejecting the global null hypothesis, $n_i^+ = 0$, is not enough in order to make a replicability claim. For r/n replicability with $r > 2$, even rejecting $H_i^{(r-1)/n,+}$ is not enough—it is necessary to reject $H_i^{r/n,+}$.

For example, in GWAS, an association between genotype and phenotype may be particular to a single study and thus nonreplicable. We provide the following documented case regarding Crohn’s disease. Franke et al. (2010) has successfully detected more than 70 genomic regions which are susceptible for association with Crohn’s disease in populations of European ancestry, but only a few have been confirmed in the East Asian population (Marigorta et al., 2018, Nakagome et al., 2012). As suggested by Nakagome et al. (2012), the lack of any association with Crohn’s disease in a specific region NOD2 in East Asians could be due to the specificity of the causal association to the population of European ancestry due to natural selection. Another possible explanation is the lack of representation of causal variants, see Section 5.2 of Li et al. (2022). Briefly, the genome has block-like patterns

of linkage disequilibrium, so that the alleles are dependent within each block, and are approximately independent across the blocks. For certain noncausal SNPs, the null hypothesis of conditional independence may be false just because they belong to the same block as the unmeasured causal SNP. The block-like patterns of linkage disequilibrium vary across populations with different ancestries (Kidd et al., 2004). Therefore, a noncausal SNP may be in the same block as the unmeasured causal SNP in the population addressed by study one, implying the falsehood of its null hypothesis in study one, but in a block containing no causal SNPs in a population addressed by study two, implying the truth of its null hypothesis in study two.

Additional examples of cases where a null hypothesis may be false in one study and true in the other study can be found in Wang et al. (2022), who search for replicated signals across different platforms or ethnic groups. Of course it may be that the signal is present only in a single study, which is carried out on a specific platform or on a specific ethnic group.

Another reason for a finding to fail to replicate is that the finding may be due to bias: the null hypothesis of no treatment effect is true, but its p -value is stochastically smaller than uniform, so it behaves as a p -value of a false null hypothesis. This is certainly a danger in observational studies. If hidden bias is present, then differing outcomes in treated and control groups may not actually be effects caused by the treatment. Suppose that this is the case in a single study, and that there is no treatment effect. Suppose also that the other studies considered do not replicate the same hidden bias. Then the prospect of replicating the finding (a ‘treatment effect’) in the single study by the other studies is reduced (Rosenbaum, 2001). This may be the case in Hughes et al. (2020), Panagiotou, Jaljuli and Heller (2020), where it is found that blood pressure lowering with antihypertensive agents is associated with a lower risk of incident dementia or cognitive impairment in at least one study, but possibly only in one study.

3. ERROR RATES FOR CONTROLLING ERRONEOUS REPLICABILITY CLAIMS

Let us first consider the case where the alternatives are right sided. In general, our goal is to identify features with effects in at least r studies, for $r \in \{2, \dots, n\}$. This goal can be achieved by testing the family $\{H_i^{r/n,+}, i = 1, \dots, m\}$. Rejection of $H_i^{r/n,+}$ leads to claiming that an effect for feature i exists in at least r studies. This replicability claim is false if $H_i^{r/n,+}$ is true, that is, if the number of studies with effects for feature i is smaller than r . Let $V^{r/n}$ and $R^{r/n}$ be the number of false replicability claims and the total number of replicability claims, respectively. The familywise error rate (FWER) for replicability analysis is the probability that one or more replicability claims

that were made are false, that is,

$$\text{FWER}^{r/n} = \mathbb{P}(V^{r/n} > 0).$$

Similarly, the false discovery rate (FDR, [Benjamini and Hochberg, 1995](#)) for replicability claims is

$$\text{FDR}^{r/n} = \mathbb{E} \left[\frac{V^{r/n}}{\max(R^{r/n}, 1)} \right].$$

Rather than identifying the features with effects in at least r studies, one can obtain a lower confidence bound for the number of such features (or equivalently, an upper confidence bound on the number of features that have an effect in at most $r - 1$ studies), in any selected set of features, with a predefined simultaneous coverage guarantee, see [Goeman and Solari \(2011\)](#) and [Blanchard, Neuvial and Roquain \(2020\)](#), among others, for methods addressing this goal.

For two-sided alternatives it is of interest to identify features with effects in the same direction in at least r studies, and to declare the common direction of effects. Let $R_i^{r/n,L}$ and $R_i^{r/n,R}$ be the indicators of whether the r/n replicability claim for feature i was made in the left and right directions, respectively. Define $R^{r/n,D} = \sum_{i=1}^m R_i^{r/n,L} + \sum_{i=1}^m R_i^{r/n,R}$ as the total number of r/n replicability claims. The directional replicability framework accounts for type III errors ([Tukey, 1991](#), [Jones and Tukey, 2000](#)), which occur when the declared common direction of effects is wrong. Specifically, the numbers of false and true directional r/n replicability claims are $V^{r/n,D}$ and $S^{r/n,D}$, respectively, where

$$\begin{aligned} S^{r/n,D} &= \sum_{i:n_i^+ \geq r} R_i^{r/n,R} + \sum_{i:n_i^- \geq r} R_i^{r/n,L}, \end{aligned}$$

and $V^{r/n,D} = R^{r/n,D} - S^{r/n,D}$. The FWER and FDR for directional r/n replicability analysis are defined with respect to the above definitions of true and false discoveries, that is,

$$\text{FWER}^{r/n,D} = \mathbb{P}(V^{r/n,D} > 0),$$

and

$$\text{FDR}^{r/n,D} = \mathbb{E} \left[\frac{V^{r/n,D}}{\max(R^{r/n,D}, 1)} \right].$$

The next section addresses the case where there is a single feature examined in each study ($m = 1$), in which the error rates above reduce to the probability of making a false r/n replicability claim, or at least one false directional r/n replicability claim. In Section 5, we address the case of multiple features, and review several methods for controlling the error rates above. In Section 6, we address the control of the Bayesian counterpart of the FDR on replicability claims, for high dimensional studies.

4. THE CASE OF A SINGLE FEATURE PER STUDY

4.1 Parametric and Naive Approaches

Systematic reviews attempt to assemble all the studies that are relevant to specific research questions. They are important for advancing good standards of care in the field of medicine, where the questions are about medical interventions. In particular, the Cochrane collaboration routinely conducts systematic reviews in order to assess the effects of health care interventions ([Higgins et al., 2022](#)).

The systematic review may carry out one, or several, meta-analyses. Each meta-analysis carried out examines a single intervention. For a meta-analysis of a specific (e.g., primary) outcome in the systematic review, let $\theta_{\mathbf{1}} = (\theta_{11}, \dots, \theta_{1n})$ denote the (unknown) effects of the intervention in the n studies meta-analysed. The meta-analysis is used to infer on whether the effect of the intervention is present in at least one study (by rejecting the test of the global null hypothesis that there is no effect in all the studies). In addition, the meta-analysis typically provides confidence intervals assuming that the effect is common to all studies (the ‘fixed effect’ model, where it is assumed that $\theta_{11} = \dots = \theta_{1n}$) or that the effects underlying the studies follow some distribution, the ‘random effects’ model.

For the ‘random effects’ model, the assumed underlying distribution is typically Gaussian, so it is assumed that $\theta_{1j} \stackrel{\text{iid}}{\sim} N(\mu, \tau^2)$. With this strong assumption on the distribution of the θ_{1j} s, it is possible to infer on the overall effect ([Higgins et al., 2022](#)), as well as provide $1 - \alpha$ level prediction intervals for the θ_{1j} ’s. [Saad et al. \(2019\)](#) view the intervention effect in a new study as independently sampled from the random effects model. By projecting the estimates and confidence regions for the population mean and standard deviation, they produce pointwise estimates and confidence intervals for the cumulative distribution function of the random effects model, which provide a quantitative method for clinical decision-making that takes into consideration the heterogeneity of the intervention effect. Replicability can also be addressed within this model: [Mathur and VanderWeele \(2019\)](#) suggest several metrics for inference, including the fraction of effects that exceed a predefined threshold.

The assumption that the true (unobserved) treatment effects come from a known distribution can have a detrimental impact on inference if it is false. This is so especially when the number of studies is small so the assumption cannot be verified using diagnostic tools. In particular, in systematic reviews of medical interventions, the number of studies is typically small, and there is no reason to believe that the unobserved treatment effects in a few different cohorts (or subgroups, or environments) have a bell-shaped histogram that concurs with the model assumption. Therefore, we would like to provide tools for assessing the extent of the evidence towards replicability, that do not rely on any parametric assumptions.

An intuitive way of declaring n/n replicability, which is also statistically valid, is if the p -values from all $n > 1$ studies are at most α (the predefined probability of a type I error). For $n = 2$, the probability of a false replicability claim is then at most α , and this bound is tight since it occurs when the p -value in one study is almost surely zero (i.e., a very strong association is detected), yet in the other study the p -value is uniform. However, for $n > 2$ the requirement that all p -values be at most α is too stringent, at least when it is enough to establish that there is signal in at least two studies—the minimal replicability requirement.

A naive approach is to declare minimal replicability if at least two studies have p -values at most α . However, the naive approach is not valid, since the probability that the replicability claim is false can be far greater than α . To see this, consider the common setting in which p -values corresponding to true null hypotheses are uniformly distributed, so the probability that each of these p -values is at most α is α . If at most one study is nonnull, then there is no replicability. However, the probability that at least two studies will have p -value at most α is lower bounded by the probability that a binomial random variable with n trials and success probability α is at least two, that is, $1 - (1 - \alpha)^n - n(1 - \alpha)^{n-1}\alpha$. For example, for $n = 10$ and $\alpha = 0.05$, it is 0.09. This lower bound is the probability of a false replicability claim (i.e., of rejecting $H_1^{2/n}$, even though $H_1^{2/n}$ is true), only when all the p -values are uniformly distributed. However, if a single study has p -value less than 0.05 with probability one, and the remaining $n - 1 = 9$ p -values are uniformly distributed, then the probability of a false replicability claim increases to $1 - (1 - \alpha)^{n-1} = 1 - 0.95^9 = 0.37$. The naive approach should not be used for establishing replicability, since the probability of falsely declaring minimal replicability is nonnegligible even for n small, and increases to one rapidly.

In Section 4.2, we briefly describe the approach in Jaljuli et al. (2022), which is completely nonparametric, and therefore it can be used for replicability analysis when n is small or large.

REMARK 4.1. Decreasing the rejection threshold from $\alpha = 0.05$ to $\alpha = 0.005$, following Benjamin et al. (2018), indeed decreases the probability of making a false replicability claim with the naive approach. However, when the number of studies n is large enough, this probability can still be close to one. In Section 4.2, we suggest a method that guarantees control of the probability of making a false replicability claim at the desired level α . Of course, as α decreases the power to detect a false no-replicability null hypothesis decreases, and the power deterioration may be large when moving from an $\alpha = 0.05$ to $\alpha = 0.005$. The suggestion to reduce α by

Benjamin et al. (2018) was motivated by the need to have a stricter statistical standard for claiming new discoveries. Since the claim of replicability is stronger than the claim of a single new discovery, it may not be necessary to raise the standards even higher by reducing α .

4.2 Replicability Analysis with Partial Conjunction p -Values

For a general method for obtaining a valid p -value for the r/n -no-replicability null hypothesis (1), see Benjamin and Heller (2008), Wang and Owen (2019), where this null hypothesis and its p -value are referred to as r/n partial conjunction (PC) null hypothesis and p -value, respectively. One example for a PC p -value is the following. Let p_{1j} be the right-sided p -value associated with H_{1j} , in study j , for $j = 1, \dots, n$. The n p -values are mutually independent, since we address the setting of n independent studies. Let $p_{1(1)} \leq p_{1(2)} \leq \dots \leq p_{1(n)}$ be the ordered sequence of right-sided p -values. The PC p -value motivated by Fisher’s p -value combining method (Fisher, 1934) for $H_1^{r/n,+}$ is

$$(3) \quad p_1^{r/n} = \mathbb{P} \left[\chi_{2(n-r+1)}^2 \geq -2 \sum_{j=r}^n \log p_{1(j)} \right].$$

Jaljuli et al. (2022) use the PC p -values in order to establish replicability in systematic reviews from the Cochrane library. The r/n PC p -value is a p -value for the r/n no replicability null hypothesis, so the smaller it is the greater the evidence towards r/n replicability. It is thus an informative quantitative measure of the extent of evidence towards replicability, just as the p -value is informative for reporting in individual studies. This measure is attractive for assessing replicability for the following reasons:

- Even if all studies are underpowered, it may convey that there is evidence towards replicability (this property is similar to the attractive property in meta-analysis of pooling evidence across the studies in order to discover that there is signal in at least one study). In particular, even if the p -values in all individual studies are larger than $\alpha = 0.05$, the PC p -value may be smaller than 0.05. For example, for $n = 3$ p -values with value 0.090, $p_1^{2/3} = 0.047$ with Fisher’s combining method.
- Even if $r - 1$ studies have tiny p -values, it may convey that there is no evidence for r/n replicability. For example, for $n = 3$ p -values with values 0.00001, 0.2, and 0.4, $p_1^{2/3} = 0.28$ with Fisher’s combining method; the global null p -value is $p_1^{1/n} = 9 \times 10^{-5}$. So in this example we can conclude that there is signal in at least one study, but we cannot rule out that the signal may be present only in one study (since we fail to establish minimal replicability at $\alpha = 0.05$, or any level $\alpha < 0.28$).

In addition, a $1 - \alpha$ confidence lower bound on the number of studies with effect is provided by testing the PC hypotheses in order starting from $r = 1$, following [Benjamini, Heller and Yekutieli \(2009\)](#). Specifically, this lower bound is given by

$$(4) \quad l_\alpha = \max\{l \in \{0, \dots, n\} : p_1^{r/n} \leq \alpha \text{ for } r = 0, \dots, l\},$$

where $p_1^{0/n}$ is defined as 0 (see [Jaljuli et al., 2022](#) for details).

For directional inference, one-sided p -values are combined in each direction in [Jaljuli et al. \(2022\)](#), following [Owen \(2009\)](#). For example, the PC p -value for $H_1^{r/n}$ using Fisher's combining method is $2 \min\{q_1^{r/n}, p_1^{r/n}\}$, where $p_1^{r/n}$ is defined in (3) and

$$q_1^{r/n} = \mathbb{P}\left[\chi_{2(n-r+1)}^2 \geq -2 \sum_{j=r}^n \log q_{1(j)}\right],$$

where $q_{1(1)} \leq q_{1(2)} \leq \dots \leq q_{1(n)}$ is the ordered sequence of left-sided p -values (for continuous test statistics, $q_{1j} = 1 - p_{1j}$, $j = 1, \dots, n$). If the PC p -value above is below α , the common direction of effects is declared to be right sided if $p_1^{r/n} < q_1^{r/n}$, and left sided otherwise.

REMARK 4.2. The methodology we describe takes care to control the probability of declaring replicability falsely in the case that exactly one study has an association, that is, $n_1^+ = 1$ when testing right-sided alternatives. We believe it is important to apply replicability analysis methods that are valid even when $n_1^+ = 1$, since such methods guard against claiming replicability due to evidence coming only from a single study, which is a realistic setting that cannot be ruled out, as explained in [Section 2](#). Of course, if it is a priori somehow known that $n_i^+ = 1$ cannot be true, then it is possible to devise more powerful criteria for replication success. The potential increase in power is due to the fact that in this case the no-replicability null hypothesis, $n_1^+ = 0$, is much narrower than (1). Specifically, for testing the 2/2 no replicability null hypothesis we require that $\max(p_{11}, p_{12}) \leq \alpha$ in order to declare replicability. The least favorable configuration, for which the probability that $\max(p_{11}, p_{12}) \leq \alpha$ is exactly α , is that $n_1^+ = 1$, and that this study has a p -value of zero (almost surely), and the other study has a p -value that is uniformly distributed. However, if it is assumed that $n_1^+ \in \{0, 2\}$, then if there is no replicability, the global null is assumed to be true and so the probability that $\max(p_{11}, p_{12}) \leq \alpha$ is α^2 (i.e., the test is too conservative). Under this strong assumption, where no replicability means that the global null is true, [Pawel and Held \(2022\)](#), [Held, Micheloud and Balabdaoui \(2022\)](#) suggested less conservative tests for replication success.

4.3 Extensions and Future Directions

The difference in power between different combining methods used for testing $H_1^{r/n}$ can be large, and the combining methods differ in the settings in which they work best. A recent comparative study ([Hoang and Dickhaus, 2022](#)) examines several combination methods of p -values as well as of e -values. Choosing a good combination method is important in order to have good power to establish replicability, but the choice depends on a priori unknown characteristics of the data generation, notably the number of studies with signal. Therefore, it can be useful to have adaptive methods that let the data guide the choice of the appropriate combination method for testing the no-replicability null hypothesis.

Meta-analysis is prone to publication bias, where only significant results (typically at p -value ≤ 0.05) are published. If the group of studies considered may be biased, a sensitivity analysis that takes the bias into account should be carried out. For example, the p -values for the replicability analysis may be adjusted for selection bias by combining for PC testing only p -values at most 0.05, after inflating each by multiplying it by $1/0.05 = 20$ ([Fithian, Sun and Taylor, 2017](#), [Zhao, Small and Su, 2019](#), [Hung and Fithian, 2020](#)). Such an analysis is more conservative, but also more robust to the bias that arises from publishing only findings that are significant at the 0.05 level.

A lower confidence bound on n_1^+ is of great interest for assessing replicability. An upper confidence bound may also be of interest, since it conveys the limit on replicability, or the specificity, to a subset of studies. [Heller and Solari \(2023\)](#) provides methods targeting upper (as well as lower) confidence bounds for n_1^+ .

Finally, the systematic reviews for health care interventions typically consider multiple primary and secondary outcomes. Therefore, they carry out multiple meta-analyses that each examines a single endpoint. It can be thus useful to develop a replicability analysis pipeline that takes the multiplicity of the outcomes into account, to guarantee a desired overall error rate. The setting of multiple outcomes fits into the methods we discuss next in its simplest form where the outcomes can be treated as exchangeable. This is not necessarily the case for medical interventions, where it may be best to use weights or have a hierarchical structure that conveys the relative importance of, or dependence between, the outcomes.

5. FREQUENTIST APPROACHES FOR MULTIPLE FEATURES PER STUDY

5.1 The Shortcomings of the Naive Approach

A common practice is to corroborate the reported findings from a study with reported findings in one or more previous studies that examine the same or a similar problem. Although this is an intuitive approach for establishing replicability, it is not a formal statistical approach.

For a single feature, that is, $m = 1$, this approach reduces to the naive approach of declaring minimal replicability when the p -value is at most α in this study as well as in at least one of the previous $n - 1$ studies. This approach is valid for $n = 2$, as discussed in Section 4, but it is not valid for $n > 2$. In this section we address the settings where $m > 1$, and characterize the scenarios in which the naive approach may be problematic in the following two aspects: it may cause an unacceptable inflation of false replicability claims; it may lack power to discover features with replicated signals (that a nonnaive replicability analysis procedure may have good power to discover).

Let us first formalize the naive approach that we address. Although the investigators often report in their papers that their findings were replicated in one or two studies examining the same features, in practice, it seems natural that the investigators consider multiple previous studies examining their features of interest, and pick for claiming replicability only those which indeed corroborate their findings. Our formal simplification is that the analyst declares as replicated the features in the set $\mathcal{D}_1 \cap (\bigcup_{j=2}^n \mathcal{D}_j)$, where \mathcal{D}_1 is the set of features discovered in the given study, call it study one, and \mathcal{D}_j , $j = 2, \dots, n$ are the sets of features discovered in $n - 1$ previous studies, where each \mathcal{D}_j is the result of a multiple testing procedure (e.g., Bonferroni). Of course in practice n may not be fixed in advance.

We illustrate the problematic aspects of the naive approach by examples below and in the Supplementary Material (SM) (Bogomolov and Heller, 2023). Briefly, some key observations are as follows. First, suppose that the proportion of nonreplicable signals in study one is non-negligible. Suppose that the discoveries in study one are those that pass the severe Bonferroni correction at level α (i.e., only features with $p_{i1} \leq \alpha/m$ are rejected). Then, for any fixed m and $n > 2$, the naive approach may lead to an inflated FWER for minimal replicability analysis, that is, the probability of making one or more false minimal replicability claims may be higher than α . The inflation may increase, and thus be of greater concern, as (1) the number of studies considered, n , increases, or (2) the proportion or strength of nonreplicable signals in study one increases, or (3) the analyst applies a multiple testing procedure within each study that is less stringent than Bonferroni, for example, if Bonferroni is replaced by Hommel’s procedure (Hommel, 1988) or the Benjamini–Hochberg (BH) procedure (Benjamini and Hochberg, 1995). The naive approach with discoveries from the BH procedure may lead to a highly inflated FDR for replicability analysis even for $n = 2$. Second, consider the case when m is large and the proportion of nonreplicable signals in study one is close to 0, which is a typical scenario in GWA studies. Then the naive approach with Bonferroni and with BH may be

valid in terms of FWER and FDR control, respectively. However, in these settings their power to discover features with replicated signals may be very low, compared to the procedures in Section 5.2 and Section 5.3, which are designed for controlling the FWER or FDR for replicability analysis. We thus recommend using the latter procedures and not the naive approach above, both for the sake of replicability error control and for the sake of power for identifying true replicated signals.

We now give more details and examples for justifying the observations above. For simplicity, we address one-sided hypotheses. We start by two examples which apply the Bonferroni procedure in each study. They serve to demonstrate the validity, yet low power, the naive approach can have for $n = 2$ and for GWAS-like problems with $n > 2$, as well as the inflation it can have in non GWAS-like problems. We provide in the SM additional examples that assess the naive approach when in each study discoveries are made using FWER-controlling procedures (including those which are less conservative than Bonferroni).

EXAMPLE 5.1. Assume one faces $n = 2$ studies examining $m > 1$ features, and the discoveries within them are obtained using the level- α Bonferroni procedure, which rejects a given hypothesis if its p -value is below α/m . In this case, the set of indices of features with replicability claims based on the naive approach is $\mathcal{D}_1 \cap \mathcal{D}_2 = \{i : \max(p_{i1}, p_{i2}) \leq \alpha/m\}$, that is, the rejection set of Bonferroni applied on $\{\max(p_{i1}, p_{i2}), i = 1, \dots, m\}$. Since for any $i \in \{1, \dots, m\}$, $\max(p_{i1}, p_{i2})$ is a valid p -value for $H_i^{2/2,+}$ defined in (1), and the Bonferroni procedure guarantees FWER control under any dependency among the p -values, in this case the naive approach guarantees FWER control with respect to replicability claims. Although this approach is valid, it may be much more conservative than the FWER-controlling Procedure 5.1 with Bonferroni in Step 2, especially when the proportion of signals within each study is small, see Remark 2 in Bogomolov and Heller (2018).

EXAMPLE 5.2. We now address the configuration in Table 1. For simplicity, we address the idealized setting where the p -values for false null hypotheses are, almost surely, equal to 0. We assume that the hypotheses within each study are tested using Bonferroni at level α . In this case, an analyst can make false minimal replicability claims for all the features except for feature m , for which the effects are replicated in all the studies. The probability of making a false minimal replicability claim for feature 1 is $g_1(\alpha) = 1 - (1 - \alpha/m)^{n-1}$, and the probability of making a false replicability claim for a given feature with index $i \in \{2, \dots, m - 1\}$ is $g_2(\alpha) = (\alpha/m)[1 - (1 - \alpha/m)^{n-1}]$. Therefore, under indepen-

TABLE 1

Illustration of the setting addressed in Example 5.2. The ij th entry indicates the distribution of p_{ij} , for each $i \in \{1, \dots, m\}$, $j \in \{1, \dots, n\}$: 0 indicates that the given p -value is 0 with probability 1, $U(0, 1)$ indicates that its distribution is uniform

	Study 1	Study 2	Study 3	...	Study n
1	0	$U(0, 1)$	$U(0, 1)$...	$U(0, 1)$
2	$U(0, 1)$	$U(0, 1)$	$U(0, 1)$...	$U(0, 1)$
...
$m-1$	$U(0, 1)$	$U(0, 1)$	$U(0, 1)$...	$U(0, 1)$
m	0	0	0	...	0

dence within each study, the FWER for minimal replicability analysis is equal to

$$\begin{aligned} & \mathbb{P}(\text{at least one false minimal replicability claim}) \\ &= 1 - [1 - g_1(\alpha)][1 - g_2(\alpha)]^{m-2} \\ &= 1 - (1 - \alpha/m)^{n-1} \\ & \quad \times \{1 - (\alpha/m)[1 - (1 - \alpha/m)^{n-1}]\}^{m-2}. \end{aligned}$$

When m is large and $n \ll m$, so that n/m is close to 0, the FWER for replicability analysis is close to 0. This conclusion suggests that the naive approach is valid for GWAS, since this setting can be viewed as a stylized GWAS example.¹ However, as will become clear in the following sections, there are more powerful approaches than the naive approach. For $\alpha = 0.05$ and $n = 10$, the FWER for replicability analysis is 5×10^{-7} for $m = 10^6$, and it increases to 0.005 for $m = 100$. Moving away from the GWAS setting, for small m and $n \gg m$, the FWER for replicability analysis may approach 1. For $\alpha = 0.05$, $m = 10$ and $n = 100$, the FWER is 0.4.

Let us now consider the case where the hypotheses within each study are tested using the level- α BH procedure. A toy example in the SM shows that even in the low-dimensional case where $m = n = 2$, the FDR for replicability analysis may be 2α . The following example shows that the replicability FDR of the naive approach can be close to 1 in settings with large m and a nonnegligible proportion of features with nonreplicable effects. For simulations showing the inflation of replicability FDR of the naive approach with BH under more realistic settings with $n = 2$ studies, see Figure 1 in Bogomolov and Heller (2018) and Figure 3 in Bogomolov and Heller (2013). Essentially, these results show that when the proportion of

¹Although in GWAS the number of SNPs examined is often higher than a million, it is commonly assumed that the effective number of independent SNPs is a million, which leads to choosing the threshold of 5×10^{-8} , corresponding to the Bonferroni correction of the target FWER level 0.05, which accounts for the multiplicity of a million features.

features with effects in exactly one study is small, the FDR of the naive approach may be controlled, however its power may be very low compared to other methods designed for controlling FDR for replicability analysis (e.g., Procedure 5.2 of Bogomolov and Heller, 2018). When the proportion of nonreplicable effects increases, the naive approach loses its control of FDR for replicability claims. Interestingly, these simulation results essentially concur with the results shown in Figure 1 of Li et al. (2022) with respect to the naive approach in the model-X framework of Candès et al. (2018), which uses a knockoff-based method rather than BH within each study.²

EXAMPLE 5.3. An investigator faces $n = 2$ studies examining m features, where m is large. Assume that within each study $j \in \{1, 2\}$, the following conditions hold: (1) the p -values are independent, (2) the p -values for features with an effect in study $j \in \{1, 2\}$ have the same cumulative distribution function, denoted by F_j , which satisfies the conditions of Theorem 1 of Genovese and Wasserman (2002) (3) the features with no effect have uniform p -values. Assume that replicability claims are made for features for which the null hypotheses were rejected by BH at level α in both studies. Let m_{00} , m_{11} , m_{01} and m_{10} denote the numbers of features with no effect in both studies, with an effect in both studies, with an effect only in study two, and with an effect only in study one, respectively. Let $\pi_{rs} = m_{rs}/m$, $r, s \in \{0, 1\}$ be the corresponding proportions among all the m features. We consider the setting where $\pi_{11} = 0$, while π_{10} and π_{01} are both positive. According to Theorem 1 of Genovese and Wasserman (2002), for large enough m , applying the BH procedure at level α within each study $j \in \{1, 2\}$ is approximately equivalent to rejecting the hypotheses with p -values less than a fixed threshold u_j^* , which is the solution to $F_j(u) = \beta_j u$, with $\beta_1^* = [1 - \alpha(\pi_{00} + \pi_{01})]/[\alpha(1 - \pi_{00} - \pi_{01})]$, and $\beta_2^* = [1 - \alpha(\pi_{00} + \pi_{10})]/[\alpha(1 - \pi_{00} - \pi_{10})]$. In this setting any replicability claim is false, and the numbers of false replicability claims for features with no effects in both studies, with an effect only in study one, and with an effect only in study two, are approximately $\text{Bin}(m_{00}, u_1^* u_2^*)$, $\text{Bin}(m_{10}, F_1(u_1^*) u_2^*)$, and $\text{Bin}(m_{01}, u_1^* F_2(u_2^*))$ random variables, respectively. Therefore, for replicability the FDR is equivalent to the FWER, and is approximately

$$1 - (1 - u_1^* u_2^*)^{m_{00}} (1 - F_1(u_1^*) u_2^*)^{m_{10}} (1 - u_1^* F_2(u_2^*))^{m_{01}}.$$

²In the simulation for Figure 1 of Li et al. (2022), the naive procedure declares n/n replicability for features with discoveries obtained in all the n studies. This approach does not control FDR for n/n replicability analysis when most conditional associations are inconsistent. It retains FDR control when most conditional associations are consistent, but in this case it is underpowered compared to knockoff-based procedures tailored for replicability analysis with FDR control (described in Procedure 5.7).

For very strong effects, corresponding to p -values which are equal to 0 almost surely, it holds $F_1(u_1^*) = F_2(u_2^*) = 1$, so $u_i^* = 1/\beta_i^*$, for $i = 1, 2$. Obviously, the expression above may be close to 1. For example, for $\alpha = 0.05$, and $(\pi_{00}, \pi_{10}, \pi_{01}) = (0.85, 0.1, 0.05)$, we obtain $u_1^* = 0.0052$, $u_2^* = 0.0026$, so $\text{FDR} \approx 0.995$ for $m = 10,000$ and $\text{FDR} \approx 0.235$ for $m = 500$. For a more sparse setting, where $(\pi_{00}, \pi_{10}, \pi_{01}) = (0.98, 0.01, 0.01)$, we obtain $u_1^* = u_2^* = 0.00053$, so $\text{FDR} \approx 0.102$ for $m = 10,000$ and $\text{FDR} \approx 0.005$ for $m = 500$.

We review below several methods tailored for controlling FWER or FDR for replicability analysis, which we recommend using instead of the naive approach discussed above.

5.2 Targeting Minimal Replicability

The goal of replicability analysis across two independent studies with summary statistics available for all the features has been considered in Bogomolov and Heller (2018), Zhao, Small and Rosenbaum (2018). Their methods are based on two steps. First, the promising features are selected from each study separately, based on the data from that study. Second, the features that are selected from both studies are tested for replicability. The methods are based on the idea that replicable features will seem to be promising in both studies if there is enough power to detect them, so (almost) all replicable features are expected to be selected for testing, while accounting for less features at the testing stage leads to power gain. Indeed, the method targeting FDR control is shown to be typically more powerful than applying the BH procedure on maximum of the two studies' p -values (in case $n = 2$, the latter is the partial conjunction approach of Benjamini, Heller and Yekutieli, 2009). For replicability with FWER^{2/2,D} control at level α , the method is the following.

PROCEDURE 5.1 (Bogomolov and Heller, 2018, Zhao, Small and Rosenbaum, 2018).

Step 1: Apply a selection rule on the set of p -values $\{(p_{i1}, q_{i1}), i = 1, \dots, m\}$ to obtain the set of selected features \mathcal{S}_1 . Similarly, apply a selection rule on the set of p -values $\{(p_{i2}, q_{i2}), i = 1, \dots, m\}$ to obtain the set of selected features \mathcal{S}_2 . If all the alternatives are two sided, define

$$p'_{i1} = \begin{cases} q_{i1} & q_{i2} \leq p_{i2} \\ p_{i1} & q_{i2} > p_{i2} \end{cases}, \quad p'_{i2} = \begin{cases} q_{i2} & q_{i1} \leq p_{i1} \\ p_{i2} & q_{i1} > p_{i1} \end{cases}.$$

If all the alternatives are right sided, define $p'_{i1} = p_{i1}$, and $p'_{i2} = p_{i2}$.

Step 2: Apply a FWER-controlling procedure (e.g. Holm's method, Holm, 1979) on the set $\{p'_{i1}, i \in \mathcal{S}_2\}$ at level $\alpha/2$, and on the set $\{p'_{i2}, i \in \mathcal{S}_1\}$ at level $\alpha/2$. Let \mathcal{R}_j be the set of indices of rejected hypotheses in set \mathcal{S}_j .

Step 3: Declare as replicated the features with indices in the set $\mathcal{R}_1 \cap \mathcal{R}_2$.

For replicability with FDR^{2/2,D} control at level α , Bogomolov and Heller (2018) suggested the following method. Throughout the paper, for any event A , $I(A)$ denotes the indicator of A , and inequality $x \leq y$ for vectors $x, y \in \mathbb{R}^n$ means that the corresponding inequality holds coordinatewise, that is, $x_i \leq y_i$ for $i = 1, \dots, n$.

PROCEDURE 5.2 (Bogomolov and Heller, 2018).

Step 1: Identical to Step 1 in Procedure 5.1.

Step 2: Compute

$$R = \max \left\{ r : \sum_{i \in \mathcal{S}_1 \cap \mathcal{S}_2} I \left[(p'_{i1}, p'_{i2}) \leq \left(\frac{r\alpha}{2|\mathcal{S}_2|}, \frac{r\alpha}{2|\mathcal{S}_1|} \right) \right] = r \right\}.$$

Step 3: Declare minimal replicability for features with indices in the set

$$\mathcal{R} = \left\{ i : (p'_{i1}, p'_{i2}) \leq \left(\frac{R\alpha}{2|\mathcal{S}_2|}, \frac{R\alpha}{2|\mathcal{S}_1|} \right), i \in \mathcal{S}_1 \cap \mathcal{S}_2 \right\}.$$

Rather than fixing the target replicability error rate α in advance, one can compute the adjusted p -value for each feature, which quantifies its evidence for replicability. The adjusted p -value for each feature is the minimal level of FWER or FDR for which the feature is declared as replicated by a given procedure. Therefore, for error control at level α , replicability is established for features with adjusted p -values at most α .

The adjusted p -values for the FWER-controlling Procedure 5.1 depend on the choice of the FWER-controlling procedure used in Step 2. If Bonferroni's method is used in Step 2, then the adjusted p -value for feature $i \in \mathcal{S}_1 \cap \mathcal{S}_2$ based on Procedure 5.1 is

$$(5) \quad p_i^{\text{Bon}} = 2 \max(|\mathcal{S}_2|p'_{i1}, |\mathcal{S}_1|p'_{i2}).$$

The adjusted p -value for feature $i \in \mathcal{S}_1 \cap \mathcal{S}_2$ based on Procedure 5.2 is

$$(6) \quad p_i^{\text{BH}} = \min_{\{j: p_j^{\text{Bon}} \geq p_i^{\text{Bon}}, j \in \mathcal{S}_1 \cap \mathcal{S}_2\}} \frac{p_j^{\text{Bon}}}{\sum_{k \in \mathcal{S}_1 \cap \mathcal{S}_2} I(p_k^{\text{Bon}} \leq p_j^{\text{Bon}})},$$

where p_i^{Bon} for $i \in \mathcal{S}_1 \cap \mathcal{S}_2$ is given in (5).

See Bogomolov and Heller (2018) for theoretical guarantees of these procedures. A straightforward generalization of these procedures for the case where $n \geq 2$ and the goal is identifying features with minimal replicability is as follows. Prior to Step 1, the n studies are divided into two groups, and the right-sided and left-sided p -values within each group are combined for obtaining global null p -values for both directions. Then the same steps are applied, where (p_{i1}, q_{i1}) and (p_{i2}, q_{i2}) are replaced by the corresponding global null p -values for each

group of studies. For both procedures, Step 3 defines the set of features with minimal replicability claims. The theoretical guarantees for the generalized Procedure 5.1 and 5.2 address $\text{FWER}^{2/n,D}$ and $\text{FDR}^{2/n,D}$, respectively. Interestingly, these procedures provide a partial identification of studies with effects for features with minimal replicability claims: for each such feature the global null in each of the two groups of studies is rejected, that is, one can claim that the feature has an effect in at least one study in each group. The lists of global null discoveries for each group come with an $\alpha/2$ -FWER control guarantee for generalized Procedure 5.1, and with an $\alpha/2$ -FDR control guarantee for generalized Procedure 5.2.

Variants of Procedures 5.1 and 5.2, obtained by incorporating plug-in estimators for the proportions of nulls in one study among the selected in the other study, were suggested by Bogomolov and Heller (2018) for gaining power. Similar methods for the case where one study is primary and the other one is a follow-up study, have been suggested in Bogomolov and Heller (2013), Heller, Bogomolov and Benjamini (2014).

5.3 General Replicability Analysis

In case where there are more than two studies, one can address the goal of identifying features with effects in at least r out of n studies, for $r \geq 2$. For $r = 2$, this reduces to minimal replicability, and the higher r is, the stronger is the requirement for claiming replicability. Benjamini and Heller (2008), Benjamini, Heller and Yekutieli (2009) suggested addressing this goal by multiple testing of r/n no replicability null hypotheses. Specifically, for $\text{FDR}^{r/n}$ control, they suggested applying the BH procedure on r/n PC p -values (see (3) for an example of a PC p -value). In some applications it may be unclear what is the appropriate choice of r in order to claim replicability. For example, while one investigator may claim that association of a genotype with a disease in at least two out of six studies is enough for claiming replicability, another investigator may claim that stronger evidence is required, for example, association with a disease in at least half of the studies (i.e., fixing $r = 3$ rather than $r = 2$). In order to provide the highest (nonzero) lower bound on the number of studies with effect for each feature, while controlling for noncoverage, Benjamini, Heller and Yekutieli (2009) suggested the following procedure for right-sided alternatives.

PROCEDURE 5.3 (Benjamini, Heller and Yekutieli, 2009).

Step 1: For each feature $i = 1, \dots, m$, compute the global null p -value $p_i^{1/n}$.

Step 2: Apply the BH procedure at level α on $\{p_i^{1/n}, i = 1, \dots, m\}$. Select the features for which rejections are made. Let \mathcal{S} be the set of indices of selected features.

Step 3: For each feature $i \in \mathcal{S}$, test the no-replicability null hypotheses in order at level $|\mathcal{S}|\alpha/m$, and define

$$l_i = \max\{l \in \{0, \dots, n\} : p_i^{r/n} \leq |\mathcal{S}|\alpha/m \text{ for } r = 0, \dots, l\}$$

to be the lower bound for the number of studies where feature i has an effect.

This procedure is a generalization of the method in (4) for $m = 1$. Benjamini, Heller and Yekutieli (2009) proved that under independence, Procedure 5.3 guarantees that the expected proportion of features with lower bounds exceeding the true number of studies where they have effects, out of all the selected features, is upper bounded by α . We refer to the error measure above as the false coverage rate (FCR, Benjamini and Yekutieli, 2005) for the selected features. This procedure was generalized in Bogomolov (2023), allowing for different rules for selecting the features for which lower bounds will be obtained. For example, the selected set \mathcal{S} can be obtained by applying a multiple testing procedure on p -values of a given study, defining the discoveries in that study. Suppose a researcher applies Bonferroni for obtaining discoveries in his or her study, referred to as study one, then the selected set of features of interest regarding their replicability extent can be $\mathcal{S} = \{i : p_{i1} \leq \alpha/m\}$. The same can be done if the researcher's study is primary and the other studies are follow-up studies, examining only the features with significant effects in the primary study. The generalized procedure is the following.

PROCEDURE 5.4 (Bogomolov, 2023).

Step 1: Apply a certain selection rule on all the p -values, for obtaining the selected set of features \mathcal{S} .

Step 2: Apply Step 3 of Procedure 5.3 with respect to the selected set of features \mathcal{S} .

It was shown in Bogomolov (2023) that the generalized procedure controls the FCR for the selected features under certain forms of positive dependencies within the studies, and under lenient conditions on the selection rule used in Step 1 of Procedure 5.4. In addition, the procedure can be conservatively adjusted for allowing arbitrary dependencies within the studies and arbitrary selection rules.

A drawback of methods which are based on applying a multiple testing procedure on PC p -values is their possibly low power. The power loss occurs due to the fact that $H_i^{r/n}$ is a composite null hypothesis. So when $H_i^{r/n}$ is true, the PC p -value gives an exact test only under the least favorable configuration, where the number of false null hypotheses for feature i is $r - 1$, their corresponding p -values are all equal to 0, and the remaining $n - r + 1$ p -values have a uniform distribution. Under other configurations, the PC p -value gives a conservative test. For example, a very conservative, yet typically prevalent, configuration among the features is that the global null is true.

For feature i with $n_i^+ = n_i^- = 0$, that is, for which the global null is true, $\mathbb{P}(p_i^{r/n} \leq x) \leq x^n$.

In order to overcome the power loss induced by the conservativeness of the PC p -values, Wang et al. (2022) developed the AdaFilter methods for testing the family of PC hypotheses $H_1^{r/n,+}, \dots, H_m^{r/n,+}$ for any given $r \in \{2, \dots, n\}$. Essentially, rather than adjusting for the multiplicity of all the m features, these methods adjust only for the number of features with evidence for effects in at least $r - 1$ studies, thereby gaining power. Specifics follow. Let $p_{i(1)} \leq p_{i(2)} \dots \leq p_{i(n)}$ be the sorted right-sided p -values for feature $i \in \{1, \dots, m\}$. For testing the family $H_1^{r/n,+}, \dots, H_m^{r/n,+}$, the method is based on assigning the filtering statistic $F_i = (n - r + 1)p_{i(r-1)}$ and the selection statistic $S_i = (n - r + 1)p_{i(r)}$ for feature $i \in \{1, \dots, m\}$. The AdaFilter Bonferroni and AdaFilter BH, which are a version of AdaFilter targeting FWER control and FDR control, respectively, for the family $H_1^{r/n,+}, \dots, H_m^{r/n,+}$, are as follows.

PROCEDURE 5.5 (Wang et al., 2022, AdaFilter-Bonferroni).

Step 1: Compute

$$\gamma_0^{\text{Bon}} = \sup \left\{ \gamma \in [0, \alpha] \mid \gamma \sum_{i=1}^m I(F_i < \gamma) \leq \alpha \right\}.$$

Step 2: Reject $H_i^{r/n}$ if $S_i < \gamma_0^{\text{Bon}}$, for $i \in \{1, \dots, m\}$.

PROCEDURE 5.6 (Wang et al., 2022, AdaFilter-BH).

Step 1: Compute

$$\gamma_0^{\text{BH}} = \sup \left\{ \gamma \in [0, \alpha] \mid \frac{\gamma \sum_{i=1}^m I(F_i < \gamma)}{\max(\sum_{i=1}^m I(S_i < \gamma), 1)} \leq \alpha \right\}.$$

Step 2: Reject $H_i^{r/n}$ if $S_i < \gamma_0^{\text{BH}}$, for $i \in \{1, \dots, m\}$.

For theoretical guarantees of the procedures, see Wang et al. (2022). The AdaFilter Bonferroni adjusted p -value for $H_i^{r/n}$ is

$$(7) \quad p_i^{\text{Bon}} = S_i \sum_{k=1}^m I(F_k \leq S_i)$$

and AdaFilter BH adjusted p -value for $H_i^{r/n}$ is

$$p_i^{\text{BH}} = \min_{\{j: p_j^{\text{Bon}} \geq p_i^{\text{Bon}}\}} \frac{p_j^{\text{Bon}}}{\sum_{k=1}^m I(p_k^{\text{Bon}} \leq p_j^{\text{Bon}})},$$

where $\{p_i^{\text{Bon}}, i = 1, \dots, m\}$ are given in (7). For any level $\alpha \in (0, 1]$, the sets of rejections of AdaFilter Bonferroni and AdaFilter BH are $\{i : p_i^{\text{Bon}} < \alpha\}$ and $\{i : p_i^{\text{BH}} < \alpha\}$, respectively, as shown in Proposition 3.4 of Wang et al. (2022).

REMARK 5.1. For two-sided hypotheses and directional r/n replicability for a certain fixed $r \in \{2, \dots, n\}$, Wang et al. (2022) suggested using their methods as follows. Apply the given method with parameter r twice, separately on $\{p_{ij}, i = 1, \dots, m, j = 1, \dots, n\}$ and on $\{q_{ij}, i = 1, \dots, m, j = 1, \dots, n\}$, each time at level $\alpha/2$. Let \mathcal{R}^R and \mathcal{R}^L be the corresponding sets of discoveries. Claim r/n replicability of positive effect for features with indices in \mathcal{R}^R , and r/n replicability of negative effect for features with indices in \mathcal{R}^L . This strategy guarantees control of directional FWER/FDR for replicability analysis as long as the corresponding method applied on one-sided p -values guarantees FWER/FDR for nondirectional replicability. The approach of Benjamini, Heller and Yekutieli (2009), which amounts to applying BH on r/n PC p -values, can be used similarly for directional r/n replicability analysis.

All the methods above address the general case where one has a p -value for each null hypothesis H_{ij} , which requires knowing the distribution of test statistics under the null. Li et al. (2022) addressed the goal of replicability in a different setting, which does not require computing p -values. In their setting one has observations (X, Y) consisting of m variables $X \in \mathcal{X}^m$ and an outcome $Y \in \mathcal{Y}$ sampled from n environments. For variable i and environment j , the null hypothesis is defined as follows:

$$H_{ij} : Y^j \perp\!\!\!\perp X_i^j \mid X_{-i}^j,$$

where Y^j , X_i^j and X_{-i}^j , respectively, denote the outcome, the i th variable, and the vector of all variables except X_i in environment j . In words, H_{ij} states that the outcome is independent of the i th explanatory variable conditional on all other explanatory variables, when considering the distribution of the data in the j th environment. The alternative is complementary to the null, so it has no one-sided and two-sided variants. Therefore, directional replicability is not defined in this case. Testing the r/n no-replicability null hypotheses for all m variables addresses the goal of identifying conditional associations which hold in at least r environments. The method of Li et al. (2022) controls FDR for the family of r/n no-replicability null hypotheses. It relies on the model-X approach of Candès et al. (2018), which does not require the model for conditional distribution of the outcome, but requires that the distribution of X within environment j is known, at least approximately. The latter assumption is realistic in genetic studies, because reliable knowledge is available regarding the distribution of the genotypes (Sesia, Sabatti and Candès, 2019). The method relies on multi-environment knockoff statistics defined below, which generalize the knockoff statistics of Candès et al. (2018).

DEFINITION 5.1 (Li et al., 2022). We say that $W \in \mathbb{R}^{m \times n}$ are multi-environment knockoff statistics if W has the same distribution as $W \cdot \epsilon$, where \cdot indicates element-wise multiplication, and ϵ is a random $m \times n$ matrix with independent entries in the set $\{1, -1\}$, satisfying $\epsilon_{ij} = \pm 1$ with probability $1/2$ if H_{ij} is true, and $\epsilon_{ij} = 1$ otherwise, for $i \in \{1, \dots, m\}$, $j \in \{1, \dots, n\}$.

See Li et al. (2022) for methods which can be used for constructing the multi-environment knockoff statistics. Given a fixed $r \in \{1, \dots, n\}$, the method of Li et al. (2022) for controlling the FDR for a family of r/n no-replicability null hypotheses is given below.

PROCEDURE 5.7 (Li et al., 2022).

Step 1: Compute the multi-environment knockoff statistics W . For each $i \in \{1, \dots, m\}$, let $t_i^- = \sum_{j=1}^n I(W_{ij} < 0)$, and $t_i^0 = \sum_{j=1}^n I(W_{ij} = 0)$.

Step 2: For each $i \in \{1, \dots, m\}$, compute

$$(8) \quad p_i^{r/n} = \mathbb{P}(T_i \leq t_i^- - 1) + U_i \mathbb{P}(T_i = t_i^-),$$

where $T_i \sim \text{Bin}(\max\{n - r + 1 - t_i^0, 0\}, 1/2)$, and U_i is a $U(0, 1)$ random variable which is independent of everything else.

Step 3: For each $i \in \{1, \dots, m\}$, consider the set of absolute values of knockoff statistics for variable i in all the environments, $\{|W_{ij}|, j = 1, \dots, n\}$. Obtain $|W_i^r|$, defined as the product of the r largest values in this set. See Li et al. (2022) for a more general definition of $|W_i^r|$.

Step 4: Filter the p -values obtained in Step 2 in decreasing order of $|W_i^r|$ using the level- α selective seqStep+ procedure of Barber and Candès (2015) with a certain choice of the tuning parameter c , and obtain the rejected r/n no-replicability null hypotheses. Specifically, compute

$$\hat{w} = \min \left\{ w : \frac{1 + |\{i : |W_i^r| \geq w, p_i^{r/n} > c\}|}{\max(|\{i : |W_i^r| \geq w, p_i^{r/n} \leq c\}|, 1)} \leq \frac{1 - c}{c} \alpha \right\},$$

and reject the r/n no-replicability null hypotheses for variables with indices in the set $\{i : |W_i^r| \geq \hat{w}, p_i^{r/n} \leq c\}$.

See Li et al. (2022) for another variant of the above procedure and its theoretical guarantees. A simpler method is suggested for the case where $r = n$, that is, when the goal is to find variables which are conditionally associated with the outcome in all the environments. Specifically, in Procedure 5.7 the definition in (8) for the PC p -value in Step 2 is replaced by the following definition: $p_i^{n/n} = 1/2$ if $\min\{\text{sign}(W_{ij})\}_{j=1}^n = +1$, and $p_i^{n/n} = 1$ otherwise. The next steps are obtained by substituting $r = n$ in Steps 3 and 4, and choosing $c = 1/2$ in Step 4.

TABLE 2

Summary for the replicability analysis procedures addressed in this section. The table addresses the generalized versions of Procedures 5.1, 5.2, which can analyze $n \geq 2$ studies. The first, second, and third column give the reference to the procedure, the possible values of r for which r/n replicability can be assessed, and the target error rate addressing erroneous replicability claims, respectively. The fourth column gives the summary statistics required as input

Proc.	r	Target error rate	Summary statistics
5.1	$r = 2$	$\text{FWER}^{2/n, D}$	p -values
5.2	$r = 2$	$\text{FDR}^{2/n, D}$	
5.3, 5.4	r is not fixed in advance, adaptive	FCR on the number of studies with signal	p -values
5.5	$r \geq 2$	$\text{FWER}^{r/n}$	p -values
5.6	$r \geq 2$	$\text{FDR}^{r/n}$	
5.7	$r \geq 2$	$\text{FDR}^{r/n}$	multi-environment knockoff statistics

5.4 Broad Guidelines and Future Directions

The properties of the replicability analysis methods addressed in this section are summarized in Table 2. Let us address the comparison of Procedures 5.1, 5.2 with Procedures 5.5, 5.6 for the case where the goal is minimal replicability, that is, $r = 2$. All these procedures test for replicability only the features that are selected as promising for assessing minimal replicability, however the selection mechanisms are different. While Procedures 5.1, 5.2 allow the researcher to choose the selection rules within each study, Procedures 5.5, 5.6 are based on fixed selection rules. In addition, Procedure 5.1 allows using any FWER-controlling procedure in Step 2, on selected one-sided p -values, while Procedure 5.5 does not have this flexibility. The generalized versions of Procedures 5.1 and 5.2 for the case $n > 2$ have an additional advantage over Procedures 5.5, 5.6: the former procedures allow partial identification of studies where the discovered features have effects, with an appropriate error rate control. The disadvantage of generalized Procedures 5.1, 5.2 is that their results depend on the division of the n studies into two groups, which can be arbitrary. These procedures aim to identify features with effects in both groups of studies. Thus, for the case where $n > 2$, generalized Procedures 5.1, 5.2 may be less powerful than Procedures 5.5 and 5.6, respectively, for the case where many features have two or more effects in only one group of studies. However, it may be possible to overcome this limitation by considering adaptively chosen groups. Moreover, it may be of great interest to further generalize Procedures 5.1, 5.2 for assessing r/n replicability across n studies for any $r \geq 2$, thus carrying over to $r > 2$ the flex-

ibility and potential power gain of these procedures observed for $r = 2$.

Procedures 5.3, 5.4 are different from the other procedures, because rather than requiring to fix r in advance, these procedures use the data for identifying what is the strongest replicability claim that can be made for a given feature. The replicability claims are made in terms of lower bounds for the number of studies with effects for the selected features, and the error rate controlled is FCR for the selected features. As shown in Table 2, the other methods control for erroneous r/n replicability claims for a fixed value of r . If the adaptivity of Procedures 5.3, 5.4 is attractive, then these procedures should be chosen for replicability analysis.

Procedures 5.3, 5.4, as well as Procedures 5.5, 5.6 are given above for the case of one-sided hypotheses. Remark 5.1 suggests a way to perform directional r/n replicability using Procedures 5.5, 5.6, or using the partial conjunction approach of Benjamini, Heller and Yekutieli (2009), which amounts to applying the BH procedure on r/n PC p -values. The latter procedure is equivalent to that in Definition 7 of Benjamini and Yekutieli (2005) for $n = 1$. Benjamini and Yekutieli (2005) show that the procedure that applies the BH procedure to the two-sided p -values and further declares the sign (their Definition 6) also controls the directional FDR, and is at least as powerful as their procedure in Definition 7. This suggests that the directional inference in Remark 5.1 may be improved for directional FDR control on false replicability claims. An additional potential improvement may address overcoming the conservatism of the PC p -values used in these procedures.

Procedure 5.7 addresses a specific setting where one is given data for m explanatory variables and a certain outcome in each of n environments (studies), and for each explanatory variable one is interested in testing the null hypothesis of conditional independence between this variable and the outcome, given the other variables. This setting does not require specifying the model of dependence between the outcome and the explanatory variables. In contrast to other procedures in Table 2, Procedure 5.7 is based on knockoff statistics and does not require the ability to compute p -values for the tested hypotheses.

All the procedures assume that the studies are independent. This assumption does not always hold: it is obviously violated for case-control studies if each study has its own disease cohort, but all are compared to the same control cohort; or if each study examines a different phenotype, but the same subject may be in several studies. In these settings, it is therefore necessary to design novel powerful replicability analysis procedures.

6. THE EMPIRICAL BAYES APPROACH FOR REPLICABILITY IN HIGH DIMENSIONAL STUDIES

6.1 Replicability Analysis with the Local False Discovery Rate

The two group model (Efron et al., 2001, Efron, 2010) is popular for analyzing high dimensional studies, where many features are simultaneously tested for association with an outcome. It has been generalized for targeting the discovery of features that are associated with the outcome in at least two studies in Chung et al. (2014), Heller and Yekutieli (2014).

The generalization of the two group model for n studies, each examining a large number m of features, is as follows. Using GWAS as our illustrative example, the indicator of the hypothesis state for feature $i \in \{1, \dots, m\}$ in study $j \in \{1, \dots, n\}$, h_{ij} , is zero if there is no association, -1 if the association is negative, and one if the association is positive. The z -score $z_{ij} = \Phi^{-1}(p_{ij})$ has density $f_{j,1}(\cdot)$ if $h_{ij} = 1$, $f_{j,0}(\cdot)$ if $h_{ij} = 0$, and $f_{j,-1}(\cdot)$ if $h_{ij} = -1$. The null density $f_{j,0}(\cdot)$ is usually assumed to be the standard normal distribution (following the assumption that p_{ij} has a uniform null distribution). Each feature has probability $\pi(\mathbf{h})$ of having the vector of hypotheses states $\mathbf{h} \in \{-1, 0, 1\}^n$, where $\{-1, 0, 1\}^n$ denotes the set of 3^n possible vectors of length n with coordinates in $\{-1, 0, 1\}$. The z -scores for each feature are assumed to be independent across studies given the vector of hypotheses states, so the joint density of (z_{i1}, \dots, z_{in}) for feature i given $\mathbf{h}_i = (h_{i1}, \dots, h_{in})$ is $\prod_{j=1}^n f(z_{ij}|h_{ij})$.

The statistic that plays a central role for inference on the hypotheses states for each feature is the *local false discovery rate*, which is the probability that the state vector is a null state vector given the feature's z -scores. For replicability analysis, we wish to test the r/n no-replicability null hypothesis defined in (2), so the null state vector is a binary vector with at most $r - 1$ positive entries and at most $r - 1$ negative entries. We use the notation $\sum_{\mathbf{h} \in H_i^{r/n}}$ to denote the summation over all state vectors \mathbf{h} that are in the null space of the r/n no-replicability null hypothesis. The local false discovery rate for feature i is

$$\begin{aligned} T_i &= \sum_{\mathbf{h}_0 \in H_i^{r/n}} \mathbb{P}(\mathbf{h}_i = \mathbf{h}_0 | z_{i1}, \dots, z_{in}) \\ &= \frac{\sum_{\mathbf{h}_0 \in H_i^{r/n}} \pi(\mathbf{h}_0) \prod_{j=1}^n f(z_{ij}|h_{0j})}{\sum_{\mathbf{h} \in \{-1, 0, 1\}^n} \pi(\mathbf{h}) \prod_{j=1}^n f(z_{ij}|h_j)}. \end{aligned}$$

With these local false discovery rates, we describe next the steps for inference on the family of the r/n no-replicability null hypotheses, for a target α level Bayes FDR. The Bayes FDR is the probability that the hypothesis state vector is in the null space, given that the vector of z -scores is in the rejection region. For independent z -scores within each study, the Bayes FDR coincides with

the positive FDR, which is the expected false discovery proportion given that at least one discovery was made (Storey, 2003).³ The steps for inference are as developed for a single study in Sun and Cai (2007).

PROCEDURE 6.1 (Heller and Yekutieli, 2014).

Step 1: Compute for each $i \in \{1, \dots, m\}$ the local false discovery rate T_i .

Step 2: Sort the T_i 's from smallest to largest, so $T_{(1)} \leq \dots \leq T_{(m)}$.

Step 3: Reject the r/n no-replicability null hypothesis for the R features with the smallest T_i 's, where

$$R = \max \left\{ k : \frac{\sum_{i=1}^k T_{(i)}}{k} \leq \alpha \right\}.$$

In practice, the parameters of the model are unknown. The empirical Bayes approach uses the data in order to estimate the parameters that are needed in order to compute the T_i 's. Methods for estimation of the mixture components, using the EM algorithm, have been suggested, for example, in Chung et al. (2014), Heller and Yekutieli (2014), Amar et al. (2018). These estimation methods assume that the test statistics (such as p -values) are independent in each study, which is obviously not true for high dimensional studies. However, if the dependence is local (i.e., each feature's test statistic is associated only with a small number of other test statistics), the signal is not too sparse, and the studies are high dimensional, the estimation method (which uses the composite likelihood rather than the true likelihood) is fairly robust to deviations from independence. See Xiang, Zhao and Cai (2019) for another estimation method for the model parameters for $n = 2$; See Sun and Wei (2011), Wang and Zhu (2019) and references within for methods that take the dependency between hypothesis states into account.

6.2 The (Bayes) FDR of a Frequentist Approach

The no-replicability null hypothesis is a composite hypothesis. The frequentist approach guards against the worst (or least favorable) setting, while the empirical Bayes approach discussed in Section 6.1 relies on the ability to estimate the prior probabilities for the different vectors of hypotheses states.

The methods in Section 5.3 do not require modeling of the z -values (or p -values). Moreover, they provide an FDR (or FWER) control guarantee for any number of features m . This is in contrast with the inference based on the (estimated) local false discovery rate, which requires a model for the z -values, and can provide only asymptotic FDR control (Heller and Yekutieli, 2014).

³The FDR is the positive FDR times the probability of making at least one discovery, so the FDR and positive FDR coincide in settings where always at least one discovery is made (Storey, 2003).

To help clarify the expected power advantage of the empirical Bayes approach of Section 6.1 over the frequentist approaches in Section 5.3, when the model for the z -values approximates well the data generation mechanism, we provide the following analysis for the frequentist approach of applying BH on the PC p -values (although all procedures in Section 5 suffer from being conservative compared to the approach of Section 6.1, showing their conservativeness is more involved so it is outside the scope of this paper). Consider a procedure that thresholds the PC p -values, so that only features with $p_i^{r/n} \leq t$ are rejected. The threshold t may be selected to be the largest that satisfies $\frac{t}{\mathbb{P}(p_i^{r/n} \leq t)} \leq \alpha$, where

$$\mathbb{P}(p_i^{r/n} \leq t) = \sum_{\mathbf{h} \in \{-1, 0, 1\}^n} \mathbb{P}(p_i^{r/n} \leq t | \mathbf{h}_i = \mathbf{h}) \pi(\mathbf{h})$$

is the marginal cumulative distribution function of the PC p -value. Using this threshold provides a procedure that is a close approximation to the BH procedure on the PC p -values, which selects the largest t so that $\frac{t}{|\{i: p_i^{r/n} \leq t\}|/m} \leq \alpha$. The positive FDR of this procedure for the family of r/n no-replicability nulls can be expressed in terms of any $i \in \{1, \dots, m\}$ as follows:

$$\begin{aligned} & \frac{\sum_{\mathbf{h} \in H_i^{r/n}} \mathbb{P}(p_i^{r/n} \leq t | \mathbf{h}_i = \mathbf{h}) \pi(\mathbf{h})}{\mathbb{P}(p_i^{r/n} \leq t)} \\ & \approx \frac{\sum_{\mathbf{h} \in H_i^{r/n}} \mathbb{P}(p_i^{r/n} \leq t | \mathbf{h}_i = \mathbf{h}) \pi(\mathbf{h})}{|\{k : p_k^{r/n} \leq t\}|/m}. \end{aligned}$$

The numerator is clearly at most t , but it may also be much smaller. For example, if the probability that the global null is true is 0.95 and that the no-replicability null is false is 0.05 (i.e., we can only observe features with $\mathbf{h} = \mathbf{0}$ or features with $\mathbf{h} \notin H_i^{r/n}$), then the numerator is $0.95 \times \mathbb{P}(p_i^{r/n} \leq t | \mathbf{h}_i = \mathbf{0})$, which is much smaller than t for $r \geq 2$ (e.g., for $r = 2$ and $n = 2$ the numerator is at most $0.95 \times t^2$). Therefore, by choosing t so that $\frac{t}{|\{k: p_k^{r/n} \leq t\}|/m} \approx \alpha$, the positive FDR is in fact much smaller than α . On the other hand, with Procedure 6.1 the level is approximately α so it is not overly conservative and it may provide many more replicable findings. See Heller and Yekutieli (2014) for more comparisons of the empirical Bayes approach versus a frequentist approach for replicability.

6.3 Future Directions

The generalization of the two group model to $n > 1$ studies examining the same m features is useful for the joint analysis of the m features as long as the individual feature's mixture distribution can be estimated well. This is not possible if the unknown dependence within each

high dimensional study is long range dependence (even when m is in the order of $10^6 - 10^9$). It is not currently known how to develop a procedure for general (unknown) dependence structures within each study.

Xie et al. (2011), Heller and Rosset (2021) considered the following generalization of the two group model for a single high dimensional study $j \in \{1, \dots, n\}$ with known (or well approximated) dependence between the z -scores: the vector (z_{1j}, \dots, z_{mj}) is sampled from the joint conditional distribution given (h_{1j}, \dots, h_{mj}) , where the h_{ij} 's are independently and identically distributed as in the two group model. They showed that in order to discover the features with a false null hypothesis, the optimal statistic for each feature is the probability that the null hypothesis is true given the entire vector of z -scores, $\mathbb{P}(h_{ij} = 0 | z_{1j}, \dots, z_{mj})$. The potential gain in power in using this statistic rather than the marginal statistic $\mathbb{P}(h_{ij} = 0 | z_{ij})$ can be very large. We thus expect the power advantage to be large when incorporating the within study local dependence in a joint analysis of n studies. Specifically, for replicability analysis, it can be useful to incorporate known or approximate local dependence within each study, to boost discovery power of replicable signals.

Xiang, Zhao and Cai (2019) considered for $n = 2$ several inferential goals, each with a predefined level of asymptotic FDR control: classification of features to those that have signal in both studies, or only in the first study, or only in the second study. Their paper inspires additional interesting formulations for replicability analysis. For example, for $n = 3$ studies, consider simultaneously the classification of features with signal (in the same direction) in all three studies, as well as those with signal (in the same direction) in exactly two studies.

7. THE CASE OF A SINGLE STUDY

Splitting the sample at random into two halves is not useful for assessing replicability, since it results in having two samples, each with reduced sample size, for testing the same hypotheses H_i , $i = 1, \dots, m$.

Importantly, splitting the sample according to an observed binary covariate is helpful for replicability. For example, examination of the effect of high fish consumption in a subgroup that lives where fish are at low cost (near the coast) and in a subgroup that lives where fish are at high cost (say in midland, so the reason for consumption may be health benefits) may be possible from a publicly available database (Zhao, Small and Rosenbaum, 2018). If an outcome (e.g., heart disease) is negatively associated with increased fish consumption in both subgroups, the evidence that the discovered association is due to fish consumption is more convincing and less susceptible to bias. In this example, the covariate indicates whether the subject lives near the sea or in midland, so it can be used

to split the sample into two subgroups in which the effect of high fish consumption is examined. Although the same family of null hypotheses are considered in each subgroup (as in splitting at random), the hypothesis states need not be the same in the two resulting subgroups (unlike with random sample splitting). For example, it may be that there is no association between fish consumption and heart disease for the subgroup that lives near the sea, but there is a negative association for the subgroup that lives in midland due to say unobserved confounding with exercise: in midland the reason for high fish consumption is health benefits, and those that are health conscious exercise more, and exercise reduces the risk of heart disease.

When splitting according to an observed binary covariate, the aim is to identify features for which the null hypothesis is false in both subgroups, and the procedures applied should take care not to declare as replicated too many features $i \in \{1, \dots, m\}$ for which the null hypothesis is false in exactly one of the subgroups. The procedures in Section 5.2 may therefore be appropriate for this purpose, see Roy et al. (2022) for an example application. If the sample is split instead according to an observed covariate with $n \geq 2$ categories, the procedures in Section 5.3 may be useful.

The procedures in Section 5.2 and Section 5.3 may also be useful for causal inference in observational studies using evidence factors (Karmakar and Small, 2020). In order to establish causality in observational studies, it is useful to decompose the test of no treatment effect into evidence factors, which are separate pieces of information that are affected by different biases and are statistically nearly independent (Rosenbaum, 2010). For example, in an observational study with a matched pairs design, where each pair has one exposed and one unexposed individual, and the level of exposure among the exposed is measured, then there are two pieces of evidence towards the (say positive) effect of exposure: evidence that the exposed have higher levels of the outcome than their unexposed matched pairs; evidence among the exposed of a positive association between exposure level and the outcome. See Karmakar and Small (2020) for examples with more than two evidence factors. Suppose there are n such pieces. Karmakar and Small (2020) suggest testing whether there is evidence that at least r of the n pieces of information reject the null hypothesis of no treatment effect. This is the test of the no-replicability null hypothesis, and the PC p -value is used to evaluate this composite null.

8. DISCUSSION

Wang et al. (2022) point out that when complicated large-scale experiments or observational studies are carried out, they can be susceptible to many sources of bias. Since many complicated large-scale studies are carried out nowadays, it is crucial to have statistical methods that

objectively evaluate the consistency of findings across the studies, while properly accounting for the fact that the parameters of interest may differ across studies. More specifically, the consistency (or replicability) across studies has to be evaluated by ruling out that the finding is particular to a single study (or more generally, to $r - 1$ studies). In this review, we provided various relevant procedures that may be useful.

In collaborative research, it is possible to go one step further in terms of the flexibility of the procedures towards establishing replicability. Roy et al. (2022) suggests for $n = 2$ that two separate teams will engage in the design of the multiple testing, where each team has access only to a single study, and by examining that study it designs the multiple testing of the other study. As long as the teams act independently, they have complete freedom in design decisions. So a more flexible version of Procedure 5.1 will allow each team not only to select hypotheses and their direction of testing, but also to select the test statistic for each hypothesis, possible weights for the different hypotheses (which may reflect the prior belief in the probability that they are nonnull in the other study), and the FWER controlling procedure to apply in the other study. For $n > 2$, each team may have access to a group of studies, and the additional flexibility of choosing how to combine the evidence from the group of studies. If there are more than two teams, the flexibility is further increased. Procedure 5.1 has to be generalized, and we expect it to be a powerful and useful tool for tailoring replicability analyses to specific applications.

The suggested procedures in this review were motivated by the need for replicable results that are less susceptible to biases. However, these methods are also useful in additional settings. For example, in mediation analysis (Sampson et al., 2018, Djordjilović et al., 2019, Liu et al., 2022), where the aim is to identify the mediators that are associated both with the exposure and with the outcome. As another example, in a group analysis of size n , where multiple features are examined for each subject, and the goal is to find the features with signal in at least r/n subjects. Specifically, consider the neuroimaging application addressed by Heller et al. (2007), where the aim is to identify the brain regions that are consistently activated in a group of subjects.

Increasingly, researchers provide access to the data used in their studies, at least in aggregate form. For example, in GWAS, summary statistics for each genotype are available in various consortia, and p -values can be easily computed from the summary level data. So the suggested p -value based replicability analyses can include studies from these consortia. We hope the facilitation of access to data will increase the practice of applying a replicability analysis to establish scientific findings.

ACKNOWLEDGMENTS

The authors would like to thank the Guest Editors of the special issue of Statistical Science on Reproducibility and Replicability and the referee for their excellent suggestions. The authors are also grateful to Yoav Benjamini, Saharon Rosset, and Daniel Yekutieli for useful discussions that helped shape the paper.

FUNDING

The first author was supported in part by Israeli Science Foundation grant no. 1726/19. The second author was supported in part by Israeli Science Foundation grant no. 2180/20.

SUPPLEMENTARY MATERIAL

Additional Examples for Section 5.1 (DOI: [10.1214/23-STS892SUPP](https://doi.org/10.1214/23-STS892SUPP); .pdf). The SM provides additional examples supporting the key observations of Section 5.1. These examples assess the FWER control of the naive approach of Section 5.1 in different settings.

REFERENCES

- AMAR, D., VIZEL, A., LEVY, C. and SHAMIR, R. (2018). ADEPTUS: A discovery tool for disease prediction, enrichment and network analysis based on profiles from many diseases. *Bioinformatics* **34** 1959–1961.
- ANDREASSEN, O. A., THOMPSON, W. K., SCHORK, A. J., RIPKE, S., MATTINGSDAL, M., KELSÖE, J. R., KENDLER, K. S., O'DONOVAN, M. C., RUJESCU, D. et al. (2013). Improved detection of common variants associated with schizophrenia and bipolar disorder using pleiotropy-informed conditional false discovery rate. *PLoS Genet.* **9**.
- BARBER, R. F. and CANDÈS, E. J. (2015). Controlling the false discovery rate via knockoffs. *Ann. Statist.* **43** 2055–2085. MR3375876 <https://doi.org/10.1214/15-AOS1337>
- BENJAMIN, D., BEGER, J., JOHANNESSEN, M. et al. (2018). Redefine statistical significance. *Nat. Hum. Behav.* **2** 6–10.
- BENJAMINI, Y. and HELLER, R. (2008). Screening for partial conjunction hypotheses. *Biometrics* **64** 1215–1222. MR2522270 <https://doi.org/10.1111/j.1541-0420.2007.00984.x>
- BENJAMINI, Y., HELLER, R. and YEKUTIELI, D. (2009). Selective inference in complex research. *Philos. Trans. R. Soc. Lond. Ser. A Math. Phys. Eng. Sci.* **367** 4255–4271. MR2546387 <https://doi.org/10.1098/rsta.2009.0127>
- BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. MR1325392
- BENJAMINI, Y. and YEKUTIELI, D. (2005). False discovery rate-adjusted multiple confidence intervals for selected parameters. *J. Amer. Statist. Assoc.* **100** 71–93. MR2156820 <https://doi.org/10.1198/016214504000001907>
- BLANCHARD, G., NEUVIAL, P. and ROQUAIN, E. (2020). Post hoc confidence bounds on false positives using reference families. *Ann. Statist.* **48** 1281–1303. MR4124323 <https://doi.org/10.1214/19-AOS1847>
- BOGOMOLOV, M. (2023). Testing partial conjunction hypotheses under dependency, with applications to meta-analysis. *Electron. J. Stat.* **17** 102–155. MR4533743 <https://doi.org/10.1214/22-ejs2100>

- BOGOMOLOV, M. and HELLER, R. (2013). Discovering findings that replicate from a primary study of high dimension to a follow-up study. *J. Amer. Statist. Assoc.* **108** 1480–1492. MR3174723 <https://doi.org/10.1080/01621459.2013.829002>
- BOGOMOLOV, M. and HELLER, R. (2018). Assessing replicability of findings across two studies of multiple features. *Biometrika* **105** 505–516. MR3842881 <https://doi.org/10.1093/biomet/asy029>
- BOGOMOLOV, M. and HELLER, R. (2023). Supplement to “Replicability across multiple studies.” <https://doi.org/10.1214/23-STS892SUPP>
- CANDÈS, E., FAN, Y., JANSON, L. and LV, J. (2018). Panning for gold: ‘model-X’ knockoffs for high dimensional controlled variable selection. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **80** 551–577. MR3798878 <https://doi.org/10.1111/rssb.12265>
- CHUNG, D., YANG, C., LI, C., GELERNTER, J. and ZHAO, H. (2014). GPA: A statistical approach to prioritizing GWAS results by integrating pleiotropy and annotation. *PLoS Genet.* **10** e1004787.
- DJORDJILOVIĆ, V., PAGE, C. M., GRAN, J. M., NØST, T. H., SANDANGER, T. M., VEIERØD, M. B. and THORESEN, M. (2019). Global test for high-dimensional mediation: Testing groups of potential mediators. *Stat. Med.* **38** 3346–3360. MR3979813 <https://doi.org/10.1002/sim.8199>
- EFRON, B. (2010). *Large-Scale Inference: Empirical Bayes Methods for Estimation, Testing, and Prediction*. Institute of Mathematical Statistics (IMS) Monographs 1. Cambridge Univ. Press, Cambridge. MR2724758 <https://doi.org/10.1017/CBO9780511761362>
- EFRON, B., TIBSHIRANI, R., STOREY, J. D. and TUSHER, V. (2001). Empirical Bayes analysis of a microarray experiment. *J. Amer. Statist. Assoc.* **96** 1151–1160. MR1946571 <https://doi.org/10.1198/016214501753382129>
- FISHER, R. A. (1934). *Statistical Methods for Research Workers*, 5th ed.
- FITHIAN, W., SUN, D. and TAYLOR, J. (2017). Optimal inference after model selection. Preprint. Available at [arXiv:1410.2597](https://arxiv.org/abs/1410.2597).
- FRANKE, A., MCGOVERN, D. P., BARRETT, J. C., WANG, K., RADFORD-SMITH, G. L., AHMAD, T., LEES, C. W., BALSCHUN, T., LEE, J. et al. (2010). Genome-wide meta-analysis increases to 71 the number of confirmed Crohn’s disease susceptibility loci. *Nat. Genet.* **42** 1118–1125.
- GENOVESE, C. and WASSERMAN, L. (2002). Operating characteristics and extensions of the false discovery rate procedure. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **64** 499–517. MR1924303 <https://doi.org/10.1111/1467-9868.00347>
- GOEMAN, J. J. and SOLARI, A. (2011). Multiple testing for exploratory research. *Statist. Sci.* **26** 584–597. MR2951390 <https://doi.org/10.1214/11-STS356>
- GOODMAN, S. N., FANELLI, D. and IOANNIDIS, J. P. (2016). What does research reproducibility mean? *Sci. Transl. Med.* **8**: 341ps12–341ps12.
- HEDGES, L. V. and SCHAUER, J. M. (2019a). Consistency of effects is important in replication: Rejoinder to Mathur and VanderWeele (2019) reply. *Psychol. Methods* **24** 576–577.
- HEDGES, L. V. and SCHAUER, J. M. (2019b). Statistical analyses for studying replication: Meta-analytic perspectives. *Psychol. Methods* **24** 557–570.
- HELD, L., MICHELOUD, C. and BALABDAOUI, F. (2022). A statistical framework for replicability. arXiv preprint. Available at [arXiv:2207.00464](https://arxiv.org/abs/2207.00464).
- HELLER, R., BOGOMOLOV, M. and BENJAMINI, Y. (2014). Deciding whether follow-up studies have replicated findings in a preliminary large-scale omics study. *Proc. Natl. Acad. Sci. USA* **111** 16262–16267.
- HELLER, R., GOLLAND, Y., MALACH, R. and BENJAMINI, Y. (2007). Conjunction group analysis: An alternative to mixed/random effect analysis. *NeuroImage* **37** 1178–1185.
- HELLER, R. and ROSSET, S. (2021). Optimal control of false discovery criteria in the two-group model. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **83** 133–155. MR4220987 <https://doi.org/10.1111/rssb.12403>
- HELLER, R. and SOLARI, A. (2023). Simultaneous directional inference. arXiv preprint. Available at [arXiv:2301.01653](https://arxiv.org/abs/2301.01653).
- HELLER, R. and YEKUTIELI, D. (2014). Replicability analysis for genome-wide association studies. *Ann. Appl. Stat.* **8** 481–498. MR3191999 <https://doi.org/10.1214/13-AOAS697>
- HIGGINS, J., THOMAS, J., CHANDLER, J., CUMPSTON, M., LI, T., PAGE, M. and WELCH, V. (2022). *Cochrane Handbook for Systematic Reviews of Interventions, Version 6.3*, (updated February 2022).
- HOANG, A.-T. and DICKHAUS, T. (2022). Combining independent p -values in replicability analysis: A comparative study. *J. Stat. Comput. Simul.* **92** 2184–2204. MR4437521 <https://doi.org/10.1080/00949655.2021.2022678>
- HOLM, S. (1979). A simple sequentially rejective multiple test procedure. *Scand. J. Stat.* **6** 65–70. MR0538597
- HOMMEL, G. (1988). A stagewise rejective multiple test procedure based on a modified Bonferroni test. *Biometrika* **75** 383–386.
- HUGHES, D., JUDGE, C., MURPHY, R., LOUGHLIN, E., COSTELLO, M., WHITELEY, W., BOSCH, J., O’DONNELL, M. J. and CANAVAN, M. (2020). Association of blood pressure lowering with incident dementia or cognitive impairment: A systematic review and meta-analysis. *JAMA* **323** 1934–1944.
- HUNG, K. and FITHIAN, W. (2020). Statistical methods for replicability assessment. *Ann. Appl. Stat.* **14** 1063–1087. MR4152124 <https://doi.org/10.1214/20-AOAS1336>
- IOANNIDIS, J. (2005). Why most published research findings are false. *PLoS Med.* **2** 696–701.
- JALJULI, I., BENJAMINI, Y., SHENHAV, L., PANAGIOTOU, O. A. and HELLER, R. (2022). Quantifying replicability and consistency in systematic reviews. *Stat. Biopharm. Res.* 1–14.
- JONES, L. V. and TUKEY, J. W. (2000). A sensible formulation of the significance test. *Psychol. Methods* **5** 411.
- KARMAKAR, B. and SMALL, D. S. (2020). Assessment of the extent of corroboration of an elaborate theory of a causal hypothesis using partial conjunctions of evidence factors. *Ann. Statist.* **48** 3283–3311. MR4185809 <https://doi.org/10.1214/19-AOS1929>
- KIDD, K. K., PAKSTIS, A. J., SPEED, W. C. and KIDD, J. R. (2004). Understanding human DNA sequence variation. *J. Hered.* **95** 406–420. <https://doi.org/10.1093/jhered/esh060>
- LAWLOR, D. A., TILLING, K. and DAVEY SMITH, G. (2017). Triangulation in aetiological epidemiology. *Int. J. Epidemiol.* **45** 1866–1886.
- LI, S., SESIA, M., ROMANO, Y., CANDÈS, E. and SABATTI, C. (2022). Searching for robust associations with a multi-environment knockoff filter. *Biometrika* **109** 611–629. MR4472838 <https://doi.org/10.1093/biomet/asab055>
- LIU, Z., SHEN, J., BARFIELD, R., SCHWARTZ, J., BACCARELLI, A. A. and LIN, X. (2022). Large-scale hypothesis testing for causal mediation effects with applications in genome-wide epigenetic studies. *J. Amer. Statist. Assoc.* **117** 67–81. MR4399068 <https://doi.org/10.1080/01621459.2021.1914634>
- MARIGORTA, U. M., RODRIGUEZ, J. A., GIBSON, G. and NAVARRO, A. (2018). Replicability and prediction: Lessons and challenges from gwas. *Trends Genet.* **34** 504–517.
- MATHUR, M. B. and VANDERWEELE, T. J. (2019). New metrics for meta-analyses of heterogeneous effects. *Stat. Med.* **38** 1336–1342. MR3920618 <https://doi.org/10.1002/sim.8057>
- NAKAGOME, S., MANO, S., KOZLOWSKI, L., BUJNICKI, J. M., SHIBATA, H., FUKUMAKI, Y., KIDD, J. R., KIDD, K. K., KAWAMURA, S. et al. (2012). Crohn’s disease risk alleles on the NOD2

- locus have been maintained by natural selection on standing variation. *Mol. Biol. Evol.* **29** 1569–1585.
- NOSEK, B. A., HARDWICKE, T. E., MOSHONTZ, H., ALLARD, A., CORKER, K. S., DREBER, A., FIDLER, F., HILGARD, J., STRUHL, M. K. et al. (2022). Replicability, robustness, and reproducibility in psychological science. *Annu. Rev. Psychol.* **73** 719–748.
- OWEN, A. B. (2009). Karl Pearson's meta-analysis revisited. *Ann. Statist.* **37** 3867–3892. MR2572446 <https://doi.org/10.1214/09-AOS697>
- PANAGIOTOU, O. A., JALJULI, I. and HELLER, R. (2020). Replicability of treatment effect in study of blood pressure lowering with dementia. *JAMA* **324** 1465–1466.
- PATIL, P., PENG, R. D. and LEEK, J. T. (2019). A visual tool for defining reproducibility and replicability. *Nat. Hum. Behav.* **3** 650–652. <https://doi.org/10.1038/s41562-019-0629-z>
- PAWEL, S. and HELD, L. (2022). The sceptical Bayes factor for the assessment of replication success. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **84** 879–911. MR4460579
- ROSENBAUM, P. (2022). *Replication and Evidence Factors in Observational Studies*. Taylor & Francis, London.
- ROSENBAUM, P. R. (2001). Replicating effects and biases. *Amer. Statist.* **55** 223–227. MR1963397 <https://doi.org/10.1198/000313001317098220>
- ROSENBAUM, P. R. (2010). Evidence factors in observational studies. *Biometrika* **97** 333–345. MR2650742 <https://doi.org/10.1093/biomet/asq019>
- ROY, S., BOGOMOLOV, M., HELLER, R., CLARIDGE, A. M., BEESON, T. and SMALL, D. S. (2022). Protocol for an observational study on the effects of giving births from unintended pregnancies on later life physical and mental health. arXiv preprint. Available at arXiv:2210.05169.
- SAAD, A., YEKUTIELI, D., LEV-RAN, S., GROSS, R. and GUYATT, G. (2019). Getting more out of meta-analyses: A new approach to meta-analysis in light of unexplained heterogeneity. *J. Clin. Epidemiol.* **107** 101–106. <https://doi.org/10.1016/j.jclinepi.2018.11.023>
- SAMPSON, J. N., BOCA, S. M., MOORE, S. C. and HELLER, R. (2018). FWER and FDR control when testing multiple mediators. *Bioinformatics* **34** 2418–2424.
- SESA, M., SABATTI, C. and CANDÈS, E. J. (2019). Gene hunting with hidden Markov model knockoffs. *Biometrika* **106** 1–18. MR3912377 <https://doi.org/10.1093/biomet/asy033>
- SIMONSOHN, U. (2015). Small telescopes: Detectability and the evaluation of replication results. *Psychol. Sci.* **26** 559–569.
- SOFER, T., HELLER, R., BOGOMOLOV, M., AVERY, C. L., GRAFF, M., NORTH, K. E., REINER, A. P., THORNTON, T. A., RICE, K. et al. (2017). A powerful statistical framework for generalization testing in GWAS, with application to the HCHS/SOL. *Genet. Epidemiol.* **41** 251–258.
- STOREY, J. D. (2003). The positive false discovery rate: A Bayesian interpretation and the q -value. *Ann. Statist.* **31** 2013–2035. MR2036398 <https://doi.org/10.1214/aos/1074290335>
- SUN, W. and CAI, T. T. (2007). Oracle and adaptive compound decision rules for false discovery rate control. *J. Amer. Statist. Assoc.* **102** 901–912. MR2411657 <https://doi.org/10.1198/016214507000000545>
- SUN, W. and WEI, Z. (2011). Multiple testing for pattern identification, with applications to microarray time-course experiments. *J. Amer. Statist. Assoc.* **106** 73–88. MR2816703 <https://doi.org/10.1198/jasa.2011.ap09587>
- TUKEY, J. W. (1991). The philosophy of multiple comparisons. *Statist. Sci.* 100–116.
- WANG, J., GUI, L., SU, W. J., SABATTI, C. and OWEN, A. B. (2022). Detecting multiple replicating signals using adaptive filtering procedures. *Ann. Statist.* **50** 1890–1909. MR4474476 <https://doi.org/10.1214/21-aos2139>
- WANG, J. and OWEN, A. B. (2019). Admissibility in partial conjunction testing. *J. Amer. Statist. Assoc.* **114** 158–168. MR3941245 <https://doi.org/10.1080/01621459.2017.1385465>
- WANG, P. and ZHU, W. (2019). Replicability analysis in genome-wide association studies via Cartesian hidden Markov models. *BMC Bioinform.* **20** 146.
- XIANG, D., ZHAO, S. D. and CAI, T. T. (2019). Signal classification for the integrative analysis of multiple sequences of large-scale multiple tests. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **81** 707–734. MR3997098
- XIE, J., CAI, T. T., MARIS, J. and LI, H. (2011). Optimal false discovery rate control for dependent data. *Stat. Interface* **4** 417–430. MR2868825 <https://doi.org/10.4310/SII.2011.v4.n4.a1>
- ZHAO, Q., SMALL, D. S. and ROSENBAUM, P. R. (2018). Cross-screening in observational studies that test many hypotheses. *J. Amer. Statist. Assoc.* **113** 1070–1084. MR3862340 <https://doi.org/10.1080/01621459.2017.1407770>
- ZHAO, Q., SMALL, D. S. and SU, W. (2019). Multiple testing when many p -values are uniformly conservative, with application to testing qualitative interaction in educational interventions. *J. Amer. Statist. Assoc.* **114** 1291–1304. MR4011780 <https://doi.org/10.1080/01621459.2018.1497499>