

Compositional Covariate Importance Testing via Partial Conjunction of Bivariate Hypotheses

Ritwik Bhaduri Siyuan Ma Lucas Janson

Abstract

Compositional data (i.e., data comprising random variables that sum up to a constant) arises in many applications including microbiome studies, chemical ecology, political science, and experimental designs. Yet when compositional data serve as covariates in a regression, the sum constraint renders every covariate automatically conditionally independent of the response given the other covariates, since each covariate is a deterministic function of the others. Since essentially all covariate importance tests and variable selection methods, including parametric ones, are at their core testing conditional independence, they are all completely powerless on regression problems with compositional covariates. In fact, compositionality causes ambiguity in the very notion of relevant covariates. To address this problem, we identify a natural way to translate the typical notion of relevant covariates to the setting with compositional covariates and establish that it is intuitive, well-defined, and unique. We then develop corresponding hypothesis tests and controlled variable selection procedures via a novel connection with *bivariate* conditional independence testing and partial conjunction hypothesis testing. Finally, we provide theoretical guarantees of the validity of our methods, and through numerical experiments demonstrate that our methods are not only valid but also powerful across a range of data-generating scenarios.

1 Introduction

1.1 Background and Motivation

Compositional random variables are random vectors whose elements sum up to a known non-random constant. Compositional data arises whenever a fixed whole is divided into parts, such as in chemistry (chemical composition of solutions (Brückner and Heethoff, 2017; Baum, 1988)), material science (Pesenson et al., 2015), marketing (e.g., brand or product shares (Joueid and Coenders, 2018)), sociology (e.g., time-use surveys (Dumuid et al., 2020)), political science (e.g., voting proportions (Rodrigues and Lima, 2009)), food science (e.g., nutrient composition (Greenfield and Southgate, 2003)), or geology (e.g., rock composition (Buccianti et al., 2006)). Compositionality also arises when data values are only meaningful relative to one another, making it natural to normalize them to sum to a constant, such as relative abundance data in microbiome research (Turnbaugh et al., 2007; Gloor et al., 2017). Any categorical or factor random variable can also

be treated as compositional via one-hot encoding, resulting in a random vector of all zeros except a single entry one indicating the factor level. Such categorical data arises in many settings, e.g., clinical trials (e.g., testing multiple drugs (Festing, 2020)), social science/economics (e.g., evaluating multiple disjoint policies (Stoker and John, 2009)), genomics (e.g., gene knockout experiments (Zhang et al., 1994)) and business/marketing (e.g., conjoint analysis (Green and Srinivasan, 1978)). While their applications and usage vary, compositional data across fields all present a common problem when they are treated as covariates (i.e., explanatory variables, X) to explain or predict a response variable (Y).

Consider a specific example where the research question is to relate the gut microbiome (covariates measured, as is standard, via the compositional relative abundances of different gut microbes) with colorectal cancers, or CRC (response variable). *Fusobacterium nucleatum*, an opportunity pathogen, is considered a gut microbial risk factor for CRC, with mechanistic and epidemiological evidence supporting its role in the cancer’s tumorigenesis and development (Wang and Fang, 2023). However, the compositionality of microbiome abundances means that *F. nucleatum* cannot be independent of the rest of the microbiome. Unconditionally, this means that any marginal dependence between CRC and *F. nucleatum* guarantees marginal dependence between CRC and other microbes, even if they are non-cancer-related. Such issues with dependent (even non-compositional) covariates are the reason that regression inference is typically done in a multivariate way, where each covariate’s dependence with the response is evaluated *conditionally* on the other covariates. Yet due to compositionality, the abundance of *F. nucleatum* is *deterministic* given the abundance of the other microbes, and hence can have *no* conditional dependence on the response, wrongfully negating its biological association with CRC. This example highlights the challenge of identifying, and even defining, important covariates under compositionality, so we now formalize a simplified version of it (where *F. nucleatum* is the only important covariate) with mathematical notation and concrete distributions.

Example 1. Let $Y \in \{0, 1\}$ represent the CRC disease status and $X := (X_{\{1\}}, \dots, X_{\{p\}})$ represent the relative abundances of p microbes in a gut microbiome sample. By definition, $\sum_{j=1}^p X_{\{j\}} = 1$. To simplify things further, suppose the $X_{\{j\}}$ are i.i.d. Expo(1) random variables normalized by their sum $\sum_{j=1}^p X_{\{j\}}$ (note the normalization makes them no longer independent of one another) and the true (unknown) causal model for Y given X is $Y | X \sim \text{Bernoulli}(e^{-X_{\{1\}}})$, where $X_{\{1\}}$ denotes the relative abundance of *F. nucleatum*. Although compositionality ensures that a model that excludes $X_{\{1\}}$, namely, $Y | X \sim \text{Bernoulli}(e^{-(1-\sum_{j=2}^p X_{\{j\}})})$, explains the response as well as the causal model, selecting $\{X_{\{2\}}, \dots, X_{\{p\}}\}$ as the set of important covariates for Y (as implied by that model) leads to a very wrong scientific conclusion. And indeed, while both the sets $\{X_{\{1\}}\}$ and $\{X_{\{2\}}, \dots, X_{\{p\}}\}$ each can explain Y equally well, in most scientific problems, a domain expert presented with these two options will choose $\{X_{\{1\}}\}$ as the far more biologically plausible one; this is the ubiquitous scientific principle of parsimony, or Occam’s Razor. Thus, we argue that a good definition of an important covariate, when applied to this example, should consider only $\{X_{\{1\}}\}$ as important, yet this is not the case for existing methods as we explain in the rest of this subsection and also in Section 1.3.

As alluded to in the microbiome example above, existing definitions of covariate im-

portance consider $X_{\{j\}}$ “unimportant” (this is considered the null hypothesis \mathcal{H}_j) either when it satisfies *unconditional* independence with Y or *conditional* independence with Y given the remaining covariates $X_{\{j\}^c}$. In Example 1, the unconditional \mathcal{H}_1 will be false, correctly defining $X_{\{1\}}$ as an important covariate. But compositionality also makes every *other* $X_{\{j\}}$ unconditionally dependent on $X_{\{1\}}$ and hence also on Y , so *all* unconditional \mathcal{H}_j will be false! It is not hard to see that this situation is typical when the covariates are compositional: for almost all compositional covariate distributions and (even sparse) response conditional distributions, all unconditional hypotheses will be false and hence any hypothesis test or selection method controlling an error rate with respect to such hypotheses will be unreliable for Type I errors with respect to any sparse true model. Switching to the conditional version of \mathcal{H}_j creates the opposite problem: compositionality means that every $X_{\{j\}}$ is a deterministic function of $X_{\{j\}^c}$, which immediately implies that each $X_{\{j\}}$ is conditionally independent of Y given $X_{\{j\}^c}$ and hence gives the degenerate conclusion that *every* conditional \mathcal{H}_j is true! Thus, any hypothesis test or selection method based on identifying false conditional \mathcal{H}_j ’s will be guaranteed to have trivial power.¹

Unfortunately, the conditional hypothesis $\mathcal{H}_j : Y \perp\!\!\!\perp X_{\{j\}} \mid X_{\{j\}^c}$ forms the basis of the majority of existing parametric and nonparametric hypothesis tests and variable selection techniques. For instance, parametric hypothesis tests and variable selection methods generally test hypotheses of the form $\beta_j = 0$ (Barber and Candès, 2015; Javanmard and Javadi, 2019; Montgomery et al., 2021) where the parameter β_j represents the contribution of $X_{\{j\}}$ to Y ’s conditional model, and a zero value represents no contribution and thus conditional independence. Similarly, a standard nonparametric null hypothesis for covariate importance in causal inference is that the average treatment effect is zero (Chernozhukov et al., 2018), which is also implied by conditional independence. It has also become increasingly popular to explicitly test conditional independence (Candès et al., 2018; Shah and Peters, 2020) (or conditional mean independence (Lundborg et al., 2022) which is also implied by conditional independence) directly. Since the conditional independence hypothesis is true for all covariates in a regression problem with compositional covariates, all the aforementioned methods (and many more) are guaranteed to have trivial power in such settings, regardless of the true relationship between Y and X .

We note here that a common workaround for conditional testing is to simply drop one arbitrarily chosen covariate, rendering the remaining covariates non-compositional. This indeed fixes the problem entirely, but *only* if the dropped covariate is not part of the true model. If the dropped covariate is part of the true model, not only does it preclude identifying that true covariate, but, due to the compositionally induced dependence, it also renders false covariates non-null. In particular, in Example 1 if we were to drop $X_{\{1\}}$, then conditional testing would of course be unable to correctly discover $X_{\{1\}}$, but would also consider all of $X_{\{2\}}, \dots, X_{\{p\}}$ as non-null. Since an analyst will not know a priori which variables are in the true model, there is no clear way in general to know that dropping a covariate will actually help.

¹For instance, a level- α hypothesis test for H_j must have power at most α .

1.2 Contributions

This paper makes two core contributions:

1. **Formalizing important covariates under compositionality:** To overcome the aforementioned misalignment between hypotheses (conditional or unconditional) and true signals in a parsimonious regression model with compositional covariates, we define the set of important covariates as the minimal set of covariates that together render all other covariates unhelpful for predicting the response, namely, the Markov boundary. We argue that under compositionality of the covariates, the Markov boundary aligns with natural intuition for what covariates should be discovered, and we show that it is well-defined, nontrivial, and unique under mild conditions. Although our conditions are technical, we show they can be simplified in certain compositional cases of interest.
2. **Methods for testing and variable selection with compositional covariates:** Our results about the Markov boundary reveal a direct connection between it and the *bivariate* conditional independence hypotheses $\mathcal{H}_{i,j} := Y \perp\!\!\!\perp X_{\{i,j\}^c} \mid X_{\{i,j\}}$, which are not rendered trivially true for compositional covariates like their univariate analogues. Leveraging this connection, we frame testing for covariate importance as a partial conjunction hypothesis test and apply multiple testing methods to such tests for variable selection. We propose a number of methodological innovations specific to our problem for partially overcoming the conservativeness inherent in partial conjunction hypothesis testing, and we demonstrate that our methods are valid and powerful in a range of simulations.

1.3 Related work

Testing for and selecting important covariates in multivariate regression is a heavily studied topic. First, there are many methods for estimating covariate importance or the set of important variables, which can be useful but do not generally come with Type I error guarantees, e.g., Zhang (2008); Schwarz (1978); Tibshirani (1996); Bertsimas et al. (2016); Berger and Pericchi (1996). As our ultimate goal in this paper is to perform rigorous hypothesis tests and *controlled* variable selection, we focus our literature review on methods with Type I error guarantees, which either devise individual hypothesis tests (whose p-values can be fed into multiple testing procedures (Holm, 1979; Benjamini and Hochberg, 1995; Benjamini and Yekutieli, 2001) for controlled variable selection) or directly perform controlled variable selection without individual hypothesis testing. Either way, these methods broadly fall into parametric (e.g., Barber and Candès (2015); Javanmard and Montanari (2018); Montgomery et al. (2021); Kuchibhotla et al. (2022); Xing et al. (2023)) and nonparametric (e.g., Candès et al. (2018); Chernozhukov et al. (2018); Berrett et al. (2020); Liu et al. (2022); Lundborg et al. (2022)) approaches. But the degeneracy of conditional covariate importance mentioned in Section 1.1 applies to all of these methods, parametric and nonparametric: they all define covariate importance in such a way that $Y \perp\!\!\!\perp X_{\{j\}} \mid X_{\{j\}^c}$ implies that $X_{\{j\}}$ is unimportant/null, and hence all of these methods treat *all* compositional covariates as null and thus will have trivial power to identify any covariates of interest. It is worth noting that although prior works like Candès et al. (2018) do, like us, define the set of important covariates as the Markov

boundary, they justify this target under standard assumptions which fail to hold for compositional covariates, and in particular their inferential methods are entirely based on the identity that the Markov boundary is exactly the set of covariates $X_{\{j\}}$ for which $Y \perp\!\!\!\perp X_{\{j\}} \mid X_{\{j\}^c}$; this identity fails to hold when the covariates are compositional, and hence new theory (which we present in Section 2) is needed to handle regression with compositional covariates.

There does exist a substantial body of literature on analyzing compositional data, but much of it has been focused on analyzing properties of the distribution of the compositional data itself (see, e.g., Aitchison (1982); Egozcue et al. (2003); Filzmoser and Hron (2009); Filzmoser et al. (2009); Pawlowsky-Glahn and Buccianti (2011); Greenacre et al. (2021); Pal and Heumann (2022)), as opposed to analyzing the conditional distribution of a response variable Y conditioned on compositional covariates X . There is work on compositional outcome regression (Fiksel et al., 2022), which is regression with a compositional response variable, but this is still different from our setting where the *covariates* are compositional. We note that, primarily within the microbiome field, a popular group of methods known as the Differential Abundance (DA) analysis (Paulson et al., 2013; Mallick et al., 2021) tests for marginal independence, i.e., $Y \perp\!\!\!\perp X_{\{j\}}$ for each j , optionally adjusting for confounding covariates *but not other microbes* (e.g. any host demographics, phenotypes, or exposure variables than can be associated with both X and Y).² Within canonical DA analysis, $X_{\{j\}}$ is usually treated as the outcome in regression modeling, against Y as the exposure variable of interest and other covariates. The corresponding null hypothesis, however, is equivalent to treating individual $X_{\{j\}}$ ’s as covariates and Y as the outcome (as they are treated in our paper), due to the tests’ marginal nature. As such, DA methods do not account for the compositionality, or more generally the dependence, of microbial covariates. Hence they all suffer from the issue highlighted in Section 1.1 that unconditional testing will usually treat all compositional covariates as *non-null*, and thus cannot be expected to control any Type I error rate with respect to any sparse or parsimonious set of important covariates (see Candès et al. (2018, Appendix B) for related discussion).

Finally, there are works which take a conditional approach to testing or selecting important covariates in regression with compositional covariates, and as these are the works most closely related to ours, we review them more closely in the rest of this subsection. A number of such works assume $Y \mid X$ follows a linear or generalized linear model with respect to the log-transformed covariates (Lin et al., 2014; Shi et al., 2016; Lu et al., 2019; Sohn and Li, 2019). To address the unidentifiability of the coefficient vector induced by compositionality, these works enforce a constraint on the coefficient vector that it must sum to zero. To the best of our knowledge this constraint is purely statistically motivated and thus may impact the interpretation of the coefficient vector they treat as their inferential target. In particular, if in Example 1 we were to model $Y \mid X \sim \text{Bernoulli}(e^{-(\beta_0 + X\beta)})$ for an intercept $\beta_0 \in \mathbb{R}$ and coefficient vector $\beta \in \mathbb{R}^p$ that sums to 0, then the “correct”

²We note that “differential abundance” analysis can also refer to methods that test for marginal independence between Y and the latent *absolute* abundance of the j^{th} microbe (Zhou et al., 2022; Wang, 2023; Lin and Peddada, 2024; Zong et al., 2024). Even moreso than canonical DA analysis, these methods differ fundamentally from the approach of this paper but require more space to fully explain, so we defer such discussion to Appendix A.

value of β would be $(\frac{p-1}{p}, -\frac{1}{p}, \dots, -\frac{1}{p})$ which is not sparse, and hence inference based on zeros/non-zeros in β will scientifically incorrectly treat $X_{\{2\}}, \dots, X_{\{p\}}$ as non-null. Srinivasan et al. (2021) employs a similar linear model as these works but replaces the assumption that the coefficient vector sums to zero with an assumption that a non-empty set of variables can be screened away based on the data, and that set of variables is guaranteed to contain *no* non-nulls—this is akin to a minimum signal size condition which may not hold in many settings.

More broadly, when regression covariates contain *factors*, i.e., covariates that take k possible non-metric levels, they are often handled by *contrast coding*, which creates k binary features each representing the indicator function for one of the factor levels, but then drops one of the factor levels (called the *reference* level) to address the compositionality of these k covariates (Montgomery, 2017). Contrast coding raises the challenge of selecting the reference level, since any downstream conclusions will depend on its choice (and one is generally not allowed to choose it based on the data itself). In particular, as mentioned at the end of Section 1.1, dropping one of the covariates in Example 1 and then defining importance via conditional dependence will correctly consider only $X_{\{1\}}$ as important if one of $X_{\{2\}}, \dots, X_{\{p\}}$ is dropped, but will very wrongly consider $X_{\{2\}}, \dots, X_{\{p\}}$ as important if $X_{\{1\}}$ is dropped, and, given the analyst does not know a priori that $X_{\{1\}}$ is special, it is unclear how an analyst can avoid such a mistake via dropping a variable.

Finally, the work that is closest to ours is Ma et al. (2024), which also adapts the usual conditional definition of covariate importance via conditional independence to a new definition catered to compositional covariates that avoids the degeneracy of conditional independence. Ma et al. (2024) defines a covariate as important if and only if $Y \not\perp\!\!\!\perp X_{\{j\}} \mid \frac{X_{\{j\}}^c}{\sum_{k \neq j} X_{\{k\}}}$, which says that a covariate is important if, when we intervene on its value while keeping the *relative* proportions of the remaining covariates fixed, Y 's (conditional) distribution changes. While this interventional definition is appealing and may be appropriate in some settings, it automatically treats all levels of a one-hot-coded factor covariate as null, and furthermore it defines all covariates as important in Example 1. In contrast, the definition of important covariates we lay out in this paper, using the Markov boundary, naturally handles factor covariates and correctly identifies only $X_{\{1\}}$ as important in Example 1.

1.4 Notation

In this paper, we use $[p]$ to denote the set $\{1, \dots, p\}$ for any $p \in \mathbb{N}$. For any set $A \subseteq [p]$, denote its complement $[p] \setminus A$ by A^c . For any p -dimensional vector $X = (X_1, \dots, X_p)$ and a subset $A \subseteq [p]$, X_A denotes the subvector $(X_j : j \in A)$ of X . For an $m \times n$ matrix P and indices $i \in [m]$ and $j \in [n]$, we denote the entry in row i and column j by $P_{i,j}$ and we define $\log(P)$ as the $m \times n$ matrix with $\log(P)_{i,j} = \log(P_{i,j})$.

1.5 Outline

In Section 2, we formally define important covariates in the compositional setting via the Markov boundary and prove that it remains well-defined under mild conditions despite compositionality. Section 3 details our methods for testing and controlled variable se-

lection and their theoretical guarantees. In Section 4, we assess the validity and power of our proposed methods across diverse data generation scenarios and hyperparameter settings. Finally, Section 5 concludes the paper with contributions, limitations, and ideas for future research.

2 Defining important compositional covariates

As highlighted in the previous section, the typical ways of defining an important covariate are not well-suited to compositional covariates. In this section, we put forth the Markov boundary as a solution and argue that, unlike conditional or unconditional dependence, membership in the Markov boundary continues to capture the right notion of an important covariate even under compositionality.

2.1 The Markov boundary

Suppose we have an \mathbb{R}^p -valued covariate vector $X = (X_{\{1\}}, \dots, X_{\{p\}})$ and a response Y ; we place no restriction on the form of Y , and in particular it can be Euclidean-valued or categorical. As X is compositional, $\sum_{i=1}^p X_{\{i\}} = c$ for some constant c , and without loss of generality, we will assume $c = 1$. As has been done in previous work (Candès et al., 2018) (though not for compositional covariates), we consider Y 's Markov boundary as a natural definition of the set of important covariates:

Definition 2.1 (Markov boundary (Pearl, 1988)). *For random variables $X \in \mathbb{R}^p$ and Y defined on the same probability space, a subset \mathcal{M} of $[p]$ is called a Markov boundary of Y if:*

1. $Y \perp\!\!\!\perp X_{\mathcal{M}^c} \mid X_{\mathcal{M}}$, and
2. $Y \not\perp\!\!\!\perp X_{\mathcal{M}'^c} \mid X_{\mathcal{M}'}$ for all $\mathcal{M}' \subsetneq \mathcal{M}$.

Unfortunately, when the covariates are compositional, there will in general be multiple Markov boundaries of Y , so the remaining subsections of this section are devoted to establishing mild conditions under which the above definition provides a well-defined, unique, and nontrivial set of important covariates. But in the remainder of this subsection we will first justify why Definition 2.1 provides a natural and intuitive definition for the set of important covariates. For readability, in the remainder of this subsection, we will suppose we are in the setting established by the end of next subsection, in which there is a unique nontrivial Markov boundary, and hence refer to *the* Markov boundary instead of *a* Markov boundary.

Item 1 in Definition 2.1 says that, after accounting for the covariates in the Markov boundary, all the remaining covariates provide no further information about Y . Item 2 says that the Markov boundary is the minimal such set, in the sense that no subset of it has the property in item 1. Together, this definition informally says that the Markov boundary is the minimal set of variables that, once known, allows us to drop all other variables without losing information about Y , and removing any variable from this set would lead to a strict loss of information about Y . When the covariates are not compositional, the Markov boundary is (under very mild conditions) equivalent to the set of

important covariates defined via conditional dependence (Edwards, 2012; Candès et al., 2018), which, as mentioned in the second-to-last paragraph of Section 1.1, is the basis for covariate importance throughout the literature on parametric and nonparametric methods for identifying important covariates in regression. Thus we find it to be a natural and intuitive target for variable selection with compositional covariates *if* we can show it remains well-defined under compositionality.

2.2 Conditions for uniqueness of the Markov boundary

The reader may be surprised that we can establish any conditions under which the Markov boundary will ever be unique with compositional covariates, since an immediate consequence of covariate compositionality is that *every* set of size $p - 1$ covariates will satisfy item 1 of Definition 2.1 and hence contain a subset satisfying condition 2 (and hence there can never be a truly unique Markov boundary when the covariates are compositional). However, we will show in this subsection that under very mild conditions, all but at most one Markov boundary will be of *exactly* size $p - 1$, and it is our position that, when X is compositional, a Markov boundary of size $p - 1$ is *trivial* and not worth considering, since every set of covariates of size $p - 1$ is information theoretically equivalent to the entire set of covariates due to compositionality. In Example 1, $\{1\}$ and $\{2, \dots, p\}$ are both technically Markov boundaries, but the latter is trivial and the former is the unique nontrivial Markov boundary. Our position is that the unique (under mild conditions) *nontrivial* Markov boundary is the natural target for variable selection even though other Markov boundaries technically exist.

First, we give an example to demonstrate why it is necessary to make some sort of assumptions in order to establish uniqueness of even of a nontrivial Markov boundary.

Example 2. Suppose $p = 4$, $X_{\{1\}}, X_{\{2\}}, X_{\{3\}}$ are mutually independent, and $X_{\{4\}} = 1 - X_{\{1\}} - X_{\{2\}} - X_{\{3\}}$. If $Y = X_{\{1\}} + X_{\{2\}} + \epsilon$ for ϵ independent of X , then:

- $Y \perp\!\!\!\perp X_{\{3\}}, X_{\{4\}} \mid X_{\{1\}}, X_{\{2\}}$, since $\epsilon \perp\!\!\!\perp X_{\{3\}}, X_{\{4\}}$,
- $Y \perp\!\!\!\perp X_{\{1\}}, X_{\{2\}} \mid X_{\{3\}}, X_{\{4\}}$, since $X_{\{1\}} + X_{\{2\}} = 1 - X_{\{3\}} - X_{\{4\}}$, and
- $Y \not\perp\!\!\!\perp X_{\{1,2,3,4\} \setminus \{j\}} \mid X_{\{j\}}$ for all $j \in \{1, 2, 3, 4\}$.

Thus, both $\{1, 2\}$ and $\{3, 4\}$ are Markov boundaries of size $< p - 1$.

Such a simple example may give the impression that it is not hard to find non-unique nontrivial Markov boundaries, but in fact it is constructed very carefully and we argue it should be vanishingly rare to encounter such an example in real data. In particular, the non-uniqueness of the nontrivial Markov boundary is a product of a precise interaction between $Y \mid X$, which depends on $X_{\{1\}}$ and $X_{\{2\}}$ through exactly their unweighted sum, and the compositionality of X , whose constraint also involves $X_{\{1\}}$ and $X_{\{2\}}$ through exactly their unweighted sum. So even changing $Y \mid X$ to $Y = X_{\{1\}} + 0.999X_{\{2\}} + \epsilon$ or $Y = X_{\{1\}} + X_{\{2\}}^{0.999} + \epsilon$ would result in a unique nontrivial Markov boundary. Thus, it seems one would have to get *very* unlucky to encounter a real data set whose alignment of $Y \mid X$ with the compositional constraint on X is so perfect.

To formalize when the nontrivial Markov boundary is uniquely defined, we will represent

it as a function of identifiable and testable statistical hypotheses. Let \mathcal{M} denote any Markov boundary as defined by Definition 2.1. For any $j \in \mathcal{M}$, if X were *not* compositional, we would expect Y to be conditionally dependent of $X_{\{j\}}$ given $X_{\{j\}^c}$ since $X_{\{j\}}$ contains unique conditional information about Y not captured by $X_{\{j\}^c}$; indeed the Markov boundary is typically (when the covariates are non-compositional) uniquely identified as the set of j for which $Y \perp\!\!\!\perp X_{\{j\}} \mid X_{\{j\}^c}$ is false (Edwards, 2012; Candès et al., 2018). The reason this identification strategy fails for compositional covariates is that any one covariate is perfectly determined by all the others, so, intuitively, $X_{\{j\}}$'s unique conditional information about Y cannot shine through because it is entirely absorbed into the information in $X_{\{j\}^c}$. An intuitive fix is to shift one more covariate from $X_{\{j\}^c}$ to join $X_{\{j\}}$, i.e., for some $i \in \{j\}^c$, consider the conditional relationship between Y and $X_{\{i,j\}}$ given $X_{\{i,j\}^c}$. If $j \in \mathcal{M}$, we would expect $Y \not\perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$ to hold for any $i \in \{j\}^c$, since $X_{\{j\}}$ contains unique conditional information about Y , and hence so should $X_{\{i,j\}}$, and that unique information is no longer entirely absorbed by the covariates that are conditioned on under compositionality, since $X_{\{i,j\}} \mid X_{\{i,j\}^c}$ is *not* deterministic under compositionality. This leads us to define

$$\mathcal{S} = \{j \in [p] : \mathcal{H}_{i,j} \text{ is false for all } i \in \{j\}^c\},$$

where the $\mathcal{H}_{i,j} : Y \perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$ are *bivariate* conditional independence null hypotheses. \mathcal{S} consists of the indices j such that all the bivariate null hypotheses containing j as one of the indices are false. After a brief remark, the remainder of this subsection will establish what our intuition suggests: under mild conditions, \mathcal{S} (which is uniquely defined by construction) coincides with the unique nontrivial Markov boundary.

Remark 1. *Key to our intuitive argument that bivariate conditional independence will allow us to identify the Markov boundary is the fact that compositionality places only a single constraint on X . If there were a further constraint, e.g., $\sum_{j=1}^p jX_{\{j\}} = p/2$, then $X_{\{i,j\}} \mid X_{\{i,j\}^c}$ would again be deterministic for all i, j pairs, and we would have to instead consider trivariate conditional independence, $\mathcal{H}_{i,j,k} : Y \perp\!\!\!\perp X_{\{i,j,k\}} \mid X_{\{i,j,k\}^c}$, and so on if there were yet more constraints on X . We will detail an important exception to this intuition in Section 2.3.2, but in the remainder of this subsection it may help the reader to think of the compositional X we consider as satisfying no additional deterministic constraints beyond compositionality, and our technical conditions will for the most part preclude additional constraints.*

First, we show that without any further conditions, any nontrivial Markov boundary contains \mathcal{S} :

Lemma 2.1. *For any Markov boundary \mathcal{M} such that $|\mathcal{M}| < p - 1$, $\mathcal{M} \supseteq \mathcal{S}$.*

Proof. Since \mathcal{M} is a Markov boundary, $Y \perp\!\!\!\perp X_{\mathcal{M}^c} \mid X_{\mathcal{M}}$. The weak union property³ (Pearl, 1988) implies that for any $\{i, j\} \subseteq \mathcal{M}^c$, $\mathcal{H}_{i,j}$ is true. Since \mathcal{S} contains exactly the i such that $\mathcal{H}_{i,j}$ is false for all $j \neq i$, that means that as long as $|\mathcal{M}^c| > 1$ (which holds by assumption), $i \in \mathcal{M}^c$ implies $i \in \mathcal{S}^c$. Thus $\mathcal{M}^c \subseteq \mathcal{S}^c$, i.e., $\mathcal{M} \supseteq \mathcal{S}$. \square

³(Weak union property) For any \mathbb{R}^p -valued random variable X and four subsets A, B, C and D of $[p]$, $X_A \perp\!\!\!\perp X_{B \cup D} \mid X_C$ implies $X_A \perp\!\!\!\perp X_B \mid X_{C \cup D}$.

If we can show $Y \perp\!\!\!\perp X_{\mathcal{S}^c} \mid X_{\mathcal{S}}$, then we are done, since $\mathcal{M} \supseteq \mathcal{S}$ by Lemma 2.1 but if $\mathcal{M} \supsetneq \mathcal{S}$ it would contradict property 2 of Definition 2.1 of a Markov boundary, so that only leaves the possibility that $\mathcal{M} = \mathcal{S}$. Such a result will require some technical (yet, we argue, mild) conditions, but the intuition is as follows. If there exists an \mathcal{M} of size $< p - 1$, then since $\mathcal{S} \subseteq \mathcal{M}$ by Lemma 2.1, $|\mathcal{S}^c| \geq 2$, i.e., given some $i \in \mathcal{S}^c$, there exists by definition some other $j \in \mathcal{S}^c$ such that $Y \perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$. Suppose we have another index $k \in \mathcal{S}^c \setminus \{i, j\}$ such that $Y \perp\!\!\!\perp X_{\{j,k\}} \mid X_{\{j,k\}^c}$ (if there is no such k , then $|\mathcal{S}| = p - 2$ and Lemma 2.1 immediately proves $\mathcal{M} = \mathcal{S}$ given $|\mathcal{M}| < p - 1$). If we can apply the *intersection property*⁴ of conditional independence (Pearl, 1988) on these variables, then $Y \perp\!\!\!\perp X_{\{i,j,k\}} \mid X_{\{i,j,k\}^c}$. So, the intersection property ‘grew’ the conditional independence statement from $Y \perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$ to $Y \perp\!\!\!\perp X_{\{i,j,k\}} \mid X_{\{i,j,k\}^c}$. If all the elements of \mathcal{S}^c can be iteratively added to the conditional independence statement via the intersection property, then we arrive at $Y \perp\!\!\!\perp X_{\mathcal{S}^c} \mid X_{\mathcal{S}}$ as desired. So, the key to proving $\mathcal{M} = \mathcal{S}$ is showing that we can apply the intersection property recursively, and the rest of this subsection describes sufficient conditions for this.

The standard sufficient condition for the intersection property is that X has a continuous density with respect to the Lebesgue measure (Peters, 2015). This fails for compositional X , whose support lies on the standard $(p - 1)$ -simplex, which has Lebesgue measure zero. So, we need to find a different set of assumptions which will be appropriate for compositional covariates. To do that, we first need to reintroduce some definitions from Peters (2015) about a random variable $X \in \mathbb{R}^p$ with distribution F .

Definition 2.2 (Coordinatewise connectivity (through A and B) (Peters, 2015)). *Let A and B be two subsets of $[p]$. Then two points $x, x' \in \text{support}(F)$ are said to be coordinatewise connected (through A and B) if either $x_A = x'_A$ or $x_B = x'_B$.*

Definition 2.3 (Equivalence (through A and B) (Peters, 2015)). *Any two points $x_1, x_2 \in \text{support}(F) \subseteq \mathbb{R}^p$ are said to be equivalent (through A and B) if there exists $z_1, z_2, \dots, z_L \in \text{support}(F)$ for some finite L such that the pairs $(x_1, z_1), (z_1, z_2), \dots, (z_{L-1}, z_L), (z_L, x_2)$ are each coordinatewise connected through A and B .*

So, equivalence extends coordinatewise connectivity in a pathwise way, in the sense that any two equivalent points can be connected by a sequence of consecutively coordinatewise connected points. Note that, as the name suggests, equivalence is indeed an equivalence relation, so we can define corresponding equivalence classes. The following lemma proves a version of the intersection property (Pearl, 1988) for two sets of variables such that X has a single equivalence class through them.

Lemma 2.2. *For random variables $X \in \mathbb{R}^p$ and Y defined on the same probability space, if $A \subseteq [p]$ and $B \subseteq [p]$ are such that $X_{A \cup B} \mid X_{(A \cup B)^c}$ has only one equivalence class (through $A \setminus B$ and $B \setminus A$), then $Y \perp\!\!\!\perp X_A \mid X_{A^c}$ and $Y \perp\!\!\!\perp X_B \mid X_{B^c}$ together imply $Y \perp\!\!\!\perp X_{A \cup B} \mid X_{(A \cup B)^c}$.*

The lemma and its proof in Appendix B.1 are inspired by Peters (2015, Corollary 1).

⁴(Intersection property) For any \mathbb{R}^p -valued random variable X and four subsets A, B, C , and D of $[p]$, if $X_A \perp\!\!\!\perp X_B \mid X_{C \cup D}$ and $X_A \perp\!\!\!\perp X_C \mid X_{B \cup D}$, then $X_A \perp\!\!\!\perp X_{B \cup C} \mid X_D$.

Next, define

$$\Delta = \{ \{A, B\} : X_{A \cup B} \mid X_{(A \cup B)^c} \text{ has 1 equivalence class (through } A \setminus B \text{ and } B \setminus A) \}.$$

So Δ denotes the pairs of sets on which we can apply Lemma 2.2. Further, define \mathcal{I} as the set of distinct $\{i, j\}$ pairs whose bivariate null hypotheses are true:

$$\mathcal{I} = \{ \{i, j\} : \mathcal{H}_{i,j} \text{ is true} \}.$$

Now, let us define the following operation on sets.

Definition 2.4 (Set action). *Suppose we have a set Φ of pairs of sets (like Δ). Consider a set Q of sets (like \mathcal{I}). We define the action of Φ on Q (denoted by $\Phi \circ Q$) as*

$$\Phi \circ Q := Q \cup \{A \cup B : A \in Q, B \in Q, \{A, B\} \in \Phi\}.$$

Note that both instances of the word “sets” in Definition 2.4 will refer to sets of integers in $[p]$ in this paper. Consider any set $C \in \Delta \circ \mathcal{I}$. By Definition 2.4, there exist $A, B \in \mathcal{I}$ satisfying $\{A, B\} \in \Delta$ with $A \cup B = C$. Since $A, B \in \mathcal{I}$, we have, by Lemma 2.2, $Y \perp\!\!\!\perp X_{A \cup B} \mid X_{(A \cup B)^c}$, i.e., $Y \perp\!\!\!\perp X_C \mid X_{C^c}$. Since C was generic, we conclude that $Y \perp\!\!\!\perp X_C \mid X_{C^c}$ for all $C \in \Delta \circ \mathcal{I}$. Applying the same argument with $\Delta \circ \mathcal{I}$ replacing \mathcal{I} , we get the same conditional independence for all $C \in \Delta \circ (\Delta \circ \mathcal{I})$. Continuing in this way, we get

$$Y \perp\!\!\!\perp X_C \mid X_{C^c} \quad \forall C \in (\Delta \circ)^k \mathcal{I}, \quad \forall k \in \mathbb{Z}_{\geq 0}.$$

The above argument is just a recursive way of combining sets for which the conditional independence statement $Y \perp\!\!\!\perp X_C \mid X_{C^c}$ is true through the intersection property. We continue this until we have a big enough set C . In particular, to show $\mathcal{M} \subseteq \mathcal{S}^c$ it is enough to show that we can reach a $C \supseteq \mathcal{S}^c$ in this way. Theorem 2.1, whose proof is presented in Appendix B.2, formalizes this argument.

Theorem 2.1 (Uniqueness of nontrivial Markov boundary). *If $\mathcal{S} = [p]$ then no nontrivial Markov boundary exists. Otherwise, if $\mathcal{S}^c \in (\Delta \circ)^k \mathcal{I}$ for some $k \in \mathbb{Z}_{\geq 0}$, then \mathcal{S} is the unique nontrivial Markov boundary.*

In addition to establishing when \mathcal{S} is the unique nontrivial Markov boundary for compositional X , Theorem 2.1 also says that there is no nontrivial Markov boundary when all the $\mathcal{H}_{i,j}$ are false. We can interpret this as the setting where there is no parsimony or sparsity at all and thus it is natural to consider *all* covariates as being important in this setting. Hence, although $\mathcal{S} = [p]$ is not a Markov boundary in this case, it remains an appropriate definition of the set of important covariates.

Theorem 2.1 is applicable to different types of compositional covariates, including discrete covariates as found in many experimental designs, continuous covariates such as relative abundances of chemicals, and mixed covariate distributions such as for microbe relative abundances which have both discrete (at 0) and continuous components. Additionally, the theorem is not restricted solely to compositional data but can also apply to any covariate data satisfying at most one deterministic constraint; see Section 5 for other examples. However, the condition that $\mathcal{S}^c \in (\Delta \circ)^k \mathcal{I}$ is rather abstract, so in the next subsection we provide simpler and more interpretable corollaries for important classes of compositional covariates.

2.3 Simplifications of Theorem 2.1

2.3.1 Continuous covariates

Continuous compositional distributions such as the Dirichlet distribution are often used to model compositional covariates. These arise commonly when a continuous whole is partitioned, such as component fractions of a chemical solution or amount of time spent on different activities during a day. For continuous compositional distributions, the assumptions in Theorem 2.1 can be simplified substantially. This is formalized in the following corollary:

Corollary 2.1 (Uniqueness of nontrivial Markov boundary for continuous distributions). *If $\mathcal{S} = [p]$ then no nontrivial Markov boundary exists. Otherwise, if the following two assumptions on the distribution of (Y, X) hold:*

- (i) *For all $i, j \in \mathcal{S}^c$, there exists a finite t and a sequence l_1, l_2, \dots, l_t such that $\mathcal{H}_{i,l_1}, \mathcal{H}_{l_1,l_2}, \dots, \mathcal{H}_{l_{t-1},l_t}, \mathcal{H}_{l_t,j}$ are all true,*
- (ii) *X can be written as $X = \frac{Z}{\sum_{j \in [p]} Z_{\{j\}}}$ for a \mathbb{R}^p -valued random variable Z such that there exists a function f_Z that is a density for Z and is continuous in $[0, \infty)^p \setminus \{\mathbf{0}_p\}$ and zero elsewhere, and for all $C \subseteq [p]$ and $c \in [0, 1]^{|C|}$, the set*

$$\left\{ z : \frac{z_C}{\sum_{j \in [p]} z_{\{j\}}} = c, f_Z(z) > 0 \right\}$$

is path connected if non-empty,

then \mathcal{S} is the unique nontrivial Markov boundary.

The proof of Corollary 2.1 from Theorem 2.1 in Appendix B.3 requires a nontrivial adaptation of Peters (2015, Proposition 1) to the compositional setting. Assumption (i) places restrictions on the distribution of $Y \mid X$ and precludes pathological cases like Example 2 which had $\{\mathcal{H}_{1,2}, \mathcal{H}_{3,4}\}$ as the (non-path-connected) set of true nulls. On the other hand, assumption (ii) is a condition on the marginal distribution of X and is completely agnostic to the distribution $Y \mid X$. Its main role is to exclude the possibility of additional constraints on the distribution of X aside from the sum constraint due to compositionality. Standard continuous compositional distributions such as the Dirichlet distribution (which is a normalization of independent Gamma distributed random variables) can be easily shown to satisfy assumption (ii).

2.3.2 Categorical or factor covariates

In addition to relative abundance data, compositional data can arise when covariates represent categorical factors, which is common, for instance, in experimental design. The simplest case is when the covariates represent a single factor with p levels; such covariates arise, e.g., in one-way ANOVA experimental designs. In this case, $X_{\{j\}}$ is the indicator variable for the j^{th} level of the factor; this is also sometimes called “one-hot” encoding. Since the factor can only take one level at a time, X sums to 1. We can generalize this to K factors (as would arise in a K -factor design) by considering X as the concatenation of K single-factor versions of X , and notationally, we will say that the p entries of X

can be partitioned into sets F_1, \dots, F_K , corresponding to each factor, so that for each $k \in [K]$, $\sum_{j \in F_k} X_{\{j\}} = 1$.⁵ While this imposes K constraints, contradicting the intuition of Remark 1, the disjointness of the constraints (since each only applies to X_{F_k} and the F_k are disjoint) makes it a special case that still permits the application of Theorem 2.1.

Note that, due to compositionality within each factor, $\mathcal{H}_{i,j}$ is always true if i and j correspond to different factors. This immediately implies $\mathcal{S} = \emptyset$ when $K > 1$, so we present a modified definition of \mathcal{S} for this setting:

$$\mathcal{S}_F = \bigcup_{k \in [K]} \{j \in F_k : \mathcal{H}_{i,j} \text{ is false for all } i \in F_k \setminus \{j\}\}.$$

Note that when $K = 1$, $\mathcal{S}_F = \mathcal{S}$. We also generalize the definition of a nontrivial Markov boundary: for a Markov boundary \mathcal{M} , we define the factored Markov boundary corresponding to the factor k as $\mathcal{M}_k = \mathcal{M} \cap F_k$ and we call a factored Markov boundary \mathcal{M}_k nontrivial if $|\mathcal{M}_k| < |F_k| - 1$.

Corollary 2.2 (Uniqueness of nontrivial factored Markov boundary for regression with K factors). *Consider regression with K -factor covariates, and fix some $k \in [K]$. If $F_k \subseteq \mathcal{S}_F$, then no nontrivial factored Markov boundary corresponding to factor k exists. Otherwise, if the following two assumptions hold:*

- (i) *For all $i, j \in \mathcal{S}_F^c \cap F_k$, there exists a finite t and a sequence $l_1, l_2, \dots, l_t \in F_k$ such that $\mathcal{H}_{i,l_1}, \mathcal{H}_{l_1,l_2}, \dots, \mathcal{H}_{l_{t-1},l_t}, \mathcal{H}_{l_t,j}$ are all true,*
- (ii) *Positive probability is assigned to all levels of factor k ,*

then $\mathcal{S}_F \cap F_k$ is the unique nontrivial Markov boundary corresponding to factor k .

Assumption (i) is on $Y \mid X$ while (ii) is only on the marginal distribution of X . The proof of Corollary 2.2 in Appendix B.4 generalizes that of Corollary 2.1 to multiple compositional groups of covariates.

A generalization of K -factor covariates enforces that each observation take $L \geq 1$ levels from each factor; this arises frequently in gene knockout experiments where exactly $L > 1$ out of p genes are knocked out per individual, resulting in $\sum_{j \in F_k} X_{\{j\}} = L$ for all $k \in [K]$. Corollary C.1 in Appendix C proves uniqueness of the Markov boundary for this setting.

Remark 2. *Although the Markov boundary can usually be thought of as a property of $Y \mid X$, when X has highly restricted support $\text{supp}(X)$, the Markov boundary can also depend on $\text{supp}(X)$ as well as $Y \mid X$. This phenomenon is perhaps most intuitively clear when thinking about interactions: if $Y = X_{\{1\}}X_{\{2\}}$ and $X_{\{1\}}, X_{\{2\}} \in \{0, 1\}$, then we expect the Markov boundary will be $\{1, 2\}$ and indeed it will be if the support of $(X_{\{1\}}, X_{\{2\}})$ contains $(1, 1)$ and at least one other point. But if instead $X_{\{1\}} + X_{\{2\}} = 1$, then $Y \stackrel{\text{a.s.}}{=} 0$ and the Markov boundary will be empty (though well-defined). Thus, analysts should think carefully about the choice of $\text{supp}(X)$ in an experimental design and how it may impact the Markov boundary, especially when interactions are of interest (Dean and Voss, 1999).*

⁵While this means $\sum_{j=1}^p X_{\{j\}} = K$ instead of 1, we still refer to X as compositional.

3 Hypothesis testing and variable selection

Section 2 identified sufficient conditions for the existence of a unique nontrivial Markov boundary and argued that it provides a natural definition of important compositional covariates in regression. It also argued that when no nontrivial Markov boundary exists, the natural definition of important compositional covariates in regression is simply $[p]$. In particular, Theorem 2.1 established two possibilities: either (a) $\mathcal{S} = [p]$ and no nontrivial Markov boundary exists, or (b) $\mathcal{S} \subsetneq [p]$ and is (under the conditions of Theorem 2.1) the unique nontrivial Markov boundary. Note that the possibility that no nontrivial Markov boundary exists and $\mathcal{S} \subsetneq [p]$ is precluded by the assumption of the theorem that if $\mathcal{S} \subsetneq [p]$, then $\mathcal{S}^c \in (\Delta \circ)^k \mathcal{I}$ for some $k \in \mathbb{Z}_{\geq 0}$ and \mathcal{S} is the unique nontrivial Markov boundary. So while Section 2 was focused conceptually on the Markov boundary, what it really did was justify \mathcal{S} , which is always unique and well-defined, as an appropriate definition of the set of important covariates via its connection with the Markov boundary, under mild conditions. Thus, turning to methodology in this section, we will always assume the conclusion of Theorem 2.1:

Assumption 1. *Either (a) $\mathcal{S} = [p]$ and no nontrivial Markov boundary exists, or (b) $\mathcal{S} \subsetneq [p]$ and is the unique nontrivial Markov boundary.*

Then, we simply treat \mathcal{S} as our inferential target, and in particular we will develop hypothesis tests for

$$\mathcal{H}_{0j} : j \notin \mathcal{S}$$

as well as multiple testing procedures combining such tests for the task of controlled variable selection.

3.1 Testing for the importance of a compositional covariate

First, as is always the case in constructing hypothesis tests, we need only consider the properties of a test when the hypothesis is true (since if the hypothesis is false, it cannot be falsely rejected and hence the Type I error is automatically controlled). And the existence of a true \mathcal{H}_{0j} precludes case (a) of Assumption 1, so we only need to construct hypothesis tests that are valid under case (b) of Assumption 1. Thus, in the remainder of this subsection, we will focus solely on the case when $\mathcal{S} \subsetneq [p]$ is the unique nontrivial Markov boundary, as this is the only case we need to consider to establish valid hypothesis tests.

Just based on its definition, it may not be immediately clear how to test \mathcal{H}_{0j} . But since \mathcal{S} is defined via the hypotheses $\mathcal{H}_{i,j}$, which are standard conditional independence hypotheses which many methods exist to test, our first goal in this section is to connect \mathcal{H}_{0j} with the hypotheses $\mathcal{H}_{i,j}$. In particular, by definition, $\mathcal{H}_{0j} = \mathcal{H}_{0j}^{p-1}$, where for any integer r , \mathcal{H}_{0j}^r is defined as

$$\mathcal{H}_{0j}^r : |\{i \in \{j\}^c : \mathcal{H}_{i,j} \text{ false}\}| < r. \quad (1)$$

But in fact, letting $s := |\mathcal{S}|$, since we are only considering the case when $s < p$ (recall also that by definition, $s \neq p-1$, so $s < p \Rightarrow s < p-1$), then because \mathcal{S} is a Markov boundary ($Y \perp\!\!\!\perp X_{\mathcal{S}^c} \mid X_{\mathcal{S}}$) and by the weak union property of conditional independence, $\mathcal{H}_{i,j}$ is

true for all $i, j \in \mathcal{S}^c$. So (under Assumption 1) $\mathcal{H}_{0j} = \mathcal{H}_{0j}^{s+1}$, and combining this with the fact that $\mathcal{H}_{0j}^r \Rightarrow \mathcal{H}_{0j}^{r'}$ for any $r \leq r'$, we get two useful conclusions: first, a hypothesis test for \mathcal{H}_{0j}^r constitutes a valid test for \mathcal{H}_{0j} so long as $r > s$, and second, tests for \mathcal{H}_{0j}^r will tend to have power nonincreasing in r . Thus, if an analyst has any domain knowledge upper-bounding the size of \mathcal{S} , then they should test $\mathcal{H}_{0j}^{\bar{s}}$ where \bar{s} is the smallest (strict) upper-bound on s they are confident about, and that test will be valid for \mathcal{H}_{0j} as long as $\bar{s} > s$ (or $\bar{s} = p - 1$ remains always a valid choice, regardless of the value of s).

In fact it is not uncommon to have at least a coarse upper-bound for s (see Filzmoser et al. (2018); Chen and Li (2013); Rivera-Pinto et al. (2018) for examples of regression with compositional covariates in which such an upper-bound is available), since sparsity is a common belief in many applications. As we show in our simulations, even a relatively weak upper-bound on s of $\bar{s} = p/2$ is sufficient to gain substantial power over using just $\bar{s} = p - 1$. However, we emphasize that extra knowledge upper-bounding s is not necessary for the methods we propose, as the choice $\bar{s} = p - 1$ is always valid and can still provide good power.

Having equated our hypothesis of interest, \mathcal{H}_{0j} , with the hypothesis \mathcal{H}_{0j}^r from Equation (1) with $r = \bar{s}$, we now point out that \mathcal{H}_{0j}^r is a *partial conjunction hypothesis* (PCH) (Benjamini and Heller, 2008). For a given integer r , a PCH states that strictly fewer than r base hypotheses are false out of some collection of base hypotheses (in our case, the set of $p - 1$ hypotheses $\{\mathcal{H}_{i,j} : i \in \{j\}^c\}$), aligning exactly with Equation (1). PCH testing is a well-studied topic in the statistical literature (applications include meta-analyses (Owen, 2009), genomics (Wang and Owen, 2019), and neuroscience (Friston et al., 1999)), allowing us to apply existing PCH tests to $\mathcal{H}_{0j}^{\bar{s}}$ and thus to \mathcal{H}_{0j} .

PCH tests assume access to a p-value for each of the base hypotheses. Since our base hypotheses $\mathcal{H}_{i,j}$ are standard conditional independence hypotheses, there exist numerous methods for testing them given, say, a data set of n i.i.d. observations of the random vector (X, Y) (certain forms of dependence among observations, such as those arising from common experimental designs, are also compatible with conditional independence testing; see Section 3.4). This paper does not attempt to contribute to the substantial literature on conditional independence testing, and instead we take a conditional-independence-test-agnostic approach by simply assuming that for each $\mathcal{H}_{i,j}$, one of the many excellent available conditional independence tests has been applied to obtain a p-value $P_{i,j}$ that is valid for $\mathcal{H}_{i,j}$.⁶ In Section 3.4 we discuss a bit further the topic of testing $\mathcal{H}_{i,j}$ and describe the particular choice we make for our simulations.

In this paper, we will consider two PCH tests as ways to combine our base p-values $P_{i,j}$ into a valid test for \mathcal{H}_{0j} . Recognizing that the PCH \mathcal{H}_{0j}^r is equivalent to the statement that at least $(p - 1) - (r - 1)$ base hypotheses are null, the standard PCH testing strategy is to apply a global null hypothesis test to the $(p - 1) - (r - 1)$ largest base p-values, and the validity of such an approach follows from the super-uniformity of such p-values (Benjamini and Heller, 2008). The first PCH test we consider applies the Bonferroni global test, which, given $r = \bar{s} \leq p - 1$ as discussed earlier (which can either be $p - 1$ or

⁶While one may expect $P_{i,j} = P_{j,i}$ since $\mathcal{H}_{i,j} = \mathcal{H}_{j,i}$, we will not assume this, though we will allow it.

a known strict upper-bound on s), gives the p-value

$$P_j^B = (p - \bar{s})P_{(\bar{s}),j}, \quad (2)$$

where for any $i \in [p - 1]$, $P_{(i),j}$ denotes the i^{th} smallest p-value among $\{P_{i',j} : i' \in \{j\}^c\}$. The following result, which given our results so far is an immediate consequence of Benjamini and Heller (2008, Theorem 2), proves the validity of P_j^B for testing \mathcal{H}_{0j} .

Corollary 3.1. *Under Assumption 1, if $\bar{s} > s$ or $\bar{s} = p - 1$, then for any $\alpha \in [0, 1]$, $\mathbb{P}_{\mathcal{H}_{0j}}(P_j^B \leq \alpha) \leq \alpha$.*

The appeal of the Bonferroni approach is that it requires no assumptions about the dependence among the base p-values $P_{i,j}$. This is always a nice property to have, but is perhaps of particular value in our setting where $P_{i,j}$ and $P_{i',j}$ are testing hypotheses which are quite strongly linked, since they are computed from the same data and both contain $X_{\{j\}}$ in the pair of covariates whose conditional independence with Y is being tested. However, the Bonferroni approach tends to have low power when r is much less than the number of base hypotheses, and, perhaps surprisingly and as we discuss in Appendix E.7, P_j^B is not monotonic with respect to \bar{s} .

The other PCH test we consider applies the Simes global test (Simes, 1986), which requires some assumptions about the dependence among the $P_{i,j}$ but can have substantially more power than the Bonferroni approach. Given $r = \bar{s}$ as discussed earlier, the Simes approach to testing \mathcal{H}_{0j} gives the p-value

$$P_j^S = \min_{\bar{s} \leq i \leq p-1} \left\{ \frac{p - \bar{s}}{i - \bar{s} + 1} P_{(i),j} \right\}. \quad (3)$$

Validity of the Simes p-value requires a form of positive dependence among the null base p-values, namely, *positive regression dependency on each one from a subset* (PRDS) (Benjamini and Yekutieli, 2001). Positive dependence may be plausible (if not easily provable) for the base p-values in our setting because the shared structure of their corresponding hypotheses is a form of alignment, not disalignment, and indeed in our simulations we never observed Type I error violations for the Simes approach despite having no formal proof of PRDS for the null base p-values in any setting. Formally, we have the following validity result for P_j^S , which given our results so far follows from Benjamini and Heller (2008, Theorem 1).

Corollary 3.2. *Under Assumption 1, if either $\bar{s} > s$ or $\bar{s} = p - 1$, and the base p-values $\{P_{i,j} : i \in \mathcal{S}^c \setminus \{j\}\}$ are PRDS, then for all $\alpha \in [0, 1]$, $\mathbb{P}_{\mathcal{H}_{0j}}(P_j^S \leq \alpha) \leq \alpha$.*

Note that the power benefit of the Simes approach over the Bonferroni approach is only available when $\bar{s} < p - 1$, since when $\bar{s} = p - 1$,

$$P_j^S = P_j^B = \max_{i \in \{j\}^c} P_{i,j},$$

which is a valid p-value under no assumptions at all (not even Assumption 1).

Corollary 3.3. *For all $\alpha \in [0, 1]$, $\mathbb{P}_{\mathcal{H}_{0j}}(\max_{i \in \{j\}^c} P_{i,j} \leq \alpha) \leq \alpha$.*

Because all the methods described in this subsection apply PCH testing methods to bivariate conditional independence tests, we refer to this general approach as Bivariate Conditional PCH, or BCP for short.

3.2 Controlled variable selection

In addition to testing for the importance of an individual compositional covariate X_j , an analyst may be interested in finding a set of covariates $\hat{\mathcal{S}}$ that contains as many important covariates as possible while controlling some error rate, i.e., controlled variable selection.

3.2.1 Variable selection controlling the familywise error rate

To select a set $\hat{\mathcal{S}}$ controlling the *familywise error rate* (FWER) at level α , i.e., $\mathbb{P}(\hat{\mathcal{S}} \not\subseteq \mathcal{S}) \leq \alpha$, one option is to simply apply Holm’s procedure (Holm, 1979) to the p-values $\{P_j^{\text{B}}\}_{j=1}^p$, which will control the FWER under the conditions of Corollary 3.1. Similarly, Holm’s procedure applied to $\{P_j^{\text{S}}\}_{j=1}^p$ will control the FWER when the conditions of Corollary 3.2 hold for all $j \in \mathcal{S}^c$, but for expositional purposes we will focus on $\{P_j^{\text{B}}\}_{j=1}^p$ first before returning to $\{P_j^{\text{S}}\}_{j=1}^p$ later in this subsection.

While the above procedure is straightforward and can be powerful, the structure of the BCP p-values can be leveraged for an even more powerful procedure that controls the FWER under identical assumptions. To see this, first recall that Holm’s procedure sequentially builds a set of rejected indices $\hat{\mathcal{S}}$ by starting from $\hat{\mathcal{S}}_1 = \emptyset$ and at the j^{th} step, executes the following:

$$\text{if } \min_{j \in \hat{\mathcal{S}}_j^c} P_j^{\text{B}} \leq \frac{\alpha}{p - j + 1}, \text{ set } \hat{\mathcal{S}}_{j+1} = \hat{\mathcal{S}}_j \cup \{\arg \min_{j \in \hat{\mathcal{S}}_j^c} P_j^{\text{B}}\}; \text{ else, return } \hat{\mathcal{S}} = \hat{\mathcal{S}}_j. \quad (4)$$

The intuition behind the Holm’s procedure’s increasing (in j) p-value threshold $\frac{\alpha}{p-j+1}$ is that, at the j^{th} step, the $j-1$ hypotheses in $\hat{\mathcal{S}}_j$ have already been confidently identified as false, and thus one only needs to correct for the multiplicity of the remaining $p - (j-1) = p - j + 1$ hypotheses in $\hat{\mathcal{S}}_j^c$. But we can also apply this same intuition to the BCP p-values being thresholded to make them smaller, as follows. For a set $A \subseteq \{j\}^c$, let $P_{(i),j}(A)$ denote the i^{th} smallest p-value among $\{P_{i',j} : i' \in \{j\}^c \setminus A\}$, and define

$$P_j^{\text{B}}(A) = (p - \bar{s})P_{(\bar{s}-|A|),j}(A). \quad (5)$$

Then our proposed Holm’s procedure simply replaces the P_j^{B} ’s in (4) with $P_j^{\text{B}}(\hat{\mathcal{S}}_j)$; see Algorithm 1 in Appendix D.1 for a full statement.⁷ The intuition is that, having already rejected $\hat{\mathcal{S}}_j$ before the j^{th} step, we can not only discard its $j-1$ elements in the multiplicity correction, but we can also discard its elements from the set of base hypotheses $P_{i,j}$ we consider in constructing our PCH tests. Note that for any $A \subseteq A'$ and $i > |A'|$, $P_{(i-|A'|),j}(A') \leq P_{(i-|A|),j}(A)$, and thus also $P_j^{\text{B}}(A') \leq P_j^{\text{B}}(A)$. In particular, $P_j^{\text{B}}(\hat{\mathcal{S}}_j) \leq P_j^{\text{B}}(\emptyset) = P_j^{\text{B}}$, guaranteeing that Algorithm 1’s rejections will always contain those of the original Holm’s procedure with P_j^{B} . The following result (proved in Appendix B.6) establishes that it also controls the FWER under identical conditions as the original Holm’s procedure—in particular, under no assumptions at all on the dependence among the $P_{i,j}$ or P_j^{B} .

⁷This procedure becomes undefined at the $(\bar{s} + 1)^{\text{th}}$ step since $\bar{s} - |\hat{\mathcal{S}}_{\bar{s}+1}| = 0$, but if we assume (as we do) that $\bar{s} > s$, then the procedure can simply reject all of $[p]$ if this step is ever reached. In practice it seems more sensible (but slightly less powerful) to simply stop at the \bar{s}^{th} step if it is reached, unless \bar{s} is set to its default value $p - 1$ and is not known to be a strict upper-bound on s .

Theorem 3.1. *Under Assumption 1, if $\bar{s} > s$ or $\bar{s} = p - 1$, then Algorithm 1 controls the FWER at level α .*

As in the previous section, if a user is willing to make an assumption of PRDS among $\{P_{i,j} : i \in \mathcal{S}^c \setminus \{j\}\}$ for each j , then it is valid and more powerful to use the Simes p-values P_j^S in Holm's procedure, and this procedure also admits a uniform improvement by using

$$P_j^S(\hat{\mathcal{S}}_j) := \min_{\bar{s} - |\hat{\mathcal{S}}_j| \leq i \leq p-1 - |\hat{\mathcal{S}}_j|} \left\{ \frac{p - \bar{s}}{i - \bar{s} + 1 + |\hat{\mathcal{S}}_j|} P_{(i),j}(\hat{\mathcal{S}}_j) \right\} \quad (6)$$

in place of P_j^S at each step; see Algorithm 2 in Appendix D.1 for a full statement. The following result (proved in Appendix B.7) establishes that this modified Holm's procedure with $P_j^S(\hat{\mathcal{S}}_j)$ also controls the FWER under identical conditions as the original Holm's procedure with P_j^S .

Theorem 3.2. *Under Assumption 1, if either $\bar{s} > s$ or $\bar{s} = p - 1$, and for each $j \in [p]$ the base p-values $\{P_{i,j} : i \in \mathcal{S}^c \setminus \{j\}\}$ are PRDS, then Algorithm 2 controls the FWER at level α .*

Informally, Theorem 3.2 requires the $P_{i,j}$'s that share the same j to be positively dependent, and as was mentioned immediately after P_j^S was introduced in Equation (3), since $P_{i,j}$ and $P_{i',j}$ are testing analogous hypotheses on two pairs of covariates that overlap in one covariate (X_j), we may expect them to be at least somewhat positively aligned, and thus positively dependent. Furthermore, Theorem 3.2 places no assumptions beyond this, and in particular allows $P_{i,j}$ and $P_{j,i}$ to be arbitrarily strongly dependent, which is important because $P_{i,j}$ and $P_{j,i}$ are in fact testing the *same* hypothesis (on the same data), and hence are likely to be extremely strongly dependent, and could even be equal. Allowing for these forms of dependence seems to be important, as an alternative method for FWER control with multiple PCH p-values, AdaFilter-Bonferroni (Wang et al., 2016), assumes independence among the base p-values in order to prove validity and we found that it failed to control the FWER in our simulations; see Appendix E.8. In contrast, Algorithm 2 controlled the FWER empirically in all our simulations. We note that Bogomolov and Heller (2018) provide another method for FWER control with PCH p-values, but it is only applicable when the number of base hypotheses for each PCH is 2, making it unsuitable for our setting.

3.2.2 Variable selection controlling the false discovery rate

To select a set $\hat{\mathcal{S}}$ controlling the *false discovery rate* (FDR) at level α , i.e., $\mathbb{E} \left[\frac{|\hat{\mathcal{S}}|}{\max\{|\hat{\mathcal{S}}|, 1\}} \right] \leq \alpha$, one option is to apply the Benjamini–Yekutieli procedure (Benjamini and Yekutieli, 2001) to the Bonferroni BCP p-values $\{P_j^B\}_{j=1}^p$, which will control the FDR under the conditions of Corollary 3.1. However, it is much more powerful to instead apply the Benjamini–Hochberg procedure (Benjamini and Hochberg, 1995) to the Simes BCP p-values $\{P_j^S\}_{j=1}^p$, and it turns out that this is valid when the base p-values are PRDS.

Theorem 3.3 (FDR Control). *Under Assumption 1, if either $\bar{s} > s$ or $\bar{s} = p - 1$, and the $p(p-1)$ base p-values $\{P_{i,j}\}_{i \neq j}$ are PRDS, then Algorithm 3 controls the FDR at level α .*

Note that Theorem 3.3 would be a direct corollary of standard FDR control results (Benjamini and Yekutieli, 2001) if it assumed the *BCP* p-values were PRDS, but instead it assumes the *base* p-values are PRDS, and thus relies instead on recent results from Bogomolov (2021, Theorem 3.1) (see Appendix B.8 for complete proof). As discussed in the paragraph after Theorem 3.2, there are multiple sources of (potentially strong) positive dependence among the base p-values that are important to allow for. As with FWER, another method for FDR control with PCH p-values, AdaFilter-BH (Wang et al., 2016), assumes the base p-values $P_{i,j}$ are independent, and fails to control the FDR in our simulations (see Appendix E.8), underscoring the importance of allowing positive dependence among the p-values in the theoretical guarantees of our methods. In contrast, Algorithm 3 controls the FDR empirically in all our simulations. We note that another FDR control method for PCH tests by Heller and Yekutieli (2014) has a time complexity exponentially increasing with p , making it impractical for our setting, and the FDR control method of Bogomolov and Heller (2018) is only applicable when the number of base hypotheses for each PCH is 2, making it unsuitable for our setting.

3.3 Addressing sparsity in the covariates

An interesting phenomenon arises when some of the compositional covariates $X_{\{j\}}$ are sparse, i.e., $\mathbb{P}(X_{\{j\}} \neq 0)$ is small, as would be expected in, e.g., the microbiome, as certain microbial taxa may be completely absent from most observations. In the non-compositional setting, a high degree of sparsity in a covariate may make the power to identify *that* covariate as important quite low, but does not generally have much impact on the power to identify *other* (non-sparse) important covariates. Surprisingly, when the covariates are compositional, even having just one sparse covariate can make it nearly impossible to identify *any* important covariates via BCP, even if they are not sparse at all. To see why, let $\epsilon_j := \mathbb{P}(X_{\{j\}} \neq 0)$. If $\epsilon_j = 0$, then $X_{\{j\}} \stackrel{\text{a.s.}}{=} 0$ and for any $i \neq j$, $\mathcal{H}_{i,j} : Y \perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$ is equivalent to $\mathcal{H}_i : Y \perp\!\!\!\perp X_{\{i\}} \mid X_{\{i\}^c}$, and as discussed throughout Section 1, this \mathcal{H}_i is definitionally true for all i when the covariates are compositional, guaranteeing that $\mathcal{S} = \emptyset$ and trivializing the power of any BCP methods introduced in this section thus far. This is due to exactly the issue raised in Remark 1, and the problem with sparse covariates is that they approximate this situation by having $X_{\{j\}} = 0$ for *most* observations, effectively leading to most observations providing no evidence against \mathcal{H}_{0i} for any i and, thus, BCP methods having very low power to reject any \mathcal{H}_{0i} .

Remark 1 alluded to addressing a similar issue via higher-order conditional independence hypotheses, such as trivariate conditional independence, but this required that the extra constraint(s) (beyond compositionality) be deterministic (and known a priori), which will not typically be the case with sparse covariates for which ϵ_j is small but strictly positive. Our proposed solution is to *condition on* covariates that are expected a priori to be sparse, giving up on having any power to identify them as members of \mathcal{S} (which would likely have been low anyway, due to their sparsity), but salvaging the power to identify the rest of \mathcal{S} . Formally, suppose the covariate vector X is split into two parts: \mathcal{D} , the indices of covariates expected to be dense, and \mathcal{D}^c , the indices of covariates expected to be sparse. Define

$$\mathcal{S}_{\mathcal{D}} = \{j : \mathcal{H}_{i,j} \text{ is false for all } i \in \mathcal{D} \setminus \{j\}\},$$

as the analogue of \mathcal{S} if we simply condition on $X_{\mathcal{D}^c}$ and only consider $X_{\mathcal{D}}$ as candidate covariates for rejection. By simply treating all our random variables as conditional on $X_{\mathcal{D}^c}$ and treating $X_{\mathcal{D}}$ as our covariates (which, conditional on $X_{\mathcal{D}^c}$, remain compositional, since they sum to $1 - |X_{\mathcal{D}^c}|$), all the results and BCP methods of this paper extend directly to (multiple or single) testing $\mathcal{H}_{0i}^{\mathcal{D}} : j \in \mathcal{S}_{\mathcal{D}}$. And as long as none of the elements of \mathcal{D} are sparse, no bivariate conditional independence tests involving a sparse covariate are required (indeed, only $P_{\mathcal{D}}$, the $|\mathcal{D}| \times |\mathcal{D}|$ submatrix of P containing $P_{i,j}$ for $i, j \in \mathcal{D}$, is needed when conditioning on $X_{\mathcal{D}^c}$), so they cannot impact power.

Ultimately, however, our interest remains in our original target, \mathcal{S} , so to interpret the results of our BCP procedures conditional on $X_{\mathcal{D}^c}$, we must connect $\mathcal{S}_{\mathcal{D}}$ with \mathcal{S} . Indeed the following theorem shows that under similarly mild conditions as Theorem 2.1, $\mathcal{S}_{\mathcal{D}} = \mathcal{S} \cap \mathcal{D}$, so that BCP methods conditioning on $X_{\mathcal{D}^c}$ remain valid for our original target \mathcal{S} , though they will only have power to detect its dense elements, $\mathcal{S} \cap \mathcal{D}$.

Theorem 3.4 (Connecting $\mathcal{S}_{\mathcal{D}}$ and \mathcal{S}). *If $\mathcal{S} = [p]$, then $\mathcal{S}_{\mathcal{D}} = \mathcal{D}$ ($= \mathcal{S} \cap \mathcal{D}$). Otherwise, if $|\mathcal{S}^c \cap \mathcal{D}| \neq 1$ and $\mathcal{S}^c \in (\Delta \circ)^k \mathcal{I}$ for some $k \in \mathbb{Z}_{\geq 0}$ then $\mathcal{S}_{\mathcal{D}} = \mathcal{S} \cap \mathcal{D}$ which is the intersection of the unique nontrivial Markov boundary with \mathcal{D} .*

The proof of this theorem is presented in Appendix B.5. In addition to the conditions in Theorem 2.1, there is one additional condition in Theorem 3.4: $|\mathcal{S}^c \cap \mathcal{D}| \neq 1$. When $|\mathcal{S}^c \cap \mathcal{D}| = 1$ (call the element in this intersection j^*), the conditional compositionality of $X_{\mathcal{D}}$ results in $\mathcal{S}_{\mathcal{D}} = \mathcal{D}$ rather than the desired $\mathcal{D} \setminus \{j^*\} = \mathcal{S} \cap \mathcal{D}$, so we must assume this case does not occur.

From a methodological perspective, note that BCP procedures conditional on $X_{\mathcal{D}^c}$ can not only increase power (when $X_{\mathcal{D}^c}$ are sparse) over their unconditional counterparts, but also save on computation, as they only require computing bivariate conditional independence tests for pairs of elements of \mathcal{D} .

3.4 Obtaining base p-values $P_{i,j}$

This section so far has described methods for testing and variable selection that take as input the bivariate conditional independence test p-values $P_{i,j}$ ($i \neq j$), but has left the choice of test to the user. As the topic of conditional independence testing is well-studied, we view this as a feature of our methods that they are flexible to this choice of test, as well as the assumptions under which the test is valid. For further details on methods to test conditional independence, and the assumptions under which they are valid, see, e.g., Doran et al. (2014); Sen et al. (2017); Candès et al. (2018); Shah and Peters (2020); Berrett et al. (2020); Lundborg et al. (2022); Liu et al. (2022); Tansey et al. (2022); Kim et al. (2022); Shi et al. (2024).

While our position on which conditional independence test to use with our proposed BCP methods remains an agnostic one, we are forced to make a choice when we run our simulations; we choose the Conditional Randomization Test (Candès et al., 2018), and in particular a computationally efficient variant called the distilled conditional randomization test (dCRT) with a lasso-based test statistic (Liu et al., 2022). Appendix E.1 contains implementation details for the dCRT for bivariate conditional independence testing as it is used in our simulations.

We note here that while we originally assumed that the $P_{i,j}$ were computed from a dataset of n i.i.d. samples of (Y, X) , nothing in our theory relied on this assumption. Indeed, many conditional independence tests, including the dCRT, can be readily extended to settings in which the samples are neither independent nor identically distributed (see Appendix E.1 for more discussion of this topic for the dCRT), and hence the BCP methods proposed in this paper will immediately apply in such settings as well.

3.5 Computational speedups

BCP variable selection methods require computing p^2 (or $|\mathcal{D}|^2$ when conditioning on $X_{\mathcal{D}^c}$) base p-values, which can be quite expensive, especially in high dimensions, although it is trivially and fully parallelizable. Nevertheless, to help alleviate this issue, we discuss in this subsection some computational shortcuts for BCP methods.

Following the idea of ‘screening’ in Liu et al. (2022), any subset of the $P_{i,j}$ ’s can be set to 1 without impacting the marginal validity of (all) the $P_{i,j}$ ’s (only making them conservative), even if this subset is chosen based on the entire data set. This can lead to a substantial computational speedup if there is a relatively inexpensive way to identify which $P_{i,j}$ ’s are likely to be large, since then these $P_{i,j}$ ’s can simply be set to 1 without computing them. And if the $P_{i,j}$ ’s set to 1 were indeed going to be large anyway (e.g., > 0.1), then making them larger isn’t likely to reduce power much, if at all, since any such base p-value would have been too large to contribute to a rejection anyway, even if fully computed.

Such a subset of the $P_{i,j}$ ’s could be chosen via sparse regression of the response on the covariates, e.g., as the zero elements of the fitted coefficient vector of a single run of cross-validated LASSO. Another option, which can be used in conjunction with the aforementioned sparse regression approach, is possible if the $P_{i,j}$ ’s are being computed serially: when about to compute a given $P_{i,j}$, if enough other base p-values, $P_{i',j}$, in the same column have been computed already and are large, then there will already be no hope of the PCH p-value combining the base p-values of that column being small, and hence it is faster and no less powerful to simply set $P_{i,j}$ (and any other as-yet-uncomputed base p-values in the same column) to 1 without wasting computation on it.

In some cases, further speedups are possible within the computation of the conditional independence tests themselves; we detail one such for the dCRT in Appendix E.9 which also takes the form of setting some $P_{i,j}$ ’s to 1 without (fully) computing them, i.e., another form of data-dependent screening.

While data-dependent screening provably retains marginal validity of the $P_{i,j}$ ’s (Liu et al., 2022), it does in general change the dependence structure of the base p-values. This will have no impact on the validity of our BCP procedures that do not assume anything about the dependence among the base p-values (e.g., Holm’s procedure applied to Bonferroni PCH p-values), but could in principle impact the validity of BCP procedures that assume a form of positive dependence. However, we found all the speedups used in our simulations never led to empirical error control violations, and in fact led to nearly identical sets of rejections as the same procedures without any speedups (average Jaccard indices $> 98\%$ in each setting we tried; see Appendix E.9). Furthermore, putting these speedups together

made a big computational difference, with about a $5\times$ speedup with $n = p = 100$, and more than a $30\times$ speedup for $n = p = 1000$; see Appendix E.9 for details.

4 Numerical Experiments

In this section we demonstrate our methods’ performance in compositional settings (Section 4.1), non-compositional settings (for comparison with non-compositional methods; Section 4.2), and compositional settings with sparse covariates as described in Section 3.3 (Section 4.3). The appendix contains many more simulations, details, and plots, in particular: additional simulations to assess robustness under covariate model misspecification (Appendix E.6), evaluate the performance of the multiple PCH testing method AdaFilter (Appendix E.8), and examine the impact of computational speedups on validity and average power (Appendix E.9). For each simulation presented in the main text, we include only key, representative figures to demonstrate the main findings, with additional figures and results provided in Appendices E.2–E.5. Code to replicate all numerical experiments is available at https://github.com/Ritwik-Bhaduri/Compositional_Covariate_Importance_Testing/.

While Section 3 considered combining base p-values into PCH p-values via both Bonferroni’s and Simes’ method, the simulations in this section will focus only on PCH p-values combined via Simes’ approach (i.e., P_j^S), while we provide all the same results for Bonferroni’s approach in Appendix E.7. This is because BCP methods using the Simes’ approach have significantly more power than those using Bonferroni’s while still empirically controlling their respective error rates well. See Appendix E.7 for more details on the comparison between the Bonferroni and Simes approaches within BCP.

We compare three different values of \bar{s} for our methods: $p - 1$, $p/2$, and $s + 1$ (in all our simulations, $s + 1 < p/2$). These values represent increasing levels of knowledge about the sparsity of the data: $p - 1$ indicates no knowledge of sparsity, $s + 1$ indicates perfect knowledge, and $p/2$ indicates a moderate (and we argue realistic in many settings; see discussion in Section 3.1) level of knowledge by assuming at least half of the covariates are irrelevant. Given \bar{s} , we call the BCP hypothesis test (with Simes’ approach, i.e., P_j^S) $\text{BCP}(\bar{s})$, and for variable selection, we call our methods $\text{BCP}(\bar{s})\text{-BH}$ (Algorithm 3) and $\text{BCP}(\bar{s})\text{-Holm}$ (Algorithm 2) for FDR and FWER control, respectively. In Section 4.1 and Section 4.2, we compare BCP methods to case-specific benchmark methods described in those respective subsections.

4.1 Standard compositional covariate settings

In these simulations, the rows of the covariate matrix are generated as independent copies of a compositional random variable $X \in \mathbb{R}^{100}$. We consider two different models of X to show the validity of our method for different dependence structures among the compositional covariates. First we consider the Dirichlet distribution where $X \sim \text{Dirichlet}(\alpha_1, \dots, \alpha_{100})$ with $\alpha_1 = \dots = \alpha_{100} = 2$. The other compositional distribution we consider is the logistic-normal distribution (Xia et al., 2013), where $X_j = \frac{e^{Z_j}}{\sum_{i=1}^{100} e^{Z_i}}$ for $j \in \{1, \dots, 100\}$ and $Z \sim \mathcal{N}(\mathbf{0}, \Sigma)$ with $\Sigma_{ij} = 0.6^{|i-j|}$.


The response Y is generated from a linear model $\mathcal{N}(\log(X)\beta, 1)$. Our simulated data sets consist of $n = 100$ i.i.d. samples from (X, Y) . For single testing, the coefficient of the variable being tested is varied on the x-axis of our plots, while the non-null nuisance coefficients are i.i.d. $\mathcal{N}(0, 1)$. For variable selection, all non-null coefficients are $\mathcal{N}(0, \text{SNR}^2)$, where the scalar value SNR is varied on the x-axis, and can be thought of as a signal-to-noise ratio as it modulates the signal in the non-nulls. In all these simulations, there is a unique non-trivial Markov boundary, also equal to \mathcal{S} , which is given by the set of non-zero coefficients; its size s is always 10.

As we are not aware of any existing methods that can test for or select elements of the Markov boundary with error control guarantees when the covariates are compositional, we compare BCP methods to a benchmark that is invalid, but popular and powerful. Namely, we consider as a benchmark a Leave-One-Out (LOO) approach that first drops a column j uniformly at random from the data, and then analyzes the resulting non-compositional data via a standard (i.e., non-compositional) conditional independence test directly analogous to the bivariate conditional independence tests used within the BCP methods, namely, the dCRT with a lasso-based test statistic (see Appendix E.2 for more details). For variable selection with FWER control or FDR control, we plug the $p - 1$ p-values from this LOO approach into Holm’s procedure (LOO-Holm) or the Benjamini–Hochberg procedure (LOO-BH), respectively. This method is consistent with heuristics currently used in practice (like contrast coding) and is powerful. However, it lacks error control, as discussed near the end of Section 1.3.


The left panel of Figure 1a shows that for single testing, the LOO method fails to control type I error at the nominal level of 5%, while the BCP methods all control type I error well below 5%. Indeed, BCP methods were conservative across simulations due to their use of PCH p-values, which are well-known to be conservative (Wang et al., 2016). The right panel of Figure 1a shows the power of the BCP(\bar{s}) methods decreases with the value of \bar{s} , as expected. BCP($s + 1$), which assumes complete knowledge of s , achieves power similar to the LOO method and, in fact, beats the latter for higher signal strengths. The more realistic method—BCP($p/2$)—achieves power close to the oracle method BCP($s + 1$) and LOO, and well above that of BCP($p - 1$). Figure 1b shows analogous performance of BCP methods for variable selection: LOO-BH fails to control FDR, while BCP(\bar{s})-BH methods are conservative but have average power similar to LOO-BH. Appendix E.2 shows the results for FWER control are similar to those for FDR control. Furthermore, Appendix E.3 shows that both single testing and variable selection with logistic-normal covariates behave broadly similarly as with Dirichlet covariates. The only notable difference is a substantial drop in power for BCP($p - 1$) (but not the other BCP methods), which could be attributed to the additional dependence structure in the data leading to the base p-values $P_{i,j}$, where i is null and j is non-null, being less concentrated around zero. This effect would have more impact on the power of BCP methods with higher \bar{s} .

4.2 Non-compositional covariate settings

While BCP methods are primarily designed and recommended for regression problems with compositional covariates, nothing in their development or theory precludes their use on non-compositional covariates. In such a setting, we can benchmark them against



(a) Comparison of type I error and power for single testing with Dirichlet covariates. The target type I error is 5% and error bars correspond to ± 2 Monte Carlo standard errors.



(b) Comparison of FDR and average power for variable selection with Dirichlet covariates. The target FDR level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.

Figure 1: Comparison of methods for (a) single testing and (b) variable selection with Dirichlet covariates.

state-of-the-art non-compositional methods that are both valid and powerful (unlike in the previous subsection, where there was no available valid benchmark). In this subsection’s simulation, $X \in \mathbb{R}^{100}$ is generated from $\mathcal{N}(\mathbf{0}, \Sigma)$ where Σ is a Toeplitz matrix such that $\Sigma_{i,j} = 0.6^{|i-j|}$. The rest of the simulation setup including the conditional distribution of the response given the covariates and the sample size remain the same as in Section 4.1.

As a benchmark, we consider a standard (i.e., non-compositional) conditional independence test directly analogous to the bivariate conditional independence tests used with the BCP methods, namely, the dCRT with a lasso-based test statistic. This Univariate testing approach is valid and powerful for testing whether a covariate is a member of the Markov boundary when the covariates are non-compositional, and for variable selection with FWER control or FDR control, these test’s p-values can be plugged into Holm’s procedure (Univariate-Holm) or the Benjamini–Hochberg procedure (Univariate-BH), respectively.


Figure 2a shows that all single-testing methods control the type I error and the non-compositional Univariate test has the highest power, closely followed by the proposed BCP methods. Figure 2b shows the same pattern for FDR-controlled variable selection except the gap between average power of the BCP methods and Univariate-BH is even narrower; Appendix E.4 shows very similar results for FWER control. Thus, despite all the methodological overhead in the BCP approach that is needed to accommodate compositional covariates (and the significant corresponding conservativeness that comes with it), we find that in non-compositional settings, BCP methods are nevertheless not far behind state-of-the-art conditional independence tests in terms of power, with particularly competitive variable selection average power.

4.3 Sparse compositional covariate settings


To explore the BCP approach to sparse compositional covariates as presented in Section 3.3, we generate covariates in this section from a 100-dimensional Dirichlet-Multinomial distribution as follows: let $\alpha \in \mathbb{R}^{100}$ have entries $\alpha_j = (1 + e^{\frac{j-50}{5}})^{-1}$ and sample $\theta \sim \text{Dirichlet}(\alpha)$ and $X \sim \text{Multinomial}(200, \theta)$. This X distribution overall has about half of its entries non-zero but the sparsity is highly imbalanced, with the last entries much sparser than the first ones. The response Y is generated from a log-linear model $\mathcal{N}(\log(1 + X)\beta, 1)$, where the coefficient vector β is zero except for 10 uniformly randomly selected entries which are sampled independently from $\mathcal{N}(0, 1)$. The sample size was $n = 100$.

The goal of this simulation is to study the effect of \mathcal{D} on power, and since the density of the covariates decreases with increasing index, we choose $\mathcal{D} = [k]$ and vary k from 30 to 100. In this subsection, we will consider BCP methods with the following three values of \bar{s} : $|\mathcal{D}| - 1$, $p/4 (= 25 < 30 \leq |\mathcal{D}|)$, and $s + 1$; this is because the conditional BCP approach of Section 3.3 requires $\bar{s} < |\mathcal{D}|$.

Figure 3 shows that conditioning on sparse covariates as proposed in Section 3.3 can yield substantial gains in power over not doing so (i.e., $|\mathcal{D}^c| = 0$) in both single testing and variable selection. All methods in all settings benefited from conditioning on at least the 35 sparsest covariates (the 35th-sparsest covariate is 8% non-zero), and in some cases continued to benefit from conditioning on more than half the covariates. However,



(a) Comparison of type I error and power for single testing with multivariate normal covariates. The target type I error is 5% and error bars correspond to ± 2 Monte Carlo standard errors.



(b) Comparison of FDR and average power for variable selection with multivariate normal covariates. The target FDR level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.

Figure 2: Comparison of methods for (a) single testing and (b) multiple testing with multivariate normal covariates.

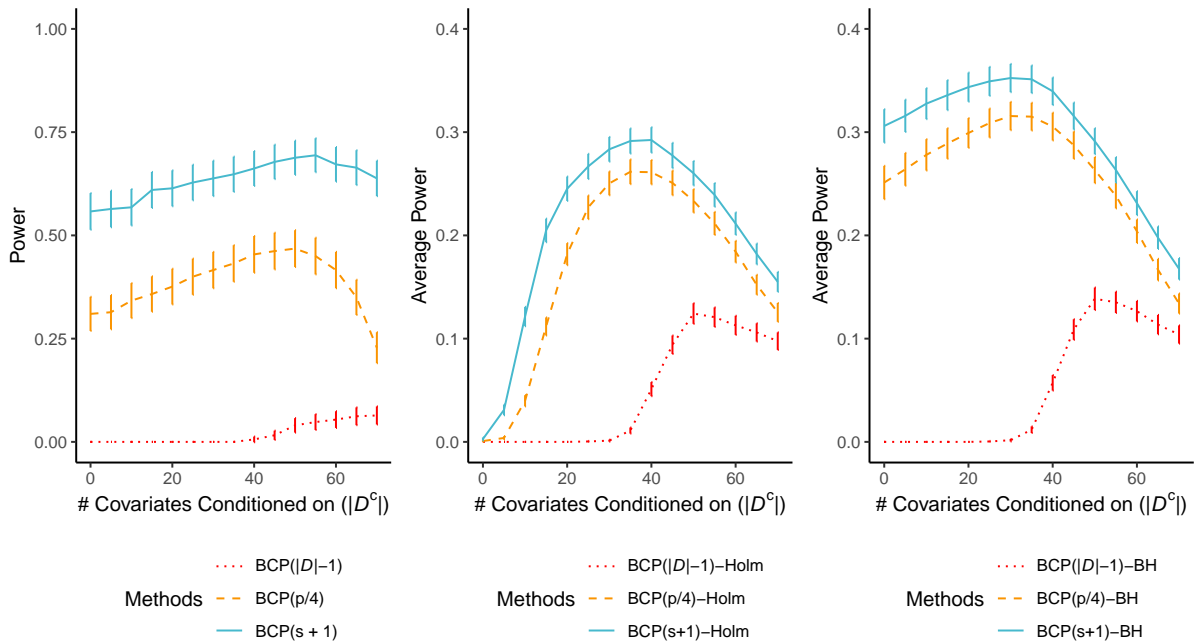


Figure 3: Effect of conditioning out the sparse covariates on power for single testing (left) and average power for FWER control (middle) and FDR control (right), respectively. Error bars correspond to ± 2 Monte Carlo standard errors.

as expected, for nearly all methods there is a sweet spot, and conditioning on too many covariates starts to harm power as covariates that are conditioned on cannot be discovered (recall that in these simulations, the probability that a covariate is non-null is independent of its sparsity). Appendix E.5 shows the corresponding type I error, FWER, and FDR plots, all of which are controlled below their nominal levels.

5 Discussion

In this work, we have defined an interpretable model-free notion of an important covariate in regression problems with compositional covariates. With this notion as an inferential target, we then developed novel, valid, and powerful procedures for testing and controlled variable selection, integrating ideas from conditional independence testing and partial conjunction hypothesis testing.

While this paper has primarily focused on compositional covariates, its contents in fact apply more generally to covariates that satisfy any single deterministic equality constraint. For example, the methods and theorems in this paper remain valid when the covariates are constrained to any $(p - 1)$ -dimensional linear subspace. Another allowable constraint is that the covariates lie on the surface of the $((p - 1)$ -dimensional) unit p -sphere (Lin, 2019); such data arises in, e.g., geophysics (Freeden et al., 1998), cosmology (Dodelson and Schmidt, 2020), plate tectonics (Chang et al., 2000) and embryology (Tyszka et al., 2005).

A natural future direction for this work is the setting of multiple constraints, i.e., when

the covariates are constrained to lie in a set of dimension $p - k$ for $k > 1$. This setting arises in, e.g., feature engineering (Zheng and Casari, 2018), basis expansion (Oreshkin et al., 2019), multi-factor experimental designs (Box and Hunter, 1957), and experimental designs with interactions (Dean and Voss, 1999). Section 2.3.2 showed that this paper’s approach can already handle multi-factor experimental design since the constraints are on disjoint subsets of the indices, but for other type of multiple constraints, naïve extension of the ideas in this paper would lead to computational intractability as discussed in Remark 1, with k constraints requiring testing of $\mathcal{O}(p^{k+1})$ k -variate conditional independence hypotheses. We hope that further insights and improvements along the lines of the computational speedups in Section 3.5 can be discovered in future work to make such approaches tractable.

Acknowledgments

RB and LJ were partially supported by DMS-2045981.

References

- Aitchison, J. (1982). The statistical analysis of compositional data. *Journal of the Royal Statistical Society: Series B (Methodological)*, 44(2):139–160.
- Barber, R. F. and Candès, E. J. (2015). Controlling the false discovery rate via knockoffs. *Annals of Statistics*.
- Baum, E. J. (1988). The use of pattern recognition to characterize air pollution regimes. *Chemometrics and Intelligent Laboratory Systems*, 3(1-2):91–98.
- Benjamini, Y. and Heller, R. (2008). Screening for partial conjunction hypotheses. *Biometrics*, 64(4):1215–1222.
- Benjamini, Y. and Hochberg, Y. (1995). Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal statistical society: series B (Methodological)*, 57(1):289–300.
- Benjamini, Y. and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. *Annals of statistics*, pages 1165–1188.
- Berger, J. O. and Pericchi, L. R. (1996). The intrinsic bayes factor for model selection and prediction. *Journal of the American Statistical Association*, 91(433):109–122.
- Berrett, T. B., Wang, Y., Barber, R. F., and Samworth, R. J. (2020). The conditional permutation test for independence while controlling for confounders. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(1):175–197.
- Bertsimas, D., King, A., and Mazumdar, R. (2016). Best subset selection via a modern optimization lens. *The Annals of Statistics*, 44(2):813–852.
- Blanchard, G. and Roquain, E. (2008). Two simple sufficient conditions for fdr control.

- Bogomolov, M. (2021). Testing partial conjunction hypotheses under dependency, with applications to meta-analysis. *arXiv preprint arXiv:2105.09032*.
- Bogomolov, M. and Heller, R. (2018). Assessing replicability of findings across two studies of multiple features. *Biometrika*, 105(3):505–516.
- Box, G. E. and Hunter, J. S. (1957). Multi-factor experimental designs for exploring response surfaces. *The Annals of Mathematical Statistics*, pages 195–241.
- Brückner, A. and Heethoff, M. (2017). A chemo-ecologists’ practical guide to compositional data analysis. *Chemoecology*, 27:33–46.
- Buccianti, A., Mateu-Figueras, G., and Pawlowsky-Glahn, V. (2006). *Compositional data analysis in the geosciences: from theory to practice*, volume 264. Geological Society of London.
- Candès, E., Fan, Y., Janson, L., and Lv, J. (2018). Panning for gold: ‘model-x’ knockoffs for high dimensional controlled variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(3):551–577.
- Chang, T., Ko, D., Royer, J.-Y., and Lu, J. (2000). Regression techniques in plate tectonics. *Statistical Science*, pages 342–356.
- Chen, J. and Li, H. (2013). Variable selection for sparse dirichlet-multinomial regression with an application to microbiome data analysis. *The annals of applied statistics*, 7(1).
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2018). Double/debiased machine learning for treatment and structural parameters.
- Dean, A. and Voss, D. (1999). *Design and analysis of experiments*. Springer.
- Dodelson, S. and Schmidt, F. (2020). *Modern cosmology*. Academic press.
- Doran, G., Muandet, K., Zhang, K., and Schölkopf, B. (2014). A permutation-based kernel conditional independence test. In *UAI*, pages 132–141.
- Dumuid, D., Pedišić, Ž., Palarea-Albaladejo, J., Martín-Fernández, J. A., Hron, K., and Olds, T. (2020). Compositional data analysis in time-use epidemiology: what, why, how. *International journal of environmental research and public health*, 17(7):2220.
- Edwards, D. (2012). *Introduction to graphical modelling*. Springer Science & Business Media.
- Egozcue, J. J., Pawlowsky-Glahn, V., Mateu-Figueras, G., and Barcelo-Vidal, C. (2003). Isometric logratio transformations for compositional data analysis. *Mathematical geology*, 35(3):279–300.
- Festing, M. F. (2020). The “completely randomised” and the “randomised block” are the only experimental designs suitable for widespread use in pre-clinical research. *Scientific reports*, 10(1):17577.

- Fiksel, J., Zeger, S., and Datta, A. (2022). A transformation-free linear regression for compositional outcomes and predictors. *Biometrics*, 78(3):974–987.
- Filzmoser, P. and Hron, K. (2009). Correlation analysis for compositional data. *Mathematical Geosciences*, 41:905–919.
- Filzmoser, P., Hron, K., and Reimann, C. (2009). Principal component analysis for compositional data with outliers. *Environmetrics: The Official Journal of the International Environmetrics Society*, 20(6):621–632.
- Filzmoser, P., Hron, K., Templ, M., Filzmoser, P., Hron, K., and Templ, M. (2018). Methods for high-dimensional compositional data. *Applied Compositional Data Analysis: With Worked Examples in R*, pages 207–225.
- Freeden, W., Gervens, T., and Schreiner, M. (1998). *Constructive approximation on the sphere: with applications to geomathematics*. Oxford University Press.
- Friston, K. J., Holmes, A. P., Price, C., Büchel, C., and Worsley, K. (1999). Multisubject fmri studies and conjunction analyses. *Neuroimage*, 10(4):385–396.
- Gloor, G. B., Macklaim, J. M., Pawlowsky-Glahn, V., and Egozcue, J. J. (2017). Microbiome datasets are compositional: and this is not optional. *Frontiers in microbiology*, 8:2224.
- Green, P. E. and Srinivasan, V. (1978). Conjoint analysis in consumer research: issues and outlook. *Journal of consumer research*, 5(2):103–123.
- Greenacre, M., Martínez-Álvarez, M., and Blasco, A. (2021). Compositional data analysis of microbiome and any-omics datasets: a validation of the additive logratio transformation. *Frontiers in microbiology*, 12:727398.
- Greenfield, H. and Southgate, D. A. (2003). *Food composition data: production, management, and use*. Food & Agriculture Org.
- Heller, R. and Yekutieli, D. (2014). Replicability analysis for genome-wide association studies. *The Annals of Applied Statistics*, 8(1):481–498.
- Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, pages 65–70.
- Javanmard, A. and Javadi, H. (2019). False discovery rate control via debiased lasso. *Electronic Journal of Statistics*, 13(1):1212 – 1253.
- Javanmard, A. and Montanari, A. (2018). Debiasing the lasso: Optimal sample size for gaussian designs.
- Joueid, A. and Coenders, G. (2018). Marketing innovation and new product portfolios. a compositional approach. *Journal of Open Innovation: Technology, Market, and Complexity*, 4(2):19.

- Kim, I., Neykov, M., Balakrishnan, S., and Wasserman, L. (2022). Local permutation tests for conditional independence. *The Annals of Statistics*, 50(6):3388–3414.
- Kuchibhotla, A. K., Kolassa, J. E., and Kuffner, T. A. (2022). Post-selection inferenced. *Annual Review of Statistics and Its Application*, 9:505–527.
- Lin, H. and Peddada, S. D. (2024). Multigroup analysis of compositions of microbiomes with covariate adjustments and repeated measures. *Nature Methods*, 21(1):83–91.
- Lin, S.-B. (2019). Nonparametric regression using needlet kernels for spherical data. *Journal of Complexity*, 50:66–83.
- Lin, W., Shi, P., Feng, R., and Li, H. (2014). Variable selection in regression with compositional covariates. *Biometrika*, 101(4):785–797.
- Liu, M., Katsevich, E., Janson, L., and Ramdas, A. (2022). Fast and powerful conditional randomization testing via distillation. *Biometrika*, 109(2):277–293.
- Lu, J., Shi, P., and Li, H. (2019). Generalized linear models with linear constraints for microbiome compositional data. *Biometrics*, 75(1):235–244.
- Lundborg, A. R., Kim, I., Shah, R. D., and Samworth, R. J. (2022). The projected covariance measure for assumption-lean variable significance testing. *arXiv preprint arXiv:2211.02039*.
- Ma, S., Huttenhower, C., and Janson, L. (2024). Compositional differential abundance testing: Defining and finding a new type of health-microbiome associations. *bioRxiv*.
- Mallick, H., Rahnavard, A., McIver, L. J., Ma, S., Zhang, Y., Nguyen, L. H., Tickle, T. L., Weingart, G., Ren, B., Schwager, E. H., et al. (2021). Multivariable association discovery in population-scale meta-omics studies. *PLoS computational biology*, 17(11):e1009442.
- Minka, T. (2000). Estimating a dirichlet distribution.
- Montgomery, D. C. (2017). *Design and analysis of experiments*. John wiley & sons.
- Montgomery, D. C., Peck, E. A., and Vining, G. G. (2021). *Introduction to linear regression analysis*. John Wiley & Sons.
- Oreshkin, B. N., Carpov, D., Chapados, N., and Bengio, Y. (2019). N-beats: Neural basis expansion analysis for interpretable time series forecasting. *arXiv preprint arXiv:1905.10437*.
- Owen, A. B. (2009). Karl pearson’s meta-analysis revisited.
- Pal, S. and Heumann, C. (2022). Clustering compositional data using dirichlet mixture model. *Plos one*, 17(5):e0268438.
- Paulson, J. N., Stine, O. C., Bravo, H. C., and Pop, M. (2013). Differential abundance analysis for microbial marker-gene surveys. *Nature methods*, 10(12):1200–1202.

- Pawlowsky-Glahn, V. and Buccianti, A. (2011). *Compositional data analysis*. Wiley Online Library.
- Pearl, J. (1988). Chapter 3 - markov and bayesian networks: Two graphical representations of probabilistic knowledge. In Pearl, J., editor, *Probabilistic Reasoning in Intelligent Systems*, pages 77–141. Morgan Kaufmann, San Francisco (CA).
- Pesenson, M. Z., Suram, S. K., and Gregoire, J. M. (2015). Statistical analysis and interpolation of compositional data in materials science. *ACS combinatorial science*, 17(2):130–136.
- Peters, J. (2015). On the intersection property of conditional independence and its application to causal discovery. *Journal of Causal Inference*, 3(1):97–108.
- Rivera-Pinto, J., Egozcue, J. J., Pawlowsky-Glahn, V., Paredes, R., Noguera-Julian, M., and Calle, M. L. (2018). Balances: a new perspective for microbiome analysis. *MSystems*, 3(4):10–1128.
- Rodrigues, P. C. and Lima, A. T. (2009). Analysis of an european union election using principal component analysis. *Statistical Papers*, 50:895–904.
- Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, pages 461–464.
- Sen, R., Suresh, A. T., Shanmugam, K., Dimakis, A. G., and Shakkottai, S. (2017). Model-powered conditional independence test. *Advances in neural information processing systems*, 30.
- Shah, R. D. and Peters, J. (2020). The hardness of conditional independence testing and the generalised covariance measure. *The Annals of Statistics*, 48(3).
- Shi, H., Drton, M., and Han, F. (2024). On azadkia–chatterjee’s conditional dependence coefficient. *Bernoulli*, 30(2):851–877.
- Shi, P., Zhang, A., and Li, H. (2016). Regression analysis for microbiome compositional data. *The Annals of Applied Statistics*.
- Simes, R. J. (1986). An improved bonferroni procedure for multiple tests of significance. *Biometrika*, 73(3):751–754.
- Sohn, M. B. and Li, H. (2019). Compositional mediation analysis for microbiome studies. *The Annals of Applied Statistics*, 13(1):661–681.
- Srinivasan, A., Xue, L., and Zhan, X. (2021). Compositional knockoff filter for high-dimensional regression analysis of microbiome data. *Biometrics*, 77(3):984–995.
- Stoker, G. and John, P. (2009). Design experiments: Engaging policy makers in the search for evidence about what works. *Political Studies*, 57(2):356–373.
- Tansey, W., Veitch, V., Zhang, H., Rabadan, R., and Blei, D. M. (2022). The holdout randomization test for feature selection in black box models. *Journal of Computational and Graphical Statistics*, 31(1):151–162.

- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 58(1):267–288.
- Turnbaugh, P. J., Ley, R. E., Hamady, M., Fraser-Liggett, C. M., Knight, R., and Gordon, J. I. (2007). The human microbiome project. *Nature*, 449(7164):804–810.
- Tyszka, J. M., Ewald, A. J., Wallingford, J. B., and Fraser, S. E. (2005). New tools for visualization and analysis of morphogenesis in spherical embryos. *Developmental dynamics: an official publication of the American Association of Anatomists*, 234(4):974–983.
- Wang, J., Gui, L., Su, W. J., Sabatti, C., and Owen, A. B. (2016). Detecting multiple replicating signals using adaptive filtering procedures. *arXiv preprint arXiv:1610.03330*.
- Wang, J. and Owen, A. B. (2019). Admissibility in partial conjunction testing. *Journal of the American Statistical Association*, 114(525):158–168.
- Wang, N. and Fang, J.-Y. (2023). Fusobacterium nucleatum, a key pathogenic factor and microbial biomarker for colorectal cancer. *Trends in Microbiology*, 31(2):159–172.
- Wang, S. (2023). Robust differential abundance test in compositional data. *Biometrika*, 110(1):169–185.
- Wright, S. P. (1992). Adjusted p-values for simultaneous inference. *Biometrics*, pages 1005–1013.
- Xia, F., Chen, J., Fung, W. K., and Li, H. (2013). A logistic normal multinomial regression model for microbiome compositional data analysis. *Biometrics*, 69(4):1053–1063.
- Xing, X., Zhao, Z., and Liu, J. S. (2023). Controlling false discovery rate using gaussian mirrors. *Journal of the American Statistical Association*, 118(541):222–241.
- Yekutieli, D. and Benjamini, Y. (1999). Resampling-based false discovery rate controlling multiple test procedures for correlated test statistics. *Journal of Statistical Planning and Inference*, 82(1-2):171–196.
- Zhang, T. (2008). Adaptive forward-backward greedy algorithm for sparse learning with linear models. In Koller, D., Schuurmans, D., Bengio, Y., and Bottou, L., editors, *Advances in Neural Information Processing Systems*, volume 21. Curran Associates, Inc.
- Zhang, Y., Proenca, R., Maffei, M., Barone, M., Leopold, L., and Friedman, J. M. (1994). Positional cloning of the mouse obese gene and its human homologue. *Nature*, 372(6505):425–432.
- Zheng, A. and Casari, A. (2018). *Feature engineering for machine learning: principles and techniques for data scientists*. ” O’Reilly Media, Inc.”.
- Zhou, H., He, K., Chen, J., and Zhang, X. (2022). Linda: linear models for differential abundance analysis of microbiome compositional data. *Genome biology*, 23(1):1–23.

Zong, Y., Zhao, H., and Wang, T. (2024). mbdecoda: a debiased approach to compositional data analysis for microbiome surveys. *Briefings in Bioinformatics*, 25(3):bbae205.

A Differential Absolute Abundance Methods

As mentioned in a footnote in Section 1.3, there is a class of differential abundance methods in the microbiome field that test for marginal independence between Y and the *absolute* abundance of the j^{th} microbe (Zhou et al., 2022; Wang, 2023; Lin and Peddada, 2024; Zong et al., 2024). We refer to such methods collectively as differential absolute abundance (DAA) to distinguish them from canonical DA analysis. To understand how DAA methods relate to this paper, suppose that the observed compositional covariates X represent the normalization of a latent noncompositional random vector Z , i.e., $X = Z / \sum_j Z_{\{j\}}$. Then DAA methods are designed to test hypotheses of the form $Y \perp\!\!\!\perp Z_{\{j\}}$. Because DAA methods resemble DA methods but with $X_{\{j\}}$ replaced by $Z_{\{j\}}$, they are in some sense a step further removed from this paper than DA methods, which is why we defer their discussion to this appendix.

First and foremost, DAA methods test a hypothesis in terms of the latent random variables $Z_{\{j\}}$. Thus, it represents fundamentally different *scientific questions* than considered by DA methods or any of the methods in this paper, which are entirely in terms of the observable random variables Y and X . In fact, Z is only a meaningful quantity in certain applications where X represents the normalization of some underlying random variable (for example, the microbiome field). Many other compositional random variables (e.g., factors in experimental design, chemical compositions, or time use within a day), however, are compositional but do not represent normalization from underlying noncompositional variables. There is thus no natural scientific question regarding the latent $Z_{\{j\}}$'s for these types of applications.

Beyond the conceptual difference, DAA methods differ from our approach in two additional, important ways. First, like DA methods, DAA methods test *marginal* independence, meaning that unlike the multivariate approach taken in this paper, they do not account for any covariates other than the j^{th} in their hypothesis. Second, the latency of Z will generally make it (and hence nearly any hypothesis involving it) unidentifiable except under specific assumptions: most DAA methods assume highly parametric distributions on $Z_{\{j\}}$ and $X_{\{j\}}$; all require some form of sparsity conditions that not too many $Z_{\{j\}}$'s can be associated with Y . As such, DAA methods also differ from this paper in taking a parametric approach to inference, necessitated in part by the latency of Z .

B Proof of Theorems

B.1 Proof of Lemma 2.2

We are first going to prove Lemma 2.2 holds for coordinatewise connected vectors (through $A \setminus B$ and $B \setminus A$). The proof for equivalent vectors will follow similarly.

Consider two vectors (u, v, w) and (u', v', w') in the support of $X_{A \cup B}$, conditioned on

$X_{(A \cup B)^c} = z$, that are coordinatewise connected through $A \setminus B$ and $B \setminus A$. Here, (u, v, w) is shorthand for $X_{A \cap B} = u$, $X_{A \setminus B} = v$, and $X_{B \setminus A} = w$.

We will use the notation $p(\cdot \mid \cdot)$ to represent the conditional measure of Y given a subvector of X . This notation is general and accommodates various scenarios:

1. Conditional probability in the discrete case, i.e., $\mathbb{P}(Y = y \mid X_{A \cup B} = (u, v, w))$.
2. Conditional density function in the continuous case, i.e., $f_Y(y \mid X_{A \cup B} = (u, v, w))$.
3. More generally, using measure-theoretic notation, for any Borel set \mathcal{B} in the sigma-algebra of Y and the event $X_{A \cup B} = (u, v, w)$, $p(\mathcal{B} \mid X_{A \cup B} = (u, v, w))$ represents the conditional expectation $\mathbb{E}[1_{\mathcal{B}} \mid \mathcal{F}_{X_{A \cup B}}]$ where $\mathcal{F}_{X_{A \cup B}}$ is the sigma-algebra generated by $X_{A \cup B}$. This accommodates cases where our random variables might not have purely discrete or continuous distributions.

Without loss of generality (WLOG), assume the vectors are coordinatewise connected through $B \setminus A$. Then, $w = w'$. So, we have,

$$\begin{aligned} p(Y \mid X_{(A \cap B)} = u, X_{(A \setminus B)} = v, X_{(B \setminus A)} = w, X_{(A \cup B)^c} = z) \\ &= p(Y \mid X_{(B \setminus A)} = w, X_{(A \cup B)^c} = z) \quad (\text{by the conditional independence } Y \perp\!\!\!\perp X_A \mid X_{A^c}) \\ &= p(Y \mid X_{(B \setminus A)} = w', X_{(A \cup B)^c} = z) \quad (\text{as } w = w') \\ &= p(Y \mid X_{(A \cap B)} = u', X_{(A \setminus B)} = v', X_{(B \setminus A)} = w', X_{(A \cup B)^c} = z). \end{aligned}$$

Instead of the vectors being coordinatewise connected through $B \setminus A$, if they were connected through $A \setminus B$, we would have used the conditional independence $Y \perp\!\!\!\perp X_B \mid X_{B^c}$. So, it follows that $p(Y \mid u, v, w, z) = p(Y \mid u', v', w', z)$ whenever the vectors (u, v, w) and (u', v', w') from the support of $X_{A \cup B}$ conditioned on $X_{(A \cup B)^c} = z$ are coordinatewise connected (via $A \setminus B$ and $B \setminus A$).

Now suppose (u, v, w) and (u', v', w') are equivalent (through $A \setminus B$ and $B \setminus A$). So there exist points (u_i, v_i, w_i) , $i = 1, \dots, L$ for some L such that $((u, v, w), (u_1, v_1, w_1))$, $((u_1, v_1, w_1), (u_2, v_2, w_2))$, \dots , $((u_{L-1}, v_{L-1}, w_{L-1}), (u_L, v_L, w_L))$ and $((u_L, v_L, w_L), (u', v', w'))$ are coordinatewise connected. From the above argument it follows that

$$p(Y \mid u, v, w, z) = p(Y \mid u_1, v_1, w_1, z_1) = \dots = p(Y \mid u_L, v_L, w_L, z_L) = p(Y \mid u', v', w', z).$$

This completes the proof.

B.2 Proof of Theorem 2.1

We begin by proving the following result first.

Proposition B.1. *Suppose there exists $i \neq j$ such that $Y \perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$. If $\mathcal{S}^c \in (\Delta \circ)^k \mathcal{I}$ for some $k \in \mathbb{Z}_{\geq 0}$ then there exists a unique set \mathcal{M} which satisfies the following:*

1. $Y \perp\!\!\!\perp X_{\mathcal{M}^c} \mid X_{\mathcal{M}}$
2. For all \mathcal{M}' such that $Y \perp\!\!\!\perp X_{\mathcal{M}'^c} \mid X_{\mathcal{M}'}$, $|\mathcal{M}| \leq |\mathcal{M}'|$

Further, this unique set is given by $\mathcal{M} = \mathcal{S}$.

Proof. Recall that we defined \mathcal{I} as $\mathcal{I} = \{\{i, j\} : \mathcal{H}_{i,j} \text{ is true}\}$ and \mathcal{S} was defined as $\mathcal{S} = \{i \in [p] : \mathcal{H}_{i,j} \text{ is false for all } j \in \{i\}^c\}$. Finally, $\Delta = \{\{A, B\} : X_{A \cup B} \perp\!\!\!\perp X_{(A \cup B)^c} \text{ has 1 equivalence class (through } A \setminus B \text{ and } B \setminus A)\}$. If $\mathcal{H}_{i,j}$ is true for some $i \neq j$, \mathcal{I} is non-empty.

Consider the case $|\mathcal{I}| = 1$, where $\mathcal{I} = \{\{i, j\}\}$. Define $\mathcal{M} = \{i, j\}^c$. \mathcal{M} satisfies condition 1 by definition of \mathcal{I} . Now for assumption 2, consider any \mathcal{M}' such that $Y \perp\!\!\!\perp X_{\mathcal{M}'^c} \mid X_{\mathcal{M}'}$. If $|\mathcal{M}'| < |\mathcal{M}| = p - 2$, for any three distinct elements of \mathcal{M}'^c , i', j' , and k' , by the weak union property (Pearl, 1988), $Y \perp\!\!\!\perp X_{\{i', j', k'\}} \mid X_{\{i', j', k'\}^c}$. Using the weak union property again, $\{i', j'\}$, $\{j', k'\}$, $\{k', i'\} \in \mathcal{I}$ which contradicts $|\mathcal{I}| = 1$. So, \mathcal{M} satisfies both conditions 1 and 2. To show that no other set satisfies these conditions, consider any other set \mathcal{M}' which satisfies them. Due to the previous argument $|\mathcal{M}'| \geq p - 2$. Also, $|\mathcal{M}'| < p - 1$ since \mathcal{M} satisfies condition 2. If $|\mathcal{M}'| = p - 2$ but $\mathcal{M}' \neq \mathcal{M}$, then this implies $\mathcal{M}'^c \in \mathcal{I}$ and $\mathcal{M}'^c \neq \{i, j\}$ which again contradicts \mathcal{I} . Therefore, when $|\mathcal{I}| = 1$ and $\mathcal{I} = \{\{i, j\}\}$, $\mathcal{M} = \{i, j\}^c$ is the unique set satisfying conditions 1 and 2. Also, in this case, $\mathcal{S} = \{i, j\}^c$ making $\mathcal{M} = \mathcal{S}$ which proves the statement of the proposition for the case $|\mathcal{I}| = 1$.

In the case $|\mathcal{I}| \geq 2$, consider any set $C \in \mathcal{I}$. Note that C is a set of two elements. For such a C , $Y \perp\!\!\!\perp X_C \mid X_{C^c}$. If we take any two sets C and D in \mathcal{I} such that $X_{C \cup D} \perp\!\!\!\perp X_{(C \cup D)^c}$ has 1 equivalence class (through $C \setminus D$ and $D \setminus C$), then Lemma 2.2 implies the following is true:

$$Y \perp\!\!\!\perp X_{C \cup D} \mid X_{(C \cup D)^c}.$$

Since Δ is defined as all the sets of sets with one equivalence class, the same logic implies that for any set $C \in \Delta \circ \mathcal{I}$, $Y \perp\!\!\!\perp X_C \mid X_{C^c}$. This argument can be applied to elements of $\Delta \circ \mathcal{I}$ also to get the following,

$$Y \perp\!\!\!\perp X_C \mid X_{C^c} \text{ for all } C \in (\Delta \circ)^2 \mathcal{I}.$$

Proceeding recursively, we get, for any $k \in \mathbb{Z}_{\geq 0}$,

$$Y \perp\!\!\!\perp X_C \mid X_{C^c} \text{ for all } C \in (\Delta \circ)^k \mathcal{I}.$$

So, if $\mathcal{S}^c \in (\Delta \circ)^k \mathcal{I}$ for some integer k , then $Y \perp\!\!\!\perp X_{\mathcal{S}^c} \mid X_{\mathcal{S}}$ and \mathcal{S} satisfies condition 1, and thus $|\mathcal{M}| \leq |\mathcal{S}|$ for any \mathcal{M} satisfying conditions 1 and 2. Now, conditions 1 and 2 imply \mathcal{M} is a Markov boundary. Hence, by Lemma 2.1, for any such \mathcal{M} , $\mathcal{S} \subseteq \mathcal{M}$. Combining them, we get under the assumptions of the proposition, $\mathcal{M} = \mathcal{S}$ is unique set satisfying conditions 1 and 2. \square

Now that we have proven the proposition, we shall use it along with Lemma 2.1 to prove Theorem 2.1. Under the assumptions of Theorem 2.1, we have two cases: $\mathcal{S} = [p]$ and $|\mathcal{S}| < p$. If $\mathcal{S} = [p]$, then $Y \not\perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$ for all $i \neq j$. For the sake of contradiction, suppose there exists a nontrivial Markov boundary, i.e., some set \mathcal{M} such that $|\mathcal{M}| < p - 1$ and $Y \perp\!\!\!\perp X_{\mathcal{M}^c} \mid X_{\mathcal{M}}$. Then there exists some $\{i, j\} \subseteq \mathcal{M}^c$ such that $i \neq j$. Then, the weak union property (Pearl, 1988) implies $Y \perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$ which contradicts $\mathcal{S} = [p]$. So, in this case there does not exist a nontrivial Markov boundary.

In the other case, that is when $|\mathcal{S}| < p$, there must exist some $i \neq j$ such that $Y \perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$. Now, consider any subset \mathcal{M}' of $[p]$ which is a Markov boundary and $|\mathcal{M}'| < p - 1$. So,

$Y \perp\!\!\!\perp X_{\mathcal{M}'^c} \mid X_{\mathcal{M}'}$. Lemma 2.1 and $|\mathcal{M}'^c| > 1$ imply $\mathcal{M}' \supseteq \mathcal{S}$. So, any nontrivial Markov boundary must be a superset of \mathcal{S} . The only thing left to show is that it is indeed equal to \mathcal{S} . But Proposition B.1 already tells us that \mathcal{S} is the smallest set such that $Y \perp\!\!\!\perp X_{\mathcal{S}} \mid X_{\mathcal{S}^c}$. If $\mathcal{M}' \supsetneq \mathcal{S}$, then \mathcal{M}' violates item 2 in Definition 2.1. Hence, \mathcal{M}' has to be equal to \mathcal{S} . So, any Markov boundary must either have size $= p - 1$ or be equal to \mathcal{S} which proves the Theorem.

B.3 Proof of Corollary 2.1

Corollary 2.1 presents an application of Theorem 2.1 for the case of continuous distributions with further assumptions on the support of the conditional distribution. First the case $\mathcal{S} = [p]$ is identical to that in Theorem 2.1 and is omitted in this proof.

For the case $|\mathcal{S}| < p$, we will apply Theorem 2.1 to prove the corollary. First, suppose that for any two subsets A and B of $[p]$ with $|B| = 2$ and $|A \cap B| = 1$, the support of the conditional distribution $X_{A \cup B} \mid X_{(A \cup B)^c} = c$ has only one equivalence class through $A \setminus B$ and $B \setminus A$; we will prove this supposition shortly. Then $\Delta = \{\{A, B\} : |A \cap B| = 1, |B| = 2\}$. Due to assumption (i), for any two (two element) sets C_1 and C_2 in \mathcal{I} there exists a sequence of sets D_1, \dots, D_k for some $k \in \mathbb{Z}_{\geq 0}$ such that $|C_1 \cap D_1| = 1$, $|D_1 \cap D_2| = 1$, $|D_2 \cap D_3| = 1$, \dots , $|D_{k-1} \cap D_k| = 1$ and $|D_k \cap C_2| = 1$. Consider all pairs of elements of \mathcal{I} . For each of them we get a sequence of sets $(D_i : i \in [k])$ connecting them for some $k \in \mathbb{Z}_{\geq 0}$. Consider k^* to be the maximum such k . Note that k^* is finite as it is the maximum over $\binom{|\mathcal{I}|}{2}$ integers. Therefore, $\mathcal{S}^c \in (\Delta \circ)^{k^*} \mathcal{I}$ which satisfies the assumptions of Theorem 2.1.

The only thing left to prove is that for any $A, B \subset [p]$ such that $|A \cap B| = 1$ and $|B| = 2$, the support of the conditional distribution $X_{A \cup B} \mid X_{(A \cup B)^c} = c$ has only one equivalence class. First note that if $c < 0$ or $c > 1$, the support is empty and when $c = 1$, the support is equal to the singleton $\{\mathbf{0}_{|A \cup B|}\}$ and hence, has a single equivalence class. So, we consider the only nontrivial case which is $0 \geq c < 1$. To do that, define the sets

$$\begin{aligned} \text{Supp}_{>0}(Z) &= \{z \in \mathbb{R}^p : f_Z(z) > 0\} \\ \text{and } \text{Supp}_{>0}(X) &= \left\{x \in \mathbb{R}^p : \sum_{j \in [p]} x_{\{j\}} = 1, f_X(x) > 0\right\}. \end{aligned}$$

Here f_Z is the density of Z and f_X is the density of X with respect to the $p-1$ -dimensional Hausdorff measure induced by f_Z . Note that $\text{Supp}_{>0}(X)$ represents a set related to the support of f_X , however with the difference that $\text{Supp}_{>0}(X)$ only includes points with positive density, as opposed to its closure which is commonly defined as the support of f_X . For some $C \subseteq [p]$ (think of C as $(A \cup B)^c$), and some $c \in \mathbb{R}^{|C|}$ such that there exists $x \in \text{Supp}_{>0}(X)$ with $x_C = c$, consider the following set

$$\text{Supp}_{>0}(X \mid X_C = c) = \{x \in \text{Supp}_{>0}(X) : x_C = c\}.$$

We want to show that $\text{Supp}_{>0}(X \mid X_C = c)$ is open and path connected in

$$\left\{x \in \mathbb{R}^p : \sum_{j \in [p]} x_{\{j\}} = 1, x_C = c\right\}.$$

To do that, we will be using the following Lemma (see Appendix B.9 for proof):

Lemma B.1. *Let Z be a p -dimensional random vector. Define the function $g : [0, \infty)^p \setminus \{\mathbf{0}_p\} \rightarrow [0, \infty)^p$ as follows*

$$x = g(z) := \left(\frac{z_{\{1\}}}{\sum_{k=1}^p z_{\{k\}}}, \frac{z_{\{2\}}}{\sum_{k=1}^p z_{\{k\}}}, \dots, \frac{z_{\{p\}}}{\sum_{k=1}^p z_{\{k\}}} \right).$$

Let f_Z denote the density of Z with respect to the Lebesgue measure and f_X that of X ($:= g \circ Z$) with respect to the $(p-1)$ -dimensional Hausdorff measure induced by f_Z . Then, if f_Z is continuous in $[0, \infty)^p \setminus \{\mathbf{0}_p\}$, $\text{Supp}_{>0}(X) = g(\text{Supp}_{>0}(Z))$.

Step 1: Path connectedness of conditional support We begin by demonstrating that the support is path-connected. Consider any two points x and x' in $\text{Supp}_{>0}(X \mid X_C = c)$. Since $x \in \text{Supp}_{>0}(X \mid X_C = c)$, it follows that $f_X(x) > 0$. By Lemma B.1, this implies there exists a $z \in \text{Supp}_{>0}(Z)$ such that $x = g(z)$, where g is defined as in the Lemma. Similarly, for x' , there exists a point $z' \in [0, \infty)^p \setminus \mathbf{0}_p$ such that $f_Z(z') > 0$ and $x' = g(z')$. Assumption (ii) implies that there exists a path $z(t)$ in $\{z \in \text{Supp}_{>0}(Z) : g(z)_C = c\}$ connecting z and z' . Formally, $z(\cdot)$ is a continuous map from $[0, 1]$ to $[0, \infty)^p \setminus \mathbf{0}_p$ such that $z(0) = z$, $z(1) = z'$, $z(t) \in \text{Supp}_{>0}(Z)$, and $g(z(t))_C = c$ for all $t \in [0, 1]$.

Define $x(t) = g(z(t))$ for $t \in [0, 1]$. For any $t \in [0, 1]$, since $z(t) \in \text{Supp}_{>0}(Z)$, by Lemma B.1, we have $x(t) \in \text{Supp}_{>0}(X)$. Additionally, given that $g(z(t))_C = c$, it follows that $x(t) \in \text{Supp}_{>0}(X \mid X_C = c)$. To demonstrate that $\text{Supp}_{>0}(X \mid X_C = c)$ is path-connected, it is sufficient to show that $x(\cdot)$ is continuous on $[0, 1]$. This holds because $x(\cdot) = g \circ z(\cdot)$ is the composition of two continuous functions. Therefore, x and x' are connected via the path $x(t)$, which proves that $\text{Supp}_{>0}(X \mid X_C = c)$ is path-connected in $\{x \in \mathbb{R}^p : \sum_{j=1}^p x_j = 1, x_C = c\}$.

Step 2: Openness of conditional support Now, we will show that $\text{Supp}_{>0}(X \mid X_C = c)$ is open in $\{x \in \Delta^{p-1} : x_C = c\}$. To show that, consider any point $x \in \text{Supp}_{>0}(X \mid X_C = c)$. By Lemma B.1, there exists $z \in \text{Supp}_{>0}(Z)$ such that $x = g(z)$. By assumption (ii), since f_Z is continuous in $[0, \infty)^p \setminus \{\mathbf{0}_p\}$, there exists an open ball $B_\epsilon(z)$ (open in $[0, \infty)^p \setminus \{\mathbf{0}_p\}$) such that $B_\epsilon(z) \subseteq \text{Supp}_{>0}(Z)$. Define

$$B_{\frac{\epsilon}{t}}(x) = \left\{ x' \in \Delta^{p-1} : x'_C = c, \|x - x'\| < \frac{\epsilon}{t} \right\}, \quad t = \sum_{j=1}^p z_j.$$

It is enough to show that $B_{\frac{\epsilon}{t}}(x) \subseteq \text{Supp}_{>0}(X \mid X_C = c)$. Consider some $x' \in B_{\frac{\epsilon}{t}}(x)$. Define $z' = t \cdot x'$. Since $\|z - z'\| = \|tx - tx'\| = t\|x - x'\| < t \cdot \frac{\epsilon}{t} = \epsilon$, $z' \in B_\epsilon(z)$ and hence $z' \in \text{Supp}_{>0}(Z)$. Therefore, by Lemma B.1, $x' \in \text{Supp}_{>0}(X)$. Finally, since $x'_C = c$, $x' \in \text{Supp}_{>0}(X \mid X_C = c)$. This shows that $\text{Supp}_{>0}(X \mid X_C = c)$ is open in $\{x \in \Delta^{p-1} : x_C = c\}$.

Step 3: Proving uniqueness of equivalence class of conditional support Now that we have shown $\text{Supp}_{>0}(X \mid X_C = c)$ is open and path connected in

$\{x \in \mathbb{R}^p : \sum_{j \in [p]} x_{\{j\}} = 1, x_C = c\}$ for all $C \subseteq [p]$, and all $c \in \mathbb{R}^{|C|}$ such that there exists $x \in \text{Supp}_{>0}(X)$ with $x_C = c$, we shall show that for any $A, B \subset [p]$ such that $|A \cap B| = 1$ and $|B| = 2$ and for all $c \in \mathbb{R}^{p-|A \cup B|}$ such that $\sum_{j \in [p-|A \cup B|]} c_{\{j\}} < 1$, $X_{A \cup B} \mid X_{(A \cup B)^c} = c$ has only one equivalence class through $A \setminus B$ and $B \setminus A$. Note that we only need to consider 2-sized sets B with one intersection with A because we want to “combine” the sets $\{i, j\}$ for $\{i, j\} \in \mathcal{I}$. Consider any two points x and x' in $\text{Supp}_{>0}(X)$ such that $x_{(A \cup B)^c} = x'_{(A \cup B)^c} = c$. We shall show that x and x' are equivalent through $A \setminus B$ and $B \setminus A$. Since $\text{Supp}_{>0}(X \mid X_{(A \cup B)^c} = c)$ is path connected in $\{x \in \mathbb{R}^p : \sum_{j \in [p]} x_{\{j\}} = 1, x_C = c\}$, there exists a path $x(t)$ in the latter set connecting x and x' . Since $\text{Supp}_{>0}(X \mid X_{(A \cup B)^c} = c)$ is open in $\{x \in \mathbb{R}^p : \sum_{j \in [p]} x_{\{j\}} = 1, x_C = c\}$, there exist open balls $B_{\epsilon_t}(x(t))$ (open in $\{x \in \mathbb{R}^p : \sum_{j \in [p]} x_{\{j\}} = 1, x_C = c\}$) for all $t \in [0, 1]$, such that $B_{\epsilon_t}(x(t)) \subseteq \text{Supp}_{>0}(X \mid X_{(A \cup B)^c} = c)$ for all $t \in [0, 1]$. Here the balls are defined with respect to the standard euclidean norm (or the L^2 norm). Now consider the sets $B_{\frac{\epsilon_t}{\sqrt{2}}}(x(t))$ for $t \in [0, 1]$. $\bigcup_{t \in [0, 1]} B_{\frac{\epsilon_t}{\sqrt{2}}}(x(t))$ is an open cover of $\{x(t) : t \in [0, 1]\}$. Also, since $[0, 1]$ is compact and $x(t)$ is continuous, $\{x(t) : t \in [0, 1]\}$ is compact in $\text{Supp}_{>0}(X \mid X_{(A \cup B)^c} = c)$. This implies there exists a finite subcover B_1, \dots, B_n of $\{x(t) : t \in [0, 1]\}$. Call the respective centers x_{t_1}, \dots, x_{t_n} . Remember that for all $t \in \{t_1, \dots, t_n\}$, the ball centered at $x(t)$ has L^2 radius $\frac{\epsilon_t}{\sqrt{2}}$. WLOG assume B_1 is centered at x and B_n is centered at x' . Call the rest of the centers y_1, \dots, y_{n-2} . If we can show that the centers of adjacent open balls are equivalent, then we will have shown x and x' are equivalent. To show that, consider any two balls B_j and B_{j+1} with centers w and w' , respectively. Since B_j and B_{j+1} are adjacent open balls, they must have a non-empty intersection. Consider a particular point $w^* \in B_j \cap B_{j+1}$. If we can show w and w^* are equivalent, we can use the same argument to show w^* and w' are equivalent, thereby showing the equivalence of w and w' , which would complete the proof. Since

$$\left(\sum_{j \in A \setminus B} w_{\{j\}} + w_{B \setminus A}^* \right) + \left(\sum_{j \in A \setminus B} w_{\{j\}}^* + w_{B \setminus A} \right) \leq \sum_{j \in [p]} w_{\{j\}} + \sum_{j \in [p]} w_{\{j\}}^* = 2,$$

either $\sum_{j \in A \setminus B} w_{\{j\}} + w_{B \setminus A}^* \leq 1$ or $\sum_{j \in A \setminus B} w_{\{j\}}^* + w_{B \setminus A} \leq 1$.

First, consider the case when $\sum_{j \in A \setminus B} w_{\{j\}} + w_{B \setminus A}^* \leq 1$. Introduce the notation $w = (w_{A \setminus B}, w_{A \cap B}, w_{B \setminus A}, c)$ and $w' = (w'_{A \setminus B}, w'_{A \cap B}, w'_{B \setminus A}, c)$ and define

$$w^\dagger = \left(w_{A \setminus B}, 1 - \sum_{j \in A \setminus B} w_{\{j\}} - w_{B \setminus A}^* - \sum_{j \in [p-|A \cup B|]} c_j, w_{B \setminus A}^*, c \right).$$

Then the L^2 distance between w and w^\dagger is as follows

$$\begin{aligned}
\|w - w^\dagger\|_2 &= \sqrt{\sum_{j \in [p]} (w_{\{j\}} - w_{\{j\}}^\dagger)^2} \\
&= \sqrt{\left(w_{A \cap B} - \left(1 - \sum_{j \in A \setminus B} w_{\{j\}} - w_{B \setminus A}^* - \sum_{j \in [p - |A \cap B|]} c_j \right) \right)^2 + \left(w_{B \setminus A} - w_{B \setminus A}^* \right)^2} \\
&= \sqrt{2 \cdot \left(w_{B \setminus A} - w_{B \setminus A}^* \right)^2} \\
&\leq \sqrt{2 \cdot \sum_{j \in [p]} \left(w_{\{j\}} - w_{\{j\}}^* \right)^2} \\
&\leq \sqrt{2 \cdot \frac{\epsilon_{t_j}^2}{2}} \quad \left(\text{Since } w^* \in B_{\frac{\epsilon}{\sqrt{2}}}(w) \right) \\
&= \epsilon_{t_j}
\end{aligned}$$

Since $\|w - w^\dagger\|_2 \leq \epsilon_{t_j}$, $w^\dagger \in B_{\epsilon_{t_j}}(x(t_j))$ which implies $f_X(w^\dagger) > 0$ or in other words, $w^\dagger \in \text{Supp}_{>0}(X \mid X_C = c)$. Since w and w^\dagger are coordinatewise connected through $A \setminus B$ and w^\dagger and w^* are connected through $B \setminus A$, w and w^* are equivalent through $A \setminus B$ and $B \setminus A$.

Now, consider the other case: $\sum_{j \in A \setminus B} w_{\{j\}}^* + w_{B \setminus A} \leq 1$. In this case, define w^\dagger as follows

$$w^\dagger = \left(w_{A \setminus B}^*, 1 - \sum_{j \in A \setminus B} w_{\{j\}}^* - w_{B \setminus A} - \sum_{j \in [p - |A \cup B|]} c_j, w_{B \setminus A}, c \right).$$

With this w^\dagger the L^2 distance can be bounded as

$$\begin{aligned}
\|w - w^\dagger\|_2 &= \sqrt{\sum_{j \in [p]} (w_{\{j\}} - w_{\{j\}}^\dagger)^2} \\
&= \sqrt{\sum_{j \in A \setminus B} \left(w_{\{j\}} - w_{\{j\}}^* \right)^2 + \left(w_{A \cap B} - \left(1 - \sum_{j \in A \setminus B} w_{\{j\}}^* - w_{B \setminus A} - \sum_{j \in [p - |A \cap B|]} c_j \right) \right)^2} \\
&= \sqrt{\sum_{j \in A \setminus B} \left(w_{\{j\}} - w_{\{j\}}^* \right)^2 + \left(w_{A \cap B} - w_{A \cap B}^* + w_{B \setminus A} - w_{B \setminus A}^* \right)^2} \\
&\leq \sqrt{\sum_{j \in A \setminus B} \left(w_{\{j\}} - w_{\{j\}}^* \right)^2 + 2(w_{A \cap B} - w_{A \cap B}^*)^2 + 2(w_{B \setminus A} - w_{B \setminus A}^*)^2} \\
&\leq \sqrt{2 \cdot \sum_{j \in [p]} \left(w_{\{j\}} - w_{\{j\}}^* \right)^2} \\
&\leq \sqrt{2 \cdot \frac{\epsilon_{t_j}^2}{2}} \quad \left(\text{Since } w^* \in B_{\frac{\epsilon}{\sqrt{2}}}(w) \right) \\
&= \epsilon_{t_j}
\end{aligned}$$

Using the same argument as in the previous case, $\|w - w^\dagger\|_2 \leq \epsilon_{t_j}$ implies $w^\dagger \in B_{\epsilon_{t_j}}(x(t_j))$ which implies $w^\dagger \in \text{Supp}_{>0}(X \mid X_C = c)$. Further, since w and w^\dagger are coordinatewise connected through $B \setminus A$ and w^\dagger and w^* are connected through $A \setminus B$, w and w^* are equivalent through $A \setminus B$ and $B \setminus A$. So, in both the cases,

$$w \text{ and } w^* \text{ are equivalent through } A \setminus B \text{ and } B \setminus A.$$

Using a similar construction, we can show that w^* and w' are equivalent through $A \setminus B$ and $B \setminus A$. Applying this argument recursively, we can show that all the centers of the balls B_1, \dots, B_n are equivalent, thereby proving x and x' are equivalent (through $A \setminus B$ and $B \setminus A$). Since x and x' are chosen to be any two points in $\text{Supp}_{>0}(X \mid X_C = c)$, this proves that $\text{Supp}_{>0}(X \mid X_C = c)$ has only one equivalence class. This proves the conditions of Theorem 2.1 which completes the proof of the Corollary.

B.4 Proof of Corollary 2.2

First, consider the case when $F_k \subseteq \mathcal{S}_F$. Suppose, for the sake of contradiction, there exists a nontrivial Markov boundary \mathcal{M} i.e., $|\mathcal{M}^c \cap F_k| \geq 2$. Then there exists $i \neq j$ such that $\{i, j\} \subseteq \mathcal{M}^c \cap F_k$. By the weak union property (Pearl, 1988), $Y \perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$. Therefore $i \notin \mathcal{S}_F$ and $j \notin \mathcal{S}_F$ which implies $|F_k \setminus \mathcal{S}_F| \geq 2$ which contradicts $F_k \subseteq \mathcal{S}_F$.

Now, in the other case, consider some nontrivial Markov boundary \mathcal{M} . To show $\mathcal{M} = \mathcal{S}_F \cap F_k$ we will follow the following steps:

1. Δ includes sets of sets of the following form:
 - (a) $\{A, B\}$ where both $A, B \subseteq F_k$ for $|A \cap B| \geq 1$.
 - (b) $\{A, B\}$ where $A \subseteq F_k$ and $B \subseteq F_k^c$.
2. For any Markov boundary \mathcal{M} with $|\mathcal{M}^c \cap F_k| \geq 2$, $\mathcal{M}_k \supseteq \mathcal{S}_F \cap F_k$.
3. If $|\mathcal{M}^c \cap F_k| \geq 2$, $\mathcal{M}_k = \mathcal{S}_F \cap F_k$.

Step 1: Δ contains sets of form $\{A, B\}$, $A, B \subseteq F_k$ Consider two subsets A and B of F_k such that $|A \cap B| \geq 1$. Given $X_{(A \cup B)^c} = c$, let us consider the support of $X_{A \cup B}$. If $\sum_{j \in (A \cup B)^c \cap F_k} c_j = 1$, then the support of $X_{A \cup B} \mid X_{(A \cup B)^c} = c$ only includes the $\mathbf{0}_{|A \cup B|}$ vector. So, it trivially contains a single equivalence class. Conversely, when $\sum_{j \in (A \cup B)^c \cap F_k} c_j = 0$, the support of $X_{A \cup B} \mid X_{(A \cup B)^c}$ contains all the canonical vectors of length $|A \cup B|$. This is true because the rows of X are randomly selected such that all treatment combinations have positive probability. Consider any two vectors x and y in the support of $X_{A \cup B} \mid X_{(A \cup B)^c}$. Note that, since both vectors x and y have a single non zero entry ($=1$), there can only be one of the two following cases, with the roles of x and y being reversible:

- (a) $x_{A \setminus B} = y_{A \setminus B} = \mathbf{0}_{|A \setminus B|}$ or $x_{B \setminus A} = y_{B \setminus A} = \mathbf{0}_{|B \setminus A|}$.
- (b) $\sum_{j \in A \setminus B} x_{\{j\}} = 1, y_{A \setminus B} = \mathbf{0}_{|A \setminus B|}$ and $x_{B \setminus A} = \mathbf{0}_{|B \setminus A|}, \sum_{j \in B \setminus A} y_{\{j\}} = 1$.

In case (a), the two vectors x and y are connected through $A \setminus B$ and $B \setminus A$ respectively. This shows that any two vectors in the support of $X_{A \cup B} \mid X_{(A \cup B)^c}$ are coordinatewise

connected through $A \setminus B$ and $B \setminus A$, thereby proving that support of $X_{A \cup B} \mid X_{(A \cup B)^c}$ has only one equivalence class (through $A \setminus B$ and $B \setminus A$).

In case (b), consider a canonical vector z with $\sum_{j \in A \cap B} z_{\{j\}} = 1$. If $|A \cap B| > 1$, then z could be any canonical vector with $\sum_{j \in A \cap B} z_{\{j\}} = 1$ and if $|A \cap B| = 1$, then z is the canonical vector with $z_{A \cap B} = 1$ and zeros in all other indices. Then, $x_{B \setminus A} = z_{B \setminus A}$ and $y_{A \setminus B} = z_{A \setminus B}$. Note that z belongs to the support of $X_{A \cup B} \mid X_{(A \cup B)^c}$. This shows that x and y are equivalent, which shows that support of $X_{A \cup B} \mid X_{(A \cup B)^c}$ has only one equivalence class (through $A \setminus B$ and $B \setminus A$).

Therefore, in both cases (a) and (b), the support of $X_{A \cup B} \mid X_{(A \cup B)^c}$ has only one equivalence class proving Δ contains all sets of the form $\{A, B\}$ where $|A \cap B| \geq 1$ and $A, B \subseteq F_k$.

Step 2: Δ contains sets of the form $\{A, B\}$, $A \subseteq F_k$, $B \subseteq F_k^c$ Consider $A \subseteq F_k$ and $B \subseteq F_k^c$. Without loss of generality, we can assume that both A and B are non-empty. If both A and B are singletons, then, the support of $X_{A \cup B} \mid X_{(A \cup B)^c}$ only has one element which is determined by $X_{(A \cup B)^c}$. Consider the case when one of A and B is a singleton. Suppose, A is a singleton. Then for all points in the support of $X_{A \cup B} \mid X_{(A \cup B)^c}$ will be either have $x_A = 0$ or $x_A = 1$. In both the cases, any two points in the support is coordinatewise connected through $A \setminus B (= A)$. The case when B is a singleton is same as then B is a proper subset of some other factor (k') and A is a subset of $F_{k'}^c$.

So, the only case left is when both $|A| \geq 2$ and $|B| \geq 2$. Consider any point x in the support of $X_{A \cup B} \mid X_{(A \cup B)^c} = c$ for some vector c in the support of $X_{(A \cup B)^c}$. If $\sum_{j \in A^c \cap F_k} x_{\{j\}} = 1$, then for all points in the support, $x_A = \mathbf{0}_{|A|}$ making the support coordinatewise connected through $A \setminus B (= A)$. If $\sum_{j \in A^c \cap F_k} x_{\{j\}} = 0$, then for any point x in the support $x_A = e_i$ for some i . Here e_i represent canonical vectors of length $|A|$. Consider any two points x and y in the support of $X_{A \cup B} \mid X_{(A \cup B)^c} = c$. Define $z = (x_A, y_B)$. Since z satisfies the criteria $\sum_{j \in F_k} z_{\{j\}} = 1$, z lies in the support. x and z are coordinatewise connected through $A \setminus B (= A)$ and y and z are coordinatewise connected through $B \setminus A (= B)$. This implies x and y are equivalent which proves the support of $X_{A \cup B} \mid X_{(A \cup B)^c}$ has only one equivalence class. Therefore Δ contains all sets of the form $\{A, B\}$ such that $A \subseteq F_k$ and $B \subseteq F_k^c$.

Step 3: For \mathcal{M} with $|\mathcal{M}^c \cap F_k| \geq 2$, $\mathcal{M}_k \supseteq \mathcal{S}_F \cap F_k$ Consider any Markov boundary such that $|\mathcal{M}^c \cap F_k| \geq 2$ and some $i \in \mathcal{M}^c \cap F_k$. Since $|\mathcal{M}^c \cap F_k| \geq 2$, there exists some $j \in (\mathcal{M}^c \cap F_k) \setminus \{i\}$. By the weak union property, since $\{i, j\} \subseteq \mathcal{M}^c$, $Y \perp\!\!\!\perp X_{\mathcal{M}^c} \mid X_{\mathcal{M}}$ implies $Y \perp\!\!\!\perp X_{\{i, j\}} \mid X_{\{i, j\}^c}$. This means $i \in \mathcal{S}_F^c \cap F_k$. Therefore, $\mathcal{M}^c \cap F_k \subseteq \mathcal{S}_F^c \cap F_k$ which proves $\mathcal{M}_k = \mathcal{M} \cap F_k \supseteq \mathcal{S}_F \cap F_k$.

Step 4: For \mathcal{M} with $|\mathcal{M}^c \cap F_k| \geq 2$, $\mathcal{M}_k = \mathcal{S}_F \cap F_k$ To show this, it is enough to show that $\mathcal{M}_k \not\supseteq \mathcal{S}_F \cap F_k$. We will first begin by showing that

$$Y \perp\!\!\!\perp X_{\mathcal{S}_F^c \cap F_k} \mid X_{(\mathcal{S}_F \cap F_k) \cup F_k^c}.$$

This is equivalent to showing $\mathcal{S}_F^c \cap F_k \in (\Delta_0)^l \mathcal{I}$ for some finite l . Consider the sets in \mathcal{I} which have both the elements in F_k . We call the collection of those sets \mathcal{I}_k . Formally,

$$\mathcal{I}_k = \mathcal{I} \cap \{\{i, j\} : i, j \in F_k\}.$$

Then $\mathcal{S}_F^c \cap F_k$ can be written as the union of all the sets in \mathcal{I}_k . Consider any $A, B \in \mathcal{I}_k$. By step 1, if $|A \cap B| \geq 1$, $\{A, B\} \in \Delta$. Since all the elements in \mathcal{I}_k are assumed to be connected by a sequence of true bivariate conditional hypotheses (of the form $\mathcal{H}_{i,j} : \{i, j\} \subseteq F_k$), we have $\mathcal{S}^c \cap F_k \in (\Delta_0)^{l_k} \mathcal{I}_k$ where $l_k = |\mathcal{I}_k|$. Therefore,

$$Y \perp\!\!\!\perp X_{\mathcal{S}_F^c \cap F_k} \mid X_{(\mathcal{S}_F \cap F_k) \cup F_k^c}.$$

Now, \mathcal{M} is a Markov boundary implies $Y \perp\!\!\!\perp X_{\mathcal{M}^c} \mid X_{\mathcal{M}}$ which further implies, by the weak union property,

$$Y \perp\!\!\!\perp X_{\mathcal{M}^c \cap F_k^c} \mid X_{\mathcal{M} \cup F_k}.$$

Define $A = \mathcal{S}_F^c \cap F_k$ and $B = \mathcal{M}^c \cap F_k^c$. Since A and B satisfy the form of sets considered in Step 2, $\{A, B\} \in \Delta$. Therefore $Y \perp\!\!\!\perp X_{A \cup B} \mid X_{A^c \cap B^c}$. Putting the expressions of A and B we get,

$$Y \perp\!\!\!\perp X_{(\mathcal{S}_F^c \cap F_k) \cup (\mathcal{M}^c \cap F_k^c)} \mid X_{(\mathcal{S}_F \cap F_k) \cup (\mathcal{M} \cup F_k^c)}.$$

So, $(\mathcal{S}_F \cap F_k) \cup (\mathcal{M} \cup F_k^c)$ is a Markov blanket. If $\mathcal{S}_F \cap F_k \subsetneq \mathcal{M} \cap F_k$, then $(\mathcal{S}_F \cap F_k) \cup (\mathcal{M} \cup F_k^c) \subsetneq \mathcal{M}$ which is a contradiction to the definition of Markov boundary. Therefore $\mathcal{S}_F \cap F_k = \mathcal{M} \cap F_k$. This completes the proof of Step 4 and hence, Corollary 2.2.

B.5 Proof of Theorem 3.4

If $\mathcal{S} = [p]$, $Y \not\perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$ for all $i \neq j$. In particular, $Y \not\perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$ for all $\{i, j\} \subseteq \mathcal{D}$. Therefore, $\mathcal{S}_{\mathcal{D}} = \mathcal{D}$ which proves first part of the Theorem. For the other case, first note that $\mathcal{S}_{\mathcal{D}} \supseteq \mathcal{S} \cap \mathcal{D}$. When $\mathcal{S}^c \cap \mathcal{D} = \emptyset$, $\mathcal{D} \subseteq \mathcal{S}$ which implies $\mathcal{S}_{\mathcal{D}} \subseteq \mathcal{S} \cap \mathcal{D}$. Combining them we get, $\mathcal{S}_{\mathcal{D}} = \mathcal{S} \cap \mathcal{D}$.

Now consider the case $|\mathcal{S}^c \cap \mathcal{D}| > 1$. For the sake of contradiction, suppose there exists $i \in \mathcal{S}_{\mathcal{D}} \setminus \mathcal{S}$. Since $|\mathcal{S}^c \cap \mathcal{D}| \geq 2$, there exists $j \neq i$ such that $j \in \mathcal{D} \setminus \mathcal{S}$. Therefore $\{i, j\} \subseteq \mathcal{S}^c$. Since the conditions of Theorem 2.1 are satisfied in this case, we have $\mathcal{S} = \mathcal{M}$ and hence, $Y \perp\!\!\!\perp X_{\mathcal{S}^c} \mid X_{\mathcal{S}}$. By the weak union property (Pearl, 1988) $Y \perp\!\!\!\perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$. This implies $i \notin \mathcal{S}_{\mathcal{D}}$ which is a contradiction. Therefore, $\mathcal{S}_{\mathcal{D}} \subseteq \mathcal{S} \cap \mathcal{D}$. Since we already showed that $\mathcal{S}_{\mathcal{D}} \supseteq \mathcal{S} \cap \mathcal{D}$, $\mathcal{S}_{\mathcal{D}} = \mathcal{S} \cap \mathcal{D}$.

Finally, as we already mentioned, the conditions of Theorem 2.1 are satisfied here. So, $\mathcal{S} = \mathcal{M}$. Combining everything, $\mathcal{S}_{\mathcal{D}} = \mathcal{S} \cap \mathcal{D} = \mathcal{M} \cap \mathcal{D}$ which proves the Theorem.

B.6 Proof of Theorem 3.1

Since we are considering Bonferroni's p-values after $h-1$ rejections (call the corresponding rejection set $\hat{\mathcal{S}}$), we have the following p-values to consider:

$$P_j^B(\hat{\mathcal{S}}) = (p - \bar{s}) P_{(\bar{s}-h+1),j}(\hat{\mathcal{S}}) \text{ for all } j \in \hat{\mathcal{S}}^c.$$

Now, these p-values are valid for the null variables j (i.e., $j \in \mathcal{S}^c$) if there are at least $p - \bar{s} - 1$ null p-values in the vector $P_{\hat{\mathcal{S}}^c \setminus \{j\},j}$. If there are no null variables in the rejection

set after the $h - 1^{\text{th}}$ step— $\hat{\mathcal{S}}$, then there are indeed $p - s - 1 (\geq p - \bar{s} - 1)$ null p-values in this vector which is ensured by our assumption.

So, $P_j^{\text{B}}(\hat{\mathcal{S}})$ are valid for all null j if and only if $\hat{\mathcal{S}}$ does not contain any null variables. Coming back to the proof of validity of the Bonferroni–Holm method (described in Algorithm 1), this method introduces an order among the hypotheses $\mathcal{H}_j : j \in [p]$ which is the order in which each column is considered for rejection (if we were to reject all the variables). Let the hypotheses arranged according to this order be represented as $\mathcal{H}_{(1)}, \mathcal{H}_{(2)}, \dots, \mathcal{H}_{(p)}$. Suppose $\mathcal{H}_{(h)}$ is the first falsely rejected null hypothesis and let us denote the corresponding variable as j^* . Since j^* is a null variable and all the $h - 1$ variables corresponding to the hypotheses $\mathcal{H}_{(1)}, \mathcal{H}_{(2)}, \dots, \mathcal{H}_{(h-1)}$ which were rejected prior to it were non-null, Bonferroni’s p-value $P_{j^*}^{\text{B}}(\hat{\mathcal{S}})$ is superuniform. So, we have,

$$P_{j^*}^{\text{B}}(\hat{\mathcal{S}}) \succeq \text{Uniform}[0, 1].$$

Let us now consider the cutoff values. Since only non-null variables are rejected before step h , $h < s + 1$. This implies $p - h + 1 \geq p - s$. Therefore, j^* is rejected implies $P_{j^*}^{\text{B}}(\hat{\mathcal{S}}) \leq \frac{\alpha}{p - h + 1} \leq \frac{\alpha}{p - s}$. Using these results, we can upper bound the FWER as follows:

$$\begin{aligned} \text{FWER} &= \mathbb{P}(V > 1) = \mathbb{P}\left(\bigcup_{j \in \mathcal{S}^c} \{j \text{ is rejected}\}\right) = \mathbb{P}\left(P_{j^*}^{\text{B}}(\hat{\mathcal{S}}) \leq \frac{\alpha}{p_0}\right) \\ &= \mathbb{P}\left(\bigcup_{j \in \mathcal{S}^c} \left\{P_j^{\text{B}}(\hat{\mathcal{S}}) \leq \frac{\alpha}{p_0}\right\}\right) \leq \sum_{j \in \mathcal{S}^c} \mathbb{P}\left(P_j^{\text{B}}(\hat{\mathcal{S}}) \leq \frac{\alpha}{p_0}\right) \\ &\leq \sum_{j \in \mathcal{S}^c} \frac{\alpha}{p - s} = \alpha. \end{aligned}$$

B.7 Proof of Theorem 3.2

The proof of this theorem follows the exact same structure as the proof of Theorem 3.1. The only difference is that, we are using Simes’ p-values here instead of Bonferroni’s p-values. For any column j , after $h - 1$ rejections and the corresponding rejection set $\hat{\mathcal{S}}$, we have the following p-values to consider:

$$P_j^{\text{S}}(\hat{\mathcal{S}}) := \min_{\bar{s}-h+1 \leq i \leq p-h} \left\{ \frac{p - \bar{s}}{i - \bar{s} + h} P_{(i),j}(\hat{\mathcal{S}}) \right\}$$

For a null variable j , this p-value is valid when there are at least $p - \bar{s} - 1$ null p-values in the vector $P_{[p] \setminus \{\hat{\mathcal{S}} \cup \{j\}\}, j}$ and $P_{[p] \setminus \{\hat{\mathcal{S}} \cup \{j\}\}, j}$ satisfy the PRDS assumption (see Appendix B.8 for definition) which is true by the assumptions of this Theorem. Therefore, we only need to show that there are at least $p - \bar{s} - 1$ p-values in $P_{[p] \setminus \{\hat{\mathcal{S}} \cup \{j\}\}, j}$. This is exactly what we have showed in Appendix B.6. So, using a similar argument, we can show that $\text{FWER} \leq \alpha$.

B.8 Proof of Theorem 3.3

Simes p-values are defined as follows:

$$P_j^S = \min_{s \leq i \leq p-1} \frac{p - \bar{s}}{i - \bar{s} + 1} P_{(i),j}.$$

Simes p-values are superuniform under the partial conjunction null assuming the base p-values are PRDS (Benjamini and Heller, 2008, Theorem 1). Formally, the PRDS assumption, defined in Benjamini and Yekutieli (2001), is as follows:

Definition B.1 (Positive Regression Dependency on the Subset (PRDS)). *(Benjamini and Yekutieli, 2001) We say that p-values $P' := (P'_1, P'_2, \dots, P'_m)$ satisfy the PRDS property on a subset $\mathbf{I}_0 \subseteq [m]$ if, for any increasing set $\mathbf{D} \in [0, 1]^m$ and each $j \in \mathbf{I}_0$, the probability $\mathbb{P}(P' \in \mathbf{D} \mid P'_j = x)$ is non-decreasing in x . Here, \mathbf{D} is called a non-decreasing set if $x \in \mathbf{D}$ and $y \geq x$ implies that $y \in \mathbf{D}$ as well. We simply say that the p-values P' are PRDS if the p-values satisfy are PRDS on the the entire set $[m]$.*

To prove the validity of our method, it is enough to show that the PRDS assumption on the base p-values implies that the Simes p-values are also PRDS (Benjamini and Yekutieli, 2001, Theorem 1.2). We will do so using Bogomolov (2021, Theorem 3.1). Specifically, we will be examining conditions outlined in item 1 of the Theorem—we restate a special case of their Theorem below in our notation, but first we need a few new definitions.

Definition B.2 (Blanchard and Roquain (2008)). *A **threshold collection** Δ is function*

$$\Delta : (j, r) \in [m] \times \mathbb{R}^+ \rightarrow \Delta(j, r) \in \mathbb{R}^+,$$

*which is non-decreasing in its second input. A **factored threshold collection** is a threshold collection of the form*

$$\Delta(j, r) = \frac{\alpha w_j \beta(r)}{m} \quad \forall (j, r) \in [m] \times \mathbb{R}^+,$$

*where w_j is a prior weight for the i^{th} hypothesis and $\beta : \mathbb{R}^+ \rightarrow \mathbb{R}^+$ is a non-decreasing function called a shape function. Given a threshold collection Δ , the **Δ -thresholding-based multiple testing procedure** at rejection volume r is a multiple testing procedure that has the following rejection set when applied to the vector of p-values $P' = \{P'_1, \dots, P'_m\}$:*

$$L_\Delta(r) = \{j \in [m] : P'_j \leq \Delta(j, r)\}.$$

Definition B.3 (Bogomolov (2021)). *Consider a multiple testing procedure applied on a set of p-values $\{P'_j : j \in [m]\}$ corresponding to hypotheses $\{\mathcal{H}_j : j \in [m]\}$ with the set of indices of rejected hypotheses $\mathcal{R} = \mathcal{R}(P') \subseteq [m]$.*

1. *The procedure is called **self consistent** with respect to threshold collection Δ if the rejection set \mathcal{R} satisfies $\mathcal{R} \subseteq L_\Delta(|\mathcal{R}|)$ almost surely.*
2. *The procedure is **non-increasing** if $|\mathcal{R}(P')|$ is non-increasing in each p-value P'_j , $j \in [m]$.*

3. The procedure is said to follow the **natural monotonicity property** if given a set of rejected hypotheses \mathcal{R}_1 , decreasing a certain p-value will result in rejecting a superset \mathcal{R}_2 .

Definition B.4 (Wright (1992)). *The adjusted p-value for a particular hypothesis within a collection of hypotheses of a certain multiple testing procedure is the smallest overall significance level at which the multiple testing procedure rejects said particular hypothesis.*

When the multiple testing procedure is Benjamini–Hochberg at level α (BH) applied on a set of p-values $\{P'_j : j \in [m]\}$, the j^{th} adjusted p-value is

$$\max_{k \in [m]: P'_k \geq P'_j} \left\{ \frac{P'_k \cdot m}{\text{rank}(P'_k; P')} \right\},$$

where $\text{rank}(P'_k; P')$ is the rank of P'_k in P' (Yekutieli and Benjamini, 1999).

Theorem B.1 (Bogomolov (2021), Theorem 3.1). *Consider M base null hypotheses with p-values P'_1, \dots, P'_M , grouped into m groups A_j , for $j \in [m]$, with $|A_j| = m_j$. Let \mathcal{P} be a self-consistent multiple testing procedure with respect to thresholds of the form $\Delta(j, r) = \alpha r / m \forall j \in [m]$. Consider a family of partial conjunction hypotheses $\{\mathcal{H}_{0,j}^{r_j/m_j}, j \in [m]\}$ for the m groups of p-values. For each $j \in [m]$, let \mathcal{P}_j be a multiple testing procedure satisfying the natural monotonicity property and for each $j \in [m]$, the partial conjunction p-value $P_j^{r_j/m_j}$ is connected to \mathcal{P}_j in the following way:*

$$P_j^{r_j/m_j} = \max \left\{ P_j^{1/(m_j - r_j + 1)}[A] : A \subseteq A_j, |A| = m_j - r_j + 1 \right\},$$

where $P_j^{1/(m_j - r_j + 1)}[A]$ is the minimum adjusted p-value according to \mathcal{P}_j applied on the p-values with indices in A . Suppose the procedure \mathcal{P} , the procedures $\{\mathcal{P}_j, j \in [m]\}$, and the base p-values satisfy the following conditions:

- (a) The p-values $\{P'_1, \dots, P'_M\}$ are PRDS on the set of true base null hypotheses.
- (b) The multiple testing procedure \mathcal{P} is non-increasing.
- (c) For each $j \in [m]$, the multiple testing procedure \mathcal{P}_j is self-consistent with respect to thresholds of form $\Delta(i, r) = \alpha r / m_j$.

Then, if \mathbf{I}_0 denotes the set of true partial conjunction hypotheses $\mathcal{H}_{0,j}^{r_j/m_j}$ and \mathcal{R} denotes the set of rejections when the procedure \mathcal{P} is applied on $\{P_j^{r_j/m_j}, j \in [m]\}$ at level α , the FDR of the resulting procedure

$$\mathbb{E} \left[\frac{|\mathcal{R} \cap \mathbf{I}_0|}{|\mathcal{R}|} \right] \leq \alpha \frac{|\mathbf{I}_0|}{m} \leq \alpha.$$

For Algorithm 3, the multiple testing procedure \mathcal{P} is BH and $m = p$. We will be showing shortly that the multiple testing procedures $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m$ connected to the partial conjunction p-values in Algorithm 3 are also BH. For all $j \in [p]$, the partial conjunction p-value $P_j^{r_j/m_j} = P_j^S$ with $r_j = \bar{s}$, $m_j = p - 1$, and $A_j = \{(i, j) : i \in \{j\}^c\}$.

We will begin the proof of validity of Algorithm 3 by proving that BH is self-consistent, non-increasing and has the natural monotonicity property. Specifically, we will consider BH applied on a vector of \tilde{m} p-values $\tilde{P} = (\tilde{P}_1, \tilde{P}_2, \dots, \tilde{P}_{\tilde{m}})$. To show self-consistency, we will show the rejection set

$$\mathcal{R} = \mathcal{R}(\tilde{P}) \subseteq \left\{ j \in [m] : \tilde{P}_j \leq \frac{\alpha|\mathcal{R}|}{m} \right\}.$$

We begin by noting the following identities about the rejection set of the BH procedure.

$$\mathcal{R} = \{j \in [\tilde{m}] : P_j \leq \tilde{P}_{(|\mathcal{R}|)}\}$$

and

$$|\mathcal{R}| = \max \left\{ j \in [\tilde{m}] : \tilde{P}_{(j)} \leq \frac{\alpha j}{\tilde{m}} \right\}.$$

Here for any $j \in [\tilde{m}]$, $\tilde{P}_{(j)}$ denotes the j^{th} ordered p-value among $\{\tilde{P}_1, \tilde{P}_2, \dots, \tilde{P}_{\tilde{m}}\}$. In particular, $\tilde{P}_{(|\mathcal{R}|)} \leq \frac{\alpha|\mathcal{R}|}{\tilde{m}}$. This, combined with the first identity implies that for any $j \in \mathcal{R}$, $\tilde{P}_j \leq \frac{\alpha|\mathcal{R}|}{\tilde{m}}$ which proves self-consistency of BH. This proves that the methods $\mathcal{P}, \mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m$ are all self-consistent.

To show that BH is non-increasing and follows the natural monotonicity property, we consider some $j^* \in [\tilde{m}]$ and define $\tilde{Q} = (\tilde{P}_1, \dots, \tilde{P}_{j^*-1}, \tilde{Q}_{j^*}, \tilde{P}_{j^*+1}, \dots, \tilde{P}_{\tilde{m}})$ where $\tilde{Q}_{j^*} \geq \tilde{P}_{j^*}$. To prove BH satisfies the natural monotonicity property, we will show that $\mathcal{R}(\tilde{P}) \supseteq \mathcal{R}(\tilde{Q})$ and to show that BH is non-increasing, we will show that $|\mathcal{R}(\tilde{P})| \geq |\mathcal{R}(\tilde{Q})|$. For the proof, first observe that the ordered vectors of \tilde{P} and \tilde{Q} satisfy $\tilde{P}_{(j)} \leq \tilde{Q}_{(j)}$ for all $j \in [\tilde{m}]$. Therefore

$$\left\{ j \in [\tilde{m}] : \tilde{P}_{(j)} \leq \frac{\alpha j}{\tilde{m}} \right\} \supseteq \left\{ j \in [\tilde{m}] : \tilde{Q}_{(j)} \leq \frac{\alpha j}{\tilde{m}} \right\}.$$

Since $|\mathcal{R}(\tilde{P})| = \max \left\{ j \in [\tilde{m}] : \tilde{P}_{(j)} \leq \frac{\alpha j}{\tilde{m}} \right\}$ and $|\mathcal{R}(\tilde{Q})| = \max \left\{ j \in [\tilde{m}] : \tilde{Q}_{(j)} \leq \frac{\alpha j}{\tilde{m}} \right\}$, we get that $|\mathcal{R}(\tilde{P})| \geq |\mathcal{R}(\tilde{Q})|$ proving BH is non-increasing. Then, the natural monotonicity property holds because of the following:

$$\begin{aligned} \mathcal{R}(\tilde{P}) &= \left\{ j \in [\tilde{m}] : \tilde{P}_j \leq \frac{\alpha|\mathcal{R}(\tilde{P})|}{\tilde{m}} \right\} \supseteq \left\{ j \in [\tilde{m}] : \tilde{P}_j \leq \frac{\alpha|\mathcal{R}(\tilde{Q})|}{\tilde{m}} \right\} \\ &\supseteq \left\{ j \in [\tilde{m}] : \tilde{Q}_j \leq \frac{\alpha|\mathcal{R}(\tilde{Q})|}{\tilde{m}} \right\} = \mathcal{R}(\tilde{Q}). \end{aligned}$$

The first inclusion (\supseteq) holds because $|\mathcal{R}(\tilde{P})| \geq |\mathcal{R}(\tilde{Q})|$, and the second inclusion follows from the fact that $\tilde{Q}_j \geq \tilde{P}_j$ for all $j \in [m]$. This proves that the BH procedure, and by extension the multiple testing procedures $\mathcal{P}_1, \mathcal{P}_2, \dots, \mathcal{P}_m$, follow the natural monotonicity property.

Next, we will establish the connection between the Simes' partial conjunction p-values and BH. For that, we again consider a vector of p-values $\tilde{P} = (\tilde{P}_1, \tilde{P}_2, \dots, \tilde{P}_{\tilde{m}})$ and some

$\tilde{r} \leq \tilde{m}$. We want to show that the Simes' PCH p-value $P^{\tilde{r}, \tilde{m}}$ (as defined in Equation (3), with $P_{\cdot, j}$ replaced by \tilde{P} , \bar{s} by \tilde{r} , and $p - 1$ by \tilde{m}) satisfies the following

$$P^{\tilde{r}, \tilde{m}} = \max \left\{ P^{1, \tilde{m} - \tilde{r} + 1}[A] : A \subseteq [\tilde{m}], |A| = \tilde{m} - \tilde{r} + 1 \right\},$$

where $P^{1, \tilde{m} - \tilde{r} + 1}[A]$ is the minimum adjusted p-value according to BH applied to the subvector \tilde{P}_A . The set of adjusted p-values for BH applied to \tilde{P}_A is

$$\left\{ \frac{\tilde{m} - \tilde{r} + 1}{\text{rank}(\tilde{P}_j; \tilde{P}_A)} \tilde{P}_j : j \in A \right\},$$

where $\text{rank}(\tilde{P}_j; \tilde{P}_A)$ is the rank of \tilde{P}_j in \tilde{P}_A . Therefore, it is enough to show

$$\min_{\tilde{r} \leq j \leq \tilde{m}} \frac{\tilde{m} - \tilde{r} + 1}{j - \tilde{r} + 1} \tilde{P}_{(j)} = \max_{A \subseteq [\tilde{m}]; |A| = \tilde{m} - \tilde{r} + 1} \min_{j \in A} \frac{\tilde{m} - \tilde{r} + 1}{\text{rank}(\tilde{P}_j; \tilde{P}_A)} \tilde{P}_j. \quad (7)$$

Now,

$$\min_{j \in A} \frac{\tilde{m} - \tilde{r} + 1}{\text{rank}(\tilde{P}_j; \tilde{P}_A)} \tilde{P}_j = \min_{j \in [\tilde{m} - \tilde{r} + 1]} \frac{\tilde{m} - \tilde{r} + 1}{j} (\tilde{P}_A)_{(j)},$$

where $(\tilde{P}_A)_{(j)}$ denotes the j^{th} ordered p-value in \tilde{P}_A . Define A_{\max} as the set of $\tilde{m} - \tilde{r} + 1$ maximum p-values in \tilde{P} . Then $\left((\tilde{P}_{A_{\max}})_{(1)}, (\tilde{P}_{A_{\max}})_{(2)}, \dots, (\tilde{P}_{A_{\max}})_{(\tilde{m} - \tilde{r} + 1)} \right)$ is coordinatewise greater than or equal to $\left((\tilde{P}_A)_{(1)}, (\tilde{P}_A)_{(2)}, \dots, (\tilde{P}_A)_{(\tilde{m} - \tilde{r} + 1)} \right)$ which implies

$$\min_{j \in [\tilde{m} - \tilde{r} + 1]} \frac{\tilde{m} - \tilde{r} + 1}{j} (\tilde{P}_{A_{\max}})_{(j)} \geq \min_{j \in [\tilde{m} - \tilde{r} + 1]} \frac{\tilde{m} - \tilde{r} + 1}{j} (\tilde{P}_A)_{(j)} \quad \text{for any } A \subseteq [\tilde{m}], |A| = \tilde{m} - \tilde{r} + 1.$$

Therefore, the RHS of Equation (7) equals

$$\begin{aligned} \max_{A \subseteq [\tilde{m}]; |A| = \tilde{m} - \tilde{r} + 1} \min_{j \in A} \frac{\tilde{m} - \tilde{r} + 1}{\text{rank}(\tilde{P}_j; \tilde{P}_A)} \tilde{P}_j &= \min_{j \in A_{\max}} \frac{\tilde{m} - \tilde{r} + 1}{\text{rank}(\tilde{P}_j; \tilde{P}_{A_{\max}})} \tilde{P}_j \\ &= \min_{j \in \{\tilde{r}, \dots, \tilde{m}\}} \frac{\tilde{m} - \tilde{r} + 1}{j - \tilde{r} + 1} \tilde{P}_{(j)} \\ &= \text{LHS of Equation (7)}. \end{aligned}$$

Therefore, the Simes' PC p-values and BH are connected as stated in Theorem B.1 which proves the connection of the Simes' PC p-values P_j^S and BH in Algorithm 3.

Finally, we assume that the p-values $\{P_{i, j} : i, j \in [p], i \neq j\}$ are PRDS. We have demonstrated that Algorithm 3 satisfies all the conditions of Theorem B.1, which implies that it controls the FDR at level α .

B.9 Proof of Lemma B.1

Define the transformation $h : \mathbb{R}^p \rightarrow \mathbb{R}^{p+1}$ as follows:

$$g(Z) = (X, T) \quad X_{\{j\}} = \frac{Z_{\{j\}}}{\sum_{k=1}^p Z_{\{k\}}} \forall j \in [p], \quad T = \sum_{k=1}^p Z_{\{k\}}$$

X lies on the $p - 1$ dimensional simplex. For any $d \in \mathbb{N}$ define H^d and λ_d as the d -dimensional Hausdorff measure and Lebesgue measure in \mathbb{R}^d respectively. So f_Z is the Radon–Nikodym derivative or simply the probability density function of Z with respect to λ_p . f_X is the Radon–Nikodym derivative of the push forward measure of X induced by h and the measure of Z with respect to H^{p-1} if and only if $P(X \in B) = \int_B f_X dH^{p-1}$ for all open sets B on the $p - 1$ dimensional simplex. Now

$$P(X \in B) = \int_{h^{-1}(B \times [0, \infty))} f_Z d\lambda_p = \int_{B \times [0, \infty)} f_Z(xt) D(x, t) dH^p$$

where $D(x, t)$ is the appropriate volume element associated with this change of variables. If $J(x, t)$ is the Jacobian of the transformation h at (x, t) , then $D(x, t) = \det(J(x, t)^T J(x, t))^{-\frac{1}{2}}$. Now,

$$\frac{\partial x_i}{\partial z_i} = \frac{\sum_{j \neq i} z_j}{(\sum_{j=1}^p z_j)^2} = \frac{1}{t} - \frac{x_i}{t}, \quad \frac{\partial x_i}{\partial z_j} = \frac{-z_i}{(\sum_{j=1}^p z_j)^2} = -\frac{x_i}{t}, \quad \frac{\partial t}{\partial z_i} = 1$$

implying $J(x, t) = \begin{pmatrix} \frac{1}{t} I_p - \frac{1}{t} x \mathbf{1}_p^T \\ \mathbf{1}_p^T \end{pmatrix}$ and hence $J(x, t)^T J(x, t) = \frac{1}{t^2} I_p + (1 + \frac{1}{t^2} x^T x) \mathbf{1}_p \mathbf{1}_p^T - \frac{1}{t^2} \mathbf{1}_p x^T - \frac{1}{t^2} x \mathbf{1}_p^T$.

$$\begin{aligned} \therefore \det(J(x, t)^T J(x, t)) &= \frac{1}{t^{2p}} \det(I_p + (t^2 + x^T x) \mathbf{1}_p \mathbf{1}_p^T - \mathbf{1}_p x^T - x \mathbf{1}_p^T) \\ &= \frac{1}{t^{2p}} \det \left(I_p + (\sqrt{t^2 + x^T x} \mathbf{1}_p \quad -\mathbf{1}_p \quad -x) \begin{pmatrix} \sqrt{t^2 + x^T x} \mathbf{1}_p^T \\ x^T \\ \mathbf{1}_p^T \end{pmatrix} \right) \\ &= \frac{1}{t^{2p}} \det \left(I_3 + \begin{pmatrix} p(t^2 + x^T x) & -p\sqrt{t^2 + x^T x} & -\sqrt{t^2 + x^T x} \\ \sqrt{t^2 + x^T x} & -1 & -x^T x \\ p\sqrt{t^2 + x^T x} & -p & -1 \end{pmatrix} \right) \\ &= \frac{1}{t^{2p}} \det \left(\begin{pmatrix} 1 + p(t^2 + x^T x) & -p\sqrt{t^2 + x^T x} & -\sqrt{t^2 + x^T x} \\ \sqrt{t^2 + x^T x} & 0 & -x^T x \\ p\sqrt{t^2 + x^T x} & -p & 0 \end{pmatrix} \right) \\ &= \frac{1}{t^{2p}} p t^2 = \frac{p}{t^{2(p-1)}}. \end{aligned}$$

So, $D(x, t) \propto t^{p-1}$ which implies the following

$$\begin{aligned} P(X \in B) &\propto \int_{B \times [0, \infty)} f_Z(xt) t^{p-1} dH^p \\ &\propto \int_B \int_{[0, \infty)} f_Z(xt) t^{p-1} d\lambda_1 dH^{p-1} \end{aligned}$$

Define

$$f_X(x) = c \int_{[0, \infty)} f_Z(xt) t^{p-1} d\lambda_1 = c \int_0^\infty f_Z(xt) t^{p-1} dt$$

where c is a constant $\left(= \int_{\Delta^{p-1}} \int_{[0, \infty)} f_Z(xt) t^{p-1} d\lambda_1 dH^{p-1} \right)$. Then by Radon–Nikodym Theorem, f_X is the density of X with respect to H^{p-1} .

Now, to show the Lemma, we first prove the following claims

Claim 1: $\text{Supp}_{>0}(\mathbf{X}) \supseteq g(\text{Supp}_{>0}(\mathbf{Z}))$ Consider a point $z \in \text{Supp}_{>0}(\mathbf{Z})$ and define $x = g(z)$. We want to prove that $x \in \text{Supp}_{>0}(\mathbf{X})$. Since f_Z is continuous in $\mathbb{R}^p \setminus \mathbf{0}_p$ and $f_Z(z) > 0$, $f_Z(z') > 0$ for all z' in some open neighborhood of z . Formally, there exists $\epsilon > 0$ such that $f_Z(z') > 0 \forall z' \in B_\epsilon(z) \cap [0, \infty)^p \setminus \mathbf{0}_p$. Then, we have the following

$$f_X(x) = c \int_0^\infty f_Z(tx) \cdot t^{p-1} dt \geq c \int_{\min(\{t: xt \in B_\epsilon(z) \cap [0, \infty)^p \setminus \mathbf{0}_p\})}^{\max(\{t: xt \in B_\epsilon(z) \cap [0, \infty)^p \setminus \mathbf{0}_p\})} f_Z(tx) \cdot t^{p-1} dt.$$

Note that for any $z \in [0, \infty)^p \setminus \mathbf{0}_p$ and $x = g(z)$, $\{t : xt \in B_\epsilon(z) \cap [0, \infty)^p \setminus \mathbf{0}_p\}$ is an open interval. This, together with the fact that $f_Z(tx) > 0$ and $t > 0$ for all $t \in \{t : xt \in B_\epsilon(z) \cap [0, \infty)^p \setminus \mathbf{0}_p\}$, imply that, $f_X(x) > 0$ which proves Claim 1.

Claim 2: $\text{Supp}_{>0}(\mathbf{X}) \subseteq g(\text{Supp}_{>0}(\mathbf{Z}))$ It is enough to show that for any point $x \in \Delta^{p-1}$ such that $f_X(x) > 0$, there exists $z \in [0, \infty)^p \setminus \mathbf{0}_p$ such that $f_Z(z) > 0$. Assume the contrapositive, i.e., $f_Z(z) = 0$ for all $z \in \mathbb{R}^n$ such that $x = g(z)$. Then, $\int_0^\infty f_Z(tx) t^{p-1} dt = 0$, which implies $f_X(x) = 0$, contradicting our assumption. Therefore, there must exist z such that $f_Z(z) > 0$ and $x = g(z)$. This proves Claim 2 and, in turn, the Lemma.

C Application of Theorem 2.1 for gene knockout experiments.

In Corollary 2.2, we considered regression problems with factor covariates where each observation is assigned a single level from each factor. We can generalize this setup where each observation receives multiple levels from a single factor. This arises frequently in gene knockout experiments where exactly L out of p genes are knocked out per individual, resulting in the sum of each row of the binary design matrix equal to L . For simplicity we consider only one factor with L levels and define \mathcal{S} and the nontrivial Markov boundary as in Section 2.1.

Corollary C.1 (Uniqueness of nontrivial Markov boundary for gene knockout experiments). *Consider a regression problem where the covariate X denotes levels of a single factor. If $\mathcal{S} = [p]$ then no nontrivial Markov boundary exists. Otherwise if (Y, X) satisfies the following two assumptions:*

(i) *For all $i, j \in \mathcal{S}^c$, there exists a finite t and a sequence l_1, l_2, \dots, l_t such that $\mathcal{H}_{i, l_1}, \mathcal{H}_{l_1, l_2}, \dots, \mathcal{H}_{l_{t-1}, l_t}, \mathcal{H}_{l_t, j}$ are all true,*

(ii) *Positive probability is assigned to all vectors x such that $\sum_{j \in [p]} X_{\{j\}} = L$,*

then \mathcal{S} is the unique nontrivial Markov boundary.

Proof of Corollary C.1. If $\mathcal{S} = [p]$, Theorem 2.1 proves there does not exist any nontrivial Markov boundary. Lets consider the other case, that is $|\mathcal{S}| < p$. By assumption (i), to prove Corollary C.1 using Theorem 2.1, using the same logic as in Appendix B.3, it is enough to show that $\{A, B\} \in \Delta$ for any $A, B \subseteq [p]$ such that $|A \cap B| = 1$ and $|B| = 2$.

So, we choose any such A, B and consider the distribution of $X_{A \cup B} \mid X_{(A \cup B)^c} = c$. Define $L_1 = L - \|c\|_1 = |X_{A \cup B}|$ and $L_2 = L - L_1$. If $L_1 = 0$, then the support of $X_{A \cup B} \mid X_{(A \cup B)^c} = c$ is equal to the set $\{\mathbf{0}_p\}$ which is path connected. Otherwise if $L_1 > 0$, consider any two points x and y in the support of $X_{A \cup B} \mid X_{(A \cup B)^c} = c$. If $x_{B \setminus A} = y_{B \setminus A}$, then they are coordinatewise connected through $B \setminus A$ implying $\{A, B\} \in \Delta$. In the other case, when $x_{B \setminus A} \neq y_{B \setminus A}$, assume WLOG that $x_{B \setminus A} = 1$ and $y_{B \setminus A} = 0$. There are three cases based on the values of $x_{A \cap B}$ and $y_{A \cap B}$:

Case 1: ($x_{A \cap B} = y_{A \cap B} = 0$)

We introduce the following notation: for any vector $w \in \mathbb{R}^p$, we denote $w = (w_{A \setminus B}, w_{A \cap B}, w_{B \setminus A}, w_{(A \cup B)^c})$. In the case $x_{A \cap B} = y_{A \cap B} = 0$, define $z = (x_{A \setminus B}, 1, y_{B \setminus A}, c)$. $\sum_{j \in [p]} z_{\{j\}} = L$ which implies z is in the support of $X_{A \cup B} \mid X_{(A \cup B)^c} = c$ by assumption (ii). x and z are coordinatewise connected through $A \setminus B$ and y and z are coordinatewise connected through $B \setminus A$ proving the equivalence of x and y through $A \setminus B$ and $B \setminus A$.

Case 2: ($y_{A \cap B} = 1$)

Define $z = (y_{A \setminus B}, 0, x_{B \setminus A}, c)$. Similar to case 1, since $\sum_{j \in [p]} z_{\{j\}} = L$, assumption (ii) implies that z is in the support of $X_{A \cup B} \mid X_{(A \cup B)^c} = c$. Also, x and z are coordinatewise connected through $B \setminus A$ and y and z are connected through $A \setminus B$ which shows that x and y are equivalent through $A \setminus B$ and $B \setminus A$.

Case 3: ($x_{A \cap B} = 1$ and $y_{A \cap B} = 0$)

Since $\sum_{j \in A \setminus B} x_{\{j\}} - \sum_{j \in A \setminus B} y_{\{j\}} = 2$, $x_{A \setminus B}$ must have at least two zeros. Define $x'_{A \setminus B}$ as $x_{A \setminus B}$ with one additional 1 in place of a zero. Then if we define $x^{(1)} = (x'_{A \setminus B}, 0, x_{B \setminus A}, c)$, $\sum_{j \in [p]} z_{\{j\}}^{(1)} = L$ which implies $x^{(1)}$ belongs to the support of $X_{A \cup B} \mid X_{(A \cup B)^c} = c$. Define another point $x^{(2)} = (x'_{A \setminus B}, 1, 0, c)$. Again, entries of $x^{(2)}$ also add up to L implying $x^{(2)}$ belongs to the support of $X_{A \cup B} \mid X_{(A \cup B)^c} = c$ by assumption (ii). Finally, we note that the following:

- x and $x^{(1)}$ are coordinatewise connected through $B \setminus A$.
- $x^{(1)}$ and $x^{(2)}$ are coordinatewise connected through $A \setminus B$.
- $x^{(2)}$ and y are coordinatewise connected through $B \setminus A$.

Together, they imply x and y are equivalent through $A \setminus B$ and $B \setminus A$.

Therefore, in all the cases, x and y are equivalent through $A \setminus B$ and $B \setminus A$ which shows that $X_{A \cup B} \mid X_{(A \cup B)^c}$ has only one equivalence class. Hence, $\{A, B\} \in \Delta$. This is true for any A, B such that $|A \cap B| = 1$ and $|B| = 2$. Hence, we can apply Theorem 2.1 which completes the proof of Corollary C.1. \square

While Corollary C.1 is stated for a single factor, one can extend it to the L -factor case similar to Corollary 2.2.

D Pseudocode of Algorithms

D.1 FWER control

Algorithm 1: Method: (for FWER control under arbitrary dependence)

Data: P-value matrix $P \in [0, 1]^{p \times p}$, desired FWER threshold α , and a strict upper bound on the number of non-nulls \bar{s}

Result: Rejected indices $\hat{\mathcal{S}}$

```
1 Initialize:  $\hat{\mathcal{S}} = \emptyset$ 
2 for  $k \in [p]$  :
3   for  $j \in [p]$  :
4     Define Bonferroni's p-value  $P_j^B(\hat{\mathcal{S}})$  as in Equation (5).
5     Define  $j^* = \underset{j \in [p]}{\operatorname{argmin}} P_j^B(\hat{\mathcal{S}})$ .
6     if  $P_{j^*}^B(\hat{\mathcal{S}}) \leq \frac{\alpha}{p-k+1}$  then
7       Add  $j^*$  to  $\hat{\mathcal{S}}$ ;
8     else
9       return  $\hat{\mathcal{S}}$ .
10 return  $\hat{\mathcal{S}}$ .
```

Algorithm 2: Method: (for FWER control under positive dependence)

Data: P-value matrix $P \in [0, 1]^{p \times p}$, desired FWER threshold α , and a strict upper bound on the number of non-nulls \bar{s}

Result: Rejected indices $\hat{\mathcal{S}}$

```
1 Initialize:  $\hat{\mathcal{S}} = \emptyset$ 
2 for  $k \in [p]$  :
3   for  $j \in [p]$  :
4     Define Simes' p-value  $P_j^S(\hat{\mathcal{S}})$  as in Equation (6).
5     Define  $j^* = \underset{j \in [p]}{\operatorname{argmin}} P_j^S(\hat{\mathcal{S}})$ .
6     if  $P_{j^*}^S(\hat{\mathcal{S}}) \leq \frac{\alpha}{p-k+1}$  then
7       Add  $j^*$  to  $\hat{\mathcal{S}}$ ;
8     else
9       return  $\hat{\mathcal{S}}$ .
10 return  $\hat{\mathcal{S}}$ .
```

D.2 FDR control

D.3 FDR control

Algorithm 3: Detecting Markov boundary while Controlling FDR

Data: P-value matrix $P \in [0, 1]^{p \times p}$, desired FDR threshold α , and a strict upper bound on the number of non-nulls \bar{s}

Result: Rejected indices $\hat{\mathcal{S}}$

1 **Initialize:** $\hat{\mathcal{S}} = \emptyset$

2 **for** $j \in [p]$:

3 \lfloor Define Simes' p-value P_j^S as in Equation (3).

4 Sort Simes' p-values in ascending order: $P_{(1)}^S \leq P_{(2)}^S \leq \dots \leq P_{(p)}^S$;

5 Define $k^* = \max \left\{ k \in [p] : P_{(k)}^S \leq \frac{k}{p} \cdot \alpha \right\}$;

6 Define $\hat{\mathcal{S}} = \left\{ j : P_j^S \in \left\{ P_{(1)}^S, \dots, P_{(k^*)}^S \right\} \right\}$;

7 **return** $\hat{\mathcal{S}}$;

E Additional simulation details

E.1 Conditional independence testing for simulations

In Section 3.4, we mentioned that any conditional independence test could be used to obtain the $P_{i,j}$'s. For our simulations, we use the distilled conditional randomization test (dCRT) (Liu et al., 2022), which is a computationally efficient special case of the conditional randomization test (Candès et al., 2018), to test the hypotheses $\mathcal{H}_{i,j} : Y \perp \perp X_{\{i,j\}} \mid X_{\{i,j\}^c}$ for $i, j \in [p]$. To describe how we implement the dCRT, we need to introduce some notation first. Let \mathbb{Y} denote the vector of observed outcomes and let \mathbb{X} denote the design matrix. $\mathbb{X}_{\{i,j\}}$ denotes the matrix of the i^{th} and j^{th} columns of \mathbb{X} and $\mathbb{X}_{\{i,j\}^c}$ denotes the matrix with all the columns of \mathbb{X} except i and j .

The first step of the dCRT is to draw samples $\tilde{\mathbb{X}}_{\{i,j\}}^{(1)}, \tilde{\mathbb{X}}_{\{i,j\}}^{(2)}, \dots, \tilde{\mathbb{X}}_{\{i,j\}}^{(K)}$ independently from the distribution of $\mathbb{X}_{\{i,j\}} \mid \mathbb{X}_{\{i,j\}^c}$. Since, in our simulations, rows of \mathbb{X} consist of independent observations of the random variable X , each row of $\tilde{\mathbb{X}}_{\{i,j\}}^{(k)}$ can be sampled independently from the conditional distribution $X_{\{i,j\}} \mid X_{\{i,j\}^c}$.

Now we define the test statistics for dCRT. Let $\hat{\mathbb{Y}} \in \mathbb{R}^n$ denote the fitted values obtained from a 5-fold cross-validated LASSO fitting \mathbb{Y} to $\mathbb{X}_{\{i,j\}^c}$. Also, for any vector $y \in \mathbb{R}^n$ and matrix $z \in \mathbb{R}^{n \times 2}$, define $T(y, z)$ as the R^2 from the regression of y on $\log(z)$. Then the dCRT p-value is defined as follows:

$$P_{i,j} = \frac{1}{K+1} \left(1 + \sum_{k=1}^K \mathbb{1}_{\{T(\mathbb{Y} - \hat{\mathbb{Y}}, \tilde{\mathbb{X}}_{\{i,j\}}^{(k)}) \geq T(\mathbb{Y} - \hat{\mathbb{Y}}, \mathbb{X}_{\{i,j\}})\} \right). \quad (8)$$

Under $\mathcal{H}_{i,j}$, the dCRT p-value $P_{i,j}$ has a super-uniform distribution because it is a special case of the conditional randomization test, which is proved to be valid in Candès et al. (2018).

E.2 Dirichlet

As mentioned in the previous subsection, we use thedCRT (Liu et al., 2022) to test the base hypotheses $\mathcal{H}_{i,j}$. For each $i, j \subset [p]$, we calculate the p-value $P_{i,j}$ as described in Appendix E.1. For the benchmark LOO methods, we select some $j \in [p]$ at random and for each $i \in \{j\}^c$, we define $d_Y^{\{i\}}$ as the prediction from a 5-fold cross-validated LASSO regression of Y on $X_{\{i,j\}^c}$. Since $X_{\{j\}^c}$ is non-compositional, we draw resamples from $X_{\{i\}} \mid X_{\{i,j\}^c}$ (which follows a Beta distribution in this case). Then, we define the r-squared of the regression of $Y - d_Y^{\{i\}}$ on $\log(X_i)$ as the test statistic T and calculate p-values as in Liu et al. (2022, Algorithm 1). In this way, we only get p-values for all $i \in \{j\}^c$.





Figure 4: Comparison of FWER and average power for variable selection with Dirichlet covariates. The target FWER level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.


Figure 4 shows that like in the case of FDR control (Figure 1b), the proposed methods— $\text{BCP}(p-1)$ -Holm, $\text{BCP}(p/2)$ -Holm and $\text{BCP}(s+1)$ -Holm—control the FWER at a level much below the nominal level of 10% while the benchmark method LOO-Holm violates the FWER level. In terms of average power, $\text{BCP}(p/2)$ -Holm and $\text{BCP}(s+1)$ -Holm achieve comparable average power to LOO-Holm.

E.3 Logistic-normal


For this distribution, the conditional distribution of $X_{\{i,j\}}$ given $X_{\{i,j\}^c}$ lacks a closed form, which is why we use MCMC to generate dCRT resamples. These resamples although not independent, are exchangeable, which is sufficient for validity. Figure 5 shows that LOO methods fail to control respective error rates while the proposed methods control the same (Figure 5). $\text{BCP}(p/2)$ and $\text{BCP}(s+1)$ methods achieve power and average power similar to the benchmarks. In contrast to the Dirichlet case, $\text{BCP}(p-1)$ methods have considerably lower average power than the rest of the methods in multiple testing scenarios.



(a) Comparison of type I error and power for single testing with Logistic-normal covariates. The target type I error is 5% and error bars correspond to ± 2 Monte Carlo standard errors.



(b) Comparison of FWER and average power for variable selection with Logistic-normal covariates. The target FWER level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.



(c) Comparison of FDR and average power for variable selection with Logistic-normal covariates. The target FDR level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.

Figure 5: Comparison of methods for (a) single testing, (b) FWER control and (c) FDR control with Logistic-normal covariates.

E.4 Multivariate normal




Figure 6: Comparison of FWER and average power for variable selection with multivariate normal covariates. The target FWER level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.

Figure 6 shows that the three proposed methods achieve average power very close to the benchmark Univariate-Holm while controlling FWER at a level much lower than the nominal level 10%.

E.5 Studying the effect of conditioning on sparse covariates

Figure 3 illustrated that we can gain substantial power by conditioning out sparse covariates. As Figure 7 shows, the error rates for the proposed methods remain well below the nominal levels, 5% for single testing and 10% for multiple testing, irrespective of the number of covariates we condition out.

E.6 Robustness study

In this simulation, we study the effect of conditioning out the sparse covariates. To do this, we generate the data from a Dirichlet distribution where the parameter vector is a vector of 2's of length 100. So, the distribution of $X_{\{i,j\}}$ given $X_{\{i,j\}^c}$ is a scaled beta distribution which can be used in the dCRT step. However, here we consider the scenario where the parameters for the Dirichlet distribution are not known. Instead, we estimate the parameters of the Dirichlet distribution using the `fit_dirichlet` function (Minka, 2000) in R. Figures 8b and 8c show that estimating parameters has little impact on the average power and the error rates of the proposed methods $\text{BCP}(p-1)$, $\text{BCP}(p/2)$, and $\text{BCP}(s+1)$.

E.7 Bonferroni p-values

All of the plots in Section 4 used PCH p-values combined via Simes' approach. In this section, we present the same results for Bonferroni's approach.





Figure 7: Effect of conditioning out the sparse covariates on type I error (left), FWER (middle), and FDR (right), respectively. The target type I error, FWER, and FDR are 5%, 10%, and 10%, respectively, and error bars correspond to ± 2 Monte Carlo standard errors.


Figures 9, 10, and 11 show a considerable decrease in power and average power of $\text{BCP}(\bar{s})$, $\text{BCP}(\bar{s})$ -Holm, and $\text{BCP}(\bar{s})$ -BH methods when we use Bonferroni p-values instead of Simes in the cases of Dirichlet, Logistic-normal, and multivariate normal covariates respectively. In figures 10 and 11 we see that BCP methods with Bonferroni p-values almost always have zero power or average power. One reason for this is that 1,500 dCRT resamples are insufficient to achieve non-zero power at these SNR values when using Bonferroni p-values, unlike when Simes p-values are used. In order to confirm that, we have increase the number of dCRT resamples to 25,000 in Figure 9 and we observe that more BCP methods achieve non-zero (average) power using Bonferroni p-values. We also observe from Figure 9b that unlike the Simes p-values, the Bonferroni p-values are not monotone with respect to \bar{s} . $\text{BCP}(p-1)$ -Bonferroni-Holm has non zero average power while $\text{BCP}(p/2)$ -Bonferroni-Holm and $\text{BCP}(s+1)$ -Holm both have zero average power in the SNR range used.

While a smaller \bar{s} should provide more information and yield smaller p-values, Bonferroni's p-values lack monotonicity with respect to \bar{s} . To illustrate, consider a simple example: for a particular $\mathcal{H}_{0j}^{\bar{s}}$ with $p = 4$ and $s = 1$, suppose the base p-values are $P_{\{j\}^c, j} = (0.2, 0.25, 0.3)$. Bonferroni's PCH p-value for $\bar{s} = 1$ is $P_{1/3}^B = 3 \times 0.2 = 0.6$, while for $\bar{s} = 2$, it is $P_{2/3}^B = 2 \times 0.25 = 0.5$. Thus, the p-value for $\bar{s} = 1$ is higher than for $\bar{s} = 2$, contrary to the expectation that a smaller \bar{s} would result in a lower p-value. Unlike Bonferroni's p-values $P_{r/n}^B$, Simes' p-values satisfy monotonicity with respect to \bar{s} .


Figure 12 shows that the power and average power of BCP methods using Bonferroni



(a) Effect of sampling from an estimated distribution on single testing methods. The target type I error is 5% and error bars correspond to ± 2 Monte Carlo standard errors.



(b) Effect of sampling from an estimated distribution on FWER control methods. The target FWER level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.



(c) Effect of sampling from an estimated distribution on FDR control methods. The target FDR level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.

Figure 8: Simulation studies showing the effects of sampling from an estimated Dirichlet distribution on (a) FWER and (b) FDR control methods.

p-values increases when we condition out more and more sparse covariates. In fact, when we do not condition out any covariate, all BCP methods have zero power in both single and multiple testing problems.

E.8 Performance of AdaFilter

AdaFilter (Wang et al., 2016) is a method designed for testing multiple PCHs. Standard PCH p-values can be quite conservative which reduces power in detecting weak signals. AdaFilter involves a filtering step which reduces the potential detected signals to a smaller set of more promising hypotheses, giving it an edge over traditional methods in terms of average power. Since we also combine multiple PCH tests, AdaFilter may seem like a natural tool to use in our framework. However, AdaFilter requires independence of the base p-values for validity and as discussed in the paragraph after Theorem 3.2, the base p-values will typically be far from independent in our framework.

In Figure 13 we denote AdaFilter–BH and AdaFilter–Bonferroni methods with $r = \bar{s}$ (Wang et al., 2016, Section 3.1) as `adafilter(\bar{s})-BH` and `adafilter(\bar{s})-Bonferroni` respectively. Figure 13 shows that, in the Dirichlet data setting (see Section 4.1), AdaFilter fails to control the FDR and the FWER at the nominal level (10%) particularly in the low SNR regime, while the proposed methods, `BCP(\bar{s})-BH` and `BCP(\bar{s})-Holm`, effectively control both the FDR and Familywise Error Rate (FWER) well below the nominal level of 10% while having comparable average power. In fact, for FWER control, `adafilter($s+1$)-Bonferroni` fails to control the FWER anywhere in our SNR range. So, AdaFilter is not a suitable method for our problem. The reason, as we stated in the previous paragraph, is the strong dependence among the p-values.

E.9 Computational speedups

In Section 3.5, we discussed two ways of significantly speeding up the proposed method. In this section, we look at simulation studies quantifying said speedups.


But first, we mention one further computational improvement we considered that is specific to the dCRT (in fact it would apply to any conditional randomization test), which is the conditional independence testing approach we used in our simulations to calculate the $P_{i,j}$. In the dCRT, resamples $\tilde{X}_{\{i,j\}}$ are drawn from the conditional distribution of $X_{\{i,j\}} \mid X_{\{i,j\}^c}$, and a test statistic $T(y, \tilde{X}_{\{i,j\}}, X_{\{i,j\}^c})$ is computed for each CRT resample to be compared with $T(y, X_{\{i,j\}}, X_{\{i,j\}^c})$. However, many resamples (we denote the number of resamples by K) are needed for a powerful test, especially for variable selection where multiple testing corrections are applied. For faster computation, one can initially draw a smaller number (e.g., $K/10$) of resamples and compute an initial p-value based on just these resamples. If that initial p-value is above, say, 0.1, then it is very unlikely that more resamples will lead to a very small p-value that would provide strong evidence against any null hypothesis, and hence we could just set the p-value to 1 without further resampling in this case. If the initial p-value is small, then we would continue resampling to the full number K and compute the p-value as normal.

As mentioned in the main text, all three proposed computational speedups take the form of data-dependent screening: they replace individual p-values by 1 in the absence of


evidence of them being non-null. This makes the resulting p-values strictly superuniform when both i and j are null variables. While this proves validity of FWER control method as in Algorithm 2 defined using Bonferroni’s combining function, proving FDR control is difficult. This is because the screening procedures can introduce complex dependence among the p-values, potentially breaking the PRDS structure.

To assess the statistical impact of these computational speedups, we perform a simulation study where we consider the same setup as in model 1 of Section 4.1. The setup consists of 100 observations of the form (Y, X) where X has 100 columns and each row of X is sampled independently from a Dirichlet distribution, making X compositional. As Figure 14 shows, the proposed methods with the speedups effectively maintain the False Discovery Rate (FDR) and Family-Wise Error Rate (FWER) well below the designated threshold of 10%. Notably, the $\text{BCP}(\bar{s})$ procedures, coupled with computational improvements, demonstrate average power levels almost identical to their counterparts without these speedups. Moreover, these optimizations reduce the average computation time for the $\text{BCP}(\bar{s})$ methods with BH to 3.22 minutes, compared to 15.55 minutes without the speedups, yielding an approximate 4.8-fold decrease in computation time. The runtimes were calculated using a single core of a second-generation Intel Xeon processor. It is also observed that the average computation time increases with the SNR level. Specifically, the average runtimes for SNR values of 0.25, 0.5, 1, 2, and 4 are 1.4, 2.3, 3.3, 3.5, and 3.8 minutes, respectively for $\text{BCP}(\bar{s})$ -BH with the computational speedups. In contrast, the run times for methods without the speedups decrease with SNR, with corresponding times of 19.1, 17.1, 14.9, 13.5, and 12.9 minutes.


While we get approximately 5 times speedup for a 100×100 covariate matrix, the computational gains are much more pronounced for higher dimensional datasets. To illustrate this we ran a similar experiment with 1000×1000 dimensional covariate matrix. In this case, we ran the proposed methods, both with and without the speedups, on a 24-core Intel Xeon CPU parallelly. The average time taken for the proposed method without any speedups was 23.99 hours while that of the method with them was 42.22 minutes. So, we achieve around 34 fold increase in speed. This shows the effectiveness of these computational improvements in real world datasets where covariates are often high dimensional. Another important point to note here is that, in this simulation, X contains no sparse columns. When X contains sparse columns, one can substantially speed up our method using the technique mentioned in Section 3.3.



(a) Comparison of type I error and power for single testing with Dirichlet covariates. The target type I error is 5% and error bars correspond to ± 2 Monte Carlo standard errors.




(b) Comparison of FWER and average power for variable selection with Dirichlet covariates. The target FWER level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.




(c) Comparison of FDR and average power for variable selection with Dirichlet covariates. The target FDR level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.


Figure 9: Comparison of BCP(\bar{s}) with Simes and Bonferroni PC p-values.



(a) Comparison of type I error and power for single testing with Logistic-normal covariates. The target type I error is 5% and error bars correspond to ± 2 Monte Carlo standard errors.




(b) Comparison of FWER and average power for variable selection with Logistic-normal covariates. The target FWER level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.




(c) Comparison of FDR and average power for variable selection with Logistic-normal covariates. The target FDR level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.


Figure 10: Comparison of $\text{BCP}(\bar{s})$ with Simes and Bonferroni PC p-values for Logistic-normal covariates.



(a) Comparison of type I error and power for single testing with multivariate normal covariates. The target type I error is 5% and error bars correspond to ± 2 Monte Carlo standard errors.




(b) Comparison of FWER and average power for variable selection with multivariate normal covariates. The target FWER level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.

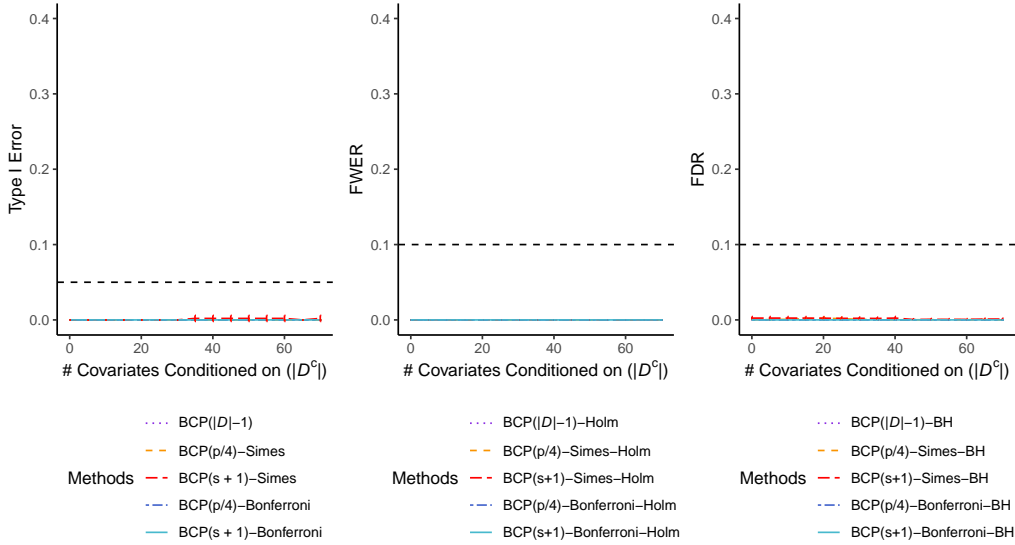


(c) Comparison of FDR and average power for variable selection with multivariate normal covariates. The target FDR level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.

Figure 11: Comparison of BCP(\bar{s}) with Simes and Bonferroni PC p-values for multivariate normal covariates.




(a) Effect of conditioning out the sparse covariates on power for single testing (left) and average power for FWER control (middle) and FDR control (right), respectively. Error bars correspond to ± 2 Monte Carlo standard errors.




(b) Effect of conditioning out the sparse covariates on type I error (left), FWER (middle), and FDR (right), respectively. The target type I error, FWER, and FDR are 5%, 10%, and 10%, respectively, and error bars correspond to ± 2 Monte Carlo standard errors.

Figure 12: Comparison of effect of conditioning out sparse covariates on BCP(\bar{s}) with Simes and Bonferroni PC p-values.




(a) FWER and average power of AdaFilter for variable selection with Dirichlet covariates. The target FWER level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.




(b) FDR and average power of AdaFilter for variable selection with Dirichlet covariates. The target FDR level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.

Figure 13: Performance of AdaFilter on Dirichlet data in controlling (a) FWER and (b) FDR.



(a) Comparison of FWER and average power for variable selection with computational speedups. The target FWER level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.



(b) Comparison of FDR and average power for variable selection with computational speedups. The target FDR level is 10% and error bars correspond to ± 2 Monte Carlo standard errors.

Figure 14: Comparison of proposed method with computational speedups discussed in Section 3.5 for variable selection while controlling (a) FWER and (b) FDR control.