

Searching for subgroup-specific associations while controlling the false discovery rate

Matteo Sesia* and Tianshu Sun*

September 15, 2023

Abstract

This paper introduces an innovative method for conducting conditional independence testing in high-dimensional data, facilitating the automated discovery of significant associations within distinct subgroups of a population, all while controlling the false discovery rate. This is achieved by expanding upon the model-X knockoff filter to provide more informative inferences. Our enhanced inferences can help explain sample heterogeneity and uncover interactions, making better use of the capabilities offered by modern machine learning models. Specifically, our method is able to leverage any model for the identification of data-driven hypotheses pertaining to interesting population subgroups. Then, it rigorously tests these hypotheses without succumbing to selection bias. Importantly, our approach is efficient and does not require sample splitting. We demonstrate the effectiveness of our method through simulations and numerical experiments, using data derived from a randomized experiment featuring multiple treatment variables.

1 Introduction

1.1 Background and motivation

Contemporary data sets are growing in both size and diversity, often incorporating observations from various environments or diverse populations characterized by a multitude of distinct patterns and behaviours. This complexity offers many opportunities to statisticians. For instance, it can help alleviate issues of collinearity [1], enhance the robustness of predictive models to distribution shifts [2], and aid in the identification of causal associations by filtering out spurious correlations [3, 4]. However, such heterogeneous data also pose new challenges, highlighting some limitations of existing analysis methods that rely on rigid population-wide hypothesis testing.

This paper focuses on the problem of identifying significant conditional associations between an outcome of interest and many possible explanatory variables. For this purpose, we adopt and expand the *model-X* framework proposed by Candès et al. [5], which does not require assuming a parametric regression model for the outcome. Within this context, we contend that population-wide hypothesis testing is not fully satisfactory when dealing with heterogeneous samples. Such testing may yield technically valid discoveries, but it leads to findings that are not necessarily the most practically valuable, particularly when the objective is to guide data-driven decisions involving diverse individuals who exhibit varying behaviours. This concern serves as the driving force behind our development of a novel methodology, one that can provide more informative model-X inferences tailored to interesting subgroups of individuals.

Our work is partly motivated by the analytical challenges faced by e-commerce and social media enterprises, which have access to extensive data [6] and are frequently able to monitor customer responses to randomized interventions [7]. These experiments often entail multiple treatment variables, and their joint

*Department of Data Sciences and Operations, University of Southern California, Los Angeles, CA 90089, USA.

distribution can be intricate, and possibly contingent on many covariates [7]. For instance, consider a scenario within social networks where the “things in common” are displayed to encourage friendship formation [8]. In this case, users who hover on another user’s profile are randomly presented with some shared features (e.g., having attended the same school, expressing similar interests, etc.), with the objective of discerning which interventions facilitate the creation of new connections. Alternatively, researchers in an e-commerce company may wish to explore the efficacy of different targeted marketing strategies, leveraging insights from the previous purchases and browsing histories of current customers to make personalized product recommendations.

In both of the aforementioned scenarios, there may exist numerous significant associations between the outcome of interest and various explanatory variables. Given a sufficiently large sample size, any sensible test should be able to identify these associations. Nevertheless, converting these findings into actionable business insights can prove to be a difficult task since these associations may hold little relevance for the majority of individuals in the population.

Consider, for instance, a subscription box service offering an array of products to its subscribers on a monthly basis. The company’s objective is to understand how various promotional strategies impact customer retention and product preferences. Their data collection encompasses a range of promotional tactics, such as discounts, exclusive items, and early access. Upon analysis, an association may emerge between providing exclusive outdoor adventure gear (e.g., backpacking food) as part of the subscription box and an increase in customer retention. Nonetheless, this association probably holds relevance only for a *minority* of users who belong to a subgroup characterized by specific demographic attributes (e.g., younger adults), geographical features (e.g., proximity to mountains) or who have purchase histories of related equipment (e.g., hiking gear). Therefore, the analysts would likely find it helpful to have at their disposal methods that can not only detect significant associations, but also automatically pinpoint the population subgroups for which the discoveries are *relevant*.

Although the potential impacts of this statistical problem are far-reaching, and wealth of related literature has already been written, we contend that there remains an unaddressed void in the current methodology that necessitates the development of a novel approach. On one hand, there are many useful data mining algorithms and machine learning models that can discover subgroups of individuals exhibiting coherent behaviours [9, 10], detect interactions [11, 12], or estimate heterogeneous effects in regression and causal inference [13, 14]. However, such methods are usually not designed to offer precise statistical guarantees in finite samples. On the other hand, there is the *model-X* framework of Candès et al. [5], which makes it possible to rigorously test conditional independence hypotheses under finite-sample type-I error control, while taking advantage of any model to powerfully discover significant associations. However, existing model-X methods are currently limited in the analysis of heterogeneous data because they can only test rigid hypotheses defined over the entire population. Such discoveries may be useful to screen out completely irrelevant predictors but are not directly informative about the particular effects of any variables across diverse subgroups of individuals, and they cannot be used to test for interactions.

1.2 Preview of our contribution

This paper presents an extension of the model-X knockoff filter [5] that enhances the informativeness of inferences derived from heterogeneous data and in the presence of interactions. Our novel method accomplishes two key objectives: (1) it harnesses the power of any machine learning model to automatically identify interesting conditional independence hypotheses specific to well-defined population subgroups, and (2) it rigorously tests these data-driven hypotheses while controlling the false discovery rate [15] in finite samples. Notably, our method efficiently utilizes the same data for both tasks, avoiding selection bias [16–18] without wasteful sample splitting. We call this method the *subgroup-selective knockoff filter*.

Figure 1 gives a preview of our method’s performance. These numerical experiments are conducted using simulated data that consist of multiple explanatory (or “treatment”) variables and an outcome generated from a parametric model, which we refer to as the true “causal” model. This underlying model is intentionally structured so that the outcome is influenced by a distinct subset of variables for different subgroups of individuals, as defined by specific observable covariates. To be more precise, the data generation model is

a linear model that incorporates interaction terms between the explanatory variables and the covariates. Consequently, each variable has a unique effect on each individual, with this effect often equating to zero and generally contingent upon certain covariates. The analyst’s objective is to discover, with the highest possible precision, which variables exert a non-zero effect within particular subgroups.

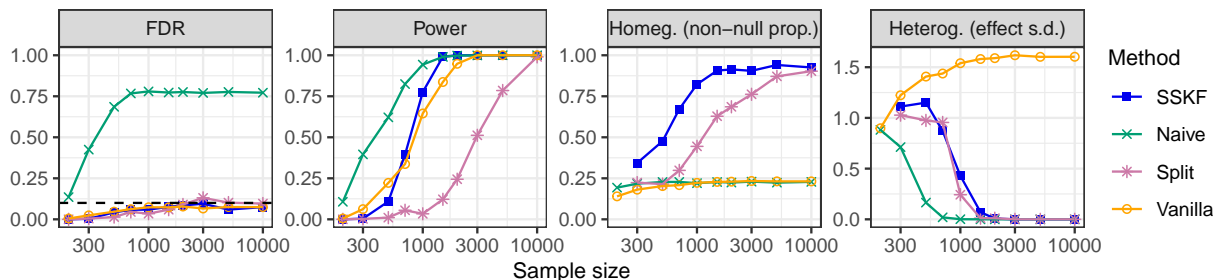


Figure 1: Performance of the subgroup-selective knockoff filter (SSKF) and benchmark methods on synthetic data. The informativeness of the discoveries is quantified by the homogeneity of the corresponding subgroups (higher is better) and the heterogeneity of the underlying true effects in the data-generating model (lower is better). The nominal false discovery rate (FDR) level is 0.1.

As the ground-truth model is known, it is possible to verify which discovered associations are not spurious. Further, it is also feasible to gauge the informativeness of any findings. This can be achieved by comparing the sub-populations identified by the rejected hypotheses and the subgroups of individuals characterized by non-zero treatment effects within the data-generating model. To be more specific, we introduce two distinct metrics for assessing informativeness: one metric quantifies the homogeneity of the true effects within the reported subgroups, while the other metric quantifies their heterogeneity in a complementary manner. We refer to Section 4.1 for further information about the setup of these numerical experiments.

The first alternative (benchmark) approach considered in Figure 1 is the vanilla model-X knockoff filter [5]. This method can be powerful in discovering population-wide conditional associations and it reliably controls the false discovery rate. However, it is unable to pinpoint the subgroups in which the significant variables indeed possess non-zero effects. The limitations of this approach are evident in Figure 1, where the informativeness metrics remain relatively unchanged as the sample size increases.

In contrast, the subgroup-selective knockoff filter progressively uncovers more informative discoveries as the data set grows in size, ultimately enabling precise identification of subgroups with non-zero effects. Additionally, Figure 1 illustrates the strength of our method by comparing it to an intuitive “data splitting” benchmark. This benchmark involves applying the two distinct modules of the subgroup-selective knockoff filter (i.e., learning and testing) using separate subsets of the data obtained through random sample splitting. While this simpler approach also provides valid and informative inferences, it is notably less powerful than our method.

Finally, Figure 1 evaluates the performance of the subgroup-selective knockoff filter against a naive benchmark that differs from our method in that it does not carefully protect against possible selection bias. Unsurprisingly, the naive benchmark results in an excess of spurious discoveries.

1.3 Related work

Our work builds upon the model-X version [5] of the knockoff filter [19]. Prior research studied the problem of constructing (approximately) valid *knockoff* variables for different types of data [20–24]; therefore, in this paper we will assume that suitable knockoffs are available for the application of interest. Others have studied the robustness of the knockoff filter to approximations in construction of the knockoffs [25, 26] and its power [27–30], providing additional support for this framework.

This is the first paper in the model-X literature to seek subgroup-specific inferences for heterogeneous data. Our work is most closely related to Li et al. [31], which proposed an extension of the knockoff filter to analyze data collected from different environments, and argued that robust conditional associations in that

context can be seen as relatively good proxies for causal inferences. While we take some inspiration from Li et al. [31], we study a different problem. In fact, [31] considered each individual as belonging to a distinct and *known* population, and their methods cannot test hypotheses corresponding to adaptively discovered subgroups.

The subgroup-selective knockoff filter is flexible and can learn which subgroups are likely to be affected by each variable using any model. This connects our work to the rich literature on subgroup analysis or the estimation of heterogeneous effects in regression and causal inference [13, 32–39]; see also Caron, Baio, and Manolopoulou [14] for a recent review. The learning module of our method can take advantage of any such models, which may be either parametric or non-parametric, before the subsequent testing module translates their output into precise statistical inferences.

For simplicity, this paper focuses on learning using generalized linear models with interaction terms, as opposed to more complex approaches which may for example involve random forests [18], Bayesian additive regression trees [40], or neural networks [41]. Interaction terms are relatively simple to explain in the context of (generalized) linear models, and they are an important source of sample heterogeneity because they can result in some variables having much larger effects in subgroups with certain covariates. There are numerous works relating interactions to sample heterogeneity in regression or causal inference [13, 35, 37, 38], and several techniques have been developed to discover interactions using high-dimensional data [11, 12, 42–45]. We will demonstrate the use of the subgroup-selective knockoff filter in combination with some of these techniques. However, on their own, the latter would rely on different assumptions and provide different guarantees compared to our method.

1.4 Outline of this paper

Section 2 introduces useful notation and recalls the relevant technical details of the model-X knockoff filter from Candès et al. [5]. Section 3 states our problem and presents our method. Section 4 demonstrates the use of our method empirically. Section 5 concludes with some ideas for future work. The Appendices contains additional methodological details, numerical results, and mathematical proofs. The Appendices also extends our novel method to test *partial conjunctions* [46, 47] of our subgroup-specific hypotheses, complementing the types of inferences considered in this paper.

2 Technical preliminaries

2.1 Basic notation and setup

We consider a data set comprising n individual observations of triplets (X^i, Y^i, Z^i) . For each individual $i \in [n] = \{1, \dots, n\}$, $Y^i \in \mathbb{R}$ is the outcome of interest (which may be either numerical or categorical), $X^i = (X_1^i, \dots, X_p^i) \in \mathbb{R}^p$ describes p explanatory variables, and $Z^i = (Z_1^i, \dots, Z_m^i) \in \mathbb{R}^m$ represents m additional covariates. Let us denote as $\mathbf{X} \in \mathbb{R}^{n \times p}$, $\mathbf{Y} \in \mathbb{R}^{n \times 1}$, and $\mathbf{Z} \in \mathbb{R}^{n \times m}$ the data matrices collecting the observations for all individuals. Without loss of generality, the data distribution can be factored as $P_{X,Y,Z}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}) = P_Z(\mathbf{Z}) \cdot P_{X|Z}(\mathbf{X} | \mathbf{Z}) \cdot P_{Y|X,Z}(\mathbf{Y} | \mathbf{X}, \mathbf{Z})$. In the following, $P_{X|Z}$ will be assumed to be known in order to generate knockoffs, while P_Z and $P_{Y|X,Z}$ may remain completely arbitrary and unknown.

The assumption that $P_{X|Z}$ is known may not be justified in every application, but it is often a reasonable approximation and is well-suited for the analysis of data from randomized experiments [48] or genome-wide association studies [20]. Further, it is feasible to verify the appropriateness of this assumption in practice [23].

Intuitively, the goal is to shed some light onto how Y depends on X given Z ; this aim will be stated more precisely later in terms of testing suitable conditional independence hypotheses. For simplicity, we will assume that different individuals are independent of one another.

2.2 Relevant background on conditional testing with knockoffs

The model-X problem studied by Candès et al. [5] can be thought of as testing, for all $j \in [p]$, whether X_j is associated with Y given Z and all variables excluding X_j (i.e., X_{-j}); that is, whether the following null hypothesis holds true:

$$\mathcal{H}_{0,j} : Y \perp\!\!\!\perp X_j \mid X_{-j}, Z. \quad (1)$$

We refer to Appendix A1 for further details regarding the interpretation of (1) and its connections to parametric inference for generalized linear models and causal inference.

The model-X knockoff filter [5] is designed to test $\mathcal{H}_{0,j}$ in (1) for all $j \in [p]$ while controlling the false discovery rate—the expected proportion of falsely rejected hypotheses. This is an especially useful error rate for analyses in which many discoveries are expected.

The key ingredient for the knockoff filter are the knockoffs, which we denote as \tilde{X} . These are synthetic variables constructed by the statistician as a function of X and Z without looking at Y , so that $\tilde{X} \perp\!\!\!\perp Y \mid X, Z$. Further, the knockoffs are designed to be exchangeable with X in the joint distribution of (X, \tilde{X}) conditional on Z . That is, if $[\mathbf{X}, \tilde{\mathbf{X}}] \in \mathbb{R}^{n \times 2p}$ is the matrix obtained by concatenating \mathbf{X} with the corresponding $\tilde{\mathbf{X}} \in \mathbb{R}^{n \times p}$ and, for any $j \in [p]$, the matrix $[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(j)}$ is obtained by swapping the j -th columns of \mathbf{X} and $\tilde{\mathbf{X}}$, then

$$[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(j)} \mid \mathbf{Z} \stackrel{d}{=} [\mathbf{X}, \tilde{\mathbf{X}}] \mid \mathbf{Z}, \quad \forall j \in [p], \quad (2)$$

where the symbol $\stackrel{d}{=}$ denotes equality in distribution. Thus, the only meaningful difference between X_j and \tilde{X}_j may be the lack of conditional association of the latter with Y .

Constructing knockoffs that satisfy (2) requires knowledge of $P_{X|Z}$, as anticipated in Section 2.1, and it can be computationally involved. However, the problem is well-studied; see Section 1.3. Therefore, this paper assumes that suitable knockoffs are available for the data at hand.

The purpose of the knockoffs is to serve as negative control variables within a predictive model for Y given X, \tilde{X} and Z . This model is utilized to compute importance measures T_j and \tilde{T}_j for each X_j and \tilde{X}_j . Any model can be employed for this purpose, as long as swapping X_j with \tilde{X}_j only results in T_j being swapped with \tilde{T}_j . A typical choice is to fit a (generalized) lasso model [49] and define the importance measures as the absolute values of the scaled regression coefficients for X and \tilde{X} , after tuning the regularization via cross-validation. Then, T_j and \tilde{T}_j are combined pairwise for each j into anti-symmetric statistics W_j ; i.e., $W_j = T_j - \tilde{T}_j$. This ensures that the signs of the W_j are mutually independent coin flips for all j corresponding to a true $\mathcal{H}_{0,j}$, while larger and positive values provide evidence against the null. Letting $\epsilon \in \{-1, +1\}^p$ be an independent random vector such that $\epsilon_j = +1$ if $\mathcal{H}_{0,j}$ is false and $\mathbb{P}[\epsilon_j = +1] = 1/2$ otherwise, then it can be proved that W satisfies the *flip-sign* property:

$$W \mid \mathbf{Z} \stackrel{d}{=} W \odot \epsilon \mid \mathbf{Z}, \quad (3)$$

where \odot indicates element-wise multiplication. Intuitively, the sign of W_j gives rise to a conservative binary p-value p_j for $\mathcal{H}_{0,j}$. That is, $p_j = 1/2$ if $W_j > 0$ and $p_j = 1$ otherwise, as long as $\mathcal{H}_{0,j}$ is true. Finally, a rejection rule for (1) can be obtained by applying a sequential testing procedure to these p-values, in the order defined by the absolute values of W . As each of these p-values contains a single bit of information, an appropriate sequential testing procedure is the knockoff filter (or selective SeqStep+ test) of [19], which can compute an adaptive threshold for W_j controlling the false discovery rate below any desired threshold.

3 Methodology

This section is organized as follows. Section 3.1 defines our subgroup-specific conditional hypotheses, extending the population-wide hypotheses defined in (1). Section 3.2 gives general guidelines for learning adaptive subgroups that can lead to informative hypotheses without introducing selection bias in our subsequent

tests. Section 3.3 outlines a concrete approach for learning data-driven subgroups using parametric regression models with interaction terms. Section 3.4 explains the core of our method, which can test adaptively discovered subgroup-specific conditional hypotheses while controlling the false discovery rate. Section 3.5 provides a detailed example of how to implement our method using powerful test statistics.

3.1 Subgroup-specific conditional hypotheses

Standard conditional testing is not fully satisfactory when analyzing heterogeneous data. In fact, rejecting $\mathcal{H}_{0,j}$ in (1) informs us that X_j may have *some* conditional association with Y in the population, but it sheds no light on the possible heterogeneity of this relation.

Consider for example a simple mixture population with $p = 2$ and $m = 1$, in which $Z_1 = 1$ for half of the individuals and $Z_1 = 0$ for the others. Suppose X_1 and X_2 are independent standard normal and $Y \sim \mathcal{N}(Z_1 X_2 \beta_2, 1)$, for some coefficient $\beta_2 \in \mathbb{R}$. Then, (1) is true if and only if $\beta_2 = 0$. Thus, testing $\mathcal{H}_{0,j}$ for $j \in \{1, 2\}$ may be helpful to distinguish X_2 from X_1 , but it does not inform us about the interaction with Z_1 . This motivates the more flexible hypothesis testing framework presented below, which will allow us to automatically discover whether a significant association holds only within a specific population subgroup. Continuing with the previous example, our method will be able to discover that X_2 is important within the population subgroup with $Z_1 = 1$, and it will report no evidence of such association among individuals with $Z_1 = 0$. Crucially, there will be no need to know in advance that X_2 interacts with Z_1 .

For any variable index $j \in [p]$ and constant $G \in \mathbb{N}$, consider a *fixed* function $\psi_j : \mathbb{R}^m \mapsto [G]$ that partitions the covariate space \mathbb{R}^m into G disjoint regions, each identifying a distinct population *subgroup*. For example, in the case of the toy model described above, the most informative partition function would intuitively be $\psi_1 = 1$ and $\psi_2 = 1 + \mathbb{1}[Z_1 = 1]$, where $\mathbb{1}[\cdot]$ is the indicator function. Note that the number of possible subgroups can generally vary implicitly across different variables, and G simply refers to the largest partition size across all $j \in [p]$. In fact, the function ψ_1 in the previous example defines a single trivial subgroup encompassing the whole population, which is a reasonable choice in this case since X_1 is unimportant for all individuals.

To extend the standard population-wide conditional testing framework (1), we propose to study *subgroup-specific* conditional hypotheses of the form

$$\mathcal{H}_{0,j}^{(g)} : Y^{(g)} \perp\!\!\!\perp X_j^{(g)} \mid X_{-j}^{(g)}, Z^{(g)}, \quad (4)$$

for each $j \in [p]$ and $g \in [G]$, where $(X^{(g)}, Y^{(g)}, Z^{(g)})$ denotes a random sample from the distribution $P_{X,Y,Z}$ restricted to $\{z \in \mathbb{R}^m : \psi_j(z) = g\}$, which we call $P_{X,Y,Z}^{(g)}$. Note that the dependence of the group g on the variable index j is not shown explicitly to shorten the notation. Considering (4) for all $j \in [p]$ and $g \in [G]$ gives a multiple testing problem with mG hypotheses, if the ψ_j are fixed. In the special case of $G = 1$, (4) reduces to (1). In general, however, a rejection of (4) is more informative than one of (1), because (4) is implied by (1) but the converse is not true. In particular, rejecting (4) can help us to pinpoint the effect of X_j to subgroup g .

An appropriate choice of ψ_j is therefore essential to ensure we test interesting hypotheses. In fact, (4) can be more informative than (1) only if the subgroups are relatively homogeneous in their associations with Y . Thus, unless prior information about $P_{Y|X,Z}$ is available, ψ_j should be data-driven. As mentioned in Section 1, there exist many subgroup discovery algorithms, but one must be careful that allowing the hypotheses to be random may create the risk of selection bias in the subsequent tests. In particular, if the same data utilized to select the subgroups were re-used naively to test the hypotheses in (4), the type-I errors might be inflated [18], consistently with Figure 1. Sample splitting could of course circumvent this issue, but it is unnecessarily wasteful.

3.2 Knockoff-invariant subgroups

We now begin to present a novel method that can powerfully test subgroup-specific conditional hypotheses (4) adaptively learnt from the data, while controlling the false discovery rate. The key idea is that one can select

informative subgroups and then test the corresponding hypotheses while avoiding an uncontrollable inflation of the type-I errors, as long as the learning component of the analysis satisfies a suitable invariance property. This section defines the necessary invariance property and explains how to achieve it in practice.

Let $\tilde{\mathbf{X}} \in \mathbb{R}^{n \times p}$ be a matrix of knockoffs for \mathbf{X} , generated with existing techniques. For convenience, we refer to the knockoff-augmented data as $\mathcal{D} = (\mathbf{X}, \tilde{\mathbf{X}}, \mathbf{Y}, \mathbf{Z})$. For any $j \in [p]$, let $\hat{\psi}_j : \mathbb{R}^m \mapsto \mathbb{N}$ be a random function whose parameters may depend on \mathcal{D} as well as on an independent random matrix \mathbf{V} generated by the analyst, which is specified precisely below; i.e., $\hat{\psi}_j(z; [\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}, \mathbf{Z}, \mathbf{V}) \in \mathbb{N}$. In the following, we will focus on a special class of invariant partition functions that are allowed to depend on the data in \mathcal{D} but cannot make use of any knowledge about which variables are real and which are knockoffs.

Definition 1. *Let $\mathbf{V} \in \{0, 1\}^{n \times p}$ be i.i.d. Bernoulli random variables generated by the analyst. Denote as $[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})} \in \mathbb{R}^{n \times 2p}$ the concatenation of \mathbf{X} and $\tilde{\mathbf{X}}$, with the i -th observation of X_j swapped with its knockoff if and only if $V_{ij} = 1$. Then, we say that a vector-valued partition function $\hat{\psi} : \mathbb{R}^m \mapsto [G]^p$ is knockoff-invariant if it can be written as*

$$\hat{\psi}_j(z; [\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}, \mathbf{Z}, \mathbf{V}) = \hat{\psi}_j^0(z; [\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Y}, \mathbf{Z}),$$

for all $j \in [p]$, $z \in \mathbb{R}^m$, and for some fixed $\hat{\psi}^0 : \mathbb{R}^m \rightarrow [G]^p$ whose parametrization may depend on the data with swapped knockoffs. Above $\hat{\psi}_j$ and $\hat{\psi}_j^0$ are the j -th components of $\hat{\psi}$ and $\hat{\psi}^0$.

We pause for a moment to explain Definition 1 in generality. The intuition is that the partition function $\hat{\psi}$ can only be learnt by looking at a modified data set in which the identities of the real variables are masked by random swapping with their corresponding knockoffs. Although this (deliberately) introduces some noise, such constraint does not prevent the analyst from learning how to partition the covariate space into useful subgroups because the masked data in $[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}$ still carries valuable information. In fact, each column of this matrix contains approximately $n/2$ real observations, and any machine learning algorithm $\hat{\psi}^0$ may be utilized to extract useful knowledge from those data. Further, Definition 1 is highly versatile because it allows one to employ any function $\hat{\psi}^0$, which may involve sophisticated statistical or machine learning algorithms, to discover meaningful subgroups.

3.3 Subgroup learning using interaction-based models

An intuitive approach to discovering informative subgroups while satisfying Definition 1 involves implementing the function $\hat{\psi}^0$ using parametric regression models with interaction terms. Suppose that we suspect some of the variables denoted as X may interact with certain covariates Z , akin to the earlier illustration in Section 3.1. For simplicity we assume here that the covariates take binary values, although one could also similarly handle more general categorical covariates, or even continuous-valued covariates, through standard discretization techniques.

In this scenario, the analyst first generates a matrix \mathbf{V} of i.i.d. Bernoulli random variables, independent of the data. Then, the analyst proceeds by fitting a sparse generalized linear model (i.e., the lasso) to predict \mathbf{Y} given $[[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Z}]$, after augmenting the data matrix with all possible pairwise interactions between \mathbf{Z} and $[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}$. Several algorithms exist for fitting such models [11, 43, 45], making this approach practical and effective.

After tuning the lasso regularization via cross-validation, let $\hat{\beta}_j$ and $\hat{\beta}_{j+p}$ indicate the estimated main effects for the (randomly swapped) variables X_j and \tilde{X}_j , for any $j \in [p]$. Let also $\hat{\gamma}_{l,j}$ and $\hat{\gamma}_{l,j+p}$ indicate the corresponding interaction coefficients involving Z_l and the (randomly swapped) variables X_j and \tilde{X}_j , respectively, for all $l \in [m]$ and $j \in [p]$. Given an upper bound G_{\max} on the number of interactions per variable (e.g., $G_{\max} = 2$), for all $j \in [p]$ let $\hat{\mathcal{I}}_j \subset [m]$ indicate the subset of G_{\max} covariates with the strongest (nonzero) interactions involving X_j or \tilde{X}_j . Above, we understand that any ties are broken at random and $|\hat{\mathcal{I}}_j| < G_{\max}$ if the number of covariates with nonzero interaction coefficients is too small. In other words, we compute

$$\hat{\mathcal{I}}_j := \{l \in [m] : |\hat{\gamma}_{l,j}| + |\hat{\gamma}_{l,j+p}| > 0, \sum_{l' \neq l} \mathbb{1}[|\hat{\gamma}_{l',j}| + |\hat{\gamma}_{l',j+p}| > |\hat{\gamma}_{l,j}| + |\hat{\gamma}_{l,j+p}|] > m - G_{\max}\}.$$

This model thus links X_j to at most G_{\max} covariates. The corresponding functions $\hat{\psi}_j$ take values in $\{1, \dots, 2^{|\hat{\mathcal{Z}}_j|}\}$, which indexes all possible configurations of the covariates in $\hat{\mathcal{Z}}_j$.

Table 1 helps visualize this approach by highlighting a thought experiment involving 3 explanatory variables and 3 binary covariates. In this toy example, we imagine that the lasso selects Z_1, Z_2 as potentially interacting with X_1 , and Z_3 for X_2 . No interactions with X_3 are detected. Thus, the population is partitioned into 4 subgroups for X_1 , 2 subgroups for X_2 , and 1 trivial subgroup for X_3 , yielding a total of 7 data-driven hypotheses.

Table 1: Subgroups selected by fitting a lasso model with interactions, and corresponding interpretations of the possible discoveries obtained with our method, in an imaginary toy example with 3 variables and 3 binary covariates. The term ‘‘influences’’ is utilized loosely here to indicate a (possibly non-causal) significant conditional association.

Variable	Covariates	Partition		Interpretation of the findings (upon rejection)
		Label	Definition	
X_1	$\{Z_1, Z_2\}$	1	$Z_1 = 0, Z_2 = 0$	X_1 influences Y if $Z_1 = 0$ and $Z_2 = 0$
		2	$Z_1 = 0, Z_2 = 1$	X_1 influences Y if $Z_1 = 0$ and $Z_2 = 1$
		3	$Z_1 = 1, Z_2 = 0$	X_1 influences Y if $Z_1 = 1$ and $Z_2 = 0$
		4	$Z_1 = 1, Z_2 = 1$	X_1 influences Y if $Z_1 = 1$ and $Z_2 = 1$
X_2	$\{Z_3\}$	1	$Z_3 = 0$	X_2 influences Y if $Z_3 = 0$
		2	$Z_3 = 1$	X_2 influences Y if $Z_3 = 1$
X_3	\emptyset	1	All individuals	X_3 influences Y

Note that the parameter G_{\max} generally controls an important trade-off between the power of our method and the interpretability of the findings. A larger value of G_{\max} tends to lead to more specific hypotheses corresponding to smaller subgroups, but it also makes it more difficult to reject those hypotheses, as it will become clear soon. In fact, partitioning roughly reduces the effective number of samples available during the testing phase by a factor of up to $2^{G_{\max}}$, in the case of binary covariates. Concretely, all demonstrations presented in this paper will utilize $G_{\max} \in \{1, 2\}$, but the optimal value of G_{\max} may generally depend on the data.

3.4 The subgroup-selective knockoff filter

For any knockoff-invariant partition function $\hat{\psi}$ and any variable index $j \in [p]$, let $\hat{G}_j([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}, \mathbf{Z}) = \max_{z \in \mathbb{R}^m} \hat{\psi}_j(z) \in \mathbb{N}$ be the number of disjoint subgroups induced by $\hat{\psi}_j$. Further, let $\hat{G} = \sum_{j=1}^p \hat{G}_j$. The dependence of $\hat{\psi}(x)$, \hat{G}_j , and \hat{G} on the data will not be shown explicitly hereafter unless needed to avoid ambiguity. Let $[\mathbf{T}, \tilde{\mathbf{T}}] \in \mathbb{R}^{2\hat{G}}$ denote a vector of importance measures for all variables and knockoffs in each region of the covariate space determined by $\hat{\psi}$. We will explain later how to compute these importance measures. For now, note that \mathbf{T} (resp. $\tilde{\mathbf{T}}$) is the concatenation of p sub-vectors T_j^g (resp. \tilde{T}_j^g) for all $j \in [p]$ and $g \in [\hat{G}_j]$, whose elements quantify the importance of X_j (resp. \tilde{X}_j) in predicting Y within subgroup g .

Without loss of generality, the importance measures $[\mathbf{T}, \tilde{\mathbf{T}}]$ can be written as the output of a (possibly randomized) function τ applied to a knockoff-augmented data set $[\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}, \mathbf{Z}$; i.e.,

$$[\mathbf{T}, \tilde{\mathbf{T}}] = \tau([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}, \mathbf{Z}) = [\mathbf{t}([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}, \mathbf{Z}), \tilde{\mathbf{t}}([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}, \mathbf{Z})]. \quad (5)$$

Above, \mathbf{t} (resp. $\tilde{\mathbf{t}}$) are defined in terms of τ , as its first (resp. last) \hat{G} elements. In analogy with the familiar case of the model-X knockoff filter, however, achieving false discovery rate control requires imposing some constraints on the function τ .

For any $\mathcal{S} \subseteq [\hat{G}]$, whose elements uniquely identify pairs (j, g) for $j \in [p]$ and $g \in [\hat{G}_j]$, let $[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathcal{S})}$ be the matrix obtained from $[\mathbf{X}, \tilde{\mathbf{X}}]$ after swapping the sub-column $\mathbf{X}_j^{(g)}$, which contains all observations of X_j in subgroup g , with the corresponding knockoffs, for all $(j, g) \in \mathcal{S}$. Note the change of notation compared

to (2), where swapping was simultaneous for an entire column. Then, we ask that swapping a real variable with its knockoff within any subgroup should have the only effect of swapping the corresponding importance measures in that subgroup:

$$\boldsymbol{\tau}(\mathbf{Y}, [\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(S)}, \mathbf{Z}) = [\mathbf{t}(\mathbf{Y}, [\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Z}), \tilde{\mathbf{t}}(\mathbf{Y}, [\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Z})]_{\text{swap}(S)}. \quad (6)$$

In truth, it would be sufficient for this invariance to hold in distribution if the function $\boldsymbol{\tau}$ contained additional independent randomization (e.g., cross-validation to tune some hyper-parameters). However, we can safely ignore this detail here to avoid complicating the notation.

For our method to be powerful, the function $\boldsymbol{\tau}$ should be such that a larger value in position (j, g) of \mathbf{T} indicates evidence that Z_j is associated with Y in subgroup g . By contrast, a larger value in position (j, g) of $\tilde{\mathbf{T}}$ should point to an association of \tilde{Z}_j with Y , which we know must be spurious. An example of a valid $\boldsymbol{\tau}$ based on the (generalized) lasso is described in Section 3.5. That approach is inspired by Li et al. [31] and it is both intuitive and relatively inexpensive, as all of its components can be computed in parallel, but it is not the only possible one.

In general, given any importance measures $[\mathbf{T}, \tilde{\mathbf{T}}]$ obtained as described above, we assemble test statistics $\mathbf{W} \in \mathbb{R}^{\hat{G}}$ for all selected hypotheses as usual, by computing $W_j^g = T_j^g - \tilde{T}_j^g$ for each pair (j, g) . Finally, we vectorize \mathbf{W} and apply the selective SeqStep+ procedure of Barber and Candès [19] (i.e., the final component of the standard knockoff filter) in order to compute an adaptive significance threshold that determines which hypotheses should be rejected. Figure A2 in Appendix A2 provides a schematic visualization of this procedure.

The following theorem establishes that our method controls the false discovery rate. It is worth emphasizing that this result is far from trivial and does not stem directly from Candès et al. [5] or Li et al. [31]. While our proof strategy is inspired by those earlier works, our scenario is inherently more complex. In fact, our hypotheses are not predetermined but instead automatically selected using the adaptive partition function $\hat{\psi}$, which depends on the same data used to compute the test statistics. This is a delicate situation that can generally introduce selection bias [16], as previewed earlier in Figure 1. Nevertheless, our method is meticulously designed to circumvent such selection bias through the incorporation of independent randomization within the matrix \mathbf{V} , which is used by the analyst to partially mask the training data for $\hat{\psi}$.

Theorem 1. *Consider a knockoff-augmented data set $[\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}, \mathbf{Z}$, and let $\mathbf{V} \in \{0, 1\}^{n \times p}$ be i.i.d. Bernoulli random variables generated by the analyst, independent of the data. Let $\hat{\psi}$ be a knockoff-invariant partition function with total size $\hat{G} = \sum_{j=1}^p \hat{G}_j$, trained looking only at $[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Y}, \mathbf{Z}$. Let also $\mathbf{W} \in \mathbb{R}^{\hat{G}}$ indicate test statistics, for the hypotheses defined by $\hat{\psi}$, computed as described above using a function $\boldsymbol{\tau}$ that satisfies (6). Suppose $\mathbf{U} \in \{\pm 1\}^{\hat{G}}$ is a random vector with independent entries such that: $U_j^g = \pm 1$ with probability $1/2$ if $\mathcal{H}_{0,j}^{(g)}$ in (4) is true and $U_j^g = +1$ otherwise, for all $j \in [p]$ and $g \in [\hat{G}_j]$. Then, $\mathbf{W} \mid \hat{\psi} \stackrel{d}{=} \mathbf{W} \odot \mathbf{U} \mid \hat{\psi}$.*

The proof of Theorem 1 is in Appendix A5. An immediate corollary is that the standard knockoff filter applied to statistics \mathbf{W} computed as described above will control the false discovery rate for (4) conditional on $\hat{\psi}$; e.g., see Section 2.2 and Candès et al. [5].

3.5 Computing powerful test statistics for the subgroup-selective knockoff filter

We describe here an implementation of our method based on generalized linear models. As the property in (6) requires swapping X_j and \tilde{X}_j in subgroup g to have the only effect of swapping the corresponding T_j^g and \tilde{T}_j^g , an intuitive approach inspired by [31] is the following. First, a sparse generalized linear model (e.g., the lasso) is trained to predict \mathbf{Y} given $[[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Z}]$, tuning the regularization via cross-validation. The absolute values of the regression coefficients serve as “prior” importance measures, T_j^{prior} and $\tilde{T}_j^{\text{prior}}$, for all variables and knockoffs indexed by $j \in [p]$. These are combined pairwise into a weight π_j for each j ; for example, as $\pi_j = \zeta(T_j^{\text{prior}} + \tilde{T}_j^{\text{prior}})$, where ζ is a positive and decreasing function such as $\zeta(t) = 1/(0.05 + t)$. Larger values of π_j suggest the j -th variable is more likely to have a significant effect among some individuals. Similar weights can also be defined for the covariates. Then, separately for each (j, g) , a new model is fitted

using only the data from the g -th subgroup to predict $\mathbf{Y}^{(g)}$ given $[[\mathbf{X}^{(g)}, \tilde{\mathbf{X}}^{(g)}]_{\text{swap}(\mathbf{V})}, \mathbf{Z}^{(g)}]$, after restoring the true variable and knockoff identities in the j -th and $(j+p)$ -th columns. See Figure A3 for a schematic visualization of this algorithm.

This model utilizes feature-specific regularization that depends on two hyper-parameters, $\lambda^{(j,g)} > 0$ and $\xi^{(j,g)} \in [0, 1]$, both tuned by cross-validation, and on π . Specifically, the penalty for the l -th variable is $\lambda_l^{(j,g)} = \lambda^{(j,g)}(1 - \xi^{(j,g)}) + \xi^{(j,g)}\pi_l$, for all $l \in [p]$. If $\xi^{(j,g)} = 0$, this reduces to a standard lasso looking only at subgroup g . But, in general, larger values of $\xi^{(j,g)}$ can make our approach more powerful, as we gather strength from the data in all subgroups. In fact, null variables will tend to receive smaller values of π and will thus be less likely to be incorrectly selected by the final model, thereby reducing the noise in the test statistics. Of course, the weights π may not always be informative, hence why $\xi^{(j,g)}$ is tuned by cross-validation. For that purpose, an expensive two-dimensional grid search can be avoided by tuning first $\lambda^{(j,g)} > 0$ and then $\xi^{(j,g)}$.

Finally, the importance measures for the j -th variable and knockoff in subgroup g are defined as the absolute values of the regression coefficients for X_j and \tilde{X}_j , respectively. It is easy to prove this procedure satisfies (6). In fact, conditional on $\mathbf{Y}, [\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Z}$, swapping all observations of X_j in subgroup g with the corresponding \tilde{X}_j for any pair (j, g) simply results in swapping $T_j^{(g)}$ with $\tilde{T}_j^{(g)}$, as π is unperturbed because it only depends on $[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}$.

While it is convenient to think about the lasso, the main idea behind our approach is applicable with any model. The key to achieving (6) is that the importance measures indexed by (j, g) are computed by a model that sees the other variables and the observations in other subgroups only through the lenses of the data with randomly swapped knockoffs. The weights π must be estimated based on $[[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Z}]$ to make the test statistics mutually independent, which is required for false discovery rate control; see the proof of Theorem 1. Further, revealing the identities of X_j and \tilde{X}_j within subgroup g is essential to achieve non-trivial power—otherwise, it would be impossible to tell important variables apart from knockoffs. A generalization of the above solution would involve estimating different π weights in different subgroups, as sketched by Figure A3 in Appendix A2. That approach still satisfies (6) if the models utilized to compute the priors only look at the randomly swapped data in $[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}$.

4 Empirical demonstrations

4.1 Demonstrations with synthetic data

We generate synthetic data with 20 binary variables and 80 covariates. The first 20 covariates are independently sampled from a Bernoulli(0.5) distribution. The remaining 60 covariates follow a Gaussian autoregressive model of order one with correlation parameter 0.5. The outcome is generated from a linear model with heterogeneous effects and homoscedastic standard Gaussian noise. The linear coefficients in this model are zero for half of the variables, for all of the binary covariates, and for half of the 60 continuous covariates. That is, the expected outcome for the i -th individual is $\mathbb{E}[Y^i | X^i, Z^i] = \sum_{j=1}^p X_j^i \beta_j^i + \sum_{j=21}^m Z_j^i \gamma_j$, with $\beta_j^i = \bar{\beta}_j Z_{l_{j,1}}^i Z_{l_{j,2}}^i$, for $p = 20$ and $m = 80$. Above, the β coefficients are individual-specific, while the γ are constant. The parameters $\bar{\beta}_j$ and γ_j are initialized with absolute value equal to 4 and independent random signs, for each $j \in [p]$. Then, separately for each variable j and individual i , the individualized coefficient β_j^i is set equal to the product of $\bar{\beta}_j$ and the two binary covariates $Z_{l_{j,1}}$ and $Z_{l_{j,2}}$, for some $l_{j,1}, l_{j,2} \in [20]$. Thus, each variable X_j has an effect on Y only within the subgroups with $Z_{l_{j,1}} = 1$ and $Z_{l_{j,2}} = 1$, which contains approximately 1/4 of all individuals if $l_{j,1} \neq l_{j,2}$. The indices $l_{j,1}, l_{j,2}$ are sampled with replacement from $\{1, \dots, 20\}$, independently for each j . Our goal is to discover which variables are non-null within which subgroups, as powerfully (seeking more numerous findings) and precisely (seeking more informative subgroups) as possible.

The subgroup-selective knockoff filter is compared to the three benchmarks outlined in Section 1.2. The first benchmark is the vanilla knockoff filter of [5], which controls the false discovery rate but can only test population-wide hypotheses. The second benchmark is the data-splitting version of our method, which utilizes half of the samples to learn the partition and the other half to test the conditional hypotheses. This

tests the same hypotheses and enjoys the same statistical guarantees as the subgroup-selective knockoff filter but is less powerful. The third benchmark is a greedy version of our method that naively utilizes all data twice, first to select the partitions and then to test the conditional hypotheses, without randomly masking the identity of the knockoffs; this does not control the false discovery rate due to selection bias.

All methods are compared in terms of their average proportion of false discoveries, their power, and the informativeness of their findings. The latter is quantified by two complementary metrics: the average homogeneity of the rejected hypotheses, defined as the proportion of individuals in the reported subgroup for which the variable of interest has a non-zero coefficient within the data-generating model, and the corresponding heterogeneity of individual effects, defined as the standard deviation of the true model coefficients.

Figure 1, previewed in Section 3.4, reports the performances of all methods over 100 independent experiments. As discussed in Section 1.2, these results confirm the validity of our approach and demonstrate its advantages relative to the naive, data-splitting, and vanilla benchmarks. In summary, our method leads to discoveries that are more informative than those of the vanilla knockoff filter, more numerous than those of the sample-splitting benchmark, and more reliable than those of the naive greedy benchmark.

Appendix A4.1 presents additional numerical results, demonstrating the use of the subgroup-selective knockoff filter in a “transfer learning” setting inspired by [50]. The goal there is complementary: we wish to detect variables with a robust association across different subgroups. This can be achieved with an extension of our method described in Appendix A3, which tests partial conjunctions of the hypotheses in (4).

4.2 Demonstrations with data from a randomized experiment

We conduct numerical experiments based on a data set from a 2-week long randomized experiment involving 80,000 blood donors in China [51], which was designed to investigate the effectiveness of different donation incentives. In this experiment, 80,000 individuals were randomly divided into 7 groups: a control group of size 14,000 and 6 treated groups of size 11,000 each. The treated groups received a text message with a reminder to donate, while the control group received nothing. Treated groups 2–5 received further encouragement to donate: the promise of an individual monetary reward (equivalent to \$5.00–\$8.30, depending on the amount of blood donated) for group 2; a suggestion to bring a friend for group 3, both the individual monetary reward and the friend suggestion for group 4; a suggestion to bring a friend and the promise of a group reward (equivalent to the individual reward) for group 5; a suggestion to bring a friend, the promise of a group reward (as for the previous group), and the promise of a small extra gift for group 6. See Table A1 in Appendix A4.2 for a summary of this experimental design. In the end, 797 individuals made a donation. These data also contain covariates including age, sex, weight, blood type, marital status, education level, occupation, residency status (local or non-local), and time since the last blood donation. Missing values (less than 4%) are imputed with the corresponding sample median (or mode, if categorical).

To simplify the task of generating knockoffs, we approximate the joint distribution of the binary treatments X_1, \dots, X_5 by imagining that different individuals were assigned to one of the six possible groups listed in Table A1 independently of one another, with probability proportional to the observed group size. This model is not perfect because the treatment group sizes were in truth fixed a priori, but the simplification is useful as it leads to a manageable joint treatment distribution for which exact knockoffs can be generated by the Metropolized algorithm of [24]. Goodness-of-fit diagnostics for the knockoffs thus obtained are reported in Tables A2–A3. In principle, it would also be possible to generate knockoffs without the above independence assumption using an algorithm similar to that of Sesia et al. [52], but that would be more technically involved and seems likely unnecessary in this case.

Given that the positive outcomes from this randomized experiment were very rare (approximately 1%), we expect that any conditional testing method may not be able to make many discoveries with these data. Therefore, we find it more interesting to conduct a semi-synthetic analysis in which the true donation events are replaced by simulated outcomes from an imaginary causal model conditional on the treatments and covariates. This setup remains quite realistic, as it involves real variables, but it reduces the outcome imbalance and increases the effective sample size. Consequently, the analysis will be more informative regarding the effectiveness of the our methodology applied to large data sets. Further, as the ground truth

is known exactly in such a controlled setting, we can have access to useful diagnostics, including the type-I errors.

For each of the 80,000 individuals, an imaginary donation outcome is simulated from a logistic model based on the m real covariates $Z \in \mathbb{R}^m$ and the treatment variables $X \in \mathbb{R}^5$; i.e.,

$$\text{logit}(\mathbb{P}[Y = 1 | X, Z]) = -c + \sum_{j=1}^m Z_j b_j + a \sum_{j=1}^5 X_j \cdot g_j(Z),$$

where $c \in \mathbb{R}$ is an intercept, $b_1, \dots, b_m \in \mathbb{R}$ are linear coefficients, and each g_j is a binary-valued non-linear function of Z ; see Appendix A4.3 for more details. The first treatment (text reminder) is effective for individuals who did not donate recently and for those who are non-residents. The individual reward is effective for students, and twice as much for those who are male. The friends request is effective for female or unmarried individuals. The group reward is effective for males or students. The small gift is effective for individuals with fewer than 16 years of education. The goal is to discover which treatments are effective within which subgroups, as powerfully and precisely as possible. The subgroup-selective knockoff filter is applied as in the previous section, defining the subgroup-specific hypotheses based on the top two candidate covariates selected by a (logistic) lasso model fitted on the data with randomly swapped knockoffs.

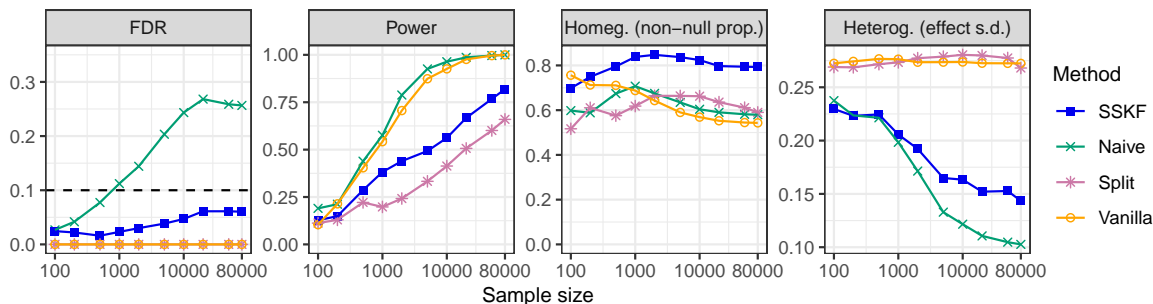


Figure 2: Performance of the subgroup-selective knockoff filter and benchmarks in simulations based on data from a randomized blood donation experiment. Other details are as in Figure 1.

Figure 2 summarizes the results of 100 experiments with independent knockoffs and outcomes, varying the sample size. Each time, a random subset of individuals is analyzed. The performance of the subgroup-selective knockoff filter is quantified in terms of false discovery rate and power, as well as in terms of the homogeneity of the selected hypotheses, as in Section 4.1. The same benchmarks as in Section 4.1 are also considered here. The subgroup-selective knockoff filter controls the false discovery rate and is more powerful than data splitting. As expected, the naive approach does not control the false discovery rate, while the vanilla knockoff filter can only discover less informative population-wide associations. The list of discoveries obtained in the first experiment with sample size 80,000 is provided in Table A4. Among the 11 findings, that of the friends request among unmarried females has the largest test statistic; indeed, we know this treatment has a causal effect for all those individuals. Although no type-I errors are made in this case, not all findings are equally informative. For example, the reminder can be effective among both male and female residents, but the subgroup-selective knockoff filter does not tell us that it is only causal for the 60% of them who are married. In fact, this interaction is not discovered because the initial model fitted on the data with randomly swapped knockoffs failed to select the corresponding covariate, likely due to the limited sample size.

5 Discussion

This paper has extended the model-X framework of [5] to enable more informative inferences in the presence of sample heterogeneity and interactions. Flexibility is the main strength of the proposed method, which can borrow strength from any statistical or machine learning model while providing rigorous guarantees in finite samples.

A promising opportunity for future research is to apply the subgroup-selective knockoff filter to large-scale data from genome-wide association studies [53], such as the UK Biobank [54] or the Millions Veteran Project [55]. The populations sampled by those studies are heterogeneous [1], and our method may be helpful to discover possible gene-environment interactions [56]. Further, the abundance of variables and observations in those studies will translate into further advantages for our methodology, which generally tends to perform better when applied to larger-scale data sets. Computing powerful test statistics based on extremely big data sets involves some computational challenges, but the model-X framework has already been demonstrated to be quite scalable [52, 57, 58].

Further, it would be interesting to explore whether our inferences could be (at least partly) derandomized using techniques inspired by Ren and Barber [59] and Ren, Wei, and Candès [60].

A software implementation of our methods is available online at <https://github.com/msesia/i-modelx>, along with the code needed to reproduce our numerical results.

Acknowledgements

M.S. is also affiliated with the Department of Computer Science at the University of Southern California. M. S. was supported by NSF grant DMS 2210637 and an Amazon Research Award. The authors are grateful to Edgar Dobriban, Yingying Fan, Nikolaos Ignatiadis, and Stefan Wager for providing insightful feedback about an earlier version of this manuscript.

References

- [1] R. E. Peterson, K. Kuchenbaecker, R. K. Walters, C.-Y. Chen, A. B. Popejoy, S. Periyasamy, M. Lam, C. Iyegbe, R. J. Strawbridge, L. Brick, et al. “Genome-wide association studies in ancestrally diverse populations: opportunities, methods, pitfalls, and recommendations”. In: *Cell* 179.3 (2019), pp. 589–603.
- [2] J. J. Berg, A. Harpak, N. Sinnott-Armstrong, A. M. Joergensen, H. Mostafavi, Y. Field, E. A. Boyle, X. Zhang, F. Racimo, J. K. Pritchard, et al. “Reduced signal for polygenic adaptation of height in UK Biobank”. In: *Elife* 8 (2019), e39725.
- [3] J. Peters, P. Bühlmann, and N. Meinshausen. “Causal inference by using invariant prediction: identification and confidence intervals”. In: *J. R. Stat. Soc. B* (2016), pp. 947–1012.
- [4] P. Bühlmann. “Invariance, causality and robustness”. In: *Stat. Science* 35.3 (2020), pp. 404–426.
- [5] E. Candès, Y. Fan, L. Janson, and J. Lv. “Panning for Gold: Model-X Knockoffs for High-dimensional Controlled Variable Selection”. In: *J. R. Stat. Soc. B.* 80 (2018), pp. 551–577.
- [6] G. George, M. R. Haas, and A. Pentland. *Big data and management*. 2014.
- [7] S. Aral and D. Walker. “Identifying influential and susceptible members of social networks”. In: *Science* 337.6092 (2012), pp. 337–341.
- [8] T. Sun and S. J. Taylor. “Displaying things in common to encourage friendship formation: A large randomized field experiment”. In: *Quant. Mark. Econ.* 18 (2020), pp. 237–271.
- [9] S. Wrobel. “An algorithm for multi-relational discovery of subgroups”. In: *European symposium on principles of data mining and knowledge discovery*. Springer. 1997, pp. 78–87.
- [10] P. K. Novak, N. Lavrač, and G. I. Webb. “Supervised descriptive rule discovery: A unifying survey of contrast set, emerging pattern and subgroup mining.” In: *Journal of Machine Learning Research* 10.2 (2009).
- [11] J. Bien, J. Taylor, and R. Tibshirani. “A lasso for hierarchical interactions”. In: *Ann. Stat.* 41.3 (2013), p. 1111.

- [12] L. Tian, A. A. Alizadeh, A. J. Gentles, and R. Tibshirani. “A simple method for estimating interactions between a treatment and a large number of covariates”. In: *J. Am. Stat. Assoc.* 109.508 (2014), pp. 1517–1532.
- [13] G. Verbeke and E. Lesaffre. “A linear mixed-effects model with heterogeneity in the random-effects population”. In: *J. Am. Stat. Assoc.* 91.433 (1996), pp. 217–221.
- [14] A. Caron, G. Baio, and I. Manolopoulou. “Estimating individual treatment effects using non-parametric regression models: A review”. In: *J. R. Stat. Soc. A.* 185.3 (2022), pp. 1115–1149.
- [15] Y. Benjamini and Y. Hochberg. “Controlling the false discovery rate: a practical and powerful approach to multiple testing”. In: *J. R. Stat. Soc. B.* 57 (1995), pp. 289–300.
- [16] S. F. Assmann, S. J. Pocock, L. E. Enos, and L. E. Kasten. “Subgroup analysis and other (mis) uses of baseline data in clinical trials”. In: *The Lancet* 355.9209 (2000), pp. 1064–1069.
- [17] D. I. Cook, V. J. Gebski, and A. C. Keech. “Subgroup analysis in clinical trials”. In: *Medical Journal of Australia* 180.6 (2004), p. 289.
- [18] S. Wager and S. Athey. “Estimation and inference of heterogeneous treatment effects using random forests”. In: *J. Am. Stat. Assoc.* 113.523 (2018), pp. 1228–1242.
- [19] R. F. Barber and E. Candès. “Controlling the false discovery rate via knockoffs”. In: *Ann. Stat.* 43.5 (2015), pp. 2055–2085.
- [20] M. Sesia, C. Sabatti, and E. Candès. “Gene hunting with hidden Markov model knockoffs”. In: *Biometrika* 106 (2019), pp. 1–18.
- [21] J. R. Gimenez, A. Ghorbani, and J. Zou. “Knockoffs for the mass: new feature importance statistics with false discovery guarantees”. In: *22nd International Conference on Artificial Intelligence and Statistics*. PMLR. 2019, pp. 2125–2133.
- [22] J. Jordon, J. Yoon, and M. van der Schaar. “KnockoffGAN: Generating knockoffs for feature selection using generative adversarial networks”. In: *International conference on learning representations*. 2018.
- [23] Y. Romano, M. Sesia, and E. Candès. “Deep Knockoffs”. In: *J. Am. Stat. Assoc.* 0.ja (2019), pp. 1–27.
- [24] S. Bates, E. Candès, L. Janson, and W. Wang. “Metropolized knockoff sampling”. In: *J. Am. Stat. Assoc.* (2020), pp. 1–15.
- [25] Y. Fan, E. Demirkaya, G. Li, and J. Lv. “RANK: large-scale inference with graphical nonlinear knockoffs”. In: *J. Am. Stat. Assoc.* (2019).
- [26] R. F. Barber, E. Candès, and R. J. Samworth. “Robust inference with knockoffs”. In: *Ann. Stat.* 48.3 (2020), pp. 1409–1431.
- [27] J. Liu and P. Rigollet. “Power analysis of knockoff filters for correlated designs”. In: *Adv. Neural Inf. Process. Syst.* (2019).
- [28] W. Wang and L. Janson. “A High-Dimensional Power Analysis of the Conditional Randomization Test and Knockoffs”. In: *Biometrika* (2021).
- [29] A. Spector and L. Janson. “Powerful knockoffs via minimizing reconstructability”. In: *Ann. Stat.* 50.1 (2022), pp. 252–276.
- [30] E. Katsevich and A. Ramdas. “On the power of conditional independence testing under model-X”. In: *Electron. J. Stat.* 16.2 (2022), pp. 6348–6394.
- [31] S. Li, M. Sesia, Y. Romano, E. Candès, and C. Sabatti. “Searching for robust associations with a multi-environment knockoff filter”. In: *Biometrika* (2021).
- [32] Y. Zhao, D. Zeng, A. J. Rush, and M. R. Kosorok. “Estimating individualized treatment rules using outcome weighted learning”. In: *J. Am. Stat. Assoc.* 107.499 (2012), pp. 1106–1118.
- [33] S. Athey and G. Imbens. “Recursive partitioning for heterogeneous causal effects”. In: *Proc. Natl. Acad. Sci. U.S.A.* 113.27 (2016), pp. 7353–7360.
- [34] L. Yao, S. Li, Y. Li, M. Huai, J. Gao, and A. Zhang. “Representation learning for treatment effect estimation from observational data”. In: *Adv. Neural Inf. Process. Syst.* 31 (2018).

- [35] R. Moore, F. P. Casale, M. Jan Bonder, D. Horta, L. Franke, I. Barroso, and O. Stegle. “A linear mixed-model approach to study multivariate gene–environment interactions”. In: *Nat. Genet.* 51.1 (2019), pp. 180–186.
- [36] S. R. Künzel, J. S. Sekhon, P. J. Bickel, and B. Yu. “Metalearners for estimating heterogeneous treatment effects using machine learning”. In: *Proceedings of the national academy of sciences* 116.10 (2019), pp. 4156–4165.
- [37] P. R. Hahn, J. S. Murray, and C. M. Carvalho. “Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion)”. In: *Bayesian Analysis* 15.3 (2020), pp. 965–1056.
- [38] X. Nie and S. Wager. “Quasi-oracle estimation of heterogeneous treatment effects”. In: *Biometrika* 108.2 (2021), pp. 299–319.
- [39] Y. Cui, M. R. Kosorok, E. Sverdrup, S. Wager, and R. Zhu. “Estimating heterogeneous treatment effects with right-censored data via causal survival forests”. In: *J. R. Stat. Soc. B.* 85.2 (2023), pp. 179–211.
- [40] J. L. Hill. “Bayesian nonparametric modeling for causal inference”. In: *J. Comput. Graph. Stat.* 20.1 (2011), pp. 217–240.
- [41] C. Louizos, U. Shalit, J. M. Mooij, D. Sontag, R. Zemel, and M. Welling. “Causal effect inference with deep latent-variable models”. In: *Adv. Neural Inf. Process. Syst.* 30 (2017).
- [42] X. Wen and M. Stephens. “Bayesian methods for genetic association analysis with heterogeneous subgroups: from meta-analyses to gene–environment interactions”. In: *Ann. Appl. Stat.* 8.1 (2014), p. 176.
- [43] M. Lim and T. Hastie. “Learning interactions via hierarchical group-lasso regularization”. In: *J. Comput. Graph. Stat.* 24.3 (2015), pp. 627–654.
- [44] G. Yu, J. Bien, and R. Tibshirani. “Reluctant interaction modeling”. In: *preprint at arXiv:1907.08414* (2019).
- [45] R. Tibshirani and J. Friedman. “A pliable lasso”. In: *J. Comput. Graph. Stat.* 29.1 (2020), pp. 215–225.
- [46] K. J. Friston, W. D. Penny, and D. E. Glaser. “Conjunction revisited”. In: *Neuroimage* 25.3 (2005), pp. 661–667.
- [47] Y. Benjamini and R. Heller. “Screening for partial conjunction hypotheses”. In: *Biometrics* 64.4 (2008), pp. 1215–1222.
- [48] Y. Nair and L. Janson. “Randomization Tests for Adaptively Collected Data”. In: *preprint at arXiv:2301.05365* (2023).
- [49] R. Tibshirani. “Regression shrinkage and selection via the lasso”. In: *J. R. Stat. Soc. Series B Stat. Methodol.* 58.1 (1996), pp. 267–288.
- [50] S. Li, Z. Ren, C. Sabatti, and M. Sesia. “Transfer learning in genome-wide association studies with knockoffs”. In: *Sankhya B* (2022), pp. 1–39.
- [51] T. Sun, G. Gao, and G. Z. Jin. “Mobile messaging for offline group formation in prosocial activities: A large field experiment”. In: *Manag. Science* 65.6 (2019), pp. 2717–2736.
- [52] M. Sesia, S. Bates, E. Candès, J. Marchini, and C. Sabatti. “False discovery rate control in genome-wide association studies with population structure”. In: *Proc. Natl. Acad. Sci. U.S.A.* 118.40 (2021).
- [53] J. D. Storey and R. Tibshirani. “Statistical significance for genomewide studies”. In: *Proc. Natl. Acad. Sci. U.S.A.* 100.16 (2003), pp. 9440–9445.
- [54] C. Bycroft, C. Freeman, D. Petkova, G. Band, L. T. Elliott, K. Sharp, A. Motyer, D. Vukcevic, O. Delaneau, J. O’Connell, A. Cortes, S. Welsh, A. Young, M. Effingham, G. McVean, S. Leslie, N. Allen, P. Donnelly, and J. Marchini. “The UK Biobank resource with deep phenotyping and genomic data”. In: *Nature* 562 (2018), pp. 203–209.
- [55] J. M. Gaziano, J. Concato, M. Brophy, L. Fiore, S. Pyarajan, J. Breeling, S. Whitbourne, J. Deen, C. Shannon, D. Humphries, et al. “Million Veteran Program: A mega-biobank to study genetic influences on health and disease”. In: *J. Clin. Epidemiol.* 70 (2016), pp. 214–223.

- [56] J. H. Moore, J. C. Gilbert, C.-T. Tsai, F.-T. Chiang, T. Holden, N. Barney, and B. C. White. “A flexible computational framework for detecting, characterizing, and interpreting statistical patterns of epistasis in genetic studies of human disease susceptibility”. In: *J. Theor. Biol.* 241.2 (2006), pp. 252–261.
- [57] M. Liu, E. Katsevich, L. Janson, and A. Ramdas. “Fast and powerful conditional randomization testing via distillation”. In: *Biometrika* (July 2021).
- [58] M. Sesia, E. Katsevich, S. Bates, E. Candès, and C. Sabatti. “Multi-resolution localization of causal variants across the genome”. In: *Nat. Commun.* 11.1 (2020), pp. 1–10.
- [59] Z. Ren and R. F. Barber. “Derandomized knockoffs: leveraging e-values for false discovery rate control”. In: *preprint at arXiv:2205.15461* (2022).
- [60] Z. Ren, Y. Wei, and E. Candès. “Derandomizing knockoffs”. In: *J. Am. Stat. Assoc.* 118.542 (2023), pp. 948–958.
- [61] S. Bates, M. Sesia, C. Sabatti, and E. Candès. “Causal inference in genetic trio studies”. In: *Proc. Natl. Acad. Sci. U.S.A.* 117.39 (2020), pp. 24117–24126.
- [62] K. A. Bollen. *Structural Equations with Latent Variables*. Vol. 210. John Wiley & Sons, 2014.
- [63] J. Pearl. *Causality*. Cambridge University Press, 2009.
- [64] A. Li and R. F. Barber. “Accumulation tests for FDR control in ordered hypothesis testing”. In: *J. Am. Stat. Assoc.* 112.518 (2017), pp. 837–849.

A1 Background on population-wide conditional testing

A1.1 The interpretation of conditional independence hypotheses

The model-X problem studied by Candès et al. [5] can be thought of as testing, for all $j \in [p]$, whether X_j is associated with Y given Z and all variables excluding X_j (i.e., X_{-j}); that is, whether the null hypothesis defined in (1) holds true: $\mathcal{H}_{0,j} : Y \perp\!\!\!\perp X_j \mid X_{-j}, Z$. For example, if one assumed a generalized linear model for $Y \mid X, Z$ (which we do not do here), then the hypothesis in (1) would reduce, under relatively mild assumptions on $P_{X|Z}$, to stating that the linear coefficient for X_j is zero [5].

More generally, without any parametric model for the distribution of the outcome, a rejection of (1) can be interpreted as a discovery that X_j has some effect on Y conditional on the other observed variables, for at least some of the individuals in the population. (Note that we are not assuming the individuals are identically distributed.) These discoveries can be generally useful to screen variables in high-dimensional data analyses [5], to help prioritize follow-up studies [58], and in some cases even to make approximate causal inferences [31]. In fact, under some additional assumptions such as the absence of unmeasured confounders, a rejection of (1) can sometimes be rigorously interpreted as stating that X_j causes Y [61]. See Appendix A1.2 and Figure A1 for details on the connections to causal inference.

A1.2 The connection to causal inference

In addition to the basic setup of Section 2.1, imagine the variables X are randomized treatments that may depend on the observed covariates Z but are independent of anything else. Further, suppose the data distribution takes the form of a structural equation model [62, 63] in which Z and X may cause Y , but not the other way around, as visualized by the directed acyclic graph in Figure A1. It is easy to see this setup implicitly rules out confounding in a test of conditional independence between X and Z , because the treatments are conditionally independent of any other unmeasured covariate C . Therefore, to test whether X_j has a causal effect on Y , it suffices to test whether the conditional independence hypothesis in (1).

Proposition A1 (From [61]). *Let C be any unmeasured covariate. If $X \perp\!\!\!\perp C \mid Z$, a valid test of the conditional independence null hypothesis $\mathcal{H}_{0,j}$ in (1) is also a valid test of the stronger null hypothesis*

$$\mathcal{H}_{0,j}^* : Y \perp\!\!\!\perp X_j \mid Z, X_{-j}, C. \quad (\text{A7})$$

Proof. The proof is analogous to that of Proposition 1 in [61]. Suppose $\mathcal{H}_{0,j}^* : Y \perp\!\!\!\perp X_j \mid Z, X_{-j}, C$ is true. Then, it follows from $C \perp\!\!\!\perp X_j \mid Z, X_{-j}$ that $(Y, C) \perp\!\!\!\perp X_j \mid Z, X_{-j}$. Therefore, $\mathcal{H}_{0,j}$ must also be true. \square

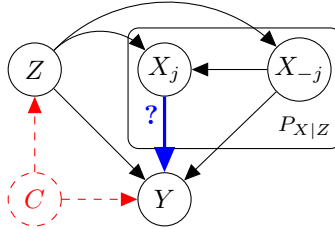


Figure A1: Graphical representation of a non-parametric causal model linking the treatment (X), the outcome (Y), the measured covariates (Z), and possibly also other unmeasured covariate (C). Our goal is to test whether a particular treatment X_j has any causal effect on the outcome. The joint distribution of $X \mid Z, C$ is assumed to be known and may depend only on Z , so that $X \perp\!\!\!\perp C \mid Z$.

A2 Method schematics

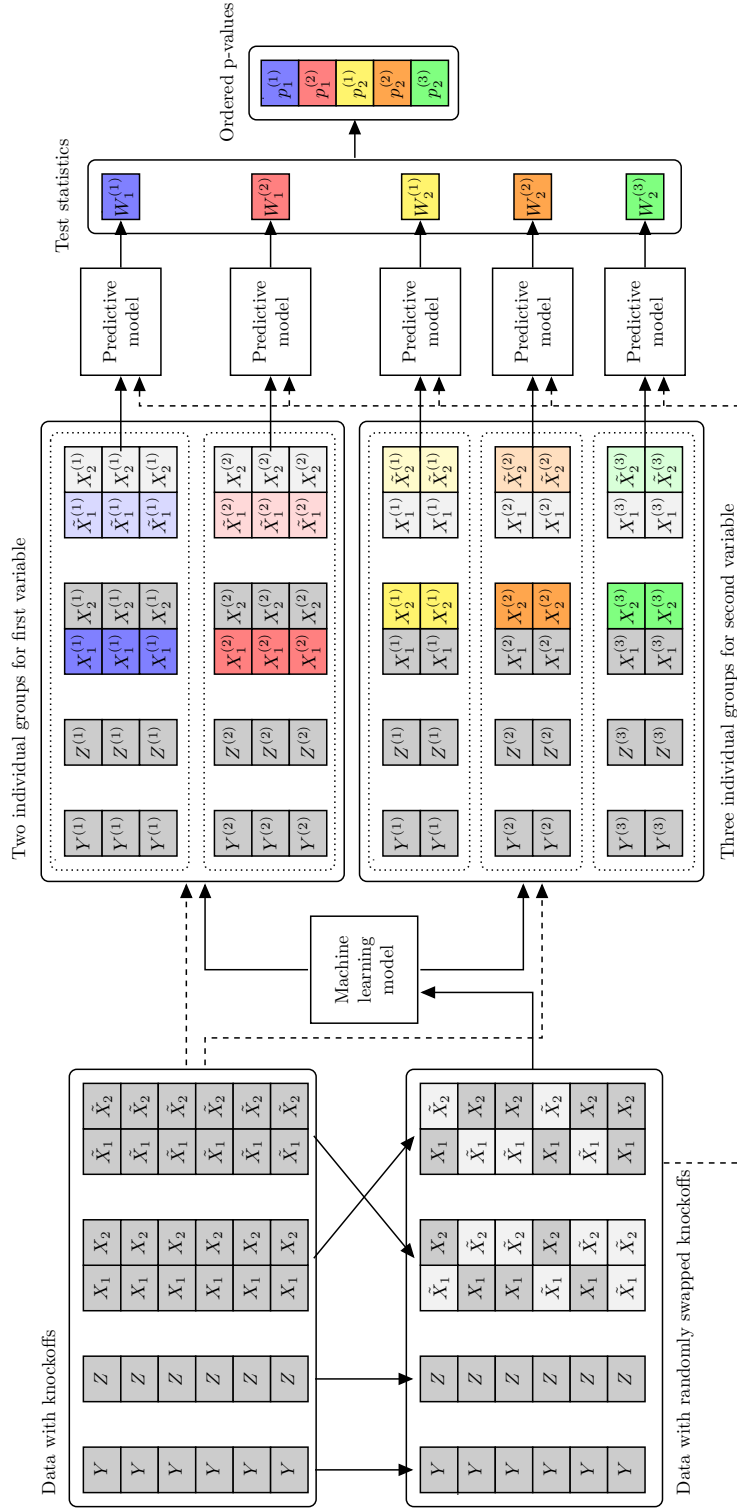


Figure A2: Schematic of the subgroup-selective knockoff filter for subgroup-specific conditional testing. The statistics for distinct combinations of variables and individual groups are shown in different colors.

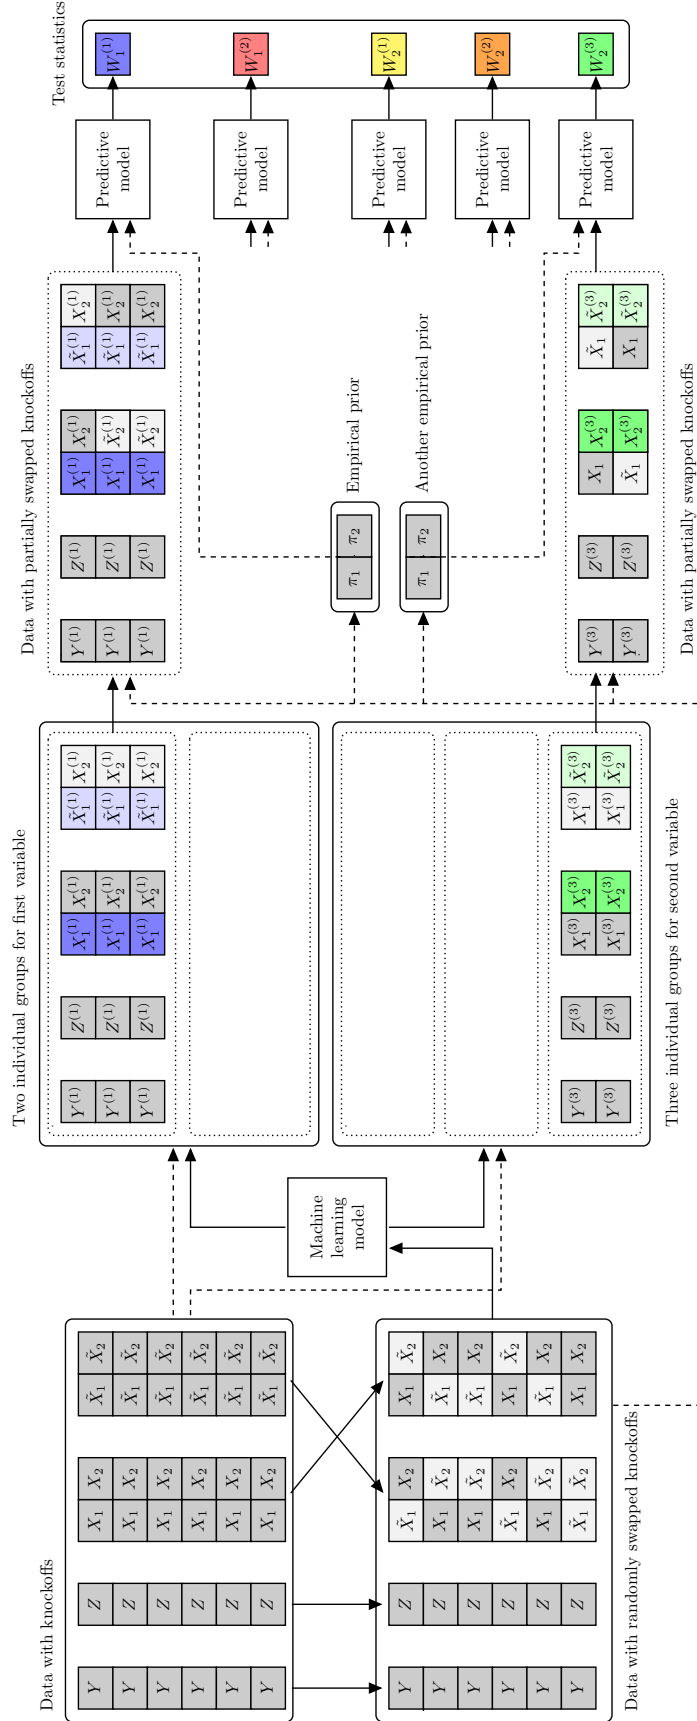


Figure A3: Schematic of the subgroup-selective knockoff filter for subgroup-specific conditional testing, focusing on the computation of the statistics for the first variable in subgroup 1 (blue) and for the second variable in subgroup 3 (green). The test statistics are based on a simple empirical prior for the two variables that is learnt from the full data set with randomly swapped knockoffs. Other details are as in Figure A2.

A3 The robust subgroup-selective knockoff filter

The subgroup-selective knockoff filter was developed to test subgroup-specific hypotheses (4) of the type

$$\mathcal{H}_{0,j}^{(g)} : Y^{(g)} \perp\!\!\!\perp X_j^{(g)} \mid X_{-j}^{(g)}, Z^{(g)},$$

for all variables j and subgroups g , conditional on a random partition $\hat{\psi}$ based on the data with randomly swapped knockoffs. Here, we present a further methodological extension that makes it possible to test the following *partial conjunction* [46, 47] hypotheses:

$$\mathcal{H}_{0,j}^{\text{pc},r} : \left| \left\{ g \in [G] : \mathcal{H}_{0,j}^{(g)} \text{ is true} \right\} \right| > G - r, \quad (\text{A8})$$

for any fixed $r \in \{1, \dots, G\}$. This is inspired by Li et al. [31]. Intuitively, the partial conjunction hypothesis states that the number of subgroups for variable j in which $\mathcal{H}_{0,j}^{(g)}$ (4) is false is strictly smaller than r . Therefore, rejecting this null suggests X_j is associated with Y in at least r distinct subgroups. If $r = 1$, the partial conjunction hypothesis reduces to the original population-wide hypothesis of Candès et al. [5], in (1); in that case, rejecting $\mathcal{H}_{0,j}^{\text{pc},r}$ suggests the existence of least one subgroup in which X_j is conditionally associated with Y . By contrast, rejecting the hypothesis in (A8) for $r = G$ requires evidence of a conditional association between X_j and Y across *all* subgroups.

One reason why it can be interesting to test $\mathcal{H}_{0,j}^{\text{pc},r}$ is that this hypothesis may point towards the discovery of variables whose association with Y is *robust* to changes in the covariate distribution [31]. Intuitively, testing $\mathcal{H}_{0,j}^{\text{pc},r}$ can be seen as a way for discovering conditional associations that do *not* involve interactions, thus taking a point of view complementary to that of the subgroup-selective knockoff filter. The p partial conjunction hypotheses defined in (A8) are relatively straightforward to test while provably controlling the false discovery rate, if one starts with a collection \mathbf{W} of subgroup-selective knockoff filter statistics satisfying the flip-sign property in Theorem 1: $\mathbf{W} \mid \hat{\psi} \stackrel{d}{=} \mathbf{W} \odot \mathbf{U} \mid \hat{\psi}$. In fact, it suffices to apply the multi-environment knockoff filter of Li et al. [31] to these statistics, as explained below.

For each $j \in \{1, \dots, p\}$, let n_j^- count the negative signs among the $W_j^{(g)}$ statistics for all subgroups g , and let n_j^0 be the number of zeros among them. For simplicity, assume $\hat{\psi}$ partitions the covariate space into G subgroups for each variable; this is without loss of generality, as it is otherwise sufficient to set the remaining undefined statistics to zero if $\hat{G}_j < G$ for some $j \in [p]$. Then, compute

$$p_j^{\text{pc},r} = \Psi \left(n_j^- - 1, (G - r + 1 - n_j^0) \vee 0, \frac{1}{2} \right) + U_j \cdot \Psi' \left(n_j^-, (G - r + 1 - n_j^0) \vee 0, \frac{1}{2} \right), \quad (\text{A9})$$

where $\Psi(\cdot, m, \pi)$ is the binomial cumulative distribution function, $\Psi'(\cdot, m, \pi)$ is the corresponding probability mass, and U_j is a uniform random variable on $[0, 1]$ independent of everything else. Further, define

$$|W_j^{\text{pc},r}| = \bar{w} \left(|W_j^{(1)}|, \dots, |W_j^{(G)}| \right), \quad (\text{A10})$$

for some symmetric function \bar{w} , such as that which multiplies the top r largest entries by absolute value:

$$w \left(|W_j^{(1)}|, \dots, |W_j^{(G)}| \right) = \prod_{g=1}^r |\bar{W}_j|_j^{(G-g+1)}.$$

Above, $|\bar{W}_j|_j^{(g)}$ are the order statistics for $\{|W_j^{(1)}|, \dots, |W_j^{(G)}|\}$. Then, the false discovery rate for (A8) can be controlled by applying the usual selective SeqStep+ sequential testing procedure of Barber and Candès [19], thanks to a relatively simple extension of Theorem 1 from Li et al. [31].

Theorem A1. *Selective SeqStep+ applied to the p -values (A9) ordered by (A10) controls the false discovery rate for (A8) if the statistics \mathbf{W} satisfy the flip-sign property in Theorem 1: $\mathbf{W} \mid \hat{\psi} \stackrel{d}{=} \mathbf{W} \odot \mathbf{U} \mid \hat{\psi}$.*

Proof. This result follows immediately from Theorem 1 in Li et al. [31] with a simple extension of Proposition 5 therein. The original statement of Proposition 5 in Li et al. [31] assumed the statistics \mathbf{W} to be computed using data collected from separate and a-priori fixed experimental settings (or *environments*), but it is easy to see that their proof only requires the flip-sign property established by our Theorem 1. \square

Alternatively, the false discovery rate for (A8) could be (approximately) controlled by applying to the p-values (A9) ordered by (A10) the accumulation test [64] instead of SeqStep+; this can be easily proved with an analogous extension of Theorem 2 from Li et al. [31].

A4 Additional results from numerical experiments

A4.1 Experiments in a “transfer learning” setting

Here, we investigate the performance of the robust subgroup-selective knockoff filter from Appendix A3 applied in a “transfer learning” setting inspired by [50]. The goal is complementary to that of the experiments in Section 4.1: we wish to detect variables with a robust association across subgroups with different covariate distributions. In other words, we want to find which variables maintain their association with Y within a future data set with covariate shift, in which the distribution of Z differs from the current one but the true model for $Y | X, Z$ is the same [50]. To simulate this scenario, we generate data from a model similar to that in Section 4.1, but now only half of the non-null variables interact with the covariates; i.e., $\beta_j^{(i)} = \bar{\beta}$ for half of the variables. Further, the data dimensions are increased to ensure sufficiently many discoveries can be made. Specifically, the number of variables is $p = 40$, while the number of covariates is $m = 160$, of which 120 are continuous. The covariate shift is imagined to be such that the 40 binary covariates are always equal to zero instead of following a symmetric Bernoulli distribution. Therefore, the goal is to identify the 20 variables whose association does not involve interactions, as the covariate shift would make the others irrelevant. These data are analyzed as in Section 4.1, with the only difference that the subgroup-selective knockoff filter is replaced by its robust version outlined in Appendix A3, which is designed to test partial conjunctions of the subgroup-specific conditional hypotheses (4).

Figure A4 reports on the results of 100 independent experiments, as a function of the sample size. The performance of the robust version of the subgroup-selective knockoff filter is quantified in terms of the empirical false discovery rate, the power, and the homogeneity of the reported findings, separately within the training population and under covariate shift. In the latter case, only the discoveries of variables with a direct conditional association with the outcome (i.e., not mediated by any interactions) are counted as true, while the others are considered false positives. Therefore, neither the naive benchmark nor the vanilla knockoff filter control this notion of false discovery rate under covariate shift because they tend to report all associated variables, including those that are non-null only thanks to interactions. By contrast, the robust subgroup-selective knockoff filter empirically controls the false discovery rate even under covariate shift. The data-splitting benchmark here is a modified version of our robust subgroup-selective knockoff filter in which half of the observations are used for partitioning the covariate space, without randomly swapping the knockoffs, and the remaining half are utilized for testing the selected hypotheses. The results demonstrate data-splitting is valid under covariate shift but it is not as powerful as our method, especially if the sample size is moderately large.

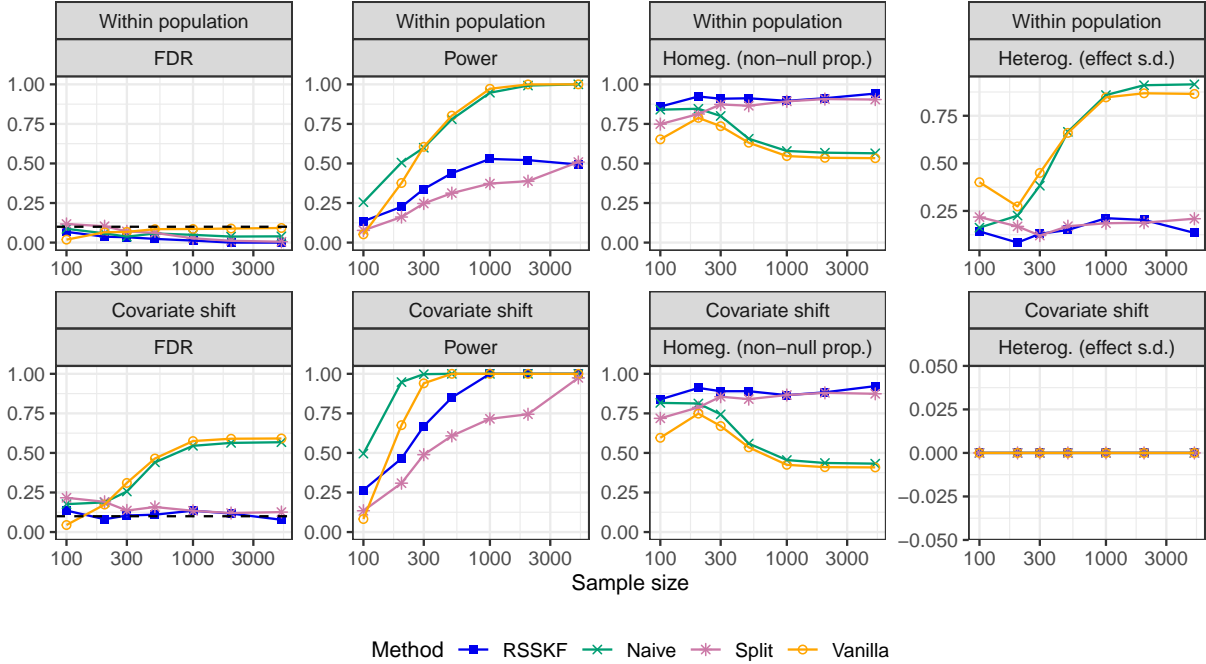


Figure A4: Performance of the robust subgroup-selective knockoff filter (RSSKF) and alternative benchmarks for subgroup-specific conditional testing under covariate shift, in numerical experiments with synthetic data. The findings summarized in the bottom panel are counted as true if and only if they report a variable whose conditional association with the outcome is robust to changes in the covariate distribution. Other details are as in Figure 1.

A4.2 Additional details about the blood donation data

Table A1: Randomized treatment-control assignments in a field experiment conducted by [51] to investigate the effectiveness of different incentives for blood donors.

Assignment	Size	Treatments (X_1, \dots, X_5)				
		Reminder	Indiv. reward	Friend	Group reward	Group gift
Control group	14k	0	0	0	0	0
Treated group 1	11k	1	0	0	0	0
Treated group 2	11k	1	1	0	0	0
Treated group 3	11k	1	0	1	0	0
Treated group 3	11k	1	1	1	0	0
Treated group 4	11k	1	1	1	0	0
Treated group 5	11k	0	1	1	1	0
Treated group 6	11k	1	1	1	0	1

Table A2: Goodness-of-fit-diagnostics for knockoff treatments, for the blood donation data. The means and standard deviations of the knockoffs approximately match those of the corresponding true treatments.

	Treatments					Knockoffs				
	X_1	X_2	X_3	X_4	X_5	\tilde{X}_1	\tilde{X}_2	\tilde{X}_3	\tilde{X}_4	\tilde{X}_5
Mean	0.825	0.412	0.550	0.138	0.138	0.822	0.411	0.549	0.139	0.140
Standard deviation	0.380	0.492	0.497	0.344	0.344	0.382	0.492	0.498	0.346	0.347

Table A3: Goodness-of-fit-diagnostics for knockoff treatments, for the blood donation data. The pairwise correlations between different treatments approximately match those of the corresponding knockoffs, as well as those between treatments and knockoffs. Smaller values of the diagonal correlation terms in the upper-right block of the correlation matrix (in red) tend to correspond to higher power.

	Treatments					Knockoffs				
	X_1	X_2	X_3	X_4	X_5	\tilde{X}_1	\tilde{X}_2	\tilde{X}_3	\tilde{X}_4	\tilde{X}_5
X_1	1	0.386	0.509	0.184	0.184	0.701	0.385	0.508	0.185	0.186
X_2	0.386	1	0.197	-0.335	0.477	0.389	0.607	0.197	-0.337	0.481
X_3	0.509	0.197	1	0.361	0.361	0.514	0.197	0.612	0.364	0.365
X_4	0.184	-0.335	0.361	1	-0.159	0.186	-0.334	0.362	0.622	-0.161
X_5	0.184	0.477	0.361	-0.159	1	0.186	0.478	0.362	-0.161	0.639
\tilde{X}_1						1	0.371	0.496	0.187	0.187
\tilde{X}_2						0.371	1	0.189	-0.319	0.464
\tilde{X}_3						0.496	0.189	1	0.347	0.347
\tilde{X}_4						0.187	-0.319	0.347	1	-0.162
\tilde{X}_5						0.187	0.464	0.347	-0.162	1

A4.3 Design of semi-synthetic experiments with blood donation data

The logistic model used to generate the imaginary donation outcomes considered in Section 4.2 is:

$$\text{logit}(\mathbb{P}[Y = 1 \mid X, Z]) = \varphi(X, Z),$$

with

$$\begin{aligned} \varphi(X, Z) = & -c + \\ & -b \cdot \mathbb{1}[\text{male}] + b \cdot \mathbb{1}[\text{married}] + b \cdot \mathbb{1}[\text{resident}] - b \cdot \mathbb{1}[\text{age} < 25] + \\ & -b \cdot \mathbb{1}[\text{student}] - b \cdot \mathbb{1}[\text{education} < 16] + b \cdot \mathbb{1}[\text{Rh}^-] - b \cdot \mathbb{1}[\text{blood type} \neq \text{O}] + \\ & + a \cdot X_1 [(1 - \mathbb{1}[\text{resident}]) + (1 - \mathbb{1}[\text{donation within 12 months}])] + \\ & + a \cdot X_2 [\mathbb{1}[\text{student}] + \mathbb{1}[\text{student}] \cdot \mathbb{1}[\text{male}]] + \\ & + a \cdot X_3 [(1 - \mathbb{1}[\text{male}]) + (1 - \mathbb{1}[\text{student}])] + \\ & + a \cdot X_4 [\mathbb{1}[\text{student}] + \mathbb{1}[\text{male}]] + \\ & + a \cdot X_5 [(1 - \mathbb{1}[\text{education} < 16])]. \end{aligned} \tag{A11}$$

Above, the coefficients are set as $a = 0.4$, $b = 0.2$, while c is such that half of the individuals in the data set on average receive a simulated $Y = 1$.

A4.4 Additional results from experiments with blood donation data

Table A4: Discoveries reported by the subgroup-selective knockoff filter in the analysis of semi-synthetic blood donation data, as in Figure 2. The sample size is 80,000 and the nominal false discovery rate is 10%.

Treatment	Subgroup	Truth	W	Homeg.	Heterog.	Samples
Reminder	Resident : 0 and Male : 1	Non-null	5.94	1.00	0.19	32,550
	Resident : 0 and Male : 0	Non-null	5.47	1.00	0.19	19,866
	Resident : 1 and Male : 1	Non-null	2.55	0.63	0.19	16,037
	Resident : 1 and Male : 0	Non-null	2.24	0.64	0.19	11,547
Individual reward	Rh-negative : 1 and Student : 1	Non-null	7.81	1.00	0.20	142
	Rh-negative : 0 and Student : 1	Non-null	7.12	1.00	0.20	26,043
Friends request	Married : 0 and Male : 0	Non-null	8.54	1.00	0.00	22,547
	Married : 0 and Male : 1	Non-null	4.37	1.00	0.00	34,240
	Married : 1 and Male : 0	Non-null	3.83	1.00	0.00	8,866
Group reward	Student : 1 and Married : 0	Non-null	4.43	1.00	0.20	25,915

A5 Proof of Theorem 1

Proof of Theorem 1. It suffices to prove the flip-sign property, $\mathbf{W} \mid \hat{\psi} \stackrel{d}{=} \mathbf{W} \odot \mathbf{U} \mid \hat{\psi}$, as the subsequent corollary about false discovery rate control follows from there immediately as in [5]. In fact, the sign-flip property is analogous to the standard one from [5], reviewed in Section 2.2, and it holds conditional on $\hat{\psi}$, which fully determines our hypotheses.

We begin by recalling that, by construction, both $\hat{\psi}$ and \hat{G} are determined by $[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Y}, \mathbf{Z}$, as well as (possibly) by some completely independent random noise (e.g., due to a cross-validation procedure). Therefore, conditional on $[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Y}, \mathbf{Z}$, we may treat $\hat{\psi}$ and \hat{G} as fixed. Similarly, the test statistics \mathbf{W} are determined by applying a (possibly randomized) function τ applied to a knockoff-augmented data set $[\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}, \mathbf{Z}$. Since the possible source of randomness in τ is also independent of everything else, we may ignore it without loss of generality, in order to simplify the notation as much as possible.

With this premise, we note that it follows from (6) that, for any fixed $\mathbf{U} \in \{\pm 1\}^{\hat{G} \times p}$,

$$\mathbf{W}([\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{U})}, \mathbf{Y}, \mathbf{Z}) = \mathbf{U} \odot \mathbf{W}([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}, \mathbf{Z}).$$

Further, it follows from Lemma A1 that, conditional on $([\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Y}, \mathbf{Z})$,

$$\mathbf{W}([\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{U})}, \mathbf{Y}, \mathbf{Z}) \stackrel{d}{=} \mathbf{W}([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}, \mathbf{Z}).$$

Combining the two equations above implies that, conditional on $([\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Y}, \mathbf{Z})$,

$$\mathbf{W} \odot \mathbf{U} \stackrel{d}{=} \mathbf{W}.$$

This concludes the proof because $\hat{\psi}$ is a function of $[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Y}, \mathbf{Z}$. □

Lemma A1. For any subset $\mathcal{S} \subseteq \{1, \dots, \hat{G}\}$ of true null hypotheses $\mathcal{H}_{0,i}^{(g)}$ in (4),

$$[\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathcal{S})} \stackrel{d}{=} [\mathbf{X}, \tilde{\mathbf{X}}]$$

conditional on $([\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathbf{V})}, \mathbf{Y}, \mathbf{Z})$, where \mathbf{V} is the random matrix used to define $\hat{\psi}$.

Proof of Lemma A1. We prove the following stronger equality in joint distribution:

$$\left([\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(S)}, [\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(V)}, \mathbf{Y}, \mathbf{Z}\right) \stackrel{d}{=} \left([\mathbf{X}, \tilde{\mathbf{X}}], [\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(V)}, \mathbf{Y}, \mathbf{Z}\right).$$

By independence, it suffices to establish the above result for a single row of the data matrices:

$$\left((X, \tilde{X})_{\text{swap}(s)}, Z, Y, (X, \tilde{X})_{\text{swap}(v)}\right) \stackrel{d}{=} \left((X, \tilde{X}), Z, Y, (X, \tilde{X})_{\text{swap}(v)}\right).$$

By construction of the knockoffs \tilde{X} , we know that

$$(X, \tilde{X})_{\text{swap}(s)} | Z \stackrel{d}{=} (X, \tilde{X}) | Z. \tag{A12}$$

Further, for any fixed swap s and random swap v it holds that $v \odot s \stackrel{d}{=} v$, which implies

$$(X, \tilde{X})_{\text{swap}(v)} | (X, \tilde{X})_{\text{swap}(s)}, Z \stackrel{d}{=} (X, \tilde{X})_{\text{swap}(v)} | (X, \tilde{X}), Z. \tag{A13}$$

Combining (A12) and (A13) yields:

$$\left((X, \tilde{X})_{\text{swap}(s)}, (X, \tilde{X})_{\text{swap}(v)}, Z\right) \stackrel{d}{=} \left((X, \tilde{X}), (X, \tilde{X})_{\text{swap}(v)}, Z\right). \tag{A14}$$

Now, recall that $Y \perp\!\!\!\perp \tilde{X} | Z, X$. Further, $Y \perp\!\!\!\perp X_s | Z, X_{-s}$ for all X_s swapped by s because the latter only involves null variables. Therefore, by the same argument as in the proof of Lemma 3.2 in [5],

$$Y | Z, (X, \tilde{X})_{\text{swap}(s)}, (X, \tilde{X})_{\text{swap}(v)} \stackrel{d}{=} Y | Z, (X, \tilde{X}), (X, \tilde{X})_{\text{swap}(v)}. \tag{A15}$$

Finally, combining (A14) with (A15) gives the desired result. \square