

Robust inference with knockoffs

Rina Foygel Barber* Emmanuel J. Candès† Richard J. Samworth‡

January, 2018

Abstract

We consider the variable selection problem, which seeks to identify important variables influencing a response Y out of many candidate features X_1, \dots, X_p . We wish to do so while offering finite-sample guarantees about the fraction of false positives—selected variables X_j that in fact have no effect on Y after the other features are known. When the number of features p is large (perhaps even larger than the sample size n), and we have no prior knowledge regarding the type of dependence between Y and X , the model- X knockoffs framework nonetheless allows us to select a model with a guaranteed bound on the false discovery rate, as long as the distribution of the feature vector $X = (X_1, \dots, X_p)$ is exactly known. This model selection procedure operates by constructing “knockoff copies” of each of the p features, which are then used as a control group to ensure that the model selection algorithm is not choosing too many irrelevant features. In this work, we study the practical setting where the distribution of X could only be estimated, rather than known exactly, and the knockoff copies of the X_j ’s are therefore constructed somewhat incorrectly. Our results, which are free of any modeling assumption whatsoever, show that the resulting model selection procedure incurs an inflation of the false discovery rate that is proportional to our errors in estimating the distribution of each feature X_j conditional on the remaining features $\{X_k : k \neq j\}$. The model- X knockoffs framework is therefore robust to errors in the underlying assumptions on the distribution of X , making it an effective method for many practical applications, such as genome-wide association studies, where the underlying distribution on the features X_1, \dots, X_p is estimated accurately but not known exactly.

1 Introduction

Our methods of data acquisition are such that we often obtain information on an exhaustive collection of possible explanatory variables. We know a priori that a large proportion of these are irrelevant for our purposes, but in the effort to cover all bases, we gather data on all what we can measure and rely on subsequent analysis to identify the relevant variables. For instance, to achieve a better understanding of biological processes behind a disease, we may evaluate variation across the entire DNA sequence and collect single nucleotide polymorphism (SNP) information, or quantify the expression level of all genes, or consider a large panel of exposures, and so on. We then expect the statistician or the scientist to sort through all these and select those important variables that truly influence a response of interest. For example, we would like the statistician to tell us which of the many genetic variations affect the risk of a specific disease, or which of the many gene expression profiles help determine the severity of a tumor.

This paper is about this variable selection problem. We consider situations where we have observations on a response Y and a large collection of variables X_1, \dots, X_p . With the goal of identifying the important variables, we want to recover the smallest set $\mathcal{S} \subseteq \{1, \dots, p\}$ such that, conditionally on $\{X_j\}_{j \in \mathcal{S}}$, the response Y is independent of all the remaining variables $\{X_j\}_{j \notin \mathcal{S}}$. In the literature on graphical models, the set \mathcal{S} would be called the *Markov blanket* of Y . Effectively, this means that the explanatory variables X_1, \dots, X_p provide

*Department of Statistics, University of Chicago

†Departments of Mathematics and of Statistics, Stanford University

‡Statistical Laboratory, University of Cambridge

information about the outcome Y only through the subset $\{X_j\}_{j \in \mathcal{S}}$. To ensure reproducibility, we are interested in methods that result in the estimation of a set $\widehat{\mathcal{S}}$ with false discovery rate (FDR) control [3], in the sense that

$$\text{FDR} = \mathbb{E} \left[\frac{\#\{j : j \in \widehat{\mathcal{S}} \setminus \mathcal{S}\}}{\#\{j : j \in \widehat{\mathcal{S}}\}} \right] \leq q,$$

i.e. a bound on the expected proportion of our discoveries $\widehat{\mathcal{S}}$ which are *not* in the smallest explanatory set \mathcal{S} .¹ (Here q is some predetermined target error rate, e.g. $q = 0.1$.)

In truth, there are not many variable selection methods that would control the FDR with finite-sample guarantees, especially when the number p of variables far exceeds the sample size n . That said, one solution is provided by the recent model-X knockoffs approach of Candès et al. [4], which is a new read on the earlier knockoff filter of Barber and Candès [1], see also [2]. One singular aspect of the method of model-X knockoffs is that it makes assumptions that are substantially different from those commonly encountered in the statistical literature. Most of the model selection literature relies on a specification of the model that links together the response and the covariates, making assumptions on $P_{Y|X}$, the distribution of Y conditional on X —for instance, assuming that the form of this distribution follows a generalized linear model or some other parametrized model. In contrast, model-X knockoffs makes no assumption whatsoever on the relation between the response Y and the variables $X = (X_1, \dots, X_p)$; in other words, the distribution $P_{Y|X}$ of Y conditional on X is “model free”. The price of this generality is that we need to be able to specify the distribution of the feature variables $X = (X_1, \dots, X_p)$, which we denote by P_X .

As argued in [4] and [8], this “shift” of the burden of knowledge is interesting because we must recall that the object of inference is on how Y relates to X , that is, on $P_{Y|X}$. It is, therefore, a strong premise to posit the form of this relationship $P_{Y|X}$ a priori—and indeed, there are many applications in which we objectively do not have any understanding of how Y depends on X . Further, the shift is also appropriate whenever we know much more about the distribution of X than on the conditional distribution of $Y | X$. For instance, it is easy to imagine applications in which we have many unlabeled samples—samples of X —whereas it may be much harder to acquire labeled data or samples with a given value of the response Y . A typical example is offered by genetic studies, where we now have available hundreds of thousands or even millions of genotypes across many different populations. At the same time, it may be difficult to recruit patients with a given phenotype (the response variable Y), and therefore, we have substantially more data with which to estimate P_X than $P_{Y|X}$.

The ease with which we can gather information about X does not imply that we know the distribution P_X exactly, but we often do have substantial information about this distribution. Returning to our genetic example, it has been shown that the joint distribution of SNPs may be accurately modeled by hidden Markov models (see [14, 16, 12, 10] for some early formulations), and there certainly is an abundance of genotype data to estimate the various model parameters; compare for instance the success of a variety of methods for genotype imputation [11, 7] based on such models. The purpose of this paper, then, is precisely to investigate common situations of this kind, namely, what happens when we run model-X knockoffs and only assume *some* knowledge of the distribution of X rather than exact knowledge.

Our contribution is a considerable extension of the original work on model-X knockoffs [4], which assumed a perfect knowledge of the distribution of X to achieve FDR control. Here, we develop a new theory, which quantifies very precisely the inflation in FDR when running the knockoff filter with estimates of the distribution of X in place of the true distribution P_X . We develop non-asymptotic bounds which show that the possible FDR inflation is well-behaved whenever the estimated distribution is reasonably close to the truth. These bounds are general and apply to all possible statistics that the researcher may want to use to tease out the signal from the noise. We also develop converse results for some settings, showing that our bounds are fairly sharp in that it is impossible to obtain tighter FDR control bounds in full generality. Thus, our theory offers finite-sample guarantees that hold for any algorithm that the analyst decides to employ, assuming no knowledge of the form of the relationship between Y and X and only an estimate of the distribution of X itself. On the other hand, since our bounds are worst-case, they may be pessimistic in the sense that the realized FDR in any practical situation may be much lower than that achieved in the worst possible case.

¹As is standard in the FDR literature, in this expected value we treat $0/0$ as 0 , to incur no penalty in the event that no variables are selected, i.e. when $\widehat{\mathcal{S}} = \emptyset$.

Underlying our novel model-X knockoffs theory is a completely new mathematical analysis and understanding of the knockoffs inferential machine. The technical innovation here is essentially twofold. First, with only partial knowledge of the distribution of X , we can no longer achieve a perfect exchangeability between the test statistics for the null variables and for their knockoffs. Hence, we need tools that can deal with only a form of approximate exchangeability. Second, our methods to prove FDR control no longer rely on martingale arguments, and rather, involve leave-one-out type of arguments. These new arguments are likely to have applications far outside the scope of the present paper.

2 Robust inference with knockoffs

To begin with, imagine we have data consisting of n i.i.d. draws from a joint distribution on (X, Y) , where $X = (X_1, \dots, X_p) \in \mathbb{R}^p$ is the feature vector while $Y \in \mathbb{R}$ is the response variable. We will gather the n observed data points into a matrix $\mathbf{X} \in \mathbb{R}^{n \times p}$ and vector $\mathbf{Y} \in \mathbb{R}^n$ —that is, the pairs $(\mathbf{X}_{i,*}, \mathbf{Y}_i)$ are i.i.d. copies of the pair (X, Y) . The joint distribution of (X, Y) is unknown—specifically, we do not assume any information about the conditional distribution of Y given X as discussed above. We work under the assumption that P_X , the marginal distribution of X , is known only approximately.

Since the Markov blanket of Y may be ill-defined (e.g. if two features are identical then the choice of the minimal set \mathcal{S} may not be unique), we follow [4] and define X_j to be a null variable if $X_j \perp\!\!\!\perp Y \mid X_{-j}$, that is if X_j is independent from the response Y conditionally on all the other variables. (We use the terms “features” and “variables” interchangeably.) Under very mild identifiability conditions, the set of non-nulls is nothing other than the Markov blanket of Y . Writing \mathcal{H}_0 to denote the set of indices corresponding to null variables, we can then reformulate the error we would like to control as $\mathbb{E} \left[\frac{|\hat{\mathcal{S}} \cap \mathcal{H}_0|}{|\hat{\mathcal{S}}|} \right] \leq q$.

2.1 Exact model-X knockoffs

Consider first an ideal setting where the distribution P_X is known. The model-X knockoffs method [4] is defined by constructing knockoff features satisfying the following conditions. \tilde{X} is drawn conditional on the feature vector X without looking at the response Y (i.e. $\tilde{X} \perp\!\!\!\perp Y \mid X$), such that the joint distribution of (X, \tilde{X}) satisfies a pairwise exchangeability condition,

$$(X, \tilde{X})_{\text{swap}(\mathcal{A})} \stackrel{d}{=} (X, \tilde{X}) \quad (1)$$

for any subset $\mathcal{A} \subseteq \{1, \dots, p\}$, where $\stackrel{d}{=}$ denotes equality in distribution. (In fact, to achieve FDR control, this condition only needs to hold for subsets $\mathcal{A} \subseteq \mathcal{H}_0$ containing only null variables.) Above, the family $(X, \tilde{X})_{\text{swap}(\mathcal{A})}$ is obtained from (X, \tilde{X}) by swapping the entries X_j and \tilde{X}_j for each $j \in \mathcal{A}$; for example, with $p = 3$ and $\mathcal{A} = \{2, 3\}$,

$$(X_1, X_2, X_3, \tilde{X}_1, \tilde{X}_2, \tilde{X}_3)_{\text{swap}(\{2,3\})} = (X_1, \tilde{X}_2, \tilde{X}_3, \tilde{X}_1, X_2, X_3).$$

As a consequence of the pairwise exchangeability property (1), we see that the null knockoff variables $\{\tilde{X}_j\}_{j \in \mathcal{H}_0}$ are distributed in exactly the same way as the original nulls $\{X_j\}_{j \in \mathcal{H}_0}$ but some dependence is preserved: for instance, for any pair $j \neq k$ where k is a null, we have that $(X_j, \tilde{X}_k) \stackrel{d}{=} (X_j, X_k)$.

Given knowledge of the true distribution P_X of the features X , our first step to implement the method of model-X knockoffs is to construct a distribution for drawing \tilde{X} conditional on X such that the pairwise exchangeability property (1) holds for all subsets of features \mathcal{A} . We can think of this mechanism as constructing some probability distribution $P_{\tilde{X}|X}(\cdot|x)$, which is a conditional distribution of \tilde{X} given $X = x$, chosen so that the resulting joint distribution of (X, \tilde{X}) , which is equal to

$$P_X(x) P_{\tilde{X}|X}(\tilde{x}|x),$$

is symmetric in the pairs (x_j, \tilde{x}_j) , and thus will satisfy the exchangeability property (1). Now, when working with data (\mathbf{X}, \mathbf{Y}) , we will treat each data point $(\mathbf{X}_{i,*}, \mathbf{Y}_i)$ independently. Specifically, after observing the data

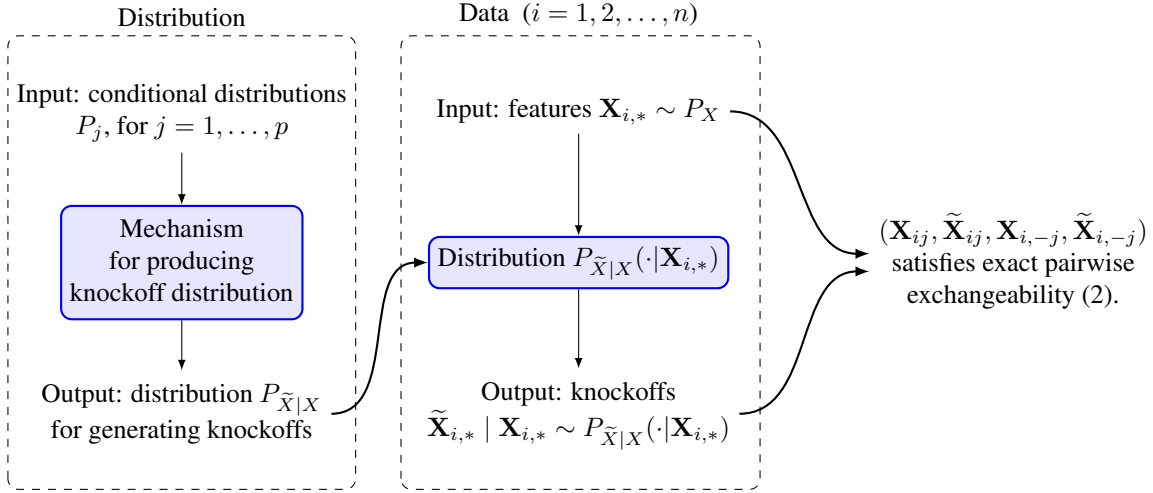


Figure 1: Schematic representation of the exact model-X knockoffs construction.

$(\mathbf{X}, \mathbf{Y}) \in \mathbb{R}^{n \times p} \times \mathbb{R}^n$, the rows $\tilde{\mathbf{X}}_{i,*}$ of the knockoff matrix are drawn from $P_{\tilde{X}|X}(\cdot | \mathbf{X}_{i,*})$, independently for each i and also independently of \mathbf{Y} . Figure 1 shows a schematic representation of the exact model-X knockoffs construction.

It is important to point out that mechanisms for producing the pairwise exchangeability property (1) do exist and can be very concrete: Candès et al. [4] develop a general abstract mechanism termed the Sequential Conditional Independent Pairs (SCIP) “algorithm”, which always produces exchangeable knockoff copies and can be applied to any distribution P_X . There are also fast algorithms for the case where X follows either a Markov or a hidden Markov model [13], and simple algorithms for sampling knockoff copies of Gaussian features [4]. Looking ahead, all of these algorithms can be used in the case where P_X is known only approximately, where the exchangeability property (1) will be required to hold only with reference to the *estimated* distribution of X , discussed in Section 2.2 below.

For assessing a model selection algorithm, the knockoff feature vectors $\tilde{\mathbf{X}}_j$ can be used as a “negative control”—a control group for testing the algorithm’s ability to screen out false positives, since $\tilde{\mathbf{X}}_j$ is known to have no real effect on \mathbf{Y} . Although details are given in Section 2.3, it is helpful to build some intuition already at this stage. Imagine for simplicity that we wish to assess the importance of a variable by measuring the strength of the marginal correlation with the response, i.e. we compute $Z_j = |\mathbf{X}_j^\top \mathbf{Y}|$. Then we can compare Z_j with $\tilde{Z}_j = |\tilde{\mathbf{X}}_j^\top \mathbf{Y}|$, the marginal correlation for the corresponding knockoff variable. The crucial point is that the pairwise exchangeability property (1) implies that if j is null (recall that this means that X_j and Y are conditionally independent), then

$$(Z_j, \tilde{Z}_j) \stackrel{d}{=} (\tilde{Z}_j, Z_j).$$

This holds without any assumptions on the form of the relationship $P_{Y|X}$ between Y and X [4]. In particular, this means that the test statistic $W_j = Z_j - \tilde{Z}_j$ is equally likely to be positive or negative. Thus to reject the null, we would need to observe a large positive value of W_j . As we will see in Section 2.3, this way of reasoning extends to any choice of statistic Z_j ; whatever statistic we choose, knockoff variables obeying (1) offer corresponding values of the statistic which can be used as “negative controls” for calibration purposes.

Throughout this paper, we will pay close attention to the distribution we obtain when swapping only one variable and its knockoff (and do not swap any of the other variables). In this context, we can reformulate the broad exchangeability condition (1) in terms of single variable swaps.

Proposition 1 (Candès et al. [4, Prop. 3.5]). *The pairwise exchangeability property (1) holds for a subset*

$\mathcal{A} \subseteq \{1, \dots, p\}$ if and only if

$$(X_j, \tilde{X}_j, X_{-j}, \tilde{X}_{-j}) \stackrel{d}{=} (\tilde{X}_j, X_j, X_{-j}, \tilde{X}_{-j}) \quad (2)$$

holds for all $j \in \mathcal{A}$.

In other words, we can restrict our attention to the question of whether a single given feature X_j and its knockoff \tilde{X}_j are exchangeable with each other (in the joint distribution that also includes X_{-j} and \tilde{X}_{-j}).

2.2 Approximate model-X knockoffs and pairwise exchangeability

Now we will work towards constructing a version of this method when the true distribution P_X of the feature vector X is not known exactly. Here, we need to relax the pairwise exchangeability assumption, since choosing a useful mechanism $P_{\tilde{X}|X}$ that satisfies this condition would generally require a very detailed knowledge of P_X , which is typically not available. This section builds towards Definition 1 in two steps.

2.2.1 Exchangeability with respect to an input distribution Q_X

We are provided with data \mathbf{X} and conditional distributions $Q_j(\cdot|x_{-j})$ for each j . We assume that the Q_j 's are fixed, i.e. do not depend on the data set (\mathbf{X}, \mathbf{Y}) . As a warm-up, assume first that these conditionals are mutually compatible in the sense that there is a joint distribution Q_X over \mathbb{R}^p that matches these p estimated conditionals—we will relax this assumption very soon. Then as shown in Figure 2, we repeat the construction from Figure 1, only with the Q_j 's as inputs. In words, the algorithm constructs knockoffs, which are samples from $P_{\tilde{X}|X}$, a conditional distribution whose construction is based on the conditionals Q_j or, equivalently, the joint distribution Q_X . In place of requiring that pairwise exchangeability of the features X_j and their knockoffs \tilde{X}_j holds relative to the true distribution P_X as in (1) and (2), we instead require that the knockoff construction mechanism satisfy pairwise exchangeability conditions relative to the joint distribution Q_X that it receives as input:

$$\begin{aligned} &\text{If } (X, \tilde{X}) \text{ is drawn as } X \sim Q_X \text{ and } \tilde{X} | X \sim P_{\tilde{X}|X}(\cdot|X), \text{ then} \\ &(X, \tilde{X})_{\text{swap}(\mathcal{A})} \stackrel{d}{=} (X, \tilde{X}), \quad \text{for any subset } \mathcal{A} \subseteq \{1, \dots, p\}. \end{aligned} \quad (3)$$

When only estimated compatible conditionals are available, original and knockoff features are required to be exchangeable with respect to the distribution Q_X , which is provided as input (but not with respect to the true distribution of X , which is unknown). To rephrase, if the distribution of X were in fact equal to Q_X , then we would have exchangeability.

2.2.2 Exchangeability with respect to potentially incompatible conditionals Q_j

We wish to provide an extension of (3) to cover the case where the conditionals may not be compatible; that is, when a joint distribution with the Q_j 's as conditionals may not exist. To understand why this is of interest, imagine we have unlabeled data that we can use to estimate the distribution of X . Then we may construct Q_j by regressing the j th feature X_j onto the $p - 1$ remaining features X_{-j} . For instance, we may use a regression technique promoting sparsity or some other assumed structure. In such a case, it is easy to imagine that such a strategy may produce incompatible conditionals. It is, therefore, important to develop a framework adapted to this setting. To address this, we shall work throughout the paper with the following definition:

Definition 1. $P_{\tilde{X}|X}$ is pairwise exchangeable with respect to Q_j if it satisfies the following property:

$$\begin{aligned} &\text{For any distribution } D^{(j)} \text{ on } \mathbb{R}^p \text{ with } j\text{th conditional } Q_j, \text{ if } (X, \tilde{X}) \text{ is drawn as } X \sim D^{(j)} \\ &\text{and } \tilde{X} | X \sim P_{\tilde{X}|X}(\cdot|X), \text{ then } (X_j, \tilde{X}_j, X_{-j}, \tilde{X}_{-j}) \stackrel{d}{=} (\tilde{X}_j, X_j, X_{-j}, \tilde{X}_{-j}). \end{aligned} \quad (4)$$

Above, $D^{(j)}$ is the product of an arbitrary marginal distribution for X_{-j} and of the conditional Q_j .

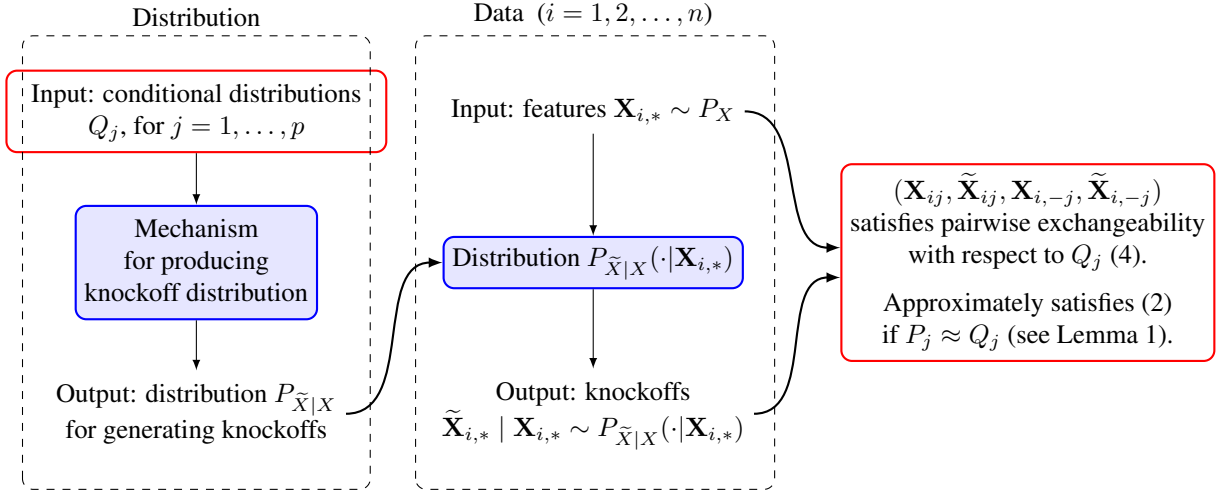


Figure 2: Schematic representation of the approximate model-X knockoffs construction. The two differences relative to Figure 1 are circled in red.

In words, with estimated conditionals Q_j , we choose $P_{\tilde{X}|X}$ to satisfy pairwise exchangeability with respect to these Q_j 's, for every j . (As before, we remark that this only needs to hold for $j \in \mathcal{H}_0$ to ensure FDR control, but since in practice we do not know which features are null, we require (4) to hold for every j .)

To see why this is an extension of (3), note that if the Q_j 's are mutually compatible, then any algorithm operating such that (4) holds for each j , obeys (3) as well. In particular, in the exact model-X framework setting where (1) is satisfied, Proposition 1 implies that the distribution $P_{\tilde{X}|X}$ is pairwise exchangeable with respect to P_j , i.e. with respect to the true conditional distribution of X_j , for each j .

As far as generating knockoff copies obeying (4), imagine we have a mechanism, which, when we input the distribution P_X of X , will produce exchangeable knockoffs obeying (1). Supposing that the estimated conditionals Q_j are all compatible with some joint distribution Q_X on X , then if we instead provide Q_X as input to this mechanism, it instead produces knockoff copies with the desired property. Hence, if the Q_j 's are mutually compatible, then all the mechanisms producing valid knockoffs under exact knowledge of P_X —we mentioned a few in the previous section—can be readily used for our purposes. Later in Section 3.3, we will also give an example of a mechanism producing valid knockoffs satisfying (4) under incompatible Q_j 's.

From this point on, we will assume without comment that for each j , either X_j and \tilde{X}_j are both discrete variables or are both continuous variables, and abusing notation, in these two settings we will use $P_j(\cdot | x_{-j})$ and $Q_j(\cdot | x_{-j})$ to denote the conditional probability mass function or conditional density, respectively, for the true and estimated conditional distribution of X_j given $X_{-j} = x_{-j}$. Furthermore, we assume that $P_j(\cdot | x_{-j})$ and $Q_j(\cdot | x_{-j})$ are supported on the same (discrete or continuous) set for any x_{-j} . Our theory can be generalized to the setting of mixed distributions and/or varying supports, but for clarity of the results we do not present these generalizations here.

The construction of the knockoff features as in Figure 2 yields the following approximate pairwise exchangeability result.

Lemma 1. *Fix any feature index j such that pairwise exchangeability with respect to Q_j (4) is satisfied. If*

X_j, \tilde{X}_j are discrete, then for any² a, b ,

$$\frac{\mathbb{P}\left\{X_j = a, \tilde{X}_j = b \mid X_{-j}, \tilde{X}_{-j}\right\}}{\mathbb{P}\left\{X_j = b, \tilde{X}_j = a \mid X_{-j}, \tilde{X}_{-j}\right\}} = \frac{P_j(a|X_{-j})Q_j(b|X_{-j})}{Q_j(a|X_{-j})P_j(b|X_{-j})}.$$

Furthermore, if index j corresponds to a null feature (i.e. $X_j \perp\!\!\!\perp Y \mid X_{-j}$) and we additionally assume that $\tilde{X} \mid X$ is drawn from $P_{\tilde{X}|X}$ independently of Y , then the same result holds when we also condition on Y :

$$\frac{\mathbb{P}\left\{X_j = a, \tilde{X}_j = b \mid X_{-j}, \tilde{X}_{-j}, Y\right\}}{\mathbb{P}\left\{X_j = b, \tilde{X}_j = a \mid X_{-j}, \tilde{X}_{-j}, Y\right\}} = \frac{P_j(a|X_{-j})Q_j(b|X_{-j})}{Q_j(a|X_{-j})P_j(b|X_{-j})}.$$

The conclusion in the continuous case is identical except with ratios of probabilities replaced with ratios of densities.

This lemma gives a useful formula for computing the ratio between the likelihoods of the two configurations $(X_j, \tilde{X}_j) = (a, b)$ and $(X_j, \tilde{X}_j) = (b, a)$ (after conditioning on the remaining data). It is important to observe that if we are working in the exact model-X framework, where $Q_j = P_j$, in this case the lemma yields

$$\frac{\mathbb{P}\left\{X_j = a, \tilde{X}_j = b \mid X_{-j}, \tilde{X}_{-j}, Y\right\}}{\mathbb{P}\left\{X_j = b, \tilde{X}_j = a \mid X_{-j}, \tilde{X}_{-j}, Y\right\}} = 1$$

for each null j ; that is, the two configurations are equally likely. This result for the exact model-X setting is proved in Candès et al. [4, Lemma 3.2] and is critical for establishing FDR control properties.

To see why this is so, let us return to our marginal correlation example from Section 2.1. In the exact model-X setting, we see from the above display that for a null feature, the variables X_j and \tilde{X}_j are exchangeable conditionally on everything else, namely, on X_{-j}, \tilde{X}_{-j} , and Y . Likewise, \mathbf{X}_j and $\tilde{\mathbf{X}}_j$ (the two columns of the data matrix) are also conditionally exchangeable. This implies that $|\mathbf{X}_j^\top \mathbf{Y}|$ and $|\tilde{\mathbf{X}}_j^\top \mathbf{Y}|$ are conditionally exchangeable as well and, therefore, marginally exchangeable as we claimed. When we use estimates Q_j rather than the true conditionals P_j , however, this property is no longer true, since Lemma 1 shows that the ratio is no longer equal to 1 in general. We can no longer use the knockoff statistics as exact negative controls; only as approximate controls. This is where the major difficulty comes in: if a knockoff statistic is only approximately distributed like its corresponding null, what is the potential inflation of the type-I error that this could cause?

2.3 The knockoff filter

After constructing the variables $\tilde{\mathbf{X}}_j$, we apply the knockoff filter to select important variables. We here quickly rehearse the main ingredients of this filter and refer the reader to [1] and [4] for additional details; our exposition borrows from [2]. Suppose that for each variable \mathbf{X}_j (resp. each knockoff variable $\tilde{\mathbf{X}}_j$), we compute a score statistic Z_j (resp. \tilde{Z}_j),

$$(Z_1, \dots, Z_p, \tilde{Z}_1, \dots, \tilde{Z}_p) = z([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}),$$

with the idea that Z_j (resp. \tilde{Z}_j) measures the importance of X_j (resp. \tilde{X}_j) in explaining Y . Assume that the scores are “knockoff agnostic” in the sense that switching a variable with its knockoff simply switches the components of Z in the same way. This means that

$$z([\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathcal{A})}, \mathbf{Y}) = z([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y})_{\text{swap}(\mathcal{A})}, \quad (5)$$

²Formally, this result holds only for a, b lying in the support of $P_j(\cdot|X_{-j})$, which is assumed to be equal to the support of $Q_j(\cdot|X_{-j})$, as otherwise the ratio is 0/0; we ignore this possibility here and throughout the paper since these results will be applied only in settings where a, b do lie in this support.

i.e. swapping X_1 and \tilde{X}_1 before calculating Z has the effect of swapping Z_1 and \tilde{Z}_1 , and similarly swapping X_2 and \tilde{X}_2 swaps Z_2 and \tilde{Z}_2 , and so on. Here, we emphasize that Z_j may be an arbitrarily complicated statistic. For instance, it can be defined as the absolute value of a lasso coefficient, or some random forest feature importance statistic; or, we may fit both a lasso model and a random forest, and choose whichever one has the lowest cross-validated error.

These scores are then combined in a single importance statistic for the variable X_j as

$$W_j = f_j(Z_j, \tilde{Z}_j) =: w_j([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}),$$

where f_j is any anti-symmetric function, meaning that $f_j(v, u) = -f_j(u, v)$. As an example, we may have $W_j = Z_j - \tilde{Z}_j$, where the Z_j 's and \tilde{Z}_j 's are the magnitudes of regression coefficients estimated by the lasso at a value of the regularization parameter given by cross-validation, say. Again, any choice of anti-symmetric function f_j and score statistic Z_j , no matter how complicated, is allowed. By definition, the statistics W_j obey the *flip-sign property*, which says that swapping the j th variable with its knockoff has the effect of changing the sign of W_j (since, by (5) above, if we swap feature vectors \mathbf{X}_j and $\tilde{\mathbf{X}}_j$ then Z_j and \tilde{Z}_j get swapped):

$$w_j([\mathbf{X}, \tilde{\mathbf{X}}]_{\text{swap}(\mathcal{A})}, \mathbf{Y}) = \begin{cases} w_j([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}), & j \notin \mathcal{A}, \\ -w_j([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y}), & j \in \mathcal{A}. \end{cases} \quad (6)$$

The W_j 's are the statistics that the knockoff filter will use. The idea is that large positive values of W_j provide evidence against the hypothesis that the distribution of Y is conditionally independent of X_j , while in contrast, if $j \in \mathcal{H}_0$, then W_j has a symmetric distribution and, therefore, is equally likely to take on positive or negative values.

In fact, it is equally valid for us to define $W_j = w_j([\mathbf{X}, \tilde{\mathbf{X}}], \mathbf{Y})$ for any function w_j satisfying the flip-sign property (5), without passing through the intermediate stage of defining Z_j 's and \tilde{Z}_j 's, and from this point on we do not refer to the feature importance scores Z_j, \tilde{Z}_j in our theoretical results. However, for better understanding of the intuition behind the method, we should continue to think of W_j as comparing the apparent importance of the feature \mathbf{X}_j versus its knockoff $\tilde{\mathbf{X}}_j$ for modeling the response \mathbf{Y} .

Now that we have test statistics for each variable, we need a selection rule. For the knockoff filter, we choose a threshold $T_0 > 0$ by setting³

$$T_0 = \min \left\{ t > 0 : \frac{\#\{j : W_j \leq -t\}}{\#\{j : W_j \geq t\}} \leq q \right\}, \quad (7)$$

where q is the target FDR level. The output of the procedure is the selected model $\hat{\mathcal{S}} = \{j : W_j \geq T_0\}$. In [1], it is argued that the ratio appearing in the right-hand side of (7) is an estimate of the false discovery proportion (FDP) if we were to use the threshold t —this is true because $\mathbb{P}\{W_j \geq t\} = \mathbb{P}\{W_j \leq -t\}$ for any null feature $j \in \mathcal{H}_0$, and so we would roughly expect

$$\begin{aligned} (\# \text{ false positives at threshold } t) &= \#\{j \in \mathcal{H}_0 : W_j \geq t\} \approx \#\{j \in \mathcal{H}_0 : W_j \leq -t\} \\ &\leq \#\{j : W_j \leq -t\}, \end{aligned} \quad (8)$$

that is, the numerator in (7) is an (over)estimate of the number of false positives selected at the threshold t . Hence, the selection rule can be interpreted as a step-up rule stopping the first time our estimate falls below our target level. A slightly more conservative procedure, the knockoff+ filter, is given by incrementing the number of negatives by one, replacing the threshold in (7) with the choice

$$T_+ = \min \left\{ t > 0 : \frac{1 + \#\{j : W_j \leq -t\}}{\#\{j : W_j \geq t\}} \leq q \right\}, \quad (9)$$

and setting $\hat{\mathcal{S}} = \{j : W_j \geq T_+\}$. Formalizing the intuition of our rough calculation (8), the false discovery rate control properties of these two procedures are studied in [1] under an exact pairwise exchangeability setting.

³We want T_0 to be positive and the formal definition is that the minimum in (7) is taken over all $t > 0$ taking on values in the set $\{|W_1|, \dots, |W_p|\}$.

3 FDR control results

3.1 Measuring errors in the distribution

If the knockoff features are generated using a mechanism designed to mimic the estimated conditionals Q_j rather than the true conditional distributions P_j , when can we hope for error control? Intuitively, if the conditional distributions P_j and Q_j are similar, then we might hope that the knockoff feature $\tilde{\mathbf{X}}_j$ is a reasonably good control group for the original feature \mathbf{X}_j .

In order to quantify this, we begin by measuring the discrepancy between the true conditional P_j and its estimate Q_j . Define the random variable

$$\widehat{\text{KL}}_j := \sum_i \log \left(\frac{P_j(\mathbf{X}_{ij} | \mathbf{X}_{i,-j}) \cdot Q_j(\tilde{\mathbf{X}}_{ij} | \mathbf{X}_{i,-j})}{Q_j(\mathbf{X}_{ij} | \mathbf{X}_{i,-j}) \cdot P_j(\tilde{\mathbf{X}}_{ij} | \mathbf{X}_{i,-j})} \right), \quad (10)$$

where the notation $\widehat{\text{KL}}_j$ suggests the KL divergence. In fact, $\widehat{\text{KL}}_j$ is the *observed* KL divergence between $(\mathbf{X}_j, \tilde{\mathbf{X}}_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j})$ and $(\tilde{\mathbf{X}}_j, \mathbf{X}_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j})$. To prove this, working in the discrete case for simplicity, Lemma 1 tells us that

$$\sum_i \log \left(\frac{P_j(\mathbf{x}_{ij} | \mathbf{x}_{i,-j}) \cdot Q_j(\tilde{\mathbf{x}}_{ij} | \mathbf{x}_{i,-j})}{Q_j(\mathbf{x}_{ij} | \mathbf{x}_{i,-j}) \cdot P_j(\tilde{\mathbf{x}}_{ij} | \mathbf{x}_{i,-j})} \right) = \log \left(\frac{\mathbb{P} \left\{ (\mathbf{X}_j, \tilde{\mathbf{X}}_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}) = (\mathbf{x}_j, \tilde{\mathbf{x}}_j, \mathbf{x}_{-j}, \tilde{\mathbf{x}}_{-j}) \right\}}{\mathbb{P} \left\{ (\tilde{\mathbf{X}}_j, \mathbf{X}_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}) = (\mathbf{x}_j, \tilde{\mathbf{x}}_j, \mathbf{x}_{-j}, \tilde{\mathbf{x}}_{-j}) \right\}} \right)$$

for any $\mathbf{x}_j, \tilde{\mathbf{x}}_j, \mathbf{x}_{-j}, \tilde{\mathbf{x}}_{-j}$. Therefore, we see that

$$\mathbb{E} \left[\widehat{\text{KL}}_j \right] = d_{\text{KL}} \left((\mathbf{X}_j, \tilde{\mathbf{X}}_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}) \parallel (\tilde{\mathbf{X}}_j, \mathbf{X}_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}) \right),$$

where d_{KL} is the usual KL divergence between distributions.

In the exact model-X setting, where the knockoff construction mechanism $P_{\tilde{\mathbf{X}}|X}$ satisfies the pairwise exchangeability property (1), Proposition 1 immediately implies that $(\mathbf{X}_j, \tilde{\mathbf{X}}_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}) \stackrel{d}{=} (\tilde{\mathbf{X}}_j, \mathbf{X}_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j})$ and, thus, $\mathbb{E} \left[\widehat{\text{KL}}_j \right] = 0$. (In fact, since we are using the true conditionals P_j , or in other words $Q_j = P_j$, we would have $\widehat{\text{KL}}_j = 0$ always.) In the approximate model-X framework, we can interpret $\widehat{\text{KL}}_j$ as measuring the extent to which the pairwise exchangeability property (2) is violated for a specific feature j . We will see in our results below that controlling the $\widehat{\text{KL}}_j$'s is sufficient to ensure control of the false discovery rate for the approximate model-X knockoffs method. More precisely, we will be able to bound the false positives coming from those null features which have small $\widehat{\text{KL}}_j$.

3.2 FDR control guarantee

We now present our guarantee for robust error control with the model-X knockoffs filter.

Theorem 1. *Under the definitions above, for any $\epsilon \geq 0$, consider the null variables for which $\widehat{\text{KL}}_j \leq \epsilon$. If we use the knockoff+ filter, then the fraction of the rejections that correspond to such nulls obeys*

$$\mathbb{E} \left[\frac{|\{j : j \in \hat{\mathcal{S}} \cap \mathcal{H}_0 \text{ and } \widehat{\text{KL}}_j \leq \epsilon\}|}{|\hat{\mathcal{S}}| \vee 1} \right] \leq q \cdot e^\epsilon. \quad (11)$$

In particular, this implies that the false discovery rate is bounded as

$$\text{FDR} \leq \min_{\epsilon \geq 0} \left\{ q \cdot e^\epsilon + \mathbb{P} \left(\max_{j \in \mathcal{H}_0} \widehat{\text{KL}}_j > \epsilon \right) \right\}. \quad (12)$$

Similarly, for the knockoff filter, for any $\epsilon \geq 0$, a slightly modified fraction of the rejections that correspond to nulls with $\widehat{\text{KL}}_j \leq \epsilon$ obeys

$$\mathbb{E} \left[\frac{|\{j : j \in \widehat{\mathcal{S}} \cap \mathcal{H}_0 \text{ and } \widehat{\text{KL}}_j \leq \epsilon\}|}{|\widehat{\mathcal{S}}| + q^{-1}} \right] \leq q \cdot e^\epsilon,$$

and therefore, we obtain a bound on a modified false discovery rate:

$$\mathbb{E} \left[\frac{|\widehat{\mathcal{S}} \cap \mathcal{H}_0|}{|\widehat{\mathcal{S}}| + q^{-1}} \right] \leq \min_{\epsilon \geq 0} \left\{ q \cdot e^\epsilon + \mathbb{P} \left(\max_{j \in \mathcal{H}_0} \widehat{\text{KL}}_j > \epsilon \right) \right\}.$$

In Section 3.3, we will see concrete examples where $\max_{j=1, \dots, p} \widehat{\text{KL}}_j$ is small with high probability, yielding a meaningful result on FDR control.

A noteworthy aspect of this result is that it makes no modeling assumption whatsoever. Indeed, our FDR control guarantees hold in any setting—no matter the relationship $P_{Y|X}$ between Y and X , no matter the distribution P_X of the feature vector X , and no matter the test statistics W the data analyst has decided to employ (as long as W obeys the flip-sign condition). What the theorem says is that when we use estimated conditionals Q_j , if the Q_j 's are close to the true conditionals P_j in the sense that the quantities $\widehat{\text{KL}}_j$ are small, then the FDR is well under control. In the ideal case where we use the true conditionals, then $\widehat{\text{KL}}_j = 0$ for all $j \in \mathcal{H}_0$, and we automatically recover the FDR-control result from Candès et al. [4]; that is, we get FDR control at the nominal level q since we can take $\epsilon = 0$.

It worth pausing to unpack our main result a little. Clearly, we cannot hope to have error control over *all* nulls if we have done a poor job in constructing some of their knockoff copies, because our knockoff “negative controls” may be completely off. Having said this, (11) tells us that that if we restrict our definition of false positives to only those nulls for which we have a reasonable “negative control” via the knockoff construction, then the FDR is controlled. Since we do not make any assumptions, this type of result is all one can really hope for. In other words, exact model-X knockoffs make the assumption that the knockoff features provide exact controls for each null, thus ensuring control of the false positives; our new result removes this assumption, and provides a bound on the false positives when counting only those nulls for which the corresponding knockoff feature provides an approximate control.

In a similar fashion, imagine running a multiple comparison procedure, e.g. the Benjamini–Hochberg procedure, with p-values that are not uniformly distributed under the null. Then in such a situation, we cannot hope to achieve error control over *all* nulls if some of the null p-values follow grossly incorrect distributions. However, we may still hope to achieve reasonable control over those nulls for which the p-value is close to uniform.

As we remarked earlier, our FDR bounds apply universally (in a sense, they are worst-case bounds) and it is expected that in any practical scenario, the achieved FDR would be lower than that suggested by our upper bounds. In particular, our theorem applies to any construction of the statistics W , including adversarial constructions that might be chosen deliberately to try to detect the differences between the X_j 's and the \tilde{X}_j 's. In practice, W would instead be chosen to try to identify strong correlations with Y , and we would not expect that this type of statistic is worst-case in terms of finding discrepancies between the distributions of X_j and \tilde{X}_j . In fact, empirical studies [4, 13] have already reported on the robustness of model-X knockoffs vis-à-vis possibly large model misspecifications when W is chosen to identify a strong dependence between X and Y .

Finally, we close this section by emphasizing that the proof of Theorem 1, which is presented in Section 4, employs arguments that are completely different from those one finds in the existing literature on knockoffs. We discuss the novelties in our techniques in Section 4.

3.3 Bounding $\widehat{\text{KL}}_j$ to control FDR

To make our FDR control results more concrete, we will consider settings where accurate estimates Q_j of the conditionals P_j ensure that the $\widehat{\text{KL}}_j$'s are bounded near zero. Examining the definition (10) of $\widehat{\text{KL}}_j$, we see that $\widehat{\text{KL}}_j$ is a sum of n i.i.d. terms, and we can therefore expect that large deviation bounds such as Hoeffding's

inequality can be used to provide an upper bound uniformly across all p features. Here we give two specific examples.

3.3.1 Bounded errors in the likelihood ratio

First, suppose that our estimates Q_j of the conditional distribution P_j satisfy a likelihood ratio bound uniformly over any values for the variables:

$$\log \left(\frac{P_j(x_j | x_{-j}) \cdot Q_j(x'_j | x_{-j})}{Q_j(x_j | x_{-j}) \cdot P_j(x'_j | x_{-j})} \right) \leq \delta \quad (13)$$

for all j , all x_j, x'_j , and all x_{-j} . In this setting, the following lemma, proved via Hoeffding's inequality, gives a bound on the $\widehat{\text{KL}}_j$'s:

Lemma 2. *If the condition (13) holds uniformly for all j and all x_j, x'_j, x_{-j} , then with probability at least $1 - \frac{1}{p}$,*

$$\max_{j=1, \dots, p} \widehat{\text{KL}}_j \leq \frac{n\delta^2}{2} + 2\sqrt{n \log(p)} \cdot \delta.$$

In other words, if Q_j satisfies (13) for some $\delta = o\left(\frac{1}{\sqrt{n \log(p)}}\right)$, then with high probability every $\widehat{\text{KL}}_j$ will be small. By Theorem 1, then, the FDR for model-X knockoffs in this setting is controlled near the target level q .

3.3.2 Gaussian knockoffs

For a second example, suppose that the distribution of the feature vector X is mean zero and has covariance Θ^{-1} , where Θ is some unknown precision matrix. (We assume zero mean for simplicity, but these results can of course be generalized to an arbitrary mean.) Suppose that we have estimated Θ with some approximation $\tilde{\Theta}$, and let Θ_j and $\tilde{\Theta}_j$ denote the j th columns of these matrices. Our results below will assume that the error in estimating each column of Θ is small, i.e. $\tilde{\Theta}_j - \Theta_j$ is small for all j .

Following Candès et al. [4, eqn. (3.2)], the Gaussian knockoff construction consists of drawing \tilde{X} from the distribution

$$\tilde{\mathbf{X}}_{i,*} | \mathbf{X}_{i,*} \sim \mathcal{N}_p((\mathbf{I}_p - D\tilde{\Theta})\mathbf{X}_{i,*}, (2D - D\tilde{\Theta}D)), \quad (14)$$

independently for each i , where $D = \text{diag}\{d_j\}$ is a nonnegative diagonal matrix chosen to satisfy $2D - D\tilde{\Theta}D \succeq 0$, or equivalently, $D \preceq 2\tilde{\Theta}^{-1}$. Another formulation is that the distribution $P_{\tilde{X}|X}(\cdot|X)$ for \tilde{X} conditional on X is obtained by setting

$$\tilde{X} | X \sim \mathcal{N}_p((\mathbf{I}_p - D\tilde{\Theta})X, 2D - D\tilde{\Theta}D). \quad (15)$$

If the true precision matrix of X were given by $\tilde{\Theta}$ (assumed to be positive definite), then we can calculate that the joint distribution of the pair (X, \tilde{X}) has first and second moments given by

$$\mathbb{E} \left[\begin{pmatrix} X \\ \tilde{X} \end{pmatrix} \right] = \begin{pmatrix} 0 \\ 0 \end{pmatrix}, \quad \text{Var} \left(\begin{pmatrix} X \\ \tilde{X} \end{pmatrix} \right) = \begin{pmatrix} \tilde{\Theta}^{-1} & \tilde{\Theta}^{-1} - D \\ \tilde{\Theta}^{-1} - D & \tilde{\Theta}^{-1} \end{pmatrix}.$$

In other words, for every j , X_j and \tilde{X}_j are exchangeable if we only look at the first and second moments of the joint distribution.

If the true distribution of X is in fact Gaussian, again with mean zero and covariance $\tilde{\Theta}^{-1}$, then a stronger claim follows—the joint distribution of (X, \tilde{X}) is then multivariate Gaussian and therefore $(X, \tilde{X})_{\text{swap}(\mathcal{A})} \stackrel{d}{=} (X, \tilde{X})$ for every subset $\mathcal{A} \subseteq [p]$. In other words, the knockoff construction determined by $P_{\tilde{X}|X}$ satisfies pairwise exchangeability, as defined in (3), with respect to the distribution $Q_X = \mathcal{N}_p(0, \tilde{\Theta}^{-1})$. To frame this

property in terms of conditionals, Q_j , rather than an estimated joint distribution, Q_X , we can calculate the estimated conditional distributions $Q_j(\cdot|X_{-j})$ as

$$X_j | X_{-j} \sim \mathcal{N} \left(X_{-j}^\top \begin{pmatrix} -\tilde{\Theta}_{-j,j} \\ \tilde{\Theta}_{jj} \end{pmatrix}, \frac{1}{\tilde{\Theta}_{jj}} \right), \quad (16)$$

where $\tilde{\Theta}_{-j,j} \in \mathbb{R}^{p-1}$ is the column $\tilde{\Theta}_j$ with entry $\tilde{\Theta}_{jj}$ removed.

As noted in Section 2.2, we may want to work with estimated precision matrices, which are not positive semidefinite (PSD). The rationale is that if $\tilde{\Theta}$ is fitted by regressing each X_j on the remaining features X_{-j} to produce the j th column, $\tilde{\Theta}_j$, then the result will not be PSD in general. If $\tilde{\Theta}$ is not PSD, although there is no corresponding joint distribution, the conditionals Q_j (16) are still well-defined as long as $\tilde{\Theta}_{jj} > 0$ for all j ; they are just not compatible. (Note that symmetry is a far easier constraint to enforce, e.g. by simply replacing our initial estimate $\tilde{\Theta}$ with $(\tilde{\Theta} + \tilde{\Theta}^\top)/2$, which preserves desirable features such as sparsity that might be present in the initial $\tilde{\Theta}$; in contrast, projecting to the PSD cone while enforcing sparsity constraints may be computationally challenging in high dimensions.)

Our first result verifies that this construction of $P_{\tilde{X}|X}$ satisfies pairwise exchangeability with respect to the conditional distributions Q_j given in (16):

Lemma 3. *Let $\tilde{\Theta} \in \mathbb{R}^{p \times p}$ be a symmetric matrix with a positive diagonal, and let $P_{\tilde{X}|X}$ be defined as in (15). Then, for each $j = 1, \dots, p$, $P_{\tilde{X}|X}$ is pairwise exchangeable with respect to the conditional distribution Q_j given in (16)—that is, the exchangeability condition (4) is satisfied.*

In practice, we would construct Gaussian knockoffs in situations where the distribution of X might be well approximated by a multivariate normal. The lemma below gives a high probability bound on the $\widehat{\text{KL}}_j$'s in the case where the features are indeed Gaussian but with an unknown covariance matrix Θ^{-1} . Here, Gaussian concentration results can be used to control the $\widehat{\text{KL}}_j$'s, which then yields FDR control. (We note that recent work by Fan et al. [5] also studies the Gaussian model-X knockoffs procedure with an estimated precision matrix $\tilde{\Theta}$, under a different framework.)

Lemma 4. *Let $\Theta, \tilde{\Theta} \in \mathbb{R}^{p \times p}$ be any matrices, where Θ is positive definite and $\tilde{\Theta}$ is symmetric with a positive diagonal. Suppose that $\mathbf{X}_{i,*} \stackrel{\text{iid}}{\sim} \mathcal{N}_p(0, \Theta^{-1})$, while $\tilde{\mathbf{X}} | \mathbf{X}$ is drawn according to the distribution $P_{\tilde{X}|X}$ given in (15). Define*

$$\delta_\Theta = \max_{j=1, \dots, p} (\Theta_{jj})^{-1/2} \cdot \|\Theta^{-1/2}(\tilde{\Theta}_j - \Theta_j)\|_2.$$

Then with probability at least $1 - \frac{1}{p}$,

$$\max_{j=1, \dots, p} \widehat{\text{KL}}_j \leq 4\delta_\Theta \sqrt{n \log(p)} \cdot (1 + o(1)),$$

where the $o(1)$ term refers to terms that are vanishing when we assume that $\frac{\log(p)}{n} = o(1)$ and that this upper bound is itself bounded by a constant.

(A formal bound making the $o(1)$ term explicit is provided in the proof.) In particular, comparing to our FDR control result, Theorem 1, we see that as long as the columnwise error in estimating the precision matrix Θ satisfies $\delta_\Theta = o\left(\frac{1}{\sqrt{n \log(p)}}\right)$, the FDR will be controlled near the target level q .

3.4 A lower bound on FDR

Next, we ask whether it is possible to prove a converse to our FDR control result, Theorem 1. We are interested in knowing whether bounding the $\widehat{\text{KL}}_j$'s is in fact necessary for FDR control—or is it possible to achieve an FDR control guarantee even when the $\widehat{\text{KL}}_j$'s are large? Our next result proves that, if there is a feature j for which $\widehat{\text{KL}}_j$ does not concentrate near zero, then we can construct an honest model selection method that, when

assuming that the conditional distribution of $X_j \mid X_{-j}$ is given by Q_j , *fails* to control FDR at the desired level if the true conditional distribution is in fact P_j . By “honest”, we mean that the model selection method would successfully control FDR at level q if Q_j were the true conditional distribution. Our construction does not run a knockoff filter on the data; it is instead a hypothesis testing based procedure, meaning that the $\widehat{\mathbf{KL}}_j$ ’s govern whether it is possible to control FDR in a general sense. Hence, our converse is information-theoretic in nature and not specific to the knockoff filter.

Theorem 2. *Fix any distribution P_X , any feature index j , and any estimated conditional distribution Q_j . Suppose that there exists a knockoff sampling mechanism $P_{\tilde{X}|X}$ that is pairwise exchangeable with respect to Q_j (4), such that*

$$\mathbb{P} \left\{ \widehat{\mathbf{KL}}_j \geq \epsilon \right\} \geq c$$

for some $\epsilon, c > 0$ when $(\mathbf{X}, \tilde{\mathbf{X}})$ is drawn from $P_X \times P_{\tilde{X}|X}$. Then there exists a conditional distribution $P_{Y|X}$, and a testing procedure $\widehat{\mathcal{S}}$ that maps data $(\mathbf{X}, \mathbf{Y}) \in \mathbb{R}^{n \times p} \times \mathbb{R}^n$ to a selected set of features $\widehat{\mathcal{S}}(\mathbf{X}, \mathbf{Y}) \subseteq \{1, \dots, p\}$, such that:

- If the data points $(\mathbf{X}_{i,*}, \mathbf{Y}_i)$ are i.i.d. draws from the distribution $Q_X \times P_{Y|X}$, where Q_X is any distribution whose j th conditional is Q_j (i.e. our estimated conditional distribution Q_j for feature X_j is correct), then

$$\text{FDR}(\widehat{\mathcal{S}}) = q.$$

- On the other hand, if the data points $(\mathbf{X}_{i,*}, \mathbf{Y}_i)$ are i.i.d. draws from the distribution $P_X \times P_{Y|X}$ (i.e. our estimated conditional distribution Q_j is not correct, as the true conditional distribution is P_j), then

$$\text{FDR}(\widehat{\mathcal{S}}) \geq q(1 + c(1 - e^{-\epsilon})).$$

For the last case (where P_X is the true distribution), if $c \approx 1$ (i.e. $\widehat{\mathbf{KL}}_j \geq \epsilon$ with high probability) then $\text{FDR}(\widehat{\mathcal{S}}) \gtrsim q(2 - e^{-\epsilon})$; when $\epsilon \approx 0$ is small, we have $2 - e^{-\epsilon} \approx 1 + \epsilon \approx e^\epsilon$, thus matching the upper bound in Theorem 1.

4 Proof of Theorem 1

Whereas all proofs of FDR control for the knockoff methods thus far have relied on martingale arguments (see [1, 2, 4]), here we will prove our main theorem using a novel leave-one-out argument. Before we begin, we would like to draw a loose analogy. To prove FDR controlling properties of the Benjamini–Hochberg procedure under independence of the p-values, Storey et al. [15] developed a very elegant martingale argument. Other proof techniques, however, operate by removing or leaving out one hypothesis (or one p-value); see Ferreira and Zwinderman [6] for example. At a very high level, our own methods are partially inspired by the latter approach.

Fix $\epsilon > 0$ and for any threshold $t > 0$, define

$$R_\epsilon(t) := \frac{\sum_{j \in \mathcal{H}_0} \mathbb{1} \left\{ W_j \geq t, \widehat{\mathbf{KL}}_j \leq \epsilon \right\}}{1 + \sum_{j \in \mathcal{H}_0} \mathbb{1} \{W_j \leq -t\}}.$$

Then, for the knockoff+ filter with threshold T_+ , we can write

$$\begin{aligned} \frac{|\{j : j \in \widehat{\mathcal{S}} \cap \mathcal{H}_0 \text{ and } \widehat{\mathbf{KL}}_j \leq \epsilon\}|}{|\widehat{\mathcal{S}}| \vee 1} &= \frac{\sum_{j \in \mathcal{H}_0} \mathbb{1} \left\{ W_j \geq T_+, \widehat{\mathbf{KL}}_j \leq \epsilon \right\}}{1 \vee \sum_j \mathbb{1} \{W_j \geq T_+\}} \\ &= \frac{1 + \sum_j \mathbb{1} \{W_j \leq -T_+\}}{1 \vee \sum_j \mathbb{1} \{W_j \geq T_+\}} \cdot \frac{\sum_{j \in \mathcal{H}_0} \mathbb{1} \left\{ W_j \geq T_+, \widehat{\mathbf{KL}}_j \leq \epsilon \right\}}{1 + \sum_j \mathbb{1} \{W_j \leq -T_+\}} \leq q \cdot R_\epsilon(T_+), \end{aligned}$$

where the last step holds by definition of R_ϵ and the construction of the knockoff+ filter. If we instead use the knockoff filter (rather than knockoff+), then we use the threshold T_0 and obtain

$$\frac{|\{j : j \in \widehat{\mathcal{S}} \cap \mathcal{H}_0 \text{ and } \widehat{\text{KL}}_j \leq \epsilon\}|}{q^{-1} + |\widehat{\mathcal{S}}|} = \frac{1 + \sum_j \mathbb{1}\{W_j \leq -T_0\}}{q^{-1} + \sum_j \mathbb{1}\{W_j \geq T_0\}} \cdot \frac{\sum_{j \in \mathcal{H}_0} \mathbb{1}\{W_j \geq T_0, \widehat{\text{KL}}_j \leq \epsilon\}}{1 + \sum_j \mathbb{1}\{W_j \leq -T_0\}} \leq q \cdot R_\epsilon(T_0),$$

where the last step holds by definition of R_ϵ and the construction of the knockoff filter. Either way, then, it is sufficient to prove that $\mathbb{E}[R_\epsilon(T)] \leq e^\epsilon$, where T is either T_+ or T_0 .

For each null feature j we will be conditioning on observing $\mathbf{X}_{-j}, \widetilde{\mathbf{X}}_{-j}, \mathbf{Y}$, and on observing the *unordered* pair $\{\mathbf{X}_j, \widetilde{\mathbf{X}}_j\}$ —that is, we observe both the original and knockoff features but do not know which is which. It follows from the flip-sign property that having observed all this, we know all the knockoff statistics W except for the sign of the j th component W_j . Put differently, W_{-j} and $|W_j|$ are both functions of the variables we are conditioning on, but $\text{sign}(W_j)$ is not. Without loss of generality, label the unordered pair of feature vectors $\{\mathbf{X}_j, \widetilde{\mathbf{X}}_j\}$, as $\mathbf{X}_j^{(0)}$ and $\mathbf{X}_j^{(1)}$, such that

- if $\mathbf{X}_j = \mathbf{X}_j^{(0)}$ and $\widetilde{\mathbf{X}}_j = \mathbf{X}_j^{(1)}$, then $W_j \geq 0$;
- if $\mathbf{X}_j = \mathbf{X}_j^{(1)}$ and $\widetilde{\mathbf{X}}_j = \mathbf{X}_j^{(0)}$, then $W_j \leq 0$.

Next, define

$$\widehat{\text{KL}}'_j = \sum_{i=1}^n \log \left(\frac{P_j(\mathbf{X}_{ij}^{(0)} | \mathbf{X}_{i,-j}) \cdot Q_j(\mathbf{X}_{ij}^{(1)} | \mathbf{X}_{i,-j})}{Q_j(\mathbf{X}_{ij}^{(0)} | \mathbf{X}_{i,-j}) \cdot P_j(\mathbf{X}_{ij}^{(1)} | \mathbf{X}_{i,-j})} \right) = \text{sign}(W_j) \cdot \widehat{\text{KL}}_j,$$

where the second equality holds whenever $|W_j| \neq 0$ —this follows from the definition of $\widehat{\text{KL}}_j$ and $(\mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)})$. Furthermore, for a threshold rule $T = T(W)$ mapping statistics $W \in \mathbb{R}^p$ to a threshold $T \geq 0$ (i.e. the knockoff or knockoff+ filter threshold), for each index $j = 1, \dots, p$ we define

$$T_j = T\left((W_1, \dots, W_{j-1}, |W_j|, W_{j+1}, \dots, W_p)\right) \geq 0,$$

i.e. the threshold that we would obtain if $\text{sign}(W_j)$ were set to +1. The following lemma establishes a property of the T_j 's in the context of the knockoff filter:

Lemma 5. *Let $T = T(W)$ be the threshold for either the knockoff or the knockoff+. For any j, k ,*

$$\text{if } W_j \leq -\min\{T_j, T_k\} \text{ and } W_k \leq -\min\{T_j, T_k\}, \text{ then } T_j = T_k. \quad (17)$$

(More generally, this result holds for any function $T = T(W)$ that satisfies a “stopping time condition” with respect to the signs of the W_j 's, defined as follows: for any $t \geq 0$, the event $\mathbb{1}\{T \leq t\}$ depends on W only through (1) the magnitudes $|W|$, (2) $\text{sign}(W_j)$ for each j with $|W_j| < t$, and (3) $\sum_{j: |W_j| \geq t} \text{sign}(W_j)$.)

Now with T being either the knockoff or knockoff+ thresholding rule, we have

$$\begin{aligned}
\mathbb{E}[R_\epsilon(T)] &= \mathbb{E} \left[\frac{\sum_{j \in \mathcal{H}_0} \mathbb{1} \{W_j \geq T, \widehat{\mathbf{KL}}_j \leq \epsilon\}}{1 + \sum_{j \in \mathcal{H}_0} \mathbb{1} \{W_j \leq -T\}} \right] \\
&= \sum_{j \in \mathcal{H}_0} \mathbb{E} \left[\frac{\mathbb{1} \{W_j \geq T, \widehat{\mathbf{KL}}_j \leq \epsilon\}}{1 + \sum_{k \in \mathcal{H}_0, k \neq j} \mathbb{1} \{W_k \leq -T\}} \right] \quad (\text{since } T > 0 \text{ by definition, so if } W_j \geq T \text{ then } W_j \not\leq -T) \\
&= \sum_{j \in \mathcal{H}_0} \mathbb{E} \left[\frac{\mathbb{1} \{W_j \geq T_j, \widehat{\mathbf{KL}}'_j \leq \epsilon\}}{1 + \sum_{k \in \mathcal{H}_0, k \neq j} \mathbb{1} \{W_k \leq -T_j\}} \right] \quad (\text{since if } W_j > 0 \text{ then } T = T_j \text{ and } \widehat{\mathbf{KL}}_j = \widehat{\mathbf{KL}}'_j) \quad (18) \\
&= \sum_{j \in \mathcal{H}_0} \mathbb{E} \left[\mathbb{E} \left[\frac{\mathbb{1} \{W_j \geq T_j, \widehat{\mathbf{KL}}'_j \leq \epsilon\}}{1 + \sum_{k \in \mathcal{H}_0, k \neq j} \mathbb{1} \{W_k \leq -T_j\}} \middle| \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, \mathbf{Y} \right] \right] \\
&= \sum_{j \in \mathcal{H}_0} \mathbb{E} \left[\frac{\mathbb{P} \{W_j > 0 \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, \mathbf{Y}\} \cdot \mathbb{1} \{ |W_j| \geq T_j, \widehat{\mathbf{KL}}'_j \leq \epsilon \}}{1 + \sum_{k \in \mathcal{H}_0, k \neq j} \mathbb{1} \{W_k \leq -T_j\}} \right], \quad (19)
\end{aligned}$$

where the last step holds since W_{-j} and $|W_j|$ and $\widehat{\mathbf{KL}}'_j$ are all determined after observing the variables $\mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, \mathbf{Y}$; after conditioning on these variables, the only unknown is the sign of W_j .

Next, fix any null j . Suppose that we observe $\mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, \mathbf{Y}$, so that we now know W_{-j} and $|W_j|$ but not $\text{sign}(W_j)$. Our goal is now to bound the conditional probability $\mathbb{P} \{W_j > 0 \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, \mathbf{Y}\}$. If $|W_j| = 0$ then this probability is zero. If not, then by definition of $\mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}$,

$$\frac{\mathbb{P} \{W_j > 0 \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, \mathbf{Y}\}}{\mathbb{P} \{W_j < 0 \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, \mathbf{Y}\}} = \frac{\mathbb{P} \{(\mathbf{X}_j, \tilde{\mathbf{X}}_j) = (\mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}) \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, \mathbf{Y}\}}{\mathbb{P} \{(\mathbf{X}_j, \tilde{\mathbf{X}}_j) = (\mathbf{X}_j^{(1)}, \mathbf{X}_j^{(0)}) \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, \mathbf{Y}\}}.$$

Since the observations $i = 1, \dots, n$ are independent, this can be rewritten as

$$\begin{aligned}
&\prod_{i=1}^n \frac{\mathbb{P} \{(\mathbf{X}_{ij}, \tilde{\mathbf{X}}_{ij}) = (\mathbf{X}_{ij}^{(0)}, \mathbf{X}_{ij}^{(1)}) \mid \mathbf{X}_{ij}^{(0)}, \mathbf{X}_{ij}^{(1)}, \mathbf{X}_{i,-j}, \tilde{\mathbf{X}}_{i,-j}, \mathbf{Y}_i\}}{\mathbb{P} \{(\mathbf{X}_{ij}, \tilde{\mathbf{X}}_{ij}) = (\mathbf{X}_{ij}^{(1)}, \mathbf{X}_{ij}^{(0)}) \mid \mathbf{X}_{ij}^{(0)}, \mathbf{X}_{ij}^{(1)}, \mathbf{X}_{i,-j}, \tilde{\mathbf{X}}_{i,-j}, \mathbf{Y}_i\}} \\
&= \prod_{i=1}^n \frac{P_j(\mathbf{X}_{ij}^{(0)} \mid \mathbf{X}_{i,-j}) \cdot Q_j(\mathbf{X}_{ij}^{(1)} \mid \mathbf{X}_{i,-j})}{Q_j(\mathbf{X}_{ij}^{(0)} \mid \mathbf{X}_{i,-j}) \cdot P_j(\mathbf{X}_{ij}^{(1)} \mid \mathbf{X}_{i,-j})} = \exp(\widehat{\mathbf{KL}}'_j),
\end{aligned}$$

where the first equality holds by Lemma 1. Rearranging terms, we have

$$\mathbb{P} \{W_j > 0 \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, \mathbf{Y}\} = \frac{\exp(\widehat{\mathbf{KL}}'_j)}{1 + \exp(\widehat{\mathbf{KL}}'_j)}.$$

In passing, note that we here recover a key result from the knockoff literature: when $Q_j = P_j$ (i.e. we are using the correct conditionals) then $\widehat{\mathbf{KL}}'_j = 0$ and the conditional probability above is equal to $1/2$. That is, for nulls, the W_j 's are conditionally equally likely to be positive and negative (and also unconditionally).

Returning to the step (19) above, our calculation implies that

$$\begin{aligned}
\mathbb{E}[R_\epsilon(T)] &\leq \sum_{j \in \mathcal{H}_0} \mathbb{E} \left[\frac{\frac{\exp(\widehat{\mathbf{KL}}'_j)}{1 + \exp(\widehat{\mathbf{KL}}'_j)} \cdot \mathbb{1} \left\{ |W_j| \geq T_j, \widehat{\mathbf{KL}}'_j \leq \epsilon \right\}}{1 + \sum_{k \in \mathcal{H}_0, k \neq j} \mathbb{1} \{W_k \leq -T_j\}} \right] \\
&\leq \frac{e^\epsilon}{1 + e^\epsilon} \cdot \sum_{j \in \mathcal{H}_0} \mathbb{E} \left[\frac{\mathbb{1} \left\{ |W_j| \geq T_j, \widehat{\mathbf{KL}}'_j \leq \epsilon \right\}}{1 + \sum_{k \in \mathcal{H}_0, k \neq j} \mathbb{1} \{W_k \leq -T_j\}} \right] \\
&\leq \frac{e^\epsilon}{1 + e^\epsilon} \cdot \left(\sum_{j \in \mathcal{H}_0} \mathbb{E} \left[\frac{\mathbb{1} \left\{ W_j \geq T_j, \widehat{\mathbf{KL}}'_j \leq \epsilon \right\}}{1 + \sum_{k \in \mathcal{H}_0, k \neq j} \mathbb{1} \{W_k \leq -T_j\}} \right] + \sum_{j \in \mathcal{H}_0} \mathbb{E} \left[\frac{\mathbb{1} \{W_j \leq -T_j\}}{1 + \sum_{k \in \mathcal{H}_0, k \neq j} \mathbb{1} \{W_k \leq -T_j\}} \right] \right) \\
&= \frac{e^\epsilon}{1 + e^\epsilon} \cdot \left(\mathbb{E}[R_\epsilon(T)] + \mathbb{E} \left[\sum_{j \in \mathcal{H}_0} \frac{\mathbb{1} \{W_j \leq -T_j\}}{1 + \sum_{k \in \mathcal{H}_0, k \neq j} \mathbb{1} \{W_k \leq -T_j\}} \right] \right),
\end{aligned}$$

by our earlier calculation (18) for R_ϵ . Finally, the sum in this last expression can be simplified as follows: if for all null j , $W_j > -T_j$, then the sum is equal to zero, while otherwise, we can write

$$\begin{aligned}
\sum_{j \in \mathcal{H}_0} \frac{\mathbb{1} \{W_j \leq -T_j\}}{1 + \sum_{k \in \mathcal{H}_0, k \neq j} \mathbb{1} \{W_k \leq -T_j\}} &= \sum_{j \in \mathcal{H}_0} \frac{\mathbb{1} \{W_j \leq -T_j\}}{1 + \sum_{k \in \mathcal{H}_0, k \neq j} \mathbb{1} \{W_k \leq -T_k\}} \\
&= \sum_{j \in \mathcal{H}_0} \frac{\mathbb{1} \{W_j \leq -T_j\}}{\sum_{k \in \mathcal{H}_0} \mathbb{1} \{W_k \leq -T_k\}} = 1,
\end{aligned}$$

where the first step applies Lemma 5. Combining everything, we have shown that

$$\mathbb{E}[R_\epsilon(T)] \leq \frac{e^\epsilon}{1 + e^\epsilon} \cdot (\mathbb{E}[R_\epsilon(T)] + 1).$$

Rearranging terms gives $\mathbb{E}[R_\epsilon(T)] \leq e^\epsilon$, which proves the theorem.

5 Discussion

In this paper, we established that the method of model-X knockoffs is robust to errors in the underlying assumptions on the distribution of the feature vector X , making it an effective method for many practical applications, such as genome-wide association studies, where the underlying distribution on the features X_1, \dots, X_p can be estimated accurately. One notable aspect is that our theory is free of any modeling assumptions, since our theoretical guarantees hold no matter the data distribution or the statistics that the data analyst wishes to use, even if they are designed to exploit some weakness in the construction of knockoffs. Looking forward, it would be interesting to develop a theory for fixed statistics. For instance, if the researcher commits to using a pre-specified random forest feature importance statistic, or some statistic based on the magnitudes of lasso coefficients (perhaps calculated at a data-dependent value of the regularization parameter), then what can be said about FDR control? In other words, what can we say when the statistics W only probe the data in certain directions? We leave such interesting questions for further research.

A Additional proofs

A.1 Proof of Lemma 1

We prove the lemma in the case where all features are discrete; the case where some of the features may be continuous is proved analogously. First, consider any null feature index j . By definition of the nulls, we know

that $X_j \perp\!\!\!\perp Y \mid X_{-j}$. Furthermore, $\tilde{X} \perp\!\!\!\perp Y \mid X$ by construction. Therefore, the distribution of $Y \mid (X, \tilde{X})$ depends only on X_{-j} , and in particular,

$$Y \perp\!\!\!\perp (X_j, \tilde{X}_j) \mid (X_{-j}, \tilde{X}_{-j}).$$

This proves that

$$\frac{\mathbb{P}\left\{X_j = a, \tilde{X}_j = b \mid X_{-j}, \tilde{X}_{-j}, Y\right\}}{\mathbb{P}\left\{X_j = b, \tilde{X}_j = a \mid X_{-j}, \tilde{X}_{-j}, Y\right\}} = \frac{\mathbb{P}\left\{X_j = a, \tilde{X}_j = b \mid X_{-j}, \tilde{X}_{-j}\right\}}{\mathbb{P}\left\{X_j = b, \tilde{X}_j = a \mid X_{-j}, \tilde{X}_{-j}\right\}}, \quad (20)$$

because the numerator and denominator are each unchanged whether we do or do not condition on Y . Thus, for null features j , it is now sufficient to prove only the first claim of the lemma, namely that the right-hand side above is equal to $\frac{P_j(a|X_{-j})Q_j(b|X_{-j})}{Q_j(a|X_{-j})P_j(b|X_{-j})}$.

From this point on, let j be any feature (null or non-null). We will now prove the first claim in the lemma. Recalling the assumption that $P_{\tilde{X}|X}$ is pairwise exchangeable with respect to Q_j (4), we introduce a pair of random variables drawn as follows: first, draw $X'_{-j} \sim P_{X_{-j}}$, where $P_{X_{-j}}$ is the distribution of X_{-j} ; then draw $X'_j \mid X'_{-j} \sim Q_j(\cdot|X'_{-j})$; and finally, draw $\tilde{X}' \mid X' \sim P_{\tilde{X}|X}(\cdot|X')$. Then by (4),

$$(X'_j, \tilde{X}'_j, X'_{-j}, \tilde{X}'_{-j}) \stackrel{d}{=} (\tilde{X}'_j, X'_j, X'_{-j}, \tilde{X}'_{-j}). \quad (21)$$

The following schematic compares the distributions of (X, \tilde{X}) and (X', \tilde{X}') :

$$\begin{array}{c} \text{Distr. of } (X, \tilde{X}) \\ \left\{ \begin{array}{l} X_{-j} \sim P_{X_{-j}} \\ X_j \mid X_{-j} \sim P_j(\cdot|X_{-j}) \\ \tilde{X} \mid X \sim P_{\tilde{X}|X}(\cdot|X) \end{array} \right. \end{array} \quad \begin{array}{c} \text{Distr. of } (X', \tilde{X}') \\ \left\{ \begin{array}{l} X'_{-j} \sim P_{X_{-j}} \\ X'_j \mid X'_{-j} \sim Q_j(\cdot|X'_{-j}) \\ \tilde{X}' \mid X' \sim P_{\tilde{X}|X}(\cdot|X') \end{array} \right. \end{array}$$

We emphasize that $P_{\tilde{X}|X}$ is the same distribution on both sides since it is generated by the same mechanism taking on as inputs the conditionals Q_j 's (see Figure 2).

The joint distribution of (X, \tilde{X}) is given by

$$\mathbb{P}\left\{X = x, \tilde{X} = \tilde{x}\right\} = P_{X_{-j}}(x_{-j})P_j(x_j|x_{-j})P_{\tilde{X}|X}(\tilde{x}|x)$$

defined over all pairs (x, \tilde{x}) in the support of (X, \tilde{X}) , and an identical expression holds for the joint distribution of (X', \tilde{X}') except that P_j is replaced with Q_j . This gives

$$\frac{\mathbb{P}\left\{X_j = a, \tilde{X}_j = b, X_{-j} = x_{-j}, \tilde{X}_{-j} = \tilde{x}_{-j}\right\}}{\mathbb{P}\left\{X'_j = a, \tilde{X}'_j = b, X'_{-j} = x_{-j}, \tilde{X}'_{-j} = \tilde{x}_{-j}\right\}} = \frac{P_j(a|x_{-j})}{Q_j(a|x_{-j})},$$

and of course the analogous statement holds with a and b reversed. Therefore,

$$\begin{aligned} & \frac{\mathbb{P}\left\{X_j = a, \tilde{X}_j = b, X_{-j} = x_{-j}, \tilde{X}_{-j} = \tilde{x}_{-j}\right\}}{\mathbb{P}\left\{X_j = b, \tilde{X}_j = a, X_{-j} = x_{-j}, \tilde{X}_{-j} = \tilde{x}_{-j}\right\}} \\ &= \frac{P_j(a|x_{-j})}{Q_j(a|x_{-j})} \cdot \frac{Q_j(b|x_{-j})}{P_j(b|x_{-j})} \cdot \frac{\mathbb{P}\left\{X'_j = a, \tilde{X}'_j = b, X'_{-j} = x_{-j}, \tilde{X}'_{-j} = \tilde{x}_{-j}\right\}}{\mathbb{P}\left\{X'_j = b, \tilde{X}'_j = a, X'_{-j} = x_{-j}, \tilde{X}'_{-j} = \tilde{x}_{-j}\right\}} \\ &= \frac{P_j(a|x_{-j})}{Q_j(a|x_{-j})} \cdot \frac{Q_j(b|x_{-j})}{P_j(b|x_{-j})} \cdot 1, \end{aligned}$$

where the last step holds by the property (21). This proves the lemma.

A.2 Proof of Theorem 2

First, we will show that our statement can be reduced to a binary hypothesis testing problem. We will work under the global null hypothesis where $Y \perp\!\!\!\perp X$, and our test will be constructed independently of Y . More formally, let $P_{Y|X}$ be any fixed distribution, e.g. $\mathcal{N}(0, 1)$. Since all features are null, this means that the false discovery proportion is 1 whenever $\widehat{\mathcal{S}}(\mathbf{X}, \mathbf{Y}) \neq \emptyset$, that is,

$$\text{FDR}(\widehat{\mathcal{S}}) = \mathbb{P} \left\{ \widehat{\mathcal{S}}(\mathbf{X}, \mathbf{Y}) \neq \emptyset \right\}.$$

Therefore, in order to prove the theorem, it is sufficient to construct a *binary* test $\psi(\mathbf{X}) \in \{0, 1\}$ such that

$$\mathbb{P}_{\mathbf{X}_{i,*} \stackrel{\text{iid}}{\sim} P_X} \{ \psi(\mathbf{X}) = 1 \} \geq q(1 + c(1 - e^{-\epsilon})), \quad \mathbb{P}_{\mathbf{X}_{i,*} \stackrel{\text{iid}}{\sim} Q_X} \{ \psi(\mathbf{X}) = 1 \} = q, \quad (22)$$

i.e. a test ψ that has better-than-random performance for testing whether the conditional distribution of X_j is given by P_j or Q_j . Once ψ is constructed, then this is sufficient for the FDR result, e.g. setting

$$\widehat{\mathcal{S}}(\mathbf{X}, \mathbf{Y}) = \begin{cases} \{j\}, & \psi(\mathbf{X}) = 1, \\ \emptyset, & \psi(\mathbf{X}) = 0. \end{cases}$$

Note that the existence of such a test ψ is essentially equivalent to proving a lower bound on

$$d_{\text{TV}} \left((P_X)^{\otimes n}, (Q_X)^{\otimes n} \right)$$

uniformly over all distributions Q_X whose j th conditional is Q_j , by the well-known equivalence between total variation distance and hypothesis testing [9]. In fact, our ψ will be given by a randomized procedure (to be fully formal, we can use the independent random Y values as a source of randomness, if needed). First, we draw $\widetilde{\mathbf{X}} | \mathbf{X}$, independently of \mathbf{Y} and drawn from the rule $P_{\widetilde{X}|X}$ as specified in the theorem, and independently we also draw $B \sim \text{Bernoulli}(2q)$ and $B' \sim \text{Bernoulli}(q)$. Next, defining $\widehat{\text{KL}}_j$ as in (10), we let

$$\psi(\mathbf{X}, \widetilde{\mathbf{X}}, B, B') = \mathbb{1} \left\{ B = 1 \text{ and } \widehat{\text{KL}}_j > 0 \right\} + \mathbb{1} \left\{ B' = 1 \text{ and } \widehat{\text{KL}}_j = 0 \right\}.$$

Clearly, by definition of B and B' , we have

$$\mathbb{P} \left\{ \psi(\mathbf{X}, \widetilde{\mathbf{X}}, B, B') = 1 \right\} = 2q \cdot \mathbb{P} \left\{ \widehat{\text{KL}}_j > 0 \right\} + q \cdot \mathbb{P} \left\{ \widehat{\text{KL}}_j = 0 \right\}, \quad (23)$$

where $\mathbb{P} \left\{ \widehat{\text{KL}}_j > 0 \right\}$ and $\mathbb{P} \left\{ \widehat{\text{KL}}_j = 0 \right\}$ are taken with respect to the joint distribution of $(\mathbf{X}, \widetilde{\mathbf{X}})$.

Next, we check that the test ψ satisfies the properties (22), as required for the FDR bounds in this theorem. We first prove the second bound in (22). Suppose $\mathbf{X}_{i,*} \stackrel{\text{iid}}{\sim} Q_X$ —that is, Q_j is indeed the correct conditional distribution for $X_j | X_{-j}$. The knockoff generating mechanism $P_{\widetilde{X}|X}$ was defined to satisfy pairwise exchangeability with respect to Q_j (4), meaning that \mathbf{X}_j and $\widetilde{\mathbf{X}}_j$ are exchangeable conditional on the other variables in this scenario. Examining the form of $\widehat{\text{KL}}_j$, we see that swapping \mathbf{X}_j and $\widetilde{\mathbf{X}}_j$ has the effect of changing the sign of $\widehat{\text{KL}}_j$. The exchangeability of the pair $(\mathbf{X}_j, \widetilde{\mathbf{X}}_j)$ implies that the distribution of $\widehat{\text{KL}}_j$ is symmetric around zero, and so under $(\mathbf{X}_{i,*}, \widetilde{\mathbf{X}}_{i,*}) \stackrel{\text{iid}}{\sim} Q_X \times P_{\widetilde{X}|X}$,

$$\mathbb{P} \left\{ \widehat{\text{KL}}_j > 0 \right\} + 0.5 \cdot \mathbb{P} \left\{ \widehat{\text{KL}}_j = 0 \right\} = 0.5.$$

Checking (23), this proves that $\mathbb{P}_{\mathbf{X}_{i,*} \stackrel{\text{iid}}{\sim} Q_X} \{ \psi(\mathbf{X}) = 1 \} = q$, which ensures FDR control for the case that the estimated conditional Q_j is in fact correct.

Finally we turn to the first part of (22), where now we assume that $(\mathbf{X}_{i,*}, \widetilde{\mathbf{X}}_{i,*}) \stackrel{\text{iid}}{\sim} P_X \times P_{\widetilde{X}|X}$. From this point on, we will condition on the observed values of \mathbf{X}_{-j} and $\widetilde{\mathbf{X}}_{-j}$. By assumption in the theorem, under

this distribution we have $\mathbb{P}\left\{\widehat{\mathbf{KL}}_j \geq \epsilon\right\} \geq c$. As in the proof of Theorem 1, we consider the unordered pair $\{\mathbf{X}_j, \tilde{\mathbf{X}}_j\}$ —that is, we see the two vectors \mathbf{X}_j and $\tilde{\mathbf{X}}_j$ but do not know which is which. Note that, with this information, we are able to compute $|\widehat{\mathbf{KL}}_j|$ but not $\text{sign}(\widehat{\mathbf{KL}}_j)$. Without loss of generality, we can label the unordered pair of feature vectors $\{\mathbf{X}_j, \tilde{\mathbf{X}}_j\}$, as $\mathbf{X}_j^{(0)}$ and $\mathbf{X}_j^{(1)}$, such that

- if $\mathbf{X}_j = \mathbf{X}_j^{(0)}$ and $\tilde{\mathbf{X}}_j = \mathbf{X}_j^{(1)}$, then $\widehat{\mathbf{KL}}_j \geq 0$;
- if $\mathbf{X}_j = \mathbf{X}_j^{(1)}$ and $\tilde{\mathbf{X}}_j = \mathbf{X}_j^{(0)}$, then $\widehat{\mathbf{KL}}_j \leq 0$.

Define $C = \text{sign}(\widehat{\mathbf{KL}}_j)$, so that $\widehat{\mathbf{KL}}_j = C \cdot |\widehat{\mathbf{KL}}_j|$. By definition of the distribution of $(\mathbf{X}, \tilde{\mathbf{X}})$, it follows from Lemma 1 that

$$\frac{\mathbb{P}\left\{(\mathbf{X}_j, \tilde{\mathbf{X}}_j) = (\mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}) \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}\right\}}{\mathbb{P}\left\{(\mathbf{X}_j, \tilde{\mathbf{X}}_j) = (\mathbf{X}_j^{(1)}, \mathbf{X}_j^{(0)}) \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}\right\}} = \prod_i \frac{P_j(\mathbf{X}_{ij}^{(0)} \mid \mathbf{X}_{i,-j})Q_j(\mathbf{X}_{ij}^{(1)} \mid \mathbf{X}_{i,-j})}{Q_j(\mathbf{X}_{ij}^{(0)} \mid \mathbf{X}_{i,-j})P_j(\mathbf{X}_{ij}^{(1)} \mid \mathbf{X}_{i,-j})}.$$

In other words, if $|\widehat{\mathbf{KL}}_j| \neq 0$, then

$$\frac{\mathbb{P}\left\{C = +1 \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}\right\}}{\mathbb{P}\left\{C = -1 \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}\right\}} = \prod_i \frac{P_j(\mathbf{X}_{ij}^{(0)} \mid \mathbf{X}_{i,-j})Q_j(\mathbf{X}_{ij}^{(1)} \mid \mathbf{X}_{i,-j})}{Q_j(\mathbf{X}_{ij}^{(0)} \mid \mathbf{X}_{i,-j})P_j(\mathbf{X}_{ij}^{(1)} \mid \mathbf{X}_{i,-j})} = \exp\left\{|\widehat{\mathbf{KL}}_j|\right\},$$

where the last step holds by our choice of which vector to label as $\mathbf{X}^{(0)}$ and which to label as $\mathbf{X}^{(1)}$.

Therefore, we can write

$$\begin{aligned} c &\leq \mathbb{P}\left\{\widehat{\mathbf{KL}}_j \geq \epsilon\right\} = \mathbb{P}\left\{C = +1 \text{ and } |\widehat{\mathbf{KL}}_j| \geq \epsilon\right\} \\ &= \mathbb{E}\left[\mathbb{P}\left\{C = +1 \mid \mathbf{X}_j^{(0)}, \mathbf{X}_j^{(1)}, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}\right\} \cdot \mathbb{1}\left\{|\widehat{\mathbf{KL}}_j| \geq \epsilon\right\}\right] \\ &= \mathbb{E}\left[\frac{e^{|\widehat{\mathbf{KL}}_j|}}{1 + e^{|\widehat{\mathbf{KL}}_j|}} \cdot \mathbb{1}\left\{|\widehat{\mathbf{KL}}_j| \geq \epsilon\right\}\right]. \end{aligned} \tag{24}$$

We can similarly calculate $\mathbb{P}\left\{\widehat{\mathbf{KL}}_j > 0\right\} = \mathbb{E}\left[\frac{e^{|\widehat{\mathbf{KL}}_j|}}{1 + e^{|\widehat{\mathbf{KL}}_j|}} \cdot \mathbb{1}\left\{|\widehat{\mathbf{KL}}_j| > 0\right\}\right]$. Therefore,

$$\begin{aligned} \frac{1}{2}\mathbb{P}\left\{\widehat{\mathbf{KL}}_j = 0\right\} + \mathbb{P}\left\{\widehat{\mathbf{KL}}_j > 0\right\} &= \mathbb{E}\left[\frac{e^0}{1 + e^0} \cdot \mathbb{1}\left\{|\widehat{\mathbf{KL}}_j| = 0\right\}\right] + \mathbb{E}\left[\frac{e^{|\widehat{\mathbf{KL}}_j|}}{1 + e^{|\widehat{\mathbf{KL}}_j|}} \cdot \mathbb{1}\left\{|\widehat{\mathbf{KL}}_j| > 0\right\}\right] \\ &= \mathbb{E}\left[\frac{e^{|\widehat{\mathbf{KL}}_j|}}{1 + e^{|\widehat{\mathbf{KL}}_j|}}\right]. \end{aligned}$$

To continue, observe that for $t \geq 0$, $e^t/(1 + e^t) \geq 1/2$. Hence,

$$\begin{aligned} \mathbb{E}\left[\frac{e^{|\widehat{\mathbf{KL}}_j|}}{1 + e^{|\widehat{\mathbf{KL}}_j|}}\right] &\geq \frac{1}{2} + \mathbb{E}\left[\left(\frac{e^{|\widehat{\mathbf{KL}}_j|}}{1 + e^{|\widehat{\mathbf{KL}}_j|}} - \frac{1}{2}\right) \cdot \mathbb{1}\left\{|\widehat{\mathbf{KL}}_j| \geq \epsilon\right\}\right] \\ &\geq \frac{1}{2} + \min_{t \geq \epsilon} \frac{e^t}{1 + e^t} - \frac{1}{2} \cdot \underbrace{\mathbb{E}\left[\frac{e^{|\widehat{\mathbf{KL}}_j|}}{1 + e^{|\widehat{\mathbf{KL}}_j|}} \cdot \mathbb{1}\left\{|\widehat{\mathbf{KL}}_j| \geq \epsilon\right\}\right]}_{\geq c \text{ by (24)}} \\ &\geq \frac{1}{2} (1 + c(1 - e^{-\epsilon})), \end{aligned}$$

where for the last step we check that the minimum is attained at $t = \epsilon$. This proves that, when $\mathbf{X}_{i,*} \stackrel{\text{iid}}{\sim} P_X$, we have $\psi(\mathbf{X}, \tilde{\mathbf{X}}, B, B') = 1$ with probability at least $q(1 + c(1 - e^{-\epsilon}))$, and so the first part of (22) is satisfied, as desired.

A.3 Proof of Lemma 2

We will in fact prove a more general result, which will be useful later on:

Lemma 6. Fix any $\delta \geq 0$, and define the event

$$\mathcal{E}_\delta = \left\{ \sum_i \left[\log \left(\frac{P_j(\mathbf{X}_{ij} | \mathbf{X}_{i,-j}) Q_j(\tilde{\mathbf{X}}_{ij} | \mathbf{X}_{i,-j})}{Q_j(\mathbf{X}_{ij} | \mathbf{X}_{i,-j}) P_j(\tilde{\mathbf{X}}_{ij} | \mathbf{X}_{i,-j})} \right) \right]^2 \leq n\delta^2 \text{ for all } i, j \right\}.$$

Then

$$\mathbb{P} \left\{ \max_{j=1, \dots, p} \widehat{\text{KL}}_j \leq \frac{n\delta^2}{2} + 2\delta\sqrt{n \log(p)} \right\} \geq 1 - \frac{1}{p} - \mathbb{P} \{(\mathcal{E}_\delta)^c\}.$$

In order to prove Lemma 2, then, we simply observe that if the universal bound (13) holds for the likelihood ratios, then the event \mathcal{E}_δ occurs with probability 1.

Now we prove the general result, Lemma 6. Fix any j . Suppose that we condition on $\mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}$, and on the unordered pair $\{\mathbf{X}_{ij}, \tilde{\mathbf{X}}_{ij}\} = \{a_{ij}, b_{ij}\}$ for each i —that is, after observing the unlabeled pair, we arbitrarily label them as a and b . Write $a_j = (a_{1j}, \dots, a_{nj})$ and same for b_j . Let $C_{ij} = 0$ if $a_{ij} = b_{ij}$, and otherwise let

$$C_{ij} := \begin{cases} +1, & \text{if } (\mathbf{X}_{ij}, \tilde{\mathbf{X}}_{ij}) = (a_{ij}, b_{ij}), \\ -1, & \text{if } (\mathbf{X}_{ij}, \tilde{\mathbf{X}}_{ij}) = (b_{ij}, a_{ij}). \end{cases}$$

Then we have

$$\begin{aligned} \widehat{\text{KL}}_j &= \sum_i \log \left(\frac{P_j(\mathbf{X}_{ij} | \mathbf{X}_{i,-j}) \cdot Q_j(\tilde{\mathbf{X}}_{ij} | \mathbf{X}_{i,-j})}{Q_j(\mathbf{X}_{ij} | \mathbf{X}_{i,-j}) \cdot P_j(\tilde{\mathbf{X}}_{ij} | \mathbf{X}_{i,-j})} \right) \\ &= \sum_i C_{ij} \log \left(\frac{P_j(a_{ij} | \mathbf{X}_{i,-j}) \cdot Q_j(b_{ij} | \mathbf{X}_{i,-j})}{Q_j(a_{ij} | \mathbf{X}_{i,-j}) \cdot P_j(b_{ij} | \mathbf{X}_{i,-j})} \right) =: \sum_i C_{ij} \widehat{\text{KL}}_{ij}. \end{aligned}$$

By Lemma 1, for each i with $a_{ij} \neq b_{ij}$ we have

$$\begin{aligned} \frac{\mathbb{P} \left\{ C_{ij} = +1 \mid a_j, b_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j} \right\}}{\mathbb{P} \left\{ C_{ij} = -1 \mid a_j, b_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j} \right\}} &= \frac{\mathbb{P} \left\{ (\mathbf{X}_{ij}, \tilde{\mathbf{X}}_{ij}) = (a_{ij}, b_{ij}) \mid a_j, b_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j} \right\}}{\mathbb{P} \left\{ (\mathbf{X}_{ij}, \tilde{\mathbf{X}}_{ij}) = (b_{ij}, a_{ij}) \mid a_j, b_j, \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j} \right\}} \\ &= \frac{P_j(a_{ij} | \mathbf{X}_{i,-j}) Q_j(b_{ij} | \mathbf{X}_{i,-j})}{Q_j(a_{ij} | \mathbf{X}_{i,-j}) P_j(b_{ij} | \mathbf{X}_{i,-j})} = e^{\widehat{\text{KL}}_{ij}}. \quad (25) \end{aligned}$$

Note that this binary outcome is independent for each i . From this point on we treat $\mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, a_j, b_j$ as fixed (where $a_j = (a_{1j}, \dots, a_{nj})$ and same for b_j), and only the C_{ij} 's as random. Since $\widehat{\text{KL}}_j$ depends only on $\mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, a_j, b_j$ (i.e. on the variables that we are conditioning on), and is therefore treated as fixed, while $|C_{ij}| \leq 1$ by definition, we see that

$$\mathbb{P} \left\{ \widehat{\text{KL}}_j - \mathbb{E} \left[\widehat{\text{KL}}_j \mid \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, a_j, b_j \right] \geq 2\sqrt{\log(p)} \sqrt{\sum_i (\widehat{\text{KL}}_{ij})^2} \mid \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, a_j, b_j \right\} \leq \frac{1}{p^2}$$

by Hoeffding's inequality. Next we work with the conditional expectation of $\widehat{\text{KL}}_j$. For any i with $a_{ij} \neq b_{ij}$, we use (25) to calculate

$$\left| \mathbb{E} \left[C_{ij} \mid \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, a_j, b_j \right] \right| = \left| \frac{e^{\widehat{\text{KL}}_{ij}} - 1}{e^{\widehat{\text{KL}}_{ij}} + 1} \right| \leq \frac{|\widehat{\text{KL}}_{ij}|}{2}.$$

Then

$$\left| \mathbb{E} \left[\widehat{\text{KL}}_j \mid \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, a_j, b_j \right] \right| = \left| \sum_i \mathbb{E} \left[C_{ij} \mid \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, a_j, b_j \right] \cdot \widehat{\text{KL}}_{ij} \right| \leq \frac{1}{2} \sum_i (\widehat{\text{KL}}_{ij})^2.$$

Therefore, combining everything,

$$\mathbb{P} \left\{ \widehat{\text{KL}}_j \geq \frac{1}{2} \sum_i (\widehat{\text{KL}}_{ij})^2 + 2\sqrt{\log(p)} \sqrt{\sum_i (\widehat{\text{KL}}_{ij})^2} \mid \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, a_j, b_j \right\} \leq \frac{1}{p^2}.$$

Now, under the event \mathcal{E}_δ we must have $\sum_i (\widehat{\text{KL}}_{ij})^2 \leq n\delta^2$, and so we can write

$$\mathbb{P} \left\{ \widehat{\text{KL}}_j \cdot \mathbb{1} \{ \mathcal{E}_\delta \} \geq \frac{n\delta^2}{2} + 2\delta\sqrt{n \log(p)} \mid \mathbf{X}_{-j}, \tilde{\mathbf{X}}_{-j}, a_j, b_j \right\} \leq \frac{1}{p^2}.$$

Marginalizing over all the conditioned variables, and taking a union bound over all j , we have proved that

$$\mathbb{P} \left\{ \max_{j=1, \dots, p} \widehat{\text{KL}}_j \cdot \mathbb{1} \{ \mathcal{E}_\delta \} \geq \frac{n\delta^2}{2} + 2\delta\sqrt{n \log(p)} \right\} \leq \frac{1}{p}.$$

This proves the lemma.

A.4 Proof of Lemma 3

Fix any feature index j , and consider any distribution $D^{(j)}$ on \mathbb{R}^p with j th conditional equal to Q_j , as defined in (16). For simplicity, from this point on, we will perform calculations treating $D^{(j)}$ as a joint density, but the result is valid without this assumption. Drawing $X \sim D^{(j)}$ and $\tilde{X} \mid X \sim P_{\tilde{X}|X}(\cdot|X)$, then the joint density of (X, \tilde{X}) is given by

$$D^{(j)}(x) \cdot P_{\tilde{X}|X}(\tilde{x} \mid x) = \underbrace{D_{-j}^{(j)}(x_{-j})}_{\text{Term 1}} \cdot \underbrace{\left(\frac{Q_j(x_j \mid x_{-j})}{\exp \left\{ -\frac{1}{2} x^\top \tilde{\Theta} x \right\}} \right)}_{\text{Term 2}} \cdot \underbrace{\left(P_{\tilde{X}|X}(\tilde{x} \mid x) \cdot \exp \left\{ -\frac{1}{2} x^\top \tilde{\Theta} x \right\} \right)}_{\text{Term 3}},$$

where $D_{-j}^{(j)}$ is the marginal distribution of X_{-j} under the joint distribution $X \sim D^{(j)}$. In order to check that (4) holds, i.e. that $(X_j, \tilde{X}_j, X_{-j}, \tilde{X}_{-j}) \stackrel{d}{=} (\tilde{X}_j, X_j, X_{-j}, \tilde{X}_{-j})$ under this distribution, we therefore need to check that this joint density is exchangeable in the variables x_j and \tilde{x}_j ; that is, swapping x_j and \tilde{x}_j does not change the value of the joint density $D^{(j)}(x) \cdot P_{\tilde{X}|X}(\tilde{x} \mid x)$. We check this by considering each of the three terms separately. Term 1 clearly does not depend on either x_j or \tilde{x}_j . Next, using the calculation of Q_j in (16), we can simplify Term 2 to obtain

$$\begin{aligned} \text{Term 2} &\propto \exp \left\{ -\frac{1}{2\tilde{\Theta}_{jj}} \left(x_j + x_{-j}^\top \tilde{\Theta}_{-j,j} / \tilde{\Theta}_{jj} \right)^2 + \frac{1}{2} x^\top \tilde{\Theta} x \right\} \\ &= \exp \left\{ \frac{1}{2} x_{-j}^\top \left(\tilde{\Theta}_{-j,-j} - \frac{\tilde{\Theta}_{-j,j} \tilde{\Theta}_{-j,j}^\top}{\tilde{\Theta}_{jj}} \right) x_{-j} \right\}, \end{aligned}$$

which also does not depend on either x_j or \tilde{x}_j . Finally, Term 3 is exchangeable in the pair x_j, \tilde{x}_j by the construction of the knockoff distribution $P_{\tilde{X}|X}$. More concretely, using the definition of $P_{\tilde{X}|X}$ given in (15), we can calculate

$$\begin{aligned} \text{Term 3} &\propto \exp \left\{ -\frac{1}{2} (\tilde{x} - (\mathbf{I} - D\tilde{\Theta})x)^\top (2D - D\tilde{\Theta}D)^{-1} (\tilde{x} - (\mathbf{I} - D\tilde{\Theta})x) - \frac{1}{2} x^\top \tilde{\Theta}x \right\} \\ &= \exp \left\{ -\frac{1}{2} (x + \tilde{x})^\top (2D - D\tilde{\Theta}D)^{-1} (x + \tilde{x}) + x^\top D^{-1}\tilde{x} \right\}, \end{aligned}$$

which is clearly exchangeable in the pair x_j, \tilde{x}_j (note that the exchangeability of x_j, \tilde{x}_j in the term $x^\top D^{-1}\tilde{x}$ follows from the fact that D is a diagonal matrix).

A.5 Proof of Lemma 4

We will apply Lemma 6 to prove this result. We first recall the conditional distributions P_j for the joint distribution $P_X = \mathcal{N}_p(0, \Theta)^{-1}$, which can be computed as

$$P_j(\cdot | x_{-j}) = \mathcal{N}(x_{-j}^\top (-\Theta_{-j,j} / \Theta_{jj}), 1 / \Theta_{jj}),$$

and the conditionals Q_j , calculated earlier in (16) as

$$Q_j(\cdot | x_{-j}) = \mathcal{N}(x_{-j}^\top (-\tilde{\Theta}_{-j,j} / \tilde{\Theta}_{jj}), 1 / \tilde{\Theta}_{jj}).$$

Then we can calculate

$$\begin{aligned} &\sum_i \left[\log \left(\frac{P_j(\mathbf{X}_{ij} | \mathbf{X}_{i,-j}) Q_j(\tilde{\mathbf{X}}_{ij} | \mathbf{X}_{i,-j})}{Q_j(\mathbf{X}_{ij} | \mathbf{X}_{i,-j}) P_j(\tilde{\mathbf{X}}_{ij} | \mathbf{X}_{i,-j})} \right) \right]^2 \\ &= \sum_i \left[-(\mathbf{X}_{ij} - \tilde{\mathbf{X}}_{ij}) \cdot \frac{\tilde{\Theta}_{jj} - \Theta_{jj}}{2} + \mathbf{X}_{i*}^\top (\tilde{\Theta}_j - \Theta_j) \right]^2 \cdot [\mathbf{X}_{ij} - \tilde{\mathbf{X}}_{ij}]^2 \\ &\leq \frac{1}{2} \sum_i \underbrace{\left[-(\mathbf{X}_{ij} - \tilde{\mathbf{X}}_{ij}) \cdot \frac{\tilde{\Theta}_{jj} - \Theta_{jj}}{2} + \mathbf{X}_{i*}^\top (\tilde{\Theta}_j - \Theta_j) \right]^4}_{\sim \mathcal{N}(0, v_j^2) \text{ for each } i} + \frac{1}{2} \sum_i \underbrace{\left[\mathbf{X}_{ij} - \tilde{\mathbf{X}}_{ij} \right]^4}_{\sim \mathcal{N}(0, w_j^2) \text{ for each } i}. \end{aligned}$$

Using standard tail bounds on Gaussian and χ^2 random variables, and computing the variances v_j^2 and w_j^2 , after some calculations we can show that the quantity above is bounded as

$$\sum_i \left[\log \left(\frac{P_j(\mathbf{X}_{ij} | \mathbf{X}_{i,-j}) Q_j(\tilde{\mathbf{X}}_{ij} | \mathbf{X}_{i,-j})}{Q_j(\mathbf{X}_{ij} | \mathbf{X}_{i,-j}) P_j(\tilde{\mathbf{X}}_{ij} | \mathbf{X}_{i,-j})} \right) \right]^2 \leq 4 \left[\left(\frac{\delta_\Theta}{1 - \delta_\Theta} \right)^2 + \left(\frac{\delta_\Theta}{1 - \delta_\Theta} \right)^4 \right] \cdot (\sqrt{n} + 2\sqrt{\log(np)})^2,$$

with probability at least $1 - \frac{1}{p}$, and therefore, $\mathbb{P}\{\mathcal{E}_\delta\} \geq 1 - \frac{1}{p}$ when we take

$$\delta = 2\sqrt{\left(\frac{\delta_\Theta}{1 - \delta_\Theta} \right)^2 + \left(\frac{\delta_\Theta}{1 - \delta_\Theta} \right)^4} \cdot \left(1 + 2\sqrt{\frac{\log(np)}{n}} \right) = 2\delta_\Theta \cdot (1 + o(1)),$$

where the last step holds as long as $\frac{\log(p)}{n} = o(1)$ and $\delta_\Theta = o(1)$. Applying Lemma 6 then proves that

$$\mathbb{P} \left\{ \max_{j=1, \dots, p} \widehat{\text{KL}}_j \leq \frac{n\delta^2}{2} + 2\delta\sqrt{n\log(p)} \right\} \geq 1 - \frac{2}{p}.$$

Assuming this upper bound on the $\widehat{\text{KL}}_j$'s is bounded by a constant, the dominant term is therefore $2\delta\sqrt{n\log(p)}$, which proves the lemma.

A.6 Proof of Lemma 5

First, recall that $T = T(W)$ is defined as follows:

$$T = \min \left\{ t \geq \epsilon(W) : \frac{\sum_{\ell=1}^p \mathbb{1}\{W_\ell \geq t\}}{c + \underbrace{\sum_{\ell=1}^p \mathbb{1}\{W_\ell \leq -t\}}_{=: f(W,t)}} \leq q \right\},$$

where $\epsilon(W) > 0$ is chosen to be the smallest magnitude of the W statistics, i.e. $\epsilon(W) = \min\{|W_\ell| : |W_\ell| > 0\}$, and where $c = 0$ for knockoff or $c = 1$ for knockoff+. Next, define

$$W^j := (W_1, \dots, W_{j-1}, |W_j|, W_{j+1}, \dots, W_p) \quad \text{and} \quad W^k := (W_1, \dots, W_{k-1}, |W_k|, W_{k+1}, \dots, W_p),$$

so that $T_j = T(W^j)$ and $T_k = T(W^k)$. Note that $|W^j| = |W^k| = |W|$, and so $\epsilon(W^j) = \epsilon(W^k) = \epsilon(W)$ since $\epsilon(W)$ depends on W only through $|W|$.

Without loss of generality, assume $T_j \leq T_k$, so that by assumption we have $W_j \leq -T_j$ and $W_k \leq -T_j$. Then we have

$$f(W^k, T_j) = \frac{\sum_{\ell=1}^p \mathbb{1}\{W_\ell^k \geq T_j\}}{c + \sum_{\ell=1}^p \mathbb{1}\{W_\ell^k \leq -T_j\}},$$

which can be written as

$$\frac{\sum_{\ell=1}^p \mathbb{1}\{W_\ell^j \geq T_j\} + \left(\mathbb{1}\{W_j^k \geq T_j\} - \mathbb{1}\{W_j^j \geq T_j\} + \mathbb{1}\{W_k^k \geq T_j\} - \mathbb{1}\{W_k^j \geq T_j\} \right)}{c + \sum_{\ell=1}^p \mathbb{1}\{W_\ell^j \leq -T_j\} + \left(\mathbb{1}\{W_j^k \leq -T_j\} - \mathbb{1}\{W_j^j \leq -T_j\} + \mathbb{1}\{W_k^k \leq -T_j\} - \mathbb{1}\{W_k^j \leq -T_j\} \right)},$$

since W^j and W^k differ only on entries j, k . Next, we know from our assumptions and definitions that $W_j^k = W_j \leq -T_j$; $W_j^j = |W_j| \geq T_j$; $W_k^k = |W_k| \geq T_j$; and $W_k^j = W_k \leq -T_j$. Therefore,

$$\left(\mathbb{1}\{W_j^k \geq T_j\} - \mathbb{1}\{W_j^j \geq T_j\} + \mathbb{1}\{W_k^k \geq T_j\} - \mathbb{1}\{W_k^j \geq T_j\} \right) = 0 - 1 + 1 - 0 = 0$$

and

$$\left(\mathbb{1}\{W_j^k \leq -T_j\} - \mathbb{1}\{W_j^j \leq -T_j\} + \mathbb{1}\{W_k^k \leq -T_j\} - \mathbb{1}\{W_k^j \leq -T_j\} \right) = 1 - 0 + 0 - 1 = 0.$$

Returning to our expression for $f(W^k, T_j)$ above, then,

$$f(W^k, T_j) = \frac{\sum_{\ell=1}^p \mathbb{1}\{W_\ell^j \geq T_j\} + 0}{c + \sum_{\ell=1}^p \mathbb{1}\{W_\ell^j \leq -T_j\} + 0} = f(W^j, T_j) \leq q,$$

where the last step holds by definition of T_j . Therefore, since $T_j \geq \epsilon(W^j) = \epsilon(W^k)$, we see from the definition of T_k that we must have $T_k \leq T_j$. This proves that $T_j = T_k$, as desired.

Acknowledgements

R. F. B. was partially supported by the National Science Foundation via grant DMS 1654076, and by an Alfred P. Sloan fellowship. E. C. was partially supported by the Office of Naval Research under grant N00014-16-1-2712, by the National Science Foundation via DMS 1712800, by the Math + X Award from the Simons Foundation and by a generous gift from TwoSigma. E. C. would like to thank Chiara Sabatti for useful conversations related to this project. R. J. S. was partially supported by Engineering and Physical Sciences Research Council Fellowships EP/J017213/1 and EP/P031447/1, and by grant RG81761 from the Leverhulme Trust.

References

- [1] R. F. Barber and E. J. Candès. Controlling the false discovery rate via knockoffs. *Ann. Statist.*, 43(5): 2055–2085, 2015.
- [2] R. F. Barber and E. J. Candès. A knockoff filter for high-dimensional selective inference. *ArXiv e-prints*, February 2016.
- [3] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B*, 57(1):289–300, 1995.
- [4] E. Candès, Y. Fan, L. Janson, and J. Lv. Panning for Gold: Model-X Knockoffs for High-dimensional Controlled Variable Selection. *ArXiv e-prints*, October 2016.
- [5] Y. Fan, E. Demirkaya, G. Li, and J. Lv. RANK: large-scale inference with graphical nonlinear knockoffs. *arXiv preprint arXiv:1709.00092*, 2017.
- [6] J. A. Ferreira and A. H. Zwinderman. On the Benjamini–Hochberg method. *Ann. Statist.*, 34(4):1827–1849, 2006.
- [7] B. Howie, C. Fuchsberger, M. Stephens, J. Marchini, and G. R. Abecasis. Fast and accurate genotype imputation in genome-wide association studies through pre-phasing. *Nat. Genet.*, 44(8):955–959, 2012.
- [8] L. Janson. *A model-free approach to high-dimensional inference*. PhD thesis, Stanford University, 2017.
- [9] E.L. Lehmann and J.P. Romano. *Testing Statistical Hypotheses*. Springer Texts in Statistics. Springer New York, 2008.
- [10] N. Li and M. Stephens. Modeling linkage disequilibrium and identifying recombination hotspots using single-nucleotide polymorphism data. *Genetics*, 165(4):2213–2233, 2003.
- [11] J. Marchini and B. Howie. Genotype imputation for genome-wide association studies. *Nat. Rev. Genet.*, 11(7):499–511, 2010.
- [12] Z. S. Qin, T. Niu, and J. S. Liu. Partition-ligation-expectation-maximization algorithm for haplotype inference with single-nucleotide polymorphisms. *Am. J. Hum. Genet.*, 71(5):1242–1247, Nov 2002.
- [13] M. Sesia, C. Sabatti, and E. J. Candès. Gene Hunting with Knockoffs for Hidden Markov Models. *ArXiv e-prints*, June 2017.
- [14] M. Stephens, N. J. Smith, and P. Donnelly. A new statistical method for haplotype reconstruction from population data. *Am. J. Hum. Genet.*, 68(4):978–989, Apr 2001.
- [15] J. D. Storey, J. E. Taylor, and D. Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 66(1):187–205, 2004.
- [16] K. Zhang, M. Deng, T. Chen, M. S. Waterman, and F. Sun. A dynamic programming algorithm for haplotype block partitioning. *Proc. Natl. Acad. Sci. U.S.A.*, 99(11):7335–7339, May 2002.