

A KNOCKOFF FILTER FOR HIGH-DIMENSIONAL SELECTIVE INFERENCE

BY RINA FOYGEL BARBER¹ AND EMMANUEL J. CANDÈS²

University of Chicago and Stanford University

This paper develops a framework for testing for associations in a possibly high-dimensional linear model where the number of features/variables may far exceed the number of observational units. In this framework, the observations are split into two groups, where the first group is used to screen for a set of potentially relevant variables, whereas the second is used for inference over this reduced set of variables; we also develop strategies for leveraging information from the first part of the data at the inference step for greater power. In our work, the inferential step is carried out by applying the recently introduced knockoff filter, which creates a knockoff copy—a fake variable serving as a control—for each screened variable. We prove that this procedure controls the *directional* false discovery rate (FDR) in the reduced model controlling for all screened variables; this says that our high-dimensional knockoff procedure “discovers” important variables as well as the directions (signs) of their effects, in such a way that the expected proportion of wrongly chosen signs is below the user-specified level (thereby controlling a notion of Type S error averaged over the selected set). This result is nonasymptotic, and holds for any distribution of the original features and any values of the unknown regression coefficients, so that inference is not calibrated under hypothesized values of the effect sizes. We demonstrate the performance of our general and flexible approach through numerical studies, showing more power than existing alternatives. Finally, we apply our method to a genome-wide association study to find locations on the genome that are possibly associated with a continuous phenotype.

1. Introduction. Many modern studies in the sciences take the following form: collect a large amount of data, and then ask which of the potentially many measured variables are possibly associated with a response of interest. A typical example would be a genome-wide association study (GWAS), where one wishes to “discover” which genetic variants are possibly associated with a trait. In such studies, it is common to read genetic variants using single-nucleotide polymorphism (SNP) arrays and then look for associations between the trait and the hundreds of

Received February 2016; revised July 2018.

¹Supported in part by the NSF Grant DMS-1654076 and by an Alfred P. Sloan fellowship.

²Supported in part by the NSF Grant CCF-0963835 and by the Math + X Award from the Simons Foundation.

MSC2010 subject classifications. 62F03, 62J05.

Key words and phrases. Knockoffs, variable selection, false discovery rate (FDR), high-dimensional regression.

thousands of SNPs. In these types of studies and others, we typically have many more variables or features (SNPs) than observations (individuals in the study).

Finding statistically significant associations between a response and a large set of potentially explanatory variables require the specification of a statistical model; in this work, we consider the classical linear regression model, which takes the form

$$(1) \quad Y = \sum_{j=1}^p \beta_j X_j + \epsilon,$$

where Y is the response variable of interest, the X_j 's are the explanatory variables and ϵ is a (stochastic) error term. In the GWAS example, Y may be the level of HDL cholesterol and X_j the number of recessive alleles at a given location on the genome. In such applications, it is important to keep in mind that the linear relationship (1) is just an approximation of a more complicated and unknown model.

In classical statistics, when we speak of finding associations, we think of testing each of the p hypotheses $H_j : \beta_j = 0$, and one can see that there are a few complications with this viewpoint. The first is that for fixed designs, the linear model (1) with p parameters is not identifiable whenever the number n of samples is smaller than this number p of features. Formally, testing an unidentifiable model may seem like a problematic proposition. A second complication is that in our GWAS example, it may very well be the case that causal mutations affecting levels of HDL cholesterol, for instance, have not been typed (i.e., the array of SNPs is a subsample of the genome, and does not include the causal mutation). In such studies, one is, therefore, led to search for indirect effects. In other words, while the causal factors have not been measured, we nevertheless hope to have measured variables that are sufficiently close or correlated to be picked up by association. In these settings, it is understood that even if there is a true sparse "causal" model, we may end up working with an approximate and nonsparse statistical model of the form (1) in which few, if any, of the features have precisely zero effect (i.e., for most features j , $\beta_j = 0$ is not exactly true); this phenomenon occurs because the causal effects have been folded onto the observed variables. There are two consequences of this: (1) sparsity assumptions, which are traditionally assumed in the literature to relieve ourselves from the identifiability problem, are not applicable³ and (2) when few or none of the regression coefficients take on the value zero, controlling a classical Type I error typically makes little sense since nearly any selected feature is technically a true positive, and we need a different measure of performance to capture the idea that we would like to correctly identify the meaningful large effects while screening out features with near-zero effects. To summarize our discussion, we have at least two problems:

³The literature also typically assumes low correlations between variables, which is also not applicable in the GWAS example.

- Lack of identifiability (at least when $p > n$).
- Possibly correlated features and lack of model sparsity so that few, if any, of the null hypotheses actually hold.

These complications should not occlude the relatively simple goal of a statistical selection procedure, which in the GWAS context can be stated as follows: *understanding that SNPs may be proxies for causal mutations, we would like to find as many proxies or locations on the genome without too many false positives—a false positive being a reported location (SNP) with no causal mutation sitting nearby.* Being guaranteed a sufficiently small fraction of false positives would assure the scientist that most of her discoveries are indeed true and, in some sense, replicable. We are thus interested in reliable methods for finding the large effects, those SNPs with large coefficient magnitudes. This paper develops a framework and methods for getting close to this goal. At a high level, we propose a two-stage procedure: first, using a portion of the data set to screen for a large set of potentially relevant features (which will likely include a high proportion of false positives), and second, applying the recently developed *knockoff filter* to the second portion of the data in order to perform inference within the submodel selected at the screening stage.

1.1. *Errors of Type S and directional FDR.* In a series of papers, Gelman and his collaborators introduce a useful measure they call *error of Type S* (S stands for sign). In their work [15], they argue that “in classical statistics, the significance of comparisons (e.g., $\theta_1 - \theta_2$) is calibrated using Type I error rate, relying on the assumption that the true difference is zero, which makes no sense in many applications”. They then propose “a more relevant framework in which a true comparison can be positive or negative, and, based on the data, you can state “ $\theta_1 > \theta_2$ with confidence,” “ $\theta_2 > \theta_1$ with confidence” or “no claim with confidence.” In this framework, a Type S error (sometimes called a Type III error in other works) occurs when we claim with confidence that the comparison goes one way when, in fact, it goes the other way. This echoes an earlier point made by Tukey and we quote from [35]: “All we know about the world teaches us that the effects of A and B are always different—in some decimal place—for any A and B . Thus asking “Are the effects different” is foolish. What we should be answering first is ‘Can we tell the direction in which the effects of A differ from the effects of B ?.’” We refer to [35] as well as [21] and [5] for additional discussion.

This point of view is extremely relevant to our (possibly high-dimensional) regression problem and our goal of finding the important effects. When we choose to report a variable, we should be able to state with confidence its direction of effect, that is, whether $\beta_j > 0$ or $\beta_j < 0$. To press this point, we would surely have little faith in a procedure that would control errors of Type I but would not be able to tell the directions of the effects. Another useful aspect of Type S errors is that small effects $\beta_j \approx 0$ are generally those for which we cannot have much certainty

about their direction, since the size of the effect is comparable to or below the noise level of our estimate. Therefore, if we were concerned by signed errors, we would probably not report these effects as discoveries (even if we believe that no effect is exactly zero)—arguably a wise thing to do.

To measure our success at selecting only those effects that are large enough to be meaningfully distinguished from noise, we might like to control the mixed directional false discovery rate (FDR_{dir}) [7, 29] defined as follows: letting $\widehat{S} \subset \{1, \dots, p\}$ be the set of selected features together with estimates $\widehat{\text{sign}}_j \in \{\pm 1\}$ of the direction of effect,

$$FDR_{dir} = \mathbb{E}[FDP_{dir}],$$

where the mixed directional false discovery proportion (FDP_{dir}) is given by

$$(2) \quad FDP_{dir} = \frac{|\{j \in \widehat{S} : \widehat{\text{sign}}_j \neq \text{sign}(\beta_j)\}|}{|\widehat{S}| \vee 1}$$

with the convention that $\text{sign}(0) = 0$. This definition subsumes Type I and Type S errors: a false discovery or error occurs either when a zero effect is selected, or when a nonzero effect is selected but with the incorrect sign. (In the denominator, $|\widehat{S}| \vee 1$ denotes $\max\{|\widehat{S}|, 1\}$; this maximum is taken so that if we make zero discoveries, then the FDP is calculated as zero.) The term “mixed” comes from this combination of two error types; for our purposes, we consider declaring that $\beta_j > 0$ to be a sign error whether the truth is that $\beta_j = 0$ or $\beta_j < 0$, and will not distinguish between these two types of errors, and so we will drop the term “mixed” and refer to it simply as the “directional FDR.” The directional FDP is then the total number of errors of this (combined) type averaged over the selected set. In regression settings in which many of the β_j ’s are approximately zero but perhaps not exactly zero, controlling the directional FDR may be more appropriate than the “classical” FDR (which counts Type I errors only):

$$FDR = \mathbb{E}[FDP], \quad FDP = \frac{|\{j \in \widehat{S} \text{ and } \beta_j = 0\}|}{|\widehat{S}| \vee 1}.$$

Note that controlling the directional FDR has a different flavor than classical FDR, since paraphrasing Gelman et al., the significance is not calibrated using Type I errors which implicitly assumes that some of the β_j ’s are zero. In contrast, we are after a form of inference holding no matter the values of the β_j ’s. In classical settings with independent statistics, where $y_i \stackrel{\text{ind}}{\sim} \mathcal{N}(\mu_i, 1)$, say, and we wish to test for the means, [7] establishes that the Benjamini–Hochberg procedure [6] achieves directional FDR control.

In earlier work [1], we introduced *the knockoff filter*, a new variable selection procedure, which rigorously controls the classical FDR, in the low-dimensional setting where $p \leq n$. It should be clear that the equivalent notion (2) associated

with Type S errors is in principle more difficult to control since by definition we have

$$\text{FDR}_{\text{dir}} \geq \text{FDR},$$

due to the fact that the classical FDR does not record an error when the j th feature is correctly included in the model but with the incorrect sign, while the directional FDR does. A first contribution of this paper is to show that the knockoff filter offers more than advertised in our earlier work [1]: indeed, we will find that this filter, with no modifications, actually controls the directional FDR as well as the FDR. In particular, this implies that the knockoff filter is applicable in settings where we do not expect any sparsity but nonetheless wish to test whether our conclusions regarding effect signs are reliable. This addresses one of the complications discussed earlier.

1.2. *High dimensionality.* To address the other complication, namely, the high-dimensionality issue ($p > n$), we propose the following general strategy:

- *Screening step.* In a first step, screen all the features $\{X_j\}$, $j = 1, \dots, p$, as to identify a set $\widehat{S}_0 \subset \{1, \dots, p\}$ of potentially relevant features with $|\widehat{S}_0| < n$. In this paper, we shall regard this step as being quite liberal in the sense that it typically produces a long list, hopefully containing most of the important features (with large effect sizes) but also possibly many features with zero effects $\beta_j = 0$ or nearly vanishing effects $\beta_j \approx 0$.
- *Inference/selection step.* The screening step yields a reduced model

$$(3) \quad Y = \sum_{j \in \widehat{S}_0} \beta_j^{\text{partial}} X_j + \epsilon,$$

where we use the notation β_j^{partial} to indicate that both the definition and the meaning of the regression coefficients has changed (see below). Then test for associations in this reduced model by controlling the directional FDR.

In our GWAS example, this means that we would first screen for promising SNPs, and then extract from this set a final and smaller list of SNPs that we are confident are associated with the phenotype under study (recall that reporting a SNP only means that we are confident a mutation lies nearby). This approach—a screening step, followed by inference on the screened submodel—has been studied by many researchers; see [37, 39] for a nonexhaustive list of important works among those lines.

We would like to emphasize the importance of considering the directional FDR rather than the classical (unsigned) FDR in the framework of a screening step followed by an inference step. Imagine that our sampled observations $\mathbf{y} \in \mathbb{R}^n$ have mean $\mathbf{X}\boldsymbol{\beta}$ and uncorrelated errors. In the reduced model, we wish to provide inference about the *partial regression coefficients*, $\boldsymbol{\beta}^{\text{partial}} \in \mathbb{R}^{|\widehat{S}_0|}$, when we regress the

response \mathbf{y} onto the features $\mathbf{X}_{\hat{\mathcal{S}}_0}$ only. Even if the original coefficient vector $\boldsymbol{\beta}$ for the full regression model is exactly sparse with many entries exactly equal to zero, the vector of partial regression coefficients $\boldsymbol{\beta}^{\text{partial}}$ is nonetheless likely to be dense unless there is special structure within \mathbf{X} (e.g., if columns of \mathbf{X} are orthogonal to each other).

Now, consider a feature \mathbf{X}_j not appearing in the original full model, that is, $\beta_j = 0$. In the partial regression, we may find one of two scenarios:

- The coefficient in the partial regression may be large, that is, $\beta_j^{\text{partial}} \not\approx 0$. This typically occurs if \mathbf{X}_j is highly correlated with some strong signal \mathbf{X}_k that was missed by the screening step. (For instance, we may have missed a causal SNP but included a close neighbor, instead, in the screened set.) Here, we would generally prefer to include \mathbf{X}_j in the selected model in any case, since it is a good proxy for the missed relevant feature \mathbf{X}_k .
- Alternately, the coefficient might remain near zero, $\beta_j^{\text{partial}} \approx 0$. This is likely whenever \mathbf{X}_j is not a proxy for any missed feature. Here, the sign of β_j^{partial} cannot be estimated with much certainty, and so with directional FDR control, we are likely to exclude \mathbf{X}_j from our final model.

These considerations hopefully make clear that careful screening followed by directional FDR control might yield a valuable selection procedure, which in our GWAS example would likely be able to report SNPs (locations on the genome) associated with a trait while controlling a very natural notion of false positive rate. On the other hand, a setting where this framework does not apply is that of causal inference, in which one would like to report true causal effects and would prefer not to select variables that, due to correlation, act as proxies to the true causal features.

1.3. *Comparison with other works.* There is of course a vast literature on inference in the regression setting. When $n > p$, we have available the sampling distribution of least-squares estimates, which makes the construction of p-values and confidence intervals possible, although selective inference procedures controlling the FDR and similar criteria are more tricky because of the correlations between such test statistics. For an overview of some of the work when $n > p$, we refer the reader to the companion article [1].

There has also been much recent work on the high-dimensional setting $p > n$, which is the focus of our work here. One type of approach to high-dimensional inference is to avoid the identifiability problem by making assumptions on the data distribution such as (1) a highly sparse coefficient vector $\boldsymbol{\beta}$, and/or (2) a “beta-min” condition which requires nonzero entries of $\boldsymbol{\beta}$ to not be too close to zero, and/or (3) assumptions on the design matrix such as low-pairwise correlations or placing a random distribution on the features. These types of conditions ensure that, with high probability (asymptotically), a sure screening property will hold:

that is, at any place in the method requiring a model selection step, with high probability all the relevant features (i.e., all \mathbf{X}_j with $\beta_j \neq 0$) will be selected. Under these conditions, it is possible to produce an asymptotically normal estimator for each coefficient β_j which can then yield a p-value or confidence interval which is asymptotically valid. Methods in this category include the work of Voorman, Shojaie and Witten [36], Belloni, Chernozhukov and Hansen [3] and Huang et al. [16]. A different type of result, relying on rapid decay of the error in estimating $\boldsymbol{\beta}$ in an asymptotic setting instead of requiring a sure screening property, is the debiasing work of Zhang and Zhang [42] and Javanmard and Montanari [20]; see also Voorman, Shojaie and Witten [36] for a different but related approach. Lockhart et al. [24] propose a sequential testing procedure, which moves along the Lasso path as variables are added into the model, under the assumption that the signals appear before the false positives along the Lasso path (which holds with high probability asymptotically under similar assumptions, i.e., highly sparse $\boldsymbol{\beta}$, the beta-min condition, and a well-behaved design matrix). To contrast these lines of work with our method presented here, empirically we observe reliable performance in terms of FDR control at even relatively small sample sizes, and our theoretical framework avoids strong assumptions on signal strength or sparsity level or the design matrix.

In contrast with the asymptotic setting described above, other approaches focus on finite-sample p-values and confidence intervals which do not rely on “high probability” selection events but instead are calculated by conditioning on the screened or selected model. Lee et al. [22] develops exact post-selection inference for the Lasso when applied with a fixed penalty parameter, with methodology that extends to other model selection procedures which can be expressed as linear constraints on \mathbf{y} ; Tian, Loftus and Taylor [32] extend this methodology to the setting of the square-root-Lasso, where the noise level σ is no longer assumed to be known. In contrast, Tibshirani et al. [34] propose methods for sequential selection procedures, such as the path of selected models obtained by tuning the Lasso penalty parameter; in this setting, FDR control can be obtained with sequential hypothesis testing procedures, for example [14], and the tests can be viewed as asking whether the most recent added feature has a true effect in the partial model fitted so far. Fithian, Sun and Taylor [12] develop general theory for this type of inference after conditioning on (sequential or nonsequential) selection events, termed “data carving.” In contrast to these methods, which perform inference on the outcome of specific model selection methods (such as the Lasso), the knockoff framework gives the flexibility of choosing (nearly) any test statistics for model selection.

In earlier work, Leeb and Pötscher [23] studied the impossibility of characterizing the (unconditional) distribution of post-selection estimators, and the PoSI (“post-selection inference”) framework of Berk et al. [8] proposed handling the range of possible selected models by taking a maximum over the set of possibilities. In practice, this type of method may have high power when considering extremely small submodels, but lose power when the selected model being considered grows larger.

Finally, we have called attention on Type S errors as a useful error metric for controlled variable selection.⁴ We believe this is a correct framework for the model selection problem in the high-dimensional setting, where strict sparsity assumptions may be violated whenever a reduced model is considered.

2. Knockoffs. Since this paper builds upon and extends the knockoff methodology introduced in [1], we first review the essentials of this method. The literature offers many tools for estimating a sparse linear model in a high-dimensional setting, that is, $\mathbf{y} \sim \mathcal{N}(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$ where $\mathbf{X} \in \mathbb{R}^{n \times p}$ is a fixed design matrix and $\boldsymbol{\beta}$ is believed to be sparse.⁵ For instance, we may use a forward stepwise selection method such as orthogonal matching pursuit (OMP) [26], where features are sequentially added to the model, choosing at each step the feature \mathbf{X}_j that offers the greatest reduction in the size of the residual. Alternately, we might choose a convex optimization method such as the Lasso [33], where

$$(4) \quad \hat{\boldsymbol{\beta}}_{\text{Lasso}}(\lambda) = \arg \min_{\mathbf{b} \in \mathbb{R}^p} \left\{ \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{b}\|_2^2 + \lambda \|\mathbf{b}\|_1 \right\}.$$

A critical question is this: can we estimate the number of false discoveries in the resulting coefficient vector $\hat{\boldsymbol{\beta}}_{\text{OMP}}$ or $\hat{\boldsymbol{\beta}}_{\text{Lasso}}$? The knockoff filter provides answers to questions of this kind by constructing a knockoff companion for each variable, which acts as a *control* for the possibility that we choose variable j when it is in fact null, that is, $\beta_j = 0$.

2.1. *Knockoff variables.* For each feature \mathbf{X}_j , a knockoff copy $\tilde{\mathbf{X}}_j$ is constructed to satisfy

$$(5) \quad [\mathbf{X} \quad \tilde{\mathbf{X}}]^\top [\mathbf{X} \quad \tilde{\mathbf{X}}] = \begin{bmatrix} \mathbf{X}^\top \mathbf{X} & \mathbf{X}^\top \tilde{\mathbf{X}} \\ \tilde{\mathbf{X}}^\top \mathbf{X} & \tilde{\mathbf{X}}^\top \tilde{\mathbf{X}} \end{bmatrix} = \begin{bmatrix} \boldsymbol{\Sigma} & \boldsymbol{\Sigma} - \text{diag}\{\mathbf{s}\} \\ \boldsymbol{\Sigma} - \text{diag}\{\mathbf{s}\} & \boldsymbol{\Sigma} \end{bmatrix} \triangleq \mathbf{G},$$

for some vector $\mathbf{s} \geq \mathbf{0}$.⁶ Thus the knockoffs exhibit the same correlation structure as the original variables, and the cross-correlations are also preserved, in the sense that $\mathbf{X}_j^\top \mathbf{X}_k = \tilde{\mathbf{X}}_j^\top \mathbf{X}_k$ for all $j \neq k$. While past methods in the statistics literature have considered adding fake variables in order to control false positives in regression constructed through other methods, such as by generating independent features or by permuting entries of existing features (see, e.g., Miller [25], Wu,

⁴Relatedly, in a setting where the aim is to test individual variables rather than control FDR, methods such as Lee et al. [22]’s post-selection inference for the Lasso can provide p-values testing a one-sided, that is, signed, hypothesis.

⁵Recent work by the second author and collaborators [9] treats this problem in the alternate setting where the design matrix \mathbf{X} is random, with a known distribution.

⁶If \mathbf{s} is chosen so that the Gram matrix is positive semidefinite, then $\tilde{\mathbf{X}}$ can be constructed to satisfy this matrix equality as long as $n \geq 2p$.

Boos and Stefanski [40]), the knockoffs construction is unique in that it yields an exchangeability property: for any subset S of nulls, we have

$$(6) \quad [\mathbf{X} \quad \tilde{\mathbf{X}}]_{\text{swap}(S)}^\top \mathbf{y} \stackrel{d}{=} [\mathbf{X} \quad \tilde{\mathbf{X}}]^\top \mathbf{y},$$

where the matrix $[\mathbf{X} \quad \tilde{\mathbf{X}}]_{\text{swap}(S)}$ is obtained from $[\mathbf{X} \quad \tilde{\mathbf{X}}]$ by swapping the columns \mathbf{X}_j and $\tilde{\mathbf{X}}_j$ for each $j \in S$. For example, (6) implies that $|\mathbf{X}_j^\top \mathbf{y}|$ is equally likely to be larger or smaller than $|\tilde{\mathbf{X}}_j^\top \mathbf{y}|$ for any null j .

Now imagine that we run the OMP or the Lasso on the original design augmented with knockoffs, that is, on the response \mathbf{y} and the design matrix $[\mathbf{X} \quad \tilde{\mathbf{X}}]$, rather than on the smaller matrix \mathbf{X} . Let $[\hat{\boldsymbol{\beta}} \quad \tilde{\boldsymbol{\beta}}] \in \mathbb{R}^{2p}$ denote the resulting coefficient estimates. Due to (6), we can see that $\hat{\beta}_j$ and $\tilde{\beta}_j$ have identical distributions whenever j is null and, therefore, a false positive on feature j (i.e., $\hat{\beta}_j \neq 0$) is equally likely to selecting the knockoff of feature j (i.e., $\tilde{\beta}_j \neq 0$). Therefore, the number of knockoff variables that are in our estimated coefficient vector provides an estimate of the number of false positives.

2.2. *The knockoff filter.* After constructing knockoff copies of the features, we then use them to create a feature selection rule that will achieve exact FDR control. (One can also control other types of error such as the family-wise error rate (FWER) or the k -FWER as in [18]). For each variable $j \in \{1, \dots, p\}$, we work with a general statistic W_j obeying a *sufficiency* and an *antisymmetry* property. The sufficiency property requires that W_j depends upon the design matrix only through its covariance, and upon the response \mathbf{y} through marginal correlations: for some function w_j ,

$$(7) \quad W_j = w_j([\mathbf{X} \quad \tilde{\mathbf{X}}]^\top [\mathbf{X} \quad \tilde{\mathbf{X}}], [\mathbf{X} \quad \tilde{\mathbf{X}}]^\top \mathbf{y}).$$

The antisymmetry property requires that swapping \mathbf{X}_j and $\tilde{\mathbf{X}}_j$ has the effect of changing the sign of W_j ,

$$(8) \quad w_j([\mathbf{X} \quad \tilde{\mathbf{X}}]_{\text{swap}(S)}, \mathbf{y}) = \begin{cases} w_j([\mathbf{X} \quad \tilde{\mathbf{X}}], \mathbf{y}) & j \notin S, \\ -w_j([\mathbf{X} \quad \tilde{\mathbf{X}}], \mathbf{y}) & j \in S. \end{cases}$$

In [1], a much discussed example of such a statistic concerns knots at which variables enter the Lasso path: for each feature \mathbf{X}_j and each knockoff $\tilde{\mathbf{X}}_j$, record the first time that this feature or its knockoff enters the Lasso path, that is, the largest penalty parameter value λ such that $\hat{\beta}_j \neq 0$ or $\tilde{\beta}_j \neq 0$; then set

$$(9) \quad W_j = (\text{largest } \lambda \text{ such that } \mathbf{X}_j \text{ or } \tilde{\mathbf{X}}_j \text{ enters Lasso path}) \times \begin{cases} +1 & \text{if } \mathbf{X}_j \text{ enters before } \tilde{\mathbf{X}}_j, \\ -1 & \text{if } \mathbf{X}_j \text{ enters after } \tilde{\mathbf{X}}_j. \end{cases}$$

We can analogously record the first time that a variable \mathbf{X}_j or $\tilde{\mathbf{X}}_j$ is selected by the OMP method or another forward stepwise algorithm, or any penalized least-squares method (replacing the ℓ_1 penalty of the Lasso with another penalty, such as ℓ_q). Many other statistics are of course possible, for instance $W_j = |\hat{\beta}_j| - |\tilde{\beta}_j|$.

In all of these examples, a large positive W_j indicates that \mathbf{X}_j was strongly preferred over $\tilde{\mathbf{X}}_j$ by the algorithm, while a large negative W_j indicates $\tilde{\mathbf{X}}_j$ was preferred. The sufficiency and antisymmetry properties of W , together with the exchangeability property (6) of the variables and knockoffs, ensure that $W_j > 0$ and $W_j < 0$ are equally likely for each null j , and in fact,

$$(10) \quad \#\{\text{null } j : W_j \leq -t\} \stackrel{d}{=} \#\{\text{null } j : W_j \geq t\}.$$

We use this property to choose a threshold for the test statistics W_j . Define a threshold $T > 0$ by setting⁷

$$(11) \quad T = \min \left\{ t > 0 : \frac{\#\{j : W_j \leq -t\}}{\#\{j : W_j \geq t\}} \leq q \right\},$$

where q is the target FDR level. By (10), the numerator in (11) is an (over)estimate for the number of false discoveries among all reported variables with $W_j \geq t$, and so the ratio appearing in (11) is an (over)estimate of the false discovery proportion (FDP) at the threshold t . The output of the procedure the selected model

$$(12) \quad \hat{S} = \{j : W_j \geq T\},$$

which has FDP estimated to be $\leq q$. A slightly more conservative procedure, the knockoff+ filter, instead uses the threshold

$$(13) \quad T_+ = \min \left\{ t > 0 : \frac{1 + \#\{j : W_j \leq -t\}}{\#\{j : W_j \geq t\}} \leq q \right\},$$

and accordingly sets

$$(14) \quad \hat{S} = \{j : W_j \geq T_+\}.$$

The key results in [1] state that knockoff+ controls the FDR, and the knockoff controls a slight modification of the FDR, at the level q . These results hold without any assumptions on the design or the value of the unknown regression coefficients. Crucially, we emphasize that neither the knockoff procedure nor its FDR controlling property assume any knowledge of the noise level σ .

⁷More formally, we minimize over $t \in \mathcal{W}_+$, where $\mathcal{W}_+ = \{|W_j| : |W_j| > 0, j = 1, \dots, p\}$, to avoid ever choosing $T = 0$.

3. Controlling errors of Type S. Although they are more difficult to control, we have argued that Type S errors are more meaningful for regression problems since we would like to be able to tell the direction of effect with confidence (and should avoid reporting those variables whose direction we cannot reliably determine). Consider the setting of Section 2 where $n \geq 2p$ and $\mathbf{y} \sim \mathcal{N}(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$. Then to estimate the direction of effect—the sign of β_j —note that

$$(\mathbf{X}_j - \tilde{\mathbf{X}}_j)^\top \mathbf{y} \stackrel{\perp}{\sim} \mathcal{N}(s_j \beta_j, 2s_j \sigma^2) \quad \text{for } j = 1, \dots, p,$$

where $\mathbf{s} \geq \mathbf{0}$ is from the knockoff construction (5). Hence, a natural estimate for the sign of β_j is

$$(15) \quad \widehat{\text{sign}}_j = \text{sign}((\mathbf{X}_j - \tilde{\mathbf{X}}_j)^\top \mathbf{y}).$$

With this, the knockoff filter with no adjustment whatsoever controls the directional FDR as well as the FDR. To understand intuitively why this error measure, which in principle is harder to control than FDR, is still bounded by the same method, note that for a nonnegative effect ($\beta_j \geq 0$), the probability of choosing a negative sign, that is, $\widehat{\text{sign}}_j = -1$, is in fact highest when $\beta_j = 0$; this error is less likely if $\beta_j > 0$ since in that case we would have $(\mathbf{X}_j - \tilde{\mathbf{X}}_j)^\top \mathbf{y}$ be normally distributed with a positive mean.

THEOREM 1. *Assume $\mathbf{y} \sim \mathcal{N}(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$ and fix any desired FDR level q . With the estimated direction effects (15), knockoff+ controls the directional FDR (2); that is, $\text{FDR}_{\text{dir}} \leq q$ for the selected set \widehat{S} defined in (14). If we use knockoff instead, then $\text{mFDR}_{\text{dir}} \leq q$, where mFDR_{dir} is the slightly modified version defined as*

$$(16) \quad \text{mFDR}_{\text{dir}} = \mathbb{E} \left[\frac{|\{j \in \widehat{S} : \widehat{\text{sign}}_j \neq \text{sign}(\beta_j)\}|}{|\widehat{S}| + q^{-1}} \right],$$

where the selected set \widehat{S} is defined as in (12).

This result (proved in Appendix A.1) is more powerful than Theorems 1 and 2 from [1], which established that knockoff+ controls the FDR at level q , that is, $\text{FDR} \leq q$, and knockoff controls the modified FDR, that is, the expected ratio between the number of (unsigned) false discoveries $|\{j \in \widehat{S} \text{ and } \beta_j = 0\}|$ and $|\widehat{S}| + q^{-1}$ (this denominator is as in (16)).

Our new Theorem 1 is also more delicate to prove. The reason is that when we work with Type I errors as in the original result, the probability of claiming that $\beta_j \neq 0$ is calculated under the assumption that $\beta_j = 0$. When working with Type S errors, however, there is no such calibration since there is no null, and each feature \mathbf{X}_j might have a nonzero coefficient β_j , although it may be close to zero. In that case, the test statistic W_j will have a different probability of appearing with a positive or negative sign, and we are no longer dealing with i.i.d. and unbiased signs for the null W_j 's.

Thus in the proof of Theorem 1, we will see that the result is ultimately a consequence of the following martingale inequality regarding a sequence of Bernoulli variables which have varying probabilities of success.

LEMMA 1. *Suppose that B_1, \dots, B_n are independent variables, with $B_i \sim \text{Bernoulli}(\rho_i)$ for each i , where $\min_i \rho_i \geq \rho > 0$. Let J be a stopping time in reverse time with respect to the filtration $\{\mathcal{F}_j\}$, where*

$$\mathcal{F}_j = \{B_1 + \dots + B_j, B_{j+1}, \dots, B_n\}.$$

Then

$$\mathbb{E} \left[\frac{1 + J}{1 + B_1 + \dots + B_J} \right] \leq \rho^{-1}.$$

The proof of this lemma is given in the Supplementary Material [2]. In the simpler setting where $\rho_1 = \dots = \rho_n$, this result is proved (in a slightly different form) in [1], with the proof relying heavily on the exchangeability of the B_j 's (when the B_j 's are determined by the signs of the null W_j 's, each has distribution $\text{Bernoulli}(0.5)$). The result stated here is more subtle due to the lack of exchangeability.

To see the connection between Lemma 1 and (directional) FDR control for the knockoff+ method, observe that

$$\begin{aligned} \text{FDP} &= \frac{\#\{j \text{ null} : W_j \geq T_+\}}{1 \vee \#\{j : W_j \geq T_+\}} \\ &= \frac{\#\{j \text{ null} : W_j \geq T_+\}}{1 + \#\{j \text{ null} : W_j \leq -T_+\}} \cdot \frac{1 + \#\{j \text{ null} : W_j \leq -T_+\}}{1 \vee \#\{j : W_j \geq T_+\}} \\ &\leq q \cdot \frac{\#\{j \text{ null} : W_j \geq T_+\}}{1 + \#\{j \text{ null} : W_j \leq -T_+\}}, \end{aligned}$$

where the inequality follows from the definition of T_+ (13). Now reorder the indices of the null W_j 's so that $|W_{(1)}| \geq |W_{(2)}| \geq \dots$, where $|W_{(1)}|$ is the null with the largest magnitude, $|W_{(2)}|$ the second largest, and so on, and let $B_j = \mathbb{1}_{W_{(j)} < 0}$. Let J be the index such that $|W_{(1)}| \geq \dots \geq |W_{(J)}| \geq T_+ > |W_{(J+1)}| \geq \dots$, in other words, $J = \sum_{\text{null } j} \mathbb{1}\{|W_j| \geq T_+\}$. Then for each j , we can calculate

$$W_{(j)} \geq T_+ \iff |W_{(j)}| \geq T_+ \text{ and } W_{(j)} > 0 \iff j \leq J \text{ and } B_j = 0,$$

and similarly,

$$W_{(j)} \leq -T_+ \iff |W_{(j)}| \geq T_+ \text{ and } W_{(j)} < 0 \iff j \leq J \text{ and } B_j = 1.$$

Therefore, we have

$$\begin{aligned} \frac{\#\{j \text{ null} : W_j \geq T_+\}}{1 + \#\{j \text{ null} : W_j \leq -T_+\}} &= \frac{(1 - B_1) + \dots + (1 - B_J)}{1 + B_1 + \dots + B_J} \\ &= \frac{1 + J}{1 + B_1 + \dots + B_J} - 1. \end{aligned}$$

Note that this last expression is the quantity in Lemma 1. If each W_j , for a null feature, is equally likely to be positive or negative, then the B_j 's are i.i.d. Bernoulli(0.5) variables, and we obtain FDR control at level q by applying Lemma 1 with $\rho_i = 0.5$. However, when considering directional FDR control, the β_j 's may be nonzero and so the W_j 's will typically not be symmetric; therefore, we may need to apply Lemma 1 with varying ρ_i 's. We refer to the Supplementary Material [2] for details.

4. Knockoffs in high dimensions. In high dimensions, where $p > n$, the knockoff construction is no longer possible—in fact, the Gram matrix condition (5) would be possible only if $\mathbf{s} = \mathbf{0}$, that is, if $\tilde{\mathbf{X}}_j = \mathbf{X}_j$ for each feature $j = 1, \dots, p$, so that the knockoff procedure would have zero power. In this setting, one straightforward approach would be to use part of the data to reduce the number of features, and a disjoint part of the data to run the knockoff filter. In the next section, we develop this data splitting approach, and then find that we can gain substantial power by using a subtle way of “recycling” the data.

4.1. *Feature screening and reusing data.* Consider splitting the n observations into two disjoint groups of size n_0 and $n_1 = n - n_0$, used for the purpose of first screening for a smaller set of potentially relevant features, then running a model selection procedure over this reduced list of features, respectively. We denote the two disjoint portions of the data as $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)}) \in \mathbb{R}^{n_0 \times p} \times \mathbb{R}^{n_0}$ and $(\mathbf{X}^{(1)}, \mathbf{y}^{(1)}) \in \mathbb{R}^{n_1 \times p} \times \mathbb{R}^{n_1}$, and then follow the next two steps:

- *Screening step:* using $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$, we identify a subset $\widehat{S}_0 \subset [p]$ of potentially relevant features, such that $|\widehat{S}_0| < n_1$.
- *Selection step (splitting):* ignoring any features that were discarded in the screening step, we run the knockoff procedure on the remaining data, that is, on $(\mathbf{X}_{\widehat{S}_0}^{(1)}, \mathbf{y}^{(1)})$.

This straightforward data splitting approach is a natural extension of the low-dimensional knockoff filter, and it is clear that the approach will control the directional FDR in the final model selection step as long as the screening step correctly captures all the relevant features—those with nonvanishing regression coefficients—a property often referred to as *sure screening* in the literature [11]. (False positives in the screening step, i.e., including too many null features, do not pose a problem). If the sure screening property fails, we can still obtain inference guarantees relative to the new, screened submodel; see Section 4.3 below.

There is an inherent loss of power due to the split of the data, since the model selection step uses n_1 rather than n observations. However, using disjoint data sets for the screening and selection steps is critical, since the distribution of $\mathbf{y}^{(0)}$ cannot be treated as a Gaussian linear model *after* the screening step has taken place. The screening step is a function of the random variable $\mathbf{y}^{(0)}$, and so the following

modification of the selection step *would not control the FDR*: ignoring any features that were discarded in the screening step, we run the knockoff procedure on the full data set, that is, on $(\mathbf{X}_{\widehat{S}_0}, \mathbf{y})$. The loss of FDR control is not merely theoretical: a null feature \mathbf{X}_j that is chosen by the screening step is generally more likely to appear as a false positive when running the knockoff filter, leading to a much higher FDR.

In light of this, we propose two mechanisms for increasing power relative to the data splitting procedure described so far:

1. *data recycling*, where the first portion of the split data can be reused to some extent without losing *any* of the guaranteed FDR control,
2. and *signed statistics*, where we can decide adaptively to focus our search on a positive effect only or a negative effect only.

4.1.1. *Increasing power with data recycling.* Surprisingly, data recycling retains the FDR control properties of the split-data procedure, while raising power substantially to approach the sensitivity of the full-data procedure.

- *Selection step (with recycling)*: we begin by constructing a knockoff matrix on the remaining data $\widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}$ for $\mathbf{X}_{\widehat{S}_0}^{(1)}$ just as we would in the data splitting version. The difference now is that we concatenate the original design matrix on the first n_0 observations with the knockoff matrix on the next n_1 observations,

$$(17) \quad \widetilde{\mathbf{X}}_{\widehat{S}_0} = \begin{bmatrix} \mathbf{X}_{\widehat{S}_0}^{(0)} \\ \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)} \end{bmatrix}.$$

(Note that on the first part of the split data, knockoffs are exact copies.) We then run the knockoff filter on the entire data set with all n samples, using the design matrix $\mathbf{X}_{\widehat{S}_0}$, the knockoff matrix $\widetilde{\mathbf{X}}_{\widehat{S}_0}$, and the original response \mathbf{y} .

The term “recycling” refers to the way that we incorporate the data $(\mathbf{X}_{\widehat{S}_0}^{(0)}, \mathbf{y}^{(0)})$, which was already used in the screening step, for the selection step. Here, we think of the mechanism for choosing \widehat{S}_0 as being completely arbitrary as opposed to having some kind of “preregistered” screening method. For this reason, we need to treat $\mathbf{y}^{(0)}$ as a fixed vector (or rather, to condition on its value) in the construction and analysis of the knockoff filter for the selection step (see Section 4.1.3 for additional discussion about this point).

Of course, knockoffs with data splitting are actually a special case of knockoffs with data recycling. Any knockoff statistics W_j that are a function of $\mathbf{X}_{\widehat{S}_0}^{(1)}, \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}$, and $\mathbf{y}^{(1)}$ (i.e., data splitting) can trivially be expressed as a function of $\mathbf{X}_{\widehat{S}_0}, \widetilde{\mathbf{X}}_{\widehat{S}_0}$ and \mathbf{y} by simply ignoring the first n_0 many data points.

Increased power. Why do we expect the knockoff method to exhibit higher power when using data recycling rather than data splitting, when we see in (17) that the knockoff matrix differs from the original features only on the second part of the split data? Indeed, later on in Section 5 we will see empirically that the power gain can be quite substantial.

To understand this phenomenon, we return to the mechanism of the knockoff filter itself. Recall that we construct a statistic W_j for each feature/knockoff pair $(\mathbf{X}_j, \tilde{\mathbf{X}}_j)$, where $|W_j|$ is large if *either* \mathbf{X}_j or $\tilde{\mathbf{X}}_j$ appear highly significant, and $\text{sign}(W_j)$ indicates which of the two appears most significant. For instance, if we are using a forward stepwise selection method as described in Section 2, $|W_j|$ will be large if either \mathbf{X}_j or $\tilde{\mathbf{X}}_j$ was one of the first variables to enter the model; we will have $W_j > 0$ if \mathbf{X}_j enters before $\tilde{\mathbf{X}}_j$, and $W_j < 0$ if \mathbf{X}_j enters after $\tilde{\mathbf{X}}_j$.

If the early portion of the knockoff path (where we order the features in order of magnitudes $|W_j|$) has many positive W_j 's and few negative W_j 's, the stopping rule (11) will allow us to choose a threshold T that is not too large, and we will make many rejections. If instead the high-magnitude $|W_j|$'s have many negative signs, though, we will be forced to stop early, choosing a high threshold T and making few rejections. For any null feature j (i.e., $\beta_j = 0$), the sign of W_j is equally likely to be negative as positive, so in order to obtain high power, we need two things:

1. Good separation in the feature ordering—most of the features that appear early in the path, that is, most j with $|W_j|$ large, are nonnulls;
2. Good power in the nonnull signs—for most nonnulls, the statistics have a positive sign, $W_j > 0$.

Next, we ask how data splitting and data recycling compare in light of these two considerations. For the second, both methods suffer from power loss relative to an unscreened method—specifically, obtaining $W_j > 0$ rather than $W_j < 0$ depends on the model selection method being able to distinguish between feature j and its knockoff copy. For data splitting, by the sufficiency (7) and antisymmetry (8) properties, we see that if

$$(18) \quad \mathbf{X}_j^{(1)\top} \mathbf{y}^{(1)} \quad \text{and} \quad \tilde{\mathbf{X}}_j^{(1)\top} \mathbf{y}^{(1)}$$

are equal in distribution, then $\text{sign}(W_j)$ is equally likely to be $+1$ or -1 ; in other words, it is only if these two inner products have substantially *different* distributions, that we can hope to see $\text{sign}(W_j) = +1$ with high probability. For data recycling, the same statement holds, except instead we compare the inner products

$$(19) \quad \mathbf{X}_j^\top \mathbf{y} = \mathbf{X}_j^{(0)\top} \mathbf{y}^{(0)} + \mathbf{X}_j^{(1)\top} \mathbf{y}^{(1)} \quad \text{and} \quad \tilde{\mathbf{X}}_j^\top \mathbf{y} = \mathbf{X}_j^{(0)\top} \mathbf{y}^{(0)} + \tilde{\mathbf{X}}_j^{(1)\top} \mathbf{y}^{(1)}.$$

Examining (18) versus (19), we see that these two comparisons are equivalent, as (19) simply adds the same constant (i.e., $\mathbf{X}_j^{(0)\top} \mathbf{y}^{(0)}$) to each term of (18). In other words, these two tasks are equally difficult.

However, for the first component of power—that is, good separation in the feature ordering—we can expect that knockoffs with data recycling will have the

advantage, since the first n_0 data points carry substantial information separating signals from nulls. Of course, if some nonnull feature j has strong correlation with the response, then the information from the first n_0 data points will push *both* \mathbf{X}_j and $\tilde{\mathbf{X}}_j$ towards the top of the path—since $\tilde{\mathbf{X}}_j^{(0)} = \mathbf{X}_j^{(0)}$ by construction—but the second portion of the data, where $\tilde{\mathbf{X}}_j^{(1)} \neq \mathbf{X}_j^{(1)}$, will hopefully enable \mathbf{X}_j to enter the path before $\tilde{\mathbf{X}}_j$ (i.e., $W_j > 0$).

To summarize, the knockoff filter’s power relies on the nonnull W_j ’s ability to *appear early* and *with positive sign* along the path of statistics; data recycling helps the nonnull W_j ’s appear early, thus increasing power.

4.1.2. *Increasing power with signed statistics.* Since the knockoff filter with data splitting or with data recycling treats the first portion of the data, $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$, as fixed once the screening step has been performed, we are free to use this data in any way we wish, to try to gather information about the true signals. While the screening step identifies the indices of the variables which are likely to contain signals, we can also use this step to identify probable signs, that is, if feature $\mathbf{X}_j^{(0)}$ is selected in the screening step, which is run on the data set $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$, we can at the same time specify the sign of the estimated effect of this feature, denoted $\widehat{\text{sign}}_j^{(0)}$. We then use this information to look specifically for an effect consistent with this sign in the second phase of our procedure. This will in general increase power, since we are extracting more information from the first part of the data before running the selection procedure on the second part, $(\mathbf{X}^{(1)}, \mathbf{y}^{(1)})$.

Example: Penalized least-squares. To illustrate the use of the sign information in practice, suppose that we are using a penalized least-squares regression method for computing the statistics W_j , for instance, as described in Section 2, the W_j ’s may indicate the time at which features \mathbf{X}_j and $\tilde{\mathbf{X}}_j$ entered the Lasso path (or the path of some other sequential selection procedure), or may be given by $W_j = |\hat{\beta}_j| - |\tilde{\beta}_j|$ where $[\hat{\beta} \quad \tilde{\beta}]$ is the solution to a penalized least squares optimization problem, with some penalty function $P(\mathbf{b})$ (with the Lasso as a special case if $P(\mathbf{b}) = \lambda \|\mathbf{b}\|_1$). In either case, if we are not using sign information, we might begin by computing the penalized least-squares solution $[\hat{\beta} \quad \tilde{\beta}]$ via either

$$\min_{\mathbf{b} \in \mathbb{R}^{2|\tilde{S}_0|}} \frac{1}{2} \|\mathbf{y}^{(1)} - [\mathbf{X}_{\tilde{S}_0}^{(1)} \quad \tilde{\mathbf{X}}_{\tilde{S}_0}^{(1)}] \mathbf{b}\|_2^2 + P(\mathbf{b})$$

for data splitting, or

$$\min_{\mathbf{b} \in \mathbb{R}^{2|\tilde{S}_0|}} \frac{1}{2} \|\mathbf{y} - [\mathbf{X}_{\tilde{S}_0} \quad \tilde{\mathbf{X}}_{\tilde{S}_0}] \mathbf{b}\|_2^2 + P(\mathbf{b})$$

for data recycling. Instead, to make use of the sign information gathered at the screening phase, we can consider a sign-restricted version of this penalized least-squares problem: in the data splitting version we would be interested in the solution

to

$$\begin{aligned} \min_{\mathbf{b} \in \mathbb{R}^{2|\widehat{S}_0|}} & \frac{1}{2} \|\mathbf{y}^{(1)} - [\mathbf{X}_{\widehat{S}_0}^{(1)} \quad \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}] \mathbf{b}\|_2^2 + P(\mathbf{b}) \\ \text{s.t.} & \quad b_j \cdot \text{sign}_j^{(0)} \geq 0 \\ & \quad b_{j+|\widehat{S}_0|} \cdot \text{sign}_j^{(0)} \geq 0 \end{aligned}$$

and similarly for the recycling version. In other words, we are running the same penalized least-squares optimization, but with the added restriction that we will only select the j th feature or j th knockoff feature if its estimated effect direction agrees with the sign information gathered at the screening stage.

Increased power. Why should we expect the knockoff method to exhibit higher power when leveraging the sign information, $\text{sign}_j^{(0)}$ for each selected feature $j \in \widehat{S}_0$? Intuitively, we think of the screening step as choosing from p possible hypotheses, where hypothesis H_j is described by the question, “Does \mathbf{X}_j appear in the true model?” However, we can instead frame this step as considering twice as many hypotheses: hypothesis H_j^+ is the question, “Does \mathbf{X}_j appear in the true model with a positive effect?” and hypothesis H_j^- is the question “Does \mathbf{X}_j appear in the true model with a negative effect?” With this in mind, if the screening step fits a model to the first part of the data $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$, and this fitted model includes the j th feature with a positive coefficient, then this indicates that hypothesis H_j^+ is more likely to contain a signal. Therefore, it makes sense that in our selection step we should focus our attention on hypothesis H_j^+ . Of course, it is possible that the true effect is in fact negative, but since this is less likely, we would generally be increasing the noise if we give equal attention to hypotheses H_j^+ and H_j^- .

4.1.3. *Relation to existing work.* We pause here to compare our data recycling technique to the data carving methods of Fithian, Sun and Taylor [12]. In that work, as in ours, some part of the data $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$ is used to perform an initial screening step, and is then reused for inference along with the remaining data $(\mathbf{X}^{(1)}, \mathbf{y}^{(1)})$. However, the mechanism behind reusing the first part of the data is highly different. In Fithian, Sun and Taylor [12]’s work, the first part of the response $\mathbf{y}^{(0)}$ is reused by leveraging the “remaining randomness” after the screening step: $\mathbf{y}^{(0)}$ is regarded as a random vector, with distribution given by conditioning on the outcome of the screening step, that is, $(\mathbf{y}^{(0)} \mid \widehat{S}_0)$. In contrast, in our “recycling” procedure, $\mathbf{y}^{(0)}$ can be used in an arbitrary way for the screening step, and so there is no “remaining randomness” in $\mathbf{y}^{(0)}$. Instead, we reuse $\mathbf{y}^{(0)}$ by treating it as a fixed vector; $\mathbf{y}^{(1)}$ is the only variable treated as random for the purpose of our FDR control results.

More broadly, the knockoff filter with either data splitting or data recycling, can be considered as similar in flavor to the “screen and clean” methodology of Wasserman and Roeder [37], where an initial screening step (performed via a high-dimensional Lasso, with a tuning parameter chosen by validation), is followed by

an inference step on an independent portion of the data, with inference performed via a least-squares regression. This methodology has been applied to genome-wide association studies (GWAS) by Wu et al. [39]. At a high level, our method follows this same overall framework of screening for a low-dimensional submodel, then using new data for low-dimensional inference; however, we will gain power by “recycling” the first part of the data, and more significantly, by using the knockoff filter rather than least-squares for the inference step, which in the low-dimensional setting gives substantial gains in power as it leverages the sparse structure of the true model [1].

4.2. *Sure screening and directional FDR control.* In order to demonstrate the correctness of the selection step with recycling, we first assume that the screening step is highly likely to find all true signal, that is, $\widehat{S}_0 \supseteq \text{support}(\boldsymbol{\beta})$ with high probability. Note that we do not assume that the screening step *exactly* selects the true support—while this type of sparsistency result holds for the Lasso in the asymptotic setting (see, e.g., [43]), in practice, even with ideal simulated data, it is generally the case that we cannot perfectly separate true signals from false positives using the limited data available, unless the signal is extremely strong and extremely sparse [30]. Instead, this assumption merely requires that there is some sufficiently liberal screening procedure that is likely to capture all the true signals, along with many false positives. From a theoretical standpoint, this property is known to hold under conditions far weaker than those needed for exact recovery of the true support [11, 41].

Suppose we fix some chosen method for the screening step, which is constrained only in that it must be a function of the first portion of the data, $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$. Define the sure screening event as

$$\mathcal{E} = \{\widehat{S}_0 \supseteq \text{support}(\boldsymbol{\beta}) \text{ and } |\widehat{S}_0| \leq n_1/2\},$$

and note that $\mathbb{1}_{\mathcal{E}}$ is a function of $\mathbf{y}^{(0)}$ when we treat \mathbf{X} and $\boldsymbol{\beta}$ as fixed. Conditioning on this event, directional FDR control of the selection step holds by the following theorem (proved in Appendix A.1).

THEOREM 2. *Suppose $\mathbf{y} \sim \mathcal{N}(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$. Then the knockoff procedure (using either data splitting or data recycling), with estimated signs $\widehat{\text{sign}}_j$ as in (15), controls the modified directional FDR at the level*

$$\mathbb{E}[\text{mFDR}_{\text{dir}} \mid \mathcal{E}] \leq q,$$

while if knockoff+ is used, then the directional FDR is controlled as

$$\mathbb{E}[\text{FDR}_{\text{dir}} \mid \mathcal{E}] \leq q.$$

In particular, if \mathcal{E} occurs with probability near 1, then the various forms of the directional FDR are controlled even without conditioning on \mathcal{E} , by reformulating the results as $\text{FDR}_{\text{dir}} \leq q + \mathbb{P}\{\mathcal{E}^c\}$ in the case of knockoff+ and $\text{mFDR}_{\text{dir}} \leq q + \mathbb{P}\{\mathcal{E}^c\}$ for knockoff.

4.3. *Directional FDR control in the reduced model.* If the screening step misses variables from the model, a bias is introduced which creates a fundamental difficulty. To discuss this, recall that knockoffs are constructed to guarantee the crucial pairwise exchangeability for the nulls (6). Among other things, we have seen that exchangeability has the consequence that null variables are equally likely to be selected as their knockoff companions. When nonzero effects have not been screened, this may no longer be the case. Concretely, imagine that \mathbf{X}_k has a nonzero effect but was not identified in the screening step. Then since the knockoff construction used only the screened set \widehat{S}_0 , we cannot guarantee that $\mathbf{X}_j^\top \mathbf{X}_k = \widetilde{\mathbf{X}}_j^\top \mathbf{X}_k$ for $j \in \widehat{S}_0$, which in turn may lead to

$$\mathbf{X}_j^\top \mathbf{y} \stackrel{d}{\neq} \widetilde{\mathbf{X}}_j^\top \mathbf{y}$$

even when j is a null feature, that is, $\beta_j = 0$. For example, if \mathbf{X}_j were strongly correlated with a true signal \mathbf{X}_k , then we could easily imagine that \mathbf{X}_j would be more likely to be selected than its knockoff $\widetilde{\mathbf{X}}_j$. In general, and in contrast to the lower dimensional setting, the bias implies that when we count the number of knockoffs selected at some threshold level, we cannot be sure that this number is a good (over)estimate of the number of selected nulls. (As explained before, however, it may be desirable to select \mathbf{X}_j in this case since it is a proxy for \mathbf{X}_k .)

We now reframe this issue in terms of the partial regression coefficients, that is, we will perform inference on the coefficients in the reduced model defined by the subset of features \widehat{S}_0 selected in the screening step, a point of view in line with much of the recent literature discussed earlier [8, 12, 13, 22]. Thus we are interested in the partial regression coefficients—the coefficients of \mathbf{y} regressed onto only the selected features. Throughout this section, we will consider a random design model, and will define the population-level partial regression coefficients as

$$(20) \quad \boldsymbol{\beta}^{\text{partial}} = (\mathbf{X}_{\widehat{S}_0}^{(1)\top} \mathbf{X}_{\widehat{S}_0}^{(1)})^{-1} \mathbf{X}_{\widehat{S}_0}^{(1)\top} \mathbb{E}[\mathbf{y}^{(1)} \mid \mathbf{X}_{\widehat{S}_0}^{(1)}; \mathbf{X}^{(0)}, \mathbf{y}^{(0)}],$$

that is, the *expected* coefficients when $\mathbf{y}^{(1)}$ regressed on the selected features $\mathbf{X}_{\widehat{S}_0}^{(1)}$. In other words, this is the quantity for which the least-squares solution in the reduced model is an unbiased estimate. Here, expectation is taken conditional on the first part of the data (i.e., $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$, the data used for the screening step), and on the selected features $\mathbf{X}_{\widehat{S}_0}^{(1)}$.⁸

⁸This definition is slightly different from the notion of partial regression coefficients found in much of the selective inference literature discussed here, that is, [8, 12, 13, 22], which treats the design \mathbf{X} as fixed and studies $\boldsymbol{\beta}^{\text{partial}} = (\mathbf{X}_{\widehat{S}_0}^\top \mathbf{X}_{\widehat{S}_0})^{-1} \mathbf{X}_{\widehat{S}_0}^\top \mathbb{E}[\mathbf{y} \mid \mathbf{X}]$. We could alternately consider a population-level version of these coefficients, $\boldsymbol{\beta}^{\text{partial}} = \mathbb{E}[\mathbf{X}_{\widehat{S}_0}^\top \mathbf{X}_{\widehat{S}_0}]^{-1} \mathbb{E}[\mathbf{X}_{\widehat{S}_0}^\top \mathbf{y}]$. The difference is between these various definitions is typically not substantial, and indeed under many random models for the design \mathbf{X} , can be proved to be vanishingly small.

FDR control with a Gaussian design. In this paper, we give results for a random feature model \mathbf{X} , which assumes that the rows $\mathbf{X}_{[i]}$ of \mathbf{X} are drawn i.i.d. from a multivariate Gaussian distribution,

$$(21) \quad \mathbf{X}_{[i]} \stackrel{\text{i.i.d.}}{\sim} \mathcal{N}(\mathbf{v}, \Psi);$$

here, the parameters (i.e., the mean vector $\mathbf{v} \in \mathbb{R}^p$ and the covariance matrix $\Psi \in \mathbb{R}^{p \times p}$) are arbitrary and completely unknown to us. The following theorem (proved in Appendix A.1) guarantees directed FDR control in this setting.

THEOREM 3. *Assume that the rows of \mathbf{X} follow (21), where $\mathbf{v} \in \mathbb{R}^p$ and $\Psi \in \mathbb{R}^{p \times p}$ are arbitrary and unknown. Define the expected partial regression coefficients β^{partial} as in (20). Then the knockoff-with-recycling procedure controls the modified directional FDR at the level*

$$\text{mFDR}_{\text{dir}} \leq q,$$

while if knockoff+ is used, then the directional FDR is controlled as

$$\text{FDR}_{\text{dir}} \leq q;$$

here, the estimated signs $\widehat{\text{sign}}_j$ are defined as in (15), and a false discovery is any $j \in \widehat{S}$ such that $\widehat{\text{sign}}_j \neq \text{sign}(\beta_j^{\text{partial}})$.

While the Gaussian design assumption is strong, we believe that FDR control would be maintained across a far broader range of models. We emphasize that we do not assume knowledge of the mean and covariance parameters of the Gaussian distribution of \mathbf{X} , and do not need to estimate these parameters in the method; in fact, the Gaussian assumption is used only for the technical details of the proof, to show that we can model the response $\mathbf{y}^{(1)}$ as

$$\mathbf{y}^{(1)} = \mathbf{X}_{\widehat{S}_0}^{(1)} \beta_{\widehat{S}_0} + \mathbf{X}_{\widehat{S}_0^c}^{(1)} \beta_{\widehat{S}_0^c} + \boldsymbol{\epsilon}^{(1)} = \mathbf{X}_{\widehat{S}_0}^{(1)} \beta_{\widehat{S}_0}^{\text{partial}} + \boldsymbol{\epsilon}'^{(1)},$$

where $\boldsymbol{\epsilon}'^{(1)}$ is i.i.d. Gaussian noise (with a larger variance than the original noise $\boldsymbol{\epsilon}$, due to the noise added from the true signals missed by the screening step). This is where the fact that knockoffs do not need the value of the noise level (in low dimensions) is immensely useful: the knockoff method will provide valid inference no matter the value of the “new noise level.”

5. Simulations. We now examine the performance of our method on simulated high-dimensional data sets, and compare to several other techniques for inference in low-dimensional and high-dimensional regression. One data set below uses real SNP features.

5.1. *Data.* To simulate data that exhibits challenging features that we might expect to see in practice, we work in two different settings for the matrix \mathbf{X} of covariates:

- *AR matrix* We consider an autoregressive model, where \mathbf{X}_j is highly correlated with \mathbf{X}_{j+1} and \mathbf{X}_{j-1} , and the correlation decays for columns farther apart in the design matrix. We work with sample size $n = 2000$ and create $p = 2500$ features by drawing the rows of \mathbf{X} independently from a $\mathcal{N}(\mathbf{0}, \Sigma)$ model, where $\Sigma_{jk} = \rho^{|j-k|}$ for a correlation parameter $\rho \in \{0, 0.25, 0.5, 0.75\}$. We then normalize the columns of \mathbf{X} to have unit norm.
- *GWAS matrix* The matrix \mathbf{X} is here taken from a genome-wide association study (GWAS); see Section 6 for further information regarding this data set. We begin with SNP data from the first three chromosomes, then cluster the SNPs using pairwise correlations and choose only one SNP from each cluster, in such a way that the resulting pairwise correlation between any two SNPs in the remaining data set is at most 0.5, and finally regress out the top five principal components (Section 6 discusses the rationale behind this extra step). The size of the data set is $p = 12,752$ features (SNPs) and sample size $n = 4682$.

We then generate the coefficient vector β by choosing $k_0 = 50$ many coordinates at random to be the locations of the strong signals, and k_1 many coordinates to be the location of the weak signals, where $k_1 = 250$ for the AR setting and $k_1 = 1250$ for the GWAS setting. Strong signals are chosen from $\{\pm 4.5\}$ for the AR setting or from $\{\pm 5\}$ for the GWAS setting (the higher amplitude is due to the larger size of the problem). Weak signals are drawn as $\beta_j \sim \mathcal{N}(0, 0.5)$ for both settings. The response is then generated as $\mathbf{y} = \mathbf{X}\beta + \epsilon$, where ϵ has independent $\mathcal{N}(0, 1)$ entries. In each setting (AR design with each value of ρ , and the GWAS design), the design matrix \mathbf{X} and the coefficient vector β are generated only once, while the noise vector ϵ is generated independently for each trial. Results are averaged over 100 trials. All procedures are run with target FDR level $q = 0.2$.

5.2. *Methods.* Our simulated experiments compare the following methods:

- The Benjamini–Hochberg (BH) [6] procedure (applied to the least-squares regression coefficients, after screening for a low-dimensional submodel).
- Exact selective inference for the square-root Lasso, developed by Tian, Loftus and Taylor [32]—this method builds on the earlier work of Lee et al. [22] developing selective inference for the Lasso.
- Knockoff filter with data splitting or data recycling (applied after screening for a low-dimensional submodel).

We now give details for the implementation of each method.

Benjamini–Hochberg. We use n_0 many data points for an initial screening step, where $n_0 = 750$ for the AR setting and $n_0 = 1800$ for the GWAS setting. The screening step is carried out by solving the Lasso across the entire range of λ values, from $\lambda = \infty$ to $\lambda = 0$. We then choose a screened set \widehat{S}_0 consisting of the first k_{\max} many features entering the path, where $k_{\max} = 450$ for the AR setting and $k_{\max} = 720$ for the GWAS setting, and then record the sign $s_j \in \{\pm 1\}$ for feature $j \in \widehat{S}_0$ (i.e., the sign of its estimated coefficient, at the time when it first enters the Lasso path). We next use the remaining $n_1 = n - n_0$ data points for inference. For both settings, we are using roughly 35–40% of the total sample size n as the sample size n_0 for the screening stage, and setting the number of screened features k_{\max} be around 25% of the remaining sample size n_1 . (These ratios appear to give a favorable balance between finding a sufficiently good screened set at the first stage, and retaining a large enough sample size for powerful inference in the second stage. It is of great interest to study this question in future work to see whether we can determine theoretically the optimal split for making the most discoveries.)

We calculate the least-squares regression coefficients $\widehat{\beta}_{\text{LS}} = (\mathbf{X}_{\widehat{S}_0}^{(1)\top} \mathbf{X}_{\widehat{S}_0}^{(1)})^{-1} \times \mathbf{X}_{\widehat{S}_0}^{(1)\top} \mathbf{y}^{(1)}$, and compute t-scores

$$t_j = \frac{(\widehat{\beta}_{\text{LS}})_j}{\widehat{\sigma} \cdot \sqrt{(\mathbf{X}_{\widehat{S}_0}^{(1)\top} \mathbf{X}_{\widehat{S}_0}^{(1)})_{jj}^{-1}}}$$

for each $j \in \widehat{S}_0$. Above, $\widehat{\sigma}^2$ is the classical estimate of variance,

$$\widehat{\sigma}^2 = \frac{\|\mathbf{y}^{(1)} - \mathbf{X}_{\widehat{S}_0}^{(1)} \widehat{\beta}_{\text{LS}}\|_2^2}{n_1 - |\widehat{S}_0|}.$$

Under the sure screening property (where \widehat{S}_0 contains all true signals, as discussed in Section 4.2), for each null feature j , t_j would follow a t -distribution with $n_1 - |\widehat{S}_0|$ degrees of freedom. We thus convert each t-score to a one-sided p-value P_j (with the side chosen based on the sign s_j from the screening step). Finally, we run the Benjamini–Hochberg procedure [6] at the level q to select a model.

Selective inference for the square-root Lasso. The square-root Lasso [4] is given by

$$\widehat{\beta} = \arg \min_{\mathbf{b} \in \mathbb{R}^p} \{ \|\mathbf{y} - \mathbf{X}\mathbf{b}\|_2 + \lambda \cdot \|\mathbf{b}\|_1 \}$$

(22) where $\lambda = \kappa \cdot \mathbb{E}_{\mathbf{g} \sim \mathcal{N}(0, \mathbf{I}_n)} \left[\frac{\|\mathbf{X}^\top \mathbf{g}\|_\infty}{\|\mathbf{g}\|_2} \right]$

for a constant κ typically chosen in $[0.5, 1]$. The solution $\widehat{\beta}$ is in fact *exactly* equal to the solution to the Lasso (4) at some different penalty parameter value; the

benefit of the square-root Lasso is that the penalty parameter λ can be chosen without reference to the unknown noise level σ^2 .

Tian, Loftus and Taylor [32] derive exact post-selection confidence intervals and p-values for the coefficients β_j , after conditioning on the selection event, that is, after observing the signed support of $\hat{\beta}$, the solution to the square-root Lasso, based on the truncated t distribution, without knowledge of the noise level σ . (This work builds on the polytope method developed in Lee et al. [22], Tibshirani et al. [34] for the Lasso and other procedures.) An implementation of this method, in Python, is available in [31], where the inference is performed approximately by estimating the noise σ and then using truncated normal distributions, as in Lee et al. [22]’s inference for the Lasso. We use this available code for our experiments, in order to calculate a one-sided selective p-value P_j for each feature selected by the square-root Lasso (based on the sign s_j with which feature j was selected). We then apply the Benjamini–Hochberg procedure at level q to the resulting set of p-values in order to select a final model. The only tuning parameter in this method is the constant κ appearing in (22). We test $\kappa \in \{0.5, 0.6, \dots, 1\}$ and find the highest power at $\kappa = 0.6$; only this value of κ is displayed in our results.⁹

Knockoff filter. The screening step is carried out exactly as for the BH procedure, with the same values of n_0 and k_{\max} , resulting in a submodel \widehat{S}_0 and signs s_j for each $j \in \widehat{S}_0$. We next use the remaining $n_1 = n - n_0$ data points for the knockoff filter, using sign information as described in Section 4.1.2, and either data splitting or data recycling. Our statistics W_j are formed using the square-root Lasso. Specifically, for data splitting, we solve a sign-restricted square-root Lasso,

$$\begin{aligned}
 (\widehat{\beta} \quad \widetilde{\beta}) = \arg \min_{\substack{\mathbf{b} \in \mathbb{R}^{2|\widehat{S}_0|} \\ b_j \cdot s_j \geq 0 \text{ for all } j = 1, \dots, p \\ b_{j+|\widehat{S}_0} \cdot s_j \geq 0 \text{ for all } j = 1, \dots, p}} \{ \|\mathbf{y}^{(1)} - [\mathbf{X}_{\widehat{S}_0}^{(1)} \quad \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}] \mathbf{b}\|_2 \\
 + \lambda \cdot \|\mathbf{b}\|_1 \},
 \end{aligned}
 \tag{23}$$

where the penalty parameter λ is defined as in (22) (but with $[\mathbf{X}_{\widehat{S}_0}^{(1)} \quad \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}]$ in place of \mathbf{X}). For data recycling, we optimize the same problem but with $[\mathbf{X}_{\widehat{S}_0} \quad \widetilde{\mathbf{X}}_{\widehat{S}_0}]$ and \mathbf{y} in place of $[\mathbf{X}_{\widehat{S}_0}^{(1)} \quad \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}]$ and $\mathbf{y}^{(1)}$. In each case, we then define $W_j = |\widehat{\beta}_j| - |\widetilde{\beta}_j|$ for each $j \in \widehat{S}_0$, and then apply the knockoff filter.

We run each method with $\kappa \in \{0.1, 0.2, \dots, 1\}$ and find the highest power at $\kappa = 0.5$ for data splitting and $\kappa = 0.7$ for data recycling; only these values of κ are displayed in our results.

⁹The software [31] sets λ as κ times the 95th percentile of $\frac{\|\mathbf{X}^\top \mathbf{g}\|_\infty}{\|\mathbf{g}\|_2}$ in lieu of what is displayed in (22).

5.3. *Results.* The knockoff filter and the BH procedure both work within the reduced model defined by the subset of features \widehat{S}_0 chosen at the initial screening step; the inference guarantees of these methods are with reference to this large submodel. On the other hand, the p-values computed by Tian, Loftus and Taylor [32]’s method are calculated with reference to a relatively small submodel, given by features selected by the square-root Lasso. Therefore, in comparing these three types of methods against each other, it does not make sense to compare their FDRs over the partial regression coefficients, since these partial regression coefficients are defined within vastly different submodels. Instead, for each correlation setting and for each method tested, we compute the following measures of performance relative to the full model (averaged over 100 trials):

- False discovery rate and directional false discovery rate relative to the true full model: the average proportion of the selected features that were selected incorrectly (for FDR), or selected with an incorrect sign (for FDR_{dir}); and
- Power relative to the true full model: the average proportion of the true support, the $k_0 + k_1$ many features in $\text{support}(\boldsymbol{\beta})$, which is selected by the method. Since this true support is split into strong and weak signals, and the weak signals are barely detectible, we also report the “restricted power,” that is, the power restricted to only the size- k_0 set of strong signals.

For all methods, the full-model FDR and directional FDR, along with the power and the restricted power, are displayed in Table 1 for both the AR design and the GWAS design settings. For both settings, we see that the methods all successfully control the FDR at the desired level, and indeed some are conservative, except that for large correlation values ρ in the AR setting, the least squares + BH method does show overly high FDR (however, this can be explained by the fact that, for a highly correlated design, the FDR in the reduced model, i.e., for the partial regression coefficients $\boldsymbol{\beta}^{\text{partial}}$, will be substantially different from the FDR in the full model). Across all settings and all methods, the directional FDR is slightly higher than the FDR, since the k_1 weak signals are extremely close to zero and, if selected, can easily be selected with the wrong sign. The (nonrestricted) power is quite low for all methods, since the k_1 weak signals are extremely difficult to distinguish from zero. Turning instead to the restricted power, indicating each method’s ability to detect strong signals, we see that the knockoff filter with data recycling shows somewhat higher power than the other methods, and in particular, we see a clear gain when using data recycling rather than data splitting, as expected.

6. Real data experiments. We next implement our high-dimensional knock-off method on a genome wide association study (GWAS) data set. The data comes from the Northern Finland Birth Cohort 1966 (NFBC1966) [19, 28], made available through the dbGaP database (accession number phs000276.v2.p1). Here, 5402 SNP arrays from subjects born in Northern Finland in 1966 are recorded, as well as a number of phenotype variables measured when the subjects were 31 years old.

TABLE 1
Simulation results. (Target FDR level is 20%.) Estimated standard errors are in parentheses

		Least-sq. + BH	Selective inf. [32]	Knockoff/ split	Knockoff/ recycle
AR design ($\rho = 0$)	FDR	17.14 (0.7)	6.27 (0.5)	7.52 (0.5)	12.65 (0.8)
	Dir. FDR	17.48 (0.7)	6.36 (0.5)	7.63 (0.5)	12.84 (0.8)
	Power	9.47 (0.2)	9.67 (0.3)	8.85 (0.2)	11.08 (0.2)
	Restr. power	43.42 (1.1)	56.34 (2.0)	51.58 (1.0)	62.60 (1.0)
AR design ($\rho = 0.25$)	FDR	16.93 (0.8)	6.46 (0.5)	8.72 (0.6)	12.11 (0.7)
	Dir. FDR	17.31 (0.8)	6.51 (0.5)	8.84 (0.6)	12.27 (0.7)
	Power	9.17 (0.2)	8.97 (0.4)	8.60 (0.2)	10.60 (0.2)
	Restr. power	41.66 (0.9)	51.98 (2.1)	49.96 (1.3)	60.56 (0.9)
AR design ($\rho = 0.5$)	FDR	19.42 (0.9)	7.71 (0.6)	9.62 (0.8)	13.51 (1.0)
	Dir. FDR	19.66 (0.9)	7.73 (0.6)	9.78 (0.8)	13.74 (1.0)
	Power	7.81 (0.2)	8.75 (0.3)	6.65 (0.3)	8.96 (0.3)
	Restr. power	34.60 (1.0)	50.54 (1.9)	38.22 (1.7)	50.36 (1.4)
AR design ($\rho = 0.75$)	FDR	29.99 (1.2)	13.72 (1.6)	14.15 (1.4)	18.25 (1.4)
	Dir. FDR	30.43 (1.2)	14.02 (1.6)	14.28 (1.4)	18.40 (1.4)
	Power	5.21 (0.2)	3.63 (0.3)	3.43 (0.3)	4.99 (0.3)
	Restr. power	20.10 (0.8)	21.00 (1.7)	19.38 (1.7)	27.90 (1.7)
GWAS design	FDR	17.97 (1.15)	8.62 (0.81)	12.28 (0.88)	15.05 (1.15)
	Dir. FDR	18.36 (1.18)	9.04 (0.80)	12.56 (0.90)	15.41 (1.13)
	Power	1.25 (0.05)	0.90 (0.05)	1.11 (0.04)	1.61 (0.05)
	Restr. power	22.08 (0.99)	24.24 (1.19)	26.76 (1.03)	37.20 (0.95)

6.1. *Methods. Data preprocessing.* Before applying knockoffs, the data needs to be prepared and our early preprocessing follows Sabatti et al. [28] and Janson, Barber and Candès [17]. Genotype features from the original data set were removed if they met any of the following conditions:

- Not a SNP (some features were, e.g., copy number variations),
- Greater than 5% of values were missing,
- All nonmissing values belonged to the same nucleotide,
- SNP location could not be aligned to the genome,
- A χ^2 test rejected Hardy–Weinberg equilibrium at the 0.01% level,
- On chromosome 23 (sex chromosome).

The remaining missing values were assumed to take the population frequency of the major allele. In the end, we have a total of $p = 328,934$ SNP features.

For each phenotype, we performed further processing on the response variables. Triglycerides were log-transformed. C-reactive protein (CRP) was also log-transformed after adding 0.002 mg/l (half the detection limit) to any values that

were recorded as zero. Subjects were excluded from the triglycerides, HDL cholesterol and LDL cholesterol if they were on diabetic medication or had not fasted before blood collection (or if either value was missing), or if they were pregnant, or if their respective phenotype measurement was more than three standard deviations from the mean, after correcting for sex, oral contraceptive use and pregnancy. For the four phenotypes of interest to us, namely, CRP, HDL, LDL and triglycerides, the sample sizes are respectively 5290, 4700, 4682 and 4644. Due to space limitations, we report our findings on HDL and LDL only.

Finally, in order to correct for population stratification, we regressed the response and the features (SNPs) on the top five principal components of the design matrix in exactly the same fashion as suggested in Price et al. [27].

Data processing with knockoffs. Now that we have a response/design pair (\mathbf{y}, \mathbf{X}) , we discuss the implementation of our full GWAS processing pipeline.

1. *Data splitting with random rotations.* In Section 4.1, we discuss splitting the data (\mathbf{X}, \mathbf{y}) into two parts, $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$ and $(\mathbf{X}^{(1)}, \mathbf{y}^{(1)})$, with the first part used for screening and the second for inference; in that section, our discussion centers on splitting by partitioning the n observations into two groups. In practice, for settings such as GWAS where the design matrix \mathbf{X} is naturally somewhat sparse, this may be problematic as splitting the set of observations can dramatically increase conditioning problems caused by sparsity. Here, we take a different approach to splitting. Under the Gaussian linear model, since the i.i.d. Gaussian noise term $\boldsymbol{\epsilon} \in \mathbb{R}^n$ has a distribution that is unchanged by rotation, we can rotate the data with a randomly chosen orthonormal matrix $\mathbf{U} \in \mathbb{R}^{n \times n}$, to obtain $(\mathbf{X}', \mathbf{y}') = (\mathbf{U}\mathbf{X}, \mathbf{U}\mathbf{y})$. The original linear model is preserved; as long as the rotation was chosen independently of the data, our model is now

$$\mathbf{y}' = \mathbf{X}'\boldsymbol{\beta} + \boldsymbol{\epsilon}',$$

where $\boldsymbol{\epsilon}' = \mathbf{U}\boldsymbol{\epsilon} \sim \mathcal{N}(0, \sigma^2\mathbf{I}_n)$. (Note that $\boldsymbol{\beta}$ is unchanged from the original model.) After rotation, the new design matrix \mathbf{X}' is now dense, and we are able to split the rotated data into a set of n_0 data points (for screening) and $n_1 = n - n_0$ data points (for inference) without matrix conditioning issues. From this point on, the partitioned data sets $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$ and $(\mathbf{X}^{(1)}, \mathbf{y}^{(1)})$ are assumed to be taken from the rotated data, without further mention. We take $n_0 = 1900$.

2. *Feature prescreening with correlations.* The original number of features (SNPs) is 328,934, which is a bit large to easily run the Lasso for feature screening. As suggested in [39] (where the screen-and-clean methodology is applied to GWAS), we begin with a prescreening step that looks only at marginal correlations: using the first part of the data $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$, we select the set $\widehat{\mathcal{S}}_{\text{pre}} \subset [p]$ of 26,300 features with the largest magnitude correlations $|\mathbf{X}_j^{(0)\top} \mathbf{y}^{(0)}|$.

3. *Feature screening with Lasso.* We then run the Lasso on $(\mathbf{X}_{\widehat{\mathcal{S}}_{\text{pre}}}^{(0)}, \mathbf{y}^{(0)})$ using the 26,300 prescreened features as covariates, and define the screened set of features $\widehat{\mathcal{S}}_0$ to be the first $n_1/4$ features, which enter the Lasso path on this reduced data set.

4. *Knockoff filter.* Next we bring back the held out part of the data, $(\mathbf{X}^{(1)}, \mathbf{y}^{(1)})$. We apply the knockoff filter with recycling as defined in Section 4.1, with statistics determined by the sign-restricted Lasso introduced in Section 4.1.2. Specifically,, we fit the sign-restricted Lasso path

$$\min_{\mathbf{b} \in \mathbb{R}^{2|\hat{S}_0|}} \frac{1}{2} \|\mathbf{y} - [\mathbf{X}_{\hat{S}_0} \quad \tilde{\mathbf{X}}_{\hat{S}_0}] \mathbf{b}\|_2^2 + \lambda \|\mathbf{b}\|_1$$

s.t. $b_j \cdot \text{sign}_j^{(0)} \geq 0$

$b_{j+|\hat{S}_0|} \cdot \text{sign}_j^{(0)} \geq 0$

and set W_j as in (9). (That is, we record the largest λ such that \mathbf{X}_j or $\tilde{\mathbf{X}}_j$ enters the Lasso path, and set $W_j = +\lambda$ if \mathbf{X}_j enters first, and $W_j = -\lambda$ otherwise.) Throughout, we use target FDR level $q = 0.2$.

5. *Repeat.* We then repeat all these steps, with a new random rotation matrix $\mathbf{U} \in \mathbb{R}^{n \times n}$, for a total of 10 repetitions. In particular, we are interested in examining whether any selected SNPs appear consistently across these 10 repetitions of our method.

6.2. *Results.* Our results for the LDL phenotype are displayed in Figure 1. (The full details of the results, labeled according to each individual SNP discovered in the analysis, are reported in the Supplementary Material [2].) In total, over the 10 trials (the random splits of the data), 44 different SNPs are selected. However, as nearby SNPs are extremely correlated, we consider SNPs whose positions are within 10^6 base pairs of each other, to be in the same “region”; at this level, in total over the 10 trials, there are 29 distinct regions discovered. In Figure 1, we show, for each region, the selection frequency of any SNP in that region, that is, the number of trials out of 10 for which at least one SNP in this region was selected (we also display the selection frequency for the individual SNPs). We also show any




FIG. 1. *Results of GWAS experiment for the LDL phenotype.* The heights of the lines show the number of trials (out of 10 trials) for which at least one SNP in this region was selected; the line type shows whether this same region was identified in [38] and/or [28]. Regions where more than one SNP was selected (across the 10 trials) are marked with a “*”. (See Section 6.2 for details.)




FIG. 2. Results of GWAS experiment for the HDL phenotype (same interpretation as in Figure 1).

SNPs in the same region (defined again as a distance of $\leq 10^6$ base pairs) which were identified as associated with LDL in the meta-analysis of Willer et al. [38]; this meta-analysis works with an extremely large sample size and, therefore, can be viewed as providing a form of “ground truth.” Comparing against the findings of this meta-analysis, we can try to estimate the FDR of our method as follows: for each trial, we count each selected SNP as a discovery, and label it as a false discovery if there is no SNP in the meta-analysis from the same region. We then average this false discovery proportion over the 10 trials, and find an estimated FDR of 32.61%. We also compare against the findings of Sabatti et al. [28], which analyzes the same data set that we use; this comparison therefore should not be viewed as independent validation, but rather as checking that our analysis agrees with existing work on the same data set.

We note that since our analysis is run treating each SNP as an individual predictor, the FDR control guarantee is at the level of the SNPs rather than the regions; that is, our method ensures bounded FDR among the list of 44 discovered SNPs, but not necessarily among the list of 29 discovered regions. In practice, it may be preferable to group the SNPs into regions ahead of time and run the knockoff filter over the list of feature groups (see, e.g., [10]).

Turning to the HDL phenotype, our results are shown in Figure 2 (and reported in full detail in the Supplementary Material [2]). There are 13 SNPs discovered at least once in the 10 trials; once grouped into regions, there are 10 distinct regions. Again estimating the FDR by comparing against Willer et al. [38]’s meta-analysis, we estimate an FDR of 4.29%. This low number of discoveries, and perhaps conservative FDR, can probably be attributed to fact that this data set has a relatively small sample size, and fairly weak signal; we would not expect a large number of discoveries in this setting.

7. Summary. We have discussed a “screen + knockoff” approach to controlled variable selection, which should appear very natural to practitioners; in addition, we develop the “recycling” technique to reuse data rather than discarding the data used for the screening step, thus improving power. A natural question is

thus what sort of inference properties would such a procedure offer? By focusing on Type I + Type S errors, we have shown that screen + knockoff is guaranteed to have a form of reproducibility in that the directional FDR is rigorously under control. This result makes no assumption on the design matrix nor on the values of the regression coefficients, and holds in finite samples; our contribution is, therefore, very different in spirit from most of the results published in the literature on high-dimensional inference (the work of Taylor and colleagues referenced earlier being a notable exception), and employs mathematical arguments centered in martingale theory which are naturally also very different from the techniques used in the related literature.

APPENDIX: PROOFS

A.1. A general form of directional FDR control. Our results, namely, Theorems 1–3, are all special cases of the general statement below. Throughout, $\widehat{\text{sign}}_j$ is defined as in (15), i.e. $\widehat{\text{sign}}_j = \text{sign}((\mathbf{X}_j - \widetilde{\mathbf{X}}_j)^\top \mathbf{y})$. Also, we shall say that a positive semidefinite matrix $\mathbf{M} \in \mathbb{R}^{2p \times 2p}$ satisfies the *pairwise exchangeability condition* if $M_{j,k} = M_{j,k+p} = M_{j+p,k} = M_{j+p,k+p}$ for all $j \neq k \in \{1, \dots, p\}$ and $M_{j,j} = M_{j+p,j+p}$ for all $j \in \{1, \dots, p\}$.

THEOREM 4. *Suppose $\mathbf{y} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Theta})$ for some mean vector $\boldsymbol{\mu} \in \mathbb{R}^n$ and some covariance $\boldsymbol{\Theta} \in \mathbb{R}^{n \times n}$. Let $\mathbf{X}, \widetilde{\mathbf{X}} \in \mathbb{R}^{n \times p}$ be any fixed matrices such that $[\mathbf{X} \ \widetilde{\mathbf{X}}]^\top \cdot \boldsymbol{\Theta} \cdot [\mathbf{X} \ \widetilde{\mathbf{X}}]$ obeys the pairwise exchangeability condition. Let $\mathbf{W} = (W_1, \dots, W_p)$ be any statistic satisfying the sufficiency and antisymmetry properties (7) and (8). Finally, let $V_j = \mathbb{1}\{\widehat{\text{sign}}_j \neq \text{sign}((\mathbf{X}_j - \widetilde{\mathbf{X}}_j)^\top \boldsymbol{\mu})\}$ be the indicator of a sign error on feature j . Then the knockoff+ method controls a directional FDR defined as*

$$\text{FDR}_{\text{dir}} = \mathbb{E} \left[\frac{|\{j \in \widehat{S} \text{ and } V_j = 1\}|}{1 \vee |\widehat{S}|} \right] \leq q.$$

If the knockoff method is used instead, then a modified directional FDR is controlled:

$$\text{mFDR}_{\text{dir}} = \mathbb{E} \left[\frac{|\{j \in \widehat{S} \text{ and } V_j = 1\}|}{|\widehat{S}| + q^{-1}} \right] \leq q.$$

We prove this theorem in the Supplementary Material. The rest of this section shows how all of our results follow as special cases.

Proof of Theorem 1. This is straightforward since the result is obtained by taking $\boldsymbol{\mu} = \mathbf{X}\boldsymbol{\beta}$ and $\boldsymbol{\Theta} = \sigma^2 \mathbf{I}_n$ in the distribution of \mathbf{y} , and noting that $\text{sign}((\mathbf{X}_j - \widetilde{\mathbf{X}}_j)^\top \boldsymbol{\mu}) = \text{sign}(\beta_j)$ for all j such that $\mathbf{X}_j \neq \widetilde{\mathbf{X}}_j$. (We ignore the trivial case where $s_j = 0$ implying $\mathbf{X}_j = \widetilde{\mathbf{X}}_j$, for in this case we would get $W_j = 0$ and thus feature j could never be selected.)

The more general form of Theorem 4 also allows us to consider other settings, such as those involving variable screening in high dimensions. For instance, our next proof concerns the testing of regression coefficients in the full model under the sure screening property.

Proof of Theorem 2. We prove directional FDR control by conditioning on $\mathbf{y}^{(0)}$. For any $\mathbf{y}^{(0)}$ such that the event \mathcal{E} holds, we will show that for the knockoff procedure,

$$(24) \quad \mathbb{E}[\text{mFDR}_{\text{dir}}|\mathbf{y}^{(0)}] \cdot \mathbb{1}_{\mathcal{E}} \leq q,$$

and for the knockoff+,

$$(25) \quad \mathbb{E}[\text{FDR}_{\text{dir}}|\mathbf{y}^{(0)}] \cdot \mathbb{1}_{\mathcal{E}} \leq q.$$

From this point on we treat $\mathbf{y}^{(0)}$ as fixed. The conditional distribution of the response \mathbf{y} is given by

$$(26) \quad (\mathbf{y} | \mathbf{y}^{(0)}) \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{y}^{(0)} \\ \mathbf{X}^{(1)} \boldsymbol{\beta} \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 0 & \sigma^2 \mathbf{I}_{n_1} \end{bmatrix} \right).$$

If $\mathbf{y}^{(0)}$ is such that the event \mathcal{E} occurs, then we can simplify this to

$$(\mathbf{y} | \mathbf{y}^{(0)}) \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Theta}),$$

$$(27) \quad \text{where } \boldsymbol{\mu} = \begin{bmatrix} \mathbf{y}^{(0)} \\ \mathbf{X}_{\widehat{S}_0}^{(1)} \boldsymbol{\beta}_{\widehat{S}_0} \end{bmatrix} \text{ and } \boldsymbol{\Theta} = \begin{bmatrix} 0 & 0 \\ 0 & \sigma^2 \mathbf{I}_{n_1} \end{bmatrix},$$

since $\beta_j = 0$ for all $j \notin \widehat{S}_0$ when \mathcal{E} occurs. Define $\boldsymbol{\Sigma}_{\widehat{S}_0}^{(1)} = \mathbf{X}_{\widehat{S}_0}^{(1)\top} \mathbf{X}_{\widehat{S}_0}^{(1)}$. Recall that $\widetilde{\mathbf{X}}_{\widehat{S}_0}^{(0)} = \mathbf{X}_{\widehat{S}_0}^{(0)}$ by (17), and that $\widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}$ satisfies

$$(28) \quad [\mathbf{X}_{\widehat{S}_0}^{(1)} \quad \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)\top}] [\mathbf{X}_{\widehat{S}_0}^{(1)} \quad \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}] = \begin{bmatrix} \boldsymbol{\Sigma}_{\widehat{S}_0}^{(1)} & \boldsymbol{\Sigma}_{\widehat{S}_0}^{(1)} - \text{diag}\{\mathbf{s}\} \\ \boldsymbol{\Sigma}_{\widehat{S}_0}^{(1)} - \text{diag}\{\mathbf{s}\} & \boldsymbol{\Sigma}_{\widehat{S}_0}^{(1)} \end{bmatrix}$$

for some vector $\mathbf{s} \geq 0$, by the knockoff construction. We then see that

$$[\mathbf{X}_{\widehat{S}_0} \quad \widetilde{\mathbf{X}}_{\widehat{S}_0}]^\top \boldsymbol{\Theta} [\mathbf{X}_{\widehat{S}_0} \quad \widetilde{\mathbf{X}}_{\widehat{S}_0}] = \sigma^2 [\mathbf{X}_{\widehat{S}_0}^{(1)} \quad \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)\top}] [\mathbf{X}_{\widehat{S}_0}^{(1)} \quad \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}],$$

is a pairwise exchangeable matrix. At this point, we have satisfied the conditions of Theorem 4. Now what is a sign error in Theorem 4? We calculate that for any $j \in \widehat{S}_0$,

$$\begin{aligned} (\mathbf{X}_j - \widetilde{\mathbf{X}}_j)^\top \boldsymbol{\mu} &= (\mathbf{X}_j^{(1)} - \widetilde{\mathbf{X}}_j^{(1)})^\top \mathbf{X}_{\widehat{S}_0}^{(1)} \boldsymbol{\beta}_{\widehat{S}_0} \\ &= \sum_{k \in \widehat{S}_0} (\mathbf{X}_j^{(1)} - \widetilde{\mathbf{X}}_j^{(1)})^\top \mathbf{X}_k^{(1)} \cdot \beta_k = s_j \cdot \beta_j, \end{aligned}$$

where the first equality holds on the event \mathcal{E} (i.e. $\beta_k = 0$ for $k \notin \widehat{S}_0$), and the third equality uses the Gram matrix condition (28). Hence, since $s_j > 0$ whenever $\mathbf{X}_j \neq \widetilde{\mathbf{X}}_j$, an error $\widehat{\text{sign}}_j \neq \text{sign}((\mathbf{X}_j - \widetilde{\mathbf{X}}_j)^\top \boldsymbol{\mu})$ is the same as $\widehat{\text{sign}}_j \neq \text{sign}(\beta_j)$, which completes the proof.

We now move on to our result concerning regression coefficients from the reduced model.

Proof of Theorem 3. This last proof is more subtle and for this, we treat \mathbf{X} as random in addition to \mathbf{y} , but will condition on $\mathbf{X}^{(0)}$ and on the screened features $\mathbf{X}_{\widehat{S}_0}^{(1)}$ and knockoffs $\widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}$ as well as on $\mathbf{y}^{(0)}$. By the assumption that the rows of \mathbf{X} are i.i.d. draws from $\mathcal{N}(0, \boldsymbol{\Psi})$, we can write

$$\mathbf{X}_{\widehat{S}_0^c}^{(1)} = \mathbf{X}_{\widehat{S}_0}^{(1)} \cdot \boldsymbol{\Gamma} + \mathbf{G} \cdot \boldsymbol{\Lambda}^{1/2},$$

where $\boldsymbol{\Gamma}$ and $\boldsymbol{\Lambda}$ are the fixed unknown matrices given by

$$\boldsymbol{\Gamma} = (\boldsymbol{\Psi}_{\widehat{S}_0, \widehat{S}_0})^{-1} \cdot \boldsymbol{\Psi}_{\widehat{S}_0, \widehat{S}_0^c}$$

and the Schur complement

$$\boldsymbol{\Lambda} = \boldsymbol{\Psi}_{\widehat{S}_0^c, \widehat{S}_0^c} - \boldsymbol{\Psi}_{\widehat{S}_0^c, \widehat{S}_0} \cdot (\boldsymbol{\Psi}_{\widehat{S}_0, \widehat{S}_0})^{-1} \cdot \boldsymbol{\Psi}_{\widehat{S}_0, \widehat{S}_0^c},$$

and where $\mathbf{G} \in \mathbb{R}^{n_1 \times (p - |\widehat{S}_0|)}$ has i.i.d. standard normal entries drawn independently from $(\mathbf{X}^{(0)}, \mathbf{y}^{(0)})$ and from $(\mathbf{X}_{\widehat{S}_0}^{(1)}, \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}, \boldsymbol{\epsilon}^{(1)})$. Then

$$\mathbf{y}^{(1)} = \mathbf{X}^{(1)} \boldsymbol{\beta} + \boldsymbol{\epsilon}^{(1)} = \mathbf{X}_{\widehat{S}_0}^{(1)} \cdot (\boldsymbol{\beta}_{\widehat{S}_0} + \boldsymbol{\Gamma} \boldsymbol{\beta}_{\widehat{S}_0^c}) + \mathbf{G} \cdot \boldsymbol{\Lambda}^{1/2} \boldsymbol{\beta}_{\widehat{S}_0^c} + \boldsymbol{\epsilon}^{(1)},$$

and note that

$$\boldsymbol{\beta}_{\widehat{S}_0} + \boldsymbol{\Gamma} \boldsymbol{\beta}_{\widehat{S}_0^c} = \boldsymbol{\beta}^{\text{partial}}$$

by definition of the expected partial regression coefficients. This is because from (20), we get that

$$\begin{aligned} \boldsymbol{\beta}^{\text{partial}} &= \mathbf{X}_{\widehat{S}_0}^{(1)\dagger} (\mathbf{X}_{\widehat{S}_0}^{(1)} \boldsymbol{\beta}_{\widehat{S}_0} + \mathbb{E}[\mathbf{X}_{\widehat{S}_0^c}^{(1)} \boldsymbol{\beta}_{\widehat{S}_0^c} | \mathbf{X}_{\widehat{S}_0}^{(1)}; \mathbf{X}^{(0)}, \mathbf{y}^{(0)}]) \\ &= \mathbf{X}_{\widehat{S}_0}^{(1)\dagger} (\mathbf{X}_{\widehat{S}_0}^{(1)} \boldsymbol{\beta}_{\widehat{S}_0} + \mathbf{X}_{\widehat{S}_0}^{(1)} \boldsymbol{\Gamma} \boldsymbol{\beta}_{\widehat{S}_0^c}) \\ &= \boldsymbol{\beta}_{\widehat{S}_0} + \boldsymbol{\Gamma} \boldsymbol{\beta}_{\widehat{S}_0^c}, \end{aligned}$$

where $\mathbf{X}_{\widehat{S}_0}^{(1)\dagger}$ is a short hand for $(\mathbf{X}_{\widehat{S}_0}^{(1)\top} \mathbf{X}_{\widehat{S}_0}^{(1)})^{-1} \mathbf{X}_{\widehat{S}_0}^{(1)\top}$. Therefore,

$$(\mathbf{y} | \mathbf{X}^{(0)}, \mathbf{y}^{(0)}, \mathbf{X}_{\widehat{S}_0}^{(1)}, \widetilde{\mathbf{X}}_{\widehat{S}_0}^{(1)}) \sim \mathcal{N} \left(\begin{bmatrix} \mathbf{y}^{(0)} \\ \mathbf{X}_{\widehat{S}_0}^{(1)} \boldsymbol{\beta}^{\text{partial}} \end{bmatrix}, \begin{bmatrix} 0 & 0 \\ 0 & (\boldsymbol{\beta}_{\widehat{S}_0^c}^\top \boldsymbol{\Lambda} \boldsymbol{\beta}_{\widehat{S}_0^c} + \sigma^2) \mathbf{I}_{n_1} \end{bmatrix} \right).$$

We are now in the same setting as in the proof of Theorem 2 (see (26)), except with $\boldsymbol{\beta}^{\text{partial}}$ in place of $\boldsymbol{\beta}$ and with a new variance level that has increased due to the randomness in the missed signals. The remainder of the proof thus proceeds identically.

Acknowledgements. E. C. would like to thank Lucas Janson, Chiara Sabatti and Jonathan Taylor for many inspiring discussions about topics ranging from high-dimensional inference to GWAS, and useful comments about an early version of the manuscript. We are indebted to Chiara Sabatti and Christine Peterson for all their help with the NFBC data set and its preprocessing, and to Pragya Sur for her editorial comments and assistance. E. C. would like to thank Mert Pilanci and Zhimei Ren for their help in running simulations on the Sherlock cluster.

SUPPLEMENTARY MATERIAL

Supplement to “A knockoff filter for high-dimensional selective inference” (DOI: [10.1214/18-AOS1755SUPP](https://doi.org/10.1214/18-AOS1755SUPP); .pdf). We provide details for the proofs of several theoretical results in the paper, and report detailed results on the GWAS real data experiment.

REFERENCES

- [1] BARBER, R. F. and CANDÈS, E. J. (2015). Controlling the false discovery rate via knockoffs. *Ann. Statist.* **43** 2055–2085. [MR3375876](#)
- [2] BARBER, R. F. and CANDÈS, E. J. (2019). Supplement to “A knockoff filter for high-dimensional selective inference.” DOI:[10.1214/18-AOS1755SUPP](https://doi.org/10.1214/18-AOS1755SUPP).
- [3] BELLONI, A., CHERNOZHUKOV, V. and HANSEN, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *Rev. Econ. Stud.* **81** 608–650. [MR3207983](#)
- [4] BELLONI, A., CHERNOZHUKOV, V. and WANG, L. (2011). Square-root Lasso: Pivotal recovery of sparse signals via conic programming. *Biometrika* **98** 791–806. [MR2860324](#)
- [5] BENJAMINI, Y. and BRAUN, H. (2002). John Tukey’s contributions to multiple comparisons. ETS Research Report Series.
- [6] BENJAMINI, Y. and HOCHBERG, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *J. Roy. Statist. Soc. Ser. B* **57** 289–300. [MR1325392](#)
- [7] BENJAMINI, Y. and YEKUTIELI, D. (2005). False discovery rate-adjusted multiple confidence intervals for selected parameters. *J. Amer. Statist. Assoc.* **100** 71–93. [MR2156820](#)
- [8] BERK, R., BROWN, L., BUJA, A., ZHANG, K. and ZHAO, L. (2013). Valid post-selection inference. *Ann. Statist.* **41** 802–837. [MR3099122](#)
- [9] CANDÈS, E., FAN, Y., JANSON, L. and LV, J. (2018). Panning for gold: ‘Model-X’ knockoffs for high dimensional controlled variable selection. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **80** 551–577. [MR3798878](#)
- [10] DAI, R. and BARBER, R. (2016). The knockoff filter for fdr control in group-sparse and multi-task regression. In *Proceedings of the 33rd International Conference on Machine Learning* 1851–1859.
- [11] FAN, J. and LV, J. (2008). Sure independence screening for ultrahigh dimensional feature space. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **70** 849–911. [MR2530322](#)
- [12] FITHIAN, W., SUN, D. and TAYLOR, J. (2014). Optimal inference after model selection. Preprint. Available at [arXiv:1410.2597](https://arxiv.org/abs/1410.2597).
- [13] G’SSELL, M. G., HASTIE, T. and TIBSHIRANI, R. (2013). False variable selection rates in regression. Preprint. Available at [arXiv:1302.2303](https://arxiv.org/abs/1302.2303).

- [14] G'SELL, M. G., WAGER, S., CHOULDECHOVA, A. and TIBSHIRANI, R. (2016). Sequential selection procedures and false discovery rate control. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **78** 423–444. [MR3454203](#)
- [15] GELMAN, A. and TUERLINCKX, F. (2000). Type S error rates for classical and Bayesian single and multiple comparison procedures. *Comput. Statist.* **15** 373–390.
- [16] HUANG, J., MA, S., ZHANG, C.-H. and ZHOU, Y. (2013). Semi-penalized inference with direct false discovery rate control in high-dimensions. Preprint. Available at [arXiv:1311.7455](#).
- [17] JANSON, L., BARBER, R. F. and CANDÈS, E. (2017). EigenPrism: Inference for high dimensional signal-to-noise ratios. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **79** 1037–1065. [MR3689308](#)
- [18] JANSON, L. and SU, W. (2016). Familywise error rate control via knockoffs. *Electron. J. Stat.* **10** 960–975. Available at [arXiv:1505.06549](#). [MR3486422](#)
- [19] JÄRVELIN, M.-R., SOVIO, U., KING, V., LAUREN, L., XU, B., MCCARTHY, M. I., HARTIKAINEN, A.-L., LAITINEN, J., ZITTING, P. et al. (2004). Early life factors and blood pressure at age 31 years in the 1966 northern Finland birth cohort. *Hypertension* **44** 838–846.
- [20] JAVANMARD, A. and MONTANARI, A. (2014). Confidence intervals and hypothesis testing for high-dimensional regression. *J. Mach. Learn. Res.* **15** 2869–2909. [MR3277152](#)
- [21] JONES, L. V. and TUKEY, J. W. (2000). A sensible formulation of the significance test. *Psychol. Methods* **5** 411.
- [22] LEE, J. D., SUN, D. L., SUN, Y. and TAYLOR, J. E. (2016). Exact post-selection inference, with application to the lasso. *Ann. Statist.* **44** 907–927. Available at [arXiv:1311.6238](#). [MR3485948](#)
- [23] LEEB, H. and PÖTSCHER, B. M. (2006). Can one estimate the conditional distribution of post-model-selection estimators? *Ann. Statist.* **34** 2554–2591. [MR2291510](#)
- [24] LOCKHART, R., TAYLOR, J., TIBSHIRANI, R. J. and TIBSHIRANI, R. (2014). A significance test for the Lasso. *Ann. Statist.* **42** 413–468. [MR3210970](#)
- [25] MILLER, A. (2002). *Subset Selection in Regression*, 2nd ed. *Monographs on Statistics and Applied Probability* **95**. CRC Press/CRC, Boca Raton, FL. [MR2001193](#)
- [26] PATI, Y. C., REZAIIFAR, R. and KRISHNAPRASAD, P. S. (1993). Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition. In *Proceedings of 27th Asilomar Conference on Signals, Systems and Computers* 40–44. IEEE, New York.
- [27] PRICE, A. L., PATTERSON, N. J., PLENGE, R. M., WEINBLATT, M. E., SHADICK, N. A. and REICH, D. (2006). Principal components analysis corrects for stratification in genome-wide association studies. *Nat. Genet.* **38** 904–909.
- [28] SABATTI, C., SERVICE, S. K., HARTIKAINEN, A.-L., POUTA, A., RIPATTI, S., BRODSKY, J., JONES, C. G., ZAITLEN, N. A., VARILO, T. et al. (2009). Genome-wide association analysis of metabolic traits in a birth cohort from a founder population. *Nat. Genet.* **41** 35–46.
- [29] SHAFFER, J. P. (2002). Multiplicity, directional (type III) errors, and the null hypothesis. *Psychol. Methods* **7** 356–369.
- [30] SU, W., BOGDAN, M. and CANDÈS, E. (2017). False discoveries occur early on the Lasso path. *Ann. Statist.* **45** 2133–2150. [MR3718164](#)
- [31] TAYLOR, J. T. (2017). Selective-inference. Available at <https://github.com/jonathan-taylor/selective-inference>.
- [32] TIAN, X., LOFTUS, J. R. and TAYLOR, J. E. (2018). Selective inference with unknown variance via the square-root Lasso. *Biometrika* **105** 755–768. Available at [arXiv:1504.08031](#). [MR3877864](#)

- [33] TIBSHIRANI, R. (1996). Regression shrinkage and selection via the Lasso. *J. Roy. Statist. Soc. Ser. B* **58** 267–288. [MR1379242](#)
- [34] TIBSHIRANI, R. J., TAYLOR, J., LOCKHART, R. and TIBSHIRANI, R. (2016). Exact post-selection inference for sequential regression procedures. *J. Amer. Statist. Assoc.* **111** 600–620. [MR3538689](#)
- [35] TUKEY, J. W. (1991). The philosophy of multiple comparisons. *Statist. Sci.* **6** 100–116.
- [36] VOORMAN, A., SHOJAIE, A. and WITEN, D. (2014). Inference in high dimensions with the penalized score test. Preprint. Available at [arXiv:1401.2678](#).
- [37] WASSERMAN, L. and ROEDER, K. (2009). High-dimensional variable selection. *Ann. Statist.* **37** 2178–2201. [MR2543689](#)
- [38] WILLER, C. J., SCHMIDT, E. M., SENGUPTA, S., PELOSO, G. M., GUSTAFSSON, S., KANONI, S., GANNA, A., CHEN, J., BUCHKOVICH, M. L. et al. (2013). Discovery and refinement of loci associated with lipid levels. *Nat. Genet.* **45** 1274–1283.
- [39] WU, J., DEVLIN, B., RINGQUIST, S., TRUCCO, M. and ROEDER, K. (2010). Screen and clean: A tool for identifying interactions in genome-wide association studies. *Genet. Epidemiol.* **34** 275–285.
- [40] WU, Y., BOOS, D. D. and STEFANSKI, L. A. (2007). Controlling variable selection by the addition of pseudovariables. *J. Amer. Statist. Assoc.* **102** 235–243. [MR2345541](#)
- [41] ZHANG, C.-H. and HUANG, J. (2008). The sparsity and bias of the Lasso selection in high-dimensional linear regression. *Ann. Statist.* **36** 1567–1594. [MR2435448](#)
- [42] ZHANG, C.-H. and ZHANG, S. S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **76** 217–242. [MR3153940](#)
- [43] ZHAO, P. and YU, B. (2006). On model selection consistency of Lasso. *J. Mach. Learn. Res.* **7** 2541–2563. [MR2274449](#)

DEPARTMENT OF STATISTICS
UNIVERSITY OF CHICAGO
5747 S. ELLIS AVENUE
CHICAGO, ILLINOIS 60637
USA
E-MAIL: rina@uchicago.edu

DEPARTMENTS OF STATISTICS
AND MATHEMATICS
STANFORD UNIVERSITY
STANFORD, CALIFORNIA 94305
USA
E-MAIL: candes@stanford.edu