

Lasso adjustments of treatment effect estimates in randomized experiments

Adam Bloniarz^{a,1}, Hanzhong Liu^{a,1}, Cun-Hui Zhang^b, Jasjeet S. Sekhon^{a,c}, and Bin Yu^{a,d,2}

^aDepartment of Statistics, University of California, Berkeley, CA 94720; ^bDepartment of Statistics and Biostatistics, Rutgers University, Piscataway, NJ 08854; ^cDepartment of Political Science, University of California, Berkeley, CA 94720; and ^dDepartment of Electrical Engineering and Computer Science, University of California, Berkeley, CA 94720

Edited by Richard M. Shiffrin, Indiana University, Bloomington, IN, and approved December 1, 2015 (received for review June 3, 2015)

We provide a principled way for investigators to analyze randomized experiments when the number of covariates is large. Investigators often use linear multivariate regression to analyze randomized experiments instead of simply reporting the difference of means between treatment and control groups. Their aim is to reduce the variance of the estimated treatment effect by adjusting for covariates. If there are a large number of covariates relative to the number of observations, regression may perform poorly because of overfitting. In such cases, the least absolute shrinkage and selection operator (Lasso) may be helpful. We study the resulting Lasso-based treatment effect estimator under the Neyman–Rubin model of randomized experiments. We present theoretical conditions that guarantee that the estimator is more efficient than the simple difference-of-means estimator, and we provide a conservative estimator of the asymptotic variance, which can yield tighter confidence intervals than the difference-of-means estimator. Simulation and data examples show that Lasso-based adjustment can be advantageous even when the number of covariates is less than the number of observations. Specifically, a variant using Lasso for selection and ordinary least squares (OLS) for estimation performs particularly well, and it chooses a smoothing parameter based on combined performance of Lasso and OLS.

randomized experiment | Neyman–Rubin model | average treatment effect | high-dimensional statistics | Lasso

Randomized experiments are widely used to measure the efficacy of treatments. Randomization ensures that treatment assignment is not influenced by any potential confounding factors, both observed and unobserved. Experiments are particularly useful when there is no rigorous theory of a system's dynamics, and full identification of confounders would be impossible. This advantage was cast elegantly in mathematical terms in the early 20th century by Jerzy Neyman, who introduced a simple model for randomized experiments, which showed that the difference of average outcomes in the treatment and control groups is statistically unbiased for the average treatment effect (ATE) over the experimental sample (1).

However, no experiment occurs in a vacuum of scientific knowledge. Often, baseline covariate information is collected about individuals in an experiment. Even when treatment assignment is not related to these covariates, analyses of experimental outcomes often take them into account with the goal of improving the accuracy of treatment effect estimates. In modern randomized experiments, the number of covariates can be very large—sometimes even larger than the number of individuals in the study. In clinical trials overseen by regulatory bodies like the Food and Drug Administration and the Medicines and Healthcare products Regulatory Agency, demographic and genetic information may be recorded about each patient. In applications in the tech industry, where randomization is often called A/B testing, there is often a huge amount of behavioral data collected on each user. However, in this “big data” setting, much of these data may be irrelevant to the outcome being studied or there may be more potential covariates than observations, especially once interactions are taken into account. In these cases, selection

of important covariates or some form of regularization is necessary for effective regression adjustment.

To ground our discussion, we examine a randomized trial of the pulmonary artery catheter (PAC) that was carried out in 65 intensive care units in the United Kingdom between 2001 and 2004, called PAC-man (2). The PAC is a monitoring device commonly inserted into critically ill patients after admission to intensive care, and it provides a continuous measurement of several indicators of cardiac activity. However, insertion of PAC is an invasive procedure that carries some risk of complications (including death), and it involves significant expenditure both in equipment costs and personnel (3). Controversy over its use came to a head when an observational study found that PAC had an adverse effect on patient survival and led to increased cost of care (4). This led to several large-scale randomized trials, including PAC-man.

In the PAC-man trial, randomization of treatment was largely successful, and a number of covariates were measured about each patient in the study. If covariate interactions are included, the number of covariates exceeds the number of individuals in the study; however, few of them are predictive of the patient's outcome. As it turned out, the (pretreatment) estimated probability of death was imbalanced between the treatment and control groups ($P = 0.005$, Wilcoxon rank sum test). Because the control group had, on average, a slightly higher risk of death, the unadjusted difference-in-means estimator may overestimate the benefits of receiving a PAC. Adjustment for this imbalance seems advantageous in this case, because the pretreatment probability of death is clearly predictive of health outcomes posttreatment.

In this paper, we study regression-based adjustment, using the least absolute shrinkage and selection operator (Lasso) to select relevant covariates. Standard linear regression based on ordinary least squares (OLS) suffers from overfitting if a large number of covariates and interaction terms are included in the model. In such cases, researchers sometimes perform model selection based on observing which covariates are unbalanced given the realized randomization. This generally leads to misleading inferences because of incorrect test levels (5). The Lasso (6) provides researchers with an alternative that can mitigate these problems and still perform model selection. We define an estimator, $\widehat{ATE}_{\text{Lasso}}$, which is based on running an l_1 -penalized linear

This paper results from the Arthur M. Sackler Colloquium of the National Academy of Sciences, “Drawing Causal Inference from Big Data,” held March 26–27, 2015, at the National Academies of Sciences in Washington, DC. The complete program and video recordings of most presentations are available on the NAS website at www.nasonline.org/Big-data.

Author contributions: A.B., H.L., C.-H.Z., J.S.S., and B.Y. designed research, performed research, analyzed data, and wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

¹A.B. and H.L. contributed equally to this work.

²To whom correspondence should be addressed. Email: binyu@stat.berkeley.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1510506113/-DCSupplemental.

regression of the outcome on treatment, covariates, and, following the method introduced in ref. 7, treatment by covariate interactions. Because of the geometry of the l_1 penalty, the Lasso will usually set many regression coefficients to 0, and is well defined even if the number of covariates is larger than the number of observations. The Lasso's theoretical properties under the standard linear model have been widely studied in the last decade; consistency properties for coefficient estimation, model selection, and out-of-sample prediction are well understood (see ref. 8 for an overview).

In the theoretical analysis in this paper, instead of assuming that the standard linear model is the true data-generating mechanism, we work under the aforementioned nonparametric model of randomization introduced by Neyman (1) and popularized by Donald Rubin (9). In this model, the outcomes and covariates are fixed quantities, and the treatment group is assumed to be sampled without replacement from a finite population. The treatment indicator, rather than an error term, is the source of randomness, and it determines which of two potential outcomes is revealed to the experimenter. Unlike the standard linear model, the Neyman–Rubin model makes few assumptions not guaranteed by the randomization itself. The setup of the model does rely on the stable unit treatment value assumption, which states that there is only one version of treatment, and that the potential outcome of one unit should be unaffected by the particular assignment of treatments to the other units; however, it makes no assumptions of linearity or exogeneity of error terms. OLS (7, 10, 11), logistic regression (12), and poststratification (13) are among the adjustment methods that have been studied under this model.

To be useful to practitioners, the Lasso-based treatment effect estimator must be consistent and yield a method to construct valid confidence intervals. We outline conditions on the covariates and potential outcomes that will guarantee these properties. We show that an upper bound for the asymptotic variance can be estimated from the model residuals, yielding asymptotically conservative confidence intervals for the ATE, which can be substantially narrower than the unadjusted confidence intervals. Simulation studies are provided to show the advantage of the Lasso-adjusted estimator and to show situations where it breaks down. We apply the estimator to the PAC-man data, and compare the estimates and confidence intervals derived from the unadjusted, OLS-adjusted, and Lasso-adjusted methods. We also compare different methods of selecting the Lasso tuning parameter on these data.

Framework and Definitions

We give a brief outline of the Neyman–Rubin model for a randomized experiment; the reader is urged to consult refs. 1, 9, and 14 for more details. We follow the notation introduced in refs. 7 and 10. For concreteness, we illustrate the model in the context of the PAC-man trial.

For each individual in the study, the model assumes that there exists a pair of quantities representing his/her health outcomes under the possibilities of receiving and not receiving the catheter. These are called the potential outcomes under treatment and control, and are denoted as a_i and b_i , respectively. In the course of the study, the experimenter observes only one of these quantities for each individual, because the catheter is either inserted or not. The causal effect of the treatment on individual i is defined, in theory, to be $a_i - b_i$, but this is unobservable. Instead of trying to infer individual-level effects, we will assume that the intention is to estimate the average causal effect over the whole population, as outlined in the next section.

In the mathematical specification of this model, we consider the potential outcomes to be fixed, nonrandom quantities, even though they are not all observable. The only randomness in the model comes from the assignment of treatment, which is

controlled by the experimenter. We define random treatment indicators T_i , which take on a value 1 for a treated individual, or 0 for an untreated individual. We will assume that the set of treated individuals is sampled without replacement from the full population, where the size of the treatment group is fixed beforehand; thus, the T_i are identically distributed but not independent. The model for the observed outcome for individual i , defined as Y_i , is thus as follows:

$$Y_i = T_i a_i + (1 - T_i) b_i.$$

This equation simply formalizes the idea that the experimenter observes the potential outcome under treatment for those who receive the treatment, and the potential outcome under control for those who do not.

Note that the model does not incorporate any covariate information about the individuals in the study, such as physiological characteristics or health history. However, we will assume we have measured a vector of baseline, preexperimental covariates for each individual i . These might include, for example, age, gender, and genetic makeup. We denote the covariates for individual i as the column vector $\mathbf{x}_i = (x_{i1}, \dots, x_{ip})^T \in \mathbb{R}^p$ and the full design matrix of the experiment as $X = (\mathbf{x}_1, \dots, \mathbf{x}_n)^T$. In *Theoretical Results*, we will assume that there is a correlational relationship between an individual's potential outcomes and covariates, but we will not assume a generative statistical model.

Define the set of treated individuals as $A = \{i \in \{1, \dots, n\} : T_i = 1\}$, and similarly define the set of control individuals as B . Define the number of treated and control individuals as $n_A = |A|$ and $n_B = |B|$, respectively, so that $n_A + n_B = n$. We add a line on top of a quantity to indicate its average and a subscript A or B to label the treatment or control group. Thus, for example, the average values of the potential outcomes and the covariates in the treatment group are as follows:

$$\bar{a}_A = n_A^{-1} \sum_{i \in A} a_i, \quad \bar{\mathbf{x}}_A = n_A^{-1} \sum_{i \in A} \mathbf{x}_i,$$

respectively. Note that these are random quantities in this model, because the set A is determined by the random treatment assignment. Averages over the whole population are denoted as

$$\bar{a} = n^{-1} \sum_{i=1}^n a_i, \quad \bar{b} = n^{-1} \sum_{i=1}^n b_i, \quad \bar{\mathbf{x}} = n^{-1} \sum_{i=1}^n \mathbf{x}_i.$$

Note that the averages of potential outcomes over the whole population are not considered random, but are unobservable.

Treatment Effect Estimation

Our main inferential goal will be average effect of the treatment over the whole population in the study. In a trial such as PAC-man, this represents the difference between the average outcome if everyone had received the catheter, and the average outcome if no one had received it. This is defined as follows:

$$ATE = \bar{a} - \bar{b}.$$

The most natural estimator arises by replacing the population averages with the sample averages:

$$\widehat{ATE}_{\text{unadj}} = \bar{a}_A - \bar{b}_B,$$

The subscript “unadj” indicates an estimator without regression adjustment. The foundational work in ref. 1 points out that, under a randomized assignment of treatment, $\widehat{ATE}_{\text{unadj}}$ is unbiased for ATE , and derives a conservative procedure for estimating its variance.

Although $\widehat{ATE}_{\text{unadj}}$ is an attractive estimator, covariate information can be used to make adjustments in the hope of reducing variance. A commonly used estimator is as follows:

$$\widehat{ATE}_{\text{adj}} = \left[\bar{a}_A - (\bar{\mathbf{x}}_A - \bar{\mathbf{x}})^T \hat{\boldsymbol{\beta}}^{(a)} \right] - \left[\bar{b}_B - (\bar{\mathbf{x}}_B - \bar{\mathbf{x}})^T \hat{\boldsymbol{\beta}}^{(b)} \right],$$

where $\hat{\boldsymbol{\beta}}^{(a)}, \hat{\boldsymbol{\beta}}^{(b)} \in \mathbb{R}^p$ are adjustment vectors for the treatment and control groups, respectively, as indicated by the superscripts. The terms $\bar{\mathbf{x}}_A - \bar{\mathbf{x}}$ and $\bar{\mathbf{x}}_B - \bar{\mathbf{x}}$ represent the fluctuation of the covariates in the subsample relative to the full sample, and the adjustment vectors fit the linear relationships between the covariates and potential outcomes under treatment and control. For example, in the PAC-man trial, this would help alleviate the imbalance in the pretreatment estimated probability of death: the corresponding element of $\bar{\mathbf{x}}_B - \bar{\mathbf{x}}$ would be positive (due to the higher average probability of death in the control group), the corresponding element of $\hat{\boldsymbol{\beta}}^{(b)}$ would be negative (a higher probability of death correlates with worse health outcomes), so the overall treatment effect estimate would be adjusted downward. This procedure is equivalent to imputing the unobserved potential outcomes; if we define

$$\hat{a}_B = \bar{a}_A + (\bar{\mathbf{x}}_B - \bar{\mathbf{x}})^T \hat{\boldsymbol{\beta}}^{(a)}, \quad \hat{b}_A = \bar{b}_B + (\bar{\mathbf{x}}_A - \bar{\mathbf{x}})^T \hat{\boldsymbol{\beta}}^{(b)},$$

we can form the equivalent estimator:

$$\widehat{ATE}_{\text{adj}} = n^{-1} (n_A \hat{a}_A + n_B \hat{a}_B) - n^{-1} (n_B \hat{b}_B + n_A \hat{b}_A).$$

If we consider these adjustment vectors to be fixed (nonrandom), or if they are derived from an independent data source, then this estimator is still unbiased, and may have substantially smaller asymptotic and finite-sample variance than the unadjusted estimator. This allows for construction of tighter confidence intervals for the true treatment effect.

In practice, the “ideal” linear adjustment vectors, leading to a minimum-variance estimator of the form of $\widehat{ATE}_{\text{adj}}$, cannot be computed from the observed data. However, they can be estimated, possibly at the expense of introducing modest finite-sample bias into the treatment effect estimate. In the classical setup, when the number of covariates is relatively small, OLS regression can be used. The asymptotic properties of this kind of estimator are explored under the Neyman–Rubin model in refs. 7, 11, and 12. We will follow a particular scheme that is studied in ref. 7 and shown to have favorable properties: we regress the outcome on treatment indicators, covariates, and treatment by covariate interactions. This is equivalent to running separate regressions in the treatment and control groups of outcome against an intercept and covariates. If we define $\hat{\boldsymbol{\beta}}_{\text{OLS}}^{(a)}$ and $\hat{\boldsymbol{\beta}}_{\text{OLS}}^{(b)}$ as the coefficients from the separate regressions, then the estimator is as follows:

$$\widehat{ATE}_{\text{OLS}} = \left[\bar{a}_A - (\bar{\mathbf{x}}_A - \bar{\mathbf{x}})^T \hat{\boldsymbol{\beta}}_{\text{OLS}}^{(a)} \right] - \left[\bar{b}_B - (\bar{\mathbf{x}}_B - \bar{\mathbf{x}})^T \hat{\boldsymbol{\beta}}_{\text{OLS}}^{(b)} \right].$$

This has some finite-sample bias, but ref. 7 shows that it vanishes quickly at the rate of $1/n$ under moment conditions on the potential outcomes and covariates. Moreover, for a fixed p , under regularity conditions, the inclusion of interaction terms guarantees that it never has higher asymptotic variance than the unadjusted estimator, and asymptotically conservative confidence intervals for the true parameter can be constructed.

In modern randomized trials, where a large number of covariates are recorded for each individual, p may be comparable to or even larger than n . In this case, OLS regression can overfit the data badly, or may even be ill posed, leading to estimators with large finite-sample variance. To remedy this, we propose estimating

the adjustment vectors using the Lasso (6). The adjustment vectors would take the following form:

$$\hat{\boldsymbol{\beta}}_{\text{Lasso}}^{(a)} = \arg \min_{\boldsymbol{\beta}} \left[\frac{1}{2n_A} \sum_{i \in A} \left(a_i - \bar{a}_A - (\mathbf{x}_i - \bar{\mathbf{x}}_A)^T \boldsymbol{\beta} \right)^2 + \lambda_a \sum_{j=1}^p |\beta_j| \right], \quad [1]$$

$$\hat{\boldsymbol{\beta}}_{\text{Lasso}}^{(b)} = \arg \min_{\boldsymbol{\beta}} \left[\frac{1}{2n_B} \sum_{i \in B} \left(b_i - \bar{b}_B - (\mathbf{x}_i - \bar{\mathbf{x}}_B)^T \boldsymbol{\beta} \right)^2 + \lambda_b \sum_{j=1}^p |\beta_j| \right], \quad [2]$$

and the proposed Lasso-adjusted ATE estimator is as follows:

$$\widehat{ATE}_{\text{Lasso}} = \left[\bar{a}_A - (\bar{\mathbf{x}}_A - \bar{\mathbf{x}})^T \hat{\boldsymbol{\beta}}_{\text{Lasso}}^{(a)} \right] - \left[\bar{b}_B - (\bar{\mathbf{x}}_B - \bar{\mathbf{x}})^T \hat{\boldsymbol{\beta}}_{\text{Lasso}}^{(b)} \right].$$

[To simplify the notation, we omit the dependence of $\hat{\boldsymbol{\beta}}_{\text{Lasso}}^{(a)}, \hat{\boldsymbol{\beta}}_{\text{Lasso}}^{(b)}$, λ_a , and λ_b on the population size n .] Here, λ_a and λ_b are regularization parameters for the Lasso, which must be chosen by the experimenter; simulations show that cross-validation (CV) works well. In the next section, we study this estimator under the Neyman–Rubin model and provide conditions on the potential outcomes, the covariates, and the regularization parameters under which $\widehat{ATE}_{\text{Lasso}}$ enjoys similar asymptotic and finite-sample advantages as $\widehat{ATE}_{\text{OLS}}$.

It is worth noting that, when two different adjustments are made for the treatment and control groups as in ref. 7 and here, the covariates do not have to be the same for the two groups. However, when they are not the same, the Lasso- or OLS-adjusted estimators are no longer guaranteed to have smaller or equal asymptotic variance than the unadjusted one, even in the case of fixed p . In practice, one may still choose between the adjusted and unadjusted estimators based on the widths of the corresponding confidence intervals.

Theoretical Results

Notation. For a vector $\boldsymbol{\beta} \in \mathbb{R}^p$ and a subset $S \subset \{1, \dots, p\}$, let β_j be the j th component of $\boldsymbol{\beta}$, $\boldsymbol{\beta}_S = (\beta_j : j \in S)^T$, S^c be the complement of S , and $|S|$ the cardinality of the set S . For any column vector $\mathbf{u} = (u_1, \dots, u_m)^T$, let $\|\mathbf{u}\|_2^2 = \sum_{i=1}^m u_i^2$, $\|\mathbf{u}\|_1 = \sum_{i=1}^m |u_i|$, $\|\mathbf{u}\|_\infty = \max_{i=1, \dots, m} |u_i|$, and $\|\mathbf{u}\|_0 = |\{j : u_j \neq 0\}|$. For a given $m \times m$ matrix D , let $\lambda_{\min}(D)$ and $\lambda_{\max}(D)$ be the smallest and largest eigenvalues of D , respectively, and D^{-1} , the inverse of the matrix D . Let \xrightarrow{d} and \xrightarrow{p} denote convergence in distribution and in probability, respectively.

Decomposition of the Potential Outcomes. The Neyman–Rubin model does not assume a linear relationship between the potential outcomes and the covariates. To study the properties of adjustment under this model, we decompose the potential outcomes into a term linear in the covariates and an error term. Given vectors of coefficients $\boldsymbol{\beta}^{(a)}, \boldsymbol{\beta}^{(b)} \in \mathbb{R}^p$, we write for $i = 1, \dots, n$,

$$a_i = \bar{a} + (\mathbf{x}_i - \bar{\mathbf{x}})^T \boldsymbol{\beta}^{(a)} + e_i^{(a)}, \quad [3]$$

$$b_i = \bar{b} + (\mathbf{x}_i - \bar{\mathbf{x}})^T \boldsymbol{\beta}^{(b)} + e_i^{(b)}. \quad [4]$$

[Again, we omit the dependence of $\boldsymbol{\beta}^{(a)}, \boldsymbol{\beta}^{(b)}$, λ_a , λ_b , $e_i^{(a)}$, and $e_i^{(b)}$ on n .]

Note that we have not added any assumptions to the model; we have simply defined unit-level residuals, $e_i^{(a)}$ and $e_i^{(b)}$, given the vectors $\boldsymbol{\beta}^{(a)}, \boldsymbol{\beta}^{(b)}$. All of the quantities in [3] and [4] are fixed, deterministic numbers. It is easy to verify that $\bar{e}^{(a)} = \bar{e}^{(b)} = 0$. To pursue a theory for the Lasso, we will add assumptions on the populations of a_i 's, b_i 's, and \mathbf{x}_i 's, and we will assume the

existence of $\beta^{(a)}, \beta^{(b)}$ such that the error terms satisfy certain assumptions.

Conditions. We will need the following to hold for both the treatment and control potential outcomes. The first set of assumptions (1–3) are similar to those found in ref. 7.

Condition 1: Stability of treatment assignment probability.

$$n_A/n \rightarrow p_A, \text{ as } n \rightarrow \infty, \quad [5]$$

for some $p_A \in (0, 1)$.

Condition 2: The centered moment conditions. There exists a fixed constant $L > 0$ such that, for all $n = 1, 2, \dots$ and $j = 1, \dots, p$,

$$n^{-1} \sum_{i=1}^n (x_{ij} - (\bar{x}_j))^4 \leq L; \quad [6]$$

$$n^{-1} \sum_{i=1}^n (e_i^{(a)})^4 \leq L; \quad n^{-1} \sum_{i=1}^n (e_i^{(b)})^4 \leq L. \quad [7]$$

Condition 3: The means $n^{-1} \sum_{i=1}^n (e_i^{(a)})^2$, $n^{-1} \sum_{i=1}^n (e_i^{(b)})^2$, and $n^{-1} \sum_{i=1}^n e_i^{(a)} e_i^{(b)}$ converge to finite limits.

Because we consider the high-dimensional setting where p is allowed to be much larger than n , we need additional assumptions to ensure that the Lasso is consistent for estimating $\beta^{(a)}$ and $\beta^{(b)}$. Before stating them, we define several quantities.

Definition 1: Given $\beta^{(a)}$ and $\beta^{(b)}$, the sparsity measures for treatment and control groups, $s^{(a)}$ and $s^{(b)}$, are defined as the number of nonzero elements of $\beta^{(a)}$ and $\beta^{(b)}$, i.e.,

$$s^{(a)} = \left| \left\{ j : \beta_j^{(a)} \neq 0 \right\} \right|, \quad s^{(b)} = \left| \left\{ j : \beta_j^{(b)} \neq 0 \right\} \right|, \quad [8]$$

respectively. We will allow $s^{(a)}$ and $s^{(b)}$ to grow with n , although the notation does not explicitly show this.

Definition 2: Define δ_n to be the maximum covariance between the error terms and the covariates.

$$\delta_n = \max_{\omega=a,b} \left\{ \max_j \left| \frac{1}{n} \sum_{i=1}^n (x_{ij} - (\bar{x}_j)) (e_i^{(\omega)} - \bar{e}^{(\omega)}) \right| \right\}. \quad [9]$$

The following conditions will guarantee that the Lasso consistently estimates the adjustment vectors $\beta^{(a)}, \beta^{(b)}$ at a fast enough rate to ensure asymptotic normality of $\widehat{ATE}_{\text{Lasso}}$. It is an open question whether a weaker form of consistency would be sufficient for our results to hold.

Condition 4: Decay and scaling. Let $s = \max\{s^{(a)}, s^{(b)}\}$.

$$\delta_n = o\left(\frac{1}{s\sqrt{\log p}}\right). \quad [10]$$

$$(s \log p) / \sqrt{n} = o(1). \quad [11]$$

Condition 5: Cone invertibility factor. Define the Gram matrix as $\Sigma = n^{-1} \sum_{i=1}^n (x_i - \bar{x})(x_i - \bar{x})^T$. There exist constants $C > 0$ and $\xi > 1$ not depending on n , such that

$$\|\mathbf{h}_S\|_1 \leq Cs \|\Sigma \mathbf{h}\|_\infty, \quad \forall \mathbf{h} \in \mathcal{C}, \quad [12]$$

with $\mathcal{C} = \{\mathbf{h} : \|\mathbf{h}_{S^c}\|_1 \leq \xi \|\mathbf{h}_S\|_1\}$, and

$$S = \left\{ j : \beta_j^{(a)} \neq 0 \text{ or } \beta_j^{(b)} \neq 0 \right\}. \quad [13]$$

Condition 6: Let $\tau = \min\{1/70, (3p_A)^2/70, (3-3p_A)^2/70\}$. For constants $0 < \eta < (\xi - 1)/(\xi + 1)$ and $1/\eta < M < \infty$, assume the

regularization parameters of the Lasso belong to the following sets:

$$\lambda_a \in \left(\frac{1}{\eta}, M\right] \times \left(\frac{2(1+\tau)L^{1/2}}{p_A} \sqrt{\frac{2 \log p}{n}} + \delta_n\right), \quad [14]$$

$$\lambda_b \in \left(\frac{1}{\eta}, M\right] \times \left(\frac{2(1+\tau)L^{1/2}}{p_B} \sqrt{\frac{2 \log p}{n}} + \delta_n\right). \quad [15]$$

Denote, respectively, the population variances of $e^{(a)}$ and $e^{(b)}$ and the population covariance between them by the following:

$$\sigma_{e^{(a)}}^2 = n^{-1} \sum_{i=1}^n (e_i^{(a)})^2, \quad \sigma_{e^{(b)}}^2 = n^{-1} \sum_{i=1}^n (e_i^{(b)})^2, \\ \sigma_{e^{(a)}e^{(b)}} = n^{-1} \sum_{i=1}^n e_i^{(a)} e_i^{(b)}.$$

Theorem 1. Assume Conditions 1–6 hold for some $\beta^{(a)}$ and $\beta^{(b)}$. Then,

$$\sqrt{n}(\widehat{ATE}_{\text{Lasso}} - ATE) \xrightarrow{d} \mathcal{N}(0, \sigma^2), \quad [16]$$

where

$$\sigma^2 = \lim_{n \rightarrow \infty} \left[\frac{1-p_A}{p_A} \sigma_{e^{(a)}}^2 + \frac{p_A}{1-p_A} \sigma_{e^{(b)}}^2 + 2\sigma_{e^{(a)}e^{(b)}} \right]. \quad [17]$$

The proof of *Theorem 1* is given in *SI Appendix*. It is easy to show, as in the following corollary of *Theorem 1*, that the asymptotic variance of $\widehat{ATE}_{\text{Lasso}}$ is no worse than $\widehat{ATE}_{\text{unadj}}$ when $\beta^{(a)}$ and $\beta^{(b)}$ are defined as coefficients of regressing potential outcomes on a subset of covariates. More specifically, suppose there exists a subset $J \subset \{1, \dots, p\}$, such that

$$\beta^{(a)} = \left((\beta_J^{(a)})^T, \mathbf{0} \right)^T, \quad \beta^{(b)} = \left((\beta_J^{(b)})^T, \mathbf{0} \right)^T, \quad [18]$$

where $\beta_J^{(a)}$ and $\beta_J^{(b)}$ are the population-level OLS coefficients for regressing the potential outcomes a and b on the covariates in the subset J with intercept, respectively.

Corollary 1. For $\beta^{(a)}$ and $\beta^{(b)}$ defined in [18] and some λ_a and λ_b , assume Conditions 1–6 hold. Then the asymptotic variance of $\sqrt{n} \widehat{ATE}_{\text{Lasso}}$ is no greater than that of the $\sqrt{n} \widehat{ATE}_{\text{unadj}}$. The difference is $(1/(p_A(1-p_A)))\Delta$, where

$$\Delta = - \lim_{n \rightarrow \infty} \|X\beta_E\|_2^2 \leq 0, \quad \beta_E = (1-p_A)\beta^{(a)} + p_A\beta^{(b)}. \quad [19]$$

Remark 1: If, instead of *Condition 6*, we assume that the covariates are uniformly bounded, i.e., $\max_{i,j} |x_{ij}| \leq L$, then the fourth moment condition on the error terms, given in [7], can be weakened to a second moment condition. Although we do not prove the necessity of any of our conditions, our simulation studies show that the distributions of the unadjusted and the Lasso-adjusted estimator may be nonnormal when (i) the covariates are generated from Gaussian distributions and the error terms do not satisfy second moment condition, e.g., being generated from a t distribution with one degree of freedom; or (ii) the covariates do not have bounded fourth moments, e.g., being generated from a t distribution with three degrees of freedom. See the histograms in Fig. 1, where the corresponding p values of Kolmogorov–Smirnov testing for normality are less than $2.2e - 16$. These findings indicate that our moment conditions cannot be dramatically weakened




Fig. 1. Histograms of the unadjusted estimator and the Lasso-adjusted estimator when the moment conditions do not hold. We select the tuning parameters for Lasso using 10-fold CV. The potential outcomes are simulated from linear regression model and then kept fixed (see more details in [SI Appendix](#)). For the upper two subplots, the error terms are generated from t distribution with one degree of freedom and therefore do not satisfy second moment condition; whereas for the lower two subplots, the covariates are generated from t distribution with three degrees of freedom and thus violate fourth moment condition.

for asymptotic normality. However, we also find that the Lasso-adjusted estimator still has smaller variance and mean squared error than the unadjusted estimator, even when these moment conditions do not hold. In practice, when the covariates do not have bounded fourth moments, one may perform some transformation—e.g., a logarithm transformation—to ensure that the transformed covariates have bounded fourth moments while having a sufficiently large variance so as to retain useful information. We leave it as future work to explore the properties of different transformations.

Remark 2: Statement [11], typically required in debiasing the Lasso (15), is stronger by a factor of $\sqrt{\log p}$ than the usual requirement for l_1 consistency of the Lasso.

Remark 3: *Condition 5* is slightly weaker than the typical restricted eigenvalue condition for analyzing the Lasso.

Remark 4: If we assume $\delta_n = O(\sqrt{\log p/n})$, which satisfies [10], then *Condition 6* requires that the tuning parameters are proportional to $\sqrt{\log p/n}$, which is typically assumed for the Lasso in the high-dimensional linear regression model.

Remark 5: For fixed p , $\delta_n = 0$ in [9], *Condition 4* holds automatically, and *Condition 5* holds when the smallest eigenvalue of Σ is uniformly bounded away from 0. In this case, *Corollary 1* reverts to corollary 1.1. in ref. 7. When these conditions are not satisfied, we should set λ_a and λ_b to be large enough to cause the Lasso-adjusted estimator to revert to the unadjusted one.

Neyman-Type Conservative Variance Estimate. We note that the asymptotic variance in *Theorem 1* involves the cross-product term $\sigma_{e^{(a)}e^{(b)}}$, which is not consistently estimable in the Neyman–Rubin model as a_i and b_i are never simultaneously observed. However, we can give a Neyman-type conservative estimate of the variance. Let

$$\hat{\sigma}_{e^{(a)}}^2 = \frac{1}{n_A - df^{(a)}} \sum_{i \in A} \left(a_i - \bar{a}_A - (\mathbf{x}_i - \bar{\mathbf{x}}_A)^T \hat{\boldsymbol{\beta}}_{\text{Lasso}}^{(a)} \right)^2,$$

$$\hat{\sigma}_{e^{(b)}}^2 = \frac{1}{n_B - df^{(b)}} \sum_{i \in B} \left(b_i - \bar{b}_B - (\mathbf{x}_i - \bar{\mathbf{x}}_B)^T \hat{\boldsymbol{\beta}}_{\text{Lasso}}^{(b)} \right)^2,$$

where $df^{(a)}$ and $df^{(b)}$ are degrees of freedom defined by the following:

$$df^{(a)} = \hat{s}^{(a)} + 1 = \left\| \hat{\boldsymbol{\beta}}_{\text{Lasso}}^{(a)} \right\|_0 + 1; \quad df^{(b)} = \hat{s}^{(b)} + 1 = \left\| \hat{\boldsymbol{\beta}}_{\text{Lasso}}^{(b)} \right\|_0 + 1.$$

Define the variance estimate of $\sqrt{n}(\widehat{ATE}_{\text{Lasso}} - ATE)$ as follows:

$$\hat{\sigma}_{\text{Lasso}}^2 = \frac{n}{n_A} \hat{\sigma}_{e^{(a)}}^2 + \frac{n}{n_B} \hat{\sigma}_{e^{(b)}}^2. \quad [20]$$

We will show in [SI Appendix, Theorem S1](#), that the limit of $\hat{\sigma}_{\text{Lasso}}^2$ is greater than or equal to the asymptotic variance of $\sqrt{n}(\widehat{ATE}_{\text{Lasso}} - ATE)$, and therefore can be used to construct a conservative confidence interval for the ATE.

Related Work. The Lasso has already made several appearances in the literature on treatment effect estimation. In the context of observational studies, ref. 15 constructs confidence intervals for preconceived effects or their contrasts by debiasing the Lasso-adjusted regression, ref. 16 employs the Lasso as a formal method for selecting adjustment variables via a two-stage procedure that concatenates features from models for treatment and outcome, and similarly, ref. 17 gives very general results for estimating a wide range of treatment effect parameters, including the case of instrumental variables estimation. In addition to the Lasso, ref. 18 considers nonparametric adjustments in the estimation of ATE. In works such as these, which deal with observational studies, confounding is the major issue. With confounding, the naive difference-in-means estimator is biased for the true treatment effect, and adjustment is used to form an unbiased estimator. However, in our work, which focuses on a randomized trial, the difference-in-means estimator is already unbiased; adjustment reduces the variance while, in fact, introducing a small amount of finite-sample bias. Another major difference between this prior work and ours is the sampling framework: we operate within the Neyman–Rubin model with fixed potential outcomes for a finite population, where the treatment group is sampled without replacement, whereas these papers assume independent sampling from a probability distribution with random error terms.

Our work is related to the estimation of heterogeneous or subgroup-specific treatment effects, including interaction terms to allow the imputed individual-level treatment effects to vary according to some linear combination of covariates. This is pursued in the high-dimensional setting in ref. 19; this work advocates solving the Lasso on a reduced set of modified covariates, rather than the full set of covariate by treatment interactions, and includes extensions to binary outcomes and survival data. The recent work in ref. 20 considers the problem of designing multiple-testing procedures for detecting subgroup-specific treatment effects; they pose this as an optimization over testing procedures where constraints are added to enforce guarantees on type I error rate and power to detect effects. Again, the sampling framework in these works is distinct from ours; they do not use the Neyman–Rubin model as a basis for designing the methods or investigating their properties.

PAC Data Illustration and Simulations. We now return to the PAC-man study introduced earlier. We examine the data in more

ATE estimates for PAC data

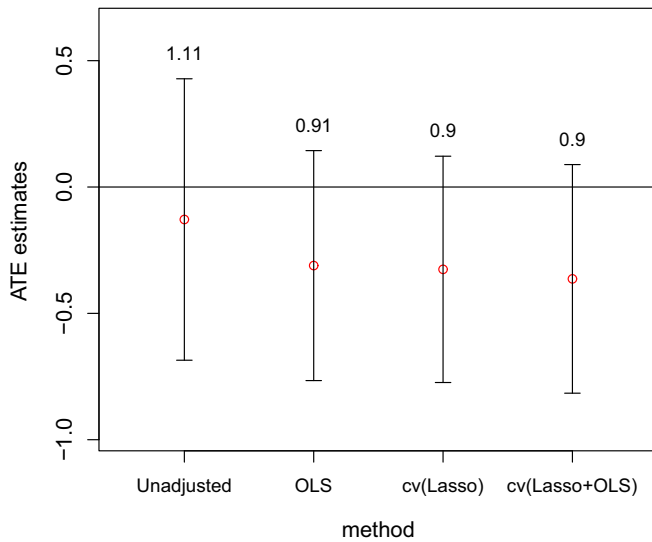


Fig. 2. ATE estimates (red circles) and 95% confidence intervals (bars) for the PAC data. The numbers above each bar are the corresponding interval lengths.

detail and explore the results of several adjustment procedures. There were 1,013 patients in the PAC-man study: 506 treated (managed with PAC) and 507 control (managed without PAC, but retaining the option of using alternative devices). The outcome variable is quality-adjusted life years (QALYs). One QALY represents 1 year of life in full health; in-hospital death corresponds to a QALY of zero. We have 59 covariates about each individual in the study; we include all main effects as well as 1,113 two-way interactions, and form a design matrix \mathbf{X} with 1,172 columns and 1,013 rows. See *SI Appendix* for more details on the design matrix.

The assumptions that underpin the theoretical guarantees of the ATE_{Lasso} estimator are, in practice, not explicitly checkable, but we attempt to inspect the quantities that are involved in the conditions to help readers make their own judgement. The uniform bounds on the fourth moments refer to a hypothetical sequence of populations; these cannot be verified given that the investigator has a single dataset. However, as an approximation, the fourth moments of the data can be inspected to ensure that they are not too large. In this dataset, the maximum fourth moment of the covariates is 37.3, which is indicative of a heavy-tailed and potentially destabilizing covariate; however, it occurs in an interaction term not selected by the Lasso, and thus does

not influence the estimate. [The fourth moments of the covariates are shown in *SI Appendix*, Fig. S9. The covariates with the largest two fourth moments (37.3 and 34.9, respectively) are quadratic term $interactnew^2$ and interaction term $IMscorecrct : systemnew$. Neither of them is selected by the Lasso to do the adjustment.] Checking the conditions for high-dimensional consistency of the Lasso would require knowledge of the unknown active set S , and moreover, even if it were known, calculating the cone invertibility factor would involve an infeasible optimization. This is a general issue in the theory of sparse linear high-dimensional estimation. To approximate these conditions, we use the bootstrap to estimate the active set of covariates S and the error terms $e^{(a)}$ and $e^{(b)}$. See *SI Appendix* for more details. Our estimated S contains 16 covariates and the estimated second moments of $e^{(a)}$ and $e^{(b)}$ are 11.8 and 12.0, respectively. The estimated maximal covariance δ_n equals 0.34 and the scaling $(s \log p)/\sqrt{n}$ is 3.55. Although this is not close to zero, we should mention that the estimation of δ_n and $(s \log p)/\sqrt{n}$ can be unstable and less accurate because it is based on a subsample of the population. As an approximation to *Condition 5*, we examine the largest and smallest eigenvalues of the sub-Gram matrix $(1/n)\mathbf{X}_S^T \mathbf{X}_S$, which are 2.09 and 0.18, respectively. Thus, the quantity in *Condition 5* seems reasonably bounded away from zero.

We now estimate the ATE using the unadjusted estimator, the Lasso-adjusted estimator, and the OLS-adjusted estimator, which is computed based on a subdesign matrix containing only the 59 main effects. We also present results for the two-step estimator $\widehat{ATE}_{Lasso+OLS}$, which adopts the Lasso to select covariates and then uses OLS to refit the regression coefficients. In the next paragraph and in *SI Appendix*, *Algorithm 1*, we show how we adapt the CV procedure to select the tuning parameter for $\widehat{ATE}_{Lasso+OLS}$ based on a combined performance of Lasso and OLS, or $cv(Lasso+OLS)$.

We use the R package “glmnet” to compute the Lasso solution path and select the tuning parameters λ_a and λ_b by 10-fold CV. To indicate the method of selecting tuning parameters, we denote the corresponding estimators as $cv(Lasso)$ and $cv(Lasso+OLS)$, respectively. We should mention that for the $cv(Lasso+OLS)$ -adjusted estimator, we compute the CV error for a given value of λ_a (or λ_b) based on the whole Lasso+OLS procedure instead of just the Lasso estimator (*SI Appendix*, *Algorithm 1*). Therefore, the $cv(Lasso+OLS)$ and the $cv(Lasso)$ may select different covariates to do the adjustment. This type of CV requires more computation than the CV based on just the Lasso estimator because it needs to compute the OLS estimator for each fold and each given λ_a (or λ_b), but it can give better prediction and model selection performance.

Fig. 2 presents the ATE estimates along with 95% confidence intervals. The interval lengths are shown on top of each interval bar. All of the methods give confidence intervals containing 0; hence, this experiment failed to provide sufficient evidence to

Table 1. Selected covariates for adjustment

Method	Treatment	Covariates
$cv(Lasso+OLS)$	Treated	Age, p_death, age · age, age:p_death
$cv(Lasso+OLS)$	Control	Age, p_death, age · age, age:p_death, p_death:mch_vent
$cv(Lasso)$	Treated	Pac_rate, age, p_death, age · age, p_death · p_death, region:im_score, region:systemnew, pac_rate:age, pac_rate:p_death, pac_rate:systemnew, im_score:interactnew, age:p_death, age:glasgow, age:systemnew, interactnew:systemnew, pac_rate:creatinine, age:mch_vent, age:respiratory, age:creatinine, interactnew:mch_vent, interactnew:male, glasgow:organ_failure, p_death:mch_vent, systemnew:male
$cv(Lasso)$	Control	Age, p_death, age · age, unitsize:p_death, pac_rate:systemnew, age:p_death, interactnew:mch_vent, p_death:mch_vent*

*Covariate definitions: age, patient’s age; p_death, baseline probability of death; mch_vent, mechanical ventilation at admission; region, geographic region; pac_rate, PAC rate in unit; creatinine, respiratory, glasgow, interactnew, organ_failure, systemnew, and im_score, various physiological indicators.




Fig. 3. Selection stability comparison of $cv(Lasso)$ and $cv(Lasso+OLS)$ for treatment group.

reject the hypothesis that PAC did not have an effect on patient QALYs (either positive or negative). Because the caretakers of patients managed without PAC retained the option of using less invasive cardiac output monitoring devices, such an effect may have been particularly hard to detect in this experiment.

However, it is interesting to note that, compared with the unadjusted estimator, the OLS-adjusted estimator causes the ATE estimate to decrease (from -0.13 to -0.31), and shortens the confidence interval by about 20%. This is due mainly to the imbalance in the pretreatment probability of death, which was highly predictive of the posttreatment QALYs. The $cv(Lasso)$ -adjusted estimator yields a comparable ATE estimate and confidence interval, but the fitted model is more interpretable and parsimonious than the OLS model: it selects 24 and 8 covariates for treated and control, respectively. The $cv(Lasso+OLS)$ estimator selects even fewer covariates: 4 and 5 for treated and control, respectively, but performs a similar adjustment as the $cv(Lasso)$ (see the comparison of fitted values in *SI Appendix, Fig. S8*). We also note that these adjustments agree with the one performed in ref. 13, where the treatment effect was adjusted downward to -0.27 after stratifying into four groups based on predicted probability of death.

The covariates selected by Lasso for adjustment are shown in Table 1, where “A·A” denotes quadratic term of the covariate A, and “A:B” denotes two-way interaction between two covariates A and B. Among them, patient’s age and estimated probability of death (p_{death}), together with the quadratic term “age·age” and interactions “age: p_{death} ” and “ p_{death} :mech_vent” (mechanical ventilation at admission), are the most important covariates for the adjustment. The patients in control group are slightly older and have slightly higher risk of death. These covariates are important predictors of the outcome. Therefore, the unadjusted estimator may overestimate the benefits of receiving PAC.

Because not all of the potential outcomes are observed, we cannot know the true gains of adjustment methods. However, we can estimate the gains via building a simulated set of potential outcomes by matching treated units to control units on observed covariates. We use the matching method described in ref. 21,

which gives 1,013 observations with all potential outcomes imputed. We match on the 59 main effects only. The ATE is -0.29 . We then use this synthetic dataset to calculate the biases, SDs, and root-mean-square errors (\sqrt{MSE}) of different ATE estimators based on 25,000 replicates of a completely randomized experiment, which assigns 506 subjects to the treated group and the remainders to the control group.

SI Appendix, Table S5, shows the results. For all of the methods, the bias is substantially smaller (by a factor of 100) than the SD. The SD and \sqrt{MSE} of the OLS-adjusted estimator are both 10.2% smaller than those of the unadjusted estimator, whereas the $cv(Lasso)$ - and $cv(Lasso+OLS)$ -adjusted estimators further improve the SD and \sqrt{MSE} of the OLS-adjusted estimator by $\sim 4.7\%$. Moreover, all these methods provide conservative confidence intervals with coverage probabilities higher than 99%. However, the interval lengths of the OLS-, $cv(Lasso)$ -, and $cv(Lasso+OLS)$ -adjusted estimator are comparable and are $\sim 10\%$ shorter than that of the unadjusted estimator. The $cv(Lasso+OLS)$ -adjusted estimator is similar to the $cv(Lasso)$ -adjusted estimator in terms of mean squared error, confidence interval length, and coverage probability, but outperforms the latter with much fewer and more stable covariates in the adjustment (see Figs. 3 and 4 for the selection frequency of each covariate for treatment group and control group, respectively). We also show in *SI Appendix, Fig. S10*, that the sampling distribution of the estimates is very close to Normal.

We conduct additional simulation studies to evaluate the finite sample performance of \widehat{ATE}_{Lasso} and compare it with that of the OLS-adjusted estimator and the unadjusted estimator. A qualitative analysis of these simulations yields the same conclusions as presented above; however, for the sake of brevity, we defer the simulation details in *SI Appendix*.

Discussion

We study the Lasso-adjusted ATE estimate under the Neyman–Rubin model for randomization. Our purpose in using the Neyman–Rubin model is to investigate the performance of the Lasso under a realistic sampling framework that does not impose strong assumptions on the data. We provide conditions that ensure

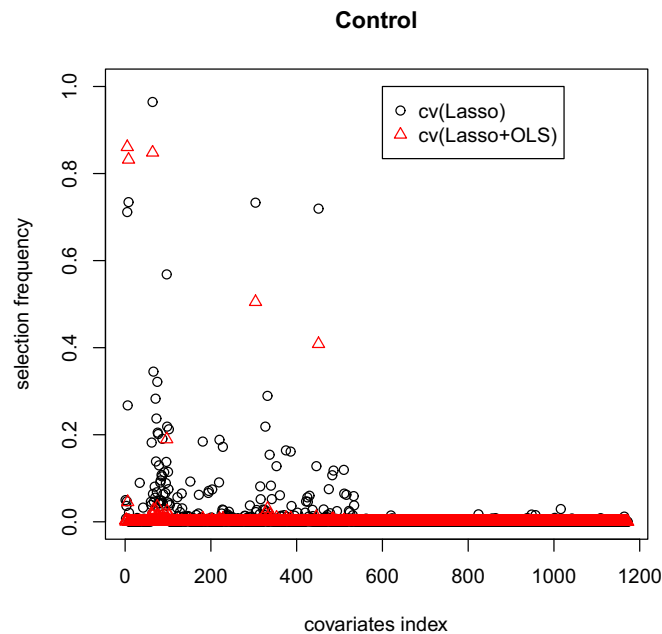


Fig. 4. Selection stability comparison of $cv(Lasso)$ and $cv(Lasso+OLS)$ for control group.

asymptotic normality, and provide a Neyman-type estimate of the asymptotic variance that can be used to construct a conservative confidence interval for the ATE. Although we do not require an explicit generative linear model to hold, our theoretical analysis requires the existence of latent “adjustment vectors” such that moment conditions of the error terms are satisfied, and that the cone invertibility condition of the sample covariance matrix is satisfied in addition to moment conditions for OLS adjustment as in ref. 7. Both assumptions are difficult to check in practice. In our theory, we do not address whether these assumptions are necessary for our results to hold, although simulations indicate that the moment conditions cannot be substantially weakened. As a by-product of our analysis, we extend Massart’s concentration inequality for sampling without replacement, which is useful for theoretical analysis under the Neyman–Rubin model. Simulation studies and the real-data illustration show the advantage of the Lasso-adjusted estimator in terms of estimation accuracy and model interpretation. In practice, we recommend a variant of Lasso, $cv(\text{Lasso}+\text{OLS})$, to select covariates and perform the adjustment, because it gives similar coverage probability and confidence interval length compared with $cv(\text{Lasso})$, but with far fewer covariates selected. In future work, we plan to extend our analysis to other popular methods in high-dimensional statistics such as Elastic-Net and ridge regression, which may be more appropriate for estimating adjusted ATE under different assumptions.

The main goal of using Lasso in this paper is to reduce the variance (and overall mean squared error) of ATE estimation. Another important task is to estimate heterogenous treatment effects and provide conditional treatment effect estimates for subpopulations. When the Lasso models of treatment and

control outcomes are different, both in variables selected and coefficient values, this could be interpreted as modeling treatment effect heterogeneity in terms of covariates. However, reducing variance of the ATE estimate and estimating heterogenous treatment effects have completely different targets. Targeting heterogenous treatment effects may result in more variable ATE estimates. Moreover, our simulations show that the set of covariates selected by the Lasso is unstable, and this may cause problems when interpreting them as evidence of heterogenous treatment effects. How best to estimate such effects is an open question that we would like to study in future research.

Materials and Methods

We did not conduct the PAC-man experiment, and we are analyzing secondary data without any personal identifying information. As such, this study is exempt from human subjects review. The original experiments underwent human subjects review in the United Kingdom (2).

ACKNOWLEDGMENTS. We thank David Goldberg for helpful discussions, Rebecca Barter for copyediting and suggestions for clarifying the text, and Winston Lin for comments. We thank Richard Grieve [London School of Hygiene and Tropical Medicine (LSHTM)], Sheila Harvey (LSHTM), David Harrison [Intensive Care National Audit and Research Centre (ICNARC)], and Kathy Rowan (ICNARC) for access to data from the PAC-Man Cost Effectiveness Analysis and the ICNARC Case Mix Programme database. This research was partially supported by NSF Grants DMS-11-06753, DMS-12-09014, DMS-1107000, DMS-1129626, DMS-1209014, Computational and Data-Enabled Science and Engineering in Mathematical and Statistical Sciences 1228246, DMS-1160319 (Focused Research Group); AFOSR Grant FA9550-14-1-0016; NSA Grant H98230-15-1-0040; the Center for Science of Information, a US NSF Science and Technology Center, under Grant Agreement CCF-0939370; Department of Defense for Office of Naval Research Grant N00014-15-1-2367; and the National Defense Science and Engineering Graduate Fellowship Program.

- Splawa-Neyman J, Dabrowska DM, Speed TP (1990) On the application of probability theory to agricultural experiments. Essay on principles. Section 9. *Stat Sci* 5(4):465–472.
- Harvey S, et al.; PAC-Man Study Collaboration (2005) Assessment of the clinical effectiveness of pulmonary artery catheters in management of patients in intensive care (PAC-Man): A randomised controlled trial. *Lancet* 366(9484):472–477.
- Dalen JE (2001) The pulmonary artery catheter—friend, foe, or accomplice? *JAMA* 286(3):348–350.
- Connors AF, Jr, et al.; SUPPORT Investigators (1996) The effectiveness of right heart catheterization in the initial care of critically ill patients. *JAMA* 276(11):889–897.
- Permutt T (1990) Testing for imbalance of covariates in controlled experiments. *Stat Med* 9(12):1455–1462.
- Tibshirani R (1994) Regression selection and shrinkage via the Lasso. *J R Stat Soc B* 58(1):267–288.
- Lin W (2013) Agnostic notes on regression adjustments to experimental data: Re-examining Freedman’s critique. *Ann Appl Stat* 7(1):295–318.
- Bühlmann P, Van De Geer S (2011) *Statistics for High-Dimensional Data: Methods, Theory and Applications* (Springer Science and Business Media, Berlin).
- Rubin DB (1974) Estimating causal effects of treatments in randomized and non-randomized studies. *J Educ Psychol* 66(5):688–701.
- Freedman DA (2008) On regression adjustments to experimental data. *Adv Appl Math* 40(2):180–193.
- Freedman DA (2008) On regression adjustments in experiments with several treatments. *Ann Appl Stat* 2(1):176–196.
- Freedman DA (2008) Randomization does not justify logistic regression. *Stat Sci* 23(2): 237–249.
- Miratrix LW, Sekhon JS, Yu B (2013) Adjusting treatment effect estimates by post-stratification in randomized experiments. *J R Stat Soc Series B Stat Methodol* 75(2): 369–396.
- Holland PW (1986) Statistics and causal inference. *J Am Stat Assoc* 81(396):945–960.
- Zhang CH, Zhang SS (2014) Confidence intervals for low dimensional parameters in high dimensional linear models. *J R Stat Soc Series B Stat Methodol* 76(1):217–242.
- Belloni A, Chernozhukov V, Hansen C (2013) Inference on treatment effects after selection among high-dimensional controls. *Rev Econ Stud* 81(2):608–650.
- Belloni A, Chernozhukov V, Fernández-Val I, Hansen C (2013) . Program evaluation with high-dimensional data. arXiv:1311.2645.
- Li L, Tchetgen ET, van der Vaart A, Robins JM (2011) Higher order inference on a treatment effect under low regularity conditions. *Stat Probab Lett* 81(7):821–828.
- Tian L, Alizadeh A, Gentles A, Tibshirani R (2014) A simple method for detecting interactions between a treatment and a large number of covariates. *J Am Stat Assoc* 109(508):1517–1532.
- Rosenblum M, Liu H, Yen EH (2014) Optimal tests of treatment effects for the overall population and two subpopulations in randomized trials, using sparse linear programming. *J Am Stat Assoc* 109(507):1216–1228.
- Diamond A, Sekhon JS (2013) Genetic matching for estimating causal effects: A general multivariate matching method for achieving balance in observational studies. *Rev Econ Stat* 95(3):932–945.