

AGNOSTIC NOTES ON REGRESSION ADJUSTMENTS TO EXPERIMENTAL DATA: REEXAMINING FREEDMAN'S CRITIQUE

BY WINSTON LIN

University of California, Berkeley

Freedman [*Adv. in Appl. Math.* **40** (2008) 180–193; *Ann. Appl. Stat.* **2** (2008) 176–196] critiqued ordinary least squares regression adjustment of estimated treatment effects in randomized experiments, using Neyman's model for randomization inference. Contrary to conventional wisdom, he argued that adjustment can lead to worsened asymptotic precision, invalid measures of precision, and small-sample bias. This paper shows that in sufficiently large samples, those problems are either minor or easily fixed. OLS adjustment cannot hurt asymptotic precision when a full set of treatment-covariate interactions is included. Asymptotically valid confidence intervals can be constructed with the Huber–White sandwich standard error estimator. These results hold under the same assumptions used by Freedman. The strongest reasons to support his preference for unadjusted estimates are transparency and the dangers of specification search.

Data from Angrist, Lang, and Oreopoulos's [*Am. Econ. J.: Appl. Econ.* **1:1** (2009) 136–163] evaluation of strategies to improve college students' achievement are used to illustrate checks on the asymptotic approximations.

1. Introduction. One of the attractions of randomized experiments is that, ideally, the strength of the design reduces the need for statistical modeling. Simple comparisons of means can be used to estimate the average effects of assigning subjects to treatment. Nevertheless, many researchers use linear regression models to adjust for random differences between the baseline characteristics of the treatment groups. The usual rationale is that adjustment tends to improve precision if the sample is large enough and the covariates are correlated with the outcome; this argument, which assumes that the regression model is correct, stems from Fisher (1932) and is taught to applied researchers in many fields. At research firms that conduct randomized experiments to evaluate social programs, adjustment is standard practice.¹

In an important and influential critique, Freedman (2008ab) analyzes the behavior of ordinary least squares regression-adjusted estimates without assuming a regression model. He uses Neyman's (1923) model for randomization inference:

Keywords and phrases: Analysis of covariance, covariate adjustment, randomization inference, sandwich estimator, robust standard errors, social experiments, program evaluation.

¹Cochran (1957), Cox and McCullagh (1982), Raudenbush (1997), and Klar and Darlington (2004) discuss precision improvement. Greenberg and Shroder (2004) document the use of regression adjustment in many randomized social experiments.

treatment effects can vary across subjects, linearity is not assumed, and random assignment is the source of variability in estimated average treatment effects. Freedman shows that (i) adjustment can actually worsen asymptotic precision, (ii) the conventional OLS standard error estimator is inconsistent, and (iii) the adjusted treatment effect estimator has a small-sample bias. He writes [Freedman (2008a)], “The reason for the breakdown is not hard to find: randomization does not justify the assumptions behind the OLS model.”

This paper offers an alternative perspective. Although I agree with Freedman’s (2008b) general advice (“Regression estimates . . . should be deferred until rates and averages have been presented”), I argue that in sufficiently large samples, the statistical problems he raised are either minor or easily fixed. Under the Neyman model with Freedman’s regularity conditions, I show that (i) OLS adjustment cannot hurt asymptotic precision when a full set of treatment \times covariate interactions is included, and (ii) the Huber–White sandwich standard error estimator is consistent or asymptotically conservative (regardless of whether the interactions are included). I also briefly discuss the small-sample bias issue and the distinction between unconditional and conditional unbiasedness.

Even the traditional OLS adjustment has benign large-sample properties when subjects are randomly assigned to two groups of equal size. Freedman (2008a) shows that in this case, adjustment (without interactions) improves or does not hurt asymptotic precision, and the conventional standard error estimator is consistent or asymptotically conservative. However, Freedman and many excellent applied statisticians in the social sciences have summarized his papers in terms that omit these results and emphasize the dangers of adjustment. For example, Berk et al. (2010) write: “Random assignment does not justify any form of regression with covariates. If regression adjustments are introduced nevertheless, there is likely to be bias in any estimates of treatment effects and badly biased standard errors.”

One aim of this paper is to show that such a negative view is not always warranted. A second aim is to help provide a more intuitive understanding of the properties of OLS adjustment when the regression model is incorrect. An “agnostic” view of regression [Angrist and Imbens (2002); Angrist and Pischke (2009, ch. 3)] is adopted here: without taking the regression model literally, we can still make use of properties of OLS that do not depend on the model assumptions.

1.1. *Precedents.* Similar results on the asymptotic precision of OLS adjustment with interactions are proved in interesting and useful papers by Yang and Tsiatis (2001), Tsiatis et al. (2008), and Schochet (2010), under the assumption

that the subjects are a random sample from an infinite superpopulation.² These results are not widely known, and Freedman was apparently unaware of them. He did not analyze adjustment with interactions, but conjectured, “Treatment by covariate interactions can probably be accommodated too” [Freedman (2008b, p. 186)].

Like Freedman, I use the Neyman model, in which random assignment of a finite population is the sole source of randomness; for a thoughtful philosophical discussion of finite- vs. infinite-population inference, see Reichardt and Gollob (1999, pp. 125–127). My purpose is not to advocate finite-population inference, but to show just how little needs to be changed to address Freedman’s major concerns. The results may help researchers understand why and when OLS adjustment can backfire. In large samples, the essential problem is omission of treatment \times covariate interactions, not the linear model. With a balanced two-group design, even that problem disappears asymptotically, because two wrongs make a right (underadjustment of one group mean cancels out overadjustment of the other).

Neglected parallels between regression adjustment in experiments and regression estimators in survey sampling turn out to be very helpful for intuition.

2. Basic framework. For simplicity, the main results in this paper assume a completely randomized experiment with two treatment groups (or a treatment group and a control group), as in Freedman (2008a). Results for designs with more than two groups are discussed informally.

2.1. The Neyman model with covariates. The notation is adapted from Freedman (2008b). There are n subjects, indexed by $i = 1, \dots, n$. We assign a simple random sample of fixed size n_A to treatment A and the remaining $n - n_A$ subjects to treatment B . For each subject, we observe an outcome Y_i and a row vector of covariates $\mathbf{z}_i = (z_{i1}, \dots, z_{iK})$, where $1 \leq K < \min(n_A, n - n_A) - 1$. Treatment does not affect the covariates.

Assume that each subject has two potential outcomes [Neyman (1923); Rubin (1974, 2005); Holland (1986)], a_i and b_i , which would be observed under treatments A and B , respectively.³ Thus, the observed outcome is $Y_i = a_i T_i + b_i (1 - T_i)$, where T_i is a dummy variable for treatment A .

Random assignment is the sole source of randomness in this model. The n subjects are the population of interest; they are not assumed to be randomly drawn from a superpopulation. For each subject, a_i , b_i , and \mathbf{z}_i are fixed, but T_i and thus Y_i are random.

²Although Tsiatis et al. write that OLS adjustment *without* interactions “is generally more precise than ... the difference in sample means” (p. 4661), Yang and Tsiatis’s asymptotic variance formula correctly implies that this adjustment may help or hurt precision.

³Most authors use notation such as $Y_i(1)$ and $Y_i(0)$, or Y_{1i} and Y_{0i} , for potential outcomes. Freedman’s (2008b) choice of a_i and b_i helps make the finite-population asymptotics more readable.

Let \bar{a} , \bar{a}_A , and \bar{a}_B denote the means of a_i over the population, treatment group A , and treatment group B :

$$\bar{a} = \frac{1}{n} \sum_{i=1}^n a_i, \quad \bar{a}_A = \frac{1}{n_A} \sum_{i \in A} a_i, \quad \bar{a}_B = \frac{1}{n - n_A} \sum_{i \in B} a_i.$$

Use similar notation for the means of b_i , Y_i , \mathbf{z}_i , and other variables.

Our goal is to estimate the average treatment effect of A relative to B :

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

2.2. Estimators of average treatment effect. The unadjusted or difference-in-means estimator of ATE is

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

The usual OLS-adjusted estimator of ATE is the estimated coefficient on T_i from the OLS regression of Y_i on T_i and \mathbf{z}_i . (All regressions described in this paper include intercepts.) Let $\widehat{ATE}_{\text{adj}}$ denote this estimator.

A third estimator, $\widehat{ATE}_{\text{interact}}$, can be computed as the estimated coefficient on T_i from the OLS regression of Y_i on T_i , \mathbf{z}_i , and $T_i(\mathbf{z}_i - \bar{\mathbf{z}})$. Section 3 motivates this estimator by analogy with regression estimators in survey sampling. In the context of observational studies, Imbens and Wooldridge (2009, pp. 28–30) give a theoretical analysis of $\widehat{ATE}_{\text{interact}}$, and a related method is known as the Peters–Belson or Oaxaca–Blinder estimator.⁴ When \mathbf{z}_i is a set of indicators for the values of a categorical variable, $\widehat{ATE}_{\text{interact}}$ is equivalent to subclassification or poststratification [Miratrix, Sekhon, and Yu (2012)].

3. Connections with sampling. Cochran (1977, ch. 7) gives a very readable discussion of regression estimators in sampling.⁵ In one example [Watson (1937)], the goal was to estimate \bar{y} , the average surface area of the leaves on a plant. Measuring a leaf’s area is time-consuming, but its weight can be found quickly. So the researcher weighed all the leaves, but measured area for only a small sample. In simple random sampling, the sample mean area \bar{y}_S is an unbiased estimator of \bar{y} . But \bar{y}_S ignores the auxiliary data on leaf weights. The sample and population mean weights (\bar{z}_S and \bar{z}) are both known, and if $\bar{z} > \bar{z}_S$, then we expect that $\bar{y} > \bar{y}_S$. This motivates a “linear regression estimator”

$$(3.1) \quad \widehat{y}_{\text{reg}} = \bar{y}_S + q(\bar{z} - \bar{z}_S)$$

⁴See Cochran (1969), Rubin (1984), and Kline (2011). Hansen and Bowers (2009) analyze a randomized experiment with a variant of the Peters–Belson estimator derived from logistic regression.

⁵See also Fuller (2002, 2009).

where q is an adjustment factor. One way to choose q is to regress leaf area on leaf weight in the sample.

Regression adjustment in randomized experiments can be motivated analogously under the Neyman model. The potential outcome a_i is measured for only a simple random sample (treatment group A), but the covariates \mathbf{z}_i are measured for the whole population (the n subjects). The sample mean \bar{a}_A is an unbiased estimator of \bar{a} , but it ignores the auxiliary data on \mathbf{z}_i . If the covariates are of some help in predicting a_i , then another estimator to consider is

$$(3.2) \quad \widehat{a}_{\text{reg}} = \bar{a}_A + (\bar{\mathbf{z}} - \bar{\mathbf{z}}_A)\mathbf{q}_a$$

where \mathbf{q}_a is a $K \times 1$ vector of adjustment factors. Similarly, we can consider using

$$(3.3) \quad \widehat{b}_{\text{reg}} = \bar{b}_B + (\bar{\mathbf{z}} - \bar{\mathbf{z}}_B)\mathbf{q}_b$$

to estimate \bar{b} and then $\widehat{a}_{\text{reg}} - \widehat{b}_{\text{reg}}$ to estimate $ATE = \bar{a} - \bar{b}$.

The analogy suggests deriving \mathbf{q}_a and \mathbf{q}_b from OLS regressions of a_i on \mathbf{z}_i in treatment group A and b_i on \mathbf{z}_i in treatment group B —in other words, separate regressions of Y_i on \mathbf{z}_i in the two treatment groups. The estimator $\widehat{a}_{\text{reg}} - \widehat{b}_{\text{reg}}$ is then just $\widehat{ATE}_{\text{interact}}$. If, instead, we use a pooled regression of Y_i on T_i and \mathbf{z}_i to derive a single vector $\mathbf{q}_a = \mathbf{q}_b$, then we get $\widehat{ATE}_{\text{adj}}$.

Connections between regression adjustment in experiments and regression estimators in sampling have been noted but remain underexplored.⁶ All three of the issues that Freedman raised have parallels in the sampling literature. Under simple random sampling, when the regression model is incorrect, OLS adjustment of the estimated mean still improves or does not hurt asymptotic precision [Cochran (1977)], consistent standard error estimators are available [Fuller (1975)], and the adjusted estimator of the mean has a small-sample bias [Cochran (1942)].

4. Asymptotic precision.

4.1. *Precision improvement in sampling.* This subsection gives an informal argument, adapted from Cochran (1977), to show that in simple random sampling, OLS adjustment of the sample mean improves or does not hurt asymptotic precision, even when the regression model is incorrect. Regularity conditions and other technical details are omitted; the purpose is to motivate the results on completely randomized experiments in Section 4.2.

⁶Connections are noted by Fienberg and Tanur (1987), Hansen and Bowers (2009), and Middleton and Aronow (2012) but are not mentioned by Cochran despite his important contributions to both literatures. He takes a design-based (agnostic) approach in much of his work on sampling, but assumes a regression model in his classic overview of regression adjustment in experiments and observational studies [Cochran (1957)].

First imagine using a “fixed-slope” regression estimator, where q in Eq. (3.1) is fixed at some value q_0 before sampling:

$$\widehat{y}_f = \bar{y}_S + q_0(\bar{z} - \bar{z}_S).$$

If $q_0 = 0$, \widehat{y}_f is just \bar{y}_S . More generally, \widehat{y}_f is the sample mean of $y_i - q_0(z_i - \bar{z})$, so its variance follows the usual formula with a finite-population correction:

$$\text{var}(\widehat{y}_f) = \frac{N-n}{N-1} \frac{1}{n} \frac{1}{N} \sum_{i=1}^N [(y_i - \bar{y}) - q_0(z_i - \bar{z})]^2$$

where N is the population size and n is the sample size.

Thus, choosing q_0 to minimize the variance of \widehat{y}_f is equivalent to running an OLS regression of y_i on z_i in the population. The solution is the “population least squares” slope,

$$q_{\text{PLS}} = \frac{\sum_{i=1}^N (z_i - \bar{z})(y_i - \bar{y})}{\sum_{i=1}^N (z_i - \bar{z})^2},$$

and the minimum-variance fixed-slope regression estimator is

$$\widehat{y}_{\text{PLS}} = \bar{y}_S + q_{\text{PLS}}(\bar{z} - \bar{z}_S).$$

Since the sample mean \bar{y}_S is a fixed-slope regression estimator, it follows that \widehat{y}_{PLS} has lower variance than the sample mean, unless $q_{\text{PLS}} = 0$ (in which case $\widehat{y}_{\text{PLS}} = \bar{y}_S$).

The actual OLS regression estimator is almost as precise as \widehat{y}_{PLS} in sufficiently large samples. The difference between the two estimators is

$$\widehat{y}_{\text{OLS}} - \widehat{y}_{\text{PLS}} = (\widehat{q}_{\text{OLS}} - q_{\text{PLS}})(\bar{z} - \bar{z}_S)$$

where \widehat{q}_{OLS} is the estimated slope from a regression of y_i on z_i in the sample. The estimation errors $\widehat{q}_{\text{OLS}} - q_{\text{PLS}}$, $\bar{z}_S - \bar{z}$, and $\widehat{y}_{\text{PLS}} - \bar{y}$ are of order $1/\sqrt{n}$ in probability. Thus, the difference $\widehat{y}_{\text{OLS}} - \widehat{y}_{\text{PLS}}$ is of order $1/n$, which is negligible compared to the estimation error in \widehat{y}_{PLS} when n is large enough.

In sum, in large enough samples,

$$\text{var}(\widehat{y}_{\text{OLS}}) \approx \text{var}(\widehat{y}_{\text{PLS}}) \leq \text{var}(\bar{y}_S)$$

and the inequality is strict unless y_i and z_i are uncorrelated in the population.

4.2. *Precision improvement in experiments.* The sampling result naturally leads to the conjecture that in a completely randomized experiment, OLS adjustment with a full set of treatment \times covariate interactions improves or does not hurt asymptotic precision, even when the regression model is incorrect. The adjusted estimator $\widehat{ATE}_{\text{interact}}$ is just the difference between two OLS regression estimators from sampling theory, while $\widehat{ATE}_{\text{unadj}}$ is the difference between two sample means.

The conjecture is confirmed below. To summarize the results:

1. $\widehat{ATE}_{\text{interact}}$ is consistent and asymptotically normal (as are $\widehat{ATE}_{\text{unadj}}$ and $\widehat{ATE}_{\text{adj}}$, from Freedman's results).
2. Asymptotically, $\widehat{ATE}_{\text{interact}}$ is at least as efficient as $\widehat{ATE}_{\text{unadj}}$, and more efficient unless the covariates are uncorrelated with the weighted average

$$\frac{n - n_A}{n} a_i + \frac{n_A}{n} b_i.$$

3. Asymptotically, $\widehat{ATE}_{\text{interact}}$ is at least as efficient as $\widehat{ATE}_{\text{adj}}$, and more efficient unless (a) the two treatment groups have equal size or (b) the covariates are uncorrelated with the treatment effect $a_i - b_i$.

4.2.1. *Assumptions for asymptotics.* Finite-population asymptotic results are statements about randomized experiments on (or random samples from) an imaginary infinite sequence of finite populations, with increasing n . The regularity conditions (assumptions on the limiting behavior of the sequence) may seem vacuous, since one can always construct a sequence that contains the actual population and still satisfies the conditions. But it may be useful to ask whether a sequence that preserves any relevant "irregularities" (such as the influence of gross outliers) would violate the regularity conditions. See also Lumley (2010, pp. 217–218).

The asymptotic results in this paper assume Freedman's (2008b) regularity conditions, generalized to allow multiple covariates; the number of covariates K is constant as n grows. One practical interpretation of these conditions is that in order for the results to be applicable, the size of each treatment group should be sufficiently large (and much larger than the number of covariates), the influence of outliers should be small, and near-collinearity in the covariates should be avoided.

As Freedman (2008a) notes, in principle, there should be an extra subscript to index the sequence of populations: for example, in the population with n subjects, the i th subject has potential outcomes $a_{i,n}$ and $b_{i,n}$, and the average treatment effect is ATE_n . Like Freedman, I drop the extra subscripts.

CONDITION 1. There is a bound $L < \infty$ such that for all $n = 1, 2, \dots$ and $k = 1, \dots, K$,

$$\frac{1}{n} \sum_{i=1}^n a_i^4 < L, \quad \frac{1}{n} \sum_{i=1}^n b_i^4 < L, \quad \frac{1}{n} \sum_{i=1}^n z_{ik}^4 < L.$$

CONDITION 2. Let \mathbf{Z} be the $n \times (K + 1)$ matrix whose i th row is $(1, \mathbf{z}_i)$. Then $n^{-1}\mathbf{Z}'\mathbf{Z}$ converges to a finite, invertible matrix. Also, the population means of a_i , b_i , a_i^2 , b_i^2 , $a_i b_i$, $a_i \mathbf{z}_i$, and $b_i \mathbf{z}_i$ converge to finite limits. For example, $\lim_{n \rightarrow \infty} n^{-1} \sum_{i=1}^n a_i \mathbf{z}_i$ exists and is a finite vector.

CONDITION 3. The proportion n_A/n converges to a limit p_A , with $0 < p_A < 1$.

4.2.2. *Asymptotic results.* Let \mathbf{Q}_a denote the limit of the vector of slope coefficients in the population least squares regression of a_i on \mathbf{z}_i . That is,

$$\mathbf{Q}_a = \lim_{n \rightarrow \infty} \left[\left(\sum_{i=1}^n (\mathbf{z}_i - \bar{\mathbf{z}})' (\mathbf{z}_i - \bar{\mathbf{z}}) \right)^{-1} \sum_{i=1}^n (\mathbf{z}_i - \bar{\mathbf{z}})' (a_i - \bar{a}) \right].$$

Define \mathbf{Q}_b analogously.

Now define the prediction errors

$$a_i^* = (a_i - \bar{a}) - (\mathbf{z}_i - \bar{\mathbf{z}})\mathbf{Q}_a, \quad b_i^* = (b_i - \bar{b}) - (\mathbf{z}_i - \bar{\mathbf{z}})\mathbf{Q}_b$$

for $i = 1, \dots, n$.

For any variables x_i and y_i , let σ_x^2 and $\sigma_{x,y}$ denote the population variance of x_i and the population covariance of x_i and y_i . For example,

$$\sigma_{a^*,b^*} = \frac{1}{n} \sum_{i=1}^n (a_i^* - \bar{a}^*) (b_i^* - \bar{b}^*) = \frac{1}{n} \sum_{i=1}^n a_i^* b_i^*.$$

Theorem 1 and its corollaries are proved in the supplement [Lin (2012)].

THEOREM 1. Assume Conditions 1–3. Then $\sqrt{n}(\widehat{ATE}_{\text{interact}} - ATE)$ converges in distribution to a Gaussian random variable with mean 0 and variance

$$\frac{1-p_A}{p_A} \lim_{n \rightarrow \infty} \sigma_{a^*}^2 + \frac{p_A}{1-p_A} \lim_{n \rightarrow \infty} \sigma_{b^*}^2 + 2 \lim_{n \rightarrow \infty} \sigma_{a^*,b^*}.$$

COROLLARY 1.1. Assume Conditions 1–3. Then $\widehat{ATE}_{\text{unadj}}$ has at least as much asymptotic variance as $\widehat{ATE}_{\text{interact}}$. The difference is

$$\frac{1}{np_A(1-p_A)} \lim_{n \rightarrow \infty} \sigma_E^2$$

where $E_i = (\mathbf{z}_i - \bar{\mathbf{z}})\mathbf{Q}_E$ and $\mathbf{Q}_E = (1-p_A)\mathbf{Q}_a + p_A\mathbf{Q}_b$. Therefore, adjustment with $\widehat{ATE}_{\text{interact}}$ helps asymptotic precision if $\mathbf{Q}_E \neq \mathbf{0}$ and is neutral if $\mathbf{Q}_E = \mathbf{0}$.

REMARKS. (i) \mathbf{Q}_E can be thought of as a weighted average of \mathbf{Q}_a and \mathbf{Q}_b , or as the limit of the vector of slope coefficients in the population least squares regression of $(1 - p_A)a_i + p_A b_i$ on \mathbf{z}_i .

(ii) The weights may seem counterintuitive at first, but the sampling analogy and Eqs. (3.2–3.3) can help. Other things being equal, adjustment has a larger effect on the estimated mean from the smaller treatment group, because its mean covariate values are further away from the population mean. The adjustment added to \bar{a}_A is

$$(\bar{\mathbf{z}} - \bar{\mathbf{z}}_A)\widehat{\mathbf{Q}}_a = \frac{n - n_A}{n}(\bar{\mathbf{z}}_B - \bar{\mathbf{z}}_A)\widehat{\mathbf{Q}}_a$$

while the adjustment added to \bar{b}_B is

$$(\bar{\mathbf{z}} - \bar{\mathbf{z}}_B)\widehat{\mathbf{Q}}_b = -\frac{n_A}{n}(\bar{\mathbf{z}}_B - \bar{\mathbf{z}}_A)\widehat{\mathbf{Q}}_b,$$

where $\widehat{\mathbf{Q}}_a$ and $\widehat{\mathbf{Q}}_b$ are OLS estimates that converge to \mathbf{Q}_a and \mathbf{Q}_b .

(iii) If the covariates' associations with a_i and b_i go in opposite directions, it is possible for adjustment with $\widehat{ATE}_{\text{interact}}$ to have no effect on asymptotic precision. Specifically, if $(1 - p_A)\mathbf{Q}_a = -p_A\mathbf{Q}_b$, the adjustments to \bar{a}_A and \bar{b}_B tend to cancel each other out.

(iv) In designs with more than two treatment groups, estimators analogous to $\widehat{ATE}_{\text{interact}}$ can be derived from a separate regression in each treatment group, or equivalently a single regression with the appropriate treatment dummies, covariates, and interactions. The resulting estimator of (for example) $\bar{a} - \bar{b}$ is at least as efficient as $\bar{Y}_A - \bar{Y}_B$, and more efficient unless the covariates are uncorrelated with both a_i and b_i . The supplement [Lin (2012)] gives a proof.

COROLLARY 1.2. *Assume Conditions 1–3. Then $\widehat{ATE}_{\text{adj}}$ has at least as much asymptotic variance as $\widehat{ATE}_{\text{interact}}$. The difference is*

$$\frac{(2p_A - 1)^2}{np_A(1 - p_A)} \lim_{n \rightarrow \infty} \sigma_D^2$$

where $D_i = (\mathbf{z}_i - \bar{\mathbf{z}})(\mathbf{Q}_a - \mathbf{Q}_b)$. Therefore, the two estimators have equal asymptotic precision if $p_A = 1/2$ or $\mathbf{Q}_a = \mathbf{Q}_b$. Otherwise, $\widehat{ATE}_{\text{interact}}$ is asymptotically more efficient.

REMARKS. (i) $\mathbf{Q}_a - \mathbf{Q}_b$ is the limit of the vector of slope coefficients in the population least squares regression of the treatment effect $a_i - b_i$ on \mathbf{z}_i .

(ii) For intuition about the behavior of $\widehat{ATE}_{\text{adj}}$, suppose there is a single covariate, z_i , and the population least squares slopes are $Q_a = 10$ and $Q_b = 2$. Let \widehat{Q} denote the estimated coefficient on z_i from a pooled OLS regression of Y_i on T_i and z_i . In sufficiently large samples, \widehat{Q} tends to fall close to $p_A Q_a + (1 - p_A)Q_b$. Consider two cases:

- If the two treatment groups have equal size, then $\bar{z} - \bar{z}_B = -(\bar{z} - \bar{z}_A)$, so when $\bar{z} - \bar{z}_A = 1$, the ideal linear adjustment would add 10 to \bar{a}_A and subtract 2 from \bar{b}_B . Instead, $\widehat{ATE}_{\text{adj}}$ uses the pooled slope estimate $\widehat{Q} \approx 6$, so it tends to underadjust \bar{a}_A (adding about 6) and overadjust \bar{b}_B (subtracting about 6). Two wrongs make a right: the adjustment adds about 12 to $\bar{a}_A - \bar{b}_B$, just as $\widehat{ATE}_{\text{interact}}$ would have done.
- If group A is 9 times larger than group B , then $\bar{z} - \bar{z}_B = -9(\bar{z} - \bar{z}_A)$, so when $\bar{z} - \bar{z}_A = 1$, the ideal linear adjustment adds 10 to \bar{a}_A and subtracts $9 \cdot 2 = 18$ from \bar{b}_B , thus adding 28 to the estimate of ATE . In contrast, the pooled adjustment adds $\widehat{Q} \approx 9.2$ to \bar{a}_A and subtracts $9\widehat{Q} \approx 82.8$ from \bar{b}_B , thus adding about 92 to the estimate of ATE . The problem is that the pooled regression has more observations of a_i than of b_i , but the adjustment has a larger effect on the estimate of \bar{b} than on that of \bar{a} , since group B 's mean covariate value is further away from the population mean.

(iii) The example above suggests an alternative regression adjustment: when group A has nine-tenths of the subjects, give group B nine-tenths of the weight. More generally, let $\tilde{p}_A = n_A/n$. Run a weighted least squares regression of Y_i on T_i and \mathbf{z}_i , with weights of $(1 - \tilde{p}_A)/\tilde{p}_A$ on each observation from group A and $\tilde{p}_A/(1 - \tilde{p}_A)$ on each observation from group B . This “tyranny of the minority” estimator is asymptotically equivalent to $\widehat{ATE}_{\text{interact}}$ (the supplement [Lin (2012)] outlines a proof). It is equal to $\widehat{ATE}_{\text{adj}}$ when $\tilde{p}_A = 1/2$.

(iv) The tyranny estimator can also be seen as a one-step variant of Rubin and van der Laan’s (2011) two-step “targeted ANCOVA.” Their estimator is equivalent to the difference in means of the residuals from a weighted least squares regression of Y_i on \mathbf{z}_i , with the same weights as in remark (iii).

(v) When is the usual adjustment worse than no adjustment? Eq. (23) in Freedman (2008a) implies that with a single covariate z_i , for $\widehat{ATE}_{\text{adj}}$ to have higher asymptotic variance than $\widehat{ATE}_{\text{unadj}}$, a necessary (but not sufficient) condition is that either the design must be so imbalanced that more than three-quarters of the subjects are assigned to one group, or z_i must have a larger covariance with the treatment effect $a_i - b_i$ than with the expected outcome $p_A a_i + (1 - p_A) b_i$. With multiple covariates, a similar condition can be derived from Eq. (14) in Schochet (2010).

(vi) With more than two treatment groups, the usual adjustment can be worse than no adjustment even when the design is balanced [Freedman (2008b)]. All the groups are pooled in a single regression without treatment \times covariate interactions, so group B 's data can affect the contrast between A and C .

4.2.3. *Example.* This simulation illustrates some of the key ideas.

1. For $n = 1,000$ subjects, a covariate z_i was drawn from the uniform distribu-

TABLE 1
Simulation (1,000 subjects; 40,000 replications)

Estimator	Proportion assigned to treatment A				
	0.75	0.6	0.5	0.4	0.25
SD (asymptotic) $\times 1,000$					
Unadjusted	93	49	52	78	143
Usual OLS-adjusted	171	72	46	79	180
OLS with interaction	80	49	46	58	98
Tyranny of the minority	80	49	46	58	98
SD (empirical) $\times 1,000$					
Unadjusted	93	49	53	78	142
Usual OLS-adjusted	171	73	47	80	180
OLS with interaction	81	50	47	59	99
Tyranny of the minority	81	50	47	59	99
Bias (estimated) $\times 1,000$					
Unadjusted	0	0	0	0	-2
Usual OLS-adjusted	-3	-3	-3	-3	-5
OLS with interaction	-5	-3	-3	-4	-6
Tyranny of the minority	-5	-3	-3	-4	-6

tion on $[-4, 4]$. The potential outcomes were then generated as

$$a_i = \frac{\exp(z_i) + \exp(z_i/2)}{4} + v_i,$$

$$b_i = \frac{-\exp(z_i) + \exp(z_i/2)}{4} + \varepsilon_i$$

with v_i and ε_i drawn independently from the standard normal distribution.

2. A completely randomized experiment was simulated 40,000 times, assigning $n_A = 750$ subjects to treatment A and the remainder to treatment B.
3. Step 2 was repeated for four other values of n_A (600, 500, 400, and 250).

These are adverse conditions for regression adjustment: z_i covaries much more with the treatment effect $a_i - b_i$ than with the potential outcomes, and the population least squares slopes $Q_a = 1.06$ and $Q_b = -0.73$ are of opposite signs.

Table 1 compares $\widehat{ATE}_{\text{unadj}}$, $\widehat{ATE}_{\text{adj}}$, $\widehat{ATE}_{\text{interact}}$, and the “tyranny of the minority” estimator from remark (iii) after Corollary 1.2. The first panel shows the asymptotic standard errors derived from Freedman’s (2008b) Theorems 1 and 2 and this paper’s Theorem 1 (with limits replaced by actual population values). The second and third panels show the empirical standard deviations and bias estimates from the Monte Carlo simulation.

The empirical standard deviations are very close to the asymptotic predictions, and the estimated biases are small in comparison. The usual adjustment hurts precision except when $n_A/n = 0.5$. In contrast, $\widehat{ATE}_{\text{interact}}$ and the tyranny estimator

improve precision except when $n_A/n = 0.6$. [This is approximately the value of p_A where $\widehat{ATE}_{\text{interact}}$ and $\widehat{ATE}_{\text{unadj}}$ have equal asymptotic variance; see remark (iii) after Corollary 1.1.]

Randomization does not “justify” the regression model of $\widehat{ATE}_{\text{interact}}$, and the linearity assumption is far from accurate in this example, but the estimator solves Freedman’s asymptotic precision problem.

5. Variance estimation. Eicker (1967) and White (1980ab) proposed a covariance matrix estimator for OLS that is consistent under simple random sampling from an infinite population. The regression model assumptions, such as linearity and homoskedasticity, are not needed for this result.⁷ The estimator is

$$(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\text{diag}(\hat{\varepsilon}_1^2, \dots, \hat{\varepsilon}_n^2)\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}$$

where \mathbf{X} is the matrix of regressors and $\hat{\varepsilon}_i$ is the i th OLS residual. It is known as the sandwich estimator because of its form, or as the Huber–White estimator because it is the sample analog of Huber’s (1967) formula for the asymptotic variance of a maximum likelihood estimator when the model is incorrect.

Theorem 2 shows that under the Neyman model, the sandwich variance estimators for $\widehat{ATE}_{\text{adj}}$ and $\widehat{ATE}_{\text{interact}}$ are consistent or asymptotically conservative. Together, Theorems 1 and 2 in this paper and Theorem 2 in Freedman (2008b) imply that asymptotically valid confidence intervals for ATE can be constructed from either $\widehat{ATE}_{\text{adj}}$ or $\widehat{ATE}_{\text{interact}}$ and the sandwich standard error estimator.

The vectors \mathbf{Q}_a and \mathbf{Q}_b were defined in Section 4.2.2. Let \mathbf{Q} denote the weighted average $p_A\mathbf{Q}_a + (1 - p_A)\mathbf{Q}_b$. As shown in Freedman (2008b) and the supplement [Lin (2012)], \mathbf{Q} is the probability limit of the vector of coefficients on \mathbf{z}_i in the OLS regression of Y_i on T_i and \mathbf{z}_i .

Mimicking Section 4.2.2, define the prediction errors

$$a_i^{**} = (a_i - \bar{a}) - (\mathbf{z}_i - \bar{\mathbf{z}})\mathbf{Q}, \quad b_i^{**} = (b_i - \bar{b}) - (\mathbf{z}_i - \bar{\mathbf{z}})\mathbf{Q}$$

for $i = 1, \dots, n$.

Theorem 2 is proved in the supplement [Lin (2012)].

THEOREM 2. *Assume Conditions 1–3. Let \widehat{v}_{adj} and $\widehat{v}_{\text{interact}}$ denote the sandwich variance estimators for $\widehat{ATE}_{\text{adj}}$ and $\widehat{ATE}_{\text{interact}}$. Then $n\widehat{v}_{\text{adj}}$ converges in probability to*

$$\frac{1}{p_A} \lim_{n \rightarrow \infty} \sigma_{a^{**}}^2 + \frac{1}{1 - p_A} \lim_{n \rightarrow \infty} \sigma_{b^{**}}^2,$$

⁷See, e.g., Chamberlain (1982, pp. 17–19) or Angrist and Pischke (2009, pp. 40–48). Fuller (1975) proves a finite-population version of the result.

which is greater than or equal to the true asymptotic variance of $\sqrt{n}(\widehat{ATE}_{\text{adj}} - ATE)$. The difference is

$$\lim_{n \rightarrow \infty} \sigma_{(a-b)}^2 = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n [(a_i - b_i) - ATE]^2.$$

Similarly, $n\widehat{v}_{\text{interact}}$ converges in probability to

$$\frac{1}{p_A} \lim_{n \rightarrow \infty} \sigma_{a^*}^2 + \frac{1}{1 - p_A} \lim_{n \rightarrow \infty} \sigma_{b^*}^2,$$

which is greater than or equal to the true asymptotic variance of $\sqrt{n}(\widehat{ATE}_{\text{interact}} - ATE)$. The difference is

$$\lim_{n \rightarrow \infty} \sigma_{(a^*-b^*)}^2 = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n [(a_i - b_i) - ATE - (\mathbf{z}_i - \bar{\mathbf{z}})(\mathbf{Q}_a - \mathbf{Q}_b)]^2.$$

REMARKS. (i) Theorem 2 generalizes to designs with more than two treatment groups.

(ii) With two treatment groups of equal size, the conventional OLS variance estimator for $\widehat{ATE}_{\text{adj}}$ is also consistent or asymptotically conservative [Freedman (2008a)].

(iii) Freedman (2008a) shows analogous results for variance estimators for the difference in means; the issue there is whether to assume $\sigma_a^2 = \sigma_b^2$. Reichardt and Gollob (1999) and Freedman, Pisani, and Purves (2007, pp. 508–511) give helpful expositions of basic results under the Neyman model. Related issues appear in discussions of the two-sample problem [Miller (1986, pp. 56–62); Stonehouse and Forrester (1998)] and randomization tests [Gail et al. (1996); Chung and Romano (2011ab)].

(iv) With a small sample or points of high leverage, the sandwich estimator can have substantial downward bias and high variability. MacKinnon (2011) discusses bias-corrected sandwich estimators and improved confidence intervals based on the wild bootstrap. See also Wu (1986), Tibshirani (1986), Angrist and Pischke (2009, ch. 8), and Kline and Santos (2012).

(v) When $\widehat{ATE}_{\text{unadj}}$ is computed by regressing Y_i on T_i , the HC2 bias-corrected sandwich estimator [MacKinnon and White (1985); Royall and Cumberland (1978); Wu (1986, p. 1274)] gives exactly the variance estimate preferred by Neyman (1923) and Freedman (2008a): $\widehat{\sigma}_a^2/n_A + \widehat{\sigma}_b^2/(n - n_A)$, where $\widehat{\sigma}_a^2$ and $\widehat{\sigma}_b^2$ are the sample variances of Y_i in the two groups.⁸

⁸For details, see Hinkley and Wang (1991), Angrist and Pischke (2009, pp. 294–304), or Samii and Aronow (2012).

(vi) When the n subjects are randomly drawn from a superpopulation, $\widehat{v}_{\text{interact}}$ does not take into account the variability in \bar{z} [Imbens and Wooldridge (2009, pp. 28–30)]. In the Neyman model, \bar{z} is fixed.

(vii) Freedman’s (2006) critique of the sandwich estimator does not apply here, as $\widehat{ATE}_{\text{adj}}$ and $\widehat{ATE}_{\text{interact}}$ are consistent even when their regression models are incorrect.

(viii) Freedman (2008a) associates the difference in means and regression with heteroskedasticity-robust and conventional variance estimators, respectively. His rationale for these pairings is unclear. The pooled-variance two-sample t -test and the conventional F -test for equality of means are often used in difference-in-means analyses. Conversely, the sandwich estimator has become the usual variance estimator for regression in economics [Stock (2010)]. The question of whether to adjust for covariates should be disentangled from the question of whether to assume homoskedasticity.

6. Bias. The bias of OLS adjustment diminishes rapidly with the number of randomly assigned units: $\widehat{ATE}_{\text{adj}}$ and $\widehat{ATE}_{\text{interact}}$ have biases of order $1/n$, while their standard errors are of order $1/\sqrt{n}$. Brief remarks follow; see also Deaton (2010, pp. 443–444), Imbens (2010, pp. 410–411), and Green and Aronow (2011).

(i) If the actual random assignment yields substantial covariate imbalance, it is hardly reassuring to be told that the difference in means is unbiased over all possible random assignments. Senn (1989) and Cox and Reid (2000, pp. 29–32) argue that inference should be conditional on a measure of covariate imbalance, and that the conditional bias of $\widehat{ATE}_{\text{unadj}}$ justifies adjustment. Tukey (1991) suggests adjustment “perhaps as a supplemental analysis” for “protection against either the consequences of inadequate randomization or the (random) occurrence of an unusual randomization.”

(ii) As noted in Section 2.2, poststratification is a special case of $\widehat{ATE}_{\text{interact}}$. The poststratified estimator is a population-weighted average of subgroup-specific differences in means. Conditional on the numbers of subgroup members assigned to each treatment, the poststratified estimator is unbiased, but $\widehat{ATE}_{\text{unadj}}$ can be biased. Miratrix, Sekhon, and Yu (2012) give finite-sample and asymptotic analyses of poststratification and blocking; see also Holt and Smith (1979) in the sampling context.

(iii) Cochran (1977) analyzes the bias of \widehat{y}_{reg} in Eq. (3.1). If the adjustment factor q is fixed, \widehat{y}_{reg} is unbiased, but if q varies with the sample, \widehat{y}_{reg} has a bias of $-\text{cov}(q, \bar{z}_S)$. The leading term in the bias of \widehat{y}_{OLS} is

$$-\frac{1}{\sigma_z^2} \left(\frac{1}{n} - \frac{1}{N} \right) \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{i=1}^N e_i (z_i - \bar{z})^2$$

where n is the sample size, N is the population size, and e_i is the prediction error in the population least squares regression of y_i on \underline{z}_i .

(iv) By analogy, the leading term in the bias of $\widehat{ATE}_{\text{interact}}$ (with a single covariate z_i) is

$$-\frac{1}{\sigma_z^2} \left[\left(\frac{1}{n_A} - \frac{1}{n} \right) \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n a_i^* (z_i - \bar{z})^2 - \left(\frac{1}{n - n_A} - \frac{1}{n} \right) \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n b_i^* (z_i - \bar{z})^2 \right].$$

Thus, the bias tends to depend largely on n , n_A/n , and the importance of omitted quadratic terms in the regressions of a_i and b_i on z_i . With multiple covariates, it would also depend on the importance of omitted first-order interactions between the covariates.

(v) Remark (iii) also implies that if the adjustment factors \mathbf{q}_a and \mathbf{q}_b in Eqs. (3.2–3.3) do not vary with random assignment, the resulting estimator of ATE is unbiased. Middleton and Aronow's (2012) insightful paper uses out-of-sample data to determine $\mathbf{q}_a = \mathbf{q}_b$. In-sample data can be used when multiple pretests (pre-randomization outcome measures) are available: if the only covariate z_i is the most recent pretest, a common adjustment factor $q_a = q_b$ can be determined by regressing z_i on an earlier pretest.

7. Empirical example. This section suggests empirical checks on the asymptotic approximations. I will focus on the validity of confidence intervals, using data from a social experiment for an illustrative example.

7.1. Background. Angrist, Lang, and Oreopoulos (2009; henceforth ALO) conducted an experiment to estimate the effects of support services and financial incentives on college students' academic achievement. At a Canadian university campus, all first-year undergraduates entering in September 2005, except those with a high-school grade point average (GPA) in the top quartile, were randomly assigned to four groups. One treatment group was offered support services (peer advising and supplemental instruction). Another group was offered financial incentives (awards of \$1,000 to \$5,000 for meeting a target GPA). A third group was offered both services and incentives. The control group was eligible only for standard university support services (which included supplemental instruction for some courses).

ALO report that for women, the combination of services and incentives had sizable estimated effects on both first- and second-year academic achievement, even though the programs were only offered during the first year. In contrast, there was no evidence that services alone or incentives alone had lasting effects for women or that any of the treatments improved achievement for men (who were much less likely to contact peer advisors).

To simplify the example and focus on the accuracy of large-sample approximations in samples that are not huge, I use only the data for men (43 percent of the students) in the services-and-incentives and services-only groups (9 percent and 15 percent of the men). First-year GPA data are available for 58 men in the services-and-incentives group and 99 in the services-only group.

Table 2 shows alternative estimates of *ATE* (the average treatment effect of the financial incentives, given that the support services were available). The services-and-incentives and services-only groups had average first-year GPAs of 1.82 and 1.86 (on a scale of 0 to 4), so the unadjusted estimate of *ATE* is close to zero. OLS adjustment for high-school GPA hardly makes a practical difference to either the point estimate of *ATE* or the sandwich standard error estimate, regardless of whether the treatment \times covariate interaction is included.⁹ The two groups had similar average high-school GPAs, and high-school GPA was not a strong predictor of first-year college GPA.

TABLE 2
Estimates of average treatment effect on men's first-year GPA

	Point estimate	Sandwich SE
Unadjusted	-0.036	0.158
Usual OLS-adjusted	-0.083	0.146
OLS with interaction	-0.081	0.146

The finding that adjustment appears to have little effect on precision is not unusual in social experiments, because the covariates are often only weakly correlated with the outcome [Meyer (1995, pp. 100, 116); Lin et al. (1998, pp. 129–133)]. Examining eight social experiments with a wide range of outcome variables, Schochet (2010) finds R^2 values above 0.3 only when the outcome is a standardized achievement test score or Medicaid costs and the covariates include a lagged outcome.

Researchers may prefer not to adjust when the expected precision improvement is meager. Either way, confidence intervals for treatment effects typically rely on either strong parametric assumptions (such as a constant treatment effect or a normally distributed outcome) or asymptotic approximations. When a sandwich standard error estimate is multiplied by 1.96 to form a margin of error for a 95 percent confidence interval, the calculation assumes the sample is large enough that (i) the estimator of *ATE* is approximately normally distributed, (ii) the bias and variability of the sandwich standard error estimator are small relative to the true standard error (or else the bias is conservative and the variability is small), and (iii) the bias of adjustment (if used) is small relative to the true standard error.

Below I discuss a simulation to check for confidence interval undercoverage due

⁹ALO adjust for a larger set of covariates, including first language, parents' education, and self-reported procrastination tendencies. These also have little effect on the estimated standard errors.

to violations of (i) or (ii), and a bias estimate to check for violations of (iii). These checks are not foolproof, but may provide a useful sniff test.

7.2. Simulation. For technical reasons, the most revealing initial check is a simulation with a constant treatment effect. When treatment effects are heterogeneous, the sandwich standard error estimators for $\widehat{ATE}_{\text{unadj}}$ and $\widehat{ATE}_{\text{adj}}$ are asymptotically conservative,¹⁰ so nominal 95 percent confidence intervals for ATE achieve greater than 95 percent coverage in large enough samples. A simulation that overstates treatment effect heterogeneity may overestimate coverage.

Table 3 reports a simulation that assumes treatment had no effect on any of the men. Keeping the GPA data at their actual values, I replicated the experiment 250,000 times, each time randomly assigning 58 men to services-and-incentives and 99 to services-only. The first panel shows the means and standard deviations of $\widehat{ATE}_{\text{unadj}}$, $\widehat{ATE}_{\text{adj}}$, and $\widehat{ATE}_{\text{interact}}$. All three estimators are approximately unbiased, but adjustment slightly improves precision. Since the simulation assumes a constant treatment effect (zero), including the treatment \times covariate interaction does not improve precision relative to the usual adjustment.

The second and third panels show the estimated biases and standard deviations of the sandwich standard error estimator and the three variants discussed in Angrist and Pischke (2009, pp. 294–308). ALO’s paper uses HC1 [Hinkley (1977)], which simply multiplies the sandwich variance estimator by $n/(n-k)$, where k is the number of regressors. HC2 [see remark (v) after Theorem 2] and the approximate jackknife HC3 [Davidson and MacKinnon (1993); Tibshirani (1986)] inflate the squared residuals in the sandwich formula by the factors $(1-h_{ii})^{-1}$ and $(1-h_{ii})^{-2}$, where h_{ii} is the i th diagonal element of the hat matrix $\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'$. All the standard error estimators appear to be approximately unbiased with low variability.

The fourth and fifth panels evaluate thirteen ways of constructing a 95 percent confidence interval. For each of the three estimators of ATE , each of the four standard error estimators was multiplied by 1.96 to form the margin of error for a normal-approximation interval. Welch’s (1949) t -interval [Miller (1986, pp. 60–62)] was also constructed. Welch’s interval uses $\widehat{ATE}_{\text{unadj}}$, the HC2 standard error estimator, and the t -distribution with the Welch–Satterthwaite approximate degrees of freedom.

The fourth panel shows that all thirteen confidence intervals cover the true value of ATE (zero) with approximately 95 percent probability. The fifth panel shows the average widths of the intervals. (The mean and median widths agree up to three decimal places.) The regression-adjusted intervals are narrower on average than the unadjusted intervals, but the improvement is meager. In sum, adjustment appears

¹⁰By Theorem 2, the sandwich standard error estimator for $\widehat{ATE}_{\text{interact}}$ is also asymptotically conservative unless the treatment effect is a linear function of the covariates.

TABLE 3

Simulation with zero treatment effect (250,000 replications). The fourth panel shows the empirical coverage rates of nominal 95 percent confidence intervals. All other estimates are on the four-point GPA scale.

	ATE estimator		
	Unadjusted	Usual OLS-adjusted	OLS with interaction
Bias & SD of ATE estimator			
Mean (estimated bias)	0.000	0.000	0.000
SD	0.158	0.147	0.147
Bias of SE estimator			
Classic sandwich	-0.001	-0.002	-0.002
HC1	0.000	0.000	0.000
HC2	0.000	0.000	0.000
HC3	0.001	0.002	0.002
SD of SE estimator			
Classic sandwich	0.004	0.004	0.004
HC1	0.004	0.004	0.004
HC2	0.004	0.004	0.004
HC3	0.004	0.004	0.005
CI coverage (percent)			
Classic sandwich	94.6	94.5	94.4
HC1	94.8	94.7	94.7
HC2 (normal)	94.8	94.8	94.8
HC2 (Welch t)	95.1		
HC3	95.0	95.0	95.1
CI width (average)			
Classic sandwich	0.618	0.570	0.568
HC1	0.622	0.576	0.575
HC2 (normal)	0.622	0.576	0.577
HC2 (Welch t)	0.629		
HC3	0.627	0.583	0.586

to yield slightly more precise inference without sacrificing validity.

7.3. *Bias estimates.* One limitation of the simulation above is that the bias of adjustment may be larger when treatment effects are heterogeneous. With a single covariate z_i , the leading term in the bias of $\widehat{ATE}_{\text{adj}}$ is¹¹

$$-\frac{1}{n} \frac{1}{\sigma_z^2} \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n [(a_i - b_i) - ATE] (z_i - \bar{z})^2.$$

¹¹An equivalent expression appears in the version of Freedman (2008a) on his web page. It can be derived from Freedman (2008b) after correcting a minor error in Eqs. (17–18): the potential outcomes should be centered.

Thus, with a constant treatment effect, the leading term is zero (and the bias is of order $n^{-3/2}$ or smaller). Freedman (2008b) shows that with a balanced design and a constant treatment effect, the bias is exactly zero.

We can estimate the leading term by rewriting it as

$$-\frac{1}{n} \frac{1}{\sigma_z^2} \left[\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (a_i - \bar{a})(z_i - \bar{z})^2 - \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n (b_i - \bar{b})(z_i - \bar{z})^2 \right]$$

and substituting the sample variance of high-school GPA for σ_z^2 , and the sample covariances of first-year college GPA with the square of centered high-school GPA in the services-and-incentives and services-only groups for the bracketed limits. The resulting estimate of the bias of $\widehat{ATE}_{\text{adj}}$ is -0.0002 on the four-point GPA scale. Similarly, the leading term in the bias of $\widehat{ATE}_{\text{interact}}$ [Section 6, remark (iv)] can be estimated, and the result is also -0.0002 . The biases would need to be orders of magnitude larger to have noticeable effects on confidence interval coverage (the estimated standard errors of $\widehat{ATE}_{\text{adj}}$ and $\widehat{ATE}_{\text{interact}}$ in Table 2 are both 0.146).

7.4. Remarks. (i) This exercise does not prove that the bias of adjustment is negligible, since it just replaces a first-order approximation (the bias is close to zero in large enough samples) with a second-order approximation (the bias is close to the leading term in large enough samples), and the estimate of the leading term has sampling error. The checks suggested here cannot validate an analysis, but they can reveal problems.

(ii) Another limitation is that the simulation assumes the potential outcome distributions have the same shape. In Stonehouse and Forrester's (1998) simulations, Welch's t -test was not robust to extreme skewness in the smaller group when that group's sample size was 30 or smaller. That does not appear to be a serious issue in this example, however. The distribution of men's first-year GPA in the services-and-incentives group is roughly symmetric (e.g., see ALO, Fig. 1A).

(iii) The simulation check may appear to resemble permutation inference [Fisher (1935); Tukey (1993); Rosenbaum (2002)], but the goals differ. Here, the constant treatment effect scenario just gives a benchmark to check the finite-sample coverage of confidence intervals that are asymptotically valid under weaker assumptions. Classical permutation methods achieve exact inference under strong assumptions about treatment effects, but may give misleading results when the assumptions fail. For example, the Fisher–Pitman permutation test is asymptotically equivalent to a t -test using the conventional OLS standard error estimator. The test can be inverted to give exact confidence intervals for a constant treatment effect, but these intervals may undercover ATE when treatment effects are heterogeneous and the design is imbalanced [Gail et al. (1996)].

(iv) Chung and Romano (2011ab) discuss and extend a literature on permutation tests that do remain valid asymptotically when the null hypothesis is weakened. One such test is based on the permutation distribution of a heteroskedasticity-robust t -statistic. Exploration of this approach under the Neyman model (with and without covariate adjustment) would be valuable.

8. Further remarks. Freedman’s papers answer important questions about the properties of OLS adjustment. He and others have summarized his results with a “glass is half empty” view that highlights the dangers of adjustment. To the extent that this view encourages researchers to present unadjusted estimates first, it is probably a good influence. The difference in means is the “hands above the table” estimate: it is clearly not the product of a specification search, and its transparency may encourage discussion of the strengths and weaknesses of the data and research design.¹²

But it would be unwise to conclude that Freedman’s critique should always override the arguments for adjustment, or that studies reporting only adjusted estimates should always be distrusted. Freedman’s own work shows that with large enough samples and balanced two-group designs, randomization justifies the traditional adjustment. One does not need to believe in the classical linear model to tolerate or even advocate OLS adjustment, just as one does not need to believe in the Four Noble Truths of Buddhism to entertain the hypothesis that mindfulness meditation has causal effects on mental health.

From an agnostic perspective, Freedman’s theorems are a major contribution. Three-quarters of a century after Fisher discovered the analysis of covariance, Freedman deepened our understanding of its properties by deriving the regression-adjusted estimator’s asymptotic distribution without assuming a regression model, a constant treatment effect, or an infinite superpopulation. His argument is constructed with unsurpassed clarity and rigor. It deserves to be studied in detail and considered carefully.

Acknowledgments. I am grateful to Jas Sekhon, Dylan Small, Erin Hartman, Danny Hidalgo, and Terry Speed for valuable advice; to Peter Aronow, Gus Baker, Erik Beecroft, Dean Eckles, Jed Friedman, Raphael Kang, Pat Kline, Alex Mayer, Justin McCrary, David McKenzie, Tod Mijanovich, Luke Miratrix, Deb Nolan, Berk Ozler, Susan Paddock, Cyrus Samii, an anonymous associate editor, and reading groups at Berkeley and Penn for helpful discussions; and to David Freedman for his generous help earlier in my education. Any errors are my own.

¹²On transparency and critical discussion, see Ashenfelter and Plant (1990), Freedman (1991, 2008c, 2010), Moher et al. (2010), and Rosenbaum (2010, ch. 6).

SUPPLEMENTARY MATERIAL

Supplement to “Agnostic notes on regression adjustments to experimental data: Reexamining Freedman’s critique”

(doi: ???; .pdf). Proofs of theorems, corollaries, and selected remarks.

REFERENCES

- ANGRIST, J. D. and IMBENS, G. W. (2002). Comment on “Covariance adjustment in randomized experiments and observational studies” by P. R. Rosenbaum. *Statist. Sci.* **17** 304–307.
- ANGRIST, J. D., LANG, D. and OREOPOULOS, P. (2009). Incentives and services for college achievement: Evidence from a randomized trial. *Am. Econ. J.: Appl. Econ.* **1:1** 136–163.
- ANGRIST, J. D. and PISCHKE, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton Univ. Press, Princeton.
- ASHENFELTER, O. and PLANT, M. W. (1990). Nonparametric estimates of the labor-supply effects of negative income tax programs. *J. Labor Econ.* **8** S396–S415.
- BERK, R., BARNES, G., AHLMAN, L. and KURTZ, E. (2010). When second best is good enough: A comparison between a true experiment and a regression discontinuity quasi-experiment. *J. Exp. Criminol.* **6** 191–208.
- CHAMBERLAIN, G. (1982). Multivariate regression models for panel data. *J. Econometrics* **18** 5–46.
- CHUNG, E. Y. and ROMANO, J. P. (2011a). Exact and asymptotically robust permutation tests. Technical Report 2011-05, Dept. of Statistics, Stanford Univ.
- CHUNG, E. Y. and ROMANO, J. P. (2011b). Asymptotically valid and exact permutation tests based on two-sample U -statistics. Technical Report 2011-09, Dept. of Statistics, Stanford Univ.
- COCHRAN, W. G. (1942). Sampling theory when the sampling-units are of unequal sizes. *J. Amer. Statist. Assoc.* **37** 199–212.
- COCHRAN, W. G. (1957). Analysis of covariance: Its nature and uses. *Biometrics* **13** 261–281.
- COCHRAN, W. G. (1969). The use of covariance in observational studies. *J. R. Stat. Soc. Ser. C. Appl. Stat.* **18** 270–275.
- COCHRAN, W. G. (1977). *Sampling Techniques*, 3rd ed. Wiley, New York.
- COX, D. R. and MCCULLAGH, P. (1982). Some aspects of analysis of covariance (with discussion). *Biometrics* **38** 541–561.
- COX, D. R. and REID, N. (2000). *The Theory of the Design of Experiments*. CRC Press, Boca Raton.
- DAVIDSON, R. and MACKINNON, J. G. (1993). *Estimation and Inference in Econometrics*. Oxford Univ. Press, Oxford.
- DEATON, A. (2010). Instruments, randomization, and learning about development. *J. Econ. Lit.* **48** 424–455.
- EICKER, F. (1967). Limit theorems for regressions with unequal and dependent errors. *Proc. Fifth Berkeley Sympos. Math. Statist. and Probability* **1** 59–82.
- FIENBERG, S. E. and TANUR, J. M. (1987). Experimental and sampling structures: Parallels diverging and meeting. *Internat. Statist. Rev.* **55** 75–96.
- FISHER, R. A. (1932). *Statistical Methods for Research Workers*, 4th ed. Oliver and Boyd, Edinburgh.
- FISHER, R. A. (1935). *The Design of Experiments*. Oliver and Boyd, Edinburgh.
- FREEDMAN, D. A. (1991). Statistical models and shoe leather (with discussion). *Socio. Meth.* **21** 291–358.
- FREEDMAN, D. A. (2006). On the so-called “Huber sandwich estimator” and “robust standard errors.” *Amer. Statist.* **60** 299–302.
- FREEDMAN, D. A. (2008a). On regression adjustments to experimental data. *Adv. in Appl. Math.* **40** 180–193.

- FREEDMAN, D. A. (2008b). On regression adjustments in experiments with several treatments. *Ann. Appl. Stat.* **2** 176–196.
- FREEDMAN, D. A. (2008c). Editorial: Oasis or mirage? *Chance* **21:1** 59–61. Annotated references at www.stat.berkeley.edu/~census/chance.pdf.
- FREEDMAN, D. A. (2010). Survival analysis: An epidemiological hazard? In *Statistical Models and Causal Inference: A Dialogue with the Social Sciences* (D. Collier, J. S. Sekhon, and P. B. Stark, eds.) 169–192. Cambridge Univ. Press, Cambridge.
- FREEDMAN, D. A., PISANI, R. and PURVES, R. (2007). *Statistics*, 4th ed. Norton, New York.
- FULLER, W. A. (1975). Regression analysis for sample survey. *Sankhyā Ser. C* **37** 117–132.
- FULLER, W. A. (2002). Regression estimation for survey samples. *Surv. Meth.* **28** 5–23.
- FULLER, W. A. (2009). *Sampling Statistics*. Wiley, Hoboken.
- GAIL, M. H., MARK, S. D., CARROLL, R. J., GREEN, S. B. and PEE, D. (1996). On design considerations and randomization-based inference for community intervention trials. *Stat. Med.* **15** 1069–1092.
- GREEN, D. P. and ARONOW, P. M. (2011). Analyzing experimental data using regression: When is bias a practical concern? Working paper, Yale Univ.
- GREENBERG, D. and SHRODER, M. (2004). *The Digest of Social Experiments*, 3rd ed. Urban Institute Press, Washington.
- HANSEN, B. B. and BOWERS, J. (2009). Attributing effects to a cluster-randomized get-out-the-vote campaign. *J. Amer. Statist. Assoc.* **104** 873–885.
- HINKLEY, D. V. (1977). Jackknifing in unbalanced situations. *Technometrics* **19** 285–292.
- HINKLEY, D. V. and WANG, S. (1991). Efficiency of robust standard errors for regression coefficients. *Comm. Statist. Theory Methods* **20** 1–11.
- HOLLAND, P. W. (1986). Statistics and causal inference (with discussion). *J. Amer. Statist. Assoc.* **81** 945–970.
- HOLT, D. and SMITH, T. M. F. (1979). Post stratification. *J. Roy. Statist. Soc. Ser. A* **142** 33–46.
- HUBER, P. J. (1967). The behavior of maximum likelihood estimates under nonstandard conditions. *Proc. Fifth Berkeley Sympos. Math. Statist. and Probability* **1** 221–233.
- IMBENS, G. W. (2010). Better LATE than nothing: Some comments on Deaton (2009) and Heckman and Urzua (2009). *J. Econ. Lit.* **48** 399–423.
- IMBENS, G. W. and WOOLDRIDGE, J. M. (2009). Recent developments in the econometrics of program evaluation. *J. Econ. Lit.* **47** 5–86.
- KLAR, N. and DARLINGTON, G. (2004). Methods for modelling change in cluster randomization trials. *Stat. Med.* **23** 2341–2357.
- KLINE, P. M. (2011). Oaxaca–Blinder as a reweighting estimator. *Am. Econ. Rev.* **101:3** 532–537.
- KLINE, P. M. and SANTOS, A. (2012). Higher order properties of the wild bootstrap under misspecification. *J. Econometrics*. To appear.
- LIN, W. (2012). Supplement to “Agnostic notes on regression adjustments to experimental data: Reexamining Freedman’s critique.” Submitted to *Ann. Appl. Stat.*
- LIN, W., ROBINS, P. K., CARD, D., HARKNETT, K. and LUI-GURR, S. (1998). *When Financial Incentives Encourage Work: Complete 18-Month Findings from the Self-Sufficiency Project*. Social Research and Demonstration Corp., Ottawa.
- LUMLEY, T. (2010). *Complex Surveys: A Guide to Analysis Using R*. Wiley, Hoboken.
- MACKINNON, J. G. (2011). Thirty years of heteroskedasticity-robust inference. Working Paper 1268, Dept. of Economics, Queen’s Univ., Kingston, ON.
- MACKINNON, J. G. and WHITE, H. (1985). Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *J. Econometrics* **29** 305–325.
- MEYER, B. D. (1995). Lessons from the U.S. unemployment insurance experiments. *J. Econ. Lit.* **33** 91–131.
- MIDDLETON, J. A. and ARONOW, P. M. (2012). Unbiased estimation of the average treatment effect

- in cluster-randomized experiments. Working paper, Yale Univ.
- MILLER, R. G. (1986). *Beyond ANOVA: Basics of Applied Statistics*. Wiley, New York.
- MIRATRIX, L. W., SEKHON, J. S. and YU, B. (2012). Adjusting treatment effect estimates by post-stratification in randomized experiments. Working paper, Univ. of California, Berkeley.
- MOHER, D., HOPEWELL, S., SCHULZ, K. F. et al. (2010). CONSORT 2010 explanation and elaboration: Updated guidelines for reporting parallel group randomised trials. *BMJ* **340** c869.
- NEYMAN, J. (1923). On the application of probability theory to agricultural experiments. Translated and edited by D. M. Dabrowska and T. P. Speed (1990). *Statist. Sci.* **5** 463–480 (with discussion).
- RAUDENBUSH, S. W. (1997). Statistical analysis and optimal design for cluster randomized trials. *Psychol. Meth.* **2** 173–185.
- REICHARDT, C. S. and GOLLOB, H. F. (1999). Justifying the use and increasing the power of a t test for a randomized experiment with a convenience sample. *Psychol. Meth.* **4** 117–128.
- ROSENBAUM, P. R. (2002). Covariance adjustment in randomized experiments and observational studies (with discussion). *Statist. Sci.* **17** 286–327.
- ROSENBAUM, P. R. (2010). *Design of Observational Studies*. Springer, New York.
- ROYALL, R. M. and CUMBERLAND, W. G. (1978). Variance estimation in finite population sampling. *J. Amer. Statist. Assoc.* **73** 351–358.
- RUBIN, DANIEL B. and VAN DER LAAN, M. J. (2011). Targeted ANCOVA estimator in RCTs. In *Targeted Learning: Causal Inference for Observational and Experimental Data* (M. J. van der Laan and S. Rose, eds.) 201–215. Springer, New York.
- RUBIN, DONALD B. (1974). Estimating causal effects of treatments in randomized and nonrandomized studies. *J. Educ. Psychol.* **66** 688–701.
- RUBIN, DONALD B. (1984). William G. Cochran’s contributions to the design, analysis, and evaluation of observational studies. In *W. G. Cochran’s Impact on Statistics* (P. S. R. S. Rao and J. Sedransk, eds.) 37–69. Wiley, New York.
- RUBIN, DONALD B. (2005). Causal inference using potential outcomes: Design, modeling, decisions. *J. Amer. Statist. Assoc.* **100** 322–331.
- SAMII, C. and ARONOW, P. M. (2012). On equivalencies between design-based and regression-based variance estimators for randomized experiments. *Statist. Probab. Lett.* **82** 365–370.
- SCHOCHET, P. Z. (2010). Is regression adjustment supported by the Neyman model for causal inference? *J. Statist. Plann. Inference* **140** 246–259.
- SENN, S. J. (1989). Covariate imbalance and random allocation in clinical trials. *Stat. Med.* **8** 467–475.
- STOCK, J. H. (2010). The other transformation in econometric practice: Robust tools for inference. *J. Econ. Perspect.* **24:2** 83–94.
- STONEHOUSE, J. M. and FORRESTER, G. J. (1998). Robustness of the t and U tests under combined assumption violations. *J. Appl. Stat.* **25** 63–74.
- TIBSHIRANI, R. (1986). Discussion of “Jackknife, bootstrap and other resampling methods in regression analysis” by C. F. J. Wu. *Ann. Statist.* **14** 1335–1339. Correction: **16** 479.
- TSIATIS, A. A., DAVIDIAN, M., ZHANG, M. and LU, X. (2008). Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Stat. Med.* **27** 4658–4677.
- TUKEY, J. W. (1991). Use of many covariates in clinical trials. *Internat. Statist. Rev.* **59** 123–137.
- TUKEY, J. W. (1993). Tightening the clinical trial. *Contr. Clin. Trials* **14** 266–285.
- WATSON, D. J. (1937). The estimation of leaf area in field crops. *J. Agr. Sci.* **27** 474–483.
- WELCH, B. L. (1949). Appendix: Further note on Mrs. Aspin’s tables and on certain approximations to the tabled function. *Biometrika* **36** 293–296.
- WHITE, H. (1980a). Using least squares to approximate unknown regression functions. *Int. Econ. Rev.* **21** 149–170.
- WHITE, H. (1980b). A heteroskedasticity-consistent covariance matrix estimator and a direct test for

heteroskedasticity. *Econometrica* **48** 817–838.

WU, C. F. J. (1986). Jackknife, bootstrap and other resampling methods in regression analysis (with discussion). *Ann. Statist.* **14** 1261–1350.

YANG, L. and TSIATIS, A. A. (2001). Efficiency study of estimators for a treatment effect in a pretest–posttest trial. *Amer. Statist.* **55** 314–321.

DEPARTMENT OF STATISTICS
UNIVERSITY OF CALIFORNIA
BERKELEY, CALIFORNIA 94720-3860
USA
E-MAIL: winston@stat.berkeley.edu