



Controversy or Debate

Peng Ding*

Two seemingly paradoxical results in linear models: the variance inflation factor and the analysis of covariance

<https://doi.org/10.1515/jci-2019-0023>

Received Aug 18, 2019; accepted Feb 05, 2021

Abstract: A result from a standard linear model course is that the variance of the ordinary least squares (OLS) coefficient of a variable will never decrease when including additional covariates into the regression. The variance inflation factor (VIF) measures the increase of the variance. Another result from a standard linear model or experimental design course is that including additional covariates in a linear model of the outcome on the treatment indicator will never increase the variance of the OLS coefficient of the treatment at least asymptotically. This technique is called the analysis of covariance (ANCOVA), which is often used to improve the efficiency of treatment effect estimation. So we have two paradoxical results: adding covariates never decreases the variance in the first result but never increases the variance in the second result. In fact, these two results are derived under different assumptions. More precisely, the VIF result conditions on the treatment indicators but the ANCOVA result averages over them. Comparing the estimators with and without adjusting for additional covariates in a completely randomized experiment, I show that the former has smaller variance averaging over the treatment indicators, and the latter has smaller variance at the cost of a larger bias conditioning on the treatment indicators. Therefore, there is no real paradox.

Keywords: Causal inference; Conditioning; Design-based inference; Potential outcomes; Randomization; Rerandomization

2020 Mathematics Subject Classification: 62-01, 62A01, 62J10

1 Variance inflation factor

Consider the following linear regression:

$$y_i = \alpha + \tau z_i + \beta' x_i + \varepsilon_i, \quad (i = 1, \dots, n) \quad (1)$$

where z_i is a scalar and x_i is a scalar or vector. Without loss of generality, we center x so that $\bar{x} = n^{-1} \sum_{i=1}^n x_i = 0$. In a leading example, z_i is the treatment variable and x_i contains all the pre-treatment covariates. Using a standard result in linear models, we can write the OLS estimator for τ as

$$\hat{\tau}_a = \frac{\sum_{i=1}^n \check{z}_i y_i}{\sum_{i=1}^n \check{z}_i^2},$$

where \check{z}_i is the residual from the OLS fit of z_i on $(1, x_i)$. This result is also called the Frisch–Waugh–Lovell theorem in econometrics; see [1] for a recent review. If the regressors (z_i, x_i) 's are all fixed and the ε_i 's are independent and identically distributed (IID) with mean 0 and variance σ^2 as in the classic linear model, then

*Corresponding Author: Peng Ding: Department of Statistics, University of California, Berkeley, CA 94720, United States of America; Email: pengdingpku@berkeley.edu

the variance of $\hat{\tau}_a$ equals

$$\begin{aligned}\text{var}(\hat{\tau}_a) &= \frac{\sum_{i=1}^n \check{z}_i^2 \text{var}(y_i)}{(\sum_{i=1}^n \check{z}_i^2)^2} \\ &= \frac{\sigma^2}{\sum_{i=1}^n \check{z}_i^2} \\ &= \frac{\sigma^2}{\sum_{i=1}^n (z_i - \bar{z})^2} \times \frac{\sum_{i=1}^n (z_i - \bar{z})^2}{\sum_{i=1}^n \check{z}_i^2}.\end{aligned}\quad (2)$$

The first term of (2) equals the variance of

$$\hat{\tau} = \frac{\sum_{i=1}^n (z_i - \bar{z})y_i}{\sum_{i=1}^n (z_i - \bar{z})^2}, \quad (3)$$

i.e., the coefficient of z_i in the OLS fit of y_i on $(1, z_i)$ without adjusting for x_i . The second term of (2) is the VIF, which is no smaller than 1 because it is the total sum of squares divided by the residual sum of squares in the OLS fit of z_i on $(1, x_i)$. The VIF can be equivalently written as $(1 - R_{z|x}^2)^{-1}$, where $R_{z|x}^2$ is the sample R^2 between z_i and x_i . So the VIF result can also be written as

$$\text{var}(\hat{\tau}_a) = \text{var}(\hat{\tau}) \times (1 - R_{z|x}^2)^{-1}.$$

It highlights the bias-variance tradeoff: with more covariates included, the model is closer to the truth and thus leads to a smaller bias in estimating τ , but at the same time, it results in a larger variance of $\hat{\tau}_a$. See [2], [3], [4] and [5] for textbook discussions.

Thus, from (2), the variance of $\text{var}(\hat{\tau}_a)$ will never decrease with more covariates in (1), because the residual sum of squares $\sum_{i=1}^n \check{z}_i^2$ will decrease while the total sum of squares $\sum_{i=1}^n (z_i - \bar{z})^2$ is constant. An immediate result is

$$\text{var}(\hat{\tau}_a) \geq \text{var}(\hat{\tau}),$$

and the equality holds when $X'Z = 0$, where $Z = (z_1, \dots, z_n)'$ is the vector formed by regressors z_i 's and $X = (x'_1, \dots, x'_n)'$ is the matrix formed by the regressors x_i 's. The orthogonality of regressors (i.e., $X'Z = 0$) ensures that $R_{z|x}^2 = 0$ and $\hat{\tau} = \hat{\tau}_a$.

2 Analysis of covariance

Now we consider a special case of (1): the x_i 's are pre-treatment covariates, the z_i 's are the binary treatment indicators (1 for the treatment and 0 for the control), and the y_i 's are the outcomes of interest. Then (1) is the standard ANCOVA model [6], and the parameter of interest τ is the treatment effect. Let $n_1 = \sum_{i=1}^n z_i$ and $n_0 = \sum_{i=1}^n (1 - z_i)$ be the numbers of units under the treatment and control, respectively. As in Section 1, we assume that the ε_i 's are IID with mean 0 and variance σ^2 . In a completely randomized experiment, we further assume that the z_i 's are from random permutation of n_1 1's and n_0 0's.

Because z_i is binary, we can define

$$\hat{\delta}_x = n_1^{-1} \sum_{i=1}^n z_i x_i - n_0^{-1} \sum_{i=1}^n (1 - z_i) x_i, \quad \hat{\delta}_\varepsilon = n_1^{-1} \sum_{i=1}^n z_i \varepsilon_i - n_0^{-1} \sum_{i=1}^n (1 - z_i) \varepsilon_i$$

as the differences in means of x and ε , respectively. Further let $\hat{\beta}$ be the OLS estimator for β in (1). The estimator $\hat{\tau}$ in (3) without adjusting for covariates simplifies to the difference in means of the outcome:

$$\hat{\tau} = n_1^{-1} \sum_{i=1}^n z_i y_i - n_0^{-1} \sum_{i=1}^n (1 - z_i) y_i,$$

which further simplifies to

$$\begin{aligned}\hat{\tau} &= n_1^{-1} \sum_{i=1}^n z_i(\alpha + \tau z_i + \beta' x_i + \varepsilon_i) - n_0^{-1} \sum_{i=1}^n (1 - z_i)(\alpha + \tau z_i + \beta' x_i + \varepsilon_i) \\ &= \tau + \beta' \hat{\delta}_x + \hat{\delta}_\varepsilon,\end{aligned}\quad (4)$$

under (1). Based on OLS of (1), the estimator $\hat{\tau}_a$ adjusting for the covariates simplifies to

$$\begin{aligned}\hat{\tau}_a &= n_1^{-1} \sum_{i=1}^n z_i(y_i - \hat{\beta}' x_i) - n_0^{-1} \sum_{i=1}^n (1 - z_i)(y_i - \hat{\beta}' x_i) \\ &= \hat{\tau} - \hat{\beta}' \hat{\delta}_x \\ &= \tau + (\beta - \hat{\beta})' \hat{\delta}_x + \hat{\delta}_\varepsilon.\end{aligned}$$

With large samples, we can ignore the term $(\beta - \hat{\beta})' \hat{\delta}_x$ above to obtain

$$\hat{\tau}_a \approx \tau + \hat{\delta}_\varepsilon,$$

because $(\beta - \hat{\beta})' \hat{\delta}_x = O_p(n^{-1})$ is of higher order due to $\hat{\beta} - \beta = O_p(n^{-1/2})$ and $\hat{\delta}_x = O_p(n^{-1/2})$, both justified by central limit theorems under certain moment conditions. See [7] for technical details.

Under complete randomization, we can show that

$$E(\hat{\delta}_\varepsilon) = 0, \quad E(\hat{\delta}_x) = 0$$

based on a standard result for the differences in means [8],

$$\text{var}(\hat{\delta}_\varepsilon) = \frac{n}{n_1 n_0} \sigma^2, \quad \text{var}(\hat{\delta}_x) = \frac{n}{n_1 n_0} S_x^2, \quad (5)$$

where $S_x^2 = (n-1)^{-1} \sum_{i=1}^n (x_i - \bar{x})(x_i - \bar{x})'$ is the finite population covariance of the x_i 's [9], and moreover, the uncorrelatedness of the two differences in means [7]

$$\text{cov}(\hat{\delta}_\varepsilon, \hat{\delta}_x) = 0.$$

Then $E(\hat{\tau}) = \tau$ and $E(\hat{\tau}_a) \approx \tau$, i.e., $\hat{\tau}$ is unbiased and $\hat{\tau}_a$ is consistent for τ . Their variances satisfies

$$\text{var}(\hat{\tau}) - \text{var}(\hat{\tau}_a) \approx \text{var}(\beta' \hat{\delta}_x) = \frac{n}{n_1 n_0} \beta' S_x^2 \beta \geq 0.$$

Thus, if $\beta \neq 0$ then ANCOVA improves estimation efficiency, at least asymptotically. See [10], [11] and [12] for textbook discussions.

3 From conflict to unification

3.1 A unified data generating process

From the VIF result, we see that adding more covariates never decreases the variance of an OLS coefficient. In contrast, from the ANCOVA result, we see that adding more covariates never increases the variance of an OLS coefficient at least asymptotically. These two results are both standard in textbooks of linear models or experimental design. However, they seem to give opposite conclusions. Both results are derived under the linear model (1), and therefore, these two conflicting results seem paradoxical.

If we go back to check the derivations above carefully, we will find that Section 1 assumes that the z_i 's and x_i 's are both fixed, but Section 2 assumes that the z_i 's are random and the x_i 's are fixed. Therefore, the VIF and the ANCOVA results hold under different assumptions on the treatment indicators. This vaguely explains the paradox.

Technically, the settings for VIF and ANCOVA are slightly different. For example, z_i can be general and may have an arbitrary correlation structure with x_i in the VIF result, but it is binary and arises from complete randomization in the ANCOVA result. The data generating process below comes from the intersection of the settings for the two results, which allows for a more unified discussion of them.

Consider the following data generating process: for $i = 1, \dots, n$,

- (a) fix the x_i 's and center them at $\bar{x} = n^{-1} \sum_{i=1}^n x_i = 0$;
- (b) generate the potential outcomes under control as

$$y_i(0) = \alpha + \beta' x_i + \varepsilon_i,$$

where $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)$ are IID with mean 0 and variance σ^2 ;

- (c) generate the potential outcomes under treatment as

$$y_i(1) = y_i(0) + \tau,$$

i.e., the individual treatment effect $y_i(1) - y_i(0)$ is constant;

- (d) generate $Z = (z_1, \dots, z_n)$ from a random permutation of n_1 1's and n_0 0's;
- (e) obtain the observed outcome as

$$\begin{aligned} y_i &= z_i y_i(1) + (1 - z_i) y_i(0) \\ &= \tau z_i + y_i(0) \\ &= \alpha + \tau z_i + \beta' x_i + \varepsilon_i. \end{aligned} \tag{6}$$

In (b) and (c), I use the potential outcomes notation due to [8]. Readers who are uncomfortable with the notation $y_i(1)$ and $y_i(0)$ can ignore steps (b) and (c) and view (6) as the data generating process with random ε_i 's and z_i 's.

In the above data generating process, τ represents the individual and thus the average treatment effect. It is the parameter of interest.

3.2 Comparing the variances

Conditional on Z , (6) is a linear model with fixed (z_i, x_i) 's and homoskedastic errors ε_i 's. The discussion in Section 1 applies in this case. Then from the VIF result, we know that

$$\text{var}(\hat{\tau}_a | Z) \geq \text{var}(\hat{\tau} | Z), \tag{7}$$

i.e., the estimator adjusting for covariates x_i 's has larger variance. However, $\hat{\tau}_a$ is unbiased but $\hat{\tau}$ is biased. From the classic OLS theory,

$$E(\hat{\tau}_a | Z) = \tau,$$

and from the formula of (4), the bias of $\hat{\tau}$ is

$$E(\hat{\tau} | Z) - \tau = \beta' \hat{\delta}_x.$$

Therefore, the smaller conditional variance of $\hat{\tau}$ comes at the cost of having a larger conditional bias. The conditional bias of $\hat{\tau}$ vanishes only in the special cases with $\beta = 0$ or $\hat{\delta}_x = 0$, that is, the covariates are unrelated with the outcome, or the covariates are perfectly balanced in means across the treatment and control groups.

Averaging over Z , we have random potential outcomes and random treatment indicators. The discussion in Section 2 applies in this case. We have shown that

$$E(\hat{\tau}) = \tau, \quad E(\hat{\tau}_a) \approx \tau,$$

and moreover, asymptotically (ignoring higher order terms),

$$\text{var}(\hat{\tau}) \geq \text{var}(\hat{\tau}_a). \tag{8}$$

Mathematically, the efficiency reversal results in (7) and (8) do not lead to contradiction given the explicitly specified conditioning sets. Statistically, however, they form a paradox that is similar to the classic Simpson's paradox of effect reversal due to different conditioning sets. I give a simple explanation of this paradox using the following decompositions based on the law of total variance:

$$\begin{aligned}\text{var}(\hat{\tau}_a) &= E\{\text{var}(\hat{\tau}_a | Z)\} + \text{var}\{E(\hat{\tau}_a | Z)\} \\ &= E\{\text{var}(\hat{\tau}_a | Z)\} + \text{var}(\tau) \\ &= E\{\text{var}(\hat{\tau}_a | Z)\}\end{aligned}$$

and

$$\begin{aligned}\text{var}(\hat{\tau}) &= E\{\text{var}(\hat{\tau} | Z)\} + \text{var}\{E(\hat{\tau} | Z)\} \\ &= E\{\text{var}(\hat{\tau} | Z)\} + \text{var}(\tau + \beta' \hat{\delta}_x) \\ &= E\{\text{var}(\hat{\tau} | Z)\} + \text{var}(\beta' \hat{\delta}_x).\end{aligned}$$

Based on the VIF result, $E\{\text{var}(\hat{\tau}_a | Z)\} \geq E\{\text{var}(\hat{\tau} | Z)\}$, but their difference is small because the $R_{z|x}^2$ between z_i and x_i is close to zero under complete randomization. We can ignore their difference in the asymptotic analysis. More importantly, the unadjusted estimator has an additional variance term due to the conditional bias which reverses the ordering of $\text{var}(\hat{\tau}_a)$ and $\text{var}(\hat{\tau})$.

3.3 Comparing the estimated variances

Sections 1 and 2 compare the variances of $\hat{\tau}_a$ and $\hat{\tau}$ which are theoretical quantities depending on the unknown true data generating process. In practice, standard statistical software packages report the estimated variances based on OLS:

$$\hat{\text{var}}(\hat{\tau}_a) = \frac{\hat{\sigma}_{y|z,x}^2}{\sum_{i=1}^n (z_i - \bar{z})^2} \times \frac{1}{1 - R_{z|x}^2}$$

and

$$\hat{\text{var}}(\hat{\tau}) = \frac{\hat{\sigma}_{y|z}^2}{\sum_{i=1}^n (z_i - \bar{z})^2},$$

where $\hat{\sigma}_{y|z,x}^2$ equals the residual sum of squares divided by $n - 2 - \dim(x)$ in the OLS fit of y_i on $(1, z_i, x_i)$, and $\hat{\sigma}_{y|z}^2$ equals the residual sum of squares divided by $n - 2$ in the OLS fit of y_i on $(1, z_i)$. The ratio of these two variances depends on $\hat{\sigma}_{y|z,x}^2 / \hat{\sigma}_{y|z}^2$ and $R_{z|x}^2$, which can be larger or smaller than 1. Importantly, this is a numeric result regardless of whether or not we condition on Z .

In fact, under the data generating process in Section 3.1, $\hat{\text{var}}(\hat{\tau}_a)$ is often smaller than $\hat{\text{var}}(\hat{\tau})$ as long as the covariates are predictive to the outcome. This is true due to two basic facts: first, $R_{z|x}^2$ is close to 0 under complete randomization so the VIF can be ignored asymptotically; second, $\hat{\sigma}_{y|z,x}^2$ is often smaller than $\hat{\sigma}_{y|z}^2$ because the residual sum of squares decreases with an additional predictive covariate. This argument ignores the opposite impact of the degrees of freedom correction, which is reasonable when the sample size is large and the dimension of covariates is small. See [13] for the discussion of R^2 with high dimensional covariates.

The above heuristic comparison of $\hat{\text{var}}(\hat{\tau}_a)$ and $\hat{\text{var}}(\hat{\tau})$ is not in contradiction with the VIF result which concerns the true variances conditional on Z . The estimated variances can be different from the true variances, especially when the linear model of y_i on $(1, z_i)$ is misspecified.

4 Connection with randomization inference

4.1 Variances conditional on the error terms

Another conditioning scheme leads to the following discussion beyond Sections 1 and 2. Conditional on the error terms ε , we have fixed potential outcomes and completely randomized Z . Statistics under this regime is

called randomization inference, or design-based inference. The classic results from randomization inference are $E(\hat{\tau} | \mathcal{E}) = \tau$ [8], $E(\hat{\tau}_a | \mathcal{E}) \approx \tau$ [14, 15], and $\text{var}(\hat{\tau} | \mathcal{E}) \geq \text{var}(\hat{\tau}_a | \mathcal{E})$ asymptotically [14, 15]. This relative efficiency is coherent with $\text{var}(\hat{\tau}) \geq \text{var}(\hat{\tau}_a)$ asymptotically. Again we can explain the coherence using the following decompositions based on the law of total variance:

$$\begin{aligned}\text{var}(\hat{\tau}_a) &= E\{\text{var}(\hat{\tau}_a | \mathcal{E})\} + \text{var}\{E(\hat{\tau}_a | \mathcal{E})\} \\ &\approx E\{\text{var}(\hat{\tau}_a | \mathcal{E})\} + \text{var}(\tau) \\ &= E\{\text{var}(\hat{\tau}_a | \mathcal{E})\}\end{aligned}$$

and

$$\begin{aligned}\text{var}(\hat{\tau}) &= E\{\text{var}(\hat{\tau} | \mathcal{E})\} + \text{var}\{E(\hat{\tau} | \mathcal{E})\} \\ &= E\{\text{var}(\hat{\tau} | \mathcal{E})\} + \text{var}(\tau) \\ &= E\{\text{var}(\hat{\tau} | \mathcal{E})\},\end{aligned}$$

where both $\hat{\tau}_a$ and $\hat{\tau}$ are unbiased for τ at least asymptotically. These results are all in favor of ANCOVA which improves the estimation efficiency. I summarize the results in Table 1.

Table 1: Comparison under the data generating process in (a)–(e)

	mean of $\hat{\tau}_a$	mean of $\hat{\tau}$	variance comparison
unconditional	$E(\hat{\tau}_a) = \tau$	$E(\hat{\tau}) = \tau$	$\text{var}(\hat{\tau}_a) \leq \text{var}(\hat{\tau})$ asymptotically
conditional on Z	$E(\hat{\tau}_a Z) = \tau$	$E(\hat{\tau} Z) = \tau + \beta' \hat{\delta}_x$	$\text{var}(\hat{\tau}_a Z) \geq \text{var}(\hat{\tau} Z)$
conditional on \mathcal{E}	$E(\hat{\tau}_a \mathcal{E}) \approx \tau$	$E(\hat{\tau} \mathcal{E}) = \tau$	$\text{var}(\hat{\tau}_a \mathcal{E}) \leq \text{var}(\hat{\tau} \mathcal{E})$ asymptotically

4.2 More general potential outcomes model

The data generating process in (a)–(e) assumes constant treatment effect and homoskedastic errors. It yields the standard ANCOVA model (1) or (6). The literature of randomization-based causal inference often deals with more general potential outcomes, without requiring these linear model assumptions [7–9, 14, 15]. In general cases with possibly mis-specified linear models, [14] criticized ANCOVA by showing that it might increase or decrease the efficiency compared to $\hat{\tau}$. As a response to this critique, [15] proposed to use a modified ANCOVA estimator that also includes the interaction term $z_i \times x_i$ in OLS, and showed that this estimator is at least as efficient as $\hat{\tau}$ asymptotically. See [16] and [17] for related discussions.

4.3 A design issue

From the above discussion, $\hat{\delta}_x$ is a key quantity that causes the paradox. If it is zero, then $\hat{\tau}_a = \hat{\tau}$ and the paradox disappears. Complete randomization ensures $E(\hat{\delta}_x) = 0$, but in a particular allocation, $\hat{\delta}_x$ can differ from zero. From the experimental design perspective, $\hat{\delta}_x$ measures the covariate balance across the treatment and control groups. Complete randomization ensures covariate balance on average, but a particular allocation may have covariate imbalance. [18] and [19] proposed to use rerandomization to improve the data generating process (d) by forcing the treatment indicators Z to satisfy

$$\hat{\delta}_x' \text{cov}(\hat{\delta}_x)^{-1} \hat{\delta}_x = \hat{\delta}_x' \left(\frac{n}{n_1 n_0} S_x^2 \right)^{-1} \hat{\delta}_x \leq c_0,$$

where $c_0 > 0$ is a predetermined threshold. Under the randomization inference framework, [20] show that this new experimental design improves the efficiency of $\hat{\tau}$, which is, in fact, close to the efficiency of $\hat{\tau}_a$ for

small $c_0 \approx 0$. From our discussion before, rerandomization can also reduce the conditional bias of $\hat{\tau}$ given Z because it forces $\hat{\delta}_x$ to be small for any realized value of Z . Therefore, rerandomization can mitigate the paradox through experimental design. See [21] for more unified discussions.

5 Final remarks

I have shown that the seemingly paradoxical results of VIF and ANCOVA are due to different statistical assumptions. The key difference is whether or not statistical inference is conditional on the treatment indicators Z . Conditioning on Z , the unadjusted estimator has a smaller variance but larger bias. Averaging over Z , both unadjusted and adjusted estimators are consistent for τ but the variance of the adjusted estimator is no larger than that of the unadjusted estimator. In randomized experiments, we recommend using ANCOVA under a constant treatment effect model or its modified version for general settings [15, 21].

I end this note with two minor technical issues. First, I assume that the x_i 's are fixed throughout the paper. With random covariates, we can condition on them to obtain the same results. The key in the discussion is whether or not to condition on Z . Second, if the Z_i 's are IID Bernoulli random variables as in a Bernoulli experiment, we can condition on (n_1, n_0) to reduce the discussion to complete randomization.

Acknowledgement: The author thanks Ugur Yildirim for raising this question in his class of “Stat 230A Linear Models” at the University of California Berkeley, Luke Miratrix for the inspiring discussion of estimated variances in Section 3.3, and Anqi Zhao, Xinran Li, Liyun Chen, Zhichao Jiang, Jason Wu, and Yuting Ye for helpful suggestions. A reviewer made many constructive comments which helped to improve the paper significantly. This research was partially supported by the U. S. National Science Foundation (grant # 1945136).

Conflict of Interests: Prof. Peng Ding is a member of the Editorial Board of the Journal of Causal Inference although had no involvement in the final decision.

References

- [1] P. Ding. The Frisch–Waugh–Lovell theorem for standard errors. *Statistics and Probability Letters*, 168:108945, 2021.
- [2] G. James, D. Witten, T. Hastie, and R. Tibshirani. *An Introduction to Statistical Learning*. New York: Springer, 2013.
- [3] J. J. Faraway. *Linear Models with R*. Boca Raton: Chapman and Hall/CRC, 2016.
- [4] J. Fox. *Applied Regression Analysis and Generalized Linear Models*. Newbury Park, CA: Sage Publications, 2015.
- [5] A. Agresti. *Foundations of Linear and Generalized Linear Models*. New York: John Wiley & Sons, 2015.
- [6] R. A. Fisher. *Statistical Methods for Research Workers*. Edinburgh: Oliver and Boyd, 1stst edition, 1925.
- [7] X. Li and P. Ding. General forms of finite population central limit theorems with applications to causal inference. *Journal of the American Statistical Association*, 112:1759–1769, 2017.
- [8] J. Neyman. On the application of probability theory to agricultural experiments: Essay on principles, Section 9. Masters Thesis. Portions translated into english by D. Dabrowska and T. Speed (1990). *Statistical Science*, 5:465–472, 1923.
- [9] G. W. Imbens and D. B. Rubin. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*. New York: Cambridge University Press, 2015.
- [10] O. Kempthorne. *The Design and Analysis of Experiments*. New York: Wiley, 1952.
- [11] K. Hinkelmann and O. Kempthorne. *Design and Analysis of Experiments, Volume 1, Introduction to Experimental Design, 2nd Edition*. New York: John Wiley & Sons, 2007.
- [12] D. R. Cox and N. Reid. *The Theory of the Design of Experiments*. Boca Raton: Chapman and Hall/CRC, 2000.
- [13] D. A. Freedman. A note on screening regression equations. *The American Statistician*, 37:152–155, 1983.
- [14] D. A. Freedman. On regression adjustments to experimental data. *Advances in Applied Mathematics*, 40:180–193, 2008.
- [15] W. Lin. Agnostic notes on regression adjustments to experimental data: Reexamining Freedman’s critique. *Annals of Applied Statistics*, 7:295–318, 2013.
- [16] A. A. Tsiatis, M. Davidian, M. Zhang, and X. Lu. Covariate adjustment for two-sample treatment comparisons in randomized clinical trials: A principled yet flexible approach. *Statistics in Medicine*, 27:4658–4677, 2008.

- [17] A. Negi and J. M. Wooldridge. Revisiting regression adjustment in experiments with heterogeneous treatment effects. *Econometric Reviews*, page in press, 2020.
- [18] D. R. Cox. Randomization and concomitant variables in the design of experiments. In P. R. Krishnaiah G. Kallianpur and J. K. Ghosh, editors, *Statistics and Probability: Essays in Honor of C. R. Rao*, pages 197–202. North-Holland, Amsterdam, 1982.
- [19] K. L. Morgan and D. B. Rubin. Rerandomization to improve covariate balance in experiments. *The Annals of Statistics*, 40:1263–1282, 2012.
- [20] X. Li, P. Ding, and D. B. Rubin. Asymptotic theory of rerandomization in treatment-control experiments. *Proceedings of the National Academy of Sciences of the United States of America*, 115:9157–9162, 2018.
- [21] X. Li and P. Ding. Rerandomization and regression adjustment. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82:241–268, 2020.