Basel Biometric Society Meeting on Covariate Adjustment in Clinical Trials

Discussion

Tim P. Morris

25 March 2025

Senn

Senn: Compared to what?

Two clarifications, which may be useful

1. In the first of the three consequences, the ‘mean square error effect’, the comparator is an **analysis** that **ignores covariates** completely.
2. In the second consequence, the ‘Variance inflation factor’, the comparator changes to a **design** that **perfectly balances covariates**.

My proposition is then that VIF is not really to do with covariate adjustment, since the consequences of imbalance vs. perfect balance affect the estimator whether or not you adjust; adjustment is the way we acknowledge this in inference.

Senn: Two annotations (I think) to the table

		Effect		
		Estimated mean square error effect	Imbalance effect or VIF	Second -order precision
Influence	Design	No*	Yes	No
	Model	Yes	Yes	Yes
	Outcome	Yes	No	No

'Model' = **model used by the analyst**

'Outcome' = **model/mechanism that produced the data**

Senn: Using this

I think it also provides a useful perspective on

- How this works out for 'ANCOVA2' (Tsiatis et al., 2008), subsequently further justified by Lin (2013). I will return to this for Jack Kuipers' talk.
- Estimators based on weighting, for example inverse-probability-of-treatment weighting (Williamson et al., 2014), and overlap weighting (Zeng et al., 2020). Note: these papers do not 'cheat', but they do not use propensity scores in the way Stephen prescribes!

Kuipers

Kuipers

I was very interested by this talk and the accompanying paper but it is less familiar to me and so I feel less qualified to comment and throw in further interpretations.

My own reason not to adjust would be the possibility of small-sample bias (I have bonus slides if time). To be clear, I don't see this as a real problem (Tackney et al., 2023).

Kuipers: a question

An apparent contradiction. Tsiatis et al. (2008) showed that for $E(Y|do(X=1)) - E(Y|do(X=0))$ adjustment cannot harm (large-sample) relative precision with linear outcome models, two treatment arms and 1:1 randomization.

More generally, from Lin (2013):

> *'OLS adjustment cannot hurt asymptotic precision when a full set of treatment × covariate interactions is included'*

(Note: covariates must be centered.)

So my question is simply: what is different here? The estimator itself or finite sample size?

Magirr

Magirr: notes from a recovering conditionalist

I was initially a ‘conditionalist’, after reading Hauck et al. (1998). The claim is that a clinical trial should aim to ‘predict the direction and size of a treatment benefit for a patient with specific covariate values’ (personalized treatment effects), not average effects for groups of patients.

- Laudable (aim)
- Laughable (estimation)

My guess is that Dominic is too diplomatic to come out swinging for marginal measures.

Magirr: conditional estimation made easy?

An epidemiology 101 'trick' makes life simple: estimate the treatment effect after **restricting** the analysis, for example to males aged YY with stage X disease and...

'What's that? Your trial only had two such participants? Too bad. That's the information you have.' [Note: Stephen earlier described this as 'a fool's errand']

Touted solution is then to buy information with assumptions, e.g. $\text{logit}P(Y = 1) = \alpha_0 + \alpha_1 \mathbb{I}(\text{trt} = 1) + \alpha_2 X$.

This only targets the average conditional effects (for each combination of covariates) if the model is right. It's probably not right, however hard we wish.

Magirr: conditional estimation with supercovariates

Do you ever believe any risk prediction model is ‘correct’?

If a prediction model is not correct, adjustment for a prognostic score using a model for a conditional estimand will targets a sort of amorphous hybrid of conditional and marginal measures.

This is a strong nudge towards marginal measures when using supercovariates.

W. W. Hauck, S. Anderson, and S. M. Marcus. Should we adjust for covariates in nonlinear regression analyses of randomized trials? *Controlled Clinical Trials*, 19(3):249–256, 1998. doi: 10.1016/s0197-2456(97)00147-5.

W. Lin. Agnostic notes on regression adjustments to experimental data: Reexamining freedman's critique. *The Annals of Applied Statistics*, 7(1), 2013. doi: 10.1214/12-aoas583.

M. S. Tackney, T. P. Morris, I. R. White, C. Leyrat, K. Diaz-Ordaz, and E. Williamson. A comparison of covariate adjustment approaches under model misspecification in individually randomized trials. *Trials*, 24(1), 2023. doi: 10.1186/s13063-022-06967-6.

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(23):4658–4677, 2008. doi: 10.1002/sim.3113.

E. J. Williamson, A. Forbes, and I. R. White. Variance reduction in randomised trials by inverse probability weighting using the propensity score. *Statistics in Medicine*, 33(5):721–737, 2014. doi: 10.1002/sim.5991.

S. Zeng, F. Li, R. Wang, and F. Li. Propensity score weighting for covariate adjustment in randomized clinical trials. *Statistics in Medicine*, 40(4):842–858, 2020. doi: 10.1002/sim.8805.

Bias bonus

ANCOVA is small-sample biased

I want to start with a point that none of the speakers have mentioned: covariate adjustment can lead to small-sample bias.

This is mentioned in the supplementary materials of Tackney et al. (2023) but I think not widely known.

1 million data points; true average effect ≈ 2

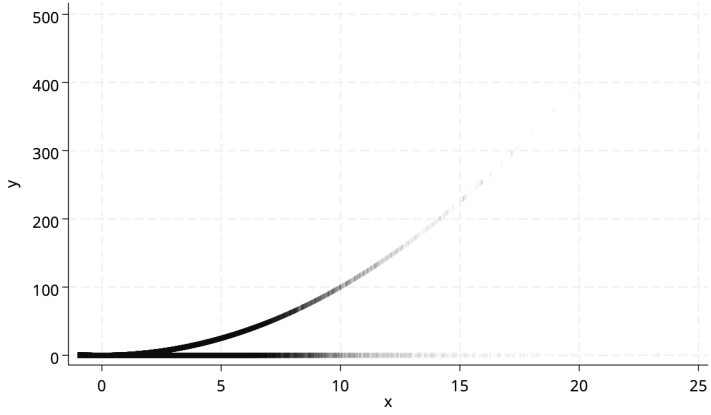

Bias of ANCOVA assuming $E[Y]$ a linear function of X