Statistical Thinking Comment Policy

Comments are welcomed. Be informed, informative, respectful, and criticize ideas, not people. Please read our Comment Policy before commenting.

Weihua Cao • 6 months ago

Great post! I used to agree with the argument "When Y is continuous and random errors have a normal distribution, it is well known that classical covariate adjustment improves power over an unadjusted analysis no matter how poorly the model fits", until I saw the paper by Freeman "ON REGRESSION ADJUSTMENTS IN EXPERIMENTS WITH SEVERAL TREATMENTS". He concludes that ANCOVA is not necessarily more efficient that ANOVA.

Frank Harrell Mod ➔ Weihua Cao • 6 months ago

Thanks for the comments. Freedman (2008) assumes that the treatment effect is different for every patient. I don't. He also has a false premise that ANCOVA exists to gain precision from adjusting for minor imbalances in covariate distributions across treatments. It doesn't. I also have issues with his 2006 (it might be the 2007 one) paper which in my mind is all about possible harms of lack of fit of the covariate adjustment model, when in fact lack of fit is not very consequential. You have to compare the lack of fit to the alternative approach, and as I exemplified in the blog article the alternative is unlikely to be better than faulty covariate adjustment.

Nicholas Lewin-Koh • a year ago

This seems to me to be an extension of the old argument about Type III sum of squares, what exactly does averaging over interactions mean? Showing my age here, Bill Venables made a very cognizant argument for treating interactions with caution, and not doing marginal inference averaging over interactions (https://www.stats.ox.ac.uk/.... It seems to me a very small jump that this situation becomes worse with the addition of a non-linear link function applied to the linear predictor.
On the other hand, I can understand why marginal models might be desirable in spite of the potential bias. In clinical trials small effect sizes are very rarely worth considering when assessing benefit for disease therapies. Ideally, a good therapeutic candidate should overcome population heterogeneity if the effect is clinically meaningful (and clinically useable).
I fall more on the side of conditional models when possible.

• Reply • Share ›

Frank Harrell Mod → Nicholas Lewin-Koh • a year ago

Nice points. Terry Therneau has pointed out that in the (nonlinear) Cox proportional hazards model, SAS got it completely wrong when implementing type III tests: htt ps://cran.r-project.org/web...

As you said, Bill Venables has been rightly critical of the idea behind type III. For example, the idea of getting a marginal treatment estimate were there to be an equal number of patients with severe as with mild disease is preposterous.

The only place I can see where marginalizing is beneficial is when you have patient types that are rare, and giving type-specific treatment effect estimates is low precision.

Related

- Commentary on Improving Precision and Power in Randomized Trials for COVID-19 Treatments Using Covariate Adjustment, for Binary, Ordinal, and Time-to-Event Outcomes
- Avoiding One-Number Summaries of Treatment Effects for RCTs with Binary Outcomes
- If You Like the Wilcoxon Test You Must Like the Proportional Odds Model
- Unadjusted Odds Ratios are Conditional
- RCT Analyses With Covariate Adjustment