Planning a covariate adjustment method for your SAP

Tim Morris, MRC Clinical Trials Unit at UCL

1 Dec 2022
PSI Webinar: Covariate Adjustment – Considerations and Examples

Smarter Studies
Global Impact
Better Health

Outline

1 Covariate adjustment

2 Three broad approaches

3 Choosing an approach to write into a statistical analysis plan

4 An example

Why adjust?

- I would usually start this talk explaining the key reasons for covariate adjustment

- Thanks to Jonathan Bartlett and Florian Voß, I don’t have to cover that today!

Planning a method

One of my colleagues really pushed me on my generic advice to use covariate adjustment: ‘What about this? What about that?’

Each *what if* was a practical concern. Today I’ll work through the more interesting ones.

The *GetTested* trial (not the ones my colleague was asking about)

Designed to assess the effectiveness of an internet-accessed sexually transmitted infection (STI) testing and results service, compared with going to a walk-in sexual health clinic.

PCR testing will mean the service should be familiar:

1. Order a test for chlamydia, gonorrhoea, HIV, and syphilis
2. Take test
3. Post it off and receive result/s

The *GetTested* trial

Covariates

Treatment allocation balanced several prognostic covariates using a ‘minimisation’ procedure:

- Gender (m/f/trans)
- 1 or 2+ sexual partners in last 12m
- Sexual orientation

Outcomes

Two binary outcomes of interest (6 weeks):

- STI test (anticipated 10%)
- Diagnosis (anticipated 0.6%)

The *GetTested* trial

Recruited 2,072 participants in two South London boroughs

Treatment effect to be summarised as a risk ratio and risk difference

Three broad approaches (for today)

These go by various names. I will use:

1. Direct adjustment
2. Standardisation
3. Inverse probability of treatment weighting

When the FDA draft guidance appeared, we* were really surprised that there was no mention of weighting approaches

*myself, Ian White (MRC CTU), Fizz Williamson (LSHTM), Brennan Kahan (LSHTM)…

Direct adjustment
‘Outcome regression’ / ‘ancova’

Fit a regression model including randomized arm and covariates

The estimated treatment effect is the $\hat{\beta}$ that pops out

Standardisation
‘g-computation’ / ‘marginalisation’

Fit some regression model, leave covariates at their natural values and predict outcome:

1. $\widehat{P_0}$ with treatment set to 0
2. $\widehat{P_1}$ with treatment set to 1

Use these to predict measures of interest, such as

$\frac{\widehat{P_1}}{\widehat{P_0}}$ (risk ratio)

$\widehat{P_1} - \widehat{P_0}$ (risk difference)

Inverse probability of treatment weighting

Fit some model to estimate $\widehat{e_i} = \Pr(\mathrm{trt} = 1 | \mathrm{X_i})$ and use estimated probabilities to weight the observations.

Research arm: $^{1}/_{1-\widehat{e_i}}$

Control arm: $^{1}/_{\widehat{e_i}}$

For example, you might fit a logistic regression model where ‘outcome’ is treatment. For each individual in the dataset, predict the probability $\widehat{e_i}$ of being randomised to ‘arm 1’.

Inverse probability of treatment weighting

This is really counterintuitive! We know the **true** probability of being randomised to each arm; why would we estimate it?

(Answer: using $\widehat{e_i}$ instead of e_i reduces the variance of the treatment effect)

Choosing between the three approaches?

Particularly for binary outcome measures

Summary measure: marginal or conditional?

- This matters for **non-collapsible** effect measures; most commonly the odds ratio and the hazard ratio

- When I previously spoke on this topic, I wrote ‘If you struggle with the idea of non-collapsibility, see the superb recent paper by Rhian Daniel and colleagues’ – lucky you, you got to hear it direct from Rhian!

What measure do they target?

Direct adjustment:	Conditional
Standardisation:	Marginal
IPTW:	Marginal

Convergence

- A **really** important practical point
- You don’t want to write a procedure into your SAP that risks non-convergence or separation

Convergence

Direct adjustment can be really risky

- Suppose you want to estimate a risk difference and use an identity-link binomial model
- This is notoriously unstable, may converge to the wrong maximum, and adding an observation can greatly increase or reduce the SE

Convergence

Standardisation can be better

- Suppose you want to estimate a risk difference/ratio
- You could (for example) use a logistic regression model adjusting for covariates, then standardise to estimate the risk difference/ratio
- Logistic regression uses a canonical link, so converges to a unique maximum, plus other magic (White, Morris, Williamson)
- The main risk is **separation**

Convergence

IPTW is safest

- Unlike direct adjustment and standardisation, IPTW is a simple contrast of means (a difference or a ratio) so this will converge
- No covariates in this model; each observation contributes (slightly) different relative weight
- The weighting model may not converge, but two useful features give it a good shot:
  1. Outcome ‘incidence’ is about 50% with 1:1 randomisation
  2. In truth all covariate effects on randomisation are null

Precision/power

- Standardisation can simply fit the direct adjustment model and then standardise. In this case, power would clearly be the same (though not necessarily precision if this changes the estimand)
- IPTW is (asymptotically) as efficient as direct adjustment using OLS and simulation studies show very little difference at small sample sizes (Williamson, 2014)

Handling balanced covariates? (no worries)

- Long history of how direct adjustment can do this
- Standardisation model can be identical to direct adjustment, so clearly also does this
- IPTW has a different estimation method but also handles designed balance: the variance formula takes the balance that it sees

Standard errors?

Direct adjustment:	[Presumably familiar]
Standardisation:	Delta method (asymptotic)
IPTW:	Robust SE's accounting for estimation of $\widehat{e_i}$ (asymptotic)

Model misspecification (i)

Consider the mean function relating covariates X to outcome Y

The true model is $Y_i = \alpha + \theta Z_i + \gamma X_i^2$
We get the data and either

1. Fit $y_i = \hat{\alpha} + \hat{\theta} z_i + \hat{\gamma} x_i$
2. Use x to predict z then use IPTW

Model misspecification (i)

- Depending on what range of x we actually see in the trial, adjustment may do very little or quite a lot
- Importantly, IPTW is exactly the same as direct adjustment – no advantage

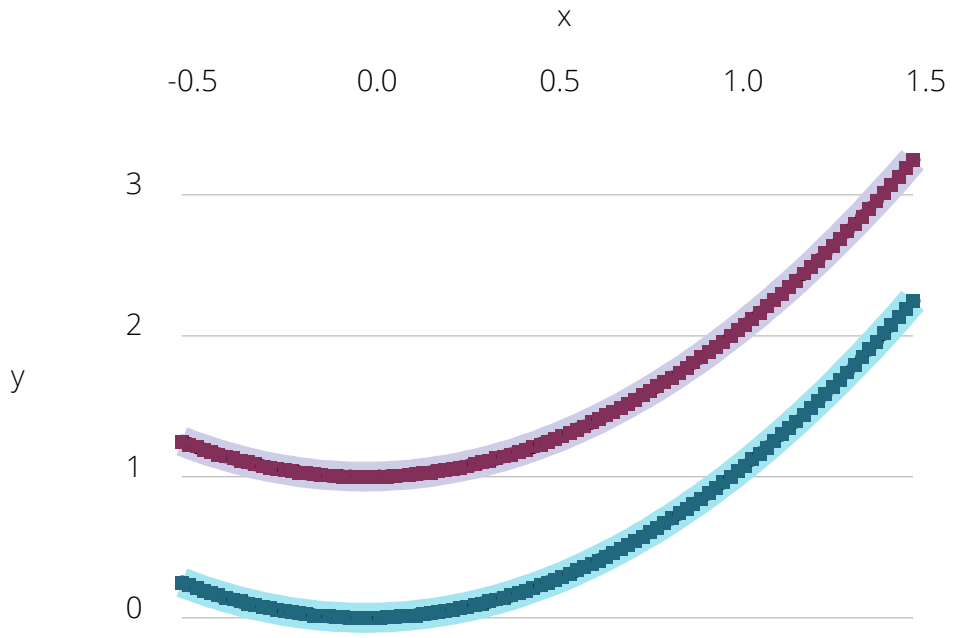

Model misspecification (i)

Depending on what range of x we see in the trial, adjustment may do very little or quite a lot

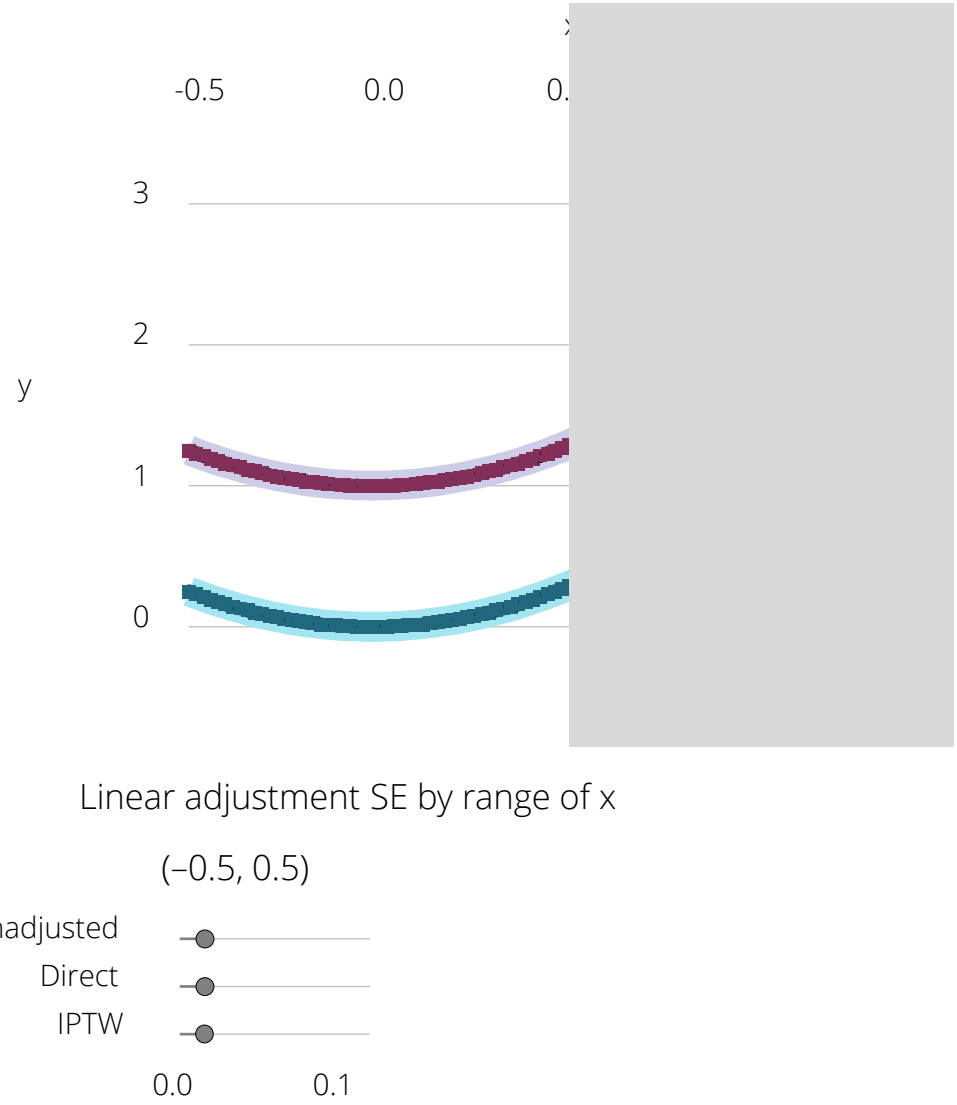

Model misspecification (i)

Depending on what range of x we see in the trial, adjustment may do very little or quite a lot

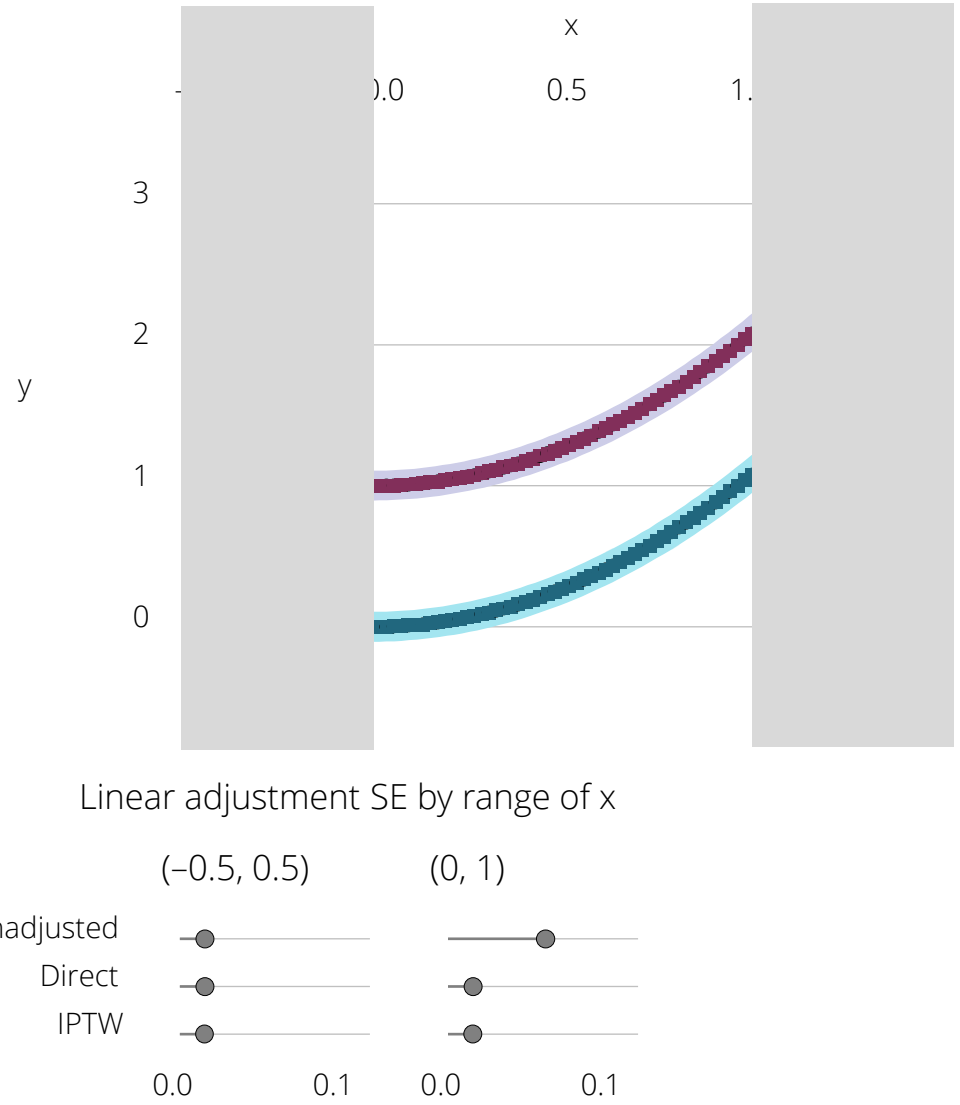

Model misspecification (i)

Depending on what range of x we see in the trial, adjustment may do very little or quite a lot
Importantly, IPTW is exactly the same as direct adjustment – so no advantage!

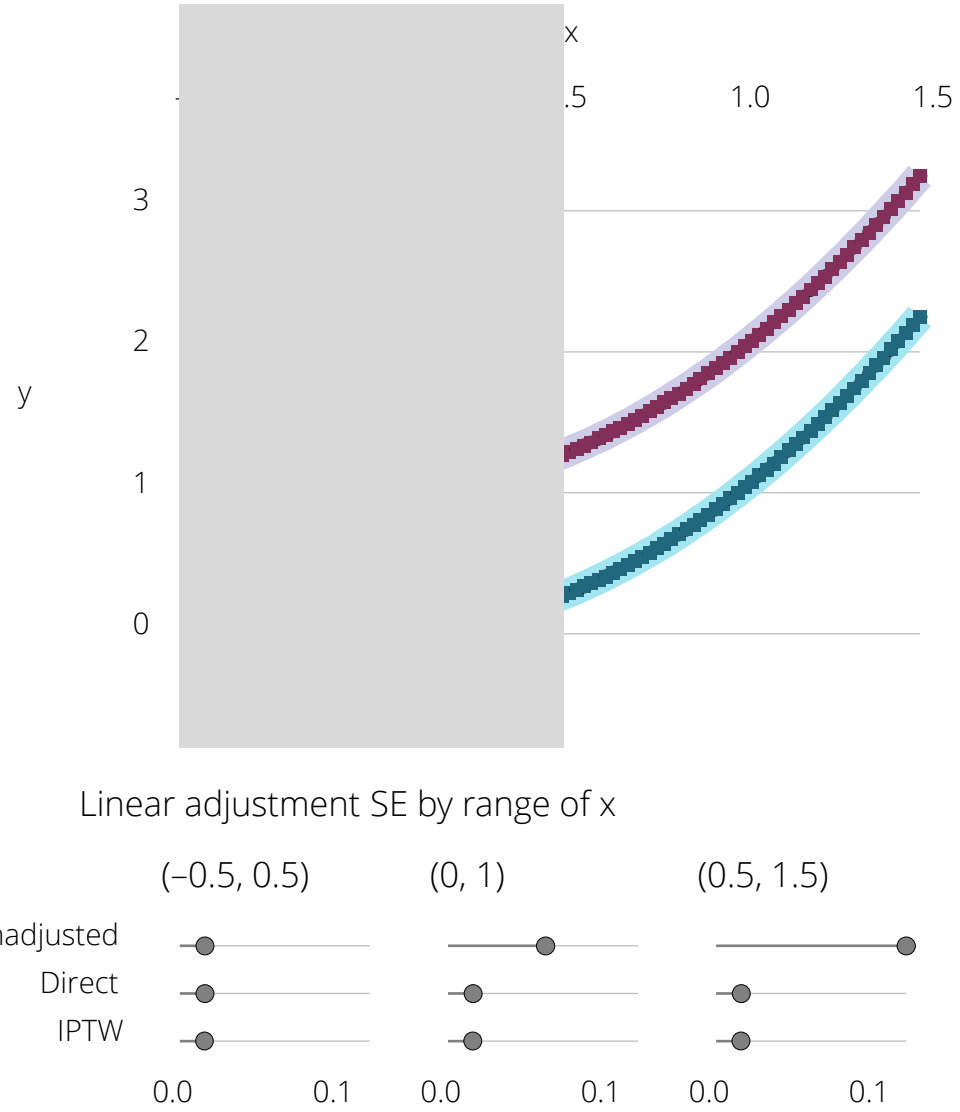

Model misspecification (ii)

I mentioned the magic of canonical link functions. These prevent misspecification.

In the figure on the next slide, I have fixed values of a binary covariate and treatment, but systematically varied the number of positive binary outcomes.

Analysed each dataset using direct adjustment with different link functions.

(For more detail see White, Morris & Williamson arxiv.org/abs/2107.07278)

Bland–Altman plot

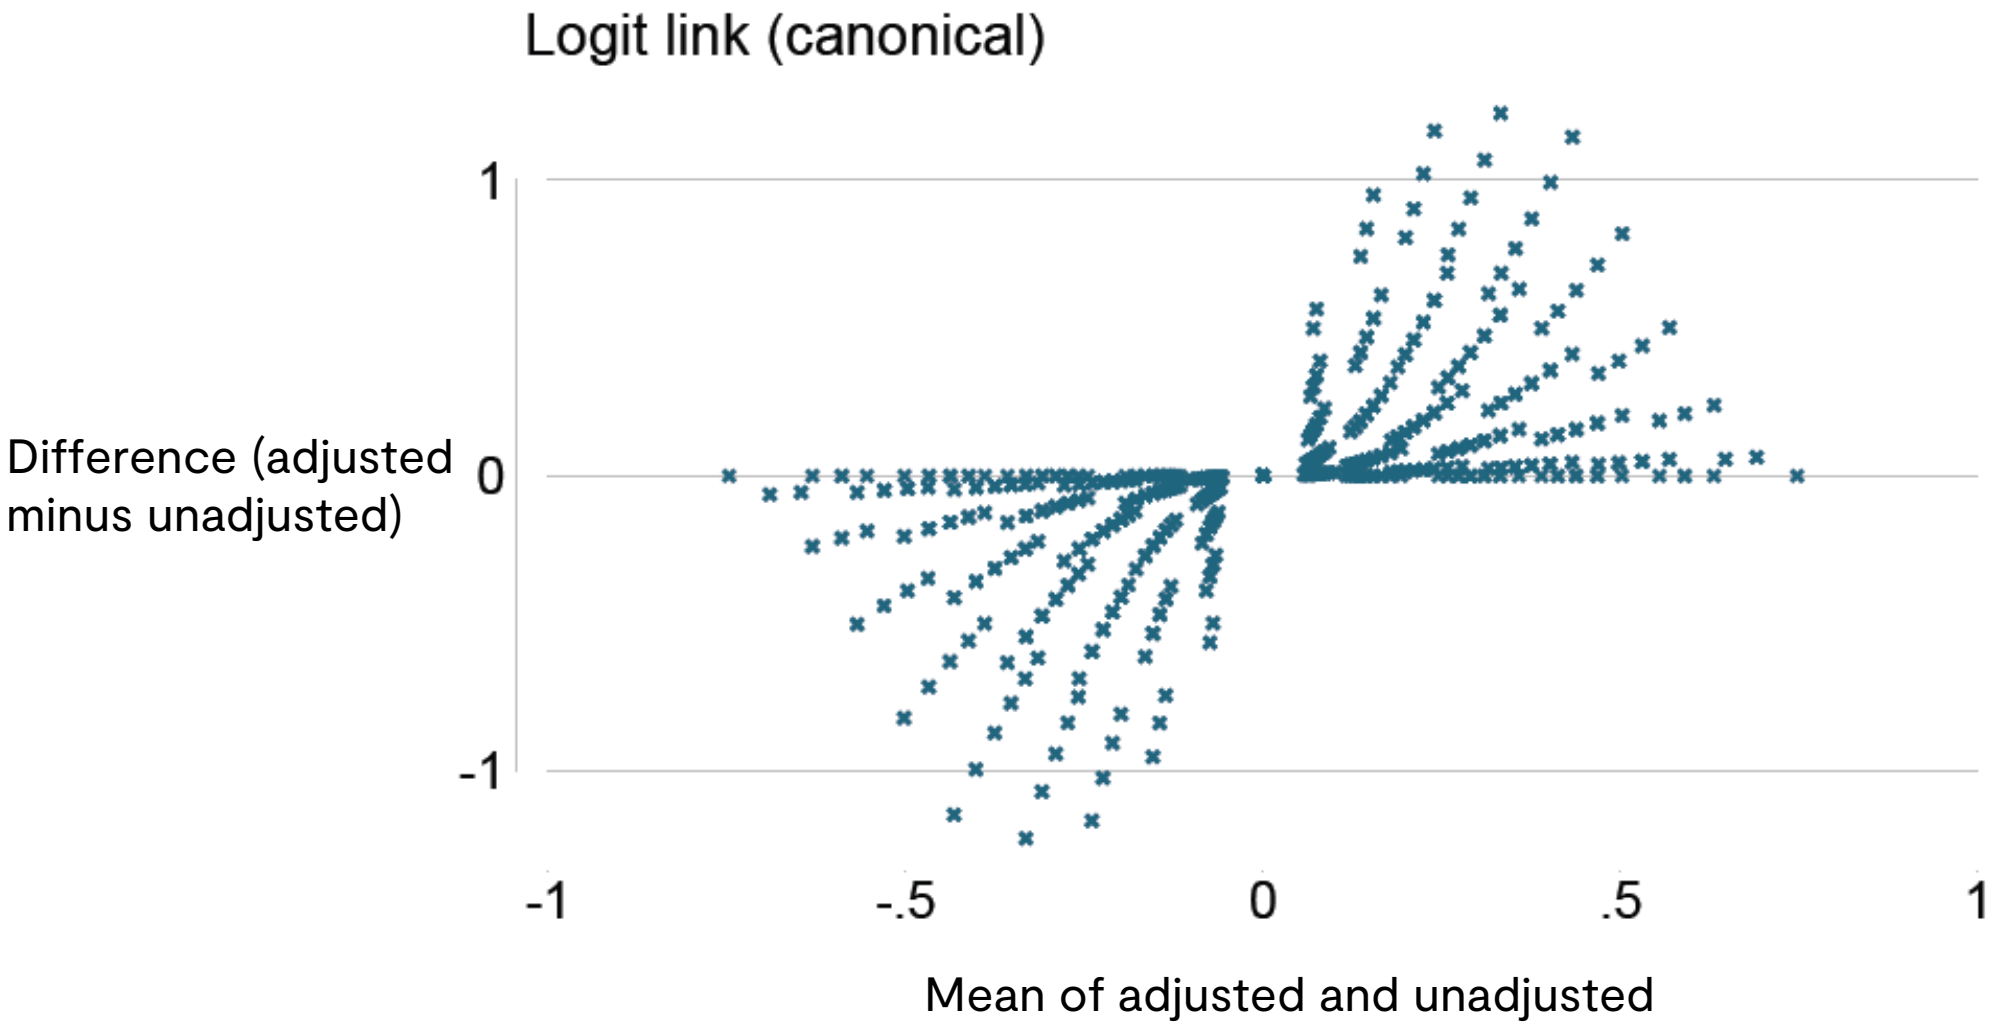

Bland–Altman plot

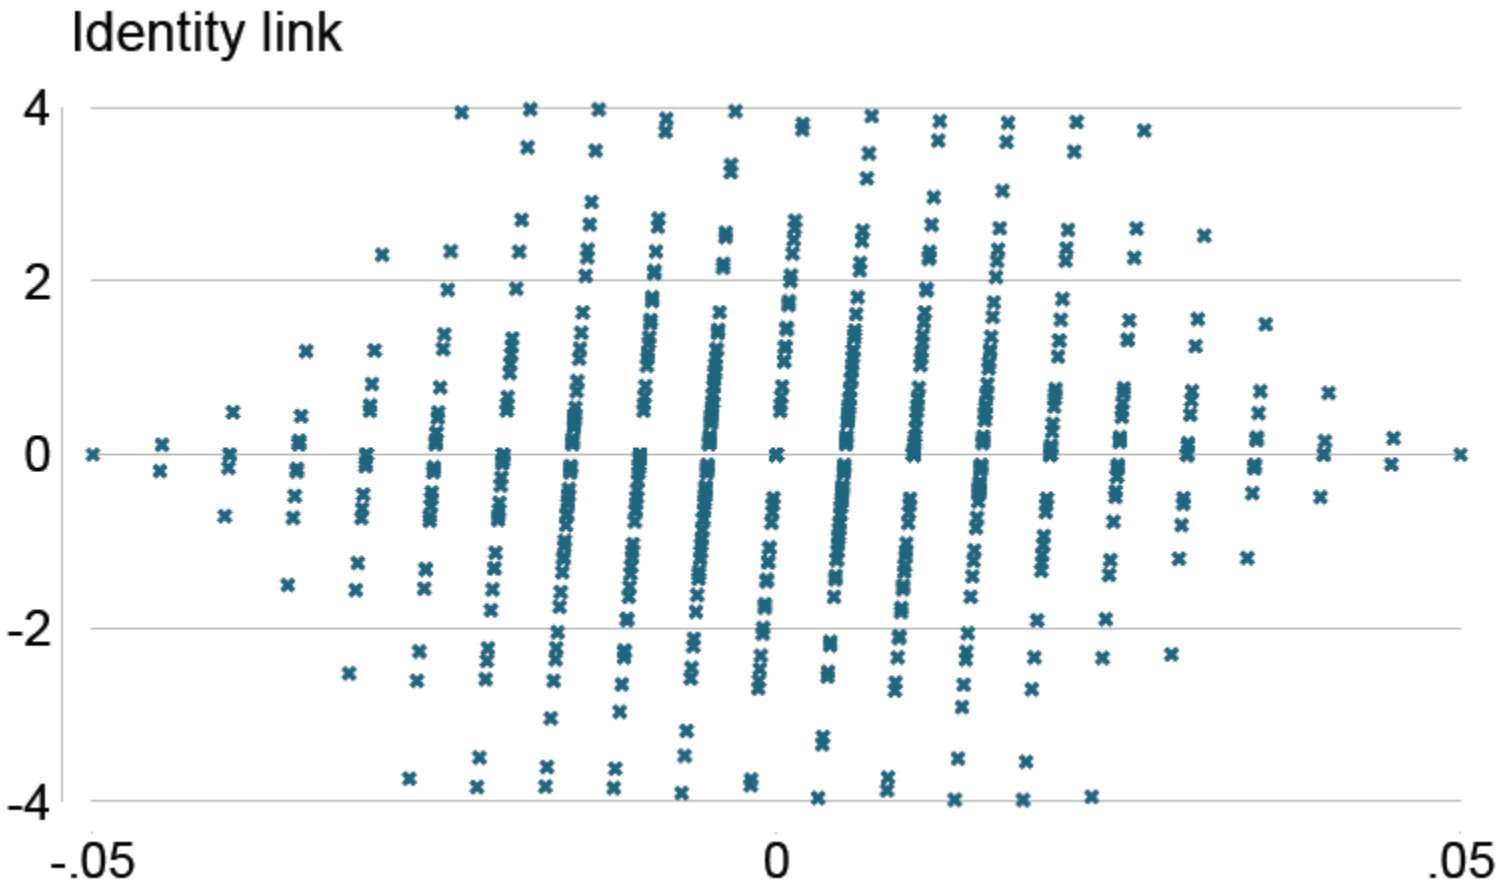

Difference (adjusted minus unadjusted)

Mean of adjusted and unadjusted

Bland–Altman plot

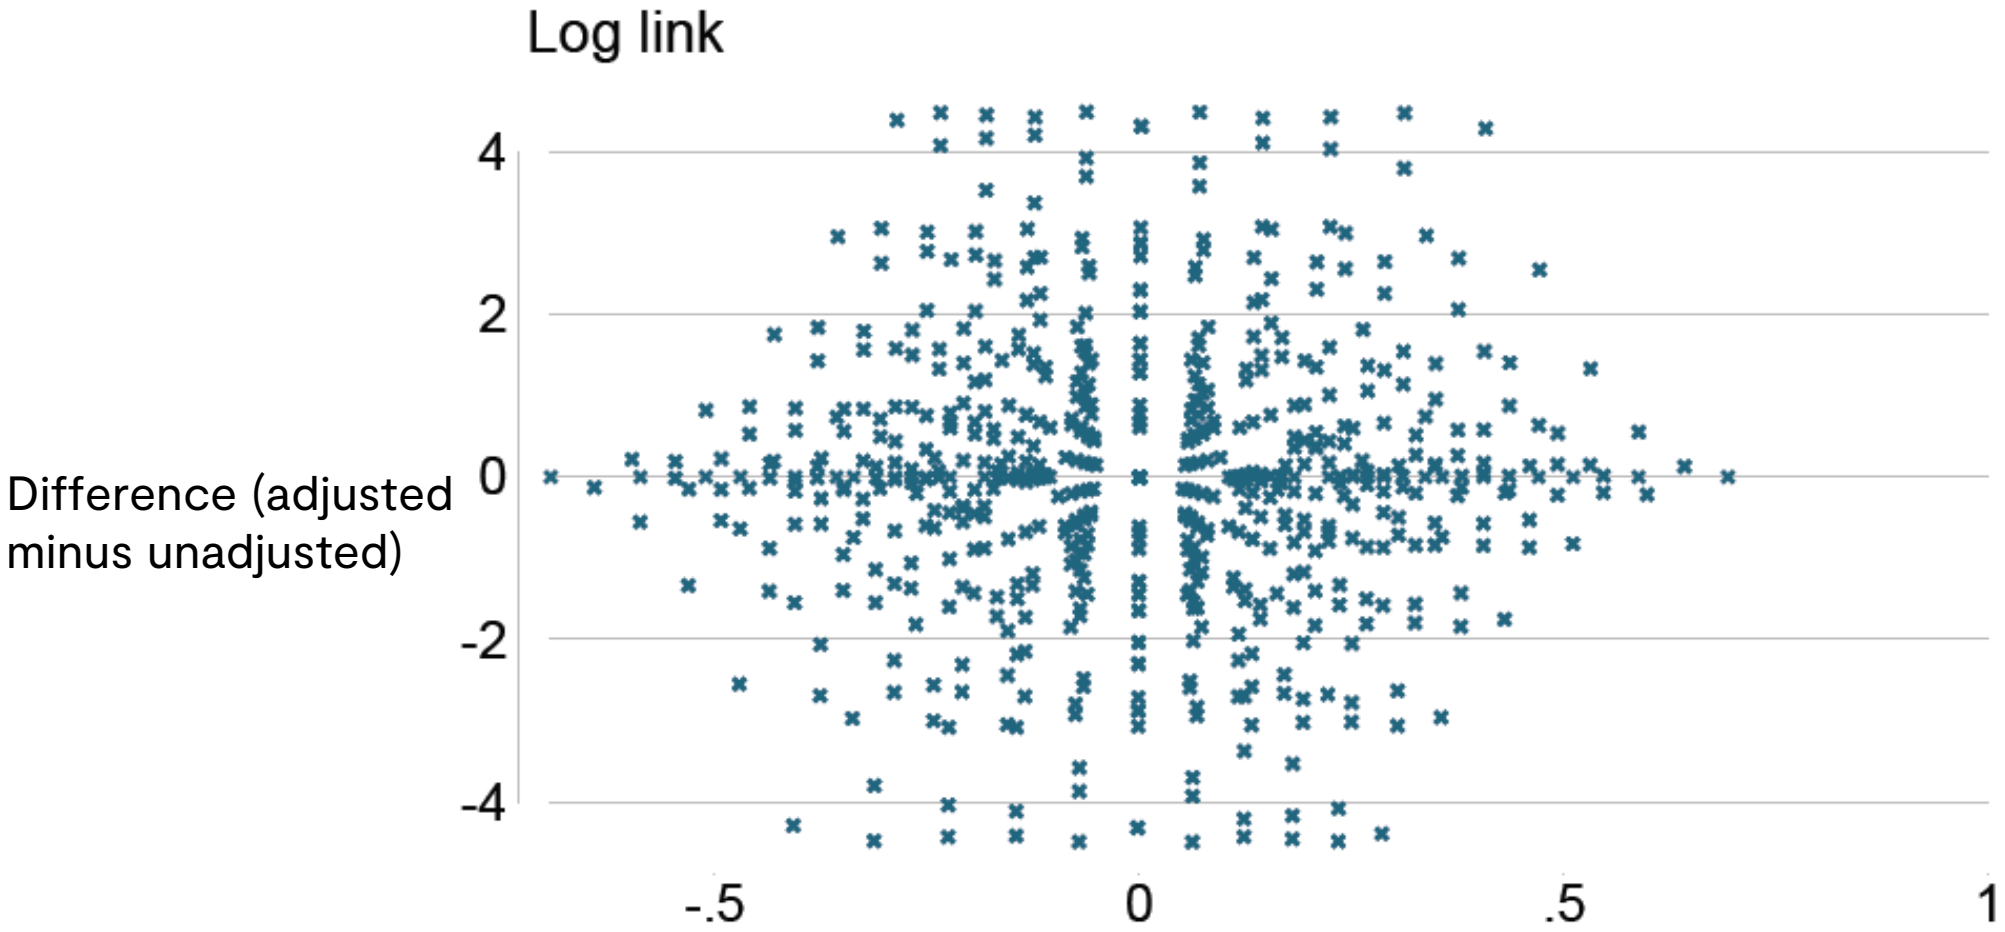

Non-canonical link functions

This is a criticism of **direct adjustment**

In the two figures where adjustment used a non-canonical link function, the unadjusted and adjusted estimates may have different signs, even though the covariate is perfectly balanced across randomised treatments.

This would be a worrying finding in the analysis of a randomised trial.

Handling missing data

A big topic that could be a whole talk!

There are some differences here, and difference ways in which we can correct bias due to missing data

Key point: when covariates predict missingness, this can change the target population of the estimand (effect among those with outcomes observed)

Handling missing data

Is this wrong-target population fixable?

In principle, yes.

Direct adjustment: multiple imputation by arm

Standardisation: predict outcomes among all, including those with missing outcomes

IPTW: Inverse probability of (non)missingness weighting!

Analyses of *GetTested*

- I am focused today on **planning**
- It is interesting to see what can happen
- Outcomes were more common than anticipated

	Anticipated (ctrl)	Observed (ctrl)
STI testing	10%	20%
Diagnosis	0.6%	1.6%

Outcome: any STI test (35% overall)

Summary measure	Adjustment method	Regression model used	Estimate (SE)
Risk difference	Direct	Identity-binomial	-
	Standardisation	Logistic (for outcome)	0.258 (0.021)
	IPTW	Logistic (for weights)	0.259 (0.021)
Risk ratio (log scale)	Direct	Log-link binomial	-
	Direct	Poisson + robust SE	0.804 (0.069)
	Standardisation	Logistic (for outcome)	0.799 (0.075)
	IPTW	Logistic (for weights)	0.802 (0.075)

Outcome: any diagnosis (1.6% overall)

Summary measure	Adjustment method	Regression model used	Estimate (SE)
Risk difference	Direct	Identity-binomial	-
	Standardisation	Logistic (for outcome)	-
	IPTW	Logistic (for weights)	0.013 (0.006)
Risk ratio (log scale)	Direct	Log-link binomial	-
	Direct	Poisson + robust SE	-
	Standardisation	Logistic (for outcome)	-
	IPTW	Logistic (for weights)	0.866 (0.414)

Summary

- When writing a SAP, we need to be precise about covariate adjustment
- People nearly always use direct adjustment
- This may be suboptimal or even a bad idea

- Start with the estimand
- There is no perfect method; each has its own advantages and drawbacks
- This may lead you to use standardisation or IPTW

Acknowledgements

- Thanks to the PSI organisers for the invite!
- Collaborators: Ian White, Fizz Williamson and Sarah Walker
- Colleagues who have contributed thoughts: Brennan Kahan, Jonathan Bartlett, Richard Riley, Leanne McCabe, Babak Choodari-Oskooei, Maarten van Smeden, Andrew Althouse.
- MRC funding: MC_UU_00004/07

Morris *et al. Trials* (2022) 23:328
https://doi.org/10.1186/s13063-022-06097-z

Trials

METHODOLOGY Open Access

Planning a method for covariate adjustment in individually randomised trials: a practical guide

Tim P. Morris[1,2*], A. Sarah Walker[1], Elizabeth J. Williamson[2] and Ian R. White[1]

Covariate adjustment in randomised trials: canonical link functions protect against model mis-specification

Ian R. White[1], Tim P Morris[1], Elizabeth Williamson[2]

[1] MRC Clinical Trials Unit at UCL, Institute of Clinical Trials and Methodology, 90 High Holborn, London WC1V 6LJ, UK.

[2] Department of Medical Statistics, LSHTM, UK

Research Article

Statistics in Medicine

Received 11 December 2012, Accepted 3 September 2013 Published online 30 September 2013 in Wiley Online Library

(wileyonlinelibrary.com) DOI: 10.1002/sim.5991

Variance reduction in randomised trials by inverse probability weighting using the propensity score

Elizabeth J. Williamson,[a,b*†] Andrew Forbes[a] and Ian R. White[c]

References (in words)

Morris, Williamson, Walker, White. Planning a method for covariate adjustment in individually randomised trials: a practical guide. *Trials*. doi:10.1186/s13063-022-06097-z

Williamson, Forbes, White. Variance reduction in randomised trials by inverse probability weighting using the propensity score. *Stat Med*. doi:10.1002/sim.5991

White, Morris, Williamson. Covariate adjustment in randomised trials: canonical link functions protect against model mis-specification. *ArXiv pre-print*. arxiv.org/abs/2107.07278