Joint PSI/EFSPI HTA SIG Webinar: Indirect treatment comparisons - choosing the right tool for the job

Antonio Remiro-Azócar, PhD

Statistics and Data Insights

Bayer, UK

Monday 26th February 2024 | 13:00 GMT

Disclaimers

- The findings and conclusions in this presentation are those of the presenter, who is responsible for its contents.
- The findings and conclusions do not necessarily represent the views of Bayer. No statement in this presentation should be construed as an official position of Bayer.

Acknowledgements

This presentation is heavily influenced by an earlier workshop presented with:

- Shannon Cope
- Harlan Campbell
- Jeroen Jansen

I gratefully acknowledge their contributions and insights

Agenda

Topic	
Part 1: Marginal versus conditional estimands, non-collapsibility	10 min
Part 2: Marginalization within the context of a single study	5 min
Part 3: Marginalization over an external target, in the context of indirect treatment comparisons	20 min
Part 4: Implications for effect modification, heterogeneity assessment	15 min
Discussion, Q&A audience	10 min

Part 1: Marginal versus conditional estimands, non-collapsibility

The HTA research question

PICO (broad question)
Population
Intervention
Comparator
Outcome

Estimand (more precise question)
Treatment
Population
Variable (endpoint)
Intercurrent event strategy
Population-level summary measure

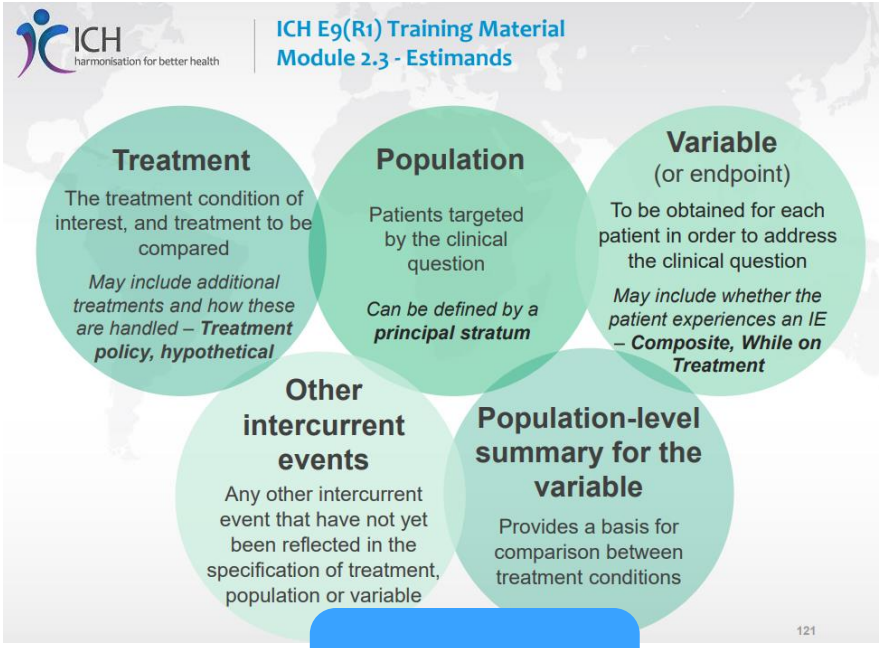

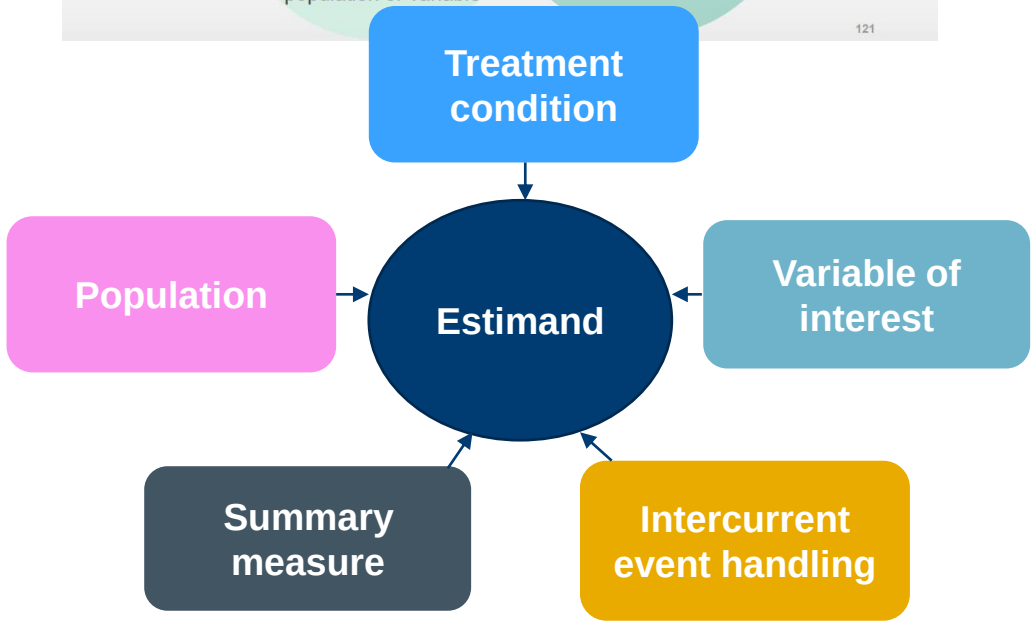

Example RCT: binary outcome

Summary effect measure: risk difference

	Stratum A: Non-smokers		Stratum B: Smokers		Population level	
	Response	No response	Response	No response	Response	No response
Treatment	90	10	50	50	140	60
Placebo	50	50	10	90	60	140
Risk difference	(90/100) – (50/100) = 0.4		(50/100) – (10/100) = 0.4		(140/200) – (60/200) = 0.4	

Smoking status is a prognostic factor but not an effect modifier on the risk difference scale!

With no confounding, perfect balance and no effect modification (treatment effect homogeneity) on the risk difference scale → ***constant stratum-level effects = population-level effect***

The risk difference is ***collapsible***

Example RCT: binary outcome

Summary effect measure: odds ratio

	Stratum A: Non-smokers		Stratum B: Smokers		Population level	
	Response	No response	Response	No response	Response	No response
Treatment	90	10	50	50	140	60
Placebo	50	50	10	90	60	140
Odds ratio	(90/10)/(50/50)=9		(50/50)/(10/90)=9		(140/60)/(60/140)≈5.4	

Smoking status is a prognostic factor but not an effect modifier on the odds ratio scale!

Despite no confounding, perfect balance and no effect modification (treatment effect homogeneity) on the odds ratio scale → ***constant stratum-level effects ≠ population-level effect***

The odds ratio is ***non-collapsible***

Marginal treatment effect

The marginal treatment effect is the average effect, at the population level, of moving an entire population from Placebo to Treatment

If all patients in the example study received Placebo, one would expect 30% response

If all patients in the example study received Treatment, one would expect 70% response

The estimated marginal odds ratio is $(0.7/0.3)/(0.3/0.7) \approx 5.4$

Smoking status does not play an explicit role in the definition of the treatment effect!

Conditional treatment effect

The conditional treatment effect is the average effect, at the subgroup level, of moving a subgroup from Placebo to Treatment.

If all smokers in the example study received Placebo, one would expect 50% response

If all smokers in the example study received Treatment, one would expect 90% response

The estimated conditional odds ratio for the smokers is (0.9/0.1)/(0.5/0.5)=9

If all non-smokers in the example study received Placebo, one would expect 10% response

If all non-smokers in the example study received Treatment, one would expect 50% response

The estimated conditional odds ratio for the non-smokers is (0.5/0.5)/(0.1/0.9)=9

Collapsibility of common summary measures

Outcome	Summary effect measure	Collapsibility
Continuous	Mean difference	**Yes, direct collapsibility**
Binary	Risk difference	**Yes, direct collapsibility**
	Risk ratio (relative risk)	**Yes**
	Odds ratio	**No**
Count	Risk ratio (relative risk)	**Yes**
	Rate ratio	**No**
	Rate difference	**No**
Time-to-event	Hazard ratio	**No**
	Restricted mean survival	**Yes**

Estimands, estimators and estimates

Estimand	**Marginal odds ratio**	**Conditional odds ratio**
Estimator	Unadjusted logistic regression $\Pr(Y = 1 \mid T = t) = \text{logit}^{-1}(\alpha_0 + \alpha_1 t)$	Logistic regression adjusted for smoking status $\Pr(Y = 1 \mid T = t, X = x) = \text{logit}^{-1}(\beta_0 + \beta_1 t + \beta_2 x)$
Estimate	$\exp(\hat{\alpha}_1) \approx 5.4$	$\exp(\hat{\beta}_1) = 9$

Part 2: Marginalization within the context of a single study

Covariate adjustment between study arms

Adjusting for prognostic factors is desirable in the estimation of marginal treatment effects:

- To correct for “chance” covariate imbalances and increase power, precision and efficiency in the analysis of randomized trials
- For valid statistical inference in trials randomized using stratified blocks (“block randomization”)
- To control for confounding in the analysis of non-randomized studies

While specific covariate adjustment methods are typically the same (computationally) with randomized and observational data, the motivation and considerations for variable selection in each scenario differ

Estimands and estimators

Marginal is not synonymous with “unadjusted”; marginal/conditional describes the estimand, adjusted/unadjusted describes the estimator and estimate

Estimand	**Marginal odds ratio**	**Conditional odds ratio**
Estimator	1. Unadjusted logistic regression 2. Inverse probability of treatment weighting (IPTW) 3. Model-based standardization (G-computation) ….	Logistic regression adjusted for smoking status (“direct adjustment”) $\Pr(Y = 1 \mid T = t, X = x) = \text{logit}^{-1}(\beta_0 + \beta_1 t + \beta_2 x)$

Inverse probability of treatment weighting (IPTW)

Creates a weighted trial sample (pseudo-population) in which covariates are balanced

1. Fit a model for the probability of treatment assignment based on participants' covariate values

$$\Pr(T = 1|X = x) = \text{logit}^{-1}(\delta_0 + \delta_1 x) = \frac{1}{1 + \exp(\delta_0 + \delta_1 x)}$$

2. Based on the fitted "treatment assignment" model, predict a propensity score for each subject

$$\hat{\pi}_i = \text{logit}^{-1}\left(\hat{\delta}_0 + \hat{\delta}_1 x_i\right)$$

3. Fit a weighted unadjusted model contrasting the treatment arms

$$\Pr(Y = 1|T = t) = \text{logit}^{-1}(\alpha_0 + \alpha_1 t),$$

$$\widehat{w}_i = \frac{t_i}{\hat{\pi}_i} + \frac{(1-t_i)}{(1-\hat{\pi}_i)},$$

with weights equal to the inverse of the estimated conditional probability of the treatment assigned in the study

Model-based standardization (G-computation)

1. Fit a covariate-adjusted model for the conditional outcome expectation

$$\Pr(Y = 1|T = t, X = x) = \text{logit}^{-1}(\beta_0 + \beta_1 t + \beta_2 x) = \frac{1}{1 + \exp(\beta_0 + \beta_1 t + \beta_2 x)}$$

2. Sum over the covariate distribution to predict marginal outcome probabilities for each treatment

$$\hat{\Pr}(Y = 1|T = 0) = \frac{1}{N}\sum_{i=1}^{N} \text{logit}^{-1}(\hat{\beta}_0 + \hat{\beta}_2 x_i)$$

"world" where all subjects assigned to T=0

$$\hat{\Pr}(Y = 1|T = 1) = \frac{1}{N}\sum_{i=1}^{N} \text{logit}^{-1}(\hat{\beta}_0 + \hat{\beta}_1 + \hat{\beta}_2 x_i)$$

"world" where all subjects assigned to T=1

3. Compute a contrast of the predictions on the odds ratio scale

$$\exp\left[\log\left(\frac{\hat{\Pr}(Y = 1|T = 1)}{1 - \hat{\Pr}(Y = 1|T = 1)}\right) - \log\left(\frac{\hat{\Pr}(Y = 1|T = 0)}{1 - \hat{\Pr}(Y = 1|T = 0)}\right)\right]$$

Covariate-adjusted estimate of the marginal odds ratio

Part 3: Marginalization in the context of indirect treatment comparisons (ITCs)

Covariate adjustment between studies (transportability): marginalizing over an external target

Background

The following setting is common in HTA submissions

- An active treatment (treatment *A*) needs to be compared against a competitor (treatment *B*)
- No head-to-head randomized trial between treatments *A* and *B*
- We have individual patient data (IPD) for study *A* but not for study *B*
- There are differences in baseline characteristics between study *A* and study *B*
- We standardize/marginalize study *A* over the covariate distribution of study *B* for a compatible ITC

ANCHORED COMPARISON

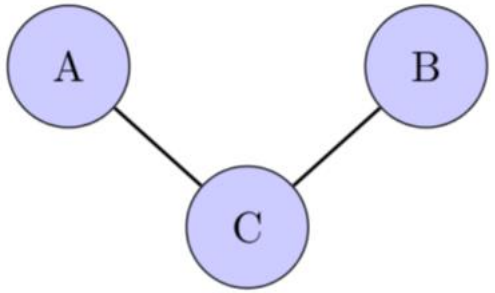

$$\hat{\Delta}_{AB}^{(B)} = \hat{\Delta}_{AC}^{(B)} - \hat{\Delta}_{BC}^{(B)}$$

UNANCHORED COMPARISON

$$\hat{\Delta}_{AB}^{(B)} = g(\hat{\mu}_A^{(B)}) - g(\hat{\mu}_B^{(B)})$$

Covariate-adjusted ITCs (2010-21)

Matching-adjusted indirect comparison (MAIC)

- Odds-weighting approach
- 164 peer-reviewed applications
- 50 NICE technology appraisals

Simulated treatment comparison (STC)

- Outcome modelling approach
- 20 peer-reviewed applications
- 9 NICE technology appraisals

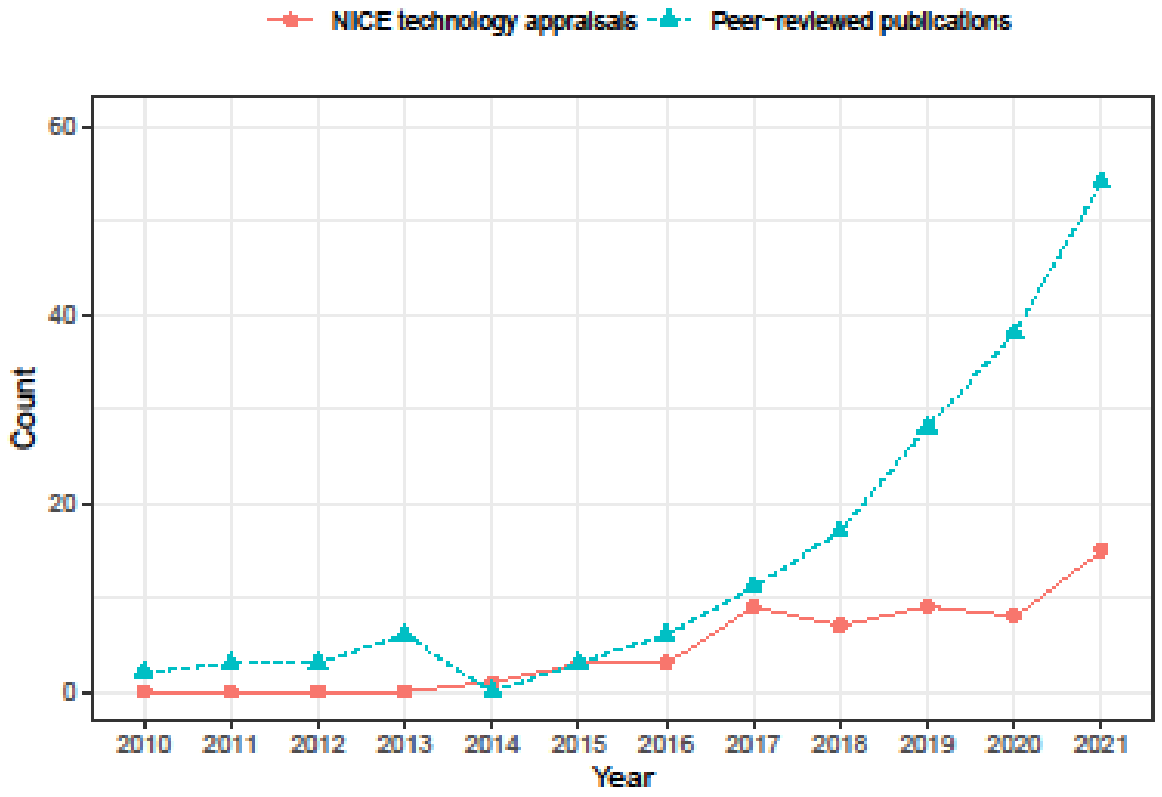

Matching-adjusted indirect comparison (MAIC)

Adjusting for between-trial differences by weighting

- Logistic model for trial assignment

$$\ln(w_i) = \ln\left[\frac{Pr(S = B \mid \boldsymbol{x}_i)}{1 - Pr(S = B \mid \boldsymbol{x}_i)}\right] = \alpha_0 + \boldsymbol{x}_i\boldsymbol{\alpha_1}$$

- Entropy balancing approach; covariate balance is viewed as a convex optimization problem

$$Q(\boldsymbol{\alpha_1}) = \sum_{i=1}^{n} \exp\left(\boldsymbol{x}_i\boldsymbol{\alpha_1}\right)$$

- The estimated weights denote the conditional odds of assignment to study *B*

$$\hat{w}_i = \exp(\boldsymbol{x}_i\hat{\boldsymbol{\alpha}}_\mathbf{1})$$

- Marginal mean outcomes and/or relative effects for study *A* treatment(s) estimated in study *B*

$$\hat{\mu}_t^{(B)} = \frac{\sum_{i=1}^{n_t} y_{i,t}\hat{w}_{i,t}}{\sum_{i=1}^{n_t} \hat{w}_{i,t}}$$

$$\hat{\Delta}_{AC}^{(B)} = g\left(\hat{\mu}_A^{(B)}\right) - g\left(\hat{\mu}_C^{(B)}\right)$$

Simulated treatment comparison (STC)

Parametric model-based standardization or G-computation

- Simulate individual-level covariates for study *B*, e.g., using a copula distribution

- Fit a multivariable regression of outcome on covariates (and treatment) to the IPD for study *A*

- $$g(\mu_i) = \beta_0 + \boldsymbol{x}_i\boldsymbol{\beta_1} + \left(\beta_t + \boldsymbol{x}_i^{(\boldsymbol{EM})}\boldsymbol{\beta_2}\right)\mathbb{1}(t_i = A)$$ (anchored case)

- Use the coefficients of the fitted model to predict hypothetical outcomes under the study *A* treatments for each simulated subject

- Consider the anchored case. We plug each treatment value into the regression fit to compute the marginal outcome means under *A* and *C*, and the corresponding relative effect.

$$\hat{\mu}_A^{(B)} = \frac{1}{m}\sum_{j=1}^{m} g^{-1}(\hat{\beta}_0 + \boldsymbol{x}_j^*\hat{\boldsymbol{\beta}}_{\mathbf{1}} + \hat{\beta}_t + \boldsymbol{x}_j^{*(\boldsymbol{EM})}\hat{\boldsymbol{\beta}}_{\mathbf{2}}) \qquad \hat{\mu}_C^{(B)} = \frac{1}{m}\sum_{j=1}^{m} g^{-1}(\hat{\beta}_0 + \boldsymbol{x}_j^*\hat{\boldsymbol{\beta}}_{\mathbf{1}})$$

$$\hat{\Delta}_{AC}^{(B)} = g\left(\hat{\mu}_A^{(B)}\right) - g\left(\hat{\mu}_C^{(B)}\right)$$

- A Bayesian implementation is also feasible; good for probabilistic sensitivity analysis

Simulated treatment comparison (STC)

Adjusting for between-trial differences by outcome modelling

- The version from NICE DSU TSD 18 targets a conditional as opposed to a marginal effect
- This leads to bias for non-collapsible effect measures, e.g. (log) hazard and odds ratios

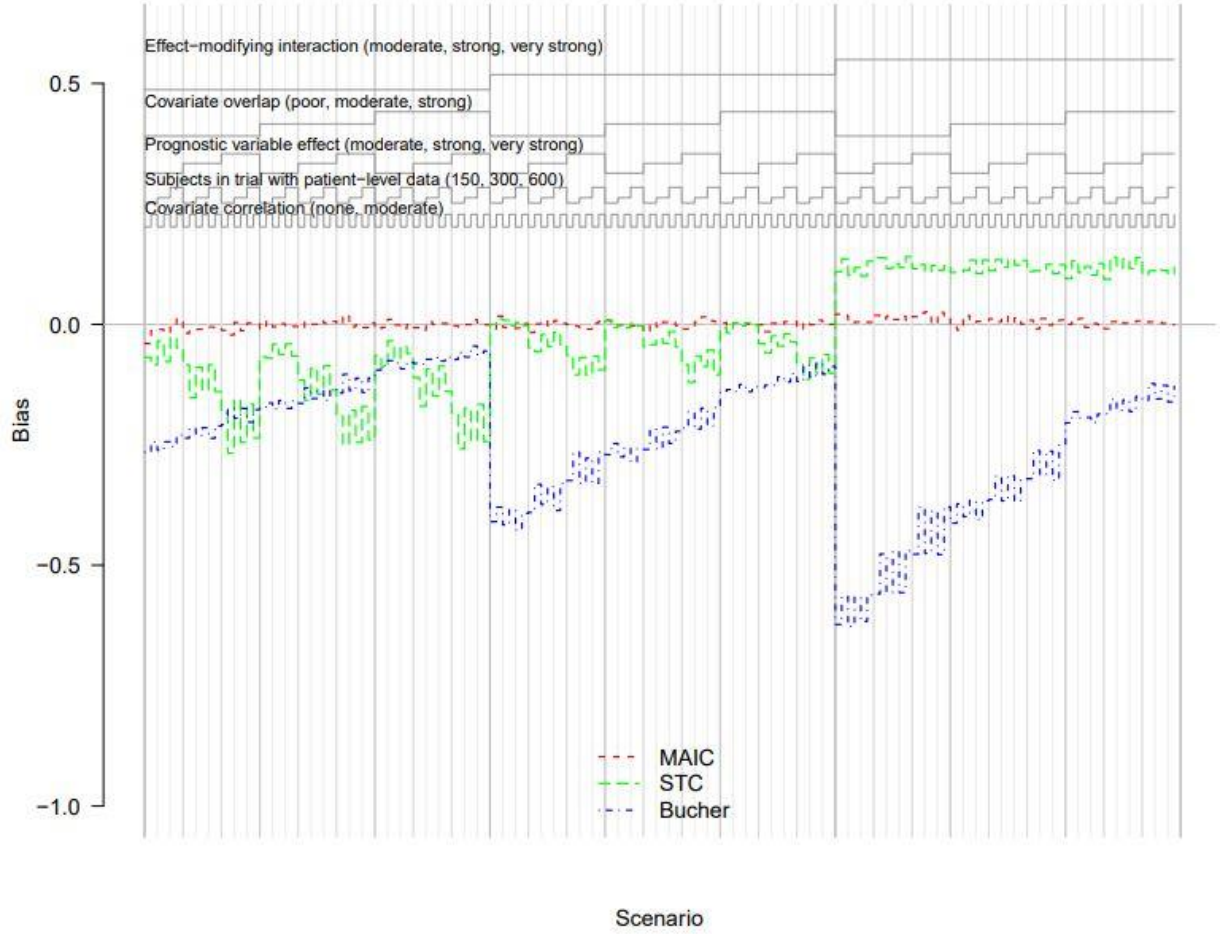

Remiro-Azócar, A., Heath, A. and Baio, G., 2021. Methods for population adjustment with limited access to individual patient data: A review and simulation study. Research synthesis methods, 12(6), pp.750-775.

Statistical performance

Outcome modelling is perceived to perform better than weighting

> Furthermore, MAICs perform poorly in simulation studies, and in some scenarios perform worse than standard NMA with no population adjustment (Phillippo, 2019).

> trial. Further details can be found in the key references, and a critique in TSD-18 (Phillippo et al., 2016). The same criticisms of MAIC apply to STC (see section 1.2.1), but with the exception that STC performs better in simulation studies than MAIC for the 2 study scenario (Phillippo, 2019).

> A recent simulation study (Phillippo, 2019) shows that ML-NMR performs similarly to STC in the 2-study scenario when the target population of interest is the population in the trial with aggregate data. However, it performs better than STC when the target population differs from that of the trial with aggregate data.

- Is MAIC biased in study *B*? ✗ If the target estimand is a conditional treatment effect, there will be bias because MAIC targets a marginal treatment effect. MAIC is unbiased if assumptions hold.

- Is MAIC potentially unprecise, therefore inefficient? ✓ Weighting methods have poor precision when the extremity of the weights is high and the effective sample size (ESS) after weighting is small.

Weighting or outcome modelling?

NICE DSU recommendations (anchored scenario)

CHTE2020 SOURCES AND SYNTHESIS OF EVIDENCE; UPDATE TO EVIDENCE SYNTHESIS METHODS. REPORT BY THE NICE DECISION SUPPORT UNIT (2020)

Recommendations

- TSD-18 (Phillippo et al., 2016) advises on circumstances under which MAIC and STC could be used in submissions, and sets out some particulars of how they should be used and presented.
- We recommend that a new TSD is prepared to show how to use ML-NMR, along with worked examples and software code, and that the Methods Guide is revised to make it clear that MAICs should not be used under any circumstances, that STCs can be can be used for two-study scenarios, and that ML-NMR is the preferred approach for anchored comparisons. This could be developed over the next 6 months.

Statistical performance

Outcome modelling is perceived to perform better than weighting

If assumptions for the methods hold, outcome modelling is more statistically precise and efficient than weighting, particularly if overlap is poor and/or the size of study *A* is small

Simulation study

- Anchored scenario, two RCTs with $N = 10000$, 1:1 randomization
- $X_k \sim \text{Normal}(0, 1)$ for study *A* and $X_k \sim \text{Normal}(-1.4, 1)$ for study *B*, $k = 1, 2, 3$ (poor overlap)
- $P(Y = 1 \mid X_1, X_2, X_3, T) = \text{expit}\left(\beta_0 + \beta_1 X_1 + \beta_2 X_2 + \beta_3 X_3 + \beta_T \mathbb{I}(T = active)\right)$
- $\beta_0 = -1$, $\beta_1 = \beta_2 = \beta_3 = 1$, $\beta_T = 1.05$
- MAIC balances the three covariate means and the outcome model is correctly specified

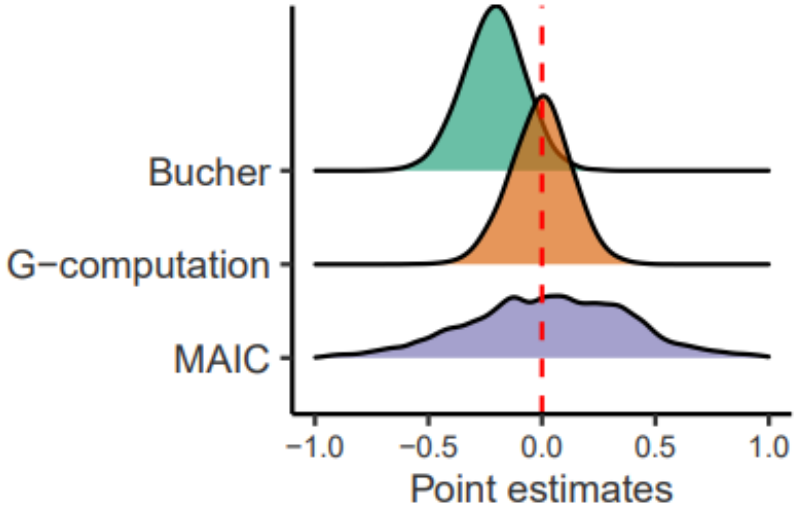

Method	Bias	Coverage	MSE
Bucher	−0.211 (0.003)	0.631 (0.011)	0.062 (0.001)
G−computation	−0.006 (0.003)	0.944 (0.005)	0.018 (0.001)
MAIC	0.034 (0.008)	0.938 (0.005)	0.137 (0.005)

Remiro-Azócar, A., 2022. Purely prognostic variables may modify marginal treatment effects for non-collapsible effect measures. arXiv preprint arXiv:2210.01757.https://doi.org/10.48550/arXiv.2210.01757

Weighting or outcome modelling?

Beyond statistical precision and efficiency under no failures of assumptions

OUTCOME MODELLING

- Relies on model-based extrapolation to improve precision and efficiency
- Susceptible to bias when extrapolating a mis-specified outcome model
- Model misspecification bias difficult to detect; an outcome model that seems approximately correct in study *A* may not fit well in extrapolated regions
- Extrapolation uncertainty not accounted for
- Can produce the treatment effect estimates that are required for HTA where there is limited overlap

WEIGHTING

- Does not extrapolate; more “honest” uncertainty quantification
- MAIC is more “bias-robust” than than the standard “inverse weighting” modelling approaches
- Model misspecification bias easier to diagnose, MAIC (entropy balancing) directly enforces balance in covariate moments
- Extreme weights explicitly manifest high uncertainty
- Feasible weighting solutions may not exist where there is limited covariate overlap, e.g. convergence failures

Standardizing with respect to "Study *B*"

Considerations about marginal estimands

- Marginal estimands for non-collapsible measures, e.g. odds and hazard ratios:
  - Depend on the full distribution of measured and unmeasured covariates, not only on covariate means
  - May change with the distribution of "purely prognostic" covariates, i.e., variables that do not induce treatment effect heterogeneity at the individual level
  - Are not identifiable with limited access to patient-level data, without making further covariate distributional assumptions

Note: *the above points also hold for measures that are not directly collapsible, e.g. log risk ratio (Part 4), in the presence of treatment effect heterogeneity at the individual level*

- In our prior simulated example:
  - The outcome-generating model only contains main effects for the covariates and lacks treatment-covariate interactions; no effect measure modification for the (log) odds ratio at the individual level
  - Nevertheless, the marginal odds ratio for active treatment versus control is 2 in study *A* and 2.46 in study *B*

Standardizing with respect to "Study *B*"

Considerations about MAIC and STC

- External validity with respect to the target population for HTA decision-making:

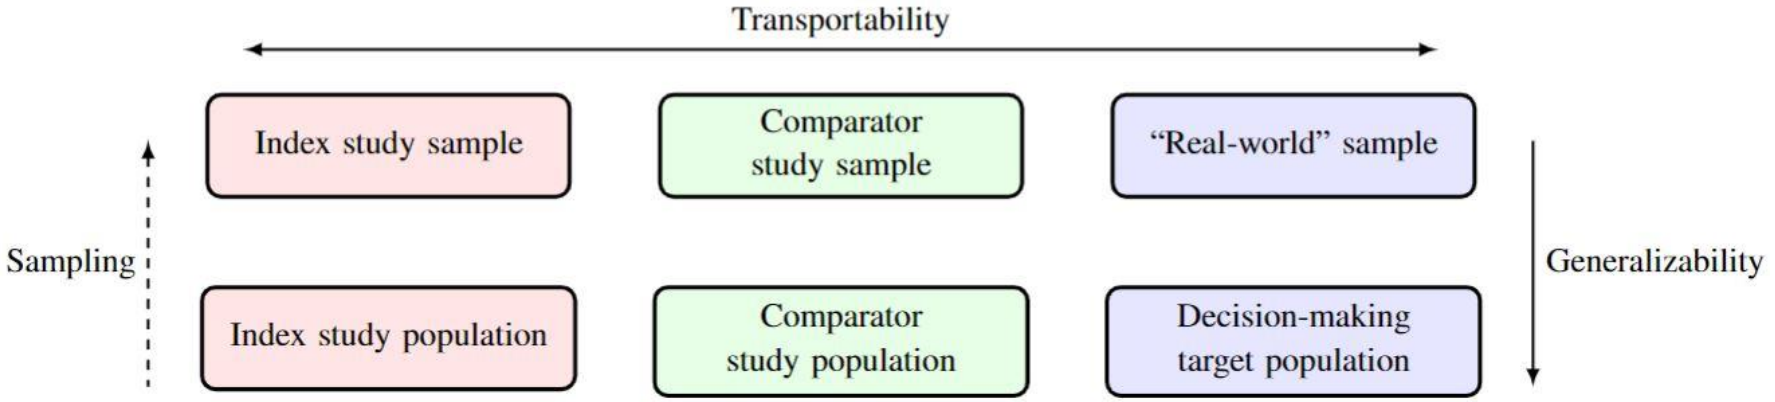

Remiro-Azócar, A., 2022. Target estimands for population-adjusted indirect comparisons. *Statistics in Medicine*, *41*(28), pp.5558-5569.

- MAIC and STC are restricted to contrast treatments in the study *B* sample
- This may not be representative of the target population of eligible patients for study *B*
- Moreover, it may differ to the target population of routine clinical practice in the jurisdiction
- A valid estimate of the marginal effect in one context is not necessarily valid in another

In the anchored scenario, multilevel network meta-regression (ML-NMR) can potentially produce marginal effect estimates in any specified target population:

- In any of the study samples included in the meta-analysis
- In an external source generated from real-world data, registries or observational studies

Multilevel network meta-regression (ML-NMR)

Define an individual-level regression model (IPD meta-regression)

$$y_{ijk} \sim \pi_{\text{Ind}}\left(\theta_{ijk}\right)$$
$$g\left(\theta_{ijk}\right) = \eta_{jk}\left(\boldsymbol{x}_{ijk}\right) = \mu_j + \boldsymbol{x}_{ijk}^T\left(\boldsymbol{\beta}_1 + \boldsymbol{\beta}_{2,k}\right) + \gamma_k$$

Average (integrate) over the target population to form the aggregate-level model

$$y_{\cdot jk} \sim \pi_{\text{Agg}}\left(\theta_{\cdot jk}\right)$$
$$\theta_{\cdot jk} = \int_X g^{-1}\left(\eta_{jk}\left(\boldsymbol{x}\right)\right) f_{jk}\left(\boldsymbol{x}\right) d\boldsymbol{x}$$

Phillippo, D.M., Dias, S., Ades, A.E., Belger, M., Brnabic, A., Schacht, A., Saure, D., Kadziola, Z. and Welton, N.J., 2020. Multilevel network meta-regression for population-adjusted treatment comparisons. Journal of the Royal Statistical Society. Series A (Statistics in Society), 183(3), p.1189.

Multilevel network meta-regression (ML-NMR)

- Parameterized on individual level-conditional treatment effects
- Conditional treatment effect at the covariate means

$$d_{ab(P)} = \int_{\boldsymbol{x}} \left(\mu_{(P)} + \boldsymbol{x}^{\mathsf{T}} \left(\boldsymbol{\beta}_1 + \boldsymbol{\beta}_{2,b}\right) + \gamma_b\right) f_{(P)}(\boldsymbol{x}) d\boldsymbol{x} - \int_{\boldsymbol{x}} \left(\mu_{(P)} + \boldsymbol{x}^{\mathsf{T}} \left(\boldsymbol{\beta}_1 + \boldsymbol{\beta}_{2,a}\right) + \gamma_a\right) f_{(P)}(\boldsymbol{x}) d\boldsymbol{x}$$
$$= \bar{\boldsymbol{x}}_{(P)}^{\mathsf{T}} \left(\boldsymbol{\beta}_{2,b} - \boldsymbol{\beta}_{2,a}\right) + \gamma_b - \gamma_a,$$

- One can obtain marginal treatment effects through integration

$$\bar{p}_{k(P)} = \int_{\boldsymbol{x}} g^{-1} \left(\mu_{(P)} + \boldsymbol{x}^{\mathsf{T}} \left(\boldsymbol{\beta}_1 + \boldsymbol{\beta}_{2,k}\right) + \gamma_k\right) f_{(P)}(\boldsymbol{x}) d\boldsymbol{x},$$

$$\Delta_{ab(P)} = g(\bar{p}_{b(P)}) - g(\bar{p}_{a(P)}).$$

Phillippo, D., Dias, S., Ades, A.E. and Welton, N.J., 2021. Target estimands for efficient decision making: Response to comments on “Assessing the performance of population adjustment methods for anchored indirect comparisons: A simulation study”. Statistics in Medicine, pp.2759-2763.

- **Open questions**
  - Application to disconnected networks (unanchored scenario)

Part 4: Implications for heterogeneity, effect modification assessment

Model-free estimands

Marginal treatment effect estimand, using the potential outcomes framework:

$$ATE = g\left(E\left(Y^1\right)\right) - g\left(E\left(Y^0\right)\right)$$

Conditional treatment effect estimand:

$$CATE = g\left(E\left(Y^1 \mid X = x\right)\right) - g\left(E\left(Y^0 \mid X = x\right)\right)$$

- The estimand definitions are formulated on an additive scale
- While the choice of the link function is often influenced by modelling preferences, the above definitions are “model-free”
- The estimands do not necessarily rely on modelling assumptions and are not necessarily encoded by the coefficients of a parametric or semi-parametric model

Model-based estimands: homogeneous model

We postulate a hypothetical **homogeneous outcome-generating model**

$$E(Y^t \mid X) = g^{-1}\left(\beta_0 + \beta_X X + \beta_T t\right)$$

- No covariate-treatment product ("interaction") term
- Covariate *X* is *purely prognostic*
- The conditional estimand, on the linear predictor scale, is the same across all subgroups or individuals regardless of their covariate value
- Treatment effect *homogeneity* because the conditional estimand is not permitted to vary with the level of *X*

$$\begin{aligned} CATE &= g\left[E\left(Y^1 \mid X = x\right)\right] - g\left[E\left(Y^0 \mid X = x\right)\right] \\ &= g\left[g^{-1}\left(\beta_0 + \beta_X x + \beta_T\right)\right] - g\left[g^{-1}\left(\beta_0 + \beta_X x\right)\right] \\ &= \beta_0 + \beta_X x + \beta_T - \left(\beta_0 + \beta_X x\right) \\ &= \beta_T \end{aligned}$$

Homogeneous model: collapsibility

1. Identity link function (linear generative model), **mean (or risk) difference**

$$
\begin{aligned}
ATE_{\text{MD}} &= E\left(Y^1\right) - E\left(Y^0\right) \\
&= E_X\left[E\left(Y^1 \mid X\right)\right] - E_X\left[E\left(Y^0 \mid X\right)\right] \\
&= E_X\left[\beta_0 + \beta_X X + \beta_T\right] - E_X\left[\beta_0 + \beta_X X\right] \\
&= \beta_0 + \beta_X E_X[X] + \beta_T - \left(\beta_0 + \beta_X E_X[X]\right) \\
&= \beta_T
\end{aligned}
$$

The treatment coefficient coincides with the model-free marginal and conditional mean difference

2. Log link function (log-linear generative model), **log risk ratio**

$$
\begin{aligned}
ATE_{\log \text{RR}} &= \ln\left[E\left(Y^1\right)\right] - \ln\left[E\left(Y^0\right)\right] \\
&= \ln\left\{E_X\left[E\left(Y^1 \mid X\right)\right]\right\} - \ln\left\{E_X\left[E\left(Y^0 \mid X\right)\right]\right\} \\
&= \ln\left\{E_X\left[\exp\left(\beta_0 + \beta_X X + \beta_T\right)\right]\right\} - \ln\left\{E_X\left[\exp\left(\beta_0 + \beta_X X\right)\right]\right\} \\
&= \ln\left\{E_X\left[\exp\left(\beta_0 + \beta_T\right)\exp\left(\beta_X X\right)\right]\right\} - \ln\left\{E_X\left[\exp\left(\beta_0\right)\exp\left(\beta_X X\right)\right]\right\} \\
&= \ln\left\{\exp\left(\beta_0 + \beta_T\right) E_X\left[\exp\left(\beta_X X\right)\right]\right\} - \ln\left\{\exp\left(\beta_0\right) E_X\left[\exp\left(\beta_X X\right)\right]\right\} \\
&= \ln\left\{\frac{\exp\left(\beta_0\right)\exp\left(\beta_T\right) E_X\left[\exp\left(\beta_X X\right)\right]}{\exp\left(\beta_0\right) E_X\left[\exp\left(\beta_X X\right)\right]}\right\} = \ln\left[\exp\left(\beta_T\right)\right] = \beta_T
\end{aligned}
$$

The treatment coefficient coincides with the model-free marginal and conditional log risk ratio

In both cases, **the marginal estimand does not depend on the purely prognostic covariate *X***

Homogeneous model: non-collapsibility

Logit link function (logistic generative model), **log odds ratio**

$$\begin{aligned} ATE_{\log \mathrm{OR}} &= \operatorname{logit}\left[E\left(Y^1\right)\right] - \operatorname{logit}\left[E\left(Y^0\right)\right] \\ &< |\beta_T| \end{aligned}$$

The treatment coefficient does not coincide with the model-free definition of the marginal log odds ratio

Because the (log) odds ratio is non-collapsible, the treatment coefficient cannot be interpreted as a population-level estimand, despite enforcing the constancy of the conditional (log) odds ratio

The marginal estimand depends on the full distribution of the purely prognostic covariate *X*

$$ATE_{\log \mathrm{OR}} = \int_x \left(\frac{\exp\left(\beta_0 + \beta_X x + \beta_T\right)}{1 + \exp\left(\beta_0 + \beta_X x + \beta_T\right)} - \frac{\exp\left(\beta_0 + \beta_X x\right)}{1 + \exp\left(\beta_0 + \beta_X x\right)} \right) P_X(dx)$$

Model-based estimands: heterogeneous model

We postulate a hypothetical **heterogeneous outcome-generating model**

$$E(Y^t \mid X) = g^{-1}\left(\beta_0 + \beta_X X + \beta_T t + \beta_{XT} X t\right)$$

- Covariate-treatment product ("interaction") term
- Covariate *X* is *prognostic* and an *effect measure modifier* (*predictive*) on the linear predictor scale
- The conditional estimand, on the linear predictor scale, varies with the covariate value
- Treatment effect *heterogeneity*: there is no longer a single conditional estimand

$$\begin{aligned} CATE &= g\left[E\left(Y^1 \mid X = x\right)\right] - g\left[E\left(Y^0 \mid X = x\right)\right] \\ &= g\left[g^{-1}\left(\beta_0 + \beta_X x + \beta_T + \beta_{XT} x\right)\right] - g\left[g^{-1}\left(\beta_0 + \beta_X x\right)\right] \\ &= \beta_0 + \beta_X x + \beta_T + \beta_{XT} x - \left(\beta_0 + \beta_X x\right) \\ &= \beta_T + \beta_{XT} x \end{aligned}$$

Heterogeneous model: direct collapsibility

A summary measure is **collapsible** if the marginal effect is always equal to a weighted average of conditional effects

For **directly collapsible** summary measures, the collapsibility weights are determined by the marginal distributions of the covariates that are conditioned on, e.g., covariate means or proportions

The mean (risk) difference is directly collapsible: for a binary/categorical covariate, the marginal mean difference is a weighted average of conditional mean differences, with weights given by covariate proportions

For example, identity link function (linear generative model):

- The marginal mean difference can be expressed simply in terms of model coefficients and the mean of *X*

$$
\begin{aligned}
ATE_{\mathrm{MD}} &= E\left(Y^1\right) - E\left(Y^0\right) \\
&= E_X\left[E\left(Y^1 \mid X\right)\right] - E_X\left[E\left(Y^0 \mid X\right)\right] \\
&= E_X\left[\beta_0 + \beta_X X + \beta_T + \beta_{XT} X\right] - E_X\left[\beta_0 + \beta_X X\right] \\
&= \beta_0 + \beta_X E_X[X] + \beta_T + \beta_{XT} E_X[X] - \left(\beta_0 + \beta_X E_X[X]\right) \\
&= \beta_T + \beta_{XT} E_X[X]
\end{aligned}
$$

Conversely, the (log) risk ratio is not directly collapsible:

- The collapsibility weights are not given by marginal covariate summary moments
- For the heterogeneous log-linear generative model, the expression for the marginal (log) risk ratio cannot simply be reduced to model coefficients and the mean of *X*
- The marginal (log) risk ratio generally depends on the full covariate distribution, including purely prognostic variables

When transporting estimates across studies…

Homogenous working model

TABLE 1 Taxonomy of model-based marginal estimands for homogeneous GLMs.

Link function	Summary effect measure	Marginal estimand
Identity	Mean difference	Does not depend on the distribution of purely prognostic covariates
Logarithmic	Log risk ratio	Does not depend on the distribution of purely prognostic covariates
Logit	Log odds ratio	Depends on the full joint distribution of purely prognostic covariates

Heterogeneous working model

TABLE 2 Taxonomy of model-based marginal estimands for heterogeneous GLMs.

Link function	Summary effect measure	Marginal estimand
Identity	Mean difference	Only depends on means of the (conditional) effect measure modifiers
Logarithmic	Log risk ratio	Depends on the full joint distribution of (conditional) effect measure modifiers and purely prognostic covariates that are associated with the former
Logit	Log odds ratio	Depends on the full joint distribution of (conditional) effect measure modifiers and purely prognostic covariates

Implications for transportability

- The constancy or transportability of marginal treatment effects depends on different covariate types for different summary measures
- The set of marginal and conditional effect measure modifiers need not coincide
- Purely prognostic variables that do not directly induce treatment effect heterogeneity at the individual level may induce treatment effect heterogeneity at the population level
- Randomization does not necessarily provide protection against between-study imbalances in purely prognostic variables
- Marginal effects can differ across settings with perfectly balanced covariate moments and identical marginal covariate distributions, but otherwise different joint covariate distributions (e.g., correlations)

Implications for indirect treatment comparisons

- Unadjusted comparisons, e.g., Bucher, are not necessarily unbiased in the absence of treatment-covariate interactions (treatment effect heterogeneity at the individual level)
- Unadjusted comparisons are not necessarily unbiased with cross-study balance in the marginal moments of influential covariates
- Covariate adjustment may be required in the absence of treatment-covariate interactions
- Covariate adjustment methods may have to account for differences in purely prognostic variables
- In the absence of subject-level data and information on the full joint covariate distribution (e.g., distributional forms and correlation structure) for the target, covariate adjustment methods may still produce bias, even if influential marginal covariate moments are perfectly balanced

Audience Q&A

Thank you!