

Principal stratum estimands in drug development – a case study

Baldur Magnusson

Presented at the Danish Society for Biopharmaceutical Statistics

29 April 2022

Outline

- Principal stratification – motivation, definition, examples
- Case study – principal stratum estimand in multiple sclerosis
- Concluding thoughts and Q&A

Motivation of principal stratification

- Notation: Z – treatment, Y – outcome, S – post-treatment variable
- Examples for S :
 - Treatment discontinuation (e.g. due to toxicity), rescue medication intake
 - Death, short-term surrogate marker
- Principal stratification is motivated by the desire to adjust for the variable S
- Naive contrast:
$$E(Y|S = s, Z = 1) - E(Y|S = s, Z = 0)$$
- Not a causal effect
 - S may be influenced by treatment → post-treatment selection bias

Principal strata – definition

Frangakis & Rubin (2002)

- Cross-classify subjects based on their joint potential outcomes ($S(0), S(1)$)
- E.g. “treatment tolerance”: $S(i) = 1$ if patient does not experience toxicity on treatment i

Principal strata in terms
of treatment tolerance

		$S(1)$	
		0	1
$S(0)$	0	Tolerate neither	Tolerate test treatment only
	1	Tolerate control only	Tolerate both

- Estimate treatment effect *within* a particular principal stratum (or strata) of interest

Examples of principal stratum estimands

- Treatment effect in
- ... in patients who would not experience toxicity regardless of treatment assignment
 $E(Y(1) - Y(0) | S(0) = S(1) = 1)$
- ... in patients who would not experience toxicity if assigned to the treatment arm
 $E(Y(1) - Y(0) | S(1) = 1)$

		$S(1)$	
		0	1
$S(0)$	0	Tolerate neither	Tolerate test treatment only
	1	Tolerate control only	Tolerate both

What's the catch?

- Suppose we are interested in the treatment effect in patients who can tolerate the active treatment
- Want to estimate $E(Y(1) | S(1) = 1) - E(Y(0) | S(1) = 1)$
- Easy to estimate $E(Y(1) | S(1) = 1)$ since we observe this on active treatment
 - Use whatever statistical model is suitable for the data at hand, e.g. regression
- Hard to estimate $E(Y(0) | S(1) = 1)$
 - We only observe $S(1)$ for patients assigned to active treatment
 - Need some way to estimate tolerance class membership among control patients **as if** they had been assigned to treatment – assumptions!

What's the catch?

- Conditioning on $(S(0), S(1))$ *does* yield a causal effect, but...
- From an **analysis perspective**, we do not observe stratum membership directly
 - Only ever observe one of $S(0)$ and $S(1)$
 - Need (untestable) identifying assumptions to link estimand & data
- From a **decision-making perspective**, how does it help to estimate a treatment effect in a population that cannot be identified at baseline?

Case study: application in the Mayzent program

DOI: 10.1002/sim.8333

RESEARCH ARTICLE


WILEY Statistics
in Medicine

Bayesian inference for a principal stratum estimand to assess the treatment effect in a subgroup characterized by postrandomization event occurrence

Baldur P. Magnusson¹ | Heinz Schmidli¹ | Nicolas Rouyre¹ |
Daniel O. Scharfstein²

<https://doi.org/10.1002/sim.8333>

Multiple sclerosis



Treatment focus:

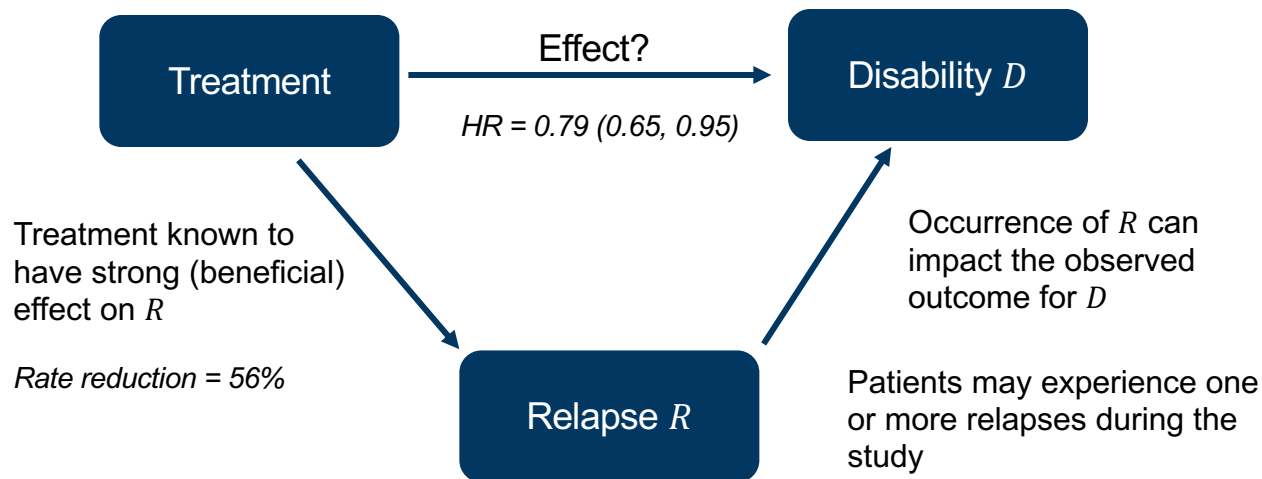
- In RRMS – reduce relapse frequency
- In SPMS – delay irreversible progression of disability

The EXPAND trial*

Placebo controlled phase 3 study in SPMS patients

*Kappos et al. (2018)


Primary question: What is the effect of Mayzent on disability D ?



Question:

What is the treatment effect among patients for whom relapses would be absent during study?

Potential estimands



- Effect in population of patients without pre-study R
 - Only useful if pre-study R is predictive of on-study R
 - Does not acknowledge the treatment effect on R
- Effect in the population of patients without on-study occurrence of R
 - Conditions on a post-randomization outcome affected by treatment
 - Estimate of treatment effect on D would not have a causal interpretation
- Effect in a world where R would not occur, hypothetical strategy in ICH E9(R1)
 - R cannot be intervened on, so may not be realistic
- Effect in the subgroup of patients who would not relapse regardless of treatment assignment

Relapse principal strata

- In a time-interval from randomization to t^* , define
 - $R(z)$ = relapse indicator in $[0, t^*]$ under treatment $z \in \{0,1\}$
 - That is, $R(z) = 1$ for patients who experience R if assigned to z
- Cross-classify patients based on $(R(0), R(1))$


		$R(1)$	
		0	1
$R(0)$	0	Non-relapser	Harmed
	1	Benefiter	Definite-relapser

- Stratum membership **not directly observable**
- Observe outcome on **actual treatment received**
- E.g. active arm patient ($z = 1$) with $R = 0$ could be either non-relapser or benefiter

Estimand of interest

- Define $D(z)$ = disability indicator in $[0, t^*]$ under treatment $z \in \{0,1\}$
- Interested in the difference in proportions of D in the **non-relapser** principal stratum

		$R(1)$	
		0	1
$R(0)$	0	Non-relapser	Harmed
	1	Benefiter	Definite-relapser





Principal stratum causal effect

$$\frac{P[D(1) = 1 \mid R(1) = 0, R(0) = 0]}{P[D(0) = 1 \mid R(1) = 0, R(0) = 0]}$$

$$= \frac{P[D(1) = 1 \mid \text{Non-relapser}]}{P[D(0) = 1 \mid \text{Non-relapser}]}$$

Identifying the estimand

Assumptions

- In practice, only **observe the margins** from this table
- Need **identifying assumptions** in order to link estimand to 'observables'
- For example: **monotonicity assumption** = no harmed patients
- Identification: (see backups for further detail)
 - Monotonicity allows estimation of estimand denominator $\frac{P[D(1) = 1 \mid \text{Non-relapser}]}{P[D(0) = 1 \mid \text{Non-relapser}]}$ 
 - Need further assumptions to identify estimand numerator 
- We proceed in the Bayesian framework, encoding assumptions through prior distributions

		R(1)		
		0	1	Sum
R(0)	0	??		✓
	1	??	??	✓
	Sum	✓	✓	

Modeling and inference

		$R(1)$	
		0	1
$R(0)$	0	Non-relapser (NR)	Harmed (H)
	1	Benefiter (B)	Definite-relapser (DR)

- Parameters:
 - G a random variable indicating principal stratum membership ($G \in \{NR, DR, B, H\}$)
 - $\pi_g = P(G = g)$
 - $\theta_g(z) = P(D(z) = 1|G = g)$
- Harmed stratum included in model specification
 - Monotonicity encoded through a **strongly informative prior**
 - Facilitates sensitivity analyses
- Estimand of interest expressed mathematically as

$$\theta_{NR}(1)/\theta_{NR}(0)$$

Bayesian model

- ω = vector of all model parameters
- In the Bayesian framework, we draw inference on ω by sampling from the **posterior distribution** $p(\omega|D, R, Z)$
- Requires a **data-generating model** $p(D, R|Z, \omega)$ and a **prior distribution** $p(\omega)$
- By Bayes' rule:


$$\begin{aligned} p(\omega|D, R, Z) &\propto p(D, R|Z, \omega) \cdot p(\omega) \\ &= \underbrace{p(D|R, Z, \omega)}_{\text{disability model}} \cdot \underbrace{p(R|Z, \omega)}_{\text{relapse model}} \cdot \underbrace{p(\omega)}_{\text{prior}} \end{aligned}$$

Prior distributions

$$\pi_g = \text{softmax}(\alpha_g) = \frac{\exp(\alpha_g)}{\sum_k \exp(\alpha_k)}$$

- Principal strata priors:
 - Work on log-odds scale with parameters α_g
 - Back-transformation ensures π_g sum to 1
 - Weakly informative priors for $\alpha_{NR}, \alpha_{DR}, \alpha_B$: independent $N(0,1)$
- Monotonicity priors:
 - Enforce **strong/weak/no monotonicity** by varying the location and scale of the α_H prior
- Disability and treatment priors:
 - $\text{logit}(\theta_g(0)) \sim N(\log(0.3), 1)$ – reflecting **expected disability rate** on placebo
 - $\text{logit}(\theta_g(1)) = \text{logit}(\theta_g(0)) + \Delta_g$, with $\Delta_g \sim N(0,1)$

Prior probability of belonging to each principal stratum



Relapse model

$$p(R|Z, \omega)$$

- Specifies the probability of relapse in $[0, t^*]$
- Underlying logic:
 - If assigned to placebo, a relapse occurs if patient is a “benefiter” or a “definite-relapser”
 - If assigned to active, a relapse occurs if patient is “harmed” or a “definite-relapser”
- Mathematical encoding:

$$p(R|Z, \omega) = 1(Z = 0) \cdot \text{Bernoulli}(\pi_B + \pi_{DR}) \\ + 1(Z = 1) \cdot \text{Bernoulli}(\pi_{DR} + \pi_H)$$

		$R(1)$	
		0	1
$R(0)$	0	Non-relapser	Harmed
	1	Benefiter	Definite-relapser

Disability model

$$p(D|R, Z, \omega)$$

- Each combination of R and Z implies two possible principal strata
 - E.g. $(R, Z) = (0, 1)$
- Mixture distribution for D when $(R, Z) = (0, 1)$:

$$\frac{\pi_{NR}}{\pi_{NR} + \pi_B} \cdot \text{Bernoulli}(\theta_{NR}(1)) + \frac{\pi_B}{\pi_{NR} + \pi_B} \cdot \text{Bernoulli}(\theta_B(1))$$


- Similarly, for $(R, Z) = (0, 0) \Rightarrow R(0) = 0$:

$$\frac{\pi_{NR}}{\pi_{NR} + \pi_H} \cdot \text{Bernoulli}(\theta_{NR}(0)) + \frac{\pi_H}{\pi_{NR} + \pi_H} \cdot \text{Bernoulli}(\theta_H(0))$$

		$R(1)$	
		0	1
$R(0)$	0	Non-relapser	Harmed
	1	Benefiter	Definite-relapser


Results

Posterior probability of belonging to each principal stratum



- Non-relapsers substantial majority
- Harmed population very small

Posterior risk ratio in non-relapser stratum



- Positive effect, though not “statistically significant”
- Wide credible intervals; effect <1 with ~80% confidence

- Not sensitive to monotonicity assumption
- Results similar for months 18 and 24 (not shown)

Conclusion

- Estimand approach to this problem was received positively by HAs
- Comments and questions focused on medical rationale, assumptions, and modeling decisions
- Work was helpful to:
 - Better understand efficacy in non-relapsing patients, and
 - To illustrate the potential utility of principal stratum estimand
- Mayzent ultimately approved for active SPMS
- Bayesian framework is appealing in this setting:
 - Straightforward to model principal strata proportions
 - Use of mixture distribution to handle lack of identifiability in the active arm
 - Structural assumptions (e.g. monotonicity) can be encoded using informative priors

Some final thoughts

- Causal inference
 - Relevant in RCTs!
 - Many pharmaceutical statistics traditions aligned with causal thinking (not all...)
 - Causal inference notation makes assumptions transparent
- Principal stratum criticized: Strong assumptions, but also relevance
 - Assumptions: Need substantive input & sensitivity analyses
 - Relevance: Clinical examples; inclusion in ICH E9(R1)
- Questions regarding post-baseline variables very subtle
 - Need to be wary of developing a “default approach” attitude
 - What is *really* the question?

References

- Bornkamp, Björn, et al. "Principal stratum strategy: Potential role in drug development." *Pharmaceutical Statistics* (2021).
- Frangakis, C. E. and Rubin, D. (2002) Principal stratification in causal inference. *Biometrics* 58: 21-29.
- Magnusson, B., Schmidli, H., Rouyrre, N. and D. Scharfstein (2019). Bayesian inference for a principal stratum estimand to assess the treatment effect in a subgroup characterized by post-randomization events. *Statistics in Medicine* 38 (23), 4761-4771. <https://doi.org/10.1002/sim.8333>

Q & A



Backup

ICH E9(R1) guideline: Intercurrent Events (IE)



INTERNATIONAL COUNCIL FOR HARMONISATION OF TECHNICAL
REQUIREMENTS FOR PHARMACEUTICALS FOR HUMAN USE

**ADDENDUM ON ESTIMANDS AND SENSITIVITY
ANALYSIS IN CLINICAL TRIALS
TO THE GUIDELINE ON STATISTICAL PRINCIPLES FOR
CLINICAL TRIALS**

E9(R1)

“... Events occurring after treatment initiation that affect either the interpretation or the existence of the measurements associated with the clinical question of interest. ...”

ICH E9(R1): Intercurrent event strategies to define estimand

- Treatment policy
 - Effect regardless of IE → IE becomes part of „treatment attribute“
- Hypothetical
 - Effect in hypothetical scenario where IE would not occur
- Composite
 - Effect on a composite variable, where IE is part of the variable
- While-on-treatment
 - Effect up to IE is considered of interest (modifies variable, i.e. observation time per patient)

ICH E9(R1): Principal Stratum Strategy

*“... The target population might be taken to be the “principal stratum” [...] in which an **intercurrent event would occur**. Alternatively, the target population might be taken to be the principal stratum in which an intercurrent event **would not occur**. The clinical question of interest relates to the treatment effect only within the principal stratum. ...”*

Example

“... a toxicity might prevent some patients from continuing the test treatment, but it would be desired to know the treatment effect among patients who are able to tolerate the test treatment. ...”

Some possible assumptions for estimation of principal stratum effects

- No assumption
 - Bounds for effect: Best/worst case scenario on correct group of control patients
- Weak assumptions
 - Bayesian model with weakly informative prior information (more later)
- Monotonicity (extrapolate S from control arm to treatment arm)
 - Patients who don't tolerate control would also not tolerate treatment + addit. assump.
- Baseline covariates (principal ignorability)
 - All patient characteristics predictive of outcome on control treatment & tolerability on test treatment → Match, adjust, weight control arm patients to find „right control group“
- Unverifiable assumptions → **Scientific understanding & Sensitivity Analyses**

Identifying the estimand

Assumptions

- In practice, only **observe the margins** from this table
- Need **identifying assumptions** in order to link estimand to 'observables'
- **Monotonicity assumption:**
There are no harmed patients
 - A patient not experiencing R on placebo **will not** experience R on active
 - That is, $R(0) = 0 \Rightarrow R(1) = 0$
 - A patient experiencing R on active **will** experience R on placebo
 - That is, $R(1) = 1 \Rightarrow R(0) = 1$

		$R(1)$		
		0	1	Sum
$R(0)$	0	??		✓
	1	??	??	✓
	Sum	✓	✓	

Identifying the estimand

Principal strata proportions

- Monotonicity allows **some patients to be classified**
 - Placebo patient with $R(0) = 0$ must be a non-relapser
 - Treated patient with $R(1) = 1$ must be a definite-relapser
- Some patients remain not classifiable
 - A treated patient who does not experience a relapse, i.e. $R(1) = 0$, could be a non-relapser or a benefiter
- We can now **estimate** the strata proportions
 - $P[\text{Definite-relapser}] = P[R = 1 | Z = 1]$
 - $P[\text{Non-relapser}] = P[R = 0 | Z = 0]$
 - $P[\text{Benefiter}] = 1 - P[\text{Definite-relapser}] - P[\text{Non-relapser}]$

		$R(1)$	
		0	1
$R(0)$	0	Non-relapser	
	1	Benefiter	Definite-relapser

Identifying the estimand

- Estimand of interest:

$$\frac{P[D(1) = 1 \mid \text{Non-relapser}]}{P[D(0) = 1 \mid \text{Non-relapser}]}$$

✗ ✓

- Randomization and monotonicity allow us to identify the denominator:

$$P[D(0) = 1 \mid \text{Non-relapser}] = P[D = 1 \mid Z = 0, R = 0]$$

- Because $R(1) = 0$ could imply non-relapser or benefiter, the numerator is not identifiable
- However, bounds on the numerator can be derived leading to a range of feasible values for the estimand

Identifying the estimand

- Using the law of total probability and **without further assumptions**:

$$P(D(1) = 1 | NR) = \underbrace{\frac{P(D(1) = 1 | NR \text{ or } B)}{P(NR | NR \text{ or } B)}}_{\text{Intercept}} - \underbrace{\frac{P(B | NR \text{ or } B)}{P(NR | NR \text{ or } B)}}_{\text{Slope}} P(D(1) = 1 | B).$$


- Intercept and slope can be calculated from the data
 - $P[D(1) = 1 | \text{Non-relapser or Benefiter}] = P[D = 1 | Z = 1, R = 0]$
 - $P[NR | NR \text{ or } B]$ is a function of strata proportions
- $P[D(1) = 1 | \text{Benefiter}]$ cannot be identified
 - Known to be between 0 and 1
 - Could make further assumptions, e.g.

$$P[D(1) = 1 | \text{Benefiter}] \leq P[D(1) = 1 | \text{Definite-relapser}]$$

Identifying the estimand


Visualizing a range of feasible values

- We can calculate the estimand of interest for a range of values of $P[D(1) = 1 \mid \text{Benefiter}]$



Full results


Posterior probability of belonging to each principal stratum for different times t^*



- Non-relapsers consistently a substantial majority
- More patients relapse as time progresses
- Harmed population very small
- Not sensitive to monotonicity

Full results, continued

Posterior risk ratios for disability in the non-relapser stratum for different times t^*



- Positive effect across all t^*
- Wide credible intervals; effect < 1 with 70-80% confidence
- Not sensitive to monotonicity

Covariates and missing data

- Due to variable follow-up time, not all patients had available (D, R) data in the time interval of interest
- Many ways to handle this – not central to the methodology
- Missing-at-random (MAR) type approach:
- Assume there exist covariates X that allow us to use available data to estimate the population parameters

Extending the model

- Use two binary baseline covariates
 - High/low measure on disability scale (EDSS)
 - Relapses within 2 years prior to study (yes/no)
- This defines four covariate strata
- Parameters indexed by x
 - $\pi_{g,x} = P(G = g|X = x)$
 - $\theta_{g,x}(z) = P(D(z) = 1|G = g, X = x)$
- Covariate-specific parameters $\pi_{g,x}$ and $\theta_{g,x}(z)$ **estimated separately** within each covariate stratum

Recovering the marginal quantities

$$\pi_g = \sum_x \underline{\pi_{g,x}} \cdot P(X = x)$$
$$\theta_g(z) = \frac{1}{\pi_g} \sum_x \underline{\theta_{g,x}(z)} \cdot \underline{\pi_{g,x}} \cdot P(X = x)$$

- Estimated from conditional models with available data
- $P(X = x)$ calculated from the empirical distribution of X

More on the conditional models

- Missingness indicator: $M = 1$ if (D, R) not observed
- Assumption: $M \perp\!\!\!\perp (D, G) \mid X, Z$
- $p(R \mid Z, X, M = 0, \omega) = (1 - Z) \cdot \text{Bernoulli}(\pi_{B,X} + \pi_{D,X}) + Z \cdot \text{Bernoulli}(\pi_{D,X} + \pi_{H,X})$
- $p(D \mid R, Z, X, M = 0, \omega)$ Bernoulli mixture, e.g. with $(R, Z) = (0, 0)$:
$$\frac{\pi_{NR,X}}{\pi_{NR,X} + \pi_{H,X}} \text{Bernoulli}(\theta_{NR,X}(0)) + \frac{\pi_{H,X}}{\pi_{NR,X} + \pi_{H,X}} \text{Bernoulli}(\theta_{H,X}(0))$$