When do analyses of randomised trials estimate causal effects?

Jonathan Bartlett
www.thestatsgeek.com
www.missingdata.org.uk

Department of Mathematical Sciences
University of Bath, UK

ISCB, 29th August 2018

Acknowledgement

Thanks to Rhian Daniel and Stijn Vansteelandt for helpful discussions.

Aalen et al 2015

Binary outcomes

Survival outcomes and the Cox model

Conclusions

Outline

Motivation - Aalen et al 2015 [4]

Lifetime Data Anal (2015) 21:579–593
DOI 10.1007/s10985-015-9335-y

CrossMark

Does Cox analysis of a randomized survival study yield a causal treatment effect?

Odd O. Aalen[1] · Richard J. Cook[2] · Kjetil Røysland[1]

Aalen et al 2015

"Despite the fact that treatment assignment is randomized, the hazard ratio is not a quantity which admits a causal interpretation in the case of unmodelled heterogeneity."

"This makes it unclear what the hazard ratio computed for a randomized survival study really means. Note, that this has nothing to do with the fit of the Cox model. The model may fit perfectly in the marginal case with X as the only covariate, but the present problem remains."

Aalen et al 's critique

They give two perspectives on the problem:

- Randomisation ensures balance (on average) of baseline risk factors (measured and unmeasured) between treatment groups.
- But at times after baseline, survivors (risk sets) in the two treatment groups in general have different distributions of baseline risk factors.
- Cox's partial likelihood consists of terms comparing these, and therefore the treatment contrast (hazard ratio) is not a valid contrast/effect.

Aalen et al 's critique

The second:

- Aalen *et al* consider individual level hazard function, which depends on all risk factor effects (frailty) and the effect of randomised treatment.
- Suppose hazards are proportional conditional on frailty and treatment. Then the coefficient of treatment in this model is true biological effect of treatment.
- If you marginalise over frailty, you lose proportional hazards (in general).
- The Cox model which adjusts only for treatment doesn't estimate this 'true' effect of treatment.

Outline

Potential outcomes and the randomised trial

Consider a population of patients.

Two binary potential outcomes (POs) Y_i^0 and Y_i^1 corresponding to control and active treatment.

We randomly allocate treatment level Z_i, $Z_i = 0$ for control, $Z_i = 1$ for active.

We observe $Y_i = Z_i Y_i^1 + (1 - Z_i) Y_i^0$.

We can calculate:

- treatment group proportions: $\bar{Y}^0$ and $\bar{Y}^1$
- risk difference (RD): $\widehat{\text{RD}} = \bar{Y}^1 - \bar{Y}^0$
- risk ratio (RR): $\widehat{\text{RR}} = \bar{Y}^1 / \bar{Y}^0$
- odds ratio (OR): $\widehat{\text{OR}} = \frac{\bar{Y}^1(1-\bar{Y}^0)}{\bar{Y}^0(1-\bar{Y}^1)}$

Stochastic vs. deterministic potential outcomes

POs could be deterministic or stochastic [1]

Stochastic POs: $Y_i^0 \sim Bernoulli(\pi_i^0)$, $Y_i^1 \sim Bernoulli(\pi_i^1)$.

'Purely' stochastic POs: $\pi_i^0 = \pi^0$, $\pi_i^1 = \pi^1$ for all i. This is implausible due to observed variation in risk between individuals.

'Partly' stochastic POs: $\pi_i^0 = g^0(X_i)$, $\pi_i^1 = g^1(X_i)$ for baseline (measured and unmeasured) variables X_i.

Deterministic POs: $\pi_i^0 = h^0(X_i) \in \{0, 1\}$, $\pi_i^1 = h^1(X_i) \in \{0, 1\}$.

Quantum physics implies (apparently) POs can't be truly deterministic. Causal inference literature tends towards deterministic, often implicitly.

Causal effect measures - stochastic POs

Stochastic POs: $Y_i^0 \sim Bernoulli(\pi^0)$, $Y_i^1 \sim Bernoulli(\pi^1)$.

$Y_i^1 - Y_i^0$, is itself stochastic, and can't be estimated.

$\bar{Y}^0$ and $\bar{Y}^1$ estimate π^0 and π^1, and $\widehat{\text{RD}}$ estimates $\pi^1 - \pi^0$.

$\pi^1 - \pi^0$ is the common individual level causal RD.

RR and OR can be interpreted as common individual level causal effects.

But we have said purely stochastic POs are implausible!

Causal effect measures - partly stochastic POs

'Partly' stochastic POs: $\pi_i^0 = g^0(X_i)$, $\pi_i^1 = g^1(X_i)$ for baseline (measured and unmeasured) variables X_i.

$Y_i^1 - Y_i^0$, is again stochastic.

$\pi_i^1 - \pi_i^0$ now varies across individuals, in general.

$\widehat{\text{RD}}$ estimates $E_i(\pi_i^1) - E_i(\pi_i^0) = E(\pi_i^1 - \pi_i^0)$.

This can be interpreted as a population level causal RD.

$\widehat{\text{RD}}$ can only be interpreted as an individual level RD if $\pi_i^1 - \pi_i^0$ did not vary across i.

Causal effect measures - partly stochastic POs

$\widehat{\text{RR}}$ estimates $\frac{E_i(\pi_i^1)}{E_i(\pi_i^0)}$.

This is a population level RR.

Again only if π_i^1/π_i^0 were identical across i could it be interpreted as an individual level RR.

If individual level RD is common, RR cannot be, and vice-versa.

Causal effect measures - partly stochastic POs

$\widehat{\text{OR}}$ estimates population level OR: $\frac{E_i(\pi_i^1)(1-E_i(\pi_i^0))}{E_i(\pi_i^0)(1-E_i(\pi_i^1))}$.

Due to non-collapsibility of the OR, the population level OR does not equal the individual level OR even when the latter is identical across individuals.

If the individual level OR were identical across individuals, we need to condition adjust (correctly) on all prognostic factors to estimate it.

No reason to think that we will ever have all the prognostic variables measured, and if we did, that we correctly model their effects.

Causal effect measures - deterministic POs

$Y_i^1 - Y_i^0$ is 0, 1, or -1, and is now fixed for each individual, although they are not identifiable.

$\widehat{\text{RD}}$ estimates population RD $E_i(Y_i^1) - E_i(Y_i^0)$, and similarly for $\widehat{\text{RR}}$ and $\widehat{\text{OR}}$.

There is no longer an individual level RD, RR, or OR.

Adjusted/conditional (e.g. logistic regression) estimates sub-population level OR, assuming this is common across sub-populations defined by covariates.

Population vs. individual effects

Some advocate adjusting for covariates and interpreting the resulting estimate as an individual level effect (e.g. Harrell [3]).

But this relies on assuming:

- the chosen effect measure is common across individuals
- we correctly model covariate effects

Because of these issues, others are more cautious, choosing instead to target population (marginal) effects (e.g. Steingrimsson *et al* [2]).

Of course one could estimate sub-population effects, allowing for the possibility/fact(!) that these will vary across sub-populations.

Outline

Cox model under purely stochastic POs

- Y_i^0 and Y_i^1 are now the failure times under control and active treatments for individual i.
- Under purely stochastic POs, we have common individual level hazard functions $\lambda^0(t)$ and $\lambda^1(t)$.
- Cox's model then assumes

$$\lambda^1(t) = \lambda^0(t)\exp(\beta)$$

- $\exp(\beta)$ then has an interpretation as the common causal effect on the individual level hazard.
- This is probably how many people interpret the hazard ratio (HR) from such a model.
- As in the binary case though, the purely stochastic model is implausible since there is always variation in risk, so such an interpretation would not generally be justified.

Cox model under partly stochastic POs

- Here we have $\lambda_i^0(t) = g^0(X_i, t)$ and $\lambda_i^1(t) = g^1(X_i, t)$ for baseline variables X_i.
- This is the view of the world Aalen *et al* assume to hold.
- As shown by Aalen *et al* , even if $\lambda_i^1(t)/\lambda_i^0(t) = \exp(\beta)$ for all i and t, then the Cox model adjusting only for treatment does not estimate $\exp(\beta)$, due to non-collapsibility.
- So, as per Aalen *et al* , the estimated HR from a Cox model with treatment as covariate cannot in general be validly interpreted as an individual level HR.
- Recall however that the same conclusion applies to ORs for binary outcomes!

Cox model under deterministic POs

- Under deterministic POs and assuming proportional hazards conditional on treatment only, the resulting HR indicates the ratio of population level hazards.
- This tells you that if you assign treatment 1 to the population, the population hazard/rate of events at any time is $\exp(\beta)$ times what it would have been had you assigned treatment 0.
- Under marginal proportional hazards we have

$$\exp(\beta) = \frac{\log\{S^1(t)\}}{\log\{S^0(t)\}} = \frac{\log[E_i\{1(Y_i^1 > t\}]}{\log[E_i\{1(Y_i^0 > t\}]}$$

- So it also has a (nasty) interpretation in terms of a ratio of log population survival probabilities.

Outline

Conclusions

- Main analyses of RCTs can't estimate individual level causal effects, unless you make the strong assumption that the effect (on your chosen scale) is identical across individuals.
- For the Cox model, even if this were the case, the Cox model adjusting only for treatment would not estimate this common effect, due to non-collapsibility.
- But the same is true for other models which are non-collapsible, e.g. logistic regression for binary outcomes.
- Assuming proportional hazards holds conditional on treatment, HRs **DO** have a valid causal interpretation as population level HRs.
- These slides, plus comments from Aalen at `www.thestatsgeek.com`

References

[1] M Hernán and J Robins.
Causal Inference.
Boca Raton: Chapman & Hall/CRC, 2018.

[2] Steingrimsson JA, Hanley DF, and Rosenblum M.
Improving precision by adjusting for prognostic baseline variables in randomized trials with binary outcomes, without regression model assumptions.
Contemporary Clinical Trials, 54:18–24, 2017.

[3] Harrell FE Jr.
Biostatistics for Biomedical Research.
`http://hbiostat.org/doc/bbr.pdf`.
Accessed: 2018-08-24.

[4] Aalen OO, Cook RJ, and Røysland K.
Does Cox analysis of a randomized survival study yield a causal treatment effect?
Lifetime Data Analysis, 21(4):579–593, 2015.