

UCSF

UC San Francisco Previously Published Works

Title

Confidence limits for the averted infections ratio estimated via the counterfactual placebo incidence rate

Permalink

<https://escholarship.org/uc/item/6rr5v5zb>

Journal

Statistical Communications in Infectious Diseases, 13(1)

ISSN

2194-6310

Authors

Dunn, David T
Stirrup, Oliver T
Glidden, David V

Publication Date

2021

DOI

10.1515/scid-2021-0002

Peer reviewed

Article

David T. Dunn*, Oliver T. Stirrup and David V. Glidden

Confidence limits for the averted infections ratio estimated via the counterfactual placebo incidence rate

<https://doi.org/10.1515/scid-2021-0002>

Received May 12, 2021; accepted November 7, 2021; published online November 24, 2021

Abstract

Objectives: The averted infections ratio (AIR) is a novel measure for quantifying the preservation-of-effect in active-control non-inferiority clinical trials with a time-to-event outcome. In the main formulation, the AIR requires an estimate of the counterfactual placebo incidence rate. We describe two approaches for calculating confidence limits for the AIR given a point estimate of this parameter, a closed-form solution based on a Taylor series expansion (delta method) and an iterative method based on the profile-likelihood.

Methods: For each approach, exact coverage probabilities for the lower and upper confidence limits were computed over a grid of values of (1) the true value of the AIR (2) the expected number of counterfactual events (3) the effectiveness of the active-control treatment.

Results: Focussing on the lower confidence limit, which determines whether non-inferiority can be declared, the coverage achieved by the delta method is either less than or greater than the nominal coverage, depending on the true value of the AIR. In contrast, the coverage achieved by the profile-likelihood method is consistently accurate.

Conclusions: The profile-likelihood method is preferred because of better coverage properties, but the simpler delta method is valid when the experimental treatment is no less effective than the control treatment. A complementary Bayesian approach, which can be applied when the counterfactual incidence rate can be represented as a prior distribution, is also outlined.

Keywords: active-control trial; averted infections; confidence limits; non-inferiority.

Introduction

In a series of papers we have considered the analysis of active-control non-inferiority trials with a time-to-event outcome in the context of HIV prevention trials (Dunn and Glidden 2019; Dunn et al. 2018; Glidden, Stirrup, and Dunn 2020). Our key conclusion is that the standard metric used in such trials, the rate ratio comparing experimental and control arms, is misleading. We further argued that clinically meaningful inference requires estimation or specification of one of two unobserved parameters: (a) the event rate that would have been

*Corresponding author: David T. Dunn, Institute for Global Health, University College London, London, UK; and MRC Clinical Trials Unit, University College London, London, UK, E-mail: d.dunn@ucl.ac.uk. <https://orcid.org/0000-0003-1836-4446>

Oliver T. Stirrup, Institute for Global Health, University College London, London, UK

David V. Glidden, Department of Epidemiology and Biostatistics, University of California San Francisco, San Francisco, CA, USA

observed in trial subjects if they had received no treatment (counterfactual placebo arm) or (b) the effectiveness of the control arm relative to the counterfactual placebo arm. With this information, in combination with the observed incidence rates in the control and experimental arms, we can estimate a measure called the averted infections ratio (AIR). The AIR is interpreted as the proportion of events that would be averted by use of the experimental treatment compared with the control treatment. In the context of non-inferiority trials, it is a natural criterion for assessing the degree to which the experimental treatment preserves the effect of the control treatment relative to no treatment (“preservation-of-effect”) (Ghosh et al. 2011; Pigeot et al. 2003; Snapinn and Jiang 2008). Non-inferiority trials using this approach typically aim to demonstrate at least 50% preservation-of-effect, although this value is context specific and higher values may be warranted (Pigeot et al. 2003; Ghosh et al. 2011). In this paper, we consider the derivation of confidence limits for the AIR when it is estimated via the counterfactual placebo incidence.

Notation and statistical formulation

Denote the hypothetical placebo, control, and experimental arms by the subscripts P, C, and E, respectively. We observe F_C person-years follow-up in control arm and F_E person-years follow-up in experimental arm. Let X_C and X_E be the random variables denoting the number of observed events, where we assume that $X_C \sim \text{Poi}(F_C \lambda_C)$ and $X_E \sim \text{Poi}(F_E \lambda_E)$. Let λ_P represent the counterfactual placebo incidence. The averted infections ratio is defined as

$$\Psi = \frac{\lambda_P - \lambda_E}{\lambda_P - \lambda_C} \quad (1)$$

Alternatively, Ψ can be expressed in terms of the counterfactual control arm effectiveness ($\theta_C = 1 - \lambda_C/\lambda_P$) rather than λ_P :

$$\Psi = \frac{1 - \lambda_E/\lambda_C(1 - \theta_C)}{\theta_C}$$

In this formulation, Ψ is a linear function of the rate ratio and confidence limits for Ψ can be obtained by direct transformation of confidence limits for the rate ratio. As the latter problem has been extensively studied (Graham, Mengersen, and Morton 2003; Li, Tang, and Wong 2014; Price and Bonett 2000; Sahai and Khurshid 1993) we focus on formulation (1).

Inference conditional on counterfactual incidence

This section considers the derivation of confidence limits for the AIR when considering a single, pre-specified value of λ_P . This allows exploration of how the confidence limits (and point estimates) vary over a range of plausible values of λ_P , which can be highly informative (Glidden, Stirrup, and Dunn 2020).

Delta method

We first apply a log transformation to the AIR, a natural procedure for any statistic that is a ratio of two variables. From Eq. (1),

$$\log_e \Psi = \log_e(\lambda_P - \lambda_E) - \log_e(\lambda_P - \lambda_C)$$

Based on a first-order Taylor series expansion (Oehlert 1992)

$$\text{var}(\log_e \hat{\Psi}) \rightarrow \frac{\text{var}(\hat{\lambda}_E)}{(\lambda_P - \hat{\lambda}_E)^2} + \frac{\text{var}(\hat{\lambda}_C)}{(\lambda_P - \hat{\lambda}_C)^2}$$

since λ_p is regarded as fixed. Thus

$$\text{var}(\log_e \hat{\Psi}) \rightarrow \frac{\hat{\lambda}_E/F_E}{(\lambda_p - \hat{\lambda}_E)^2} + \frac{\hat{\lambda}_C/F_C}{(\lambda_p - \hat{\lambda}_C)^2} \quad (2)$$

A $(1-\alpha)$ confidence interval for Ψ is obtained from

$$\exp\left(\log_e \hat{\Psi} \pm z_{\alpha/2} \sqrt{\frac{\hat{\lambda}_E/F_E}{(\lambda_p - \hat{\lambda}_E)^2} + \frac{\hat{\lambda}_C/F_C}{(\lambda_p - \hat{\lambda}_C)^2}}\right)$$

Profile-likelihood method

The log-likelihood under a Poisson model is

$$l(\lambda_C, \lambda_E) = -F_C \lambda_C + X_C \log_e(F_C \lambda_C) - F_E \lambda_E + X_E \log_e(F_E \lambda_E) \quad (3)$$

We can express (3) in terms of Ψ via Eq. (1), noting that a nuisance parameter (either λ_C , or λ_E , or a function of λ_C and λ_E) is also involved. Denoting this arbitrary nuisance parameter by ζ , the profile-likelihood confidence region for Ψ is defined by the set of values (Cole, Chu, and Greenland 2014)

$$\left\{ \Psi: 2 \left[1(\hat{\Psi}, \hat{\zeta}) - 1(\Psi, \hat{\zeta}(\Psi)) \right] < \chi_{1-\alpha}^2 \right\} \quad (4)$$

where

$$1(\hat{\Psi}, \hat{\zeta}) = -X_C + X_C \log_e(X_C) - X_E + X_E \log_e(X_E)$$

is the unconstrained maximised log-likelihood.

An alternative approach is to parameterise the problem in terms of λ_C and λ_E rather than Ψ and ζ . We therefore maximise (3) subject to the constraint implied by Eq. (1) for a specified value Ψ^* . Re-arranging,

$$g(\lambda_C, \lambda_E) = \lambda_E - \Psi^* \lambda_C + \lambda_p (\Psi^* - 1) = 0$$

Introducing a Lagrange multiplier (β), we maximise

$$l(\lambda_C, \lambda_E) + \beta g(\lambda_C, \lambda_E) \quad (5)$$

Differentiating (5) with respect to β , λ_E , and λ_C results in a set of three non-linear equations:

$$\lambda_E - \Psi^* \lambda_C + \lambda_p (\Psi^* - 1) = 0, \quad -F_E \lambda_E + X_E + \beta \lambda_E = 0, \quad -F_C \lambda_C + X_C - \beta \Psi^* \lambda_C = 0,$$

noting that Ψ^* and λ_p are constants. Using the method of elimination,

$$\lambda_C = \frac{y + \sqrt{y^2 - 4xz}}{2x}, \quad \lambda_E = \Psi^* \lambda_C - \lambda_p (\Psi^* - 1)$$

where

$$x = \Psi^* (F_C + \Psi^* F_E), \quad y = (\Psi^* - 1) \lambda_p (F_C + \Psi^* F_E) + \Psi^* (X_C + X_E), \quad z = (\Psi^* - 1) X_C \lambda_p$$

The roots of the function implied by (4) were found using the **uniroot** function in R (version 4.02), which utilises the golden-section search procedure combined with parabolic interpolation (code in Appendix).

Unconditional inference

In addition to exploring how the AIR varies over a range of values of the counterfactual incidence, we may wish to integrate over this parameter to obtain the unconditional distribution of the AIR. Bayesian inference provides a natural framework for this problem. Here we consider the case where trial investigators are able to specify a simple prior distribution for the counterfactual incidence, although more sophisticated approaches which incorporate external information are also possible (Glidden, Stirrup, and Dunn 2020).

Assume that the prior for λ_p can be specified as a Gamma distribution based on background knowledge. For λ_E and λ_C , we use weakly informative priors $\sim \text{Gamma}(0.5, 0.001)$ – this approximates to Jeffrey’s prior (Gelman et al. 1995), and also corresponds to adding 0.5 to the observed number of events as discussed in Section 6. As the Gamma distribution is the conjugate prior for the Poisson model, the posterior distributions for λ_E and λ_C are $\text{Gamma}(X_E + 0.5, F_E + 0.001)$ and $\text{Gamma}(X_C + 0.5, F_C + 0.001)$, respectively (Gelman et al. 1995). We generate samples from the distributions of λ_p , λ_E , and λ_C , to derive the posterior distribution for the AIR using Eq. (1).

The main application of the AIR is in non-inferiority trials, where it is reasonable to assume that $\lambda_C < \lambda_p$ since the effectiveness of the control drug will already have been established. Further, the AIR is uninterpretable if $\lambda_C > \lambda_p$ as this would imply there was no yardstick against which the experimental drug could be compared (nothing to preserve). In most realistic applications it is also reasonable to assume that $\lambda_E < \lambda_p$ as the experimental drug will have been selected as having some biological activity. It is therefore problematic if the sampled values $(\lambda_p^*, \lambda_C^*, \lambda_E^*)$ satisfy either

$$\lambda_C^* > \lambda_p^* \quad \text{or} \quad \lambda_E^* > \lambda_p^* \quad (6)$$

There are three possible re-sampling strategies: (a) re-sample λ_p^* only (b) re-sample $(\lambda_p^*, \lambda_C^*)$ if $\lambda_C^* > \lambda_p^*$; re-sample $(\lambda_p^*, \lambda_E^*)$ if $\lambda_E^* > \lambda_p^*$ (c) re-sample all three values $(\lambda_p^*, \lambda_C^*, \lambda_E^*)$. The best strategy is not obvious, and all are explored in the example in Section 7.

Three arm trials with a placebo arm

Trials are occasionally designed with a placebo arm in addition to the control and experimental arms, thereby providing a direct estimate of λ_p (Ghosh et al. 2011). The Taylor series approximation (Eq. (2)) requires an additional term to reflect the uncertainty in the estimate of λ_p :

$$\text{var}(\log_e \hat{\Psi}) \approx \frac{\hat{\lambda}_E/F_E}{(\hat{\lambda}_p - \hat{\lambda}_E)^2} + \frac{\hat{\lambda}_C/F_C}{(\hat{\lambda}_p - \hat{\lambda}_C)^2} + (\hat{\lambda}_p/F_p) \left[\frac{\hat{\lambda}_C - \hat{\lambda}_E}{(\hat{\lambda}_p - \hat{\lambda}_E)(\hat{\lambda}_p - \hat{\lambda}_C)} \right]^2 \quad (7)$$

The additional term is generally much smaller than the first two terms and, in expectation, (7) tends towards (2) when $\lambda_E = \lambda_C$. This leads to a paradoxical finding, namely that the sample size of the placebo group appears to be irrelevant when this equality is assumed (as is commonly the case when designing non-inferiority trials). This paradox is explained by the fact that $\Psi = 1$ when $\lambda_E = \lambda_C$ regardless of the value of λ_p . However, the placebo group needs to be large enough in order to ensure that the estimate $\hat{\lambda}_p$ is sufficiently stable. An interesting, unresolved question is the optimal relative sample size allocation to the three arms. We further note the profile-likelihood approach (Section 3.2) could, in principle, be extended to three arm trials.

Coverage probabilities

Methods

Exact coverage probabilities for the lower and upper confidence limits (at nominal coverage probabilities of $1-\alpha$, for $\alpha = 0.025, 0.05$) were computed using the delta method and profile-likelihood method described in Section 3. For the purposes of exposition we assume $F_C = F_E = 1$, so that λ_C and λ_E can be considered as the expected number of events, and λ_p the expected number of counterfactual events, in each of the two trial arms. The following parameters were examined over a grid of values: $\Psi = 0.5(0.1)1.0$; $\lambda_p = 40(20)100$; $\theta_C = 0.6(0.1)0.9$. Exact coverage probabilities were computed by

$$\sum_{X_C=0}^{\infty} \sum_{X_E=0}^{\infty} \frac{e^{-\lambda_C} \lambda_C^{X_C}}{X_C!} \frac{e^{-\lambda_E} \lambda_E^{X_E}}{X_E!} I(X_C, X_E)$$

where $\lambda_C = \lambda_p(1 - \theta_C)$, $\lambda_E = \lambda_p(1 - \psi\theta_C)$, and $I(X_C, X_E)$ equals 1 if the lower(upper) confidence limit is less(greater) than Ψ , otherwise equals 0.

The log-likelihood is undefined when either $X_C = 0$ or $X_E = 0$. However, in contrast with the rate ratio, this is a highly informative outcome in terms of the AIR (even $X_C = 0, X_E = 0$). To avoid this problem, X_C and X_E were replaced by $X_C + 0.5$ and $X_E + 0.5$ before applying the methods of Section 3.2. For consistency, this adjustment was also applied for confidence limits determined by the delta method. The addition of 0.5 resulted in improved coverage estimates under both approaches, as has previously been reported for the rate ratio (Price and Bonett 2000).

Results

The complete set of coverage probabilities for the lower and upper confidence limits are given in the Appendix. However, the lower confidence limit is of primary interest since this is the comparator for the non-inferiority margin. Also, the upper limit of Ψ may be severely constrained for large values of θ_C . Ψ can be expressed as θ_E/θ_C , so that, for example, $\Psi \leq 1.25$ if $\theta_C = 0.8$, $\Psi \leq 1.11$ if $\theta_C = 0.9$.

Figure 1 shows coverage probabilities using the delta method for the lower one-tailed $\alpha = 0.05$ confidence limit (similar patterns were observed for $\alpha = 0.025$). Coverage is generally too low for $\Psi = 0.5-0.8$, is reasonably accurate for $\Psi = 0.9$, and is too high for $\Psi = 1.0$. This pattern is explained by a negative correlation between the empirical AIR and its estimated standard error, conditional on the true AIR (Ψ). Conditional on Ψ , coverage is higher the larger the value of the control arm effectiveness (θ_C), except for $\Psi = 1.0$ when differences are minor. As expected, actual and nominal coverage are closer the larger the value of λ_p , although convergence is slow with material discrepancies even for $\lambda_p = 100$. Coverage probabilities for the upper confidence limit were consistently and substantially too high (Appendix), particularly for lower values of Ψ .

Table 1 shows coverage probabilities for the profile-likelihood-based lower confidence limit for $\lambda_p = 40$ and $\alpha = 0.05$. Coverage was close to the nominal value of 0.95 (range 0.9468–0.9615) for all permutations of Ψ and θ_C ; as expected, correspondence was even closer at higher values of λ_p (not shown). Coverage for the profile-likelihood-based upper confidence limit were also highly accurate, in contrast to the delta method (Appendix). The results of these analyses support the routine use of profile-likelihood-based confidence limits, although the delta method is valid in a conservative sense (i.e. actual coverage exceeds nominal coverage) if the true AIR is ≥ 0.9 approximately. This is reflected in larger values for the lower confidence limit using the delta method (Appendix).

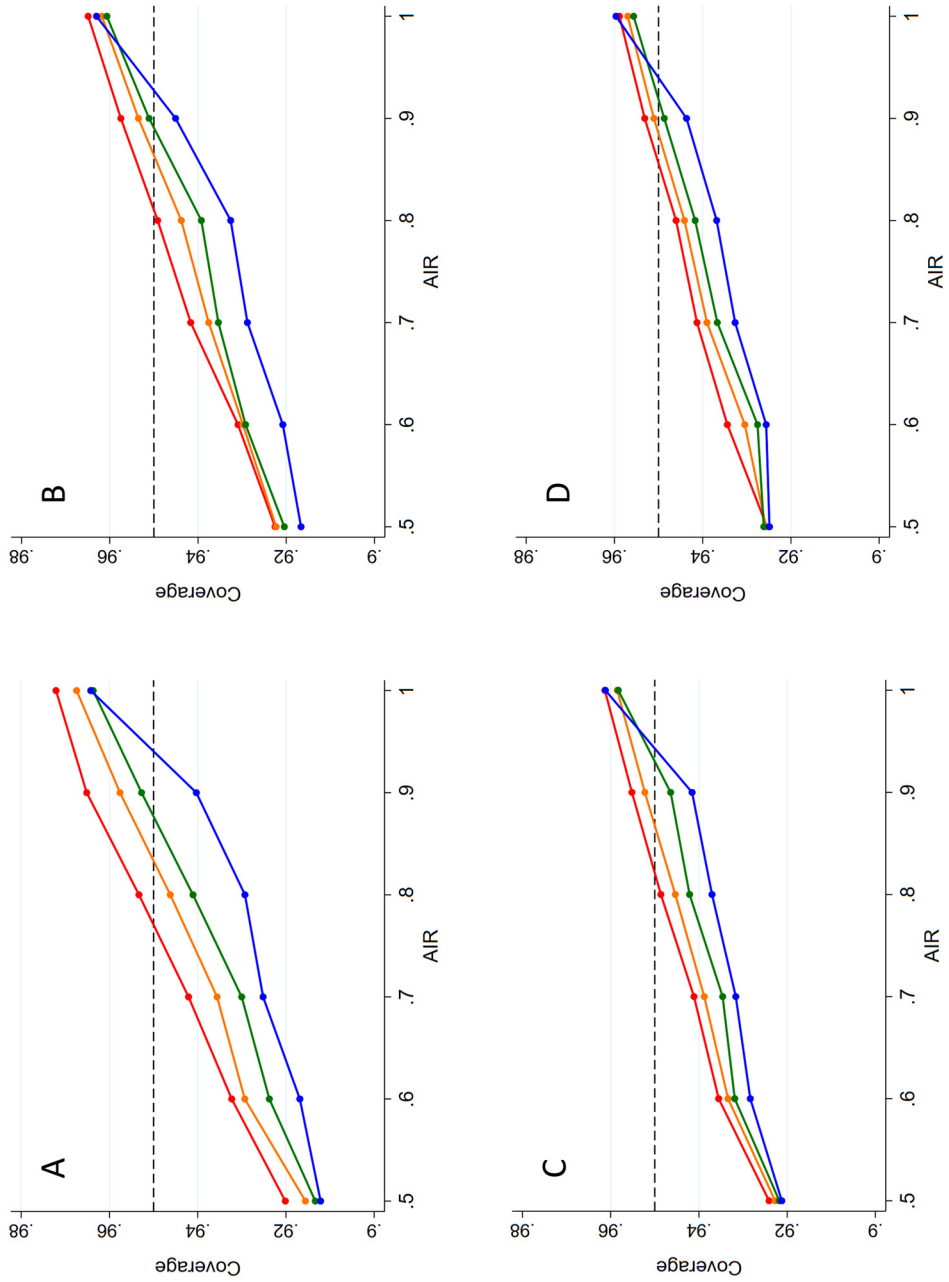


Figure 1: Coverage of lower 5% confidence limit computed using the delta method. (A) 40 expected counterfactual events per arm (B) 60 events (C) 80 events (D) 100 events. Efficacy of control treatment: 0.6, red line; 0.7, orange; 0.8, green; 0.9, blue. Horizontal line represents nominal coverage of 0.95.

Table 1: Coverage probabilities for the profile-likelihood-based lower 5% confidence limit (40 expected counterfactual events per arm).

Effectiveness of control treatment (θ_c)	AIR (Ψ)					
	0.5	0.6	0.7	0.8	0.9	1.0
0.6	0.9468	0.9521	0.9518	0.9522	0.9517	0.9502
0.7	0.9510	0.9539	0.9511	0.9522	0.9519	0.9511
0.8	0.9523	0.9522	0.9553	0.9517	0.9532	0.9518
0.9	0.9539	0.9538	0.9579	0.9489	0.9568	0.9615

Nominal coverage is 0.95.

Example

The BRIEF TB/A5279 study was a randomised, non-inferiority trial that compared two regimens for the prevention of active tuberculosis in HIV-infected patients who were living in areas of high tuberculosis prevalence or who had evidence of latent tuberculosis infection (Swindells et al. 2019). The reference regimen was 9 months of daily isoniazid alone (9-month arm) and the experimental regimen was 1 month of daily rifapentine plus isoniazid (1-month arm). The incidence of the primary endpoint (diagnosis of tuberculosis, or death from tuberculosis or unknown cause) were similar in the 1-month arm (32 endpoints, 4,926 person-years follow-up (PYFU), incidence rate 0.65 per 100 PYFU) and 9-month arm (33 endpoints, 4,896 PYFU, incidence rate 0.67 per 100 PYFU). The primary metric was the rate difference rather than the rate ratio, which is generally used in HIV prevention research. Non-inferiority was declared by the investigators because the upper 97.5% confidence limit of 0.30 per 100 PYFU was less than the pre-specified margin of 1.25 per 100 PYFU. However, this conclusion is questionable as the authors did not take the counterfactual placebo incidence into account. Notably, the observed incidence in the 9-month arm was markedly lower than the incidence rate assumed for the purposes of sample size calculation (2 per 100 PYFU).

Figure 2 shows the lower 5% and upper 95% confidence limits for the AIR as a function of the counterfactual incidence, computed using the delta and profile-likelihood methods. Consistent with results of Section 6.2, the delta method yields narrower confidence intervals. The figure also reveals the sensitive relationship between the lower confidence limit and the assumed counterfactual incidence, underlining the importance of obtaining as much information as possible about this parameter.

Figure 3 show the results of a Bayesian analysis (10,000 simulations) under two different priors for the counterfactual incidence: Gamma(10,0.001) and Gamma(10,0.002), corresponding to mean incidence rates of 1 and 2 per 100 PYFU, respectively. Without expert knowledge, we emphasise that this is an illustrative

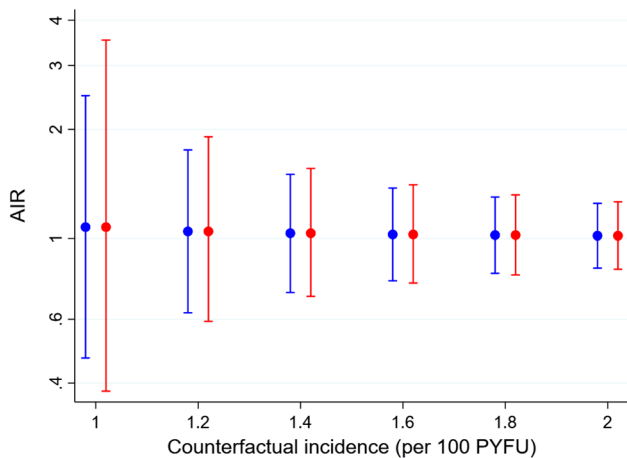


Figure 2: Lower 5% and upper 95% confidence limits for the AIR in the BRIEF TB/A5279 study. Delta method, blue line; profile-likelihood method, red line.

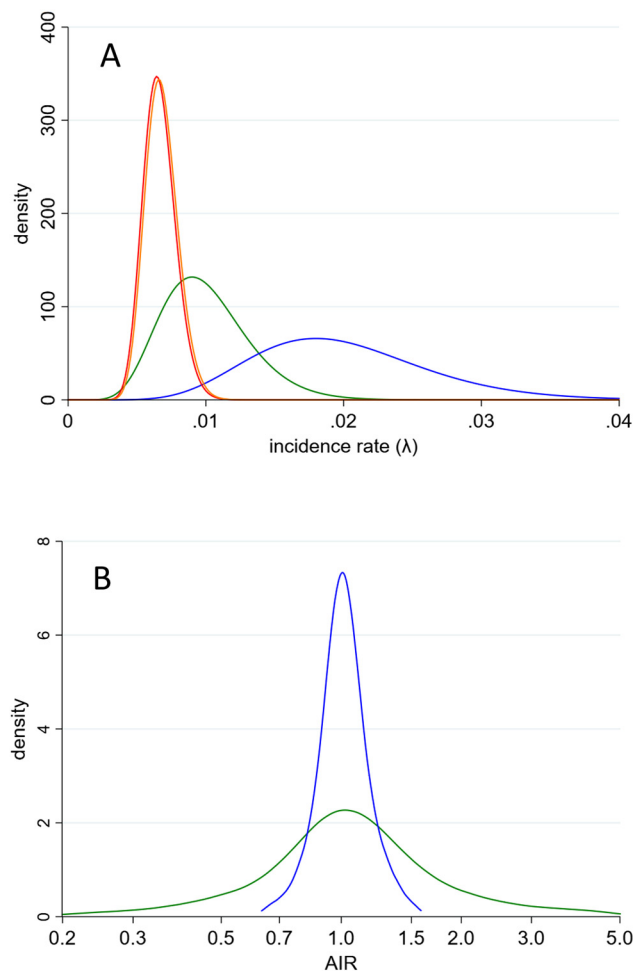


Figure 3: Bayesian analysis of BRIEF TB/A5279 study. (A) Posterior distributions of incidence rate in control regimen (orange line) and experimental regimen (red line); prior distributions of counterfactual incidence (mean = 0.01, green line; mean = 0.02, blue line). (B) Posterior distribution of AIR by prior distribution used for counterfactual incidence (mean = 0.01, green line; mean = 0.02, blue line).

rather than a definitive analysis. The lower incidence rate is broadly consistent with the overall $\sim 30\%$ efficacy of tuberculosis prophylaxis in HIV-infected patients (Ross et al. 2021); the higher value is the rate that the investigators postulated for the control regimen (post hoc, a substantial over-estimate).

For the low incidence scenario, 22.2% of initial simulations had to be re-sampled because of violation of Eq. (6). The posterior median (90% credibility interval) AIR was 1.038 (0.347, 3.627) under re-sampling strategy (a), 1.033 (0.373, 3.228) under strategy (b), and 1.031 (0.357, 3.281) under strategy (c). For the high incidence scenario, only 0.6% of initial simulations had to be re-sampled. The posterior median (90% credibility interval) AIR under re-sampling strategy (a) was 1.009 (0.760, 1.370). The values under the other re-sampling strategies were almost identical (all within ± 0.002). In general, our preference is to re-sample λ_p^* only (strategy (a)) since, lacking empirical data, λ_p is the most uncertain parameter. Figure 3B shows the posterior distributions for the AIR under this strategy, and highlights that inference is much tighter under the high incidence scenario.

Summary

We have described two approaches for calculating confidence limits for the AIR given a pre-specified value of the counterfactual incidence: a closed-form solution based on a Taylor series expansion (delta method), and an iterative method based on the profile-likelihood, for which R code is provided. The profile-likelihood method is preferred because of better coverage properties, but the delta method is valid when the experimental

treatment is no less effective than the control treatment. The difference between the two methods is minimal when the counterfactual incidence is much larger than the observed incidence in both the control and experimental arms. We also describe a simple Bayesian approach when the counterfactual incidence rate can be represented as a simple prior distribution. However, more precise inference can be achieved by harnessing other data which inform the prior distribution (Glidden, Stirrup, and Dunn 2020).

Acknowledgments: We thank Doug Taylor for discussions on three-arm trials and Trinh Duong for her comments on the paper.

Research funding: David Glidden was supported by US National Institutes of Health grant R01AI143357.

Author contribution: All authors have accepted responsibility for the entire content of this manuscript and approved its submission.

Competing interests: Authors state no conflict of interest.

Informed consent: Not applicable.

Ethical approval: Not applicable.

References

- Cole, S. R., H. Chu, and S. Greenland. 2014. "Maximum Likelihood, Profile Likelihood, and Penalized Likelihood: A Primer." *American Journal of Epidemiology* 179 (2): 252–60.
- Dunn, D. T., and D. V. Glidden. 2019. "The Connection between the Averted Infections Ratio and the Rate Ratio in Active-Control Trials of Pre-exposure Prophylaxis Agents." *Statistical Communications in Infectious Diseases* 11 (1): 20190006.
- Dunn, D. T., D. V. Glidden, O. T. Stirrup, and S. McCormack. 2018. "The Averted Infections Ratio: A Novel Measure of Effectiveness of Experimental HIV Pre-exposure Prophylaxis Agents." *Lancet HIV* 5 (6): e329–34.
- Gelman, A., J. B. Carlin, H. S. Stern, and D. B. Rubin. 1995. *Bayesian Data Analysis*. New York: Chapman & Hall.
- Ghosh, P., F. Nathoo, M. Gonen, and R. C. Tiwari. 2011. "Assessing Noninferiority in a Three-Arm Trial Using the Bayesian Approach." *Statistics in Medicine* 30 (15): 1795–808.
- Glidden, D. V., O. T. Stirrup, and D. T. Dunn. 2020. "A Bayesian Averted Infection Framework for PrEP Trials with Low Numbers of HIV Infections: Application to the Results of the DISCOVER Trial." *Lancet HIV* 7 (11): e791–6.
- Graham, P. L., K. Mengersen, and A. P. Morton. 2003. "Confidence Limits for the Ratio of Two Rates Based on Likelihood Scores: Non-iterative Method." *Statistics in Medicine* 22: 2071–83.
- Li, H.-Q., M.-L. Tang, and W.-K. Wong. 2014. "Confidence Intervals for Ratio of Two Poisson Rates Using the Method of Variance Estimates Recovery." *Computational Statistics* 29: 869–89.
- Oehlert, G. W. 1992. "A Note on the Delta Method." *The American Statistician* 46: 27–9.
- Pigeot, I., J. Schafer, J. Rohmel, and D. Hauschke. 2003. "Assessing Non-inferiority of a New Treatment in a Three-Arm Clinical Trial Including a Placebo." *Statistics in Medicine* 22 (6): 883–99.
- Price, R. M., and D. G. Bonett. 2000. "Estimating the Ratio of Two Poisson Rates." *Computational Statistics & Data Analysis* 34: 345–56.
- Ross, J. M., A. Badje, M. X. Rangaka, A. S. Walker, A. E. Shapiro, K. K. Thomas, X. Anglaret, S. Eholie, D. Gabillard, A. Boule, G. Maartens, R. J. Wilkinson, N. Ford, J. E. Golub, B. G. Williams, and R. V. Barnabas. 2021. "Isoniazid Preventive Therapy Plus Antiretroviral Therapy for the Prevention of Tuberculosis: A Systematic Review and Meta-Analysis of Individual Participant Data." *Lancet HIV* 8 (1): e8–15.
- Sahai, H., and A. Khurshid. 1993. "Confidence Intervals for the Ratio of Two Poisson Means." *The Mathematical Scientist* 18: 43–50.
- Snapinn, S., and Q. Jiang. 2008. "Preservation of Effect and the Regulatory Approval of New Treatments on the Basis of Non-inferiority Trials." *Statistics in Medicine* 27 (3): 382–91.
- Swindells, S., R. Ramchandani, A. Gupta, C. A. Benson, J. Leon-Cruz, N. Mwelase, J. Juste, J. Lama, A. Valenica, A. Omoz-Oarhe, K. Supparatpinyo, G. Masheto, L. Mohapi, R.O. da Silva Escada, S. Mawlana, P. Banda, P. Severe, J. Hakim, C. Kanyama, D. Langat, L. Moran, J. Andersen, C. V. Fletcher, E. Nuermberger, and R. E. Chaisson, BRIEF TB/A5279 Study Team. 2019. "One Month of Rifapentine Plus Isoniazid to Prevent HIV-Related Tuberculosis." *New England Journal of Medicine* 380 (11): 1001–11.

Supplementary Material: The online version of this article offers supplementary material (<https://doi.org/10.1515/scid-2021-0002>).