Bayes Bayesian Biostatistics Leiden, NL 2026

Universität Zürich[UZH]

CRS UZH

Bridging Paradigms in Drug Approval: The Optimal Bayes Procedure for Evidence Synthesis

Rubens Mita, Rachel Heyard, Leonhard Held

Center for Reproducible Science and Research Synthesis, University of Zurich

Context and Overview

— **Historically**: "at least two adequate and well-controlled studies, each convincing on its own" (FDA, 1998)

⇒ **Two-trials rule (2TR)**: one-sided $p_1, p_2 \leq \alpha = 0.025$

— **2026**: FDA's new default is **one pivotal trial + confirmatory evidence** (Makary & Prasad, NEJM 2026)

Open question: when two independent trials are available (e.g. a replication trial), **how should their evidence be combined?**

The two-trials rule and its alternatives

Same overall T1E rate $\alpha^2 = 0.025^2$: every success region has the same area

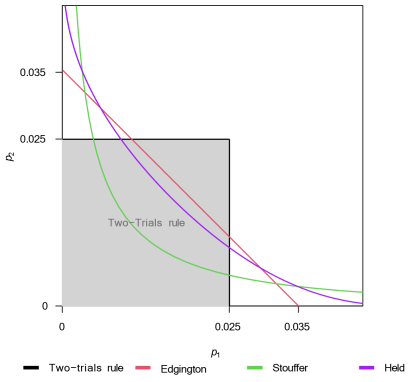
- **Intersection** null ($H_{0\cap} = H_0^1 \cap H_0^2$): both trials have null effect
  $\rightarrow$ **overall** T1E rate, $\alpha^2 = 0.025^2$ for the 2TR
- **Union** (no-replicability) null ($H_{0\cup} = H_0^1 \cup H_0^2$): at least one trial has null effect
  $\rightarrow$ **partial** T1E rate, $\alpha = 0.025$ for the 2TR

The two-trials rule and its alternatives

Same overall T1E rate $\alpha^2 = 0.025^2$: every success region has the same area

- **Intersection** null ($H_{0\cap} = H_0^1 \cap H_0^2$): both trials have null effect
  $\rightarrow$ **overall** T1E rate, $\alpha^2 = 0.025^2$ for the 2TR
- **Union** (no-replicability) null ($H_{0\cup} = H_0^1 \cup H_0^2$): at least one trial has null e
  $\rightarrow$ **partial** T1E rate, $\alpha = 0.025$ for the 2TR
- 2TR overly rigid: fails when one trial is highly significant but the other is slightly above α
- Alternative methods keep overall T1E at α^2 but allow partial T1E inflation:

Method	*Success if*	*Bound on partial T1E*
2TR	$\max\{p_1, p_2\} \leq \alpha$	2.5%
Edgington	$p_1 + p_2 \leq \sqrt{2}\alpha$	3.5%
Harmonic mean test (Held)	$\frac{1}{z_1^2} + \frac{1}{z_2^2} \leq \frac{4}{\chi^2_1(1-4\alpha^2)}$	6.5%
Stouffer	$z_1 + z_2 \geq \sqrt{2} z_{1-\alpha^2}$	100% (trivially)



A Bayes factor perspective: two null hypotheses

So far: frequentist focus on error rates. A **Bayesian** quantifies the evidence for efficacy with the Bayes factor BF_{10}

- Assume test statistics $z_i = \Phi^{-1}(1 - p_i)$, $i = 1, 2$; same design, power and sample size for both trials
- **Alternative**: $z_i \mid H_1 \sim \text{N}(\mu, 1 + \tau^2)$, with $\mu = \Phi^{-1}(1 - \alpha) + \Phi^{-1}(1 - \beta)$ from the trial power $1 - \beta$
- Select τ^2 as variance of normal prior on μ, parameterised in terms of τ/μ

A Bayes factor perspective: two null hypotheses

So far: frequentist focus on error rates. A **Bayesian** quantifies the evidence for efficacy with the Bayes factor BF_{10}

- Assume test statistics $z_i = \Phi^{-1}(1 - p_i)$, $i = 1, 2$; same design, power and sample size for both trials
- **Alternative**: $z_i \mid H_1 \sim \mathsf{N}(\mu, 1 + \tau^2)$, with $\mu = \Phi^{-1}(1 - \alpha) + \Phi^{-1}(1 - \beta)$ from the trial power $1 - \beta$
- Select τ^2 as variance of normal prior on μ, parameterised in terms of τ/μ

Intersection null

$z_1, z_2 \sim \mathsf{N}(0, 1)$

$$\mathsf{BF}_{10\cap} = \frac{f(z_1 \mid H_1) f(z_2 \mid H_1)}{f(z_1 \mid H_{0\cap}) f(z_2 \mid H_{0\cap})}$$

A Bayes factor perspective: two null hypotheses

So far: frequentist focus on error rates. A **Bayesian** quantifies the evidence for efficacy with the Bayes factor BF_{10}

- Assume test statistics $z_i = \Phi^{-1}(1 - p_i)$, $i = 1, 2$; same design, power and sample size for both trials
- **Alternative**: $z_i \mid H_1 \sim \mathsf{N}(\mu, 1 + \tau^2)$, with $\mu = \Phi^{-1}(1 - \alpha) + \Phi^{-1}(1 - \beta)$ from the trial power $1 - \beta$
- Select τ^2 as variance of normal prior on μ, parameterised in terms of τ/μ

Intersection null

$z_1, z_2 \sim \mathsf{N}(0, 1)$

$$\mathsf{BF}_{10\cap} = \frac{f(z_1 \mid H_1) f(z_2 \mid H_1)}{f(z_1 \mid H_{0\cap}) f(z_2 \mid H_{0\cap})}$$

Union null: mixture with weight w_1

$z_1,\ z_2 \sim \mathsf{N}(0, 1)$	with weight w_1
$z_1 \sim \mathsf{N}(\mu, 1 + \tau^2)$, $z_2 \sim \mathsf{N}(0, 1)$	with weight $w_2 = (1 - w_1)/2$
$z_2 \sim \mathsf{N}(\mu, 1 + \tau^2)$, $z_1 \sim \mathsf{N}(0, 1)$	with weight $w_3 = (1 - w_1)/2$

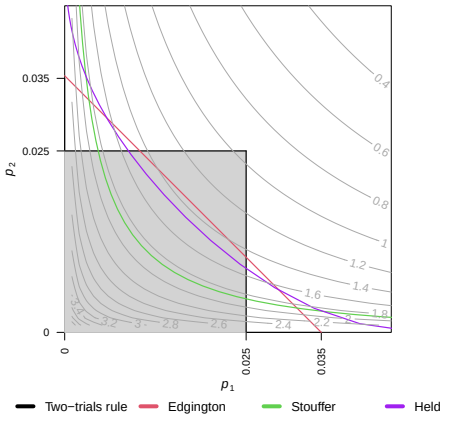
$$\mathsf{BF}_{10\cup} = \frac{f(z_1 \mid H_1) f(z_2 \mid H_1)}{\sum_{j=1}^{3} w_j \, f(z_1 \mid H_{0\cup}) f(z_2 \mid H_{0\cup})}$$

Intersection null is the special case $w_1 = 1$

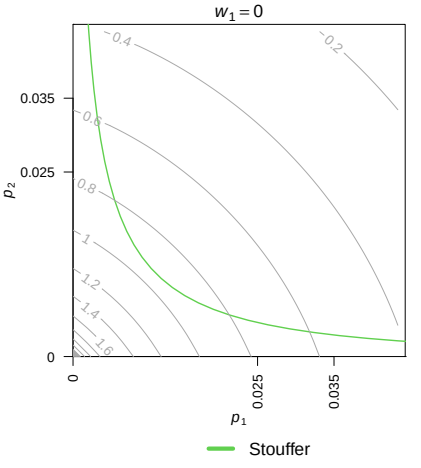
A hybrid approach: A p-value combination method based on BF

BF_{10} contour lines on $\log_{10}$ scale; $\tau = 0.25\mu \approx 0.81$ (at 90% power)

Intersection null



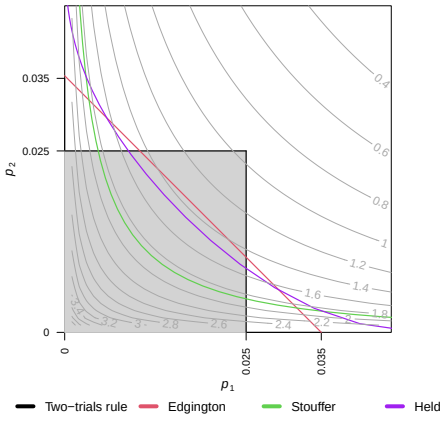
Union null



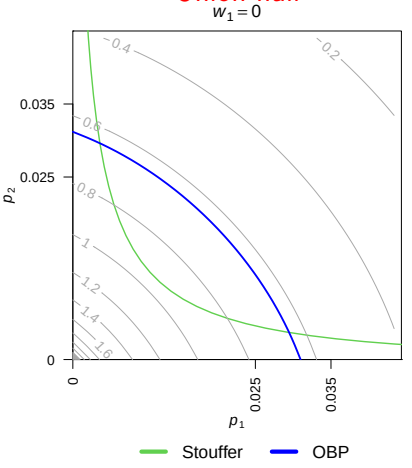
A hybrid approach: A p-value combination method based on BF

BF_{10} contour lines on $\log_{10}$ scale; $\tau = 0.25\mu \approx 0.81$ (at 90% power)

Intersection null



Union null



– **Idea**: success boundary along a BF_{10} contour, chosen so that its area is $\alpha^2 = 0.000625$

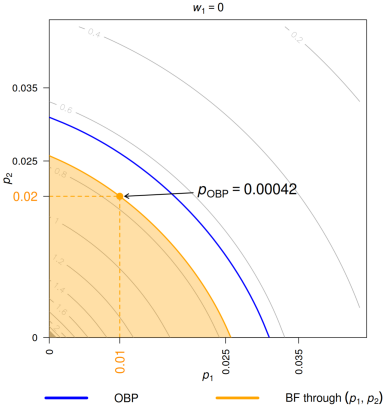
The Optimal Bayes Procedure

From Bayes factors to a combined p-value

Optimal Bayes Procedure (OBP):

→ **Success** if $\text{BF}_{10\cup}(z_1, z_2) \geq \kappa$, where κ is the BF contour with area α^2 from the previous slide

- **Combined p-value**: area under the BF contour through the observed (p_1, p_2), e.g. $p_1 = 0.01$, $p_2 = 0.02$ (shaded)
- **Success** $\iff$ combined $p_{\text{OBP}} \leq \alpha^2 = 0.000625$
- **Adaptive**: optimal for any chosen trial power, prior variance τ^2 and weight w_1



Conclusive remarks and take-home message

1. **Hybrid**: OBP is a p-value combination method based on Bayesian principles.
2. **Calibrated**: Success if $p_{\text{OBP}} \leq \alpha^2$, retaining the same overall T1E rate as the two-trials rule.

– Outlook: Extend method to more than 2 studies



rubens.mita@uzh.ch