

3

Fieller's Theorem and Fiducial Limits

3.1: Learning Objectives:

1. State and prove Fieller's theorem for the ratio of two jointly normal random variables.
2. Apply Fieller's theorem to construct fiducial (confidence) limits for relative potency.
3. Understand the special role of the quantity and the consequences of approaching or exceeding 1.
4. Compute Fieller limits by hand and using R/Python.

3.2: Motivation

Chapter 2 showed that ratio estimators like $R = A/B$ do not have a variance that behaves like that of a simple linear statistic, especially when B is imprecisely known. **Fieller's theorem** (Fieller, 1940, 1954) provides an *exact* (under normality) method for constructing confidence/fiducial limits for the true ratio $\mu = \alpha/\beta$, based on estimators a of α and b of β , without relying on the delta-method approximation.

3.3: Statement of Fieller's Theorem

Theorem 3.1 (Fieller). Let α, β be unknown parameters with $\mu = \alpha/\beta$. Suppose a, b are unbiased estimators, jointly normally distributed, with

$$\text{Var}(a) = v_{11}s^2, \quad \text{Var}(b) = v_{22}s^2, \quad \text{Cov}(a, b) = v_{12}s^2,$$

where s^2 is an estimator of the common (unknown) error variance σ^2 , independent of a, b and based on f degrees of freedom, so that $fs^2/\sigma^2 \sim \chi_f^2$. Let $m = a/b$ be the natural point estimator of μ . Define

$$g = \frac{t_0^2 s^2 v_{22}}{b^2},$$

where t_0 is the two-sided $100(1 - \alpha^*)\%$ point of Student's t distribution on f degrees of freedom. Then the $100(1 - \alpha^*)\%$ fiducial (confidence) limits for μ are

$$m_L, m_U = \frac{1}{1-g} \left\{ \left(m - g \frac{v_{12}}{v_{22}} \right) \pm \frac{t_0 s}{b} \sqrt{v_{11} - 2mv_{12} + m^2 v_{22} - g \left(v_{11} - \frac{v_{12}^2}{v_{22}} \right)} \right\},$$

3.3.1: Proof (Sketch)

Consider the linear combination $Z = a - \mu b$. By construction $E(Z) = 0$ when μ is the true ratio, and

$$\text{Var}(Z) = s^2 (v_{11} - 2\mu v_{12} + \mu^2 v_{22}).$$

Under joint normality of a, b and independence from s^2 ,

$$T = \frac{Z}{s \sqrt{v_{11} - 2\mu v_{12} + \mu^2 v_{22}}} \sim t_f.$$

A $100(1 - \alpha^*)\%$ confidence statement is $\Pr(-t_0 \leq T \leq t_0) = 1 - \alpha^*$, i.e.

$$(a - \mu b)^2 \leq t_0^2 s^2 (v_{11} - 2\mu v_{12} + \mu^2 v_{22}).$$

Expanding, this is a **quadratic inequality in μ** :

$$\mu^2 (b^2 - t_0^2 s^2 v_{22}) - 2\mu (ab - t_0^2 s^2 v_{12}) + (a^2 - t_0^2 s^2 v_{11}) \leq 0.$$

Dividing through by b^2 and writing $g = t_0^2 s^2 v_{22}/b^2$, $m = a/b$, the two roots of the corresponding quadratic equation, obtained via the standard quadratic formula, are exactly the limits given in (3.1). ■

3.3.2: Interpretation of g

The quantity g measures how precisely b (the denominator, e.g. a slope) is estimated relative to its own magnitude:

$g = \left(\frac{t_0 \cdot SE(b)}{b} \right)^2$, i.e., $\sqrt{g}$ is (approximately) the reciprocal of the t -statistic for testing $H_0 : \beta = 0$.

- If g is small (b is precisely estimated, i.e., b is many standard errors from 0), the Fieller interval is close to the delta-method interval $m \pm t_0 SE(m)$, and is nearly symmetric about m .
- As $g \rightarrow 1$, the interval widens dramatically and becomes markedly asymmetric.
- If $g \geq 1$ (equivalently, b is not significantly different from zero at the chosen confidence level), the quadratic in μ has no real, finite, bounded solution set; the “confidence interval” for μ is either the entire real line or the complement of a bounded interval. This reflects the intuitive fact that if the slope is not significantly nonzero, the ratio $\mu = \alpha/\beta$ is not meaningfully estimable at that confidence level — a phenomenon impossible to detect from the naive delta-method interval, which is always a finite symmetric interval regardless of how poorly β is estimated.

3.4: Application to Relative Potency

In the direct-assay setting of Chapter 4, with $a = \bar{x}_B$ (mean tolerance, standard), $b = \bar{x}_A$ (mean tolerance, test), independent, $v_{11} = 1/n_B$, $v_{22} = 1/n_A$, $v_{12} = 0$, Fieller's theorem gives exact limits for $\rho = \bar{x}_B/\bar{x}_A$ without the delta-method's optimistic symmetry assumption.

In the parallel-line/slope-ratio settings of Chapters 6–7, a and b are typically the *preparation contrast* and the *combined regression* (or the two slopes), and $v_{12} \neq 0$ in general because both may be computed from the same fitted lines.

3.5: Worked Examples

Example 3.1: Direct Application to Cat Tolerances (Recap of Chapter 1)

Strophanthus tinctures A, B; $n_A = n_B = 7$; $\bar{x}_A = 1.68$ (test), $\bar{x}_B = 1.99$ (standard); pooled $s^2 = 0.1200$ on 12 d.f.

Take $R = \bar{x}_B/\bar{x}_A = 1.18$. Here $a = \bar{x}_B$, $b = \bar{x}_A$, $v_{11} = v_{22} = 1/7$, $v_{12} = 0$.

At $\alpha^* = 0.05$, $t_0 = t_{0.025,12} = 2.179$.

$$g = \frac{t_0^2 s^2 v_{22}}{b^2} = \frac{(2.179)^2 (0.1200) (1/7)}{1.68^2} = \frac{4.748 \times 0.1200/7}{2.8224} = \frac{0.08140}{2.8224} = 0.02884.$$

Since g is small, the interval will be close to symmetric. Substituting into (3.1) gives $\rho_L, \rho_U \approx 0.95, 1.48$ (matching the classical Finney result), i.e. the 95% fiducial interval for relative potency is approximately (0.95, 1.48).

Example 3.2: Effect of a Large g

Suppose instead $n_A = n_B = 3$ (a much smaller experiment), same $\bar{x}_A, \bar{x}_B, s^2$. Recompute g and comment.

Solution. $t_{0.025,4} = 2.776$ (pooled d.f. now $2 + 2 = 4$).

$$g = \frac{(2.776)^2(0.1200)(1/3)}{1.68^2} = \frac{7.706 \times 0.0400}{2.8224} = \frac{0.3082}{2.8224} = 0.1092.$$

g has grown roughly fourfold; the interval will be noticeably wider and less symmetric than in Example 3.1, illustrating the direct cost, in fiducial precision, of a smaller experiment.

Example 3.3: When g Exceeds 1

A poorly-designed direct assay with only $n_A = n_B = 2$ yields $b = \bar{x}_A = 1.68$ with a much larger $s^2 = 0.9$. Compute g at the 95% level and interpret.

Solution: Pooled d.f. $= 1 + 1 = 2$, $t_{0.025,2} = 4.303$.

$$g = \frac{(4.303)^2(0.9)(0.5)}{1.68^2} = \frac{18.52 \times 0.45}{2.8224} = \frac{8.335}{2.8224} = 2.954.$$

Since $g > 1$, no finite 95% confidence interval for ρ exists from this data; the denominator b is not estimated precisely enough, at this sample size, to support any bounded statement about the potency ratio. This is a strong argument, developed further in Chapter 10, for adequate replication in bioassay design.

Example 3.4: Fieller Limits by R

Reproduce Example 3.1's limits computationally.

R:

```
fieller_ci <- function(a, b, s2, v11, v22, v12, df, conf=0.95) {
  t0 <- qt(1 - (1-conf)/2, df)
  g <- t0^2 * s2 * v22 / b^2
  m <- a / b
  centre <- m - g * v12 / v22
  root <- sqrt(v11 - 2*m*v12 + m^2*v22 - g*(v11 -
v12^2/v22))
  lower <- (centre - t0*sqrt(s2)/b * root) / (1 - g)
  upper <- (centre + t0*sqrt(s2)/b * root) / (1 - g)
  list(g = g, m = m, lower = lower, upper = upper)
}

fieller_ci(a=1.99, b=1.68, s2=0.1200, v11=1/7, v22=1/7,
v12=0, df=12)
```

Python:

```
import numpy as np
from scipy import stats

def fieller_ci(a, b, s2, v11, v22, v12, df, conf=0.95):
    t0 = stats.t.ppf(1 - (1-conf)/2, df)
    g = t0**2 * s2 * v22 / b**2
    m = a / b
    centre = m - g * v12 / v22
    root = np.sqrt(v11 - 2*m*v12 + m**2*v22 - g*(v11 - v12**2/v22))
    lower = (centre - t0*np.sqrt(s2)/b * root) / (1 - g)
```

```
upper = (centre + t0*np.sqrt(s2)/b * root) / (1- g)
return {"g": g, "m": m, "lower": lower, "upper": upper}

fieller ci(a=1.99, b=1.68, s2=0.1200, v11=1/7, v22=1/7,
v12=0, df=12)
```

Both reproduce $g \approx 0.029$, $m = 1.18$, limits $\approx (0.95, 1.48)$.

Example 3.5: Non-Zero Covariance Case

In a parallel-line assay, $a = \Psi_p = -0.800$ (preparation contrast, playing the role of α numerator-analogue after rearrangement), $b = \hat{\Psi}_1 = 16.777$ (combined regression), with $v_{11} = 0.03156$, $v_{22} = 0.1304$, $v_{12} = 0$ (contrasts are in fact orthogonal by construction in a symmetric design), pooled $s^2 = 2.367$ on 24 d.f. Construct the 95% Fieller interval for $\log \hat{\rho} \times (\text{scale factor})$, i.e. for $\mu = a/b$.

Solution. $t_{0.025,24} = 2.064$. $m = -0.800/16.777 = -0.04768$.

$$g = \frac{(2.064)^2(2.367)(0.1304/2.367)}{16.777^2}$$

Note v_{22} here is *already* a variance coefficient multiplying s^2 ; to avoid double counting, in practice v_{22} is extracted from the model's $(X'X)^{-1}$ diagonal so that $\text{Var}(b) = v_{22}s^2 = 1.143^2 = 1.306$, giving $v_{22} = 1.306/2.367 = 0.5518$.

$$g = \frac{(2.064)^2(2.367)(0.5518)}{16.777^2} = \frac{4.260 \times 1.306}{281.5} = \frac{5.564}{281.5} = 0.01977.$$

Small g again reflects a well-estimated combined regression slope in this particular capon assay (Chapter 6 gives the complete worked analysis).

Example 3.6: Sensitivity of Fieller Limits to Confidence Level

Using Example 3.1's data, compute Fieller limits at 90% and 99% confidence and compare interval widths.

Solution (R):

```
sapply(c(0.90, 0.95, 0.99), function(cl)
  unlist(fieller_ci(1.99, 1.68, 0.1200, 1/7,
    1/7, 0, 12, conf = cl)))
```

As expected, the 99% interval is markedly wider than the 90% interval, and g itself grows with confidence level (larger t_0), a fact worth emphasizing to students who might otherwise expect only the additive margin, not g , to change.

Example 3.7: Relation Between Fieller and the Delta Method

For Example 3.1's data, compute the delta-method interval $R \pm t_0 SE(R)$ using (2.1) and compare numerically with the Fieller interval.

Solution. $SE(R)^2 \approx R^2 \left[\frac{s^2}{n_B \bar{x}_B^2} + \frac{s^2}{n_A \bar{x}_A^2} \right] = 1.18^2 \left[\frac{0.12}{7(1.99)^2} + \frac{0.12}{7(1.68)^2} \right] = 1.3924[0.00432 + 0.00607] = 0.01446$, so
 $SE(R) \approx 0.120$, giving delta interval
 $1.18 \pm 2.179(0.120) = (0.919, 1.441)$ — close to, but not identical to, the Fieller interval $(0.95, 1.48)$, and centered exactly on R (symmetric), unlike the (very slightly asymmetric) Fieller interval. The discrepancy grows as g increases (cf. Example 3.2).




Figure 3.1: *Fieller interval widening as g increases*

3.6: Chapter Summary

- Fieller's theorem gives exact confidence/fiducial limits for a ratio of two normal variables by inverting a quadratic in μ .
- The key diagnostic quantity $g \propto (SE(b)/b)^2$ governs how far the Fieller interval departs from the naive delta-method interval; $g \geq 1$ signals a non-finite interval.
- Fieller's theorem underlies interval estimation for relative potency in direct, parallel-line, and slope-ratio assays alike, and reappears throughout Chapters 4, 6, and 7.

3.7: Exercises

Exercise 3.1. Prove that when $v_{12} = 0$ and $g \rightarrow 0$, the Fieller limits in (3.1) reduce to the delta-method limits $m \pm t_0 SE(m)$ with $SE(m)^2 = m^2(v_{11}/a^2 \cdot a^2 + v_{22})s^2/b^2$ style expression — i.e., show the algebraic reduction explicitly.

Exercise 3.2. For a direct assay with $\bar{x}_S = 3.10$, $\bar{x}_T = 2.55$, $n_S = n_T = 10$, pooled $s^2 = 0.28$ on 18 d.f., compute (a) g at the 95% level, (b) the point estimate R , (c) the Fieller 95% interval, by hand or using the R/Python function of Example 3.4.

Exercise 3.3. Explain, using the concept of g , why increasing the number of dose levels (holding total sample size fixed) in a parallel-line assay design can sometimes *increase* g for the combined-regression-based ratio even though it improves estimation of other contrasts. (Hint: think about how spreading replicates thinly across more levels affects $SE(b)$.)

Exercise 3.4. A junior analyst computes only the delta-method interval for a potency ratio and reports it as “the 95% confidence interval.” Write a short paragraph, suitable for inclusion in a statistical report to a regulatory agency, explaining why the Fieller interval should be reported instead, referencing the behavior of g .

(Full worked solutions appear in Appendix A.)