

Testing the Ratio of Two Poisson Rates

Michael P. Fay

June 5, 2007

1 Example

Here is a quick example of the function `rateratio.test`. Suppose you have two rates that you assume are Poisson and you want to test that they are different. Suppose you observe 2 events with time at risk of $n = 17877$ in one group and 9 events with time at risk of $m = 16660$ in another group. Here is the test:

```
> rateratio.test(c(2, 9), c(n, m))
```

```
Exact Rate Ratio Test, assuming Poisson counts
```

```
data: c(2, 9) with time of c(n, m), null rate ratio 1
p-value = 0.05011
alternative hypothesis: true rate ratio is not equal to 1
95 percent confidence interval:
 0.02177406 1.00054910
sample estimates:
Rate Ratio      Rate 1      Rate 2
0.2070941557 0.0001118756 0.0005402161
```

The result is barely non-significant at the 0.05 level. This example was chosen to make a point, that is why the p-value is so close to 0.05. See Section 5 below.

2 Assumptions and Notation

Assume that $Y \sim \text{Poisson}(n\lambda_y)$ and $X \sim \text{Poisson}(m\lambda_x)$. We are interested in the rate ratio, $\theta = \lambda_y/\lambda_x$. The parameters n and m are assumed known and represent the time spent in the Poisson process with the given rates. For example, n could be the number of person-years at risk associated with Y . We wish to test one of the three following hypotheses:

less

$$H_0 : \theta \geq \Delta$$

$$H_1 : \theta < \Delta$$

greater

$$H_0 : \theta \leq \Delta$$

$$H_1 : \theta > \Delta$$

two-sided

$$\begin{aligned} H_0 : \theta &= \Delta \\ H_1 : \theta &\neq \Delta \end{aligned}$$

For the tests using the rate ratios, we can use the uniformly most powerful (UMP) unbiased test. This test is based on conditioning on the sum $X + Y$ (see e.g., Lehmann and Romano, 2005, p. 125 or p. 152 of Lehmann, 1986). We modify Lehmann's presentation by allowing the constants m and n , representing the time in the Poisson process. We have that

$$Y|X + Y = t \sim \text{Binomial}(t, p(\theta))$$

where

$$p(\theta) = \frac{n\lambda_y}{n\lambda_y + m\lambda_x} = \frac{n\theta}{n\theta + m}. \quad (1)$$

3 Confidence Intervals

Since $p(\theta)$ is a monotonic increasing function of θ , if we have exact confidence intervals for $p(\theta)$, then we can transform them to exact confidence intervals for θ . The `R` function `binom.test` gives exact intervals for binomial observations (see Clopper and Pearson, 1934 or Leemis and Trivedi, 1996). We write the $100(1 - \alpha)\%$ one-sided lower confidence limit for p as $L_p(Y; \alpha)$ and the $100(1 - \alpha)\%$ one-sided upper confidence limit for p as $U_p(Y; \alpha)$. For the $100(1 - \alpha)\%$ two-sided confidence interval, `binom.test` and Clopper and Pearson (1934) use the central confidence interval defined as $[L_p(Y; \alpha/2), U_p(Y; \alpha/2)]$. The *central* confidence interval guarantees that

$$Pr[p < L_p(Y; \alpha/2) | p, t] \leq \alpha/2 \text{ for all } p \text{ and } t$$

and

$$Pr[p > U_p(Y; \alpha/2) | p, t] \leq \alpha/2 \text{ for all } p \text{ and } t$$

For shorter exact intervals which are not central see Blaker (2000) and the references therein.

To obtain confidence intervals for θ we set

$$L_p(Y; \alpha) = \frac{nL_\theta(Y; \alpha)}{nL_\theta(Y; \alpha) + m},$$

and perform some algebra to get

$$L_\theta(Y; \alpha) = \frac{mL_p(Y; \alpha)}{n\{1 - L_p(Y; \alpha)\}}.$$

Similarly,

$$U_\theta(Y; \alpha) = \frac{mU_p(Y; \alpha)}{n\{1 - U_p(Y; \alpha)\}}.$$

4 P-values

Just as in the last section, we can use results from the tests of p and translate them to tests of θ . Thus, for example the one-sided p-value of the test with the alternative hypothesis that $\theta > \Delta$ is equivalent to the one-sided p-value of the test that $p > p(\Delta)$. For the two-sided p-value we use the minimum of 1 or twice the minimum of the two one-sided p-values. There are other ways to define the two-sided p-value but they do not give equivalent inferences with the confidence intervals described above (see Section 5 below).

5 Relationship to Other Tests

In the R function `binom.test` (as least up until Version 2.5.0) the two-sided p-value is calculated by defining more extreme responses as those values with binomial density functions less than or equal to the observed density. This is a valid and reasonable way of defining two-sided p-values but it *does not match* with the two-sided confidence intervals. Returning to our example from Section 1 but using `binom.test` we can match the confidence intervals by using equation 1.

```
> n <- 17877
> m <- 16674
> rateratio.test(c(2, 9), c(n, m))$conf.int

[1] 0.02179236 1.00138990
attr(,"conf.level")
[1] 0.95

> b.ci <- binom.test(2, 2 + 9, p = n/(n + m))$conf.int
> theta.ci <- m * b.ci/(n * (1 - b.ci))
> theta.ci

[1] 0.02179236 1.00138990
attr(,"conf.level")
[1] 0.95
```

However, the p-values do not match for a two-sided test of $p(1) = n/(n + m)$.

```
> R.Version()$version.string

[1] "R version 2.5.0 (2007-04-23)"

> rateratio.test(c(2, 9), c(n, m))
```

Exact Rate Ratio Test, assuming Poisson counts

```
data: c(2, 9) with time of c(n, m), null rate ratio 1
p-value = 0.05027
```

```
alternative hypothesis: true rate ratio is not equal to 1
95 percent confidence interval:
 0.02179236 1.00138990
```

```
sample estimates:
```

```
Rate Ratio      Rate 1      Rate 2
0.2072681844 0.0001118756 0.0005397625
```

```
> binom.test(2, 2 + 9, p = n/(n + m))
```

```
Exact binomial test
```

```
data: 2 and 2 + 9
```

```
number of successes = 2, number of trials = 11, p-value = 0.03315
```

```
alternative hypothesis: true probability of success is not equal to 0.517409
95 percent confidence interval:
```

```
0.0228312 0.5177559
```

```
sample estimates:
```

```
probability of success
0.1818182
```

The p-values for `rateratio.test` are internally consistent, i.e., if the two-sided p-value is less than α then the $100(1 - \alpha/2)\%$ confidence interval does not contain Δ . In contrast the p-values for `binom.test` are not internally consistent as shown by the example. A similar internal inconsistency happens with `fisher.test`.

```
> fisher.test(matrix(c(2, 9, n - 2, m - 9), 2, 2))
```

```
Fisher's Exact Test for Count Data
```

```
data: matrix(c(2, 9, n - 2, m - 9), 2, 2)
```

```
p-value = 0.03312
```

```
alternative hypothesis: true odds ratio is not equal to 1
95 percent confidence interval:
```

```
0.02178155 1.00121047
```

```
sample estimates:
```

```
odds ratio
0.2072015
```

References

- Blaker, H. (2000). "Confidence curves and improved exact confidence intervals for discrete distributions" *Canadian Journal of Statistics* **28**, 783-798 (correction **29**, 681).
- Clopper, C.J. and Pearson, E.S. (1934). "The use of confidence or fiducial limits illustrated in the case of the binomial". *Biometrika*, **26**, 404-413.

- Lehmann, E.L. (1986). *Testing Statistical Hypotheses (second edition)* Wadsworth & Brooks/Cole, Pacific Grove, California.
- Lehmann, E.L., and Romano, J.P. (2005). *Testing Statistical Hypotheses (third edition)* Springer: New York.
- Leemis, L.M. and Trivedi, K.S. (1996). “A comparison of approximate interval estimators for the Bernoulli parameter” *American Statistician* **50**, 63-68.