Confidence interval for a ratio of parameters

Gilles Guillot

May 17, 2017

Abstract

We consider the problem of providing a confidence interval when the parameter of interest is $R = \theta_1 / \theta_2$ where θ_1 and θ_2 are two unknown parameters. The Delta method and the Fieller method are general and can be implemented in any context. Other methods can not be applied when only summary statistics are available. There is no guarantee that the actual coverage probability matches the nominal coverage.

Contents

We consider the problem of providing a confidence interval when the parameter of interest is $R = \theta_1/\theta_2$ where θ_1 and θ_2 are two unknown parameters. This problem is not well addressed by standard methods (because classical assumptions such as Normal or Student distribution make little sense for a ratio) and therefore received little attention in classical textbooks. We outline briefly the main methods that have been proposed earlier and comment on their usefulness in the context of the estimation of relative toxicity when data are extracted from the literature. The specificity of this context is that estimates of θ_1 and θ_2 together with estimates of the associated standard deviation (denoted below $\hat{\theta}_1$, $\hat{\theta}_2$, s_1 and s_2 are available) but neither the individual data points that have been used to obtain those estimates nor even the sample sizes are routinely available.

The ratio estimator is biased.

If $\hat{\theta}_1$ and $\hat{\theta}_2$ are estimators of θ_1 and θ_2, the most obvious estimator of R is the estimator obtained by forming the ratio of estimators, namely $\hat{R} = \hat{\theta}_1/\hat{\theta}_2$. Its expectation is $E\left[\hat{\theta}_1/\hat{\theta}_2\right] = E[\hat{\theta}_1]E[1/\hat{\theta}_2]$ (assuming independence of θ_1 and θ_2). Usually $E[1/\hat{\theta}_2] \neq 1/E[\hat{\theta}_2]$ and $\hat{R}$ is biased, even if $\hat{\theta}_1$ and $\hat{\theta}_2$ are both unbiased. Jensen's inequality implies that $E[1/\hat{\theta}_2] \geq 1/E[\hat{\theta}_2]$. $\hat{\theta}_1/\hat{\theta}_2$ over estimates the true ratio.

For example, if $\hat{\theta}_1$ and $\hat{\theta}_2$ are unbiased, independent and log-normally distributed and denoting by $\mu_.$ and $\sigma_.$ their log-mean and log-sd, we have:

$$
\begin{aligned}
R &= \theta_1/\theta_2 = \exp(\mu_{\hat{\theta}_1} + \sigma_{\hat{\theta}_1}^2/2)/\exp(\mu_{\hat{\theta}_2} + \sigma_{\hat{\theta}_2}^2/2) \\
&= \exp(\mu_{\hat{\theta}_1} - \mu_{\hat{\theta}_2} + (\sigma_{\hat{\theta}_1}^2 - \sigma_{\hat{\theta}_2}^2)/2)
\end{aligned}
\tag{1}
$$

whereas

$$
E[\hat{R}] = \exp(\mu_{\hat{\theta}_1} - \mu_{\hat{\theta}_2} + (\sigma_{\hat{\theta}_1}^2 + \sigma_{\hat{\theta}_2}^2)/2)
\tag{2}
$$

$\hat{R}$ over-estimates R by a multiplicative factor equal to $\exp(\sigma_{\hat{\theta}_2}^2)$.

Delta method

The Delta method is based on a Taylor series approximation of $\hat{R} = \hat{\theta}_1/\hat{\theta}_2$ around θ_1/θ_2 that is used to obtain estimates of the expectation and of the variance of $\hat{R}$ (Cox 1998; Casella and Berger 2002; Faraggi, Izikson, and Reiser 2003; Franz 2007; Hirschberg and Lye 2010). Assuming that $\hat{R}$ follows a normal distribution, the $1 - \alpha$ confidence interval is obtained as $\hat{R} \pm z \times$ s.d. where z is the upper $\alpha/2$ quantile of the standard normal distribution. The approximations made in this method are dubious and it is expected to behave poorly. Its main advantage are its conceptual simplicity and ease of implementation.

Fieller's method

Fieller's method (Fieller 1954) is based on the assumption that $(\hat{\theta}_1, \hat{\theta}_2)$ follows a bivariate Normal distribution For testing $\theta_1/\theta_2 = R_0$ (which amounts to testing $\theta_1 = R_0\theta_2$), the two-sided t-test based on $(\hat{\theta}_1 - R_0\hat{\theta}_2)/\text{Var}[\hat{\theta}_1 - R_0\hat{\theta}_2]^{1/2}$ is optimal in some sense. The rejection region for this test is the set of values r satisfying $(\hat{\theta}_1 - r\hat{\theta}_2) > t\text{Var}[\hat{\theta}_1 - r\hat{\theta}_2]^{1/2}$. Finding an explicit form for the confidence interval requires solving a quadratic equation in r. The confidence interval can be of the form (L, U), $(U, +\infty)$ $(0, U)$ or $(0, +\infty)$ depending on the number of solutions of the quadratic equation (Raftery and Schweder 1993; Buonaccorsi 2001; Franz 2007; Von Luxburg and Franz 2009; Hirschberg and Lye 2010)

Raftery and Schweder's Bayesian method

The method introduced by Raftery and Schweder (1993) is based on the assumption that the joint posterior distribution of $(\hat{\theta}_1, \hat{\theta}_2)$ is bivariate Normal. Assuming that estimates at hand $\hat{\theta}_1$ and $\hat{\theta}_2$ on one hand, and s_1 and s_2 on the other hand are the posterior means and variances of θ_1 and θ_2 and that θ_2 is positive (which holds most often for toxicological parameters), then
$P(R \leq r) = P(\theta_1 - r\theta_2 \leq 0) = \Phi(-(\hat{\theta}_1 - r\hat{\theta}_2)/Var[\hat{\theta}_1 - r\hat{\theta}_2|\text{Data}]^{1/2})$ where Φ denotes the cumulative distribution function (c.d.f.) of a standard normal distribution. The latter can be used to obtain numerically a Bayesian estimation interval for R. This method is based on an assumption similar to the one underlying Fieller's method (joint normality of estimators).

Bootstrap method

Choquet et al. (1999) proposed a method based on re-sampling with replacement the data that have been used to produce $\hat{\theta}_1$ and $\hat{\theta}_2$ (bootstrapping). We note that this approach is based on a biased estimate of R and that in a simulation study, Choquet et al. (1999) report that empirical coverage frequencies did not match nominal coverage probabilities. In addition, when estimates $\hat{\theta}_1$ and $\hat{\theta}_2$ have been retrieved from the scientific literature, the data that have been used to produce them are usually not available.

Likelihood ratio test method

If $\hat{\theta}_1$ and $\hat{\theta}_2$ have been obtained by maximizing a likelihood under a given model (see e.g. LD50 under a logistic regression, Faraggi, Izikson, and Reiser 2003), then a ratio of likelihood is asymptotically χ^2 distributed. This approach requires to know how many data points have been used to obtain $\hat{\theta}_1$ and $\hat{\theta}_2$ (number of degrees of freedom of the χ^2), an information usually not available when data have been extracted from the scientific literature.

Methods designed for the ratio of means of log-normal distributions

Chen and Zhou (2006) review five methods of constructing confidence intervals for the ratio of means of two log-normal distributions. (See also Wu et al. 2002; L. Jiang, Rekkas, and Wong 2014; Jafari and Abdollahnezhad 2017 for related material).

Likelihood based approaches

Chen and Zhou (2006) present three likelihood-based approach referred therein as the maximum likelihood approach, the signed log-likelihood ratio approach and the modified signed log-likelihood ratio approach. We present briefly the maximum likelihood approach.

Notation: $X = \log U$ with $U \sim N(\mu_1, \sigma_1^2)$ and $Y = \log V$ with $V \sim N(\mu_2, \sigma_2^2)$.

$\psi = \log[E(X)/E(Y)] = m_1 - m_2 + \sigma_1^2/2 - \sigma_2^2/2$ and $\theta = (\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)$

Steps

1. Compute $\hat{\theta} = (\hat{\mu}_1, \hat{\mu}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2)$, ML estimate of θ
2. $\hat{\psi} = \hat{\mu}_1 - \hat{\mu}_2 + \hat{\sigma}_1^2/2 - \hat{\sigma}_2^2/2$ is the ML estimate of ψ
3. Compute an estimate $\hat{\tau}$ of $\text{Var}[\hat{\psi}]^{1/2}$ (this estimate is a function of $\hat{\mu}_1, \hat{\mu}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2$ and of the sizes of the sample that have been used to obtain them)
4. $\hat{\psi}$ being asymptotically Normal, a $(1 - \alpha)$ confidence interval for ψ is $\hat{\psi} \pm z_{\alpha/2}\hat{\tau}$
5. A $(1 - \alpha)$ confidence interval for $E(X)/E(Y)$ is $\exp(\hat{\psi} \pm z_{\alpha/2}\hat{\tau})$.

This method requires the individual data and not only summary statistics.

Bootsrap approach

- Compute $\hat{\theta} = (\hat{\mu}_1, \hat{\mu}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2)$ as defined above

- (For j=1 to m)

    - Generate n_1 samples from a $N(\mu_1, \sigma_1^2)$ and n_2 samples from $N(\mu_2, \sigma_2^2)$
    - Calculate, using the bootstrap sample, estimate for ψ and τ as defined above denote these $\hat{\psi}_j$ and $\hat{\tau}_j$
    - Compute the test statistic $S_j = (\hat{\psi}_j - \hat{\psi})/\hat{\tau}_j$

- Find the $\alpha/2$ and $1 - \alpha/2$ quantiles of $(S_j)_{j=1,...,m}$.

A $(1 - \alpha)$ confidence interval for ψ is $[\hat{\psi} + S_{\alpha/2}\hat{\tau}; \hat{\psi} + S_{1-\alpha/2}\hat{\tau}]$

As noted by Chen and Zhou (2006) *"unlike some bootstrap approaches, the above does in fact make parametric assumptions. Specifically, it assumes that the data are lognormally distributed. The method does not however explicitly state or assume the nature of the statistic's distribution; instead, the distribution is explored computationally via the use of the bootstrap samples".*

This method requires the individual and not only summary statistics such as estimated means and variances.

Approximate log-normal method

1. It is assumed that $\hat{\theta}_i$ follows a log-normal distributions.

2. Denoting by $m_{\hat{\theta}_i}$ and $s_{\hat{\theta}_i}$ the mean and sd of $\hat{\theta}_i$ and by $\mu_{\hat{\theta}_i}$ and $\sigma_{\hat{\theta}_i}$ the mean and sd of $\ln(\hat{\theta}_i)$, we have

$$\mu_{\hat{\theta}_i} = \ln\left(\frac{m_{\hat{\theta}_i}}{\sqrt{1 + \frac{s_{\hat{\theta}_i}^2}{m_{\hat{\theta}_i}^2}}}\right) \text{ and } \sigma_{\hat{\theta}_i} = \sqrt{\ln\left(1 + \frac{s_{\hat{\theta}_i}^2}{m_{\hat{\theta}_i}^2}\right)}$$

3. We have only one replicate of $\hat{\theta}_i$, so the best estimate of its expectation is $\hat{\theta}_i$ itself. An estimate of the SD of $\hat{\theta}_i$ (the standard error) is available and denoted $\hat{s}_{\hat{\theta}_i}$

4. Substituting $\hat{\theta}_i$ to $\mu_{\hat{\theta}_i}$ and $\hat{s}_{\hat{\theta}_i}$ to $s_{\hat{\theta}_i}$ in equations above provide estimated parameters on the log scale.

5. Under the log-normal assumption for $\hat{\theta}_1$ and $\hat{\theta}_2$, $\hat{\theta}_1/\hat{\theta}_2$ is also log-normally distributed with log-mean and log-variance $\mu_{\hat{\theta}_1/\hat{\theta}_2} = \mu_{\hat{\theta}_1} - \mu_{\hat{\theta}_2}$ and $\sigma_{\hat{\theta}_1/\hat{\theta}_2}^2 = \sigma_{\hat{\theta}_1}^2 + \sigma_{\hat{\theta}_2}^2$ which can be estimated by combining parameters obtained at previous step.

6. One can then compute quantiles $q_{\hat{\theta}_1/\hat{\theta}_2}^{(\alpha/2)}$ and $q_{\hat{\theta}_1/\hat{\theta}_2}^{(1-\alpha/2)}$ for $\hat{\theta}_1/\hat{\theta}_2$.

A simulation study of coverage probability

Set-up

We simulate mortality data with a logistic regression as dose-response model. The probability of death for an individual exposed to a dose d is $1/(1 + exp(-(\beta_0 + \beta_1 d)))$. Under this model, the dose at which the probability of death is 50% (LD50) is $-\beta_0/\beta_1$. We attempt to mimic experimental conditions recommended in OECD guidelines for the testing of chemicals on honeybees: fives doses and at least three replicate test groups, each of ten bees. Under a given set of "experimental conditions", we estimate LD50 by maximum likelihood for each replicate. Then we average the various estimates and compute an estimate of the standard error as the standard deviation of the set of ML estimates divided by the square root of the number of replicates.

Results

We compute 95% confidence intervals with the Delta and the Fieller methods. The results reported in each row below are obtained each with 500 independent data-sets.

Table 1: Simulations with 3 batches of 10 individuals for each dose.

LD50.1	LD50.2	Ratio 1/2	Coverage Delta	Coverage Fieller
4.0	4.0	1.0	89.2	89.0
13.3	4.0	3.3	90.6	89.8
20.0	2.2	9.0	91.8	88.6

LD50.1	LD50.2	Ratio 1/2	Coverage Delta	Coverage Fieller
4.0	13.3	0.3	86.6	88.0
2.2	20.0	0.1	75.4	84.4

Table 2: Simulations with 10 batches of 10 individuals for each dose.

LD50.1	LD50.2	Ratio 1/2	Coverage Delta	Coverage Fieller
4.0	4.0	1.0	92.4	92.2
13.3	4.0	3.3	95.2	94.4
20.0	2.2	9.0	95.8	94.6
4.0	13.3	0.3	90.6	93.0
2.2	20.0	0.1	77.4	87.4

Conclusion

The Delta method and the Fieller method can be implemented all the time. Other methods can not be applied when only summary statistics are available. The Delta and Fieller methods are based on the assumptions that parameter estimates are Gaussian which is unlikely to make fully sense on toxicological parameter such as LD50, in particular when obtained with small sample size. Under this specific model and this limited set of simulations, one of the two methods achieves a better accuracy on the other one. There is no guarantee that the actual coverage probability matches the nominal coverage.

Appendix: R code for simulation study of coverage

```r
#######################
## functions

## Simulating data
sim.logist = function(nrep,# nb of replications
                       doses,# vector of doses
                       batch.size, # nb of individuals in a batch
                       beta0,beta1# intercept and slope in linear predictor
                       )
{
  ## simulate "mortality" data under dose-response logistic regression model
  ## nrep replicates
  ## return ML estimates of LD50 parameters for each replicate and takes the average
  LD50.tilde = rep(NA,nrep)
  for(irep in 1:nrep)
  {
  responses = rbinom(n = length(doses),
                     size = rep(batch.size,length(doses)),
                     prob = 1/(1+exp(-(beta0+beta1*doses))))
  ##  estimating parameter
  res.glm = glm(formula=cbind(responses,batch.size-responses)~doses,
                family=binomial(link='logit'),
                data = data.frame(doses=doses,responses=responses))
  LD50.tilde[irep] = (- coef(res.glm)[1] / coef(res.glm)[2])
  }
  LD50.hat = mean(LD50.tilde)
  sd.LD50.hat = sd(LD50.tilde)/sqrt(nrep)
  res = c(LD50.hat,sd.LD50.hat,-beta0/beta1)
  names(res) = c('LD50.hat','sd.LD50.hat','LD50.true')
  return(res)
}
## END function simulating data



#######################
## CI Delta method
CI.Delta = function(theta1.hat, sd1.hat,
                    theta2.hat, sd2.hat,# estimate and estimated sd of estimator
                    alpha # theoretical coverage (1-alpha)
                    )
{
```

```r
  ## CI Delta method
  ## Hirschberg, J., and J. Lye. 2010.
  ## "A Geometric Comparison of the Delta and Fieller Confidence Intervals."
  ## The American Statistician 64 (3): 234-41, Eq. 3-6
  ##  We assume here independence between theta1.hat and theta2.hat
  z = qnorm(p=1-alpha/2)
  R.hat = theta1.hat/theta2.hat
  sd.R.hat = sqrt(sd1.hat^2 + R.hat^2 * sd2.hat^2) /theta2.hat
  L = R.hat - z*sd.R.hat
  U = R.hat + z*sd.R.hat
  res = c(R.hat,L,U)
  names(res) = c('R.hat','L','U')
  return(res)
}
## END CI Delta method


##############
## CI Fieller (approximated, cf comments below)
CI.Fieller = function(theta1.hat, sd1.hat,
                      theta2.hat, sd2.hat,# estimate and estimated sd of estimator
                      alpha # theoretical coverage (1-alpha)
)
{
  ## CI Fieller
  # Raftery, A E, and T. Schweder. 1993.
  ## "Inference About the Ratio of Two Parameters, with Application to
  # Whale Censusing." The American Statistician 47 (4): 259-64. Eq. 1-2
  # We assume here independence between theta1.hat and theta2.hat
  # the size of the sample used to obtain the estimates
  # and their sd is assumed to be unknown
  # we substitute the quantile of a Normal(0,1) to that of a St(?)
  # that should be used in principle
  z = qnorm(p=1-alpha/2)
  if( (theta1.hat/sd1.hat <= z) & (theta2.hat/sd2.hat <= z) )
    {
    L = 0 ; U = Inf
    }else
    {
      rad  = (theta1.hat^2 * sd2.hat^2 + theta2.hat^2 * sd1.hat^2 -
              z^2 * sd1.hat^2 * sd2.hat^2 )
    L = (theta1.hat * theta2.hat - z*sqrt(rad) )/( theta2.hat^2 - z * sd2.hat^2 )
    U = (theta1.hat * theta2.hat + z*sqrt(rad) )/( theta2.hat^2 - z * sd2.hat^2 )
      if((theta1.hat/sd1.hat > z)  & (theta2.hat/sd2.hat <= z) ){L = U ; U = Inf}
```

```r
      if((theta1.hat/sd1.hat <= z)  & (theta2.hat/sd2.hat > z) ){L = 0}
    }
  R.hat = theta1.hat/theta2.hat
  res = c(R.hat,L,U)
  names(res) = c('R.hat','L','U')
  return(res)
}
## END CI Fieller


# ############
# ## CI  Raftery & Schweder Bayesian
# CI.RSB = function(theta1.hat, sd1.hat,
#                   theta2.hat, sd2.hat,# estimate and estimated sd of estimator
#                   alpha # theoretical coverage (1-alpha)
# ## END CI.RSB


###########
## CI Approx LN
CI.LN =  function(theta1.hat, sd1.hat,
                     theta2.hat, sd2.hat,# estimate and estimated sd of estimator
                     alpha # theoretical coverage (1-alpha)
)
{
  # parameters of underlying Normal variables
  mu1 = log(theta1.hat/sqrt(1+(sd1.hat/theta1.hat)^2))
  mu2 = log(theta2.hat/sqrt(1+(sd2.hat/theta2.hat)^2))
  sigma1 = sqrt(log(1+(sd1.hat/theta1.hat)^2))
  sigma2 = sqrt(log(1+(sd2.hat/theta2.hat)^2))
  ## parameters of  Normal underlying ratio
  mu.R =mu1 -mu2
  sigma.R =  sqrt(sigma1^2 + sigma2^2)
  L = qlnorm(meanlog=mu.R, sdlog=sigma.R, p=alpha/2)
  U = qlnorm(meanlog=mu.R, sdlog=sigma.R, p=1-alpha/2)
  R.hat = theta1.hat/theta2.hat
  res = c(R.hat,L,U)
  names(res) = c('R.hat','L','U')
  return(res)
}
## END CI Approx LN




##
```

```r
# dat1  = sim.logist(nrep=30,
#                    doses=c(2,4,8,10,20),
#                    batch.size=5,
#                    beta0 =-2,
#                    beta1 = .5) ; dat1
# dat2 = sim.logist(nrep=30,
#                    doses=c(2,4,8,10,20),
#                    batch.size=5,
#                    beta0 =-2,
#                    beta1 = .15) ; dat2
# CI.Delta(dat1[1],dat1[2],
#          dat2[1],dat2[2],
#          alpha=.05)
# CI.Fieller(dat1[1],dat1[2],
#            dat2[1],dat2[2],
#            alpha=.05)
# dat1['LD50.true']/dat2['LD50.true']


#################
## Coverage for Delta and Fieller
coverage = function(doses,nrep,bach.size,nsets,
                    beta0.1,beta1.1,beta0.2,beta1.2,alpha)
{
  ## computing coverage for Delta and Fieller's
  cover.Delta = cover.Fieller = 0
  ## Loop on datasets
  for(iset in 1:nsets)
  {
    # print(100*iset/nsets)
    dat1  = sim.logist(nrep=nrep,
                       doses=doses,
                       batch.size=batch.size,
                       beta0 = beta0.1,
                       beta1 = beta1.1)
    dat2 = sim.logist(nrep=nrep,
                      doses=doses,
                      batch.size=batch.size,
                      beta0 = beta0.2,
                      beta1 = beta1.2)
    R.true = dat1['LD50.true']/dat2['LD50.true']
    CI.D = CI.Delta(dat1[1],dat1[2],
                    dat2[1],dat2[2],
```

```r
                      alpha=alpha)[2:3]
    CI.F = CI.Fieller(dat1[1],dat1[2],
                      dat2[1],dat2[2],
                      alpha=alpha)
    if((R.true >= CI.D['L']) & (R.true <= CI.D['U']))
      { cover.Delta = cover.Delta+1 }
    if((R.true >= CI.F['L']) & (R.true <= CI.F['U']))
      { cover.Fieller = cover.Fieller+1 }
  }
  cover.Delta = cover.Delta/nsets
  cover.Fieller = cover.Fieller/nsets
  res = c(cover.Delta,cover.Fieller)
  names(res) = c('cover.Delta','cover.Fieller')
  return(res)
}
## END coverage


## END OF FILE
```

References

Beyene, J., and R. Moineddin. 2005. "Methods for Confidence Interval Estimation of a Ratio Parameter with Application to Location Quotients." *BMC Medical Research Methodology* 5 (1). BioMed Central: 32.

Buonaccorsi, J.P. 2001. "Fieller's Theorem." *Encyclopedia of Environmetrics*. Wiley Online Library.

Casella, G., and R.L. Berger. 2002. *Statistical Inference*. Vol. 2. Duxbury Pacific Grove, CA.

Chen, Y.H., and X.H. Zhou. 2006. "Interval Estimates for the Ratio and Difference of Two Lognormal Means." *Statistics in Medicine* 25 (23). Wiley Online Library: 4099–4113.

Choquet, D., P. L'ecuyer, and C. Léger. 1999. "Bootstrap Confidence Intervals for Ratios of Expectations." *ACM Transactions on Modeling and Computer Simulation (TOMACS)* 9 (4). ACM: 326–48.

Cox, C. 1998. "Delta Method." *Encyclopedia of Biostatistics*. Wiley Online Library.

Faraggi, D., P. Izikson, and B. Reiser. 2003. "Confidence Intervals for the 50 Per Cent Response Dose." *Statistics in Medicine* 22 (12). Wiley Online Library: 1977–88.

Fieller, E.C. 1954. "Some Problems in Interval Estimation." *Journal of the Royal Statistical Society. Series B (Methodological)*. JSTOR, 175–85.

Franz, V.H. 2007. "Ratios: A Short Guide to Confidence Limits and Proper Use.(2007)." *ArXiv Preprint ArXiv:0710.2024*.

Hirschberg, J., and J. Lye. 2010. "A Geometric Comparison of the Delta and Fieller Confidence Intervals." *The American Statistician* 64 (3). Taylor & Francis: 234–41.

Jafari, A. A., and K. Abdollahnezhad. 2017. "Testing the Equality Means of Several Log-Normal Distributions." *Communications in Statistics-Simulation and Computation* 46 (3). Taylor & Francis: 2311–20.

Jiang, L., M. Rekkas, and A. Wong. 2014. "Comparing the Means of Two Log-Normal Distributions: A Likelihood Approach." *Journal of Statistical and Econometric Methods* 3 (1). SCIENPRESS Ltd: 137–52.

Land, C.E. 1971. "Confidence Intervals for Linear Functions of the Normal Mean and Variance." *The Annals of Mathematical Statistics*. JSTOR, 1187–1205.

Olsson, U. 2005. "Confidence Intervals for the Mean of a Log-Normal Distribution." *Journal of Statistics Education* 13 (1).

Raftery, A E, and T. Schweder. 1993. "Inference About the Ratio of Two Parameters, with Application to Whale Censusing." *The American Statistician* 47 (4). Taylor & Francis: 259–64.

Von Luxburg, U., and V.H. Franz. 2009. "A Geometric Approach to Confidence Sets for Ratios: Fieller's Theorem, Generalizations and Bootstrap." *Statistica Sinica*. JSTOR, 1095–1117.

Wu, J., G. Jiang, ACM Wong, and X. Sun. 2002. "Likelihood Analysis for the Ratio of Means of Two Independent Log-Normal Distributions." *Biometrics* 58 (2). Wiley Online Library: 463–69.