

The Neyman-Pearson lemma for confidence distributions

Tore Schweder

University of Oslo

Box 1095 Blindern

0317 Oslo, Norway

tore.schweder@econ.uio.no

Nils Lid Hjort

University of Oslo

Box 1053 Blindern

0316 Oslo, Norway

nils@math.uio.no

Confidence distributions

Confidence intervals and p-values are the main formats of statistical reporting in the frequentist tradition. R.A. Fisher's fiducial distribution is less often used. Jerzy Neyman did not like fiducial distributions since in his view they are not proper probability distributions. They do not refer to a sample space and a stochastic experiment. Fisher, on the other hand, did not like Neyman's interpretation of the confidence interval as a stochastic interval with a controlled probability of covering the true value of the parameter.

A confidence distribution is a stochastic distribution (it depends on the stochastic data) with the property that its quantiles span confidence intervals. If \mathbf{y} is the one-dimensional parameter and $C(\mathbf{y})$ is the cumulative distribution function, then $(C^{-1}(\mathbf{a}), C^{-1}(\mathbf{b}))$ is a confidence interval with degree $\mathbf{b} - \mathbf{a}$. In other words, if \mathbf{y}_0 is the true value of the parameter, $C(\mathbf{y}_0)$ is a stochastic variable with a uniform distribution. Confidence distributions are usually constructed via pivots. For example, when $\hat{\mathbf{y}}$ and s are the mean and standard deviation in a normal sample of size n , $(\mathbf{y} - \hat{\mathbf{y}})/(s/\sqrt{n})$ is a t-distributed pivot. Thus, with T_{n-1} being the cumulative t-distribution, $C(\mathbf{y}) = T_{n-1}(\sqrt{n}(\mathbf{y} - \hat{\mathbf{y}})/s)$ is a cumulative confidence distribution for \mathbf{y} . Before data are obtained, C is stochastic. After the data are obtained, observed values of $\hat{\mathbf{y}}$ and s are inserted, and C is a cumulative distribution function for \mathbf{y} . Confidence intervals and p-values are related, and they are both obtained from the confidence distribution. The p-value of $H_0: \mathbf{y} \leq \mathbf{y}_0$ vs. $H_1: \mathbf{y} > \mathbf{y}_0$ is simply $C(\mathbf{y}_0)$. Schweder and Hjort (2001) discuss confidence distributions in more detail.

Efron (1993) does also discuss confidence distribution. In simple models, the confidence distribution is identical to the fiducial distribution. Since Neyman objected to the fiducial

distribution, he would probably not have liked the concept of confidence distribution. And Fisher, on the other hand, disliked Neyman's concept of confidence, and would probably also have hated the confidence distribution.

Nevertheless, the confidence distribution is a very useful tool in statistical reporting, and should be a competitive frequentist analogue of the Bayesian posterior distribution. After the data have been observed, the confidence distribution is a distribution of confidence and not of probability in the frequentist interpretation. Before data are observed, it is a stochastic element with a probability distribution.

The Neyman-Pearson lemma

The power function is the primary performance characteristic in hypothesis testing. In confidence interval inference, one correspondingly uses the power function defined as the probability of the confidence interval not to include an alternative value of the parameter. What should the performance characteristic be for a confidence distribution?

Since the confidence distribution by design is unbiased (the confidence median is a median-unbiased point estimator), it is reasonable to judge the confidence distribution by its spread. The smaller it's spread the better. Most measures of spread in the distribution C around \mathbf{y}_0 can be represented by a functional

$$\mathbf{g}(C) = \int_{-\infty}^0 C(\mathbf{y} + \mathbf{y}_0) \mathbf{g}(d\mathbf{y}) + \int_0^{\infty} (1 - C(\mathbf{y} + \mathbf{y}_0)) \mathbf{g}(d\mathbf{y}),$$

where \mathbf{g} is a measure on the right hand side of the equation. If this measure is chosen as the Lebesgue measure, $\mathbf{g}(C)$ is the mean absolute variation around \mathbf{y}_0 with respect to the distribution C . Other natural choices are: $\mathbf{g}(d\mathbf{y}) = |\mathbf{y}| d\mathbf{y}$ making $\mathbf{g}(C)$ the mean square variation; and point measure at $\mathbf{y}_1 < 0$ making $\mathbf{g}(C)$ the p-value when testing $H_0 : \mathbf{y} \leq \mathbf{y}_0 + \mathbf{y}_1$.

As in other frequentist inference, the performance should be measured *ex ante* data. Before data are observed, the confidence distribution is stochastic, and any measure of spread, $\mathbf{g}(C)$, is thus *ex ante* a stochastic variable.

The Neyman-Pearson lemma for confidence distributions (Schweder and Hjort 2001)

If the confidence distribution C is a function of the data through a one-dimensional statistic that is sufficient and in which the likelihood ratio is increasing, then C is uniformly most powerful in the sense that for any other confidence distribution C' and for any choice of spread functional, $\mathbf{g}(C)$ is *ex ante* stochastically less than $\mathbf{g}(C')$ when the true parameter of the probability distribution is \mathbf{y}_0 .

The proof is quite simple, and is in the reference. An extension to multi-parameter exponential models is also given there.

REFERENCES

- Efron, B. 1998. R.A. Fisher in the 21st century (with discussion). *Statistical Science*, 13:95-122.
- Schweder, T and Hjort, N.L. 2001. Confidence and likelihood. *Scandinavian Journal of Statistics*. Accepted.