DRAFT 0.2

Glen Cowan
16 June, 2011

Quick study of conditional coverage

I have done a quick study of the conditional coverage of "Diagonal Line" (DL), Power-Constrained Limit (PCL) and CLs/Bayesian intervals for our standard problem of a $x \sim \text{Gauss}(\mu, \sigma)$ with $\mu \geq 0$. Figure 1(a) shows the conditional coverage probability given that x is observed less than a given constant c for the unconstrained one-sided upper limit. This is simply

$$P(\mu_{\rm up} > \mu | x < c) = \frac{1 - \alpha - \Phi\left(\frac{\mu - c}{\sigma}\right)}{1 - \Phi\left(\frac{\mu - c}{\sigma}\right)} \,. \tag{1}$$

This essentially the same as the unconstrained limit used as the starting point of PCL, except there the point $\mu = 0$ is never excluded.

Figure 1(b) shows the corresponding curves for the Power-Constrained Limit $\mu_{\rm up}^*$, for which the conditional coverage is

$$P(\mu_{\rm up}^* > \mu | x < c) = \begin{cases} 1 & \mu < \mu_{\rm min}, \\ \frac{1 - \alpha - \Phi\left(\frac{\mu - c}{\sigma}\right)}{1 - \Phi\left(\frac{\mu - c}{\sigma}\right)} & \text{otherwise.} \end{cases} \tag{2}$$

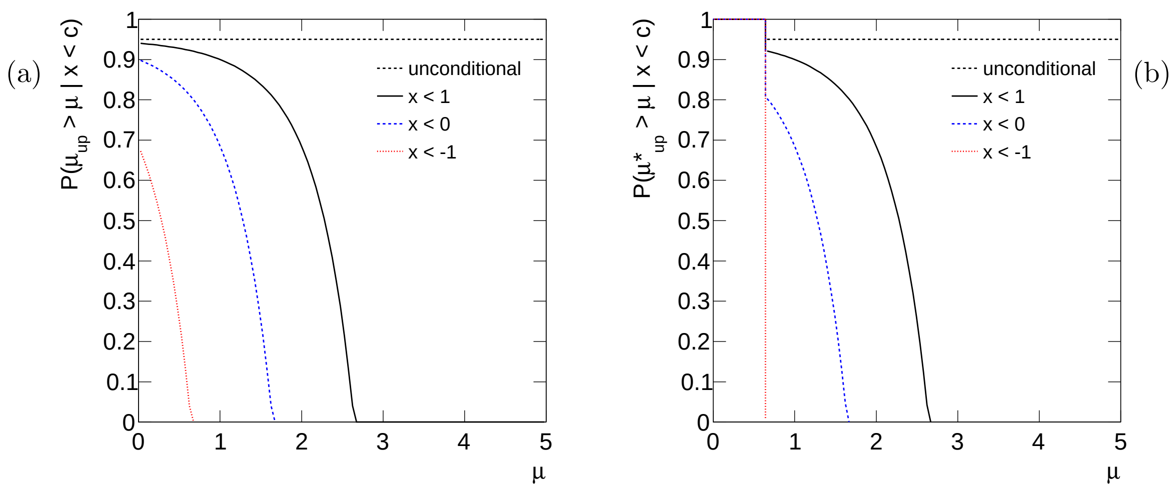

Figure 1: The conditional coverage probability given $x < c$ versus μ for several values of c for the (a) unconstrained and (b) power-constrained upper limit.

As can be seen in Fig. 1(a), the curves are bounded away from $1 - \alpha$ by an amount that depends on c, which corresponds to the negatively biased relevant subsets [1]. If the constraint $\mu > 0$ did not exist, then the curves would all come back up to $1 - \alpha$.

The fact that the curves drop down below $1 - \alpha$ for higher μ is not in itself a pathology, however, and should happen for any upper limit. That is, if one considers all data outcomes, then for a fixed non-zero μ the coverage probability is $1 - \alpha$. So if one selects a subset of data outcomes where x fluctuates lower than some amount, then the coverage probability for a given μ will in general be less.

The problem with the negatively biased relevant subsets, however, is that it allows one to have $P(\mu_{\rm up} > \mu | x < c)$ for any value of μ, and thus it is possible to set up a betting game where one could win, regardless of μ, by betting that the interval does not cover. PCL does not allow for this type of betting game, however, as its conditional coverage probabilities extend up to 100% for $\mu < \mu_{\rm min}$, as shown in Fig. 1(b).

It has been argued in [1] that since the unconditional coverage of the PCL limit is 95% for $\mu > \mu_{\rm min}$, then the conditional coverage should also be at least $1 - \alpha$ for some value of μ in this range. But to make a winning betting strategy in this way, one would have to condition the bet on the true and unknown value of μ, e.g., using an omniscient referee who says whether the bet is taken.

The unconditional coverage for PCL is 95% for $\mu > \mu_{\rm min}$, and this fact is pointed out and advertised as a useful feature of the procedure. But the full correct statement of the coverage properties of PCL is that the unconditional coverage is at least 95% for all μ, and overcovers for $\mu < \mu_{\rm min}$. And if the full range of μ is taken into account, then there are no negatively biased relevant subsets. Once it becomes allowed to condition on μ, it should always be possible to design a winning betting strategy, since curves of $P(\mu_{\rm up} > \mu | x < c)$ versus μ will always fall for any upper limit.

One could avoid negatively biased relevant subsets by only having a single value of μ for which the coverage is greater than $1 - \alpha$. This is effectively present in the unconstrained limit, since this never excludes $\mu = 0$. In PCL it could also be achieved, for example, by taking the minimum power arbitrarily close to α, which gives $\mu_{\rm min}$ arbitrarily small. But in PCL one intentionally choose the minimum power sufficiently large so that this does not happen.

The CLs upper limit [3] for the Gaussian problem is

$$\mu_{\rm up,CLs} = x + \sigma \Phi^{-1} \left(1 - \alpha \Phi(x/\sigma) \right) \ . \tag{3}$$

The conditional coverage probability for CLs is shown in Fig. 2(a), and for the (full) Feldman-Cousins interval in Fig. 2(b).

The CLs conditional probabilities are much higher, as one would expect because of the unconditional overcoverage. Nevertheless the curves all fall as a function of μ. So here as well if one were to select a given subset of μ values, one can easily find negatively biased subsets.

References

[1] Robert D. Cousins, Draft note on Negatively Biased Relevant Subsets and references therein.

[2] Glen Cowan, Kyle Cranmer, Eilam Gross and Ofer Vitells, *Power-Constrained Limits*, arXiv:1105.3166.

[3] T. Junk, Nucl. Instrum. Methods Phys. Res., Sec. A **434**, 435 (1999); A.L. Read, J. Phys. G **28**, 2693 (2002).

(a) 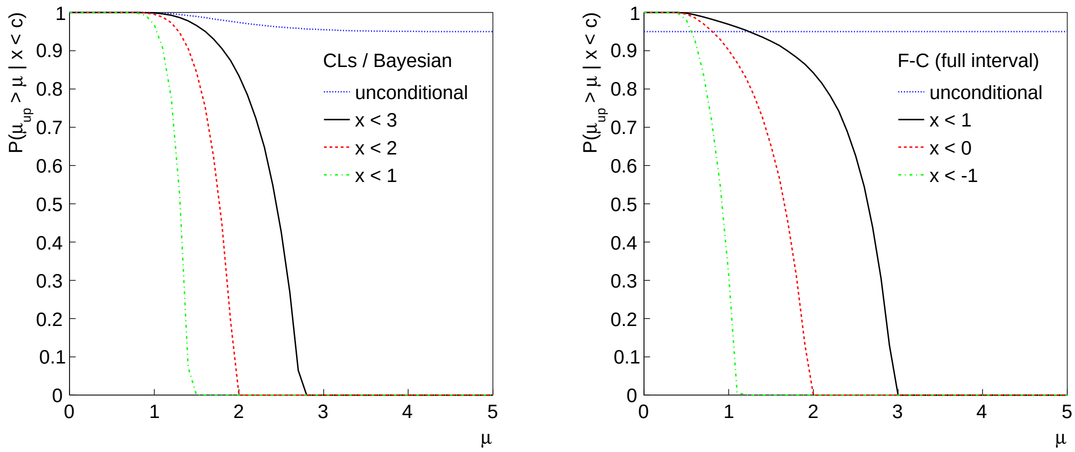 (b)

Figure 2: The conditional coverage probability given $x < c$ versus μ for several values of c for the (a) CLs/Bayesian upper limit and (b) the Feldman-Cousins (full) confidence interval.