Better Confidence Intervals for Quantiles

Michael Höhle

 m_hoehle

Department of Mathematics, Stockholm University

useR! Conference
Brussels, 04-07 July 2017

Motivation: Flint Water Crisis

Title 40 CFR Part 141

- **National Primary Drinking Water Regulations:**

> **§141.80 General requirements.**
>
> …
>
> (c) *Lead and copper action levels.* (1) The lead action level is exceeded if the concentration of lead in more than 10 percent of tap water samples collected during any monitoring period conducted in accordance with §141.86 is greater than 0.015 mg/L (*i.e.*, if the "90th percentile" lead level is greater than 0.015 mg/L).

- **Histogram of the Flint lead monitoring sample:**

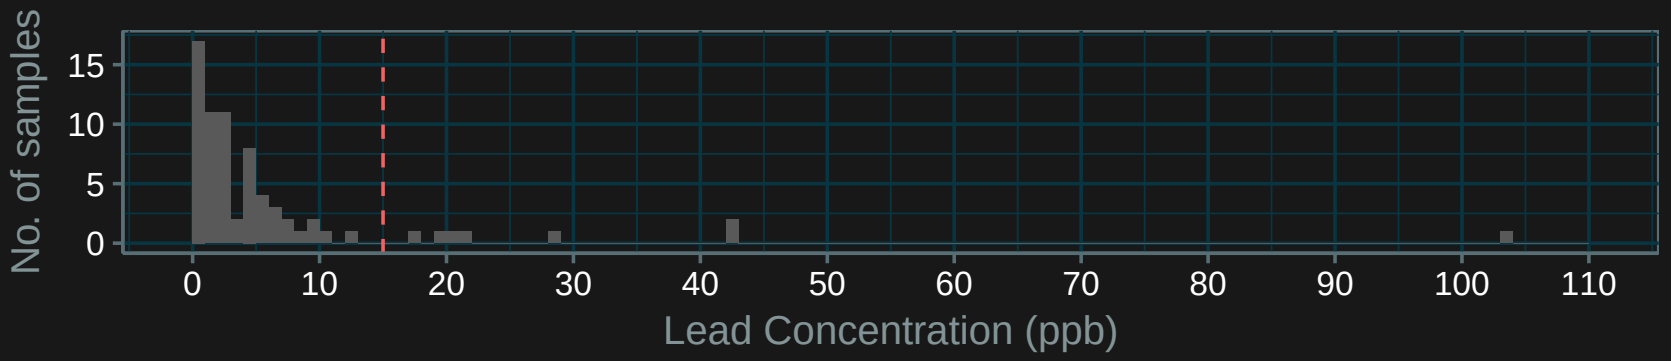

Estimating the Population Quantile

- **The duality between the proportion and the quantile only holds for `type=1` of R's `quantile` function.**

- **The `type=1` estimator for the 90% quantile is**

$$\hat{Q}(0.9) = x_{(\lceil 0.9 \cdot n \rceil)},$$

where $x_{(1)} \leq \ldots \leq x_{(n)}$ is the ordered sample.

The Monitoring Sample

```
sort(flint)
```

```
##  [1]   0   0   0   0   0   0   0   0   0   0   0   0   0   1   1   1   1
## [18]   2   2   2   2   2   2   2   2   2   2   2   3   3   3   3   3   3
## [35]   3   3   3   3   3   4   4   5   5   5   5   5   5   5   5   6   6
## [52]   6   6   7   7   7   8   8   9  10  10  11  13  18  20  21  22  29
## [69]  43  43 104
```

```
(n <- length(flint))
```

```
## [1] 71
```

```
quantile(flint, probs=0.9, type=1)
```

```
## 90%
##  18
```

```
mean(flint > 15)
```

```
## [1] 0.1126761
```

Statistical Test (1)

- **Consider the proportion π of non-conforming sites in the water system's sampling pool.**
- **Decide on compliance of π with standard based on a statistical test (at a given significance lvl.):**
  - $H_0 : \pi \geq 0.1$ **(compliance)**
  - $H_1 : \pi < 0.1$ **(compliance)**

Statistical Test (2)

- Testing is done by, e.g, an exact binomial test

$$p\text{-value} = \sum_{k=0}^{x} \binom{n}{k} 0.1^k 0.9^{n-k},$$

where x is the observed number of non-conforming sites.

Rejection Region - Proportion

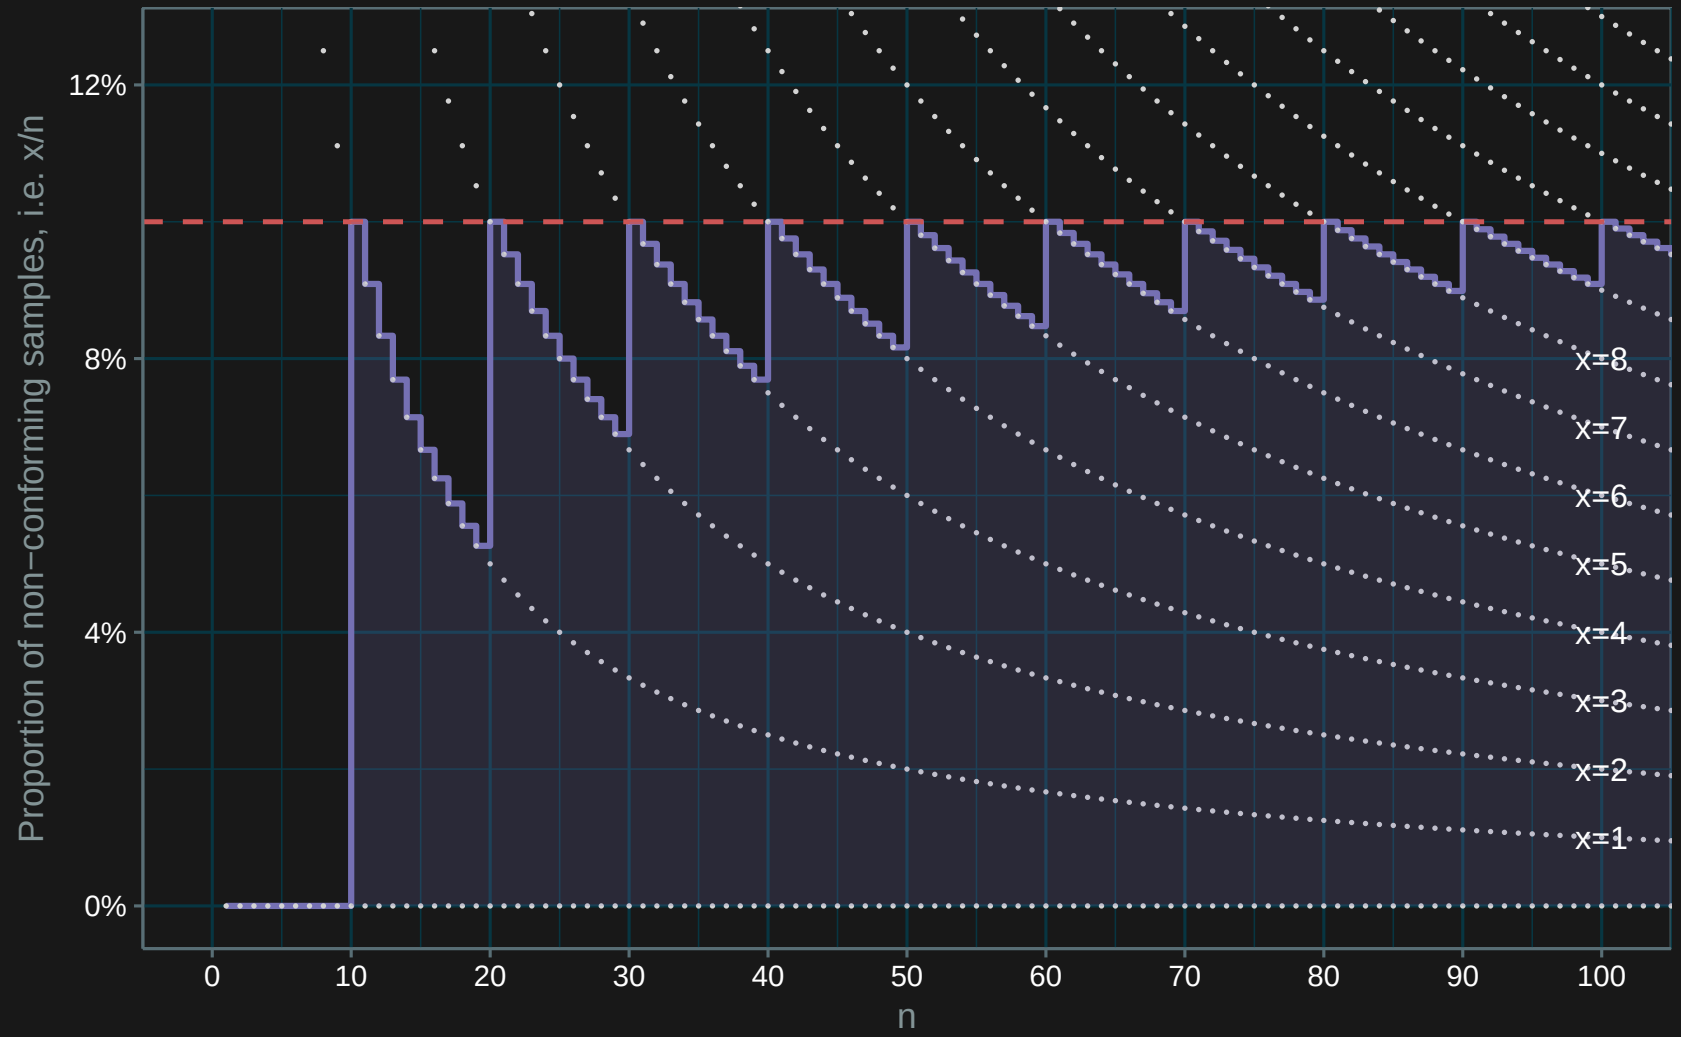

Rejection Region - Binomial Test

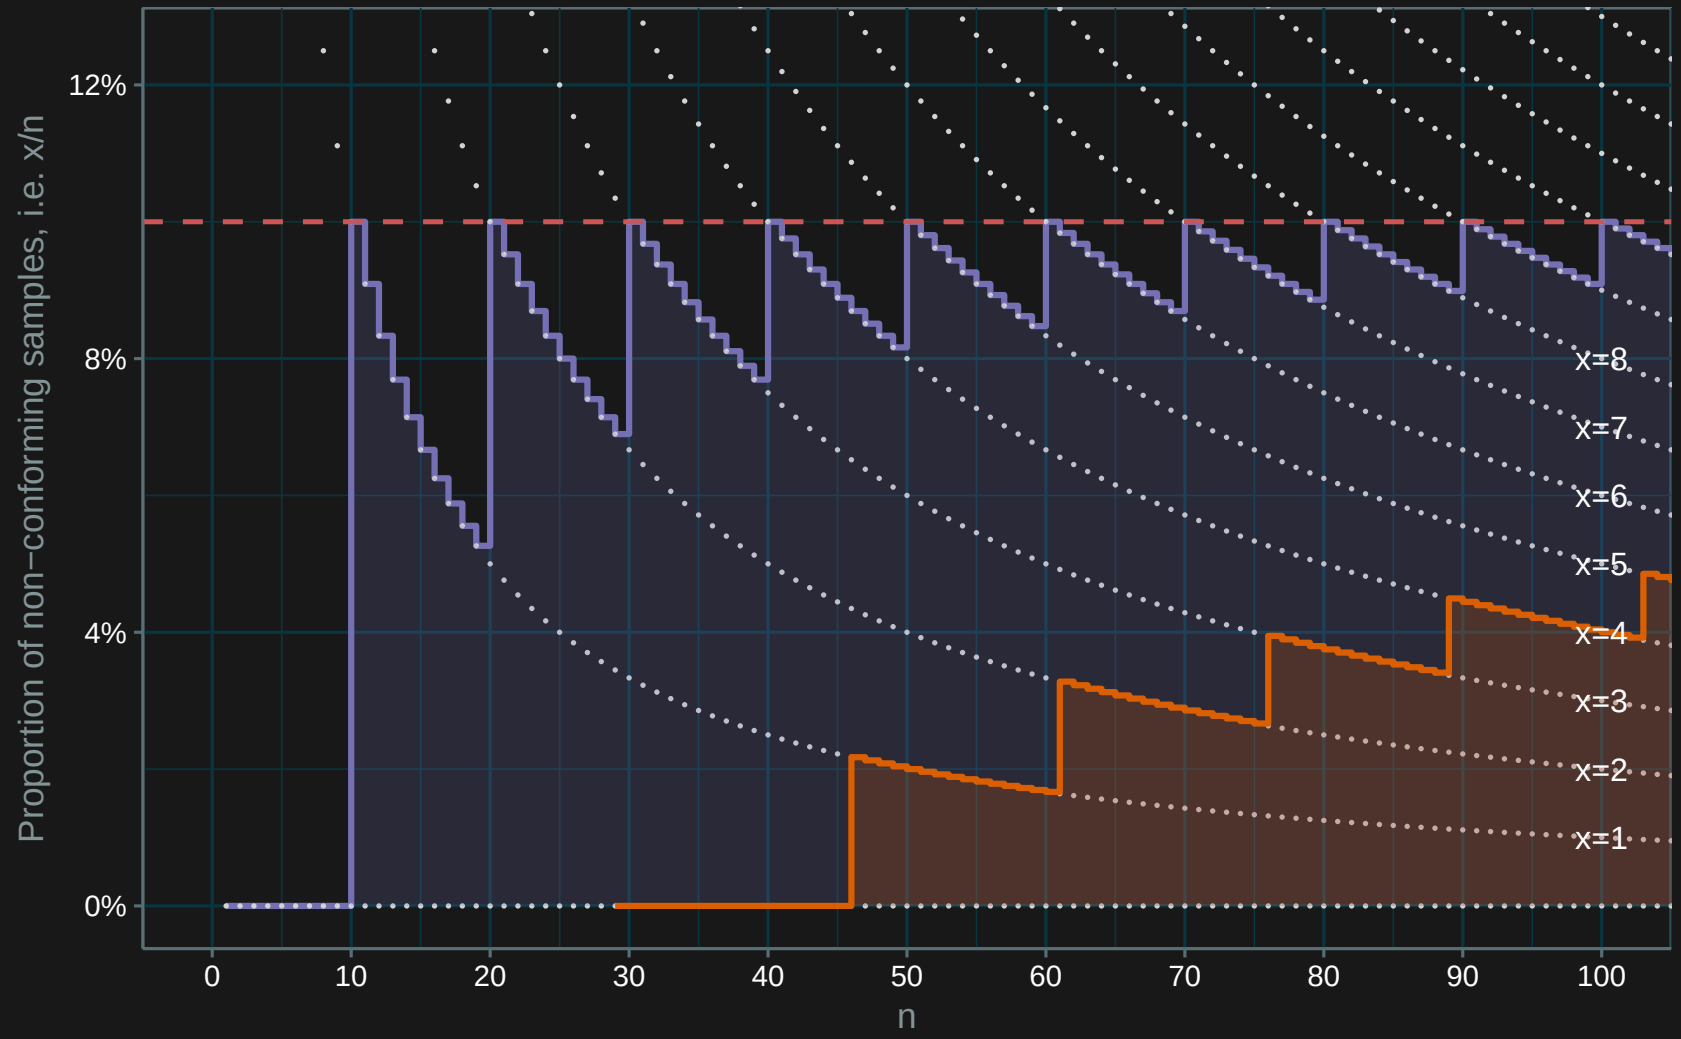

CIs for Quantiles (1)

- **Interest is in a one-sided upper 95% CI $(-\infty, x_{(e)})$ such that**

$$P\left(Q(0.9) \leq x_{(e)}\right) \geq 0.95.$$

- **Duality between such a CI and a one-sided test:**
  - $H_0 : Q(0.9) \geq$ **15 ppb vs.**
  - $H_1 : Q(0.9) <$ **15 ppb.**

CIs for Quantiles (2)

- One can show that for $1 \leq r \leq n$ we have

$$P(Q(0.9) \leq x_{(r)}) = P(\text{at least } n - r + 1 \text{ observations are larger than or equal to } Q(0.9))$$
$$= \sum_{k=n-r+1}^{n} \binom{n}{k} 0.1^k 0.9^{n-k}$$

- We can thus find e as the smallest index making the above expression ≥ 0.95:

```
qbinom(0.95, size=n, prob=0.9) + 1
```

```
## [1] 69
```

CIs for Quantiles (3)

- Note: $n - e$ is also the largest no. of non-conforming samples where we still reject H_0 with the binomial test $\rightarrow$ another duality!
- Such exact confidence intervals are, however, known to be conservative.
- As an improvement Nyblom (1992) suggests to interpolate the 2 neighbouring order statistics.

90% CI for the 90% Quantile

```
structure(sort(flint)[64:n],names=64:n)
##  64  65  66  67  68  69  70  71
##  18  20  21  22  29  43  43 104

quantile_confint_exact(x=flint, p=0.9,conf.level=0.9)
## [1] 10 43

quantile_confint_nyblom(x=flint, p=0.9,conf.level=0.9)
## [1] 10.00000 31.68225

quantile_confint_boot(x=flint, p=0.9,conf.level=0.9,
                                R=9999, type=1)
## [1]  9 22
```

Simulation Study

- **We conducted a small simulation study comparing with:**
  - `jmuOutlier::quantileCI` **(v 1.1)**
  - `envStats::eqnpar` **(v 2.1.1)**
  - `asht::quantileTest` **(v 0.6)**
- **The results confirmed the improved coverage of the Nyblom (1992) approach.**

Summary

- **Simulation study also revealed flaws in one of the R packages!**
- **R package `quantileCI` on github**
- **Details: Blog posts @ 'Theory meets practice'**

`http://staff.math.su.se/hoehle/blog/`