

CONFIDENCE INTERVALS FOR QUANTILES

M. K. Nakayama, New Jersey Institute of Technology, U.S.A., marvin@njit.edu

J.M. Calvin, New Jersey Institute of Technology, Newark, New Jersey, USA

The p -quantile of a random variable is the constant for which exactly p of the mass of its distribution lies to the left of the quantile; e.g., the median is the 0.5-quantile. Quantiles are widely used to assess risk. For example, a project manager may want to determine a time T such that the project has a 95% chance of completing by T , which is the 0.95-quantile. In finance, where a quantile is known as a value-at-risk, analysts frequently measure risk with the 0.99-quantile of a portfolio's loss; thus, there is a 1% chance that the loss will be greater than this value. For complex stochastic models, analytically computing a quantile usually is not possible, so simulation is employed. In addition to providing a point estimate for a quantile, we also want to measure the simulation estimate's error, and this is typically done by giving a confidence interval (CI) for the quantile. Indeed, the U.S. Nuclear Regulatory Commission requires that licensees of nuclear power plants demonstrate compliance using a "95/95 criterion," which entails ensuring (with 95% confidence) that a 0.95-quantile lies below a mandated limit. In this talk we present some approaches for constructing CIs for a quantile estimated via simulation. Standardized time series (STS) is a class of methods originally proposed by Schruben (1983) to construct a CI for the steady-state mean of a stochastic process. We consider applying STS techniques to construct a CI for a quantile.