Lesson 4. Sample Mean & Variance, Confidence Intervals

1 Overview

- In JaamSim, the simulations we create generate performance measures for 1 simulation run
  - e.g., the average delay for the BaristaQ in the Nimitz Coffee Bar
- The observed average delay can and typically does differ between simulation runs
- The average delay is a random variable
  - Uncertain quantity before the simulation run
  - Depends on interarrival and service times
- Can we estimate the distribution of the average delay?
  - Let's focus on estimating the mean and variance of this distribution

2 The experiment

- Run the simulation n times
- Compute performance measure (e.g. average delay) for each simulation run
  (obtaining n observations of the performance measure)
- Use the n observations to estimate the mean of the delay

3 After the experiment: observed sample mean and sample variance

- Let $X_1, \dots, X_n$ be independent and identically distributed (i.i.d.) random variables with unknown mean μ and variance σ^2
- Let $x_1, \dots, x_n$ be the observed values of $X_1, \dots, X_n$, respectively
  - Think of X_i as the average delay in the ith simulation run before the experiment
  - Think of x_i as the observed average delay in the ith simulation run after the experiment
  - Since the simulation runs replicate the same system, $X_1, \dots, X_n$ should be identically distributed
- We want to estimate μ
- The **observed sample mean** is
- The **observed sample variance** is
- The **standard error** is
- We estimate μ using the the observed sample mean

- We estimate σ^2 using the observed sample variance
- These are **point estimates** for μ and σ^2, respectively
- The standard error is a measure of the accuracy of the estimate of μ
- Why should we estimate μ and σ^2 this way?

4 Before the experiment: sample mean and sample variance

- The **sample mean** is
    - The sample mean $\overline{X}$ is a random variable: before the experiment, it is an uncertain quantity
    - $E[\overline{X}] = \mu$:
    - $\text{Var}(\overline{X}) = \sigma^2/n$:
- The **sample variance** is
    - The sample variance is also a random variable: before the experiment, it is an uncertain quantity
- The sample mean is an **unbiased estimator** of μ, and the sample variance is an **unbiased estimator** of σ^2: that is,
    - Intuitively, this indicates that using the observed sample mean to estimate μ and the observed sample variance to estimate σ^2 is not a bad idea

5 How good is the observed sample mean as an estimate?

- Is the observed sample mean $\overline{x}$ "close" to μ?
- Suppose $\overline{X}$ is normally distributed
    - This is true if $X_1, \dots, X_n$ are normally distributed
    - This is approximately true by the Central Limit Theorem if $n \geq 30$

- Then the $(1-\alpha)100\%$ **confidence interval for μ** is

  - This is an **interval estimate** for μ
  - $t_{\alpha/2,n-1}$ can be computed in R with `qt((1 - α/2), n - 1)`
  - The t-distribution with $n-1$ degrees of freedom ≈ standard Normal distribution when $n \geq 30$

- Interpretation of a confidence interval:

  - Sample mean $\overline{X}$ and sample standard deviation S^2 are random variables
  - Every experiment, we get different observed sample mean $\overline{x}$ and observed sample variance s^2

  ⇒ Every experiment, we get a different confidence interval

  - After running the experiment many times, $(1-\alpha)100\%$ of the resulting confidence intervals will contain the actual mean μ
  - We say that "we are $(1-\alpha)100\%$ <u>confident</u> that the mean μ lies within the confidence interval"
  - <u>Wrong interpretation</u>: "The mean μ lies within the given confidence interval with $(1-\alpha)100\%$ <u>probability</u>"

- Smaller confidence interval ⇒ more accurate estimate of μ

Example 1. Suppose an estimate of μ within 0.1 was desired at a confidence level of 95%. We perform a "warm-up" experiment of $n = 30$ simulation runs to compute an observed sample variance s^2, which is found to be 3.2. How many simulations runs are needed to obtain this estimate of μ?

- Example 1 is helpful in understanding the formula for computing confidence intervals but is not useful in practice when $n = 10000$ is often used
- What's wrong with the method in Example 1?