Discrete-Event Simulation:

A First Course

Section 8.1: Interval Estimation

  0-13-142917-5

Section 8.1: Interval Estimation

Theorem (Central Limit Theorem)

If X_1, X_2, . . ., X_n is an iid sequence of RVs with

- *common mean* μ
- *common standard deviation* σ

and if $\bar{X}$ *is the (sample) mean of these RVs*

$$\bar{X} = \frac{1}{n}\sum_{i=1}^{n} X_i$$

then $\bar{X}$ *approaches a Normal*$(\mu, \sigma/\sqrt{n})$ *RV as* $n \to \infty$

  0-13-142917-5

Sample Mean Distribution

- Choose one of the random variate generators in `rvgs` to generate a sequence of random variable samples with fixed sample size $n > 1$
- With the n-point samples indexed $j = 1, 2, \ldots$, the corresponding sample mean $\bar{X}_j$ and sample standard deviation s_j can be calculated using Algorithm 4.1.1

$$\underbrace{x_1, x_2, \ldots, x_n}_{\bar{x}_1, s_1}, \underbrace{x_{n+1}, x_{n+2}, \ldots, x_{2n}}_{\bar{x}_2, s_2}, \underbrace{x_{2n+1}, x_{2n+2}, \ldots, x_{3n}}_{\bar{x}_3, s_3}, x_{3n+1}, \ldots$$

- A continuous-data histogram can be created using program `cdh`

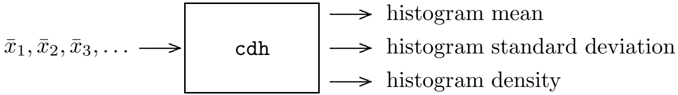

Properties of Sample Mean Histogram

- Independent of n,
  - the histogram mean is approximately μ
  - the histogram standard deviation is approximately $\sigma/\sqrt{n}$
- If n is sufficiently large,
  - the histogram density approximates the $Normal(\mu, \sigma/\sqrt{n})$ pdf

Example 8.1.2: 10000 *n*-point *Exponential*(μ) samples

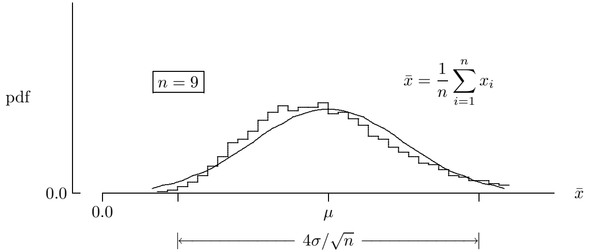

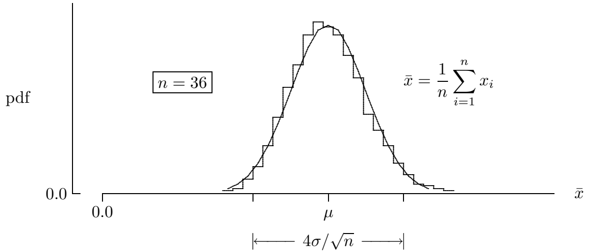

Example 8.1.2

- The histogram mean and standard deviation are approximately μ and $\sigma/\sqrt{n}$
- The histogram density corresponding to the 36-point sample means is closely matched by the pdf of a $\mathit{Normal}(\mu, \sigma/\sqrt{n})$ RV
  - For $\mathit{Exponential}(\mu)$ samples, $n = 36$ is large enough for the sample mean to be approximately $\mathit{Normal}(\mu, \sigma/\sqrt{n})$
- The histogram density corresponding to the 9-point sample means matches relatively well, but with a skew to the left
  - $n = 9$ is not large enough

  0-13-142917-5

More on Example 8.1.2

- Essentially all of the sample means are within an interval of width of $4\sigma/\sqrt{n}$ centered about μ
- Because $\sigma/\sqrt{n} \to 0$ as $n \to \infty$, if n is large, all the sample means will be close to μ
- In general:
  - The accuracy of the *Normal*$(\mu, \sigma/\sqrt{n})$ pdf approximation is dependent on the shape of a fixed population pdf
  - If the samples are drawn from a population with
    - a highly asymmetric pdf (like the *Exponential*(μ) pdf): n may need to be as large as 30 or more for good fit
    - a pdf symmetric about the mean (like the *Uniform*(a, b) pdf): n as small as 10 or less may produce a good fit

Standardized Sample Mean Distribution

- We can *standardize* the sample means $\bar{x}_1, \bar{x}_2, \bar{x}_3, \ldots$ by subtracting μ and dividing the result by $\sigma/\sqrt{n}$ to form the standardized sample means $z_1, z_2, z_3, \ldots$ defined by

$$z_j = \frac{\bar{x}_j - \mu}{\sigma/\sqrt{n}} \qquad j = 1, 2, 3, \ldots$$

- Generate a continuous-data histogram for the standardized sample means by program `cdh`

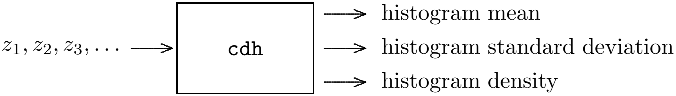

Properties of Standardized Sample Mean Histogram

- Independent of n,
  - the histogram mean is approximately 0
  - the histogram standard deviation is approximately 1
- If n is sufficiently large,
  - the histogram density approximates the $Normal(0, 1)$ pdf

Example 8.1.4

- The sample means from Example 8.1.2 were standardized

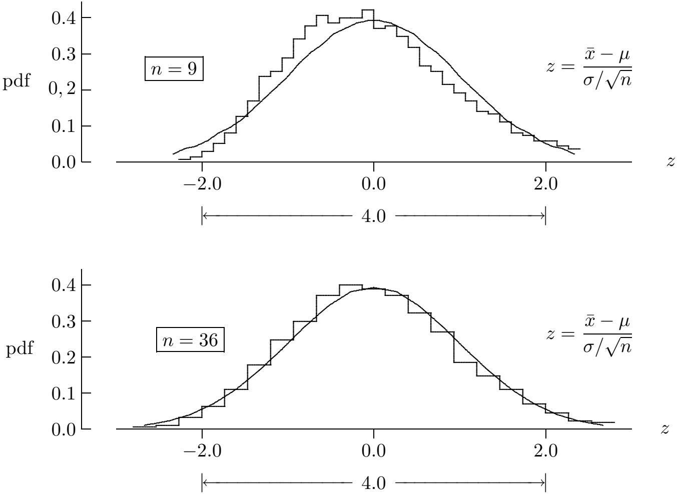

Properties of the Histogram in Example 8.1.4

- The histogram mean and standard deviation are approximately 0.0 and 1.0 respectively
- The histogram density corresponding to the 36-point sample means matches the pdf of a $\mathit{Normal}(0, 1)$ random variable almost exactly
- The histogram density corresponding to the 9-point sample means matches the pdf of a $\mathit{Normal}(0, 1)$ random variable, but not as well

  0-13-142917-5

t-Statistic Distribution

- Want to replace *population* standard deviation σ with *sample* standard deviation s_j in standardization equation

$$z_j = \frac{\bar{x}_j - \mu}{\sigma/\sqrt{n}} \qquad j = 1, 2, 3, \ldots$$

- **Definition 8.1.1**
  - Each sample mean $\bar{x}_j$ is a *point estimate* of μ
  - Each sample variance s_j^2 is a *point estimate* of σ^2
  - Each sample standard deviation s_j is a *point estimate* of σ

  0-13-142917-5

Removing Bias

- The sample mean is an *unbiased* point estimate of μ
  - The mean of $\bar{x}_1, \bar{x}_2, \bar{x}_3 \ldots$ will converge to μ
- The sample variance is a *biased* point estimate of σ^2
  - The mean of $s_1^2, s_2^2, s_3^3, \ldots$ will converge to $(n-1)\sigma^2/n$, not σ^2
- To remove this $(n-1)/n$ bias, it is conventional to multiply the sample variance by a *bias correction* $n/(n-1)$
- The point estimate of $\sigma/\sqrt{n}$ is

$$\frac{\left(\sqrt{\frac{n}{n-1}}\right) s_j}{\sqrt{n}} = \frac{s_j}{\sqrt{n-1}}$$

Example 8.1.5

- Calculate the t-statistic

$$t_j = \frac{\bar{x}_j - \mu}{s_j / \sqrt{n-1}} \qquad j = 1, 2, 3, \ldots$$

- Generate a continuous-data histogram using `cdh`

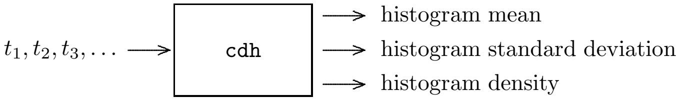

  0-13-142917-5

Properties of t-statistic Histogram

- If $n > 2$, the histogram mean is approximately 0
- If $n > 3$, the histogram standard deviation is approximately $\sqrt{(n-1)/(n-3)}$
- If n is sufficiently large, the histogram density approximates the pdf of a $Student(n-1)$ random variable

Example 8.1.6

- Generate t-statistics from Example 8.1.2

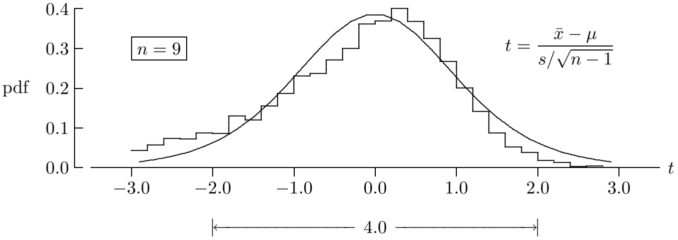

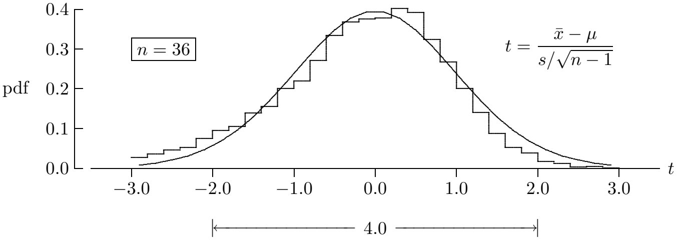

Properties of the Histogram in Example 8.1.6

- The histogram mean and standard deviation are approximately 0.0 and $\sqrt{(n-1)/(n-3)} \simeq 1.0$ respectively
- The histogram density corresponding to the 36-point sample means matches the pdf of a *Student*(35) RV relatively well
- The histogram density corresponding to the 9-point sample means matches the pdf of a *Student*(8) RV, but not as well

Interval Estimation

Theorem (8.1.2)

If $x_1, x_2, \ldots, x_n$ is an (independent) random sample from a "source" of data with unknown mean μ, if $\bar{x}$ and s are the mean and standard deviation of this sample, and if n is large, it is approximately true that

$$t = \frac{\bar{x} - \mu}{s/\sqrt{n-1}}$$

is a Student$(n-1)$ random variate

- Theorem 8.1.2 provides the justification for estimating an *interval* that is likely to contain the mean μ
- As $n \to \infty$, the *Student*$(n-1)$ distribution becomes indistinguishable from *Normal*$(0, 1)$

Interval Estimation (2)

- Suppose
  - T is a *Student*$(n-1)$ random variable
  - α is a "confidence parameter" with $0.0 < \alpha < 1.0$

Then there exists a corresponding positive real number t^*

$$\Pr(-t^* \leq T \leq t^*) = 1 - \alpha$$

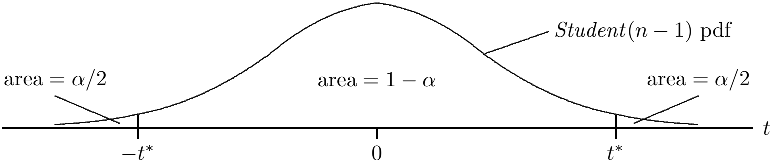

Interval Estimation (3)

- Suppose μ is *unknown*. Since $t \approx Student(n-1)$,

$$-t^* \leq \frac{\bar{x} - \mu}{s/\sqrt{n-1}} \leq t^*$$

will be approximately true with probability $1 - \alpha$

- Right inequality: $\frac{\bar{x} - \mu}{s/\sqrt{n-1}} \leq t^* \iff \bar{x} - \frac{t^* s}{\sqrt{n-1}} \leq \mu$
- Left inequality: $-t^* \leq \frac{\bar{x} - \mu}{s/\sqrt{n-1}} \iff \mu \leq \bar{x} + \frac{t^* s}{\sqrt{n-1}}$
- So, with probability $1 - \alpha$ (approximately),

$$\bar{x} - \frac{t^* s}{\sqrt{n-1}} \leq \mu \leq \bar{x} + \frac{t^* s}{\sqrt{n-1}}$$

Theorem 8.1.3

Theorem (8.1.3)

If

- $x_1, x_2, \ldots, x_n$ *is an independent random sample from a "source" of data with unknown mean* μ
- $\bar{x}$ *and* s *are the sample mean and sample standard deviation*
- n *is large*

Then, given a confidence parameter α *with* $0.0 < \alpha < 1.0$, *there exists an associated positive real number* t^* *such that*

$$\Pr\left(\bar{x} - \frac{t^* s}{\sqrt{n-1}} \leq \mu \leq \bar{x} + \frac{t^* s}{\sqrt{n-1}}\right) \cong 1 - \alpha$$

Example 8.1.7

- If $\alpha = 0.05$, we are 95% confident that μ lies *somewhere* between

$$\bar{x} - \frac{t^* s}{\sqrt{n-1}} \quad \text{and} \quad \bar{x} + \frac{t^* s}{\sqrt{n-1}}$$

- For a fixed sample size n and level of confidence $1 - \alpha$, use `rvms` to determine $t^* = \texttt{idfStudent}(n-1, 1-\alpha/2)$
- For example, if $n = 30$ and $\alpha = 0.05$, then $t^* = \texttt{idfStudent}(29, 0.975) \simeq 2.045$

Definition 8.1.2

- The interval defined by the two endpoints

$$\bar{x} \pm \frac{t^* s}{\sqrt{n-1}}$$

  is a $(1-\alpha) \times 100\%$ confidence *interval estimate* for μ

- $(1-\alpha)$ is the *level of confidence* associated with this interval estimate and t^* is the *critical value* of t

Algorithm 8.1.1

Algorithm 8.1.1

To calculate an *interval estimate* for the unknown mean μ of the *population* from which a random sample $x_1, x_2, x_3, \ldots, x_n$ was drawn:

- Pick a level of confidence $1 - \alpha$ (typically $\alpha = 0.05$)
- Calculate the sample mean $\bar{x}$ and standard deviation s (use Algorithm 4.1.1)
- Calculate the critical value $t^* = \texttt{idfStudent}(n - 1, 1 - \alpha/2)$
- Calculate the interval endpoints

$$\bar{x} \pm \frac{t^* s}{\sqrt{n - 1}}$$

If n is sufficiently large, then you are $(1 - \alpha) \times 100\%$ confident that the mean μ lies within the interval. The midpoint of the interval is $\bar{x}$.

Example 8.1.8

- The random sample of size $n = 10$:

1.051	6.438	2.646	0.805	1.505
0.546	2.281	2.822	0.414	1.307

  is drawn from a population with unknown mean μ
- $\bar{x} = 1.982$ and $s = 1.690$
- To calculate a 90% confidence interval estimate:
  - Determine $t^* = \texttt{idfStudent}(9, 0.95) \simeq 1.833$
  - Interval: $1.982 \pm (1.833)(1.690/\sqrt{9}) = 1.982 \pm 1.032$
- We are approximately 90% confident that μ is between 0.950 and 3.014

Example 8.1.8, ctd.

- To calculate a 95% confidence interval estimate:
  - Determine $t^* = \texttt{idfStudent}(9, 0.975) \simeq 2.262$
  - Interval: $1.982 \pm (2.262)(1.690/\sqrt{9}) \quad = \quad 1.982 \pm 1.274$
- We are approximately 95% confident that μ is between 0.708 and 3.256
- To calculate a 99% confidence interval estimate:
  - Determine $t^* = \texttt{idfStudent}(9, 0.995) \simeq 3.250$
  - Interval: $1.982 \pm (3.250)(1.690/\sqrt{9}) \quad = \quad 1.982 \pm 1.832$
- We are approximately 99% confident that μ is between 0.150 and 3.814
- Note: $n = 10$ is *not* large

Tradeoff - Confidence Versus Sample Size

- For a fixed sample size
  - More confidence can be achieved *only* at the expense of a larger interval
  - A smaller interval can be achieved *only* at the expense of less confidence
- The only way to make the interval smaller without lessening the level of confidence is to increase the sample size
- Good news: with simulation, we can collect more data
- Bad news: interval size decreases with $\sqrt{n}$, not n

How Much More Data Is Enough?

- How large should n be to achieve an interval estimate $\bar{x} \pm w$ where w is *user-specified*?
- Answer: Use Algorithm 4.1.1 with Algorithm 8.1.1 to iteratively collect data until a specified interval width is achieved
- Note: if n is large then t^* is essentially independent of n

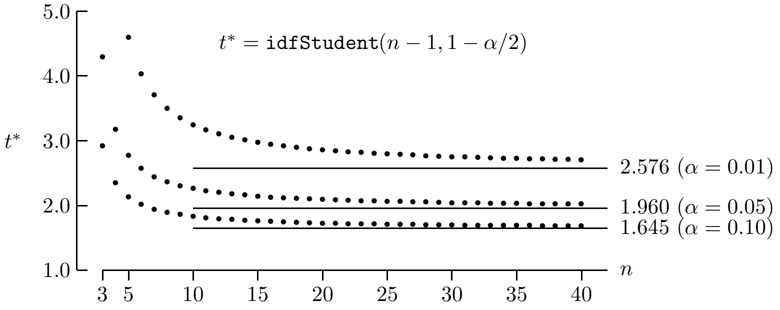

Asymptotic Value of t^*

- The asymptotic (large n) value of t^* is

$$t^*_\infty = \lim_{n\to\infty} \texttt{idfStudent}(n-1, 1-\alpha/2) = \texttt{idfNormal}(0.0, 1.0, 1-\alpha/2)$$

- Unless α is very close to 0.0, if $n > 40$, the asymptotic value t^*_∞ can be used
- If $n > 40$ and wish to construct a 95% confidence interval estimate, $t^*_\infty = 1.960$ can be used in Algorithm 8.1.1

Example 8.1.9

- Given a reasonable guess for s and a user-specified *half-width* parameter w, if t^*_∞ is used in place of t^*,

  n can be determined by solving $w = \dfrac{t^* s}{\sqrt{n-1}}$ for n:

$$n = \left\lfloor \left(\frac{t^*_\infty s}{w} \right)^2 \right\rfloor + 1$$

  provided $n > 40$

- For example, if $s = 3.0$ and want to estimate μ with 95% confidence to within ± 0.5, a value of $n = 139$ should be used

Example 8.1.10

- If a reasonable guess for s is not available, w can be specified as a proportion of s thereby eliminating s from the previous equation
- For example, if w is 10% of s and 95% confidence is desired, $n = 385$ should be used to estimate μ to within $\pm w$

  0-13-142917-5

Program Estimate

- Program `estimate` automates the interval estimation process
- A typical application: estimate the value of an unknown population mean μ by using n replications to generate an *independent* random variate sample $x_1, x_2, \ldots, x_n$
- Function `Generate()` represents a discrete-event or Monte Carlo simulation program that returns a random variate output x

Using the Generate Method

```
for (i = 1; i <= n; i++)
    x_i = Generate();
return x_1, x_2, ..., x_n;
```

- Given a level of confidence $1 - \alpha$, program `estimate` can be used with $x_1, x_2, \ldots, x_n$ to compute an interval estimate for μ

Algorithm 8.1.2

Algorithm 8.1.2

Given an interval half-width w and level of confidence $1-\alpha$, the algorithm computes the interval estimate $\bar{x} \pm w$

```
t = idfNormal(0.0, 1.0, 1-α/2); /* t*∞ */
x = Generate();
n = 1; v = 0.0; x̄ = x;
while ((n<40) or (t*sqrt(v/n) > w * sqrt(n-1)){
    x = Generate();
    n++;
    d = x - x̄;
    v = v + d * d * (n - 1) / n;
    x̄ = x̄ + d / n;
}
return n, x̄;
```

- It is important to appreciate the need for sample *independence* in Algorithms 8.1.1 and 8.1.2

The meaning of confidence

Incorrect:

- "For this 95% confidence interval, the probability that μ is within this interval is 0.95"
- Why incorrect?
  - μ is *not* a random variable; it is constant (but unknown)
  - The *interval endpoints* are random

Correct:

- "If I create many 95% confidence intervals, approximately 95% of them should contain μ"

  0-13-142917-5

Example 8.1.11

- 100 samples of size $n = 9$ drawn from *Normal*$(6, 3)$ population
- For each sample, construct a 95% confidence interval
- 95 intervals contain $\mu = 6$
- Three intervals "too low", two intervals "too high"