ST793 notes: Ch 9. Monte Carlo Study

1. Introduction Monte Carlo methods

The name *Monte Carlo* was first used in early-mid 1940's by scientists (most important ones are Stanislaw Ulam and John von Neumann) at Los Alamos National Lab, while working on nuclear weapons projects. An early variant of the technique dates since 1930's and is related to an unpublished work of the physicist Enrico Fermi (the creator of the world's first nuclear reactor). His work was later referred to as *Monte Carlo calculations*. The development of modern digital computers played a key role in the advancement of the methods.

Monte Carlo was used to describe a class of mathematical methods that allows to account for undertainty in quantitative analysis and decision making, as well as for numerical integration. Ulam described how Monte Carlo method can be used to estimate the chance of success in playing the card game solitaire: program a computer to randomize lists representing the 52 cards of a deck, prepare lists representing the different piles, and then simulate the playing of the game to completion. The computer gambling can be viewed as a faithful simulation of the real random process, namely, the card shuffling. This approach is an easy alternative to the often tedious way of mathematical calculation of such success probability, even in cases when this is based on elementary probability theory.

Why the name *Monte Carlo*?

At the time the method was developed, it was required to keep it secret; 'Monte Carlo' was given as a code name. The name was suggested by Nicholas Metropolis (a colleague of von Neumann and Ulam) based on the relation between randomness and gambling, and the connection to the Monte Carlo Casino in Monaco, where Ulam's uncle would borrow money from relatives to gamble.

What do we use Monte Carlo methods for?

Monte Carlo methods are used to:

- estimate the sampling distribution of an estimator
- estimate the mean and variance of an estimator
- estimate the expected length and actual coverage probability of a confidence interval
- estimate the percentiles of a test statistic
- estimate the percentiles of a non-standard limiting distribution

While Monte Carlo estimation plays a huge role in Bayesian analysis, in how it is used to estimate posterior distribution in a Bayesian framework, we will not discuss this aspect in this course.

In this chapter we discuss the following:

1 Examples of Monte Carlo studies

2 Principles of Monte Carlo studies

3 Importance of sample size in Monte Carlo studies

4 Analysis. Presenting the results.

This allows us to grasp what Monte Carlo experiment is and how to use it in our research. For a comprehensive study on this topic, see DP Kroese, T Talmare, ZI Botev (2011) "Handbook of Monte Carlo Methods". Other books on this topic: MH Kalos and PA Whitlock (2008), "Monte Carlo Methods"; RY Rubinstein and DP Kroese (2017) "Simulation and the Monte Carlo Method" (Wiley Series in Probability and Statistics)

2. Examples

a. Mean estimator

$\{Y_1, \ldots, Y_n\} \sim f(y)$ and $T = T(Y_1, \ldots, Y_n) = \sum_{i=1}^n Y_i/n$

So far: If n is large and f is well behaved (has finite first two moments) then

$$T \sim AN(\mu, \frac{1}{n}\sigma^2)$$

where μ and σ^2 are the mean and variance of f.

Interest: Performance of $T(Y_1, \ldots, Y_n)$ for various sample sizes. How is this affected by the underlying distribution?

Approach: Assess performance for *increasing* sample sizes (say $n = 30$, $n = 100$, $n = 1000$; WHY?) and different underlying distributions f (what kind? WHY?) by using a *Monte Carlo study.*

Set underlying distribution, say $N(0, 1)$, and sample size n, say $n = 30$. Then:

- generate $\{y_1, \ldots, y_{30}\}$ from $N(0, 1)$
- calculate $t = T(y_1, \ldots, y_{30})$

Repeat these 2 steps for a large number of times, denote it N. The values of T are

$$\{t_1, \ldots, t_N\}$$

The distribution of the estimator $T = T(Y_1, \ldots, Y_n)$ based on a sample of size n is approximated by the empirical distribution of $\{t_1, \ldots, t_N\}$, when data are drawn from f.

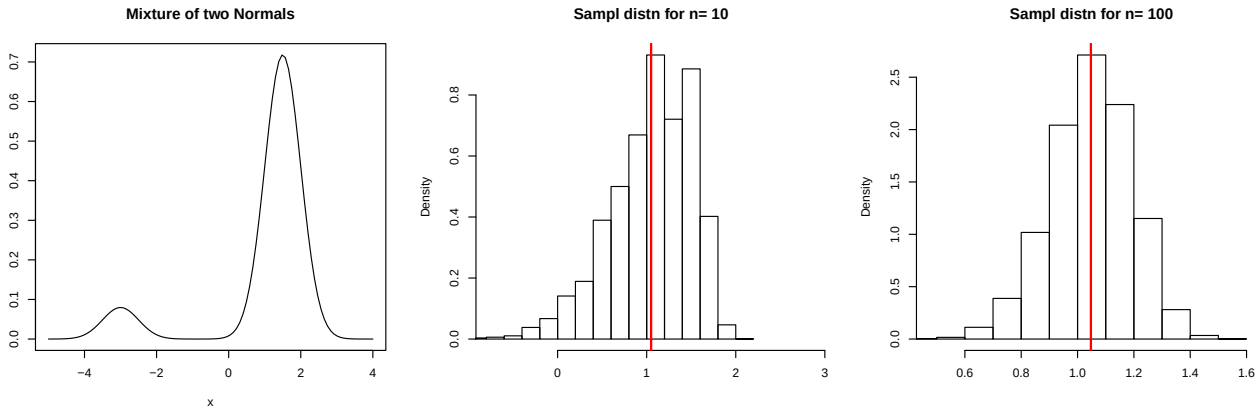

In class: Think of other measures you can use to evaluate the performance of the estimator T in estimating the parameter μ.

b. Confidence interval

Focus on: confidence interval for some parameter, say θ of an underlying distribution f. Suppose the $(1-\alpha)\%$ confidence interval has lower and upper bounds $Lo = L(Y_1, \ldots, Y_n)$ and $Up = U(Y_1, \ldots, Y_n)$, using a sample $\{Y_1, \ldots, Y_n\}$.

So far: if the sample size n is large, and the underlying distribution f is well behaved then

$$Pr\left(\theta \in (Lo, Up)\right) \approx 1 - \alpha$$

Interest: Performance of the confidence interval method for various sample sizes, when data come form different distributions.

Approach: Assess performance of the method for *increasing* sample sizes (say $n = 30$, $n = 100$, $n = 1000$; WHY?) and different underlying distributions f (what kind? WHY?) by using a *Monte Carlo study*.

Fix confidence level (say 95%). Set underlying distribution, say $N(0, 1)$, and sample size n, say $n = 30$. Then:

- generate $\{y_1, \ldots, y_{30}\}$ from $N(0, 1)$;
- calculate $lo = L(y_1, \ldots, y_{30})$ and $up = U(y_1, \ldots, y_{30})$;
- define $\delta = 1$ if $\theta \in (lo, up)$ and $\delta = 0$ otherwise.

Repeat these 3 steps for a large number of times, denote it N. Calculate the actual coverage (AC)) probability of the 95% confidence interval method by averaging all the δ- values obtained:

$$AC = \sum_{b=1}^{N} \delta_b / N$$

In class: What other measure/s could you use to evaluate the performance of this confidence interval construction?

c. Comparison among multiple estimators

Suppose that we are interested in estimating the location of a distribution, using a sample of size n. Consider three different estimators for location: sample mean, trimmed mean and sample median. The material that we studied so far is informative in that if the distribution is well behaved then we know what the asymptotic properties of these estimators are (i.e. consistency and distribution when the sample size n is very large.)

In practice, n is finite! How do we know if for the n at hand, the distribution of the estimator is approximated well by the asymptotic distribution.

We will use Monte Carlo simulation to get a sense of the distribution of these estimators for various sample sizes.

Setting: Suppose $Y_1, \ldots, Y_n \sim f(y)$ where $f(y)$ is some well behaved density function and n is the sample size. The estimators of interest are:

$$
\begin{aligned}
T_1 &= \bar{Y} &&= \frac{Y_1 + \ldots + Y_n}{n} \\
T_2 &= \bar{Y}_{.20} &&= \text{ trim 20\% from each end of the ordered sample and average the remaining data} \\
T_3 &= \tilde{Y} &&= \text{sample median, previously denoted as } \hat{\eta}_{0.5}
\end{aligned}
$$

Clearly T_1, T_2, T_3 are random (they are functions of the random Y_i). Their randomness is derermined by the randomness of Y_i's. Thus to get a sense of their distribution, we draw many samples $\{y_1, \ldots, y_n\}$ from the underlying distribution $f(\cdot)$ and for each such sample, we calculate the estimators value, say t_1 (corresponding to the estimator T_1). Then the empirican distribution of all the t_1's

We use Monte Carlo to generate many pseudo-random datasets $Y_1, \ldots, Y_n$ from $f(y)$, and for each such dataset, we calculate the three statistics. The distribution of T_1 across many datasets is called the sampling distribution of T_1. In the following we will compare the sampling distribution of the there statistics. Additionally we will assess

- the bias of the estimators in estimating the true mean
- the variance of the estimators (i.e. the standard deviation of the sampling distribution)

We are interested in how the distribution is affected by the following factors

- sample size n. Consider $n = 10$, $n = 50$ and $n = 100$
- underlying distribution $f(y)$: Normal$(0, 1)$, t_5 and Laplace$(0, 1)$. Laplace density is $f(y) = (1/2)\exp(-|y|)$

Q: Is it worthwhile to look at other distributions in these families such as Normal(0, 3) or Normal(1, 3) and so on? Justify your answer.

```
## Warning: package 'rmutil' was built under R version 3.5.2
```

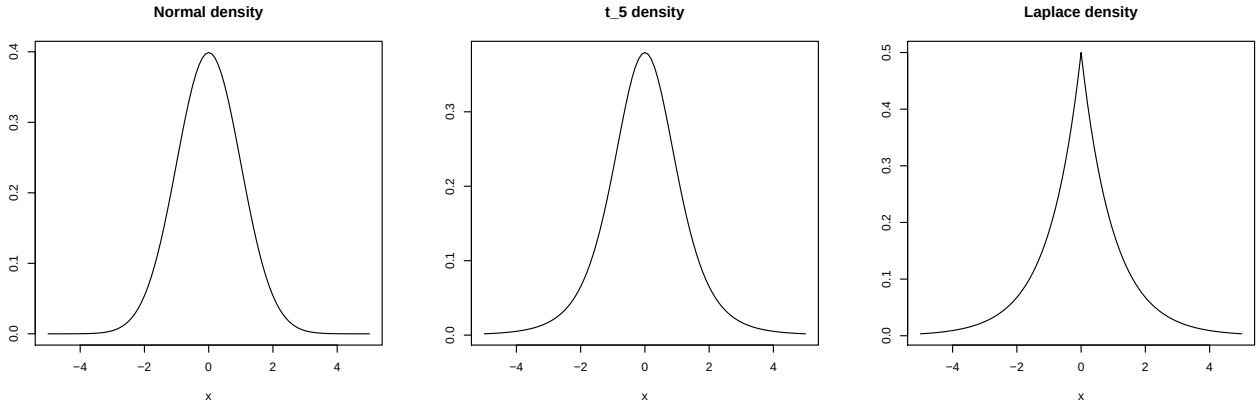


```r
# fn.stats - calculates the mean/trimmed mean/med
# data - arranged n-by-N matrix;
# N= Monte Carlo size; n=sample size
fn.stats <- function(data){

  mean_data <- apply(data, 2, mean)
  trmean_data <- apply(data, 2, mean, trim = .2)
  median_data <- apply(data, 2, median)
  out=cbind(mean_data, trmean_data, median_data)
  out
}

# fn.sampl.distn - outputs data from Normal, t_5 and Laplace
# N= Monte Carlo size; n=sample size
# sets the seed of the simulation
fn.sampl.distn<- function(N, n, seed=1234){
  set.seed(seed)

  # Normal(0,1)
  data <- matrix(rnorm(N*n), nrow=n)
  out_norm <- fn.stats(data)
  # t_5
  data <- matrix(rt(N*n, df=5), nrow=n)
  out_t5 <-   fn.stats(data)
  # Laplace(0,1)
  data <- matrix(rlaplace(N*n), nrow=n)
  out_laplace <-  fn.stats(data)

  list(norm = out_norm, t5=out_t5, laplace=out_laplace )
}


out<- fn.sampl.distn(N=100, n=10, seed=1234)
#head(out$norm);
```

```r
#head(out$t5);
#head(out$laplace)


# Simulation size
N<-10000

out.10<- fn.sampl.distn(N=N, n=10, seed=346)
out.50<- fn.sampl.distn(N=N, n=50, seed=346)
out.100<- fn.sampl.distn(N=N, n=100, seed=346)
```

Plot the sampling distribution for the three estimators in the case of Laplace

```r
library(knitr)
library(rmutil) # required for Laplace density

xlim0<-c(-7/sqrt(50), 7/sqrt(50))

par(mfrow=c(1,3))
hist(out.50$laplace[,1], probability = TRUE, xlim=xlim0,
     main="Mean, n=50", xlab="")
hist(out.50$laplace[,2], probability = TRUE,xlim=xlim0,
     main="20%Trimmed Mean, n=50", xlab="")
hist(out.50$laplace[,3], probability = TRUE, xlim=xlim0,
     main="Median, n=50", xlab="")
```

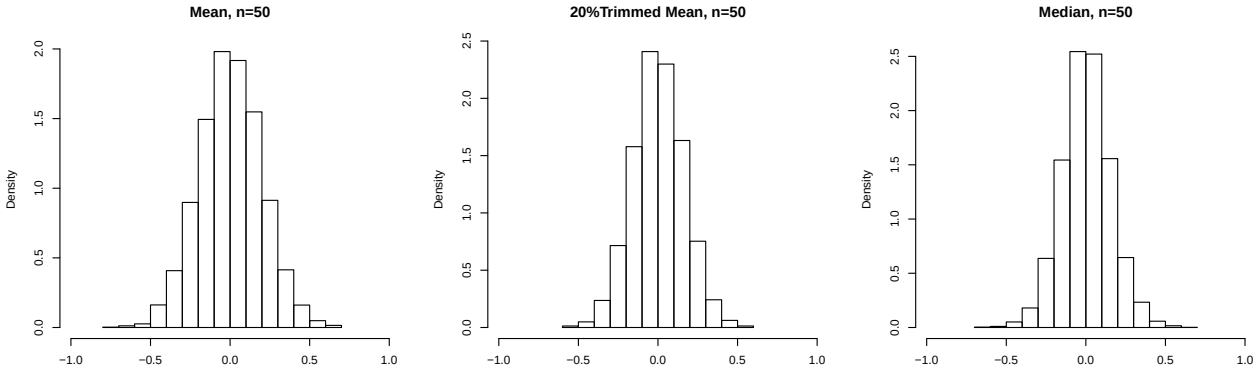

Evaluate estimators using the bias and standard error

```r
# true_mean -true mean for Normal, t_5,Laplace
true_mean <-  rep(0,3)

# Bias = Mean estimator - True mean
bias.10 <- sapply(out.10, function(x) apply(x, 2, mean), simplify = TRUE) - true_mean
bias.50 <-sapply(out.50, function(x) apply(x, 2, mean), simplify = TRUE) - true_mean
bias.100 <-sapply(out.100, function(x) apply(x, 2, mean), simplify = TRUE) -true_mean
rownames(bias.10)<-rownames(bias.50)<-rownames(bias.100)<- c( "mean",   "trmean", "media
```

```
(round(bias.10 , 3))   # n=10
```

```
##          norm    t5 laplace
## mean   -0.003 0.008  -0.004
## trmean -0.004 0.006  -0.004
## median -0.003 0.005  -0.002
```

```
(round(bias.50 , 3))   # n=50
```

```
##          norm    t5 laplace
## mean    0.000 0.000   0.002
## trmean 0.001 0.001   0.002
## median 0.002 0.001   0.003
```

```
(round(bias.100 , 3))  # n= 100
```

```
##          norm    t5 laplace
## mean    0.001 0.000   0.001
## trmean 0.001 0.001   0.001
## median 0.002 0.001   0.000
```

The standard error of the Monte Carlo based estimator is an estimate of its standard deviation.

```
se.10 <- sapply(out.10, function(x) apply(x, 2, sd), simplify = TRUE)
se.50 <- sapply(out.50, function(x) apply(x, 2, sd), simplify = TRUE)
se.100 <- sapply(out.100, function(x) apply(x, 2, sd), simplify = TRUE)
rownames(se.10)<-rownames(se.50)<-rownames(se.100)<- c( "mean",   "trmean", "median")

(round(se.10*sqrt(10) , 2))    # n=10
```

```
##         norm   t5 laplace
## mean    1.00 1.31    1.42
## trmean 1.06 1.19    1.21
## median 1.17 1.27    1.21
```

```
(round(se.50*sqrt(50) , 2))    # n=50
```

```
##         norm   t5 laplace
## mean    1.00 1.30    1.39
## trmean 1.07 1.18    1.14
## median 1.23 1.31    1.09
```

```
(round(se.100*sqrt(100) , 2))  # n=100
```

```
##         norm   t5 laplace
## mean    1.00 1.31    1.41
## trmean 1.07 1.17    1.14
## median 1.24 1.31    1.08
```

3. Principles of Monte Carlo studies

In designing a Monte Carlo study it is important to use the following basic principles:

1. A Monte Carlo experiment, just like any other experiment, involves careful planning, in addition to the carrying the experiment itself, analyzing the results and interpreting them. In particular, consideration should be given to:
   - What factors should be varied in the experiment; think of factors whose effect is relevant to study. These factors almost always include the sample size.
   - The size of the Monte Carlo simulation, which is the number of replicated datasets. This is chosen to be large enough to ensure the desired precision, while small enough to ensure the computations are feasible.
   - What information should be saved from the experiment and how it should be analyzed to accomplish the experiment objective.
2. Whenever possible, save the values of the statistics calculated on all the generated datasets.
3. When coding the simulation experiment, use a low number of replicates. Increase this number to the actual size of the Monte Carlo experiment selected, ONLY after the code is bug-free.
4. Organize your work and document both the simmulation code and the results. This proves extremely helpful when resuming the work after a break. Also whenever possible, keep track of the seed used in random number generation.

4. Importance of sample size in Monte Carlo studies

Determining the sample size in Monte Carlo is not different from other sample size calculations. Here are few examples:

- *Bias estimation.* Say you are interested in reporting the bias of an estimator $\hat{\theta}$ of a true parameter value θ_0. Using N Monte Carlo replicates (or simulations), the bias is estimated as

$$\hat{Bias}(\hat{\theta}) = \frac{1}{N}\sum_{i=1}^{N}\hat{\theta}_i - \theta_0.$$

  In class: Suppose we have access to a guess of the variance of the estimator. Can you think of a way to calculate the minimum Monte Carlo size N such that the precision (i.e. standard deviation) of the bias is within 2 decimal points?

  Work on the board.

- *Variance estimation.* In class: Suppose we have access to a guess of the variance of the estimator. Sounds a bit weird, but let's still play along.

  In class: Can you think of a way to calculate the smallest Monte Carlo size N that would allow to estimate the variance of your estimator with a desired precision?

  Recall the sample variance of the estimator $\hat{\theta}$ is

$$s^2_{N-1} = \frac{1}{N-1}\sum_{i=1}^{N}(\hat{\theta}_i - \bar{\theta})^2,$$

where $\bar{\theta} = N^{-1} \sum_i \hat{\theta}_i$ is the sample mean. Using the asymptotic results studied earlier in the course ...

Work on the board.

For large n, the distribution of $\hat{\theta}$ is approximately normal (kurtosis for normal=3), implying that $Kurtosis(\hat{\theta}) \approx 3$; this implies further that the the sample variance is approximately normal with mean equal to the true variance and variance equal to $\frac{2}{N}\sigma^4_{\hat{\theta}}$. As before we can now calculate the minimum Monte Carlo size N that ensures a specific precision of the sample variance, and is based on an apriori guess for the variance.

- *Power estimation.* Suppose we have a testing procedure described by a test statistic T and a rejection region $RR = \{T > c_\alpha\}$, where c_α is a critical value such that the corresponding to a Type I error rate of the testing procedure is α. Assuming a parameter value in the alternative hypothesis, one estimates the *power at that parameter value* by the mean rejection probability

$$\hat{power} = \frac{1}{N} \sum_{i=1}^{N} 1(T_i > c_\alpha)$$

where T_i are values of the test statistic corresponding to the ith generated dataset by using the specific parameter value included in the alternative hypothesis.

Inlass: What is the minimum Monte Carlo size N such that the power curves have a certain desired precision?

Note that the rejection indicator for some alternative, $1(T_i > c_\alpha)$, follows Bernoulli

distribution; the above estimator of the power has the mean equal to the true power value (denoted by p) and variance equal to $p(1-p)/N$ - which is maximized for $p = 1/2$. For $p = 1/2$, the variance of the estimator of the power is $1/4N$. The minimum Monte Carlo size N that ensures a certain precision in the power estimator can be easily calculated. Let d be the desired precision; the minimum sample size N is solved from the inequality (by rounding up):

$$\frac{1}{4N} \leq d^2 \quad \text{implying that} \quad N \approx \frac{1}{4d^2}.$$

- *Confidence intervals.* When constructing confidence intervals two important aspects are: 1) the actual coverage probability of the confidence intervals and 2) the expected length of the confidence intervals. To be specific, we consider intervals of the form $\hat{\theta} \pm z^*_{\alpha/2}\sqrt{Var(\hat{\theta})}$.

  In class: What is the minimum Monte Carlo size N such that the error in the coverage probability is accurate with certain precision? Alternatively, what is the minimum Monte Carlo size N such that the average expected length has a specified accuracy?

  Work on the board.

5. Analysis of the results

In any Monte Carlo study, whether of interest is estimation or testing, one should report a measure of standard deviation in the results. Additionally, if one is interested in comparison between multiple methods, one should take into account the blocking induced by simulation: each competitive estimator is evaluated on the same generated data, hence there is dependence among them which should be accounted for when comparing them.

A. Suppose you are interested in the performance of one estimator, say sample mean ($\hat{\theta}$). The quality of the variance estimator $\hat{\sigma}^2_{i,n}$ is important when constructing confidence intervals.

Assume that from large sample theory we have that the asymptotic distribution of $\hat{\theta}$ is normal with some mean and variance σ^2_n. From the ith Monte Carlo dataset we obtain empirical-based estimates $\hat{\theta}_i$ and $\hat{\sigma}^2_{i,n}$.

Thus based on the Monte Carlo study we can actually estimate the variance (σ^2_n) of the estimator $\hat{\theta}$ in two ways: 1) as the average of all the variance estimates $\hat{\sigma}^2_{i,n}$ and 2) as the variance of all the estimates $\hat{\theta}_i$. Let

$$R_N = \frac{\frac{1}{N}\sum_i \hat{\sigma}^2_{i,n}}{s^2_{N-1}};$$

the numerator, in fact, estimates $E[\hat{\sigma}^2_{i,n}]$, while the denominator is an unbiased estimator of the variance of $\hat{\theta}$.

What does it mean if we observe $R_N = 1.07$?

What does it mean if $R_N = 0.83$?

- We can use this ratio quantity to assess the quality of the variance estimator, $\hat{\sigma}^2_{i,n}$. How?

  Studying R_N we observe that it's essentially the ratio of two average estimators and thus it is a random quantity. To derive its distribution we need take into account

  i. asymptotic distribution of the numerator

  ii. asymtptotic distribution of the denominator

  iii. use Delta method to derive the asymptotic distribution of the ratio R_N (with n fixed and $N \to \infty$). In particular it's asymptotic variance is

$$\frac{1}{N}\frac{(E[\hat{\sigma}^2_n])^2}{\sigma^4_n}\{Kurtosis(\hat{\theta}) - 1 - 2\frac{Cov(\{\hat{\theta} - E[\hat{\theta}]\}^2, \hat{\sigma}^2_n)}{\sigma^2_n E[\hat{\sigma}^2_n]} + \frac{Var(\hat{\sigma}^2_n)}{(E[\hat{\sigma}^2_n])^2}\}$$

For example consider a situation when $R_N = 1.04$ and an estimated standard error of R_N (using empirical based estimates of the asymptotic variance) is 0.016. Thus it implies that the variance estimator $\hat{\sigma}^2_{i,n}$ is slightly biased upward on average.

B. Mean Squared Error (MSE) - as a performance measure criterion. MSE is commonly reported when estimating the performance of an estimator or comparing one or multiple estimators as it combines measures of bias and variance. However when it comes to assessing the variance estimation, it is recommended against it, due to the fact that it rewards underestimation too much. And, since the variance is used to construct confidence intervals, underestimation results in poor coverage probability of the confidence intervals. Read Ch 9.4.3. - for more discussion of this aspect.

C. Suppose you are interested in comparing two estimators. For simplicity consider the sample mean ($\hat{\theta}_1$) with the trimmed mean ($\hat{\theta}_2$).

- Comparing the mean of two estimators should be carried out using a paired t-test

  When comparing multiple estimators on the same generated datasets, it is important to appropriately take into account the fact that the estimators are correlated. Specifically the variance of the difference between $\hat{\theta}^{(1)}$ and $\hat{\theta}^{(2)}$ is

  $$Var\{\hat{\theta}^{(1)} - \hat{\theta}^{(2)}\}$$

  and not $Var\{\hat{\theta}^{(1)}\} + Var\{\hat{\theta}^{(2)}\}$. If further comparative analysis between the two estimators is needed, a paired t-test should be used and not a two-sample t-test.

- Comparing the variance of two estimators, when the estimators are correlated.

  Assume $s^2_{1,N-1}$ and $s^2_{2,N-1}$ are the variance estimates of the two estimators calculated on the same N generated datasets. Clearly taking their ratio $s^2_{1,N-1}/s^2_{2,N-1}$ could tell you which variance is smaller. Can you imply, based on this which estimator has smaller variance though?

  The ratio $s^2_{1,N-1}/s^2_{2,N-1}$ is indeeed random, and in order to formally conclude which estimator is more efficient for a given sample size n, one has to study the asymptotic distribution of this ratio. In establishing this take into account:

  i. the asymptotic distribution of each sample variance

  ii. the asymptotic distribution of the joint sample variances

  iii use Delta method to derive the asymptotic distribution of $s^2_{1,N-1}/s^2_{2,N-1}$; its asymptotic variance is:

  $$\frac{1}{N}\frac{\sigma^4_{1,n}}{\sigma^4_{2,n}}\{Kurtosis(\hat{\theta}_1) + Kurtosis(\hat{\theta}_2) - 2 - 2\frac{Cov(\hat{\theta}_1, \hat{\theta}_2)}{\sigma^2_{1,n}\sigma^2_{2,n}}\}.$$

  Estimate all the quantities using the sample-based versions.

6. Presentation of the results

When presenting the results, it is useful to remember the following principles:

- Use graphs whenever possible. It is fine to have both a graph and a table, subject to space restrictions.
- Always give some idea of the standard error of each estimated entry in tables, and in graphs when feasible. Preferably: report a range or average standard error in a note at the bottom of tables and in the caption of figures.
- It is best to use at most two significant digits in table entries, and seldom are more than three required.
- It makes little sense to include digits beyond the standard error of the entry. For example, suppose the computer gives .04586 for an entry but the standard error of the entry is .002. Then there is no reason to report more than three digits (in this case.046).
  One possible exception is when the difference of entries has a much smaller standard error than the individual entry standard error.

In class exercise:

Zhang and Boos (1997) discuss new methods to handle clustered binomial data appearing in a series of 2×2 tables. The details are not important here, but one of the paper's sections investigated confidence intervals for the common odds ratio. In a Monte Carlo study, a new confidence interval (CU) was compared to another confidence interval based on a competitive method (denoted CL). Table 9.4 is a portion of the first attempt at presenting the results in a table.(Boos and Stefanski, 2013)

Which of the following 3 tables presents results more clearly and why?

Table 9.4 Coverage and length of 95% confidence intervals for data from the beta-binomial(ρ) distribution with odds ratio = 1.5

Number of strata		$n_{ij}, m_{ij} =$	$\rho = 0.0$ 5	$\rho = 0.0$ 5–10	$\rho = 0.0$ 5–15	$\rho = 0.2$ 5	$\rho = 0.2$ 5–10	$\rho = 0.2$ 5–15
$k = 5$	C_L	Coverage	.977	.977	.977	.984	.980	.975
		Mean Length	1.93	1.48	1.25	5.15	2.88	3.95
	C_U	Coverage	.969	.964	.968	.977	.968	.961
		Mean Length	0.95	0.77	0.69	1.42	1.29	1.26

Table 9.5 Coverage and length of 95% confidence intervals for data from the beta-binomial(ρ) distribution with odds ratio = 1.5

Number of strata		$n_{ij}, m_{ij} =$	$\rho = 0.0$ 5	$\rho = 0.0$ 5–10	$\rho = 0.0$ 5–15	$\rho = 0.2$ 5	$\rho = 0.2$ 5–10	$\rho = 0.2$ 5–15
$k = 5$	C_L	Coverage	.98	.98	.98	.98	.98	.98
		Mean Length	1.9	1.5	1.3	5.2	2.9	4.0
	C_U	Coverage	.97	.96	.97	.98	.97	.96
		Mean Length	1.0	0.8	0.7	1.4	1.3	1.3

Table 9.6 Coverage and length of 95% confidence intervals for data from the beta-binomial(ρ) distribution with odds ratio = 1.5

Number of strata	$n_{ij}, m_{ij} =$		$\rho = 0.0$ 5	$\rho = 0.0$ 5–10	$\rho = 0.0$ 5–15	$\rho = 0.2$ 5	$\rho = 0.2$ 5–10	$\rho = 0.2$ 5–15
$k = 5$	Coverage	C_L	.98	.98	.98	.98	.98	.98
		C_U	.97	.96	.97	.98	.97	.96
	Mean Length	C_L	1.9	1.5	1.3	5.2	2.9	4.0
		C_U	1.0	0.8	0.7	1.4	1.3	1.3

Taken from Zhang and Boos (1997, p. 1193)

Figure 1: A caption