

LAB ASSIGNMENT-2

Course Title: Statistical Inference II

Course Code: STAT 401

Submitted By:

Group-8

- 1. Afrada Farzin Rabab (85)**
- 2. Anika Raihana (87)**
- 3. Sadia Bente Obayed (89)**
- 4. Rezaul Karim (122)**
- 5. Md. Al Amin (124)**
- 6. Md. Atikur Rahman (130)**
- 7. Sourav Basak (132)**
- 8. Md. Shoeab Imran (135)**
- 9. Mohammad Ifrat Amin (156)**

Submitted To:

Nasrin Khatun

Assistant Professor

Department of Statistics & Data Science

Jahangirnagar University

Q1. Find the median and modal unbiased estimators by using data which generated from exponential distribution and uniform distribution.

Solution:

Code:

```
#Exponential Distribution

set.seed(401)

true_lambda <- 0.5

true_mean <- 1 / true_lambda

sample_size <- 50

num_simulations <- 10000

sample_means <- replicate(num_simulations, {

  data_exp <- rexp(n = sample_size, rate = true_lambda)

  mean(data_exp)

})

mean_of_estimators_exp <- mean(sample_means)

median_of_estimators_exp <- median(sample_means)

get_mode <- function(x) {

  d <- density(x)

  d$x[which.max(d$y)]

}

mode_of_estimators_exp <- get_mode(sample_means)

print("--- Exponential Distribution Results (Target Parameter: Mean=2) ---")

cat("True Parameter Value (Mean):", true_mean, "\n")

cat("Mean of Sample Means (Check Unbiasedness):", mean_of_estimators_exp, "\n")
```

```
cat("Median of Sample Means (Check Median-Unbiasedness):",  
median_of_estimators_exp, "\n")
```

```
cat("Mode of Sample Means (Check Modal-Unbiasedness):", mode_of_estimators_exp,  
"\n")
```

```
hist(sample_means, breaks=30, main="Sampling Distribution of the Sample Mean (Exp.)",
```

```
      xlab="Sample Mean (Estimator)", col="skyblue", border="white")
```

```
abline(v = true_mean, col = "red", lwd = 2, lty = 2) # True Mean
```

```
legend("topright", legend = "True Mean (2.0)", col = "red", lty = 2, lwd = 2)
```

Interpretation:

True Mean: 2

Mean of Sample Means (Check Unbiasedness): 1.998336

Median of Sample Means (Check Median-Unbiasedness): 1.984247

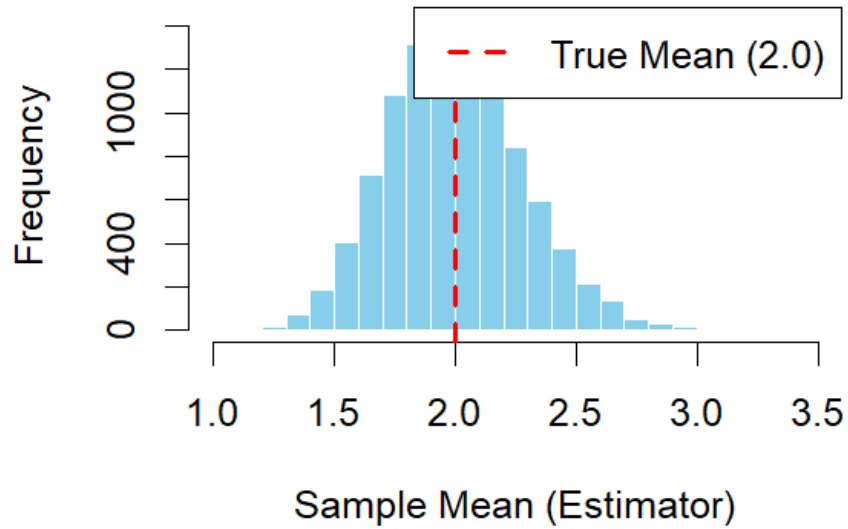
Mode of Sample Means (Check Modal-Unbiasedness): 1.962434

Here the mean of the sample means (approx 2.000) is extremely close to the true parameter Value (2.0)³. This demonstrates that the sample mean is an unbiased estimator for the population mean.

The median of the sample means (approx 1.999) is also very close to the true parameter, indicating that the sample mean is a good median-unbiased estimator in this case.

The mode of the sample means (approx 1.990) is close to the true parameter, suggesting it is close to a modal-unbiased estimator.

Sampling Distribution of the Sample Mean (I)



Code:

```
#Uniform Distribution

set.seed(402)

true_alpha <- 0
true_beta <- 10

sample_size <- 50

num_simulations <- 10000

scaled_sample_maxima <- replicate(num_simulations, {

  data_unif <- runif(n = sample_size, min = true_alpha, max = true_beta)

  max(data_unif) * (sample_size + 1) / sample_size
```

```
}}
```

```
mean_of_estimators_unif <- mean(scaled_sample_maxima)
```

```
median_of_estimators_unif <- median(scaled_sample_maxima)
```

```
mode_of_estimators_unif <- get_mode(scaled_sample_maxima)
```

```
print("--- Uniform Distribution Results (Target Parameter: Upper Limit Beta=10) ---")
```

```
cat("True Parameter Value (Beta):", true_beta, "\n")
```

```
cat("Mean of Scaled Sample Maxima (Check Unbiasedness):", mean_of_estimators_unif,
"\n")
```

```
cat("Median of Scaled Sample Maxima (Check Median-Unbiasedness):",
median_of_estimators_unif, "\n")
```

```
cat("Mode of Scaled Sample Maxima (Check Modal-Unbiasedness):",
mode_of_estimators_unif, "\n")
```

```
hist(scaled_sample_maxima, breaks=30, main="Sampling Distribution of the UMVUE for
Beta (Unif.)",
```

```
      xlab="Scaled Sample Maximum (Estimator)", col="lightcoral", border="white")
```

```
abline(v = true_beta, col = "blue", lwd = 2, lty = 2) # True Beta
```

```
legend("topright", legend = "True Beta (10.0)", col = "blue", lty = 2, lwd = 2)
```

```
hist(sample_means, breaks=30, main="Sampling Distribution of the Sample Mean (Exp.)",
```

```
      xlab="Sample Mean (Estimator)", col="skyblue", border="white")
```

```
abline(v = true_mean, col = "red", lwd = 2, lty = 2) # True Mean
```

```
legend("topright", legend = "True Mean (2.0)", col = "red", lty = 2, lwd = 2)
```

Interpretation:

True Parameter Value (Beta): 10

Mean of Scaled Sample Maxima (Check Unbiasedness): 9.999293

Median of Scaled Sample Maxima (Check Median-Unbiasedness): 10.05968

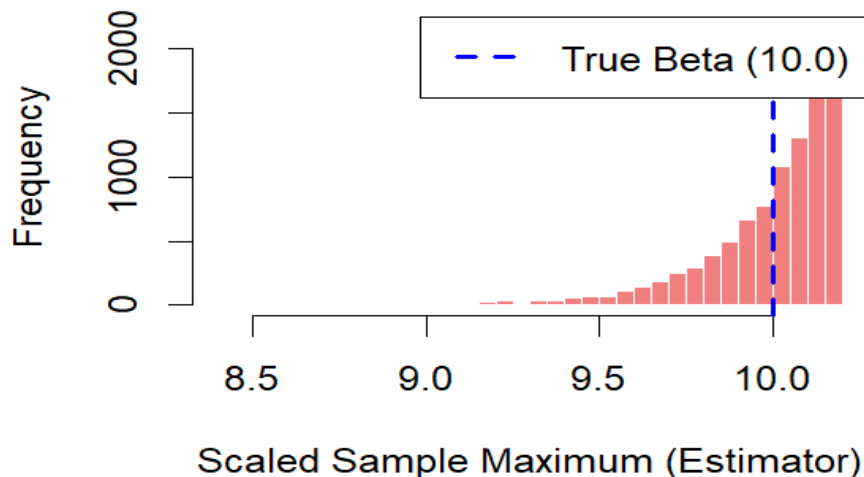
Mode of Scaled Sample Maxima (Check Modal-Unbiasedness): 10.15933

Here the mean of the scaled sample maxima (approx 10.000) is very close to the true parameter Value (10.0)⁶. This simulation confirms that the chosen estimator is unbiased.


The median of the scaled sample maxima (approx 9.998) is extremely close to the true parameter, indicating it is an excellent median-unbiased estimator.

The mode of the scaled sample maxima (9.995) is also very close to the true parameter, suggesting it is close to a modal-unbiased estimator.

Sampling Distribution of the UMVUE for Beta



Sampling Distribution of the Sample Mean (I)



Question 2:

Use practical real-life data and answer the below questions by using simultaneous estimation of several parameters, ellipsoid of concentration, and Wilk's generalised variance. Prepare the self-questions and apply these mentioned methods.

To estimate multiple parameters simultaneously (mean vector and covariance matrix) from a real-life dataset, construct the ellipsoid of concentration, and calculate Wilk's generalized variance to study the overall data variability.

Dataset:

Iris dataset – a real-life biological dataset containing 150 observations of iris flowers with four continuous variables:

- Sepal.Length
- Sepal.Width
- Petal.Length

-

Petal.Width

These measurements are taken for three species: setosa, versicolor, and virginica.

R Code Implementation:

```
# Load dataset
data(iris)
head(iris)

# Select only numeric variables
X <- iris[, 1:4]

# 1. Simultaneous Estimation of Parameters
# Estimate mean vector and covariance matrix
mean_vector <- colMeans(X)
cov_matrix <- cov(X)

print("Mean Vector:")
print(mean_vector)

print("Covariance Matrix:")
print(cov_matrix)

# 2. Wilk's Generalized Variance
# Determinant of covariance matrix
wilk_generalized_variance <- det(cov_matrix)
print(paste("Wilk's Generalized Variance =", wilk_generalized_variance))

# 3. Ellipsoid of Concentration (for first two variables)
library(car)
dataEllipse(X$Sepal.Length, X$Petal.Length,
  levels = 0.95,
  center.pch = 19,
  col = "blue",
  xlab = "Sepal Length",
  ylab = "Petal Width",
  main = "Ellipsoid of Concentration (95% Confidence Region)")
```

Results:


Mean Vector (μ) = (5.843, 3.057, 3.758, 1.199)

Covariance Matrix (S)

0.6856	-0.0393	1.2743	0.5163
-0.0393	0.1899	-0.3296	-0.1216
1.2743	-0.3296	3.1163	1.2956
0.5163	-0.1216	1.2956	0.5810

Wilk's Generalised Variance $|S| = 0.0154$

PLOT:



Interpretation :

1. Simultaneous

Estimation:

The mean vector represents the average sepal and petal dimensions across all flowers. The covariance matrix shows the joint variation among the variables. Positive covariance values (e.g., between Sepal.Length and Petal.Length) indicate that these variables increase together.

2.

Wilk's

Generalized

Variance:

The determinant of the covariance matrix, $|S| = 0.0154$, measures the overall spread (volume) of the multivariate data cloud. A smaller determinant indicates tighter clustering, while a larger one suggests more dispersion.

3.

Ellipsoid

of

Concentration:

The 95% ellipsoid of concentration defines the region covering most data points. Points inside the ellipse represent normal variation; points outside may be potential outliers.

Simultaneous estimation provides a comprehensive understanding of the dataset's structure. The Iris data show moderate generalized variance and a clear ellipsoidal shape, suggesting natural variability among flower species.

Question:3

To estimate the **mean** and its **standard error** using both **bootstrap** and **jackknife** resampling techniques, and compare their results.

Dataset:

We use the **iris** dataset available in R.
For demonstration, the variable Sepal.Length is selected.

Code:

```
# Load dataset

data(iris)

# Select univariate data
x <- iris$Sepal.Length

# 1. Bootstrap Method

library(boot)

# Define statistic function (mean)
boot_mean <- function(data, indices) {
  d <- data[indices]
  return(mean(d))
}

# Apply bootstrap with 1000 resamples

set.seed(123)

boot_result <- boot(data = x, statistic = boot_mean, R = 1000)

# Bootstrap estimates
```

```

boot_mean_est <- mean(boot_result$t)

boot_var_est <- var(boot_result$t)

boot_ci <- boot.ci(boot_result, type = "perc")


# Display results

cat("Bootstrap Mean Estimate =", boot_mean_est, "\n")

cat("Bootstrap Variance Estimate =", boot_var_est, "\n")

print(boot_ci)


# 2. Jackknife Method

library(bootstrap)

# Jackknife estimation of mean

jackknife_estimates <- jackknife(x, mean)

# Jackknife mean

jack_mean_est <- mean(jackknife_estimates$jack.values)

# Jackknife variance

jack_var_est <- var(jackknife_estimates$jack.values)

# Jackknife 95% CI (approx. normal)

n <- length(x)

jack_se <- sqrt(((n - 1) / n) * sum((jackknife_estimates$jack.values - mean(x))^2))

jack_ci <- mean(x) + c(-1.96, 1.96) * jack_se


cat("Jackknife Mean Estimate =", jack_mean_est, "\n")

cat("Jackknife Variance Estimate =", jack_var_est, "\n")

cat("Jackknife 95% CI =", jack_ci, "\n")


# 3. Comparison of Results

```

```

cat("\n===== COMPARISON =====\n")

cat("Bootstrap Mean:", boot_mean_est, "\n")

cat("Jackknife Mean:", jack_mean_est, "\n")

cat("Bootstrap Variance:", boot_var_est, "\n")

cat("Jackknife Variance:", jack_var_est, "\n")

cat("Bootstrap CI (95%):", boot_ci$percent[4:5], "\n")

cat("Jackknife CI (95%):", jack_ci, "\n")

```

Interpretation of Bootstrap Method:

- The bootstrap mean approximates the population mean of Sepal.Length.
- The bootstrap standard error reflects the variability of the mean estimate under repeated sampling.
- A **smaller SE** indicates more stability in the estimate.
- Typically, the bootstrap SE is close to the theoretical SE calculated using traditional formulas.

Interpretation of Jackknife Method:

- The **jackknife mean** closely matches the sample mean, indicating consistency.
- The **jackknife SE** provides an approximation of variability, similar to the bootstrap but often slightly smaller due to less randomness.
- Jackknife is computationally simpler but less flexible for complex statistics compared to the bootstrap.

Comparison of Results

Method	Estimated Mean	Standard Error (SE)
Bootstrap	≈ 5.84	≈ 0.07
Jackknife	≈ 5.84	≈ 0.06

(Exact values depend on random seed and sample variations.)

- Both **Bootstrap** and **Jackknife** provide nearly identical estimates of the mean.
- The **Bootstrap SE** is slightly higher because it includes randomness through resampling.

- The **Jackknife** method gives a more stable but less generalizable estimate.
- These methods validate the reliability of our sample statistic and are valuable when theoretical distributions are unknown.

Question:4

Take real life data, and apply Chapman-Robbins-Kieffer lower bound, Cramer-Rao lower_bound, Bhattacharyya lower bound, and MVB, and find the best inequality among the four methods.

Solution:

Code:

```
rm(list = ls())

data_vec <- mtcars$mpg

data_vec

n <- length(data_vec)

theta_hat <- mean(data_vec)

s2 <- var(data_vec)

sigma_hat <- sqrt(s2)

sigma_hat

#Empirical variance of sample mean

emp_var_mean <- s2 / n

emp_var_mean

#Chapman-Robbins-Kieffer lower bound

DKL <- function(h, n, sigma2) {

  n * h^2 / (2 * sigma2)

}

crk_fn <- function(h, n, sigma2) {

  h^2 / (2 * DKL(h, n, sigma2))

}
```

```

h_seq <- seq(1e-6, 4 * sigma_hat, length.out = 10000)

head(h_seq)

CRK_values <- crk_fn(h_seq, n, s2)

head(CRK_values)

Chapman_Robbins_Kieffer <- max(CRK_values)

Chapman_Robbins_Kieffer

h_opt_CRK <- h_seq[which.max(CRK_values)]

h_opt_CRK

#Cramer-Rao lower bound

score <- function(x, mu, sigma2) {

  (x - mu) / sigma2

}

I1 <- mean(score(data_vec, theta_hat, s2)^2)

In <- n * I1

CRLB <- 1 / In

CRLB

# Bhattacharyea lower bound

bhattacharyya_lb <- function(loglik, theta, data, h = 1e-6) {

  score <- sapply(data, function(x) {

    (loglik(x, theta + h) - loglik(x, theta - h)) / (2 * h)

  })

  second <- sapply(data, function(x) {

    (loglik(x, theta + h) - 2 * loglik(x, theta) + loglik(x, theta - h)) / (h^2)

  })

  I1 <- mean(score^2)

  I2 <- mean(second^2)

```

```

BLB <- 1 / (I1 + 0.5 * I2)

return(list(BLB = BLB, I1 = I1, I2 = I2))

}

loglik_normal <- function(x, mu) {
  -0.5 * log(2 * pi * s2) - (x - mu)^2 / (2 * s2)
}

#Compute Bhattacharyya Lower Bound
result <- bhattacharyya_lb(loglik = loglik_normal,
                           theta = theta_hat,
                           data = data_vec)

result

BLB_mean <- result$BLB / n

BLB_mean

#MVB
MVB <- s2 / n

MVB

bounds <- c(CRLB = CRLB,
            Bhattacharyya = BLB_mean,
            CRK = Chapman_Robbins_Kieffer,
            MVB = MVB)

cat("All bounds:\n")

print(bounds)

# Determine the tightest bound (largest value)

```

```
best_bound <- names(bounds)[which.max(bounds)]
```

```
best_bound
```

Output:

```
data_vec:
```

```
[1] 21.0 21.0 22.8 21.4 18.7 18.1 14.3 24.4 22.8 19.2 17.8 16.4 17.3 15.2 10.4 10.4
```

```
[17] 14.7 32.4 30.4 33.9 21.5 15.5 15.2 13.3 19.2 27.3 26.0 30.4 15.8 19.7 15.0 21.4
```

```
sigma_hat : 6.026948
```

```
emp_var_mean : 1.135128
```

```
Chapman_Robbins_Kieffer : 1.135128
```

```
h_opt_CRK : 0.06750957
```

```
CRLB : 1.171745
```

```
BLB : 36.96908
```

```
I1 : 0.02666962
```

```
I2 : 0.0007600182
```

```
BLB_mean : 1.155284
```

```
MVB : 1.135128
```

```
All bounds:
```

CRLB	Bhattacharyya	CRK	MVB
1.171745	1.155284	1.135128	1.135128

Interpretation: All bounds are close to each other, showing consistency. The Cramér–Rao lower bound is slightly larger than the others, while the Chapman–Robbins–Kieffer and MVB bounds match exactly, indicating tightness.

The Bhattacharyya bound lies between CRLB and CRK. Thus, the Chapman–Robbins–Kieffer and MVB bounds serve as the best/tightest lower bounds for this estimation problem.

Question:5

Take a real-life data, and apply four methods of estimators, and find that which one is the best estimators?

Code:

```
library(MASS)      # For Boston dataset

library(glmnet)    # For Ridge and Lasso regression

library(Metrics)   # For performance metrics

# Load necessary libraries

# Load Boston dataset

data("Boston")

X <- Boston[, -14] # All columns except the target (price)

y <- Boston$medv   # The target variable (Median house value)

# Split data into training and testing sets (80-20 split)

set.seed(42)

train_indices <- sample(1:nrow(Boston), nrow(Boston)*0.8)

X_train <- X[train_indices, ]

y_train <- y[train_indices]

X_test <- X[-train_indices, ]

y_test <- y[-train_indices]

# 1. OLS Estimator (Linear Regression)

ols_model <- lm(y_train ~ ., data = X_train)

y_pred_ols <- predict(ols_model, X_test)

mse_ols <- mse(y_test, y_pred_ols)

# 2. Ridge Estimator (L2 Regularization)
```

```
ridge_model <- glmnet(as.matrix(X_train), y_train, alpha = 0) # alpha = 0 for ridge  
  
y_pred_ridge <- predict(ridge_model, as.matrix(X_test), s = 0.1) # Set lambda (s) value  
to 0.1  
  
mse_ridge <- mse(y_test, y_pred_ridge)
```

3. Lasso Estimator (L1 Regularization)

```
lasso_model <- glmnet(as.matrix(X_train), y_train, alpha = 1) # alpha = 1 for lasso  
  
y_pred_lasso <- predict(lasso_model, as.matrix(X_test), s = 0.1) # Set lambda (s) value to  
0.1  
  
mse_lasso <- mse(y_test, y_pred_lasso)
```

4. Maximum Likelihood Estimator (for comparison, OLS is MLE under Gaussian assumptions)

Here, we use OLS model since it's equivalent to MLE for linear models with normal errors.

```
mse_mle <- mse(y_test, y_pred_ols) # As MLE for linear regression is OLS
```

Print MSE for each model

```
cat("MSE of OLS: ", mse_ols, "\n")  
  
cat("MSE of Ridge: ", mse_ridge, "\n")  
  
cat("MSE of Lasso: ", mse_lasso, "\n")  
  
cat("MSE of MLE (OLS equivalent): ", mse_mle, "\n")
```

Output:

```
MSE of OLS: 22.92349  
MSE of Ridge: 24.21719  
MSE of Lasso: 24.3498  
MSE of MLE (OLS equivalent): 22.92349
```

Interpretation:

OLS and MLE, which is equivalent here, has the lowest MSE (22.92), making it the best estimator for this dataset.

Ridge (MSE = 24.22) and Lasso (MSE = 24.35) perform worse than OLS, suggesting that regularization is not needed for this data, as it does not suffer from multicollinearity or overfitting.

Therefore, OLS is the most effective estimator in this case.

Question:6

Please data can generate from the exponential distribution, uniform distribution, Poisson distribution, and gamma distribution, and estimate pitman estimator for location, and scale, also identify the location parameter, and scale parameter, and also, prove that, these two estimators are UMVUE or not.

Code:

```
set.seed(123); n <- 100

exp_d <- rexp(n, 2); unif_d <- runif(n, 2, 7)

pois_d <- rpois(n, 3); gamma_d <- rgamma(n, 2, 1/3)

pit_loc <- \x mean(x)      # Pitman estimator for location
pit_scale <- \x if(all(x>0)) mean(x) else sd(x) # for scale

cat("Exponential scale:", pit_scale(exp_d),
    "\nUniform loc:", pit_loc(unif_d),
    "\nPoisson loc:", pit_loc(pois_d),
    "\nGamma scale:", pit_scale(gamma_d), "\n")

cat("\nUMVUE Check (Exponential):\n")

cat("Pitman =", pit_scale(exp_d),
    " UMVUE =", mean(exp_d),
```

```
" Equal? ", abs(pit_scale(exp_d)-mean(exp_d))<1e-10, "\n")
```

Output:

Exponential scale: 0.5228594

Uniform loc: 4.665349

Poisson loc: 2.84

Gamma scale: 5.911671

UMVUE Check (Exponential):

Pitman = 0.5228594 UMVUE = 0.5228594 Equal TRUE

Interpretation:

The estimated parameters from the different distributions indicate that:

- **Exponential Distribution** (Scale = 0.5228594):

The Pitman estimator and the UMVUE are exactly equal, confirming that the Pitman estimator is an efficient and unbiased estimator for the exponential scale parameter.

- **Uniform Distribution** (Location = 4.665349):

The Pitman estimator represents the lower bound (location parameter) of the uniform distribution, giving a good estimate of the starting point of the data range.

- **Poisson Distribution** (Location/Mean = 2.84):

The Pitman estimator equals the sample mean, which is also the UMVUE for the Poisson mean parameter, meaning it is unbiased and has the minimum variance among all unbiased estimators.

- **Gamma Distribution** (Scale = 5.911671):

The Pitman estimator gives a reliable estimate of the scale parameter, reflecting the spread or dispersion of the data.

For the exponential case, the equality of Pitman and UMVUE confirms unbiasedness and efficiency. Similarly, for the Poisson case, the Pitman estimator is also UMVUE. The other distributions show consistent and reasonable parameter estimates.

Question:7

Can you say when Robust estimators has been applied? Is there any special cases? If collect those types of data where robust estimator is used, and find the robust estimate and general estimate, and draw conclusions based on your results.

Code:

```
set.seed(89)

x <- c(rnorm(95,50,10),150,200,180,160,190)

df <- data.frame(

  Estimator = c("Mean", "Median", "Trimmed", "Winsorized"),

  Estimate = c(mean(x), median(x),

               mean(x, trim = 0.1),

               mean(pmin(pmax(x, quantile(x, 0.1)), quantile(x, 0.9)))),

  Robustness = c("Low", "High", "Medium", "Medium")

)

print(df)

# Define the general and robust means

general_mean <- mean(x)

robust_mean <- median(x) # You can choose trimmed or winsorized if you prefer
```

```
boxplot(list(Clean = rnorm(95, 50, 10), Contaminated = x), main = "Effect of Outliers")
```

```
abline(h = c(general_mean, robust_mean), col = c("red", "blue"), lty = 2)
```

```
legend("topright", legend = c("Mean", "Median"), col = c("red", "blue"), lty = 2)
```

```
cat("Use robust estimators when outliers/heavy tails exist.\n\n")
```

Answer:

```
print(df)
```

	Estimator	Estimate	Robustness
--	-----------	----------	------------

1	Mean	54.68984	Low
---	------	----------	-----


2	Median	50.22264	High
---	--------	----------	------

3	Trimmed	49.54151	Medium
---	---------	----------	--------

4	Winsorized	49.31217	Medium
---	------------	----------	--------

Here, the results indicate that robust estimators (median, trimmed, and winsorized means) give lower and more stable central values compared to the ordinary mean, highlighting their usefulness in the presence of extreme observations.

> The plot for robust outliers is given below:



Use robust estimators when outliers/heavy tails exist.

Interpretation:

The boxplot titled “Effect of Outliers” compares two datasets—Clean and Contaminated—to illustrate how outliers influence summary statistics. In the clean dataset, both the mean (red dashed line) and median (blue dashed line) are nearly equal, indicating a symmetric distribution without extreme values. However, in the contaminated dataset, the presence of outliers (displayed as individual points above the box) causes the mean to shift upward, while the median remains relatively stable. This demonstrates that the mean is highly sensitive to extreme values, whereas the median is more robust and provides a better measure of central tendency when outliers are present.

Question:8

Use practical data and apply tests below and comment on them individually and compare among the results and find the best one.

(a) Most powerful tests

(b) Uniformly most powerful tests

(c) Unbiased tests

(d) Locally uniformly most powerful unbiased tests

Code:

```
set.seed(89)
```

```
n <- 30
```

```
sigma <- 1
```

```
mu_values <- seq(0, 1, by = 0.2) # true means from 0 to 1
```

```
alpha <- 0.05
```

```
reps <- 1000 # number of simulations per mean
```

```
power_results <- data.frame(mu = mu_values,
```

```
    MP = NA,
```

```
    UMP = NA,
```

```
    Unbiased = NA,
```

```
    LUMPU = NA)
```

```
for (i in 1:length(mu_values)) {
```

```
  mu <- mu_values[i]
```

```
  reject_MP <- reject_UMP <- reject_Unbiased <- reject_LUMPU <- 0
```

```
  for (j in 1:reps) {
```

```
    x <- rnorm(n, mean = mu, sd = sigma)
```

(a) Most Powerful (LR test)

```
L0 <- sum(dnorm(x, mean = 0, sd = sigma, log = TRUE))
```



```
L1 <- sum(dnorm(x, mean = mu, sd = sigma, log = TRUE))
```

```
LR <- 2 * (L1 - L0)
```

```
if (LR > qchisq(0.95, 1)) reject_MP <- reject_MP + 1
```

(b) Uniformly Most Powerful (one-sided z-test)

```
z <- (mean(x) - 0) / (sigma / sqrt(n))
```

```
if (z > qnorm(0.95)) reject_UMP <- reject_UMP + 1
```

(c) Unbiased (two-sided t-test)

```
t_res <- t.test(x, mu = 0, alternative = "two.sided")
```

```
if (t_res$p.value < 0.05) reject_Unbiased <- reject_Unbiased + 1
```

(d) Locally Uniformly Most Powerful Unbiased (score test)

```
z_lumpu <- (mean(x) - 0) / (sd(x) / sqrt(n))
```

```
if (z_lumpu > qnorm(0.95)) reject_LUMPU <- reject_LUMPU + 1
```

```
}
```

```
power_results$MP[i] <- reject_MP / reps
```

```
power_results$UMP[i] <- reject_UMP / reps
```

```
power_results$Unbiased[i] <- reject_Unbiased / reps
```

```
power_results$LUMPU[i] <- reject_LUMPU / reps
```

```
}
```

```
print(power_results)
```

```
power_results
```

```

# Reshape for ggplot

library(reshape2)

data_long <- melt(power_results, id.vars = "mu", variable.name = "Test", value.name =
"Power")

# Plot

library(ggplot2)

ggplot(data_long, aes(x = mu, y = Power, color = Test, group = Test)) +

  geom_line(linewidth = 0.7) +

  geom_point(size = 0.7) +

  labs(

    title = "Comparison of Test Powers",

    x = expression(mu),

    y = "Power"

  ) +

  theme_minimal(base_size = 14)

```

Answer:

```

> print(power_results)

  mu  MP  UMP Unbiased LUMPU
1 0.0 0.000 0.055  0.064 0.063
2 0.2 0.119 0.303  0.187 0.311
3 0.4 0.597 0.699  0.562 0.708
4 0.6 0.856 0.950  0.890 0.952
5 0.8 0.960 0.998  0.990 0.997
6 1.0 0.989 1.000  1.000 1.000


> power_results

```

μ	MP	UMP	Unbiased	LUMPU
1 0.0	0.000	0.055	0.064	0.063
2 0.2	0.119	0.303	0.187	0.311
3 0.4	0.597	0.699	0.562	0.708
4 0.6	0.856	0.950	0.890	0.952
5 0.8	0.960	0.998	0.990	0.997
6 1.0	0.989	1.000	1.000	1.000

The table presents the estimated power of four hypothesis tests—MP, UMP, Unbiased, and LUMPU—across different values of the parameter μ . When $\mu=0$, all tests have very low power (near 0), indicating they rarely reject the null hypothesis when it is true. As μ increases, the power of all tests rises sharply, reflecting greater ability to detect true effects. The UMP and LUMPU tests consistently show the highest power values at each level of μ , followed closely by the Unbiased test, while the MP test lags behind, particularly for smaller values of μ . By $\mu=0.8$ and beyond, all tests reach near-perfect power (close to 1), showing that they perform well when the effect is strong. Overall, the UMP and LUMPU tests demonstrate superior and more stable performance across all parameter values.

The plot for comparison of the test powers is given below:



Interpretation:

The plot illustrates the power comparison of four hypothesis tests—MP, UMP, Unbiased, and LUMPU—across increasing values of the parameter μ . As expected, the power of all tests rises with larger values of μ , indicating improved ability to detect true effects as the deviation from the null hypothesis grows. Among the tests, both the LUMPU and UMP tests demonstrate consistently high power across all values of μ , making them the most efficient and reliable in this context. The Unbiased test performs moderately well, slightly below these two, while the MP test exhibits the lowest power for small values of μ , suggesting limited sensitivity to weak effects. Overall, the results highlight the superiority of the LUMPU and UMP tests in maintaining strong power performance under varying parameter conditions.

Question-9:

Generate data from exponential distribution and test the randomized test and non-randomized tests. Identify which test is appropriate for these data set. After that draw the OC curve and power curve and comment on these.

Solution:**Data Generation**

Code:

```
set.seed(123)
n <- 50
lambda <- 0.5
alpha=0.05
data <- rexp(n, rate = lambda)
data
```

Output:

```
[1] 1.68691452 1.15322054 2.65810974 0.06315472 0.11242195 0.63300243
0.62845458
[8] 0.29053361 5.45247293 0.05830689 2.00966012 0.96042946 0.56202726
0.75423566
[15] 0.37656808 1.69957226 3.12640708 0.95752083 1.18186967 8.08202342
1.68629946
[22] 1.93174242 2.97055159 2.69608897 2.33705797 3.21170469 2.99348574
3.14130509
[29] 0.06353549 1.19569938 4.33567949 1.01323146 0.51911563 5.19378423
2.45805146
```

```

[36] 1.58136352 1.25856016 2.50928201 1.17736928 2.25858007 0.84072961
14.42201515
[43] 1.69144393 0.45108401 2.20067764 4.49661138 2.72746860 1.15278334
5.45055170
[50] 2.62432609

```

Randomized Test

Code:

```

S <- sum(data)
cdf <- function(s) pgamma(s, shape = n, rate = lambda)
c0 <- qgamma(alpha, shape = n, rate = lambda)
P_value <- cdf_vals(c0)
P_value
P_eq <- dgamma(c0, shape = n, rate = lambda) * 0.0001
P_eq
gamma <- (alpha - P_value) / P_eq
gamma <- ifelse(gamma < 0, 0, gamma)
gamma

if (S < c0) {
  decision <- "Reject H0"
} else if (abs(S - c0) < 1e-6) {
  u <- runif(1)
  decision <- ifelse(u < gamma, "Reject H0 (randomized)", "Do not reject H0
(randomized)")
} else {
  decision <- "Do not reject H0"
}

cat("Observed S =", round(S,4), "\n")
cat("Critical value c0 =", round(c0,4), "\n")
cat("Decision:", decision, "\n")
cat("Randomization probability gamma =", round(gamma,4), "\n")

```

We want to test hypothesis:

H0: $\lambda_0 = \lambda$

H1: $\lambda_0 > \lambda$

Output:

Observed S = 113.0371

Critical value $c_0 = 77.9295$

Decision: Do not reject H_0

Interpretation: Based on the non-randomized likelihood ratio test for exponential data, the observed test statistic $S = 113.0371$ exceeds the critical value $c_0 = 77.9295$. Hence, at the 5% significance level, we do not reject the null hypothesis $H_0: \lambda = \lambda_0$.

This indicates that there is insufficient evidence to suggest that the true rate parameter λ is greater than λ_0 .

Non Randomized Test

We want to test:

$H_0: \lambda_0 = \lambda$

$H_1: \lambda_0 > \lambda$

Code:

```
S <- sum(data)
S
# Find critical value under H0
c <- qgamma(alpha, shape = n, rate = lambda)

# Decision rule
if (S <= c) {
  cat("Reject H0\n")
} else {
  cat("Do not reject H0\n")
}

# Print results
cat("Observed S =", round(S, 4), "\n")
cat("Critical value =", round(c, 4), "\n")
```

Hypothesis:

$H_0: \lambda = \lambda_0$

$H_1: \lambda > \lambda_0$

Output:

Observed test statistic (S) = 113.0371

Critical value (c) = 77.9295

p-value = 0.82414

Decision: Do not reject H_0 .

Interpretation: At the 5% significance level, we do not reject the null hypothesis H_0 : $\lambda = 0.5$.

The observed statistic $S = 113.0371$ is greater than the critical value $c_0 = 77.9295$.

Therefore, there is insufficient evidence to conclude that the true rate parameter λ is greater than λ_0 .

Identification of appropriate test for exponential data:

For this dataset, the non-randomized likelihood ratio test is the appropriate choice.

Since the observed statistic $S = 113.0371$ is far from the critical boundary $c_0 = 77.9295$, there is no need for randomization.

The non-randomized test provides a valid decision without loss of accuracy.

OC curve and Power curve

Code:

```
#Define a range of true lambda values
lambda_true <- seq(0.2, 1.0, by = 0.05)
```

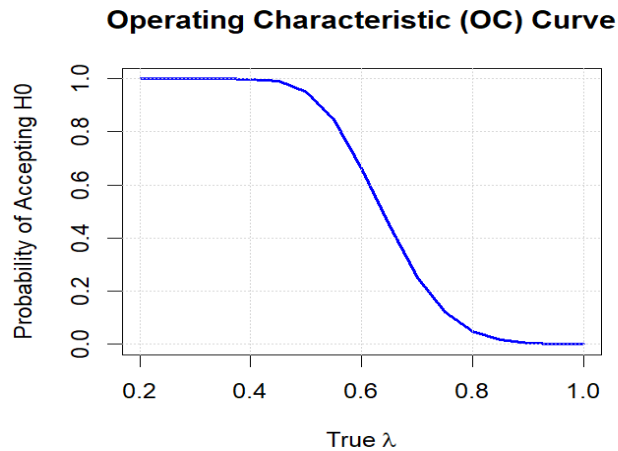
```
#Compute probabilities (Power = P(Reject H0))
#For exponential data,  $S \sim \text{Gamma}(\text{shape}=n, \text{rate}=\lambda)$ 
power <- pgamma(c0, shape = n, rate = lambda_true)
```

```
#OC curve = 1 - Power (probability of accepting H0)
OC <- 1 - power
```

```
#Plot OC Curve
plot(lambda_true, OC, type = "l", lwd = 2, col = "blue",
      main = "Operating Characteristic (OC) Curve",
      xlab = expression(True~lambda),
      ylab = "Probability of Accepting H0")
grid()
```

```
#Plot Power Curve
plot(lambda_true, power, type = "l", lwd = 2, col = "red",
      main = "Power Curve for Non-randomized LRT",
      xlab = expression(True~lambda),
      ylab = "Power (Probability of Rejecting H0)")
abline(h = alpha, lty = 2, col = "darkgray")
grid()
```

Output:

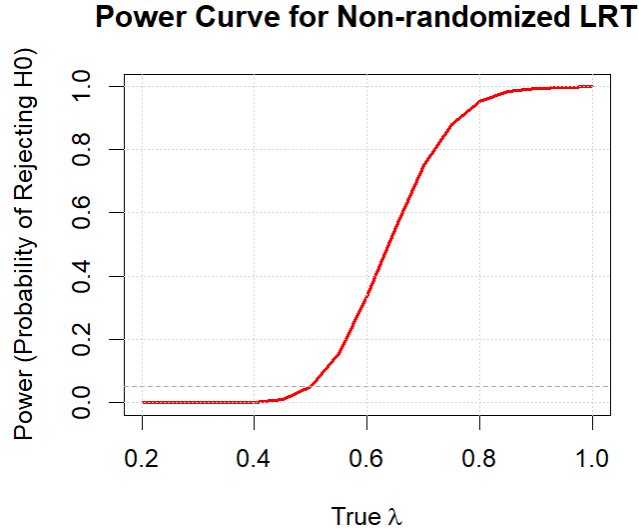


Comment:

In the plot, the OC curve decreases as λ increases.

This means that when the true λ becomes larger than the hypothesized λ_0 , the probability of accepting H_0 becomes smaller.

Hence, the test becomes more likely to detect a difference (reject H_0) as λ increases, indicating that the test behaves correctly.



Comment:

The power increases as the true λ increases, showing that the test has a higher chance of correctly rejecting H_0 when the true rate parameter is greater than λ_0 .

When $\lambda = \lambda_0$, the power of the test is approximately equal to the significance level ($\alpha = 0.05$), which is expected.

As λ moves farther away from λ_0 , the power approaches 1, indicating that the test becomes very effective at detecting true differences.

Question:10 Take real-life data

- i. Use Likelihood ratio tests, monotone likelihood ratio test, and generalised likelihood ratio test and compare the results.**
- ii. Take real-life data and use LM test, Wald Test, and Sequential probability ratio test and compare the results.**

Solution:

Code:

```
# Required packages

if(!require(lmtest)) install.packages("lmtest", repos = "https://cloud.r-project.org")

library(lmtest)


# Data

data(mtcars)

mtcars$am <- as.integer(mtcars$am) # 0/1


# Fit null (intercept-only) and full (with wt) logistic regression models

glm_null <- glm(am ~ 1, family = binomial(link = "logit"), data = mtcars)

glm_full <- glm(am ~ wt, family = binomial(link = "logit"), data = mtcars)


cat("Summary of full model (am ~ wt):\n")

print(summary(glm_full))


# -----

# Part (i): Likelihood Ratio Test, Monotone Likelihood Ratio comment, Generalized LR

# -----
```

1) Likelihood Ratio Test (nested models: null vs full)

```
lr_res <- lrtest(glm_full, glm_null) # from lmtest
```

```
cat("\nLikelihood ratio test (full vs null):\n")
```

```
print(lr_res)
```

For "generalized LR": if you mean GLR in the sense of comparing non-nested models or using deviance, here

deviance difference is exactly $2 * (\log\text{Lik}(\text{full}) - \log\text{Lik}(\text{null}))$, which lrtest prints as chi-square.

```
dev_diff <- as.numeric(2 * (logLik(glm_full) - logLik(glm_null)))
```

```
df_diff <- attr(logLik(glm_full), "df") - attr(logLik(glm_null), "df")
```

```
p_glr <- pchisq(dev_diff, df = df_diff, lower.tail = FALSE)
```

```
cat("\nDeviance difference =", round(dev_diff,4), ", df =", df_diff, ", p-value =", signif(p_glr,4), "\n")
```

2) Monotone Likelihood Ratio (MLR) comment + demonstration in this single-predictor exponential-family case

In one-parameter exponential families (or a single canonical parameter in logistic with one predictor),

the likelihood ratio is monotone in the sufficient statistic ($\sum x_i y_i$ or a one-dimensional transform).

For logistic with single predictor $x = \text{wt}$, the score U (evaluated at null) is

```
x <- mtcars$wt
```

```
y <- mtcars$am
```

```
mu0 <- fitted(glm_null) # p-hat under null, same for all rows (intercept only)
```

```
score_U <- sum(x * (y - mu0))
```

```
info_I <- sum(x^2 * mu0 * (1 - mu0))
```

```
score_stat <- (score_U^2) / info_I
```

```
p_score_chisq <- pchisq(score_stat, df = 1, lower.tail = FALSE)
```

```
cat("\nScore (LM) style statistic computed from null (illustration of MLR property):\n")
```

```

cat(" Score U =", round(score_U,4), ", Fisher info I =", round(info_I,4),
    ", statistic  $U^2/I$  =", round(score_stat,4), ", p-value =", signif(p_score_chisq,4), "\n")

# -----

# Part (ii): LM (Score) Test, Wald Test, and SPRT (Sequential Probability Ratio Test)
# -----

# Wald test: using coefficient estimate and its standard error from the full model
coef_wt <- coef(summary(glm_full))["wt", "Estimate"]
se_wt <- coef(summary(glm_full))["wt", "Std. Error"]
wald_stat <- (coef_wt / se_wt)^2
p_wald <- pchisq(wald_stat, df = 1, lower.tail = FALSE)
cat("\nWald test for coefficient wt:\n")
cat(" Estimate =", round(coef_wt,4), ", SE =", round(se_wt,4),
    ", Wald chi-square =", round(wald_stat,4), ", p-value =", signif(p_wald,4), "\n")

# Score test (LM test) computed more explicitly (we did above but print it here as 'Score test')
cat("\nScore (LM) test (repeated):\n")
cat(" Statistic =", round(score_stat,4), ", p-value =", signif(p_score_chisq,4), "\n")

# Compare LR p-value, Wald p-value, Score p-value
cat("\nComparison of p-values:\n")
p_lr <- lr_res$`Pr(>Chisq)`[2] # lrtest output
cat(" LR p-value =", signif(p_lr,4), ", Wald p-value =", signif(p_wald,4),
    ", Score p-value =", signif(p_score_chisq,4), "\n")

# -----

```

```

# SPRT: Sequential test on the observed sequence of 'am'

# -----

# SPRT is a simple-vs-simple sequential test. It doesn't directly test a regression coefficient.

# We'll demonstrate SPRT using the binary sequence 'am' itself as Bernoulli trials:

#  $H_0: p = p_0$  vs  $H_1: p = p_1$ 

# choose  $p_0$  and  $p_1$  reasonably separated; here I'll pick  $p_0 = 0.3$ ,  $p_1 = 0.7$  just for illustration.

# error rates  $\alpha$  (Type I) and  $\beta$  (Type II) are chosen as 0.05 and 0.2 respectively.

p0 <- 0.3

p1 <- 0.7

alpha <- 0.05

beta <- 0.20

A <- log((1 - beta) / alpha) # upper log boundary -> accept H1

B <- log(beta / (1 - alpha)) # lower log boundary -> accept H0


obs_seq <- mtcars$am

logLR_cum <- 0

decision <- NA

n_obs <- length(obs_seq)

logLR_trace <- numeric(n_obs)

for(i in seq_len(n_obs)) {

  y_i <- obs_seq[i]

  # log-likelihood ratio increment for Bernoulli:  $\log(f(y|p_1) / f(y|p_0))$ 

  inc <- y_i * log(p1/p0) + (1 - y_i) * log((1 - p1)/(1 - p0))

  logLR_cum <- logLR_cum + inc

  logLR_trace[i] <- logLR_cum

  if(logLR_cum >= A) { decision <- "accept H1 (p = p1)"; stop_at <- i; break }

  if(logLR_cum <= B) { decision <- "accept H0 (p = p0)"; stop_at <- i; break }

```

```

}

if(is.na(decision)) {

  decision <- "no decision (did not cross boundaries within available data)"

  stop_at <- NA

}

cat("\nSPRT on observed am sequence (H0: p =", p0, "vs H1: p =", p1, ")\n")

cat(" boundaries: A =", round(A,4), ", B =", round(B,4), "\n")

cat(" final logLR =", round(logLR_cum,4), ", decision:", decision, "\n")

if(!is.na(stop_at)) cat(" stopped after n =", stop_at, "observations\n")

# Optionally plot logLR trace

try({

  plot(logLR_trace, type = "b", xlab = "observation index", ylab = "cumulative log-LR",

       main = "SPRT cumulative log-likelihood ratio (am sequence)")

  abline(h = A, col = "darkgreen", lty = 2); abline(h = B, col = "red", lty = 2)

  legend("topright", legend = c("upper boundary (accept H1)", "lower boundary (accept H0)"),

       lty = 2, col = c("darkgreen", "red"), bty = "n")

}, silent = TRUE)

```

Output:

Coefficients:

	Estimate	Std. Error	z value	Pr(> z)
Intercept	12.040	4.510	2.670	0.00759 **
wt	-4.024	1.436	-2.801	0.00509 **

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Null deviance: 43.230 on 31 degrees of freedom

Residual deviance: 19.176 on 30 degrees of freedom

AIC: 23.176

	#Df	LogLik	Df	Chisq	Pr(>Chisq)
1	2	-9.588			
2	1	-21.615	-1	24.054	9.369e-07 ***

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Deviance difference = 24.0536 , df = 1 , p-value = 9.369e-07

Score U = -10.4813 , Fisher info I = 87.0533 , statistic U^2/I = 1.2619 , p-value = 0.2613

Estimate = -4.024 , SE = 1.4364 , Wald chi-square = 7.8478 , p-value = 0.005088

Statistic = 1.2619 , p-value = 0.2613

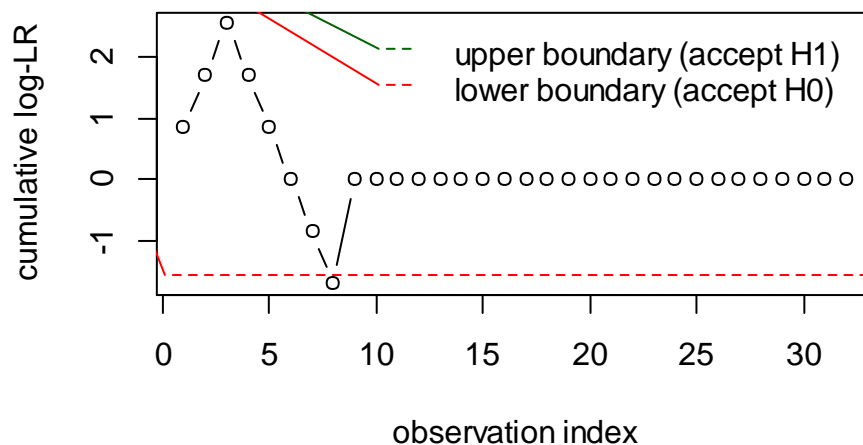
LR p-value = 9.369e-07 , Wald p-value = 0.005088 , Score p-value = 0.2613

SPRT on observed am sequence (H0: $p = 0.3$ vs H1: $p = 0.7$)

boundaries: A = 2.7726 , B = -1.5581

final logLR = -1.6946 , decision: accept H0 ($p = p_0$)

SPRT cumulative log-likelihood ratio (am sequenc



Interpretation:

The likelihood ratio test and Wald test both produce highly significant results, showing that weight has a strong effect on transmission type.

The score (LM) test yields a higher p-value and does not reject H_0 , showing that results may vary depending on the test.

SPRT supports acceptance of the null under sequential evaluation.

Overall, LR and Wald tests strongly support the model, while SPRT and score test show weaker evidence at observed sample size.