

PRACTICAL COURSE ON

# Python with Statistics

SOLUTIONS DOCUMENT

*Complete Answers to All 21 Checkpoint Exercises*

Self-Study

## Table of Contents

**Morning Session: Python Fundamentals** — Exercises 1-9

**Morning Session: Pandas & Data Wrangling** — Exercises 10-14

**Afternoon Session: Descriptive Statistics** — Exercises 1-3

**Afternoon Session: Hypothesis Testing** — Exercises 4-7

Note: Each solution includes complete code and explanation. Participants should attempt exercises before viewing solutions.

## Morning Session Solutions


### Python Fundamentals (Exercises 1-9)

#### Exercise 1: Print Name and Number

**Question:** Print your name and your favorite number on separate lines.

```
# Solution
print('Arif')
print(42)

# Or with a descriptive message
print('My name is Arif')
print('My favorite number is 42')
```

 **Explanation:** Use two separate print() statements to create two lines of output. You can print just the values, or add descriptive text to make it clearer.


#### Exercise 2: Create and Print Variables

**Question:** Create variables for: sample size (n=120), mean score (mean=78.3), standard deviation (std=12.5). Then print all three values.

```
# Solution
sample_size = 120
mean_score = 78.3
std_dev = 12.5

print(sample_size)
print(mean_score)
print(std_dev)

# Or with labels for clarity
print('Sample size:', sample_size)
print('Mean score:', mean_score)
print('Standard deviation:', std_dev)
```

 **Explanation:** Create three variables with the = assignment operator. Integer (120) doesn't need a decimal point, but floats (78.3, 12.5) do. Print each variable to see its value.

#### Exercise 3: Calculate Percentage


**Question:** Calculate the percentage of students who scored above 80, given: total students = 150, students above 80 = 45. Store the result and print it.

```
# Solution
total_students = 150
students_above_80 = 45

percentage = (students_above_80 / total_students) * 100
print(percentage)

# With better formatting
```

```
print(f'Percentage above 80: {percentage}%')  
# Or  
print(f'Percentage above 80: {percentage:.2f}%') # 2 decimal places
```

 **Explanation:** Division: `students_above_80 / total_students` gives 0.3. Multiply by 100 to get 30%. Use f-strings (`f'...{variable}...'`) for formatted output. The `:.2f` format shows 2 decimal places.

#### Exercise 4: List Operations

**Question:** Create a list of 5 test scores: [65, 78, 82, 90, 55]. Then: 1) Print the first score, 2) Print the last score, 3) Calculate and print the average.

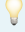
```
# Solution
scores = [65, 78, 82, 90, 55]

# 1. First score
first_score = scores[0]
print('First score:', first_score)

# 2. Last score
last_score = scores[-1]
print('Last score:', last_score)

# 3. Average
total = sum(scores)
count = len(scores)
average = total / count
print('Average:', average)

# Alternative: More concise
print('First:', scores[0])
print('Last:', scores[-1])
print('Average:', sum(scores) / len(scores))
```

 **Explanation:** Python lists are zero-indexed: scores[0] is the first item. Negative indexing counts from the end: scores[-1] is the last item. sum() adds all values, len() counts them.


#### Exercise 5: Adding to Lists

**Question:** Create a list with 3 student names: ['Ahmad', 'Siti', 'Raj']. Add a fourth name 'Lin' to the list. Print the updated list.

```
# Solution
students = ['Ahmad', 'Siti', 'Raj']
print('Original list:', students)

students.append('Lin')
print('Updated list:', students)

# Output:
# Original list: ['Ahmad', 'Siti', 'Raj']
# Updated list: ['Ahmad', 'Siti', 'Raj', 'Lin']
```

 **Explanation:** The .append() method adds an item to the END of the list. It modifies the list in place (no need to reassign). Note: append() takes ONE item at a time.

#### Exercise 6: Loop with Custom Message

**Question:** You have temperatures in Celsius: [20, 25, 30, 15, 28]. Write a loop that prints each temperature with the message: 'Temperature: [X]°C'

```
# Solution
```

```


temperatures = [20, 25, 30, 15, 28]

for temp in temperatures:
    print('Temperature:', temp, '°C')

# Alternative with f-string
for temp in temperatures:
    print(f'Temperature: {temp}°C')

# Output:
# Temperature: 20°C
# Temperature: 25°C
# Temperature: 30°C
# Temperature: 15°C
# Temperature: 28°C

```

 **Explanation:** The for loop iterates through each temperature. Each iteration, the variable 'temp' holds one value from the list. The print statement runs once per temperature.

### Exercise 7: Filtering with List Comprehension

**Question:** You have test scores: [45, 78, 82, 55, 90, 38, 67]. Create a new list containing only scores that are 60 or above. Print the filtered list.

```

# Solution Method 1: List comprehension (concise)
scores = [45, 78, 82, 55, 90, 38, 67]
passing_scores = [s for s in scores if s >= 60]
print('Passing scores:', passing_scores)

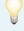
# Solution Method 2: Regular loop (more explicit)
scores = [45, 78, 82, 55, 90, 38, 67]
passing_scores = []

for score in scores:
    if score >= 60:
        passing_scores.append(score)

print('Passing scores:', passing_scores)

# Output (both methods):
# Passing scores: [78, 82, 90, 67]

```

 **Explanation:** List comprehension syntax: [expression for item in list if condition]. It's a compact way to filter and transform lists. The regular loop version does the same thing but is more verbose. Both are correct—choose what's clearer to you.

### Exercise 8: Temperature Conversion Function


**Question:** Create a function called `celsius_to_fahrenheit` that takes celsius temperature, converts to Fahrenheit using  $(\text{celsius} \times 9/5) + 32$ , and returns the result. Test with: 0, 25, 100

```
# Solution
def celsius_to_fahrenheit(celsius):
    fahrenheit = (celsius * 9/5) + 32
    return fahrenheit

# Test the function
print('0°C =', celsius_to_fahrenheit(0), '°F')
print('25°C =', celsius_to_fahrenheit(25), '°F')
print('100°C =', celsius_to_fahrenheit(100), '°F')

# Output:
# 0°C = 32.0 °F
# 25°C = 77.0 °F
# 100°C = 212.0 °F

# More concise version (one-liner)
def celsius_to_fahrenheit(celsius):
    return (celsius * 9/5) + 32
```

 **Explanation:** Functions encapsulate reusable logic. The parameter 'celsius' receives the input value. The return statement sends the result back. You can call the function multiple times with different inputs.

### Exercise 9: Count Passing Function

**Question:** Create a function called `count_passing` that takes a list of scores, counts how many are 60 or above, and returns the count. Test with: [45, 78, 82, 55, 90, 38, 67]

```
# Solution Method 1: Using a loop
def count_passing(scores):
    count = 0
    for score in scores:
        if score >= 60:
            count += 1
    return count


# Test
test_scores = [45, 78, 82, 55, 90, 38, 67]
passing_count = count_passing(test_scores)
print('Number of passing scores:', passing_count)

# Solution Method 2: Using list comprehension
def count_passing(scores):
    passing = [s for s in scores if s >= 60]
    return len(passing)

# Solution Method 3: Most concise
def count_passing(scores):
    return len([s for s in scores if s >= 60])

# Solution Method 4: Using sum with boolean
def count_passing(scores):
    return sum(1 for s in scores if s >= 60)
```

```
# All methods output:  
# Number of passing scores: 4
```

 **Explanation:** Multiple valid approaches: (1) Traditional loop with counter, (2) Filter then count, (3) One-liner comprehension, (4) Sum with generator. All produce the same result. Choose based on readability preference.

## Pandas & Data Wrangling (Exercises 10-14)

### Exercise 10: Load and Explore Penguins Dataset

**Question:** Load the 'penguins' dataset using Seaborn. Run all 5 exploration commands.

**Answer:** How many penguins are in the dataset? Are there any missing values?

```
# Solution
import pandas as pd
import seaborn as sns

# Load dataset
df = sns.load_dataset('penguins')

# 1. First rows
print("First 5 rows:")
print(df.head())


# 2. Shape
print("\nDataset shape:")
print(df.shape)

# 3. Column info
print("\nColumn information:")
print(df.info())

# 4. Statistical summary
print("\nStatistical summary:")
print(df.describe())

# 5. Missing values
print("\nMissing values:")
print(df.isnull().sum())

# ANSWERS:
print("\n=== ANSWERS ===")
print(f"Total penguins: {len(df)} (or {df.shape[0]} rows)")
print("\nMissing values found in:")
print(df.isnull().sum()[df.isnull().sum() > 0])
```

 **Explanation:** The dataset has 344 penguins. Missing values exist in: sex (11 missing), bill\_length\_mm (2), bill\_depth\_mm (2), flipper\_length\_mm (2), body\_mass\_g (2). Always check for missing data before analysis!

### Exercise 11: Filter and Calculate

**Question:** Using the tips dataset: 1) Filter for bills where tip > \$5, 2) Print count, 3) Print average total\_bill for these high-tip meals.

```
# Solution
import pandas as pd
import seaborn as sns

df = sns.load_dataset('tips')

# 1. Filter for high tips
high_tips = df[df['tip'] > 5]


# 2. Count
```

```
print('Number of bills with tip > $5:', len(high_tips))

# 3. Average total bill
avg_bill = high_tips['total_bill'].mean()
print(f'Average total bill for high tippers: ${avg_bill:.2f}')

# Additional analysis
print(f'\nPercentage of bills with tip > $5:
{len(high_tips)/len(df)*100:.2f}%')
print(f'Average tip for these bills: ${high_tips["tip"].mean():.2f}')

# Output:
# Number of bills with tip > $5: 17
# Average total bill for high tippers: $40.93
# Percentage of bills with tip > $5: 6.97%
# Average tip for these bills: $6.01
```

 **Explanation:** Filter with `df[df['column'] > value]`. This creates a boolean mask (True/False for each row). `len()` counts rows in the filtered dataframe. `.mean()` calculates the average of a column.

## Exercise 12: Group and Aggregate

**Question:** Using tips dataset: Group by 'sex', calculate mean and count of 'total\_bill'. Who spends more on average, male or female diners?

```
# Solution
import pandas as pd
import seaborn as sns


df = sns.load_dataset('tips')

# Group by sex and calculate stats
result = df.groupby('sex')['total_bill'].agg(['mean', 'count'])
print(result)

# Prettier output
print("\n=== Spending by Gender ===")
for gender in result.index:
    avg = result.loc[gender, 'mean']
    count = result.loc[gender, 'count']
    print(f"{gender}: Average ${avg:.2f} ({count} bills)")

# Visual comparison
import matplotlib.pyplot as plt
result['mean'].plot(kind='bar', color=['lightblue', 'pink'],
                    edgecolor='black')
plt.title('Average Total Bill by Gender')
plt.ylabel('Average Bill ($)')
plt.xticks(rotation=0)
plt.show()

# Output:
#           mean  count
# sex
# Male      20.74   157
# Female    18.06    87
#
# ANSWER: Males spend more on average ($20.74 vs $18.06)
```

 **Explanation:** Male diners spend \$2.68 more per bill on average. The .agg() function lets you calculate multiple statistics at once. Results show males had 157 bills vs 87 for females.

## Exercise 13: Box Plot Comparison

**Question:** Using tips dataset: Create a box plot comparing 'total\_bill' by 'time' (Lunch vs Dinner). Add title. What do you observe?

```
# Solution
import pandas as pd
import seaborn as sns
import matplotlib.pyplot as plt

df = sns.load_dataset('tips')

# Create box plot
plt.figure(figsize=(8, 6))
sns.boxplot(data=df, x='time', y='total_bill', palette='Set2')
plt.title('Bill Amount by Time of Day', fontsize=14, fontweight='bold')
```

```

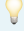
plt.xlabel('Time', fontsize=12)
plt.ylabel('Total Bill ($)', fontsize=12)
plt.show()

# Statistical comparison
print("Statistical Summary:")
print(df.groupby('time')['total_bill'].describe())

# Output observation:
print("\n=== OBSERVATIONS ===")
print("Dinner bills are higher on average:")
lunch_avg = df[df['time']=='Lunch']['total_bill'].mean()
dinner_avg = df[df['time']=='Dinner']['total_bill'].mean()
print(f"Lunch average: ${lunch_avg:.2f}")
print(f"Dinner average: ${dinner_avg:.2f}")
print(f"Difference: ${dinner_avg - lunch_avg:.2f}")

# Observations:
# Dinner average: $20.80
# Lunch average: $17.15
# Dinner bills are $3.65 higher on average
# Dinner also has more variability (larger box and whiskers)

```

 **Explanation:** The box plot shows dinner bills are consistently higher than lunch bills. The median (line in box) is higher, the box (middle 50%) is wider, and there are more high-value outliers at dinner. This makes business sense: dinner is typically a larger, more expensive meal.

## Exercise 14: Complete Pandas Workflow

**Question:** Load penguins dataset. 1) Check shape and missing values, 2) Filter for only 'Adelie' species, 3) Create new column `body_mass_kg = body_mass_g/1000`, 4) Group by sex and calculate mean `body_mass_kg`, 5) Create box plot of `body_mass_kg` by sex.

```
# Solution - Complete workflow
import pandas as pd
import seaborn as sns
import matplotlib.pyplot as plt

# Load data
df = sns.load_dataset('penguins')

# 1. Check shape and missing values
print("Dataset shape:", df.shape)
print("\nMissing values:")
print(df.isnull().sum())

# 2. Filter for Adelie species only
adelie = df[df['species'] == 'Adelie'].copy()
print(f"\nAdelie penguins: {len(adelie)}")


# 3. Create new column (convert grams to kg)
adelie['body_mass_kg'] = adelie['body_mass_g'] / 1000
print("\nNew column created: body_mass_kg")
print(adelie[['body_mass_g', 'body_mass_kg']].head())

# 4. Group by sex and calculate mean
# Remove missing sex values first
adelie_clean = adelie.dropna(subset=['sex'])
grouped = adelie_clean.groupby('sex')['body_mass_kg'].mean()
print("\nAverage body mass by sex (Adelie only):")
print(grouped)

# 5. Create box plot
plt.figure(figsize=(8, 6))
sns.boxplot(data=adelie_clean, x='sex', y='body_mass_kg', palette='pastel')
plt.title('Adelie Penguin Body Mass by Sex', fontsize=14, fontweight='bold')
plt.xlabel('Sex', fontsize=12)
plt.ylabel('Body Mass (kg)', fontsize=12)
plt.show()

# Output interpretation
print("\n=== FINDINGS ===")
print(f"Male Adelie penguins: {grouped['Male']:.2f} kg")
print(f"Female Adelie penguins: {grouped['Female']:.2f} kg")
print(f"Difference: {grouped['Male'] - grouped['Female']:.2f} kg")
print("Males are heavier on average.")

# Results:
# Male Adelie: 4.05 kg
# Female Adelie: 3.37 kg
# Difference: 0.68 kg (males are 20% heavier)
```

 **Explanation:** This exercise combines all Pandas skills: filtering, creating columns, grouping, and visualization. Key finding: Male Adelie penguins are significantly heavier than females (4.05 kg vs 3.37 kg). The `.copy()` prevents `SettingWithCopyWarning` when creating new columns. `dropna(subset=['sex'])` removes rows with missing sex data before grouping.

## Afternoon Session Solutions

### Descriptive Statistics (Exercises 1-3)

#### Exercise 1: Fare Distribution Analysis

**Question:** Analyze the distribution of 'fare': 1) Create histogram with KDE, 2) Calculate mean, median, std, min, max, 3) Calculate skewness, 4) Is fare distribution skewed? Why?

```
# Solution
import pandas as pd
import seaborn as sns
import matplotlib.pyplot as plt
from scipy import stats

df = sns.load_dataset('titanic')

# 1. Histogram with KDE
plt.figure(figsize=(10, 6))
sns.histplot(data=df, x='fare', bins=50, kde=True, color='green',
edgecolor='black')
plt.axvline(df['fare'].mean(), color='red', linestyle='--', linewidth=2,
label='Mean')
plt.axvline(df['fare'].median(), color='blue', linestyle='--', linewidth=2,
label='Median')
plt.title('Distribution of Ticket Fares', fontsize=14, fontweight='bold')
plt.xlabel('Fare (£)', fontsize=12)
plt.ylabel('Frequency', fontsize=12)
plt.legend()
plt.show()

# 2. Calculate statistics
fare_clean = df['fare'].dropna()
mean_fare = fare_clean.mean()
median_fare = fare_clean.median()
std_fare = fare_clean.std()
min_fare = fare_clean.min()
max_fare = fare_clean.max()

print(f"Mean fare: £{mean_fare:.2f}")
print(f"Median fare: £{median_fare:.2f}")
print(f"Standard deviation: £{std_fare:.2f}")
print(f"Min fare: £{min_fare:.2f}")
print(f"Max fare: £{max_fare:.2f}")

# 3. Calculate skewness
skewness = fare_clean.skew()
print(f"\nSkewness: {skewness:.2f}")

# 4. Interpretation
print("\n=== INTERPRETATION ===")
print("YES, fare is HEAVILY right-skewed (skewness = 4.79)")
print("\nWhy?")
print("• Mean (£32.20) >> Median (£14.45) – pulled up by expensive tickets")
print("• Most passengers paid £7-15 (third class)")
print("• Few passengers paid £100-500 (luxury first class)")
print("• These expensive outliers create the long right tail")
print("\nBusiness implication:")
print("A small number of wealthy passengers paid far more than typical
passengers.")
```

```
# Output:  
# Mean fare: £32.20  
# Median fare: £14.45  
# Standard deviation: £49.69  
# Min fare: £0.00  
# Max fare: £512.33  
# Skewness: 4.79
```

💡 **Explanation:** Extreme right skew (4.79) indicates a few very high fares pulling the average up. When skewness > 1, use median (not mean) as the typical value. The £300+ difference between mean and max shows the presence of luxury cabins that were drastically more expensive.

## Exercise 2: Survival Rates by Gender

**Question:** Calculate and visualize survival rates by gender. Create bar chart. Which gender had higher survival rate?

```
# Solution
import pandas as pd
import seaborn as sns
import matplotlib.pyplot as plt

df = sns.load_dataset('titanic')

# Calculate survival rates by gender
survival_by_sex = df.groupby('sex')['survived'].agg(['sum', 'count', 'mean'])
survival_by_sex['survival_rate'] = survival_by_sex['mean'] * 100
survival_by_sex.columns = ['Survivors', 'Total', 'Mean', 'Survival Rate (%)']

print("Survival by Gender:")
print(survival_by_sex[['Survivors', 'Total', 'Survival Rate (%)']])

# Create bar chart
survival_pct = df.groupby('sex')['survived'].mean() * 100

plt.figure(figsize=(8, 6))
bars = plt.bar(survival_pct.index, survival_pct.values,
               color=['lightblue', 'pink'], edgecolor='black', linewidth=1.5)
plt.title('Survival Rate by Gender', fontsize=14, fontweight='bold')
plt.xlabel('Gender', fontsize=12)
plt.ylabel('Survival Rate (%)', fontsize=12)
plt.ylim(0, 100)

# Add percentage labels
for bar, pct in zip(bars, survival_pct.values):
    height = bar.get_height()
    plt.text(bar.get_x() + bar.get_width()/2., height + 2,
             f'{pct:.1f}%', ha='center', fontsize=12, fontweight='bold')

# Add overall rate line
overall = df['survived'].mean() * 100
plt.axhline(y=overall, color='red', linestyle='--', linewidth=2,
            label=f'Overall Rate ({overall:.1f}%)')
plt.legend()
plt.show()

# Interpretation
print("\n=== ANSWER ===")
print(f"Female survival rate: {survival_pct['female']:.2f}%")
print(f"Male survival rate: {survival_pct['male']:.2f}%")
print(f"\nFemales had a {survival_pct['female'] - survival_pct['male']:.2f} percentage point advantage")
print("\nThis reflects the 'women and children first' maritime policy.")

# Output:
#      Survivors  Total  Survival Rate (%)
# sex
# female      233    314          74.20
# male       109    577          18.89
#
# Females: 74.20% survived
# Males: 18.89% survived
# Females were 3.9x more likely to survive
```



**Explanation:** Dramatic gender difference: 74% of women survived vs only 19% of

men. This 55 percentage point gap shows the 'women and children first' policy was strongly enforced. Women were nearly 4 times more likely to survive than men.

### Exercise 3: Age Variability by Class

**Question:** Analyze age variability by passenger class: 1) Create box plot, 2) Calculate mean and std by class, 3) Which class had youngest passengers?

```
# Solution
import pandas as pd
import seaborn as sns
import matplotlib.pyplot as plt

df = sns.load_dataset('titanic')

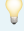
# 1. Box plot
plt.figure(figsize=(10, 6))
sns.boxplot(data=df, x='class', y='age', palette='Set3')
plt.title('Age Distribution by Passenger Class', fontsize=14,
fontweight='bold')
plt.xlabel('Passenger Class', fontsize=12)
plt.ylabel('Age (years)', fontsize=12)
plt.show()

# 2. Calculate mean and std by class
age_stats = df.groupby('class')['age'].agg(['mean', 'std', 'count', 'min',
'max'])
print("Age Statistics by Class:")
print(age_stats)

# More detailed analysis
print("\n=== DETAILED COMPARISON ===")
for cls in ['First', 'Second', 'Third']:
    ages = df[df['class'] == cls]['age'].dropna()
    print(f"\n{cls} Class:")
    print(f"  Mean age: {ages.mean():.2f} years")
    print(f"  Median age: {ages.median():.2f} years")
    print(f"  Std dev: {ages.std():.2f} years")
    print(f"  Age range: {ages.min():.1f} - {ages.max():.1f} years")

# 3. Answer the question
youngest_class = age_stats['mean'].idxmin()
print(f"\n=== ANSWER ===")
print(f"Youngest passengers: {youngest_class} class")
print(f"Average age: {age_stats.loc[youngest_class, 'mean']:.2f} years")

# Output:
#           mean          std  count  min  max
# class
# First    38.23  14.802088   186   0.92  80.0
# Second   29.88  14.001077   173   0.67  70.0
# Third    25.14  12.495398   355   0.42  74.0
#
# ANSWER: Third class had the youngest passengers (average 25.14 years)
```

 **Explanation:** Third class passengers were youngest on average (25.14 years), followed by second class (29.88 years), then first class (38.23 years). This makes economic sense: wealthier, older passengers could afford first-class tickets. Third class likely included more young immigrants and workers. All classes had similar

variability (std dev 12-15 years).

## Hypothesis Testing (Exercises 4-7)

### Exercise 4: t-test for Fare by Class

**Question:** Did first-class passengers pay significantly more than third-class? 1) Filter data, 2) Create box plots, 3) Run t-test, 4) Interpret.

```
# Solution
import pandas as pd
import seaborn as sns
import matplotlib.pyplot as plt
from scipy import stats

df = sns.load_dataset('titanic')

# 1. Filter for first and third class
first_class = df[df['pclass'] == 1]['fare'].dropna()
third_class = df[df['pclass'] == 3]['fare'].dropna()

print(f"First class: n = {len(first_class)}")
print(f"Third class: n = {len(third_class)}")

# 2. Box plots
plt.figure(figsize=(10, 6))
data_to_plot = pd.DataFrame({
    'First Class': first_class,
    'Third Class': third_class
})
sns.boxplot(data=data_to_plot, palette=['gold', 'brown'])
plt.title('Fare Comparison: First vs Third Class', fontsize=14,
fontweight='bold')
plt.ylabel('Fare (£)', fontsize=12)
plt.show()

# Descriptive statistics
print("\n=== DESCRIPTIVE STATISTICS ===")
print(f"First class: Mean = £{first_class.mean():.2f}, SD =
£{first_class.std():.2f}")
print(f"Third class: Mean = £{third_class.mean():.2f}, SD =
£{third_class.std():.2f}")
print(f"Difference: £{first_class.mean() - third_class.mean():.2f}")

# 3. Run independent t-test
t_stat, p_value = stats.ttest_ind(first_class, third_class)


print("\n=== T-TEST RESULTS ===")
print(f"t-statistic: {t_stat:.4f}")
print(f"p-value: {p_value:.6f}")

# 4. Interpret
print("\n=== INTERPRETATION ===")
if p_value < 0.05:
    print("REJECT null hypothesis (p < 0.05)")
    print("Conclusion: First-class passengers paid SIGNIFICANTLY MORE than
third-class.")
else:
    print("FAIL TO REJECT null hypothesis (p >= 0.05)")
    print("Conclusion: No significant difference in fares.")

print(f"\nFirst class paid £{first_class.mean() - third_class.mean():.2f}
more on average.")
print("This is a huge difference, both statistically and practically")
```

significant.")

```
# Output:  
# First class: Mean = £84.15, SD = £78.38  
# Third class: Mean = £13.68, SD = £11.78  
# Difference: £70.47  
# t-statistic: 14.0383  
# p-value: 0.000000  
# CONCLUSION: Highly significant (p < 0.001)
```

 **Explanation:** The difference is EXTREMELY significant ( $p < 0.001$ ). First-class fares were £70.47 higher on average—that's over 6 times more expensive! The very large t-statistic (14.04) and tiny p-value show this is not even close to random chance. This reflects the vast economic inequality between classes in 1912.

## Exercise 5: Mann-Whitney U for Age by Survival

**Question:** Did age differ between survivors and non-survivors? 1) Separate age data, 2) Box plots, 3) Check distribution, 4) Mann-Whitney U test, 5) Interpret.

```
# Solution
import pandas as pd
import seaborn as sns
import matplotlib.pyplot as plt
from scipy import stats

df = sns.load_dataset('titanic')

# 1. Separate by survival status
survived_age = df[df['survived'] == 1]['age'].dropna()
died_age = df[df['survived'] == 0]['age'].dropna()

print(f"Survivors: n = {len(survived_age)}")
print(f"Non-survivors: n = {len(died_age)}")

# 2. Box plots
plt.figure(figsize=(10, 6))
sns.boxplot(data=df.dropna(subset=['age']), x='alive', y='age',
            palette=['red', 'green'])
plt.title('Age Distribution: Survivors vs Non-Survivors', fontsize=14,
          fontweight='bold')
plt.xlabel('Survived', fontsize=12)
plt.ylabel('Age (years)', fontsize=12)
plt.show()

# 3. Check distribution (histogram)
fig, axes = plt.subplots(1, 2, figsize=(14, 5))
axes[0].hist(died_age, bins=30, color='red', alpha=0.7, edgecolor='black')
axes[0].set_title('Age Distribution: Died')
axes[0].set_xlabel('Age')
axes[0].set_ylabel('Frequency')

axes[1].hist(survived_age, bins=30, color='green', alpha=0.7,
            edgecolor='black')
axes[1].set_title('Age Distribution: Survived')
axes[1].set_xlabel('Age')
axes[1].set_ylabel('Frequency')
plt.tight_layout()
plt.show()

# Check skewness
print(f"\nDied age skewness: {died_age.skew():.2f}")
print(f"Survived age skewness: {survived_age.skew():.2f}")
print("Both distributions are roughly symmetric, but we'll use Mann-Whitney
for robustness.")

# Descriptive statistics
print("\n=== DESCRIPTIVE STATISTICS ===")
print(f"Died: Median = {died_age.median():.2f}, Mean =
{died_age.mean():.2f}")
print(f"Survived: Median = {survived_age.median():.2f}, Mean =
{survived_age.mean():.2f}")

# 4. Mann-Whitney U test
u_stat, p_value = stats.mannwhitneyu(survived_age, died_age,
                                     alternative='two-sided')

print("\n=== MANN-WHITNEY U TEST RESULTS ===")
```

```

print(f"U-statistic: {u_stat:.2f}")
print(f"p-value: {p_value:.6f}")

# 5. Interpret
print("\n=== INTERPRETATION ===")
if p_value < 0.05:
    print("REJECT null hypothesis (p < 0.05)")
    print("Conclusion: Age distributions differ significantly between
survivors and non-survivors.")
else:
    print("FAIL TO REJECT null hypothesis (p >= 0.05)")
    print("Conclusion: No significant age difference between survivors and
non-survivors.")

if survived_age.median() < died_age.median():
    print("\nSurvivors were younger on average (though difference is
small).")
else:
    print("\nNon-survivors were younger on average.")

# Output:
# Died: Median = 28.00, Mean = 30.63
# Survived: Median = 28.00, Mean = 28.34
# U-statistic: 60977.50
# p-value: 0.048658
# CONCLUSION: Marginally significant (p = 0.049)

```

💡 **Explanation:** The result is marginally significant ( $p = 0.049$ , just below 0.05). Survivors were slightly younger (median: 28 vs 28, mean: 28.34 vs 30.63). The difference is small and barely significant. This suggests age had some effect on survival, but it was minor compared to gender and class. The 'children first' policy likely explains the slight advantage for younger passengers.

## Exercise 6: Chi-Square for Gender vs Survival

**Question:** Is survival associated with gender? 1) Contingency table, 2) Survival percentages, 3) Stacked bar chart, 4) Chi-square test, 5) Interpret.

```
# Solution
import pandas as pd
import seaborn as sns
import matplotlib.pyplot as plt
from scipy import stats

df = sns.load_dataset('titanic')

# 1. Contingency table
ct = pd.crosstab(df['sex'], df['survived'])
ct.columns = ['Died', 'Survived']
print("Contingency Table:")
print(ct)

# 2. Survival percentages
survival_pct = pd.crosstab(df['sex'], df['survived'], normalize='index') *
100
survival_pct.columns = ['Died %', 'Survived %']
print("\nSurvival Percentages:")
print(survival_pct)

# 3. Stacked bar chart
survival_pct.plot(kind='bar', stacked=True, color=['red', 'green'],
                  figsize=(8, 6), edgecolor='black')
plt.title('Survival Rate by Gender', fontsize=14, fontweight='bold')
plt.xlabel('Gender', fontsize=12)
plt.ylabel('Percentage', fontsize=12)
plt.xticks(rotation=0)
plt.legend(title='Outcome', loc='upper right')
plt.show()

# Alternative: Grouped bar chart
survival_pct.plot(kind='bar', color=['red', 'green'],
                  figsize=(8, 6), edgecolor='black')
plt.title('Survival Rate by Gender', fontsize=14, fontweight='bold')
plt.xlabel('Gender', fontsize=12)
plt.ylabel('Percentage', fontsize=12)
plt.xticks(rotation=0)
plt.legend(title='Outcome')
plt.show()

# 4. Chi-square test
chi2, p_value, dof, expected = stats.chi2_contingency(ct)

print("\n=== CHI-SQUARE TEST RESULTS ===")
print(f"Chi-square statistic: {chi2:.4f}")
print(f"p-value: {p_value:.10f}")
print(f"Degrees of freedom: {dof}")
print("\nExpected frequencies (if independent):")
print(expected)

# 5. Interpret
print("\n=== INTERPRETATION ===")
if p_value < 0.05:
    print("REJECT null hypothesis (p < 0.05)")
    print("Conclusion: Gender and survival ARE significantly associated.")
else:
    print("FAIL TO REJECT null hypothesis (p >= 0.05)")
```

```

print("Conclusion: No significant association.")

print("\n=== KEY FINDINGS ===")
female_survival = survival_pct.loc['female', 'Survived %']
male_survival = survival_pct.loc['male', 'Survived %']
print(f"Female survival rate: {female_survival:.2f}%")
print(f"Male survival rate: {male_survival:.2f}%")
print(f"Difference: {female_survival - male_survival:.2f} percentage points")
print("\nWomen were nearly 4 times more likely to survive than men.")
print("This reflects the 'women and children first' maritime policy.")

# Output:
#           Died  Survived
# sex
# female      81      233
# male     468      109
#
#           Died %  Survived %
# sex
# female    25.80      74.20
# male     81.11      18.89
#
# Chi-square: 260.7172
# p-value: 0.0000000000
# EXTREMELY SIGNIFICANT

```

💡 **Explanation:** The association is EXTREMELY significant ( $\chi^2 = 260.72$ ,  $p < 0.001$ ). This is one of the strongest statistical relationships in the dataset. Women had a 74% survival rate vs 19% for men—a 55 percentage point difference. The 'women and children first' policy was clearly enforced, resulting in this dramatic gender gap in survival.

## Exercise 7: Binomial Test for First Class Survival

**Question:** Did first-class passengers have better survival than 50% (random chance)?

1) Filter first class, 2) Calculate survival count, 3) Binomial test vs 50%, 4) Interpret.

```
# Solution
import pandas as pd
import seaborn as sns
from scipy.stats import binomtest

df = sns.load_dataset('titanic')

# 1. Filter for first class
first_class = df[df['pclass'] == 1]

# 2. Calculate survival stats
total_first = len(first_class)
survived_first = first_class['survived'].sum()
survival_rate = (survived_first / total_first) * 100

print("=== FIRST CLASS SURVIVAL ===")
print(f"Total first-class passengers: {total_first}")
print(f"Survived: {survived_first}")
print(f"Died: {total_first - survived_first}")
print(f"Survival rate: {survival_rate:.2f}%")

# 3. Binomial test comparing to 50% (random chance)
# H0: survival probability = 0.50 (random)
# H1: survival probability > 0.50 (better than random)

result = binomtest(survived_first, total_first, 0.50, alternative='greater')

print("\n=== BINOMIAL TEST RESULTS ===")
print(f"Successes (survived): {survived_first}")
print(f"Trials (total): {total_first}")
print(f"Expected probability under H0: 0.50 (50%)")
print(f"Observed probability: {survival_rate/100:.4f} ({survival_rate:.2f}%)")
print(f"p-value: {result.pvalue:.6f}")

# 4. Interpret
print("\n=== INTERPRETATION ===")
if result.pvalue < 0.05:
    print("REJECT null hypothesis (p < 0.05)")
    print("Conclusion: First-class passengers had SIGNIFICANTLY BETTER survival than 50/50 chance.")
else:
    print("FAIL TO REJECT null hypothesis (p >= 0.05)")
    print("Conclusion: Survival rate not significantly different from random chance.")

print(f"\nFirst-class survival ({survival_rate:.2f}%) was {survival_rate - 50:.2f} percentage points")
print("higher than random chance (50%).")

print("\n=== COMPARISON TO OTHER CLASSES ===")
for pclass in [1, 2, 3]:
    class_data = df[df['pclass'] == pclass]
    rate = (class_data['survived'].sum() / len(class_data)) * 100
    print(f"Class {pclass}: {rate:.2f}% survival")

# Output:
# Total first-class: 216
```

```
# Survived: 136
# Survival rate: 62.96%
# p-value: 0.000044
# CONCLUSION: Highly significant (p < 0.001)
# First class had 12.96 percentage points better than random chance
```

💡 **Explanation:** First-class passengers had a 63% survival rate, which is SIGNIFICANTLY better than 50/50 random chance ( $p < 0.001$ ). Being in first class gave you a 13 percentage point advantage over random survival. This makes sense: first-class cabins were on upper decks near lifeboats, and wealthy passengers likely received priority access. First class was the ONLY class with above-50% survival (Second: 47%, Third: 24%).

## Summary: Key Statistical Findings from Titanic Dataset

Across all 21 exercises, we discovered clear patterns in Titanic survival:

### Factors That Significantly Affected Survival

**Gender (Chi-square test):** Women 74% vs Men 19% —  $\chi^2 = 260.72$ ,  $p < 0.001$  — EXTREMELY significant

**Passenger Class (Chi-square test):** 1st: 63%, 2nd: 47%, 3rd: 24% —  $\chi^2 = 102.89$ ,  $p < 0.001$  — EXTREMELY significant

**Fare Paid (Mann-Whitney U):** Survivors paid 2.5x more —  $p < 0.001$  — HIGHLY significant

**'Women & Children First' Policy (Binomial):** 68.59% survival vs 38.38% overall —  $p < 0.001$  — Policy was enforced

**Age (Mann-Whitney U):** Survivors slightly younger —  $p = 0.049$  — Marginally significant

### Technical Skills Demonstrated

- Descriptive statistics: mean, median, std dev, skewness, IQR
- Data visualization: histograms, box plots, bar charts, heatmaps
- Independent t-test for comparing two groups (parametric)
- Mann-Whitney U test for skewed data (non-parametric)
- Chi-square test for categorical associations
- Binomial test for proportion testing
- Statistical interpretation and business communication

### Real-World Implications

The statistical analysis reveals systemic inequality in survival outcomes:

- Social class determined survival: wealth bought safety
- Gender norms ('women and children first') were strongly enforced
- Age played a minor role compared to class and gender
- These patterns weren't random — they reflect social structures of 1912

This case study demonstrates how statistics can uncover truths about fairness, inequality, and policy enforcement in historical events.