

Chapter 7

Expected Value, Variance, and Samples

7.1 Expected value and variance

Previously, we determined the expected value and variance for a random variable Y , which we can think of as a single observation from a distribution. We will now extend these concepts to a linear function of Y and also the sum of n random variables. We will use these results to derive the expected value and variance of the sample mean \bar{Y} and variance s^2 , and so describe their basic statistical properties. The idea of an unbiased estimator is also expressed in terms of expected values, and we will show that \bar{Y} and s^2 are unbiased estimators of the theoretical mean and variance of Y , i.e., $E[Y]$ and $Var[Y]$. This is true regardless of the distribution of Y .

We begin by reviewing the definition of expected value and variance. Recall that if Y has a discrete distribution, the expected value (theoretical mean) of Y , or $E[Y]$, is given by the equation

$$E[Y] = \sum_y yP[Y = y] = \sum_y yf(y). \quad (7.1)$$

Here $f(y)$ is the probability distribution of Y , with the summation is taken over all possible values of y . If Y has a continuous distribution, the expected value is defined as the integral

$$E[Y] = \int_{-\infty}^{\infty} yf(y)dy, \quad (7.2)$$

where $f(y)$ is the probability density of Y . For both discrete and continuous random variables, the expected value is essentially a weighted average of all possible values of Y , with the weights being probabilities or densities.

We also defined the theoretical variance of a random variable using expectation. The variance of a random variable Y , denoted by $Var[Y]$, is defined as

$$Var[Y] = E[(Y - E[Y])^2] = \sum_y (y - E[Y])^2 P[Y = y] \quad (7.3)$$

$$= \sum_y (y - E[Y])^2 f(y). \quad (7.4)$$

The variance is a measure of the dispersion of the distribution of Y . The variance of a continuous random variable Y is similarly defined as

$$Var[Y] = E[(Y - E[Y])^2] = \int_{-\infty}^{\infty} (y - E[Y])^2 f(y) dy. \quad (7.5)$$

Table 7.1 summarizes the expected value and variance for the different distributions we have examined so far. These quantities are a function of the parameters in the distribution. Note that for the binomial, Poisson, negative binomial and uniform distributions, there is some relationship between $E[Y]$ and $Var[Y]$, because the formulas share the same parameters. For example, in the Poisson distribution the theoretical mean and variance are both equal to λ . This is not the case for the normal distribution, where the mean and variance are two separate parameters.

Table 7.1: Expected value and variance for five common probability distributions

Distribution	Parameters	$E[Y]$	$Var[Y]$
Binomial	l, p	lp	$lp(1 - p)$
Poisson	λ	λ	λ
Negative binomial	m, k	m	$m + m^2/k$
Uniform	a, b	$\frac{a+b}{2}$	$\frac{(b-a)^2}{12}$
Normal	μ, σ^2	μ	σ^2

The significance of this result is that many statistical procedures assume the mean and variance are unrelated, because they are based on the normal

distribution. If we wish to apply these procedures to other distributions, we will need to transform the observations to reduce the relationship between the mean and variance. This type of transformation is known as a **variance-stabilizing transformation** (see Chapter 15).

7.2 Linear functions and sums - expected value and variance

Before we turn to samples, we first need to determine the expected value of a linear function of Y . Let Y be a random variable with any distribution, and define a new variable $Y' = aY + b$, where a and b are constants. This is called a linear function of Y because there is a straight-line relationship between Y' and Y . What is the expected value of Y' , or $E[Y']$? It can be shown that

$$E[Y'] = E[aY + b] = aE[Y] + b. \quad (7.6)$$

Thus, multiplying a random variable by a constant and then adding another constant just shifts the theoretical mean in the same way (Mood et al. 1974). This result holds for random variables with either a discrete or continuous distribution.

Now suppose we have n random variables of any type, Y_1, Y_2, \dots, Y_n , which may or may not be independent. The random variables could also have unequal means and variances, and even different distributions. What is the expected value of the sum of these variables? One can show that

$$E[Y_1 + Y_2 + \dots + Y_n] = E[Y_1] + E[Y_2] + \dots + E[Y_n] = \sum E[Y_i]. \quad (7.7)$$

So, **the expected value of a sum is equal to the sum of the expected values** (Mood et al. 1974).

We will now examine how the theoretical variance is affected by a linear function. Let Y be a variable with any distribution with an associated variance of $Var[Y]$. Define a new random variable $Y' = aY + b$, where a and b are constants. What is the variance of Y' , or $Var[Y']$? It can be shown that

$$Var[Y'] = Var[aY + b] = a^2 Var[Y]. \quad (7.8)$$

This implies that a linear function of a random variable increases its variance by a factor of a^2 , with b playing no role in the variance. This makes intuitive

sense, because multiplying a random variable by a constant (a) should affect its breadth or dispersion, while adding a constant (b) only shifts its location and not its dispersion.

Now suppose we have n random variables of any type, Y_1, Y_2, \dots, Y_n . The random variables can have unequal means and variances, but we will assume they are independent. What is the variance of the sum of these observations? It can be shown that

$$\text{Var}[Y_1 + Y_2 + \dots + Y_n] = \text{Var}[Y_1] + \text{Var}[Y_2] + \dots + \text{Var}[Y_n] = \sum \text{Var}[Y_i]. \quad (7.9)$$

Thus, **the variance of a sum is equal to the sum of the variances** (Mood et al. 1974). As you add more and more random variables together, the variance of the sum also increases. This result only holds when the random variables are independent of each other – if they were dependent a much more complicated formula would be required. This is one advantage of working with a random sample in which the observations are independent, because it simplifies parameter estimation and other statistical procedures (see Chapter 8).

7.3 Sample mean - expected value and variance

We will now use the preceding results to find the expected value and variance of the sample mean. Suppose we have a set of observations Y_1, Y_2, \dots, Y_n drawn from some statistical population, say the body lengths of n randomly selected individuals. The random variables Y_i are independent, and because they are drawn from the same population, they also have the same expected value $E[Y_i]$ and variance $\text{Var}[Y_i]$.

The sample mean is defined using the familiar formula:

$$\bar{Y} = \frac{\sum_{i=1}^n Y_i}{n}. \quad (7.10)$$

What is the expected value of the sample mean or \bar{Y} ? Using our results for sums of variables and linear transformations, we have

$$E[\bar{Y}] = E\left[\frac{\sum Y_i}{n}\right] = \frac{E[\sum Y_i]}{n} = \frac{\sum E[Y_i]}{n} = \frac{nE[Y_i]}{n} = E[Y_i]. \quad (7.11)$$

The expected value of the mean is thus equal to the expected value of the individual variables (Mood et al. 1974).

The fact that $E[\bar{Y}] = E[Y_i]$ means that \bar{Y} is an **unbiased estimator** of the theoretical mean of the distribution of Y_i . In less technical terms, it implies that on average \bar{Y} will be equal to the underlying mean of the random variable Y_i . This is often a desirable property in an estimator, although there are useful biased estimators as well.

We also need to calculate the theoretical variance of the sample mean, written as $Var[\bar{Y}]$. Using the properties of the expected value and variance, we have

$$Var[\bar{Y}] = Var\left[\frac{\sum Y_i}{n}\right] = \frac{Var[\sum Y_i]}{n^2} = \frac{\sum Var[Y_i]}{n^2} = \frac{nVar[Y_i]}{n^2} = \frac{Var[Y_i]}{n}. \quad (7.12)$$

Thus, the variance of the sample mean is the variance of Y_i divided by n (Mood et al. 1974).

What does this result imply? **As you collect larger and larger samples, the variance of the sample mean \bar{Y} becomes smaller.** In other words, \bar{Y} becomes less variable when it includes more data. This result underlies many of the desirable effects of larger sample sizes in statistics, including better estimates of parameters (Chapter 8), smaller confidence intervals (Chapter 9), and statistical tests with more power (Chapter 10).

The standard deviation of the sample mean \bar{Y} is defined to be the square root of the above quantity:

$$\sqrt{Var[\bar{Y}]} = \sqrt{\frac{Var[Y_i]}{n}} = \frac{\sqrt{Var[Y_i]}}{\sqrt{n}}. \quad (7.13)$$

This formula makes it clear that the standard deviation of the mean is a function of the standard deviation of the individual observations and the sample size used in the mean. The common name for this quantity is the **standard error**. In general, a standard error is the standard deviation of a particular statistic, in this case the sample mean \bar{Y} .

7.4 Sample variance - expected value

Recall that the sample variance is defined using the formula

$$s^2 = \frac{\sum_{i=1}^n (Y_i - \bar{Y})^2}{n - 1}. \quad (7.14)$$

It can be shown that $E[s^2] = Var[Y_i]$, implying that the sample variance is an unbiased estimator of the underlying variance of Y_i .

It is important to note that all our results for the sample mean \bar{Y} and variance s^2 hold true for any distribution, not just the normal distribution. The basic requirement is that the observations Y_1, Y_2, \dots, Y_n are randomly drawn from some statistical population, implying they are independent and have the same expected value $E[Y_i]$ and variance $Var[Y_i]$.

7.5 Sample calculations and simulation - SAS demo

As an example of these rules of expectation and variance, suppose that Y has a normal distribution with mean $\mu = 1$ and variance $\sigma^2 = 1$, namely $Y \sim N(1, 1)$. Suppose we want to find the expected value and variance of $Y' = 2Y + 1$. Note that Y' is a linear function of Y with $a = 2$ and $b = 1$. Using the formulas for the expected value and variance of a linear function, we have $E[Y'] = aE[Y] + b = 2E[Y] + 1 = 2(1) + 1 = 3$, and also $Var[Y'] = a^2Var[Y] = 2^2Var[Y] = 4(1) = 4$.

Now suppose we have three variables Y_1, Y_2 , and Y_3 with the same distribution as above, and assumed to be independent. What is the expected value and variance of the sum of these two variables, $Y_1 + Y_2 + Y_3$? Using the formulas for sums of random variables, we have $E[Y_1 + Y_2 + Y_3] = E[Y_1] + E[Y_2] + E[Y_3] = 1 + 1 + 1 = 3$, and $Var[Y_1 + Y_2 + Y_3] = Var[Y_1] + Var[Y_2] + Var[Y_3] = 1 + 1 + 1 = 3$.

We can also calculate the expected value and variance of the sample mean \bar{Y} for Y_1, Y_2 , and Y_3 . Using the preceding results, we have $E[\bar{Y}] = E[Y_i] = 1$, and $Var[\bar{Y}] = Var[Y_i]/n = 1/3$.

We can verify that these theoretical rules for the expected value and variance have some basis in reality by conducting an experiment. Recall that the expected value for a random variable can also be thought of as the sample mean \bar{Y} for an infinite number of observations of that random variable. Similarly, its theoretical variance is the sample variance s^2 of an infinite number of observations. It is easy to generate a very large number of observations using SAS, and then compare the result predicted by these theoretical rules with the sample mean and variance of the observations. The SAS program listed below first generates 1,000 observations having the

Table 7.2: Expected value and variance

Variable	Theory		Simulation	
	$E[\cdot]$	$Var[\cdot]$	\bar{Y}	s^2
Y	1	1	1.032	0.980
Y'	3	4	3.063	3.919
$Y_1 + Y_2 + Y_3$	3	3	3.052	3.069
\bar{Y}	1	1/3	1.017	0.341
s^2	1	-	1.001	-

specified distribution [$Y, Y_i \sim N(1, 1)$] in a data step. Formulas are then used to calculate Y' , $Y_1 + Y_2 + Y_3$, \bar{Y} , and s^2 . The SAS procedure `proc univariate` is then used to calculate the sample mean and variance of these quantities. See SAS output below.

If the theory involving expected values and variances is correct, it should predict the behavior of the mean and variance in this large sample. A comparison between the results predicted using our expected value formulas and the observed simulation results is given in Table 7.2. The theoretical predictions and sample mean and variance are in close agreement.

Notice also from the SAS output that the distributions of Y' , $Y_1 + Y_2 + Y_3$, and \bar{Y} appear to be normally distributed (see Fig. 7.8 - 7.10). In fact, linear functions and sums of normal random variables are always normally distributed, as is the sample mean. This may not be the case for variables with other distributions. We also see that the variance of \bar{Y} is lower than Y (1/3 vs. 1), an important property of this statistic (see Fig. 7.8 vs. 7.10).

SAS Program

```
* Linear.sas;
title 'Demonstration of expected value and variance rules';
data linear;
  * Loop to generate 1000 random observations;
  do i = 1 to 1000;
    a = 2;
    b = 1;
    * Generate y, y1, y2, y3 with N(1,1) distribution;
    mu = 1; sig2 = 1;
    y = sqrt(sig2)*rannor(0) + mu;
    y1 = sqrt(sig2)*rannor(0) + mu;
    y2 = sqrt(sig2)*rannor(0) + mu;
    y3 = sqrt(sig2)*rannor(0) + mu;
    * Calculate a linear function of y, then sum, mean, and s2;
    yprime = a*y + b;
    ysum = y1 + y2 + y3;
    ybar = ysum/3;
    s2 = ((y1-ybar)**2+(y2-ybar)**2+(y3-ybar)**2)/(3-1);
    output;
  end;
run;
* Print simulated data, first 25 observations;
proc print data=linear(obs=25);
run;
* Descriptive statistics, histograms, and normal quantile plots;
proc univariate data=linear;
  var y yprime ysum ybar s2;
  histogram y yprime ysum ybar s2 / vscale=count normal midpoints=-6 to 12 by 0.5;
  qqplot y yprime ysum ybar s2 / normal;
run;
quit;
```

Demonstration of expected value and variance rules

Obs	i	a	b	mu	sig2	y	y1	y2	y3	yprime	ysum	ybar	s2
1	1	2	1	1	1	0.83011	1.06975	1.23689	1.14064	2.66021	3.44729	1.14910	0.00704
2	2	2	1	1	1	1.08774	0.77384	0.79274	0.70118	3.17547	2.26776	0.75592	0.00234
3	3	2	1	1	1	0.84343	0.13826	1.52898	2.23203	2.68687	3.89926	1.29975	1.13537
4	4	2	1	1	1	-0.36077	0.50470	0.84504	2.42651	0.27846	3.77626	1.25875	1.05171
5	5	2	1	1	1	2.04424	1.51003	0.89581	2.95025	5.08848	5.35609	1.78536	1.11203
6	6	2	1	1	1	0.83030	1.43820	3.22092	2.82612	2.66060	7.48524	2.49508	0.87671
7	7	2	1	1	1	0.40681	1.32073	-0.55832	0.50788	1.81361	1.27029	0.42343	0.88806
8	8	2	1	1	1	0.01667	1.08095	2.73403	-0.19628	1.03333	3.61870	1.20623	2.15845
9	9	2	1	1	1	1.50386	0.19115	0.53985	1.64232	4.00773	2.37332	0.79111	0.57383
10	10	2	1	1	1	1.74741	1.90264	2.34152	-0.95909	4.49481	3.28506	1.09502	3.21269

etc.

Figure 7.1: linear.sas - proc print

Demonstration of expected value and variance rules

The UNIVARIATE Procedure
Variable: y

Moments			
N	1000	Sum Weights	1000
Mean	1.03159375	Sum Observations	1031.59375
Std Deviation	0.98978171	Variance	0.97966782
Skewness	0.18525623	Kurtosis	0.03945178
Uncorrected SS	2042.87383	Corrected SS	978.688156
Coeff Variation	95.9468494	Std Error Mean	0.03129965

Figure 7.2: linear.sas - proc univariate

Demonstration of expected value and variance rules

The UNIVARIATE Procedure
Variable: yprime

Moments			
N	1000	Sum Weights	1000
Mean	3.06318751	Sum Observations	3063.18751
Std Deviation	1.97956341	Variance	3.91867129
Skewness	0.18525623	Kurtosis	0.03945178
Uncorrected SS	13297.8703	Corrected SS	3914.75262
Coeff Variation	64.6242976	Std Error Mean	0.06259929

Figure 7.3: linear.sas - proc univariate

Demonstration of expected value and variance rules

The UNIVARIATE Procedure
Variable: ysum

Moments			
N	1000	Sum Weights	1000
Mean	3.05218781	Sum Observations	3052.18781
Std Deviation	1.75177477	Variance	3.06871485
Skewness	-0.0287682	Kurtosis	-0.1007495
Uncorrected SS	12381.4965	Corrected SS	3065.64613
Coeff Variation	57.3940688	Std Error Mean	0.05539598

Figure 7.4: linear.sas - proc univariate

Demonstration of expected value and variance rules

The UNIVARIATE Procedure
Variable: ybar

Moments			
N	1000	Sum Weights	1000
Mean	1.01739594	Sum Observations	1017.39594
Std Deviation	0.58392492	Variance	0.34096832
Skewness	-0.0287682	Kurtosis	-0.1007495
Uncorrected SS	1375.72184	Corrected SS	340.627348
Coeff Variation	57.3940688	Std Error Mean	0.01846533

Figure 7.5: linear.sas - proc univariate

Demonstration of expected value and variance rules

The UNIVARIATE Procedure
Variable: s2

Moments			
N	1000	Sum Weights	1000
Mean	1.00997429	Sum Observations	1009.97429
Std Deviation	1.01419003	Variance	1.02858142
Skewness	1.94937753	Kurtosis	4.73366341
Uncorrected SS	2047.6009	Corrected SS	1027.55284
Coeff Variation	100.417411	Std Error Mean	0.0320715

Figure 7.6: linear.sas - proc univariate




Figure 7.7: linear.sas - proc univariate




Figure 7.8: linear.sas - proc univariate




Figure 7.9: linear.sas - proc univariate




Figure 7.10: linear.sas - proc univariate

7.6 Central limit theorem

Suppose we randomly draw a sample Y_1, Y_2, \dots, Y_n of size n from some statistical population. In this situation, the observations are independent and have a common expected value $E[Y_i]$ and variance $Var[Y_i]$. **They may have any probability distribution, known or unknown.**

The **central limit theorem** states that the distribution of the sample mean of these random variables, namely \bar{Y} , approaches a normal distribution with mean $E[Y_i]$ and variance $Var[Y_i]/n$ as the sample size n becomes large (Mood et al. 1974). In particular, we have $\bar{Y} \sim N(E[Y_i], Var[Y_i]/n)$ for large n . The central limit theorem also holds for sums of random variables, and in this case we have $\sum Y_i \sim N(nE[Y_i], nVar[Y_i])$ for large n . **These results are true for any probability distribution - \bar{Y} and $\sum Y_i$ will have a normal distribution for large sample sizes.** Note also that the variance of \bar{Y} decreases as the sample size n increases. We would also expect this from our earlier results concerning the variance of \bar{Y} .

7.6.1 Central limit theorem - SAS demo

The operation of the central limit theorem can be demonstrated in a simple experiment using a SAS program (see below). The program models Y as a Poisson random variable with $\lambda = 1$, implying $E[Y_i] = 1$ and $Var[Y_i] = 1$. Sample means are then generated for different sample sizes, ranging from $n = 1$ to $n = 50$, in a SAS `data` step. A total of 1,000 sample means are generated for each value of n in the simulation. The program then used `proc univariate` to calculate summary statistics for these data, as well as histograms and normal quantile plots (not shown). See SAS output below.

Examining the histograms, we see that as n increases the distribution of \bar{Y} approaches the normal distribution. A sample size of $n = 50$ appears sufficient to produce a distribution almost indistinguishable from normal. What is especially interesting here is that fact that the Poisson is a discrete random variable, yet the distribution of \bar{Y} approaches the normal distribution, a continuous random variable.

We also observe that the variance of \bar{Y} decreases as the sample size n increases, as predicted by the central limit theorem and our earlier results on the variance of \bar{Y} . See Table 7.3.

Table 7.3: Mean and variance of \bar{Y}

n	Theory		Simulation	
	$E[Y_i]$	$Var[Y_i]/n$	Mean of \bar{Y}	Variance of \bar{Y}
1	1.000	1.000	0.993	1.036
5	1.000	0.200	0.994	0.213
10	1.000	0.100	1.001	0.111
50	1.000	0.020	0.995	0.019

SAS Program

```

* central_limit_theorem.sas;
title 'Demonstration of central limit theorem in action';
data cntrlmt;
  * Loop to generate 1000 random observations;
  do i = 1 to 1000;
    * A single Poisson observations with lambda = 1;
    y1 = ranpoi(0,1);
    * Mean of 5 Poisson observations;
    y5 = 0;
    do j = 1 to 5;
      y5 = y5 + ranpoi(0,1);
    end;
    y5 = y5/5;
    * Mean of 10 Poisson observations;
    y10 = 0;
    do j = 1 to 10;
      y10 = y10 + ranpoi(0,1);
    end;
    y10 = y10/10;
    * Mean of 50 Poisson observations;
    y50 = 0;
    do j = 1 to 50;
      y50 = y50 + ranpoi(0,1);
    end;
    y50 = y50/50;
    * Mean of 100 Poisson observations;
    output;
  end;
  drop i j;
run;
* Print simulated data (first 25 observations);
proc print data=cntrlmt(obs=25);

```

```
run;
* Descriptive statistics, histograms, and normal quantile plots;
proc univariate data=cntrlmt;
  var y1 y5 y10 y50;
  histogram y1 y5 y10 y50 / vscale=count normal
  qqplot y1 y5 y10 y50 / normal;
  symbol1 h=3;
run;
quit;
```

Demonstration of central limit theorem in action

Obs	y1	y5	y10	y50
1	1	0.6	0.8	1.00
2	0	1.0	0.6	0.82
3	3	1.2	1.1	1.08
4	1	2.4	0.6	1.00
5	1	0.4	0.6	1.10
6	3	0.6	1.0	1.16
7	1	0.6	0.8	0.96
8	1	1.0	0.5	1.10
9	4	1.8	0.7	0.72
10	0	0.4	0.7	1.00

etc.

Figure 7.11: central_limit_theorem.sas - proc print

Demonstration of central limit theorem in action

The UNIVARIATE Procedure
Variable: y1

Moments			
N	1000	Sum Weights	1000
Mean	0.993	Sum Observations	993
Std Deviation	1.01783446	Variance	1.03598699
Skewness	1.06414174	Kurtosis	1.21139016
Uncorrected SS	2021	Corrected SS	1034.951
Coeff Variation	102.500953	Std Error Mean	0.03218675

Figure 7.12: central_limit_theorem.sas - proc univariate

Demonstration of central limit theorem in action

The UNIVARIATE Procedure
Variable: y5

Moments			
N	1000	Sum Weights	1000
Mean	0.994	Sum Observations	994
Std Deviation	0.46123393	Variance	0.21273674
Skewness	0.32859611	Kurtosis	-0.3039864
Uncorrected SS	1200.56	Corrected SS	212.524
Coeff Variation	46.4018037	Std Error Mean	0.0145855

Figure 7.13: central_limit_theorem.sas - proc univariate

Demonstration of central limit theorem in action

The UNIVARIATE Procedure
Variable: y10

Moments			
N	1000	Sum Weights	1000
Mean	1.0009	Sum Observations	1000.9
Std Deviation	0.33355727	Variance	0.11126045
Skewness	0.36056018	Kurtosis	0.23859801
Uncorrected SS	1112.95	Corrected SS	111.14919
Coeff Variation	33.3257336	Std Error Mean	0.01054801

Figure 7.14: central_limit_theorem.sas - proc univariate

Demonstration of central limit theorem in action

The UNIVARIATE Procedure
Variable: y50

Moments			
N	1000	Sum Weights	1000
Mean	0.99538	Sum Observations	995.38
Std Deviation	0.13877144	Variance	0.01925751
Skewness	0.05239973	Kurtosis	0.05542422
Uncorrected SS	1010.0196	Corrected SS	19.2382556
Coeff Variation	13.9415542	Std Error Mean	0.00438834

Figure 7.15: central_limit_theorem.sas - proc univariate




Figure 7.16: central_limit_theorem.sas - proc univariate




Figure 7.17: central_limit_theorem.sas - proc univariate




Figure 7.18: central_limit_theorem.sas - proc univariate




Figure 7.19: central_limit_theorem.sas - proc univariate

7.7 Applications of the central limit theorem

The central limit theorem provides a potential explanation why so many biological variables like the length of an organism and other continuous variables are apparently normal in distribution. These variables are often under the control of multiple genes and environmental factors that can behave like sums and means of random variables, and so their combined effect should generate a normal distribution of outcomes by the central limit theorem (Hartl & Clark 1989).

The theorem also applies to measurements of ecological variables like population density. To estimate population density, we often average the results of several quadrats (or whatever sampling units) to yield a single number for a given location. By the central limit theorem, these average densities will have a normal distribution for sufficiently large n .

Most of the statistical methods we will study are based on the assumption that the observations in a study or experiment have a normal distribution. This would seem a risky assumption, since many natural processes yield random variables that are not strictly normal, some examples being count data that are better modeled using the binomial and Poisson distributions. However, the tests themselves are often based on means that are assumed to have a normal distribution. The central limit theorem guarantees these means are normal provided sample sizes are sufficiently large. Thus, statistical tests based on normality should be valid for non-normal data given large enough sample sizes (see Stewart-Oaten 1995 for further discussion).

The central limit may not be sufficient to guarantee normality for smaller sample sizes, and so other approaches may be needed. One possibility would be a transformation of the observations to make their distribution closer to normal (Chapter 15). If that fails, there are nonparametric statistical procedures (Chapter 16) that are valid for any distribution, as well as ones that allow the use of other probability distributions.

7.8 References

- Hartl, D. L. & Clark, A. G. *Principles of Population Genetics, Second Edition*. Sinauer Associates, Inc., Sunderland, MA.
- Mood, A. M., Graybill, F. A. & Boes, D. C. (1974) *Introduction to the Theory of Statistics*. McGraw-Hill, Inc., New York, NY.
- Stewart-Oaten, A. (1995) Rules and judgments in statistics: three examples. *Ecology* 76: 2001-2009.

7.9 Problems

- Let Y_1 , Y_2 , and Y_3 be three independent random variables with $E[Y_i] = 2$ and $Var[Y_i] = 1$. Using the rules for expected value and variance, calculate the expected value and variance of the following quantities:
 - $3Y_1 + 1$.
 - $Y_1 + Y_2 + Y_3$.
 - $(Y_1 + Y_2 + Y_3)/3$.
- Suppose that Y_1 , Y_2 , and Y_3 are three independent random variables, with $E[Y_i] = 3$ and $Var[Y_i] = 2$. Using the rules for expected value and variance, calculate the expected value and variance of the following quantities:
 - $0.5Y_2 + 2$.
 - $(Y_1 + Y_2 + Y_3)/3$.
 - $2(Y_1 + Y_2) + 3$.
- The exponential distribution is often used to model the time until an event happens, such as the radioactive decay of an atom or mortality processes in population models. The probability density for the exponential distribution is defined as

$$f(y) = \frac{e^{-y/\lambda}}{\lambda} \quad (7.15)$$

for $y \geq 0$. The distribution has one parameter, λ , which is the mean decay time ($E[Y] = \lambda$). A single random observation with an exponential distribution can be generated in SAS using the expression `ranexp(0)*lambda`. Modify the program `central_limit.sas` so that it generates exponential observations instead of Poisson ones, using $\lambda = 2$. Discuss how the distribution of \bar{Y} changes as the sample size increases.

