Chapter 405

Tests for One Exponential Mean

Introduction

This program module designs studies for testing hypotheses about the mean of the exponential distribution. Such tests are often used in *reliability acceptance testing*, also called *reliability demonstration testing*.

Results are calculated for plans that are *time censored* or *failure censored*, as well as for plans that use *with replacement* or *without replacement* sampling. We adopt the basic methodology outlined in Epstein (1960), Juran (1979), Bain and Engelhardt (1991), and Schilling (1982).

Technical Details

The test procedures described here make the assumption that lifetimes follow the exponential distribution. The density of the exponential distribution is written as

$$f(t) = \frac{1}{\theta}\exp\left(-\frac{t}{\theta}\right)$$

The parameter θ is interpreted as average failure time, mean time to failure (MTTF), or mean time between failures (MTBF). Its reciprocal is the failure rate.

The reliability, or probability that a unit continues running beyond time *t*, is

$$R(t) = e^{-\frac{t}{\theta}}$$

Hypothesis Test

The relevant statistical hypothesis is $H_0: \theta_0 = \theta_1$ versus the one-sided alternative $H_1: \theta_0 > \theta_1$. Here, θ_0 represents an acceptable (high) mean life usually set from the point of view of the producer and θ_1 represents some unacceptable (low) mean life usually set from the point of view of the consumer. The test procedure is to reject the null hypothesis if the observed mean life $\hat{\theta}$ is larger than a critical value selected to meet the error rate criterion.

The error rates are often interpreted in reliability testing as *risks*. The *consumer* runs the risk that the study will fail to reject products that have a reliability less than they have specified. This *consumer risk* is β. Similarly, the *producer* runs the risk that the study will reject products that actually meet the consumer's requirements. This *producer risk* is α.

Fixed-Failure Sampling Plans

Fixed failure plans are those in which a specified number of items, *n*, are observed until a specified number of items, r_0, fail. The length of the study t_0 is random. Failed items may, or may not, be immediately replaced (*with replacement* versus *without replacement*).

The test statistic is the observed mean life $\hat{\theta}$ which is computed using

$$\hat{\theta} = \frac{\sum_{i=\text{all test items}} t_i}{r_0}$$

where t_i is the elapsed time that the *i*th item is tested, whether measured until failure or until the study is completed.

For both with-replacement and without-replacement sampling, $\hat{\theta}$ follows the two-parameter gamma distribution with density

$$g(y|r_0,\theta) = \frac{1}{(r_0-1)!}\left(\frac{r_0}{\theta}\right)^r y^{r_0-1}e^{-r_0 y/\theta}$$

This may be converted to a standard, one-parameter gamma using the transformation

$$x = r_0 y / \theta$$

However, because chi-square tables were more accessible, and because the gamma distribution may be transformed to the chi-square distribution, most results in the statistical literature are based on the chi-square distribution. That is, $2r_0\hat{\theta}/\theta$ is distributed as a chi-square random variable with $2r_0$ degrees of freedom.

Assuming that the testing of all *n* items begins at the same instant, the expected length of time needed to observe the first r_0 failures is

$$E(t_0) = \begin{cases} \theta \sum_{i=1}^{r_0} \frac{1}{n-i+1} & \text{without replacement} \\ \frac{\theta r_0}{n} & \text{with replacement} \end{cases}$$

If you choose to solve the without replacement equation for *n*, you can make use of the approximation

$$\sum_{i=1}^{r} \frac{1}{n-i+1} \approx \log_e\left(\frac{n+0.5}{n-r+0.5}\right)$$

Using the above results, sampling plans that meet the specified producer and consumer risk values may be found using the result (see Epstein (1960) page 437) that r_0 is the smallest integer such that

$$\frac{\chi^2_{\alpha,2r_0}}{\chi^2_{1-\beta,2r_0}} \geq \frac{\theta_1}{\theta_0} \quad \text{for testing } H_1: \theta_0 > \theta_1$$

and

$$\frac{\chi^2_{\beta,2r_0}}{\chi^2_{1-\alpha,2r_0}} \geq \frac{\theta_0}{\theta_1} \quad \text{for testing } H_1: \theta_0 < \theta_1$$

Note that the above formulation depends on r_0 but not *n*. An appropriate value of *n* can be found by considering $E(t_0)$. Two options are available.

1. The value of *n* is set (perhaps on economic grounds) and the value of $E(t_0)$ is calculated.
2. The value of $E(t_0)$ is set and the value of *n* is calculated.

Fixed-Time Sampling Plans

Fixed Time plans refer to those in which a specified number of items *n* are observed for a fixed length of time t_0. The number of items failing *r* is recorded. Sampling can be with or without replacement. The accept/reject decision can be based on *r* or the observed mean life $\hat{\theta}$ which is computed using

$$\hat{\theta} = \frac{\sum_{i=\text{all test items}} t_i}{r}$$

where t_i is the time that the *i*th item is being tested, whether measured until failure or until the study is completed.

With Replacement Sampling

If failed items are immediately replaced with additional items, the distribution of *r* (and $\hat{\theta}$, since $\hat{\theta} = nt_0 / r$) follows the Poisson distribution. The probability distribution of *r* is given by the Poisson probability formula

$$P(r \le r_0 | r, \theta) = \sum_{i=0}^{r} \frac{(nt_0 / \theta)^i}{i!} e^{-nt_0/\theta}$$

Thus, values of *n* and t_0 can be found which meet the α and β requirements.

Without Replacement Sampling

If failed items are not replaced, the distributions of *r* and $\hat{\theta}$ are different and thus the power and sample size calculations depend on which statistic will be used. The probability distribution of *r* is given by the binomial formula

$$P(r \le r_0 | r, \theta) = \sum_{i=0}^{r} \binom{n}{i} p^i (1-p)^{n-i}$$

where

$$p = 1 - e^{-t_0/\theta}$$

Thus, values of *n* and t_0 can be found which meet the α and β requirements. Note that this formulation ignores the actual failure times.

If $\hat{\theta}$ will be used as the test statistic, power calculations must be based on it. Bartholomew (1963) gave the following results for the case $r > 0$.

$$\Pr(\hat{\theta} \ge \theta_C) = \frac{1}{1 - e^{-nt_0/\theta}} \sum_{k=1}^{n} \binom{n}{k} \sum_{i=0}^{k} \binom{k}{i} (-1)^i \exp\left\{ -\frac{t_0}{\theta}(n-k+i) \right\} \int_W^{\infty} g(x) \mathrm{d}x$$

where $g(x)$ is the chi-square density function with $2k$ degrees of freedom and

$$W = \frac{2k}{\theta}\left\langle \theta_C - \frac{t_0}{k}(n-k+i) \right\rangle$$

$$\langle X \rangle = \begin{cases} X \text{ if } X > 0 \\ 0 \text{ otherwise} \end{cases}$$

The above equation is numerically unstable for large values of N, so we use the following approximation also given by Bartholomew (1963). This approximation is used when $N > 30$ or when the exact equation cannot be calculated. Bain and Engelhardt (1991) page 140 suggest that this normal approximation can be used when $p > 0.5$

$$z = \frac{u\sqrt{np}}{\sqrt{1 - \frac{2u(1-p)\log_e(1-p)}{p} + (1-p)u^2}}$$

where

$$u = \frac{\hat{\theta} - \theta}{\theta}$$

$$p = 1 - e^{-t_0/\theta}$$

Procedure Options

This section describes the options that are specific to this procedure. These are located on the Design tab. For more information about the options of other tabs, go to the Procedure Window chapter.

Design Tab

The Design tab contains most of the parameters and options that you will be concerned with.

Solve For

Solve For

This option specifies the parameter to be solved for from the other parameters. Under most situations, you will select either *Power* or *Sample Size*.

Select *Sample Size* when you want to calculate the sample size needed to achieve a given power and alpha level.

Select *Power* when you want to calculate the power of an experiment or test.

Test

Alternative Hypothesis

Specify the alternative hypothesis of the test. Since the null hypothesis is equality (a difference between theta0 and theta1 of zero), the alternative is all that needs to be specified. Usually, a one-tailed option is selected for these designs. In fact, the two-tailed options are only available for time terminated experiments.

Sampling Plan

Replacement Method

When failures occur, they may be immediately replaced (With Replacement) with new items or not (Without Replacement). One of the assumptions of the exponential distribution is that the probability of failure does not depend on the previous running time. That is, it is assumed that there is no wear-out. Adopting ‘with replacement’ sampling will shorten the elapsed time of an experiment that is failure terminated.

Termination Criterion

This option specifies the method used to terminate the study or experiment. There are two basic choices:

- **Fixed failures (*r*)**

  Terminate after *r* failures occur. This is also called *failure terminated* or *Type-II Censoring*.

- **Fixed time (*t*0)**

  Terminate after an elapsed time of *t*0. This is also called *time terminated* or *Type-I Censoring*. This is the most common.

In fixed failure sampling, *N* may be fixed while *t*0 varies or *t*0 may be fixed while *N* varies. All that matters is the product of these two quantities.

In fixed time sampling, two test statistics are available: *r* and theta-hat. When sampling is without replacement, tests based on theta-hat are more powerful (require smaller sample size).

r (Number of Failures)

Enter one or more values for the *rejection number* of the test. If *r* or more items fail, the null hypothesis that Theta0 = Theta1 is rejected in favor of the alternative the Theta0 > Theta1.

Note that this value is ignored for time terminated experiments, because the appropriate value is calculated. This value is also ignored in some situations in failure terminated experiments.

t0 (Test Duration Time)

Enter one or more values for the duration of the test. This value may be interpreted as the exact duration time, *t*0, or the expected duration time, E(*t*0), depending on the Termination Criterion and Replacement Method selected.

These values must be positive and in the same time units as Theta0 and Theta1.

E(t0) based on Theta1

When the experiment is failure terminated, the expected waiting time until *r* failures are observed, E(*t*0), is calculated. This value depends on the value of theta, the mean life. When checked, E(*t*0) calculations are based on Theta1. When unchecked, E(*t*0) calculations are based on Theta0. Either choice may be reasonable in a given situation.

Power and Alpha

Power = 1–Beta (Beta is Consumer’s Risk)

This option specifies one or more values for power. Power is the probability of rejecting a false null hypothesis, and is equal to one minus Beta. Beta (consumer’s risk) is the probability of a type-II error, which occurs when a false null hypothesis is not rejected. In this procedure, a type-II error occurs when you fail to reject the null hypothesis of equal probabilities of the event of interest when in fact they are different.

Values must be between zero and one. Historically, the value of 0.80 (Beta = 0.20) was used for power. Now, 0.90 (Beta = 0.10) is also commonly used.

A single value may be entered here or a range of values such as *0.8 to 0.95 by 0.05* may be entered.

Alpha (Producer's Risk)

This option specifies one or more values for the probability of a type-I error (alpha), also called the producer's risk. A type-I error occurs when you reject the null hypothesis of equal probabilities when in fact they are equal.

Values must be between zero and one. Historically, the value of 0.05 has been used for alpha. This means that about one test in twenty will falsely reject the null hypothesis. You should pick a value for alpha that represents the risk of a type-I error you are willing to take in your experimental situation.

You may enter a range of values such as *0.01 0.05 0.10* or *0.01 to 0.10 by 0.01*.

Sample Size

N (Sample Size)

Enter one or more values for the sample size N, the number of items in the study. Note that the sample size is arbitrary for sampling plans that are terminated after a fixed number of failures are observed.

You may enter a range such as 10 to 100 by 10 or a list of values separated by commas or blanks.

Effect Size

Theta0 (Baseline Mean Life)

Enter one or more values for the *mean life* under the null hypothesis. This is sometimes called the *producer's mean life*. This value is usually scaled in terms of elapsed time such as hours, days, or years. Of course, all time values must be on the same time scale.

Note that the value of theta may be calculated from the estimated probability of failure using the relationship

$$\mathrm{P}(\text{Failure}) = 1 - e^{-t0/\theta}$$

so that

$$\theta = \frac{-t0}{\ln\left(1 - P(\text{Failure})\right)}$$

Only positive values are valid. You may enter a range of values such as '10 20 30' or '100 to 1000 by 100.'

Because the exponential function is used in the calculations, try to scale the numbers so they are less than 100. For example, instead of 720 days, use 7.2 hundreds of days. This will help to avoid numerical problems during the calculations.

Theta1 (Alternative Mean Life)

Enter one or more values for the *mean life* under the alternative hypothesis. This is sometimes called the *consumer's mean life*. This value is usually scaled in terms of elapsed time such as hours, days, or years. Of course, all time values must be on the same time scale.

Note that the value of theta may be calculated from the estimated probability of failure using the relationship

$$\mathrm{P}(\text{Failure}) = 1 - e^{-t0/\theta}$$

so that

$$\theta = \frac{-t0}{\ln\left(1 - P(\text{Failure})\right)}$$

Any positive values are valid. You may enter a range of values such as '10 20 30' or '100 to 1000 by 100.'

Because the exponential function is used in the calculations, try to scale the numbers so they are less than 100. For example, instead of 720 days, use 7.2 hundreds of days. This will help to avoid numerical problems during the calculations.

Example 1 – Power for Several Sample Sizes

This example will calculate power for a time terminated, without replacement study in which the results will be analyzed using theta-hat. The study will be used to test the alternative hypothesis that Theta0 > Theta1, where Theta0 = 2.0 days and Theta1 = 1.0 days. The test duration is 1.0 days. Funding for the study will allow for a sample size of up to 40 test items. The researchers decide to look at sample sizes of 10, 20, 30, and 40. Significance levels of 0.01 and 0.05 will be considered.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Tests for One Exponential Mean** procedure window by expanding **Means**, then **One Mean**, then clicking on **Test (Inequality)**, and then clicking on **Tests for One Exponential Mean**. You may then make the appropriate entries as listed below, or open **Example 1** by going to the **File** menu and choosing **Open Example Template**.

Option	Value
Design Tab	
Solve For	**Power**
Alternative Hypothesis	**Ha: Theta0 > Theta1**
Replacement Method	**Without Replacement**
Termination Criterion	**Fixed Time using Theta-hat**
t0 (Test Duration Time)	**1**
Alpha	**0.01 0.05**
N (Sample Size)	**10 to 40 by 10**
Theta0 (Baseline Mean Life)	**2**
Theta1 (Alternative Mean Life)	**1**

Annotated Output

Click the Calculate button to perform the calculations and generate the following output.

Numeric Results

Numeric Results
Test Based on Theta-hat with Fixed Running Time t0 and Without Replacement Sampling.
H0: Theta = Theta0. Ha: Theta = Theta1 < Theta0. Reject H0 if Theta-hat ≤ ThetaC.

Power	N	Time t0	Theta0	Theta1	Target Alpha	Actual Alpha	Target Beta	Actual Beta	Theta C
0.21695	10	1.000	2.0	1.0	0.01000	0.01000		0.78305	0.7
0.45485	20	1.000	2.0	1.0	0.01000	0.01000		0.54515	1.0
0.67159	30	1.000	2.0	1.0	0.01000	0.01000		0.32841	1.1
0.80628	40	1.000	2.0	1.0	0.01000	0.01000		0.19372	1.2
0.46940	10	1.000	2.0	1.0	0.05000	0.05000		0.53060	1.0
0.71828	20	1.000	2.0	1.0	0.05000	0.05000		0.28172	1.2
0.86665	30	1.000	2.0	1.0	0.05000	0.05000		0.13335	1.3
0.93730	40	1.000	2.0	1.0	0.05000	0.05000		0.06270	1.4

Report Definitions
Power is the probability of rejecting a false null hypothesis.
N is the size of the sample drawn from the population.
Alpha is the probability of rejecting a true null hypothesis.
Beta is the probability of accepting a false null hypothesis.
Theta0 is the Mean Life under the null hypothesis.
Theta1 is the Mean Life under the alternative hypothesis.
t0 is the test duration time. It provides the scale for Theta0 and Theta1.
r is the number of failures.

Summary Statements
A sample size of 10 achieves 22% power to detect the difference between the null hypothesis mean lifetime of 2.0 and the alternative hypothesis mean lifetime of 1.0 at a 0.01000 significance level (alpha) using a one-sided test based on the elapsed time. Failing items are not replaced with new items. The study is terminated when it has run for 1.000 time units.

This report shows the power for each of the scenarios. The critical value, Theta C, is also provided.

Plots Section

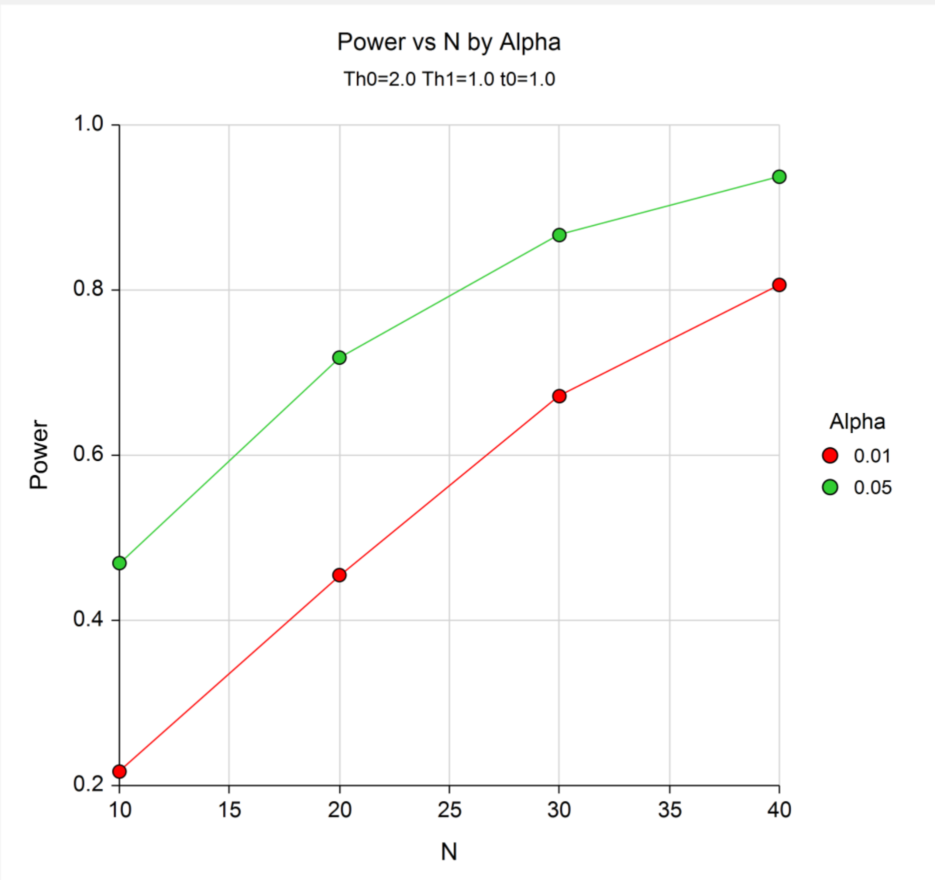

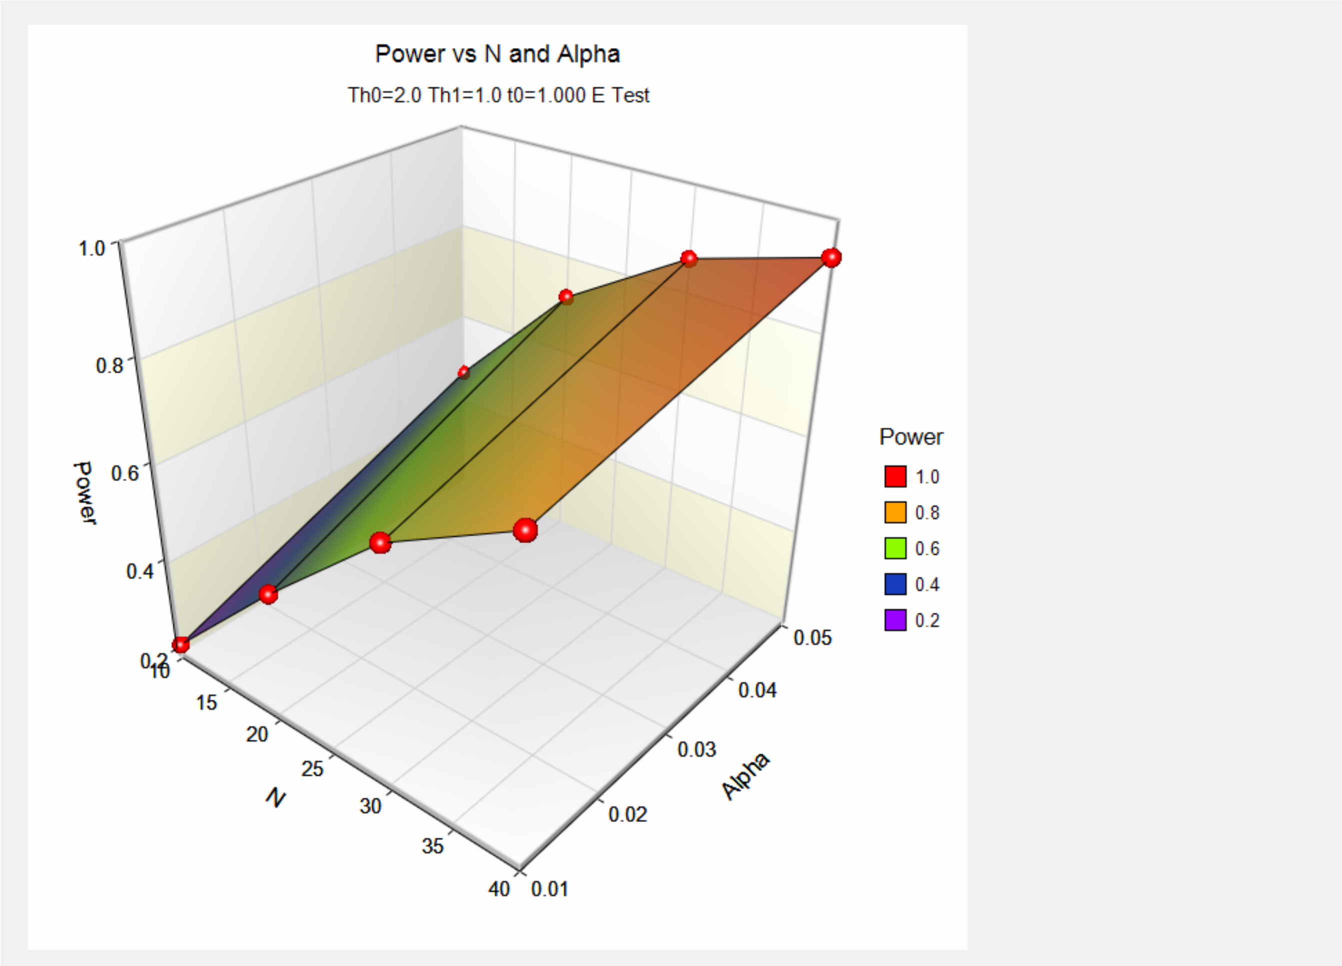

These plots show the relationship between power and sample size.

Example 2 – Validation using Epstein

Epstein (1960), page 438, presents a table giving values of *r* necessary to meet risk criteria for various values of alpha, beta, theta0, and theta1 for the fixed failures case. Specifically, when theta0 = 5, theta1 = 2, beta = 0.05, and alpha = 0.01, 0.05, and 0.10, he finds *r* = 21, 14, and 11. We will now duplicate these results.

Setup

This section presents the values of each of the parameters needed to run this example. First, from the PASS Home window, load the **Tests for One Exponential Mean** procedure window by expanding **Means**, then **One Mean**, then clicking on **Test (Inequality)**, and then clicking on **Tests for One Exponential Mean**. You may then make the appropriate entries as listed below, or open **Example 2** by going to the **File** menu and choosing **Open Example Template**.

Option	Value
Design Tab	
Solve For	**r**
Alternative Hypothesis	**Ha: Theta0 > Theta1**
Replacement Method	**Without Replacement**
Termination Criterion	**Fixed Failures, Fixed E(t0)**
t0 (Test Duration Time)	**1**
E(t0) based on Theta1	**Unchecked**
Power	**0.95**
Alpha	**0.01 0.05 0.10**
Theta0 (Baseline Mean Life)	**5**
Theta1 (Alternative Mean Life)	**2**

Output

Click the Calculate button to perform the calculations and generate the following output.

Numeric Results

Numeric Results
Test Based on Fixed Failures r, Fixed Expected Time E(t0), and Without Replacement Sampling.
H0: Theta = Theta0. Ha: Theta = Theta1 < Theta0. Reject H0 if r ≥ r0.

Power	r0 / N	Time E(t0)	Theta0	Theta1	Target Alpha	Actual Alpha	Target Beta	Actual Beta
0.95841	**21**/115	1.000	5.0	2.0	0.01000	0.01000	0.05000	0.04159
0.95956	**14**/77	1.000	5.0	2.0	0.05000	0.05000	0.05000	0.04044
0.96221	**11**/60	1.000	5.0	2.0	0.10000	0.10000	0.05000	0.03779

PASS calculates 21, 14, and 11 for *r* as in Epstein.

We should note that occasionally our results differ from those of Epstein. We have checked a few of these carefully by hand, and in every case we have found our results to be correct.