LA-UR 97-4825

Title: A CHI-SQUARE GOODNESS-OF-FIT TEST FOR NON-IDENTICALLY DISTRIBUTED RANDOM VARIABLES: WITH APPLICATION TO EMPIRICAL BAYES

Author(s): W. J. Conover, Texas Tech University
D. D. Cox, Rice University
H. F. Martz, Los Alamos National Laboratory

Submitted to: US Nuclear Regulatory Commission
Office of Nuclear Regulatory Research
Washington, DC 20555

DISTRIBUTION OF THIS DOCUMENT IS UNLIMITED

Los Alamos
NATIONAL LABORATORY

Form No. 836 R5
ST 2629 10/91

DISCLAIMER

LA-UR-97-4825

A Chi-Square Goodness-of-Fit Test for Non-Identically Distributed Random Variables: With Application to Empirical Bayes

by

W. J. Conover
Texas Tech University

D. D. Cox
Rice University

and

H. F. Martz
Los Alamos National Laboratory

Prepared for:

US Nuclear Regulatory Commission
Office of Nuclear Regulatory Research
Washington, DC 20555

Under:

NRC Contract Job Code W6505, Task 8
Los Alamos National Laboratory
Harry F. Martz, Principal Investigator

Dr. Lee R. Abramson
NRC Project Manager

December 1997

A CHI-SQUARE GOODNESS-OF-FIT TEST FOR NON-IDENTICALLY DISTRIBUTED RANDOM VARIABLES: WITH APPLICATION TO EMPIRICAL BAYES

by

W. J. Conover, D. D. Cox and H. F. Martz

Abstract

When using parametric empirical Bayes estimation methods for estimating the binomial or Poisson parameter, the validity of the assumed beta or gamma conjugate prior distribution is an important diagnostic consideration. Chi-square goodness-of-fit tests of the beta or gamma prior hypothesis are developed for use when the binomial sample sizes or Poisson exposure times vary. Nine examples illustrate the application of the methods, using real data from such diverse applications as the loss of feedwater flow rates in nuclear power plants, the probability of failure to run on demand and the failure rates of the high pressure coolant injection systems at U.S. commercial boiling water reactors, the probability of failure to run on demand of emergency diesel generators in U.S. commercial nuclear power plants, the rate of failure of aircraft air conditioners, baseball batting averages, the probability of testing positive for toxoplasmosis, and the probability of tumors in rats. The tests are easily applied in practice by means of corresponding Mathematica® computer programs which are provided.

1. Introduction

Goodness-of-fit tests usually require a random sample $X_1, X_2, \ldots, X_m$ from some distribution function F(x) to test the null hypothesis that $F(x) = F_0(x)$, where $F_0(x)$ is some hypothesized distribution function. The chi-square goodness-of-fit test, for example, may be used to test the null hypothesis that the random sample comes from a binomial distribution with known parameter n and either known or unknown parameter p. However the basic assumption is that $X_1, X_2, \ldots, X_m$ are independent and identically distributed.

There is a clear need for a goodness-of-fit test that can be used to test whether $X_1, X_2, \ldots, X_m$ comes from the family of binomial distributions, where the known parameter n_j can be different for each X_j. For example, the number of times engine j fails to start, X_j, in n_j demands may be a binomially distributed random variable. But simply knowing that engine j failed to start four times in twelve demands is not enough information to enable the comparison with a binomial distribution. However, there may be information available on the number of failures for m=40 similar engines, all believed to have the same unknown parameter p but with varying numbers of demands n_i. How can this information be combined to test the hypothesis of a binomial distribution?

This problem is not as simple as the case where all the n's are the same. One approach used in this paper is to combine the information on the basis of the number of failures. Then the

number of machines with zero failures are counted, the number with one failure are counted, etc., and compared with the expected value under the null hypothesis in a chi-square goodness-of-fit test. This approach has the weakness of counting one engine with one failure in twelve demands along with another engine with one failure in 100 demands. However, this approach does allow information from all the engines to be combined to see if the number of failures behaves in accordance with the binomial distribution probabilities.

A second approach used in this paper is to examine the estimated failure rates X_j/n_j for each engine, whose possible values have the same probabilities as X_j. These estimated failure rates are grouped into intervals with similar numerical values. Again this furnishes a method for aggregating the information from all of the engines. The drawback in this case is that an estimated probability of 0.25 may be close to a true probability of failure of 0.5 if the number of demands is twelve, but not as close if the number of demands is 100.

Although neither approach is perfect, there does not appear to be a perfect approach. In some cases in which the data are independent but non-identically distributed due to the presence of known nuisance parameter(s), a transformation exists to make the random variables identically distributed. For example, if the nuisance parameter is the varying mean of a set normally distributed random variables with equal variances, then subtraction of the mean transforms the random variables to identically distributed random variables, thus converting the sample to a random sample. However, in general this cannot always be done.

2. Parametric Empirical Bayes Prior Validation

Morris (1983), Casella (1985), Berger (1985), Maritz and Lwin (1989), Carlin and Louis (1996), and many others as well, consider the parametric empirical Bayes (PEB) compound (or two-stage) sampling model in which $X_j|\theta_j \overset{ind}{\sim} f_j(X_j|\theta_j)$ and $\theta_j \overset{iid}{\sim} G(\theta)$, $j = 1, 2, ..., m$. Here f_j denotes a specified sampling model, G is the corresponding prior distribution, and m represents the number of related situations connected by the structure of the problem (often referred to as *past or common experiments*) and for which data x_j, $j = 1, 2, ..., m$, are available. Gelman et al (1995) likewise discuss this model in the context of hierarchical Bayesian structures. Such compound sampling is quite common in a variety of practical applications (see Section 4). In PEB, a distributional family is assumed for G whose hyperparameters are subsequently estimated often using maximum likelihood or the method of moments in conjunction with the observed data x_j, $j = 1, 2, ..., m$. An important diagnostic aspect of PEB concerns the validity of G.

The validity of G is often investigated in either of two basic ways: by computing and plotting individual residuals based on the predictive distribution or by use of an omnibus goodness-of-fit test. For example, Gelfand, Dey, and Chang (1992) propose a cross-validation approach in which conditional residuals are plotted to reveal failures in the modeling assumptions regarding G. On the other hand, Gelman et al (1995) consider an omnibus goodness-of-fit test which requires calculating the corresponding Bayes p-value. Martz, Kvam, and Abramson (1996) also present an omnibus goodness-of-fit test based on the use of a randomized Kolmogorov-Smirnov test statistic. The omnibus test we consider here is an extension of common classical chi-square goodness-of-fit tests in the sense that we directly test whether or not the observed data come from the marginal distributional family corresponding to the given sampling model and assumed prior.

In particular, we consider two common PEB cases: (1) a binomial sampling model and a conjugate beta prior distribution; and (2) a Poisson sampling model and a conjugate gamma

prior distribution. In (1), X_j has a marginal beta-binomial distribution, while in (2), X_j marginally follows a gamma-Poisson distribution. However, the methods presented may be applied in more general situations in which the observed data are independent but non-identically distributed.

In the first case, suppose the number of failures X_j in n_j known demands for the *j*th system follows a binomial distribution with parameter p_j. The beta-binomial marginal distribution arises in situations where p_j is assumed to be a realization from a beta distribution with parameters α and β. The goodness-of-fit tests we propose enable testing the null hypothesis of a beta-binomial marginal distribution whose probability distribution for the *j*th system is given by

$$P(X_j = x) = \frac{n_j!}{x!(n_j - x)!} \frac{\Gamma(\alpha + x)}{\Gamma(\alpha)} \frac{\Gamma(\beta + n_j - x)}{\Gamma(\beta)} \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha + \beta + n_j)}$$

$$= \frac{n_j!}{x!(n_j - x)!} \prod_{j=0}^{x-1} (\alpha + j) \prod_{j=0}^{n_j - x - 1} (\beta + j) / \prod_{j=0}^{n_j - 1} (\alpha + \beta + j), \quad x = 1, 2, ..., n_j. \tag{1}$$

The test is based on the observed number of failures x_j from m such independent systems. If we assume that the binomial sampling model is correct, this test may be regarded as an implicit test of the assumption of a beta prior.

In the second case, suppose that the number of events (such as system failures) in known exposure (or operating) time t_j in the *j*th situation follows a Poisson distribution with parameter $\lambda_j t_j$. If λ_j has a gamma prior distribution with parameters α and β, then X_j (marginally) has a gamma-Poisson distribution whose probability distribution for the *j*th situation is given by

$$P(X_j = x) = \frac{\beta^{\alpha} t_j^{x} \Gamma(x + \alpha)}{\Gamma(\alpha)(t_j + \beta)^{(x + \alpha)} \Gamma(x + 1)}, \quad x = 1, 2, \ldots . \tag{2}$$

Using the observed number of events x_j from m such independent situations, we desire to explicitly test the hypothesis of a gamma-Poisson marginal distribution, thus implicitly testing the hypothesis of a gamma prior.

3. Chi-square Goodness-of-Fit Tests

Let X_j have a distribution function $F(x;\boldsymbol{\theta}_j)$ where $\boldsymbol{\theta}_j$ represents the parameter vector, for j = 1, 2, ..., m, and assume the X_j's are independent. In the case of the gamma-Poisson distribution, $\boldsymbol{\theta}_j = (\alpha, \beta, t_j)$, and in the case of the beta-binomial distribution, $\boldsymbol{\theta}_j = (\alpha, \beta, n_j)$. Some of the parameters may be known, such as t_j or n_j, while others may be unknown, such as α and β. Thus, in general, the X_j's are not identically distributed.

On the basis of the parameters, either known or estimated, non-overlapping adjacent intervals $I_1, I_2, \ldots, I_k$ are formed. Further, suppose h(x) is a function that maps the state space of the X's into the union of the intervals I_i and let $Z_{i,j}$ be an indicator variable for $h(X_j)$. That is,

$$Z_{i,j} = 1 \text{ if } h(X_j) \text{ is in interval } I_i$$
$$= 0 \text{ otherwise.} \quad (3)$$

Thus, the function h(x) maps each X_j into one and only one interval I_i.

Let $O_i = \Sigma_j Z_{i,j}$ be the observed number of $h(X_j)$'s mapped into interval I_i, and let E_i be the expected value of O_i given by

$$E_i = E(O_i) = \Sigma_j E(Z_{i,j}) = \Sigma_j P[h(X_j) \text{ is in interval } I_i]. \quad (4)$$

The goodness-of-fit test we propose is the usual chi-square goodness-of-fit test that uses

$$\chi^2 = \Sigma_i (O_i - E_i)^2 / E_i \quad (5)$$

as a test statistic and which asymptotically follows a chi-square distribution under the null hypothesis (see Appendix A). Thus, we use (5) in conjunction with an appropriate chi-square distribution to find the approximate p-value of the test. The expected values E_i are based on the probabilities under the null hypothesis. The unknown parameters are estimated using an efficient method, such as the method that minimizes the test statistic. If the number of estimated parameters is *s*, then under general conditions (see Appendix A), the asymptotic distribution of the test statistic (5) under the null hypothesis is chi-square with *k-s-1* degrees of freedom.

As a practical matter, in our examples involving the beta-binomial and gamma-Poisson distributions, we formed the intervals by estimating the unknown parameters α and β using the method of maximum likelihood, and then forming intervals so that the expected values were at least 0.5. Current research indicates that this is a reasonable lower bound to use on expected values in chi-square goodness-of-fit tests. Once the intervals were formed, the unknown parameters were re-estimated in order to minimize the value of the test statistic. In some cases the expected values dipped slightly below 0.5 in some cells as a result of the new parameter estimates; however, the chi-square approximation should still be valid.

One method of forming the intervals is to group the possible number of failures together until the group has an expected value of 0.5 or more. This allows a direct comparison between the number of times *i* failures is observed and the probability of getting *i* failures as determined by the probability distribution specified in the null hypothesis.

A second method of forming the intervals for the chi-square goodness-of-fit test is to group the estimated failure probabilities *(number of failures/number of demands)* in the case of a beta-binomial distribution, or the estimated failure rates *(number of failures/exposure time)* in the case of a gamma-Poisson distribution. This allows a more direct comparison between the estimated parameters and the prior distribution from which they may have come.

Both methods are used in this paper. Also, both the minimum chi-square method and the goodness-of-fit test based on the maximum likelihood estimates (MLEs) are used in this paper. Therefore, in most of the examples presented here, four tests are conducted, thus allowing the reader to compare the results of the various procedures.

4. Examples

Example 1(gamma-Poisson): Air conditioner failures.

This data set is discussed by Gaver and O'Muircheartaigh (1987) as an example of the gamma-Poisson distribution. The well-known data represent the number of failures of air conditioning equipment on 13 Boeing 720 aircraft [Proschan (1963)]. The numbers of failures and the times in service (in thousands of hours) are listed in Table 1. A test of equal failure rates concludes the failure rates are different, with a p-value of 0.027, in agreement with the conclusions of Gaver and O'Muircheartaigh (1987). Therefore, a gamma prior distribution on the failure rates may be appropriate.

Table 1. Air Conditioning Failure Data

Number of Failures	Time in Operation
2	0.623
9	1.800
14	1.832
15	1.819
12	1.297
6	0.639
23	2.201
29	2.422
6	0.493
16	1.312
27	2.074
24	1.539
30	1.788

Maximum likelihood estimates of the parameters α and β in the gamma prior distribution are $a_{MLE} = 18.40$ and $b_{MLE} = 1.73$. Using these MLEs, 18 cells are formed for the goodness-of-fit test, each having expected values of 0.5 or more using probabilities based on the MLEs. The cells group the observations based on the number of failures, and are explained in Table 2. For example, cell 1 includes 0-3 failures, which has one observation, and an expected value of 0.56 using the MLEs. The chi-square test statistic using the expected cell values based on the MLEs is 12.74, which has an associated p-value of 0.623 when compared with the approximating chi-square distribution with 15 degrees of freedom.

Table 2. Cell Summary for Air Conditioner Failure Data, Grouped by Number of Failures

Cell Number	Observed Number	Expected (MLEs)	Expected (Min χ^2)	Failures in Cell
1	1	0.56	0.42	0 - 3
2	0	0.85	0.76	4 - 5
3	2	0.96	0.90	6 - 7
4	1	0.91	0.83	8 - 9
5	0	0.93	0.80	10 - 11
6	1	1.02	0.88	12 - 13
7	1	0.55	0.49	14
8	1	0.57	0.52	15
9	1	0.58	0.55	16
10	0	0.58	0.57	17
11	0	0.56	0.57	18
12	0	0.54	0.57	19
13	0	0.52	0.56	20
14	0	0.93	1.04	21 - 22
15	2	0.78	0.90	23 - 24
16	0	0.62	0.74	25 - 26
17	2	0.66	0.81	27 - 29
18	1	0.90	1.10	30

The minimum chi-square statistic is evaluated by selecting estimates for α and β that minimize the test statistic. In this case the associated p-value (upper tail probability) is 0.660, indicating good agreement with the gamma-Poisson probabilities. The corresponding expected values are also shown in Table 2.

The test statistics (and thus the p-values) are in close agreement. The extra effort involved in finding the minimum test statistic may not be worthwhile in most cases, except perhaps when the p-value is marginal (e.g., near 0.05).

A second way of looking at these data is to directly consider the estimated failure rates per unit time. The time in service varies widely for these systems, from under 500 hours to over 2400 hours, and merely looking at numbers of failures per system ignores the time in service information. Therefore, the possible failure rates per unit time, 0/t, 1/t, 2/t, etc., where t is the time in service, are grouped into intervals with expected values, calculated using the gamma-Poisson probabilities with MLEs for α and β, of at least 0.5. This results in 19 intervals. Then the observed number of units with failure rates in each interval are compared with the expected values in a chi-square goodness-of-fit test.

The corresponding chi-square test statistic using MLEs to compute expected values is 9.15, which yields a p-value of 0.907 when compared with the chi-square distribution with 16 degrees of freedom. This indicates an excellent fit, confirming more directly the appropriateness of a gamma prior distribution for λ. Because of the large p-value, the minimum chi-square method is not necessary (it reduces the test statistic only slightly, from 9.153 to 9.152, with only a slight change in the estimates of α and β).

The details used in these goodness-of-fit tests are given in Table 3.

Table 3. Cell Summary for Air Conditioner Failure Data, Grouped by Failure Rate

Cell Number	Observed Number	Expected (MLEs)	Expected (Min χ^2)	Cell Boundaries
1	1	0.55	0.55	0 - 4.68
2	1	0.53	0.52	4.68 - 5.84
3	0	0.70	0.70	5.84 - 6.43
4	0	0.77	0.76	6.43 - 7.31
5	1	0.71	0.71	7.31 - 7.89
6	1	0.81	0.81	7.89 - 8.47
7	1	1.01	1.01	8.47 - 9.35
8	1	0.57	0.57	9.35 - 9.64
9	0	0.82	0.82	9.64 - 10.23
10	1	0.70	0.70	10.23 - 10.81
11	0	0.77	0.77	10.81 - 11.40
12	1	0.62	0.63	11.40 - 11.98
13	2	0.86	0.86	11.98 - 12.57
14	1	0.57	0.57	12.57 - 13.15
15	0	0.60	0.60	13.15 - 14.03
16	0	0.59	0.60	14.03 - 14.61
17	1	0.53	0.54	14.61 - 15.78
18	1	0.56	0.57	15.78 - 17.24
19	0	0.70	0.71	17.24

Example 2 (gamma-Poisson): Loss of feedwater flow.

This is another data set discussed by Gaver and O'Muircheartaigh (1987), and they attribute the data to Kaplan (1983).

The data represent rates of loss of feedwater flow for 23 commercial nuclear power generation systems. The number of failures and number of years the system was in operation are given in Table 4. These data represent a challenge to the goodness-of-fit test because there are some large numbers of failures, ranging up to 40, and a few systems with zero failures.

Table 4. Loss of Feedwater Flow Data

Number of Failures	Years of Operation	Number of Failures	Years of Operation
0	8	10	4
0	2	5	2
4	15	3	1
2	5	13	4
1	2	40	12
4	4	10	3
3	3	14	4
1	1	7	2
10	8	12	3
4	3	16	4
4	3	14	3
14	6		

A test of equal failure rates rejects the null hypothesis soundly, with a p-value less than 10^{-13}. Therefore, a gamma prior distribution may be appropriate. The MLEs of α and β are $a_{MLE} = 1.63$ and $b_{MLE} = 0.79$. On the basis of these MLEs, 21 intervals are formed for the goodness-of-fit test, each with an expected value of at least 0.5. The intervals are defined by numbers of failures, and the details are given in Table 5. For example, the first interval includes systems with 0 failures, which has two observed systems. The expected count using MLEs is 1.88, in good agreement.

Table 5. Cell Summary for Loss of Feedwater Flow Data, Grouped by Number of Failures

Cell Number	Observed Number	Expected (MLEs)	Expected (Min χ^2)	Failures in Cell
1	2	1.88	1.02	0
2	2	2.21	1.61	1
3	1	2.15	1.84	2
4	2	1.96	1.86	3
5	4	1.74	1.77	4
6	1	1.52	1.63	5
7	0	1.33	1.47	6
8	1	1.15	1.31	7
9	0	1.00	1.16	8
10	0	0.87	1.02	9
11	3	0.76	0.90	10
12	0	0.66	0.78	11
13	1	0.58	0.68	12
14	1	0.51	0.60	13
15	3	0.84	0.98	14 - 15
16	1	0.65	0.76	16 - 17
17	0	0.51	0.59	18 - 19
18	0	0.59	0.67	20 - 22
19	0	0.54	0.62	23 - 26
20	0	0.52	0.59	27 - 32
21	1	1.03	1.11	33

The minimum chi-square estimates of the parameters are $a_{MIN} = 2.58$ and $b_{MIN} = 1.12$, considerably different from the MLEs.

In this case, the comparison between the test statistic using MLEs and the minimized test statistic shows more difference than in Example 1. Using the MLEs the test statistic is 22.973 ($p = 0.192$), which is not very close to the minimized value of 21.04 ($p = 0.277$). However, both methods indicate good agreement with the gamma-Poisson distribution hypothesis.

When the fitting is applied to the failure rates per year instead of the numbers of failures, the results are nearly the same. Listed in Table 6 are the 25 grouped cells, the observed numbers, the expected numbers using the MLEs, the expected numbers in the minimum chi-square test, and the cell boundaries. The chi-square test statistic based on the MLEs is 28.89, with an associated p-value of 0.148 (with 22 degrees of freedom). The minimum

chi-square method produced a test statistic of 25.34, with a p-value of 0.281. This is based on parameter estimates of $a_{MIN} = 2.81$ and $b_{MIN} = 1.25$.

Table 6. Cell Summary for Loss of Feedwater Flow Data, Grouped by Failure Rate

Cell Number	Observed Number	Expected (MLEs)	Expected (Min χ^2)	Cell Boundaries
1	2	1.88	0.96	0
2	0	0.68	0.30	0 - 0.25
3	1	0.79	0.47	0.25 - 0.375
4	2	1.28	0.92	0.375 - 0.5
5	0	0.95	0.74	0.5 - 0.6875
6	0	0.53	0.44	0.6875 - 0.75
7	3	2.59	2.50	0.75 - 1.
8	1	0.77	0.81	1 -1.25
9	2	0.74	0.80	1.25 - 1.375
10	0	1.08	1.23	1.375 - 1.5
11	0	0.75	0.89	1.5 - 1.6875
12	0	0.61	0.75	1.6875 - 1.875
13	0	1.67	2.03	1.875 - 2.
14	0	0.53	0.68	2 - 2.25
15	1	0.56	0.72	2.25 - 2.4375
16	2	0.65	0.83	2.4375 - 2.5
17	0	0.74	0.96	2.5 - 2.75
18	1	1.19	1.52	2.75 - 3.
19	3	0.61	0.79	3 - 3.375
20	2	0.72	0.90	3.375 - 3.6875
21	2	0.86	1.05	3.6875 - 4.
22	0	0.60	0.70	4 - 4.5
23	1	0.66	0.73	4.5 - 5.
24	0	0.70	0.69	5 - 6.
25	0	0.84	0.58	6

Example 3 (gamma-Poisson): High pressure coolant injection (HPCI) system failures.

The numbers of failures of high pressure coolant injection (HPCI) systems at 23 nuclear power plants, along with the lengths of time the system was in operation during the calendar time period 1987-1993, are given in Table 7. The data are taken from a recent report by Grant et al (1995). A test of equal failure rates shows the plant failure rates to be unequal, with a p-value of 0.002.

Table 7. HPCI System Failure Data

Number of Failures	Years in Operation	Number of Failures	Years in Operation
2	2.25	5	5.7
10	3.83	8	3.85
11	4.59	2	6.28
2	5.64	7	3.97
2	5.09	11	3.54
3	5.42	4	3.85
9	5.63	8	5.53
9	5.55	8	5.44
8	4.49	6	5.67
9	5.89	5	6.05
8	5.97	2	6.22
6	6.15		

A gamma prior distribution is fit to the data using a gamma-Poisson marginal distribution. The MLEs of α and β are $a_{MLE} = 5.89$ and $b_{MLE} = 4.59$, resulting in 15 intervals for the goodness-of-fit test. The resulting intervals, along with the numbers of failures included in the intervals and the expected values, are given in Table 8.

Table 8. Cell Summary for HPCI System Failure Data, Grouped by Number of Failures

Cell Number	Observed Number	Expected (MLEs)	Expected (Min χ^2)	Failures in Cell
1	0	1.49	1.23	0 - 1
2	5	1.78	1.70	2
3	1	2.28	2.34	3
4	1	2.53	2.73	4
5	2	2.54	2.81	5
6	2	2.38	2.65	6
7	1	2.11	2.33	7
8	5	1.79	1.93	8
9	3	1.47	1.52	9
10	1	1.17	1.15	10
11	2	0.91	0.84	11
12	0	0.70	0.60	12
13	0	0.52	0.41	13
14	0	0.67	0.46	14 - 15
15	0	0.67	0.31	16

The minimum chi-square estimates of α and β are $a_{MIN} = 9.87$ and $b_{MIN} = 8.05$, again much larger than the MLEs. The minimized chi-square test statistic is 20.41, with associated p-value of 0.060. This is not much different than the test statistic 20.93 ($p = 0.051$) obtained using the MLEs to find the expected values. However, although the difference is small, the significance of the difference is large because they are so close to the usual cutoff value 0.05.

One reason for the poor (barely acceptable) fit to the gamma-Poisson distribution is the large number of plants (5) with exactly 2 failures and an equally large number of plants with exactly 8 failures. This appears to be a chance occurrence.

A more direct comparison with the gamma prior distribution is obtained by computing the sample failure rates per year, and comparing them with the expected values. This results in the 19 intervals listed in Table 9 along with the corresponding cell calculations. Based on the MLEs, the corresponding test statistic is 24.25, which has a p-value of 0.084.

Table 9. Cell Summary for HPCI System Failure Data, Grouped by Failure Rates

Cell Number	Observed Number	Expected (MLEs)	Expected (MIN χ^2)	Cell Boundaries
1	0	1.22	2.02	0 - 0.27
2	4	0.88	1.23	0.27 - 0.41
3	0	1.72	2.03	0.41 - 0.55
4	1	1.15	1.27	0.55 - 0.69
5	0	1.75	1.80	0.69 - 0.82
6	3	1.88	1.82	0.82 - 0.96
7	3	1.96	1.80	0.96 - 1.10
8	0	1.39	1.24	1.10 - 1.24
9	1	2.06	1.80	1.24 - 1.37
10	2	1.33	1.13	1.37 - 1.51
11	3	1.43	1.21	1.51 - 1.65
12	2	1.24	1.05	1.65 - 1.79
13	0	0.75	0.63	1.79 - 1.92
14	0	0.91	0.78	1.92 - 2.06
15	1	0.67	0.58	2.06 - 2.20
16	0	0.56	0.49	2.20 - 2.34
17	1	0.86	0.78	2.34 - 2.61
18	1	0.54	0.52	2.61 - 2.88
19	1	0.71	0.83	2.88

Finally, the minimum chi-square test statistic is computed for these same 19 intervals as 22.19, with a p-value of 0.137. This is the result of using $a_{MIN} = 3.79$ and $b_{MIN} = 3.15$, which are smaller than the MLEs of 5.89 and 4.59, respectively. The cell details are also given in Table 9.

Thus, because of these marginal p-values, a gamma prior distribution may or may not be completely appropriate for describing the underlying plant-to-plant variablity in the HPCI failure rate.

<u>Example 4 (beta-binomial)</u>: Emergency diesel generator (EDG) failures to run on demand

This example is an analysis of the failure-to-run on demand data on the emergency diesel generators (EDGs) in 63 US commercial nuclear power plants. The data are the same as those analyzed in Martz et al (1996), and are given in Table 10.

Table 10. Emergency Diesel Generator (EDG) Failure-to-Run on Demand Data

Number of Failures	Number of Demands	Number of Failures	Number of Demands	Number of Failures	Number of Demands
11	854	2	373	5	618
5	157	3	542	2	202
0	65	2	166	4	574
2	201	8	388	2	287
3	431	5	358	14	1120
1	321	3	225	8	433
7	468	1	218	9	317
0	238	1	370	2	302
3	152	2	294	0	101
0	283	2	117	0	115
1	196	0	310	4	134
2	132	0	242	4	132
4	320	0	996	1	253
13	704	2	151	2	385
5	216	0	252	1	419
0	136	3	185	7	382
6	304	2	130	1	121
2	295	1	289	0	181
2	150	0	334	2	263
0	92	2	466	2	387
1	183	5	278	0	212

There were a total of 182 failures to run in response to 19,520 demands on the system. A chi-square contingency table analysis to see if the EDG failure probability is the same for all 63 plants concludes that the probabilities are different, with a p-value of 0.0001. Therefore, a beta prior distribution may be appropriate.

The chi-square goodness-of-fit test was used to see if the data can be considered to follow the beta-binomial distribution. Cell details are given in Table 11. MLEs for the beta parameters are $a_{MLE} = 2.39$ and $b_{MLE} = 251.42$, matching those obtained by Martz, et al (1996). The chi-square goodness-of-fit test statistic, using the MLEs for α and β, is 14.56, with a corresponding p-value of 0.203.
The values $a_{MIN} = 2.03$ and $b_{MIN} = 189.1$ for α and β in the beta prior distribution produce a minimum chi-square test statistic of 13.69 which has a corresponding p-value of 0.251, slightly larger than in the MLE case. The assumption of a beta prior is thus reasonable.

Table 11. Cell Summary for EDG Failure Data, Grouped by Number of Failures

Cell Number	Number of Plants	Expected (MLEs)	Expected (Min χ^2)	Failures in Cell
1	14	13.03	12.35	0
2	9	13.42	12.50	1
3	17	10.51	9.97	2
4	5	7.57	7.41	3
5	4	5.29	5.37	4
6	5	3.68	3.88	5
7	1	2.56	2.81	6
8	2	1.80	2.05	7
9	2	1.28	1.51	8
10	1	0.92	1.13	9
11	0	0.67	0.85	10
12	1	0.88	1.14	11-12
13	2	0.51	0.68	13-14
14	0	0.88	1.35	15

Another way of looking at the distribution of failures is to group the data into intervals by similar probabilities of failur , instead of by similar numbers of failures. The intervals are formed using MLEs for α and β as before, so that the expected value in each interval is at least 0.5. This results in 32 intervals as listed in Table 12.

The chi-square test statistic as computed using the MLEs for α and β is 36.89, which has a p-value of 0.149. The minimum chi-square test statistic is somewhat smaller at 32.75, which has a p-value of 0.288, and is based on parameter estimates $a_{MIN} = 3.01$ and $b_{MIN} = 266.2$.

Table 12. Cell Summary for EDG Failure Data, Grouped by Failure Probabilities

Cell Number	Number of Plants	Expected (MLEs)	Expected (Min χ^2)	Cell Boundaries
1	14	13.03	9.84	0
2	1	1.97	1.40	0 - 0.0027
3	3	3.38	2.73	0.0027 - 0.0036
4	2	2.24	1.91	0.0036 - 0.0045
5	4	3.09	2.75	0.0045 - 0.0054
6	3	2.40	2.21	0.0054 - 0.0062
7	6	4.07	3.89	0.0062 - 0.0071
8	1	3.16	3.12	0.0071 - 0.0080
9	2	2.64	2.67	0.0080 - 0.0089
10	0	2.04	2.13	0.0089 - 0.0098
11	2	2.81	2.99	0.0098 - 0.0107
12	0	2.03	2.20	0.0107 - 0.0116
13	1	1.23	1.38	0.0116 - 0.0125
14	6	2.38	2.71	0.0125 - 0.0134
15	1	1.54	1.80	0.0134 - 0.0143
16	2	1.50	1.74	0.0143 - 0.0152
17	1	1.52	1.79	0.0152 - 0.0161
18	1	1.52	1.85	0.0161 - 0.0169
19	1	0.95	1.16	0.0169 - 0.0178
20	4	0.83	1.04	0.0178 - 0.0187
21	0	0.76	0.97	0.0187 - 0.0196
22	2	1.09	1.36	0.0196 - 0.0205
23	1	0.60	0.81	0.0205 - 0.0214
24	0	0.80	1.03	0.0214 - 0.0223
25	1	0.70	0.91	0.0223 - 0.0232
26	0	0.88	1.19	0.0232 - 0.0250
27	0	0.84	1.14	0.0250 - 0.0268
28	1	0.55	0.78	0.0268 - 0.0285
29	2	0.52	0.72	0.0285 - 0.0303
30	1	0.62	0.86	0.0303 - 0.0330
31	0	0.55	0.80	0.0330 - 0.0375
32	0	0.74	1.12	0.0375 - 1

Example 5 (beta-binomial): Baseball batting averages.

Efron and Morris (1975) used the 1970 batting averages for the first 18 major league baseball players to achieve 45 times at bat as an example of the beta-binomial distribution. They used these data to predict the batting averages for the remainder of the season for the same 18 players. Because the number of attempts is equal for all players, at 45, the early-season batting averages can be analyzed using traditional methods and are not of interest here.

The data set of interest in this paper is the number of hits these same 18 players obtained during the remainder of the season, obtained from the batting averages and number of times at bat given by Efron and Morris (1975). Presumably, this information does not include the first 45 times at bat, so it is, in a sense, independent of the previous set. The hits and times at bat are given in Table 13.

Table 13. Batting Records for 18 Baseball Players

Number of Hits	Times at Bat	Number of Hits	Times at Bat
127	367	46	200
127	426	73	277
144	521	69	270
61	275	132	435
114	418	142	538
126	466	42	186
154	586	159	558
29	138	129	408
137	510	14	70

These data do not pass the test of homogeneous hit probabilities. The 2x18 contingency table has a p-value of 0.022. Therefore, a beta prior distribution is considered as the distribution of probabilities of getting a hit for the population of players.

The fit to the beta-binomial distribution is good. The p-value is 0.280, using the MLEs a_{MLE} = 166.91 and b_{MLE} = 445.3. The chi-square test statistic is 32.950 with 29 degrees of freedom. Most of the cell counts are 0 or 1 as shown in Table 14, with two 2's and one 3. The expected values in the grouped cells ranged from 0.5 to 0.84.

Table 14. Cell Summary for Baseball Batting Records, Grouped by Number of Hits

Cell Number	Players in Cell	Expected (MLEs)	Expected (Min χ^2)	Number of Hits
1	1	0.552	0.525	0 - 19
2	1	0.521	0.528	20 - 29
3	0	0.542	0.509	30 - 38
4	1	0.534	0.516	39 - 44
5	1	0.518	0.497	45 - 49
6	0	0.577	0.581	50 - 54
7	0	0.531	0.535	55 - 59
8	1	0.591	0.550	60 - 65
9	1	0.588	0.553	66 - 70
10	1	0.543	0.549	71 - 74
11	0	0.536	0.569	75 - 78
12	0	0.557	0.596	79 - 83
13	0	0.556	0.540	84 - 90
14	0	0.548	0.464	91 - 97
15	0	0.537	0.469	98 - 102
16	0	0.544	0.500	103 - 106
17	0	0.633	0.608	107 - 110
18	0	0.517	0.511	111 - 113
19	1	0.536	0.540	114 - 116
20	0	0.540	0.550	117 - 119
21	0	0.533	0.544	120 - 122
22	0	0.519	0.526	123 - 125
23	3	0.502	0.503	126 - 128
24	2	0.645	0.639	129 - 132
25	0	0.623	0.616	133 - 136
26	1	0.605	0.607	137 - 140
27	2	0.586	0.602	141 - 144
28	0	0.556	0.588	145 - 148
29	0	0.512	0.556	149 - 152
30	1	0.557	0.619	153 - 157
31	1	0.524	0.592	158 - 163
32	0	0.835	0.917	164

The values of a and b that minimized the chi-square test statistic for these cells were a_{MIN} = 269.53 and b_{MIN} = 705.90. The resulting test statistic was 32.794 with a corresponding p-value of 0.286. The new cell expected numbers ranged from 0.464 to 0.917, with all but one being very close to 0.5. Although the estimates of α and β that produced the minimum chi-square test statistic are considerably larger than the MLEs, there is very little difference in the final values for the chi-square test statistic and the resulting p-value.

These small expected values in the cells may cause the reader to wonder how well the chi-square distribution approximates the true distribution of the test statistic. The latest research on this topic agrees that expected values can be much smaller than traditionally thought. For example, Koehler and Larntz (1980) state that when cells are approximately equi-probable, as they are in this case, the chi-square approximation may be considered sufficient when expected values are as small as 0.25 if the number of cells k is at least 3, the number of observations n is at least 10, and the ratio n^2/k is at least 10. Here k = 32 cells, n = 18

players, and all expected values are at least twice 0.25, so the conditions are met, and the chi-square approximation is likely to be sufficiently accurate.

Another way of examining the data, by batting averages instead of by numbers of hits, may seem more appropriate to baseball fans. Twenty-seven intervals are formed on the basis of possible batting averages for the given numbers of times at bat, so that the expected values are at least 0.5 using the MLEs for α and β in the beta-binomial distribution. The details of the cell intervals used in this test are given in Table 15.

Table 15. Cell Summary for Baseball Batting Records, Grouped by Batting Average

Cell Number	Number of Players	Expected (MLEs)	Batting Averages
1	2	0.51218	0 - 0.213
2	1	0.578331	0.213 - 0.225
3	2	0.601901	0.225 - 0.232
4	0	0.51657	0.232 - 0.237
5	0	0.653115	0.237 - 0.242
6	0	0.551145	0.242 - 0.245
7	0	0.526142	0.245 - 0.249
8	0	0.581705	0.249 - 0.252
9	0	0.711637	0.252 - 0.256
10	1	0.763215	0.256 - 0.259
11	0	0.758133	0.259 - 0.262
12	3	0.745399	0.262 - 0.266
13	1	0.801199	0.266 - 0.269
14	1	0.864479	0.269 - 0.273
15	1	0.775698	0.273 - 0.276
16	1	0.665845	0.276 - 0.279
17	0	0.85218	0.279 - 0.283
18	1	0.765961	0.283 - 0.286
19	0	0.583772	0.286 - 0.290
20	0	0.693266	0.290 - 0.293
21	0	0.576135	0.293 - 0.296
22	1	0.865655	0.296 - 0.302
23	1	0.616011	0.302 - 0.307
24	0	0.511862	0.307 - 0.312
25	1	0.585743	0.312 - 0.319
26	0	0.620619	0.319 - 0.329
27	1	0.722101	0.329 - 1

The chi-square test statistic using the MLEs to estimate α and β is 23.94, with a p-value of 0.465, indicating a good fit. Because the p-value is already sufficiently large to indicate a good agreement with a beta prior distribution, the minimum chi-square method was not applied in this case.

Example 6 (beta-binomial): Testing positive for toxoplasmosis.

Efron (1986) presents data on the numbers of subjects testing positive for toxoplasmosis in 34 cities in El Salvador. The data set appears in Table 16.

Table 16. Incidence of Toxoplasmosis in 34 Cities in El Salvador

No. of Cases of Toxoplasmosis	No. of Patients Examined	No. of Cases of Toxoplasmosis	No. of Patients Examined
2	4	3	54
3	10	4	9
4	5	5	18
3	10	2	12
2	2	0	1
3	5	8	11
2	8	41	77
7	19	24	51
3	6	7	16
8	10	46	82
7	24	9	13
0	1	23	43
15	30	53	75
4	22	8	13
0	1	3	10
6	11	1	6
0	1	23	37

A test of equal probabilities of testing positive fails, with a p-value on the order of 10^{-9}. Thus, the probabilities are assumed to come from a beta prior distribution. The maximum likelihood estimators for the parameters in the beta-binomial distribution are $a_{MLE} = 3.59$ and $b_{MLE} = 4.46$.

This data set illustrates the difficulty that can arise when the number of attempts (subjects) in some data pairs is smaller than the number of failures (subjects testing positive) in other data pairs. In particular, note that three cities had only one subject and one city had only two, which are hardly large enough to convey much information. Also, the chi-square distribution as an approximation depends on there being a sufficiently large number of individual contributions in each cell, and if some cells are large, say 18 or more patients testing positive, this automatically excludes most of the cities from having a positive probability in those cells because of the small numbers of patients examined in those cities.

In a case like this it makes more sense to group cities on the basis of incidence rate, and Then all cities can be included in the analysis. The test resulted in 32 cells as detailed in Table 16 yielding a chi-square value of 39.15, and a p-value of 0.099.

The minimum chi-square method results in a slight difference. The test statistic is now reduced to 38.26, which has a p-value of 0.117, based on the estimated parameters $a_{MIN} = 4.51$ and $b_{MIN} = 5.95$. These estimates can be compared with the MLEs of 3.59 and 4.46, respectively.

Table 17. Cell Summary for Toxoplasmosis Data, Grouped by Incidence Rate

Cell Number	Number of Cities	Expected (MLEs)	Expected (Min χ^2)	Incidence Rates in Cell
1	4	3.34	3.36	0
2	1	0.81	0.72	0 - 0.108
3	0	0.66	0.62	0.108 - 0.157
4	2	0.67	0.68	0.157 - 0.169
5	1	1.46	1.51	0.169 - 0.205
6	0	0.57	0.59	0.205 - 0.229
7	1	1.14	1.21	0.229 - 0.253
8	0	0.73	0.79	0.253 - 0.277
9	5	1.10	1.20	0.277 - 0.301
10	0	0.80	0.89	0.301 - 0.325
11	0	1.15	1.25	0.325 - 0.337
12	1	1.06	1.18	0.337 - 0.373
13	0	0.75	0.82	0.373 - 0.386
14	0	1.81	1.98	0.386 - 0.410
15	0	0.77	0.85	0.410 - 0.434
16	2	0.56	0.61	0.434 - 0.446
17	0	0.51	0.55	0.446 - 0.458
18	3	0.62	0.68	0.458 - 0.470
19	4	3.09	3.24	0.470 - 0.506
20	2	1.00	1.05	0.506 - 0.542
21	2	0.93	0.95	0.542 - 0.566
22	0	0.62	0.61	0.566 - 0.590
23	1	1.23	1.22	0.590 - 0.602
24	2	0.78	0.75	0.602 - 0.627
25	0	0.52	0.49	0.627 - 0.639
26	0	1.00	0.91	0.639 - 0.675
27	2	1.03	0.88	0.675 - 0.711
28	1	0.97	0.80	0.711 - 0.759
29	2	0.99	0.77	0.759 - 0.807
30	0	0.51	0.37	0.807 - 0.843
31	0	0.50	0.31	0.843 - 0.952
32	1	2.33	2.16	0.952 - 1

Note that there are many more cells using this method of counting sample proportions than there would be by counting sample frequencies. No cities had to be discarded from the analysis, and the observed cell frequencies are more spread out. For data of this type, it is preferable to look at relative rather than raw frequencies.

Example 7 (beta-binomial): Tumor incidences in rats.

Gelman et al (1995) use data on the incidence of tumors in groups of rats from Tarone (1982) to illustrate the beta-binomial distribution; however, they did not consider a goodness-of-fit test to see if the distribution is beta-binomial. Seventy groups of laboratory rats of type "F344" have been studied under control (no dosage) conditions. The number of rats in each group and the number of rats with tumors are given as follows:

Table 18. Number of Rats with Tumors in 70 Groups of Rats of Various Numbers

Number with Tumors	Number in Group	Number with Tumors	Number in Group
0	20	0	19
1	18	2	20
3	20	4	20
6	23	0	20
0	18	1	18
1	10	2	13
10	48	5	19
0	20	0	18
2	25	5	49
9	48	4	19
6	22	0	20
0	17	2	24
2	19	10	50
4	19	6	20
0	20	1	20
2	23	5	46
4	20	4	19
6	20	0	20
1	20	2	20
3	27	4	20
5	22	6	20
0	20	1	20
2	20	2	17
4	20	11	46
16	52	0	19
1	20	2	20
7	49	4	20
12	49	15	47
0	19	1	19
2	20	7	47
4	20	5	20
15	46	0	19
1	19	2	20
3	20	4	20
5	20	9	24

A test of equal tumor rates shows the tumor rates to be different from group to group, with a p-value of less than 10^{-6}. The different groups of rats were studied at different times and under different laboratory conditions, so the differences in tumor incidence rates were expected.

MLEs of the parameters in the beta prior distribution, using the beta-binomial likelihood function, result in $a_{MLE} = 2.30$ and $b_{MLE} = 14.08$. Using these MLEs in the goodness-of-fit test gives a chi-square test statistic of 16.93. The p-value is 0.110.

The minimum chi-square method results in a minimized chi-square statistic of 10[illegible] and a p-value of 0.520, based on the revised parameter values $a_{MIN} = 1.14$ and $b_{MIN} =$ [illegible]. In this case, the revised parameter values are less than half of the MLEs, and the chi-square

statistic becomes much smaller. This example shows clearly the need for using the minimum chi-square method in some cases.

Table 19. Cell Summary for Rat Tumor Data, Grouped by Numbers of Rats With Tumors

Cell Number	Groups in Cell	Expected (MLEs)	Expected (Min χ^2)	No. of Rats
1	14	8.73	12.19	0
2	9	12.13	11.17	1
3	12	11.94	9.53	2
4	3	10.09	7.88	3
5	10	7.81	6.39	4
6	6	5.71	5.10	5
7	5	4.03	4.02	6
8	2	2.78	3.13	7
9	0	1.90	2.41	8
10	2	1.31	1.84	9
11	2	0.92	1.40	10
12	1	0.66	1.06	11
13	1	0.86	1.40	12 - 13
14	3	1.12	2.50	14

When the tumor incidences are grouped by incidence rates instead of incidence numbers, the results are as follows. The chi-square test statistic is 21.60, with a corresponding p-value of 0.305. The minimum chi-square method gives a test statistic just slightly smaller, 21.19, with a p-value of 0.326, based on parameter estimates $a_{MIN} = 1.93$ and $b_{MIN} = 11.41$. The cell information is as follows:

Table 20. Cell Summary for Rat Tumor Data, Grouped by Incidence Rate

Cell Number	Groups in Cell	Expected (MLEs)	Expected (Min χ^2)	Incidence Rates
1	14	8.73	9.39	0
2	0	0.91	0.98	0 - 0.038
3	8	11.20	10.97	0.038 - 0.057
4	0	1.71	1.65	0.057 - 0.075
5	3	2.62	2.47	0.075 - 0.094
6	11	10.32	9.74	0.094 - 0.113
7	1	2.20	2.05	0.113 - 0.132
8	4	6.00	5.65	0.132 - 0.151
9	1	3.65	3.45	0.151 - 0.170
10	1	1.66	1.58	0.170 - 0.189
11	8	4.49	4.33	0.189 - 0.208
12	4	2.71	2.64	0.208 - 0.226
13	3	1.07	1.06	0.226 - 0.245
14	4	3.99	4.01	0.245 - 0.264
15	1	0.78	0.81	0.264 - 0.283
16	3	2.22	2.32	0.283 - 0.302
17	2	1.03	1.10	0.302 - 0.321
18	1	1.67	1.87	0.321 - 0.358
19	1	0.55	0.63	0.358 - 0.377
20	0	1.03	1.23	0.377 - 0.415
21	0	0.67	0.86	0.415 - 0.453
22	0	0.80	1.19	0.453 - 1

Either way of looking at the data results in the conclusion that a beta prior distribution is appropriate.

Example 8 (beta-binomial): Grant's high pressure coolant injection (HPCI) system data.

There are two sets of binomial failure data given by Grant et al (1995) based on reports submitted by nuclear power boiling water reactor plants from 1987-1993. One is on page C-12 of their report and represents the times the high-pressure coolant injection (HPCI) system in US commercial nuclear power plants failed to run for reasons other than a failure of the injection valve (FTSO) for 23 plants. The data are presented in Table 21. The other data set, presented in Example 9, represents the number of times the HPCI system failed to run (FTR) on demand for the same 23 plants, and is given on page C-14 of their report. Both data sets are characterized by varying numbers of demands from plant-to-plant. However, one is more suited to the beta-binomial model than the other.

Table 21. HPCI FTSO Data from Grant et al (1995)

Failures to Start	Number of Attempts	Failures to Start	Number of Attempts
0	3	0	5
0	5	0	13
0	13	0	8
0	11	0	5
3	8	1	8
0	7	0	6
0	6	1	5
0	6	1	7
1	6	0	5
4	14	0	5
0	13	0	6
0	14		

The 2x23 contingency table test rejects the null hypothesis of equal probabilities of failure for the 23 plants, with p less than 0.001, so the beta prior distribution is postulated for the parameter p. MLEs for the beta parameters are $a_{MLE} = 0.368$ and $b_{MLE} = 5.94$.

Cells with small expected values (less than 0.5) are combined, leaving only four cells for the goodness-of-fit test for the beta binomial, with a resulting chi-square test statistic of 2.245. Comparison with the chi-square distribution with 1 degree of freedom results in a p-value of 0.134, thus showing an adequate fit.

The minimum chi-square method resulted in $a_{MIN} = 0.271$ and $b_{MIN} = 3.31$. The minimum chi-square test statistic was 1.715, with a resulting p-value of 0.190. The cell summaries are given in Table 22.

Table 22. Cell Summary for HPCI FTSO Data, Grouped by Number of Failures

Cell Number	Plants in Cell	Expected (MLEs)	Expected (Min χ^2)	Number of Failures
1	17	16.95	16.47	0
2	4	3.62	3.31	1
3	0	1.38	1.51	2
4	2	1.05	1.71	3

If, instead of grouping by numbers of failures, the failure rates are grouped, the number of observations in the "zero" category remains unchanged, but the other cells are different than before. Now there are five cells, as described in Table 23.

Table 23. Cell Summary for HPCI FTSO Data, Grouped by Failure Rates

Cell Number	Number of Plants	Expected (MLEs)	Expected (Min χ^2)	Failure Rates
1	17	16.95	16.86	0
2	1	1.52	1.51	0 - 0.13
3	3	2.47	2.48	0.13 - 0.2
4	1	1.11	1.134	0.2 - 0.33
5	1	0.95	1.02	0.33 - 1

The chi-square test statistic, using the MLEs to estimate α and β, is 0.305, with a p-value of 0.858. The minimum chi-square method uses $a_{MIN} = 0.359$ and $b_{MIN} = 5.59$ (close to the MLEs) to get a test statistic of 0.300, which has a p-value of 0.861, not much different than before.

Example 9 (binomial): FTR HPCI systems.

The 23 plants reported a total of 7 HPCI failures to run (FTR) in response to 167 demands on the HPCI systems. Plant-by-plant data are as follows:

Table 24. HPCI FTR Data from Grant et al (1995)

Number of Failures	Number of Demands	Number of Failures	Number of Demands
0	3	1	5
1	5	0	4
2	11	0	8
0	11	1	5
0	7	0	9
0	6	0	6
0	8	0	7
0	6	0	4
0	6	0	5
0	14	1	5
1	12	0	6
0	14		

A chi-square contingency table analysis to see if the probability of failure is the same for all 23 plants concludes that the probabilities are not necessarily different, with a p-value of 0.36. This is due to the few failure-to-run occurrences consistently throughout all 23 plants.

The relative frequency of failure, 7/167=0.0419, was used as the estimate of p in the goodness-of-fit test of a binomial distribution in which the number of demands varies from plant-to-plant. The chi-square statistic calculated on three cells, all that was left after grouping so the expected values were at least 0.5, was 0.0162. The corresponding p-value was high at 0.899.

The minimum chi-square test statistic is 0.0139 when fitting the binomial distribution, and is based on an estimated p of 0.0427 The p-value is now 0.906. The grouped cell summary is given in Table 25.

Table 25. Cell Summary for HPCI FTR Data, Grouped by Number of Failures

Cell Number	Number of Plants	Expected (MLEs)	Expected (Min χ^2)	Number of Failures
1	17	16.995	16.902	0
2	5	5.115	5.179	1
3	1	0.890	0.919	2

Summary of Examples:

All of the examples satisfy the rules of thumb for small expected values suggested by Koehler and Larntz (1980) except for Example 1. There the values of n^2/k are slightly under the suggested value of 10. In most cases the minimum chi-square method did not change the p-value much, as shown in Table 26, with the notable exception of the tumor incidence in rats discussed in Example 7. The chi-square approximation is an asymptotic approximation, and it is difficult to determine how the approximation fares with small samples. In particular, it is difficult to assess whether the use of the MLEs provides a more accurate p-value or whether the asymptotically correct minimum chi-square method provides a more accurate p-value in the small sample case. Case-by-case simulation studies may be appropriate for answering this question.

Table 26. Summary of p-values in the Examples

Example Number	m	Null Hypothesis	Grouping By: Number of Failures		Failure Rate	
			MLEs	Min χ^2	MLEs	Min χ^2
1	13	gamma-Poisson	0.623	0.660	0.907	0.907
2	23	gamma-Poisson	0.192	0.277	0.148	0.281
3	23	gamma-Poisson	0.051	0.060	0.084	0.137
4	63	beta-binomial	0.203	0.251	0.149	0.288
5	18	beta-binomial	0.280	0.286	0.465	NA
6	34	beta-binomial	NA	NA	0.099	0.117
7	70	beta-binomial	0.110	0.520	0.305	0.326
8	23	beta-binomial	0.134	0.190	0.858	0.861
9	23	binomial	0.899	0.906	NA	NA

ACKNOWLEDGMENT

This work was supported by the U.S. Nuclear Regulatory Commission, Office of Nuclear Regulatory Research, under Job Code Number W6505 (Dr. Lee R. Abramson, NRC Project Manager) at the Los Alamos National Laboratory. We are grateful for this support.

REFERENCES

Agresti, A. (1990). *Categorical Data Analysis*. John Wiley & Sons, New York.

Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*, 2nd ed. Springer-Verlag, New York.

Carlin, B. P. and Louis, T. A. (1996). *Bayes and Empirical Bayes Methods for Data Analysis*. Chapman and Hall, London.

Casella, G. (1985). An introduction to empirical Bayes data analysis. The American Statistician, 39, 83-87.

Efron, B. (1986). Double exponential families and their use in generalized linear regression. Journal of the American Statistical Association, 81, 709-721.

Efron, B. and Morris, C. (1975). Data analysis using Stein's estimator and its generalizations. Journal of the American Statistical Association, 70, 311-319.

Feller, W. (1971). *An Introduction to Probability Theory and Its Applications*. John Wiley & Sons, New York.

Gaver, D. P. and O'Muircheartaigh, I. G. (1987). Robust empirical Bayes analyses of event rates. Technometrics, 29, 1-15.

Gelfand, A. E., Dey, D. K., and Chang, H. (1992). Model determination using predictive distributions with implementation via sampling-based methods (with discussion). In *Bayesian Statistics* 4, J. M. Bernardo, J. O. Berger, A. P. Dawid and A. F. M. Smith, eds., Oxford University Press, Oxford, 147-167.

Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (1995). *Bayesian Data Analysis*. Chapman and Hall, London.

Grant, G. M., Roesener, W. S., Hall, D. G., Atwood, C. L., Gentillon, C. D., and Wolf, T. R. (1995). High-Pressure Coolant Injection (HPCI) System Performance, 1987-1993, Final. Idaho National Engineering Laboratory Report Number INEL-94/0158.

Hajek, J and Sidak, Z. (1967). *Theory of Rank Tests*. Academic Press, New York.

Kaplan, S. (1983). On a two-stage Bayesian procedure for determining failure rates from experimental data. IEEE Transactions on Power Apparatus and Systems, 102, 195-202.

Koehler, K. J. and Larntz, K. (1980). An empirical investigation of goodness-of-fit statistics for sparse multinomials. Journal of the American Statistical Association, 75, 336-344.

Maritz, J. S. and Lwin, T. (1989). *Empirical Bayes Methods*, 2nd ed. Chapman and Hall, London.

Martz, H. F., Kvam, P. H., and Abramson, L. R. (1996). Empirical Bayes Estimation of the Reliability of Nuclear-Power-Plant Emergency Diesel Generators. Technometrics, 38, 11-24.

Morris, C. N. (1983). Parametric empirical Bayes inference: Theory and applications. Journal of the American Statistical Association, 78, 47-65.

Proschan, F. (1963). Theoretical explanation of observed decreasing failure rate. Technometrics, 5, 375-383.

Tarone, R. E. (1982). The use of historical control information in testing for a trend in proportions. Biometrics, 38, 215-220.

Wolfram Research, Inc.(1996). *Mathematica*®, Version 3.0, Champaign, IL.

APPENDIX A
Justification of the Chi-square Distribution as an Asymptotic Limit

As in Section 3, let X_j have a distribution function $F(x;\theta_j)$ where θ_j represents the parameter vector, for $j = 1, 2, \ldots, m$, and assume the X_j's are independent. On the basis of the parameters, known or estimated, k non-overlapping adjacent intervals $I_1, I_2, \ldots, I_k$ are formed. Let h(x) be a function that maps the state space of the X_j's into the union of the intervals I_i and let $Z_{i,j}$ be an indicator variable for $h(X_j)$. That is,

$$Z_{i,j} = 1 \text{ if } h(X_j) \text{ is in interval } I_i$$
$$= 0 \text{ otherwise}$$

Thus, the function h(x) maps each X_j into one and only one interval I_i.

Let $O_i = \Sigma_j Z_{i,j}$ be the observed number of $h(X_j)$'s mapped into interval I_i, and let E_i be the expected value of O_i which is

$$E_i = E(O_i) = \Sigma_j E(Z_{i,j}) = \Sigma_j P[h(X_j) \text{ is in interval } I_i] = \Sigma_j \pi_{i,j}$$

The variance of O_i is given by $\Sigma_j \pi_{i,j}(1-\pi_{i,j})$ and the covariance of $O_{i1,i2}$ is $-\Sigma_j \pi_{i1,j}\pi_{i2,j}$.

A necessary and sufficient condition for the asymptotic (as $m \to \infty$) multivariate normality of the vector $\{O_1, \ldots, O_k\}$ is that every linear combination of the O_i's, say $\Sigma_i \lambda_i O_i$, is asymptotically univariate normal, when properly normed [for example, see Hajek and Sidak (1967), p. 168]. However, this can be shown as follows.

The Lindeberg condition for asymptotic normality of the sum of bounded random variables is satisfied if the sum of the variances of the random variables goes to infinity [Feller (1971), p. 264]. If $\Sigma_j \pi_{i,j} \to \infty$, then $\mathrm{Var}(\Sigma_i Z_{i,j}) \to \infty$, and because the Z's are bounded by 1, the Lindeberg condition is satisfied and the O_i's are asymptotically normal. Similarly, for any given set of λ's, the linear combination $\Sigma_i \lambda_i Z_i$ is bounded by $\Sigma_i |\lambda_i|$, which implies that $\Sigma_i \lambda_i O_i$ is asymptotically normal.

Thus, the only condition for asymptotic multivariate normality of the vector $\{O_1, \ldots, O_k\}$ is that $E(O_i) = \Sigma_j \pi_{i,j} \to \infty$. This is easily satisfied for fixed k, as m goes to infinity, as long as the number of contributing variables also goes to infinity. That is, if the $\pi_{i,j}$'s are going to zero too quickly as m gets large, for any cell, then it is possible that the cell probability may be bounded. However, in the binomial or beta-binomial case, if an infinite number of n_j's, are less than M for any finite M, and only a finite number of $\pi_{i,j}$'s equal zero for each i, then an infinite number of $\pi_{i,j}$'s are bounded away from zero for each i and the expected cell size will go to infinity. In the case of finite sample sizes, however, all of these conditions are purely speculative.

Our usual practice of forming the cells so their expected values are small but equal whenever possible, as recommended by the latest studies, may seem to contradict the assumption of the cell expectations going to infinity. However, all samples are necessarily finite, and the number of X_j's never ever really goes to infinity. Thus, for practical considerations, we are forced to accept cells that, based on our experience, will lead to approximate chi-square distributions, and still have good power to detect alternative hypotheses. That happens with many cells, which necessarily have small expected values.

If the vector $\{O_1, \ldots, O_k\}$ is asymptotically multivariate normal, then the statistic

$$\chi^2 = \Sigma_i (O_i - E_i)^2 / E_i$$

has the same limiting distribution as if the O's were normal, which is chi-square, with degrees of freedom equal to $k - 1$, as the following demonstrates.

Write χ^2 as $\Sigma_i e_i^2$, where $e_i = (O_i - E_i)/E_i^{1/2}$. The rank of the covariance matrix of the vector $\{e_1, \ldots, e_k\}$ is $k-1$ because the condition $\Sigma_i O_i = m$ implies

$$\mathrm{Cov}(e_i, e_k) = -(E_k)^{-1/2} \Sigma_{\ell=1}^{k-1} (E_\ell)^{1/2} \mathrm{Cov}(e_j, e_\ell),$$

which shows that the kth column is a linear combination of the other $k - 1$ columns.

When s parameters are estimated in the minimum chi-square manner, or in any "efficient" manner, the degrees of freedom are decreased to $k - s - 1$ [for example, see Agresti (1990), p.471].

APPENDIX B
Mathematica® Programs for Use in Implementing The Goodness-of-fit Tests

These programs may be copied directly into files and used in Mathematica® [Wolfram (1996)] to perform the various analyses described in this paper.

The first group of programs B.1 through B.9 is aimed at testing the null hypothesis of a beta-binomial distribution, with parameters α and β unknown, and possibly different sample sizes n_i, all known. One program tests the null hypothesis of a binomial distribution. The details of these programs are found in the "help" program B.1 at the beginning of the group, with the exception of "beta-bin-mle-gof3.mat" which was used on some data sets but not described in this paper. It studies the grouping into intervals of the estimator (x+1)/(n+2) instead of the estimator x/n, in an attempt to subdivide the large number of observations at zero. However the results were not interesting enough to publish, so the program is not described in the "help" program. It is included in case the reader prefers to break up the large cell grouping that sometimes occurs in the "zero" cell.

The second group of programs B.10 through B.17 is aimed at testing the null hypothesis of a gamma-Poisson distribution, with parameters α and β unknown, and possibly different exposure times t_i. Again, the details of these programs are found in the "help" program B.10 at the beginning of the group.

B.1. beta-bin-help.mat

HELP FILE FOR THE MATHEMATICA® PROGRAMS TO MAKE GOODNESS-OF-FIT TESTS TO THE BETA-BINOMIAL DISTRIBUTION

INTRODUCTION

This file helps to explain the Mathematica® programs written to analyze data that express the numbers of demands and numbers of failures at several plants. One program tests the hypothesis that all plants have the same probability of failure. A second program examines how well the data fit a binomial distribution.

A third program examines how well the data fit a beta-binomial distribution, where the parameters are estimated using MLEs on the data before grouping the data into cells. There are two versions of this program. One uses a chi-square goodness-of-fit test on the counts of the number of failures. The other uses the same test but on the sample proportions grouped into intervals.

The final program refines the previous test, by adjusting the fitted parameters to obtain the minimum value of the test statistic. It, too, appears in two versions, one which examines the number of failures, and the other which examines the sample proportions. This final program sometimes takes a long time to run, and sometimes is unable to minimize the test statistic within the limits of the parameter space, due to the nature of some unusual data.

The data are assumed to be in a file where the first entry is the number of data pairs, and the subsequent entries are the pairs; number of demands, number of failures. All entries in the file are separated by spaces or line returns.

Ideally, each number of demands should be large enough so there is positive probability in each of the resulting cells in the chi-square goodness-of-fit test. Although some exceptions can be made, data with many exceptions should be analyzed by forming groups based on failure rates (programs with "gof2" in the title rather than "gof") rather than by forming groups based on numbers of failures.

First read the data file into Mathematica® with the commands:

```
ClearAll[temp]
temp=OpenRead["name of data file"]
```

SETTING UP THE DATA: beta-bin-setupfile.mat

After the data have been placed into a file called "temp," call up the file that sets up the data in a form amenable to analysis by the other Mathematica® programs, "beta-bin-setupfile.mat", with the command:

```
<<"beta-bin-setupfile.mat"
```

This file gives the following output.

m = the number of pairs of data
n[i] = number of demands, for i = 1, ... , m
f[i] = number of failures, for i = 1, ... , m
n1 = array of n's in array form
f1 = array of f's in array form

A table of the pairs (f[i], n[i]) can be obtained with the command

```
Table[{f[i],n[i]},{i,m}]
```

which can be copied directly from this file to your command file.

After using "beta-bin-setupfile.mat" any of the following six programs can be used, in any order, depending on what type of analysis is desired. The outputs from the various programs are described below, and are available on request from Mathematica®. Useful tables can also be obtained by copying the commands given here.

A TEST OF EQUAL PROBABILITIES: bin-gof.mat

This program takes the output of "beta-bin-setupfile.mat", and makes a 2xm contingency table, where the entries in the first row are the numbers of failures per unit, f[i], and the entries in the second row are the numbers of successes per unit, x[i]. The null hypothesis is that the probability of failure is the same for each unit (column). A Pearson chi-square statistic is computed and compared with the chi-square distribution with m-1 degrees of freedom, in the usual manner for testing this hypothesis.

The output of this program includes the following.

x1 = array of (n-f)'s, number of successes, in array form
totalf = total number of failures
totalx = total number of successes

totaln = total number of demands
chisq2xm = the chi-square test statistic
chisqpvalue = the p-value of the test

GOODNESS-OF-FIT TEST FOR THE BINOMIAL DISTRIBUTION: ord-bin-gof.mat

This program takes the output of "beta-bin-setupfile.mat" and performs a goodness-of-fit test of the null hypothesis that the data follow the ordinary binomial distribution, with parameter p estimated from the overall relative frequency of failures, and the various n's for each unit. Then it finds the value of p that minimizes the value of the test statistic, and finds its p-value.

The output of this program includes the following.

binp = the estimated value of the parameter p
bink = the largest number of failures, such that the probability of getting bink failures or more is as close to 0.5 as it can get without going under 0.5
obin[i] = the observed number of units having i failures, i=0,...,bink-1
obin[bink] = the observed number of units with bink or more failures
e2bin[i] = the expected number of units with i failures, i=0,...,bink-1
e2bin[bink] = the expected number of units with bink failures or more
kbin = the number of cells in the goodness-of-fit test, after combining cells with expected values less than 0.5
o1bin[i] = the observed number in cell i after grouping, i=1,...,kbin
e1bin[i] = the expected number in cell i after grouping, i=1,...,kbin
binchisq = the test statistic in the chi-square goodness-of-fit test
bindf = the degrees of freedom in the chi-square goodness-of-fit test
binpvalue = the p-value in the chi-square goodness-of-fit test
pmin = the estimate of p that minimizes the chi-square test statistic
minbinchisq = the minimum chi-square test statistic
maxbinpvalue = the corresponding maximum p-value
e1minbin[i] = the corresponding expected cell values, i=1,...,kbin

Some Mathematica® commands that bring up useful tables are as follows. To see the cell-by-cell observed and expected values, with right-tail grouping in the final cell, use:

```
Table[{i,obin[i],e2bin[i]},{i,0,bink}]
```

To see the grouped cell observed and expected values, that are used in the goodness-of-fit test, use:

```
Table[{i,o1bin[i],e1bin[i],e1minbin[i],beg[i],end[i]},{i,kbin}]
```

GOODNESS OF FIT TEST FOR THE BETA-BINOMIAL DISTRIBUTION, USING MAXIMUM LIKELIHOOD ESTIMATORS FOR THE BETA DISTRIBUTION: beta-bin-mle-gof.mat

This Mathematica® program uses the easily-obtained method of moments estimators a0 and b0 as starting points for the search for the maximum likelihood estimators a and b. It uses

the output from beta-bin-setupfile.mat, and therefore the output from that file is still available, in addition to the following terms.

a = the maximum likelihood estimator of the beta parameter alpha
b = the maximum likelihood estimator of the beta parameter beta
m1 = the average value of (f+1)/(n+2), the centroid likelihood estimator of the probability of failure for each unit
m2 = the average of the squares of the quantities averaged in m1
a0 = the estimate of alpha using the method of moments, based on m1 and m2
b0 = the estimate of beta using the method of moments, based on m1 and m2
kmax = the maximum number of failures such that the probability of kmax or more failures is greater than 0.5, using estimators a and b.
o[i] = the observed number of units with i failures, i=0,...,kmax-1
o[kmax] = the observed number of units with kmax or more failures
e[i] = the expected number of units with i failures, assuming the beta-binomial distribution with parameters a and b, i<kmax
e[kmax] = the expected number of units with kmax or more failures
k = the number of cells in the goodness-of-fit test, after grouping
beg[i] = the smallest number of failures counted in the grouped cell i
end[i] = the largest number of failures counted in the grouped cell i
o1[i] = the observed number in the grouped cell, i=1,...,k
e1[i] = the expected number in the grouped cell, i=1,...,k
df = the degrees of freedom used in the goodness-of-fit test
chisq = the test statistic obtained in the goodness-of-fit test
pvalue = the p-value obtained in the goodness-of-fit test

Some tables that present useful information can be obtained with the following commands. To obtain a table of the observed numbers of units with i failures, and the corresponding expected values using a and b in the beta-binomial distribution, with the final cells including the right-tail accumulations, use:

```
Table[{i,o[i],e[i]},{i,0,kmax}]
```

To see the grouped cells, with the cell number, the grouped observed counts, the grouped expected counts, the smallest number of failures included in the grouped cell, and the largest number of failures included in the grouped cell, use:

```
Table[{i,o1[i],e1[i],beg[i],end[i]},{i,k}]
```

A SECOND VERSION OF "beta-bin-mle-gof.mat" THAT IS BASED ON SAMPLE PROPORTIONS INSTEAD OF NUMBERS OF FAILURES: beta-bin-mle-gof2.mat

This program is similar to the previous one through the obtaining of the MLEs. Therefore the definitions of a, b, m1, m2, a0, and b0 are the same as before. Then the difference begins, for the sample proportions, f/n, are arranged into groups with expected values of at least 0.5. Usually, with highly reliable data, there is a large group with zero for a sample proportion, but the other groups will be more spread out. New notation is as follows.

k2 = the number of cells in the goodness-of-fit test, after grouping the sample proportions so the expected numbers are at least 0.5 in each cell
beg3[i] = the lower bound in the grouped cell i

end3[i] = the upper bound in the grouped cell i
o3[i] = the observed number in the grouped cell, i=1,...,k2
e3[i] = the expected number in the grouped cell, i=1,...,k2
df2 = the degrees of freedom used in the goodness-of-fit test
chisq2 = the test statistic obtained in the goodness-of-fit test
pvalue2 = the p-value obtained in the goodness-of-fit test

To see the grouped cells, with the cell number, the grouped observed sample proportions, the grouped expected cell values, and the cell boundaries, use:

Table[{i,o3[i],e3[i],beg3[i],end3[i]},{i,k2}]

GOODNESS OF FIT TEST FOR THE BETA-BINOMIAL DISTRIBUTION USING THE VALUES OF a AND b THAT MINIMIZE THE CHI-SQUARE STATISTIC: beta-bin-min-gof.mat

This program repeats everything that "beta-bin-mle-gof.mat" does, and uses a and b as starting points for the estimators of alpha and beta in the beta distribution that minimize the chi-square test statistic. The cells are the same groupings that are used in "beta-bin-min-gof.mat", so the final expected values may be slightly less than 0.5 in some cells. This "minimum chi-square method" of estimating parameters is more in accordance with the asymptotic theory that justifies subtracting one degree of freedom for each parameter estimated. The actual difference between the chi-square test statistics from the previous program and this program is usually quite small and may not be worth the extra computing time required to run this program except in special cases.

This program depends on the output from "beta-bin-setupfile.mat" and therefore the output from this program includes the output from that program, the output from "beta-bin-mle-gof.mat" and the following additional terms.

amin = the estimate of the parameter alpha that minimizes the test statistic in the goodness-of-fit test
bmin = the estimate of the parameter beta that minimizes the test statistic in the goodness-of-fit test
e1min[i] = the expected value in the grouped cell i using amin and bmin
minchisq = the minimum value of the test statistic
maxpvalue = the corresponding p-value

A useful table showing details in the grouped cells used in the goodness-of fit test can be obtained with the following command:

Table[{i,o1[i],e1[i],e1min[i],beg[i],end[i]},{i,k}]

This table shows the cell number, the total number of observations in the grouped cell, the expected cell value using a and b, the expected cell value using amin and bmin, and the smallest and largest numbers of failures included in the grouped cell.

A SECOND VERSION OF "beta-bin-min-gof.mat" THAT IS BASED ON SAMPLE PROPORTIONS INSTEAD OF NUMBERS OF FAILURES: beta-bin-min-gof2.mat

This program includes everything in "beta-bin-mle-gof2.mat" plus it finds the minimum chi-square test statistic. It takes a long time to run, and the difference in test statistics is

usually small, so it is not needed unless the p-value using the MLEs needs to be adjusted more precisely. This program is similar to the previous one through the obtaining of the MLEs, therefore the definitions of a, b, m1, m2, a0, and b0 are the same as before. Then the difference begins, for the sample proportions, f/n, are arranged into groups with expected values of at least 0.5. Usually, with highly reliable data, there is a large group with zero for a sample proportion, but the other groups will be more spread out. New notation is as follows.

This program depends on the output from "beta-bin-setupfile.mat" and therefore the output from this program includes the output from that program, the output from "beta-bin-mle-gof2.mat" and the following additional terms.

amin2 = the estimate of the parameter alpha that minimizes the test statistic in the goodness-of-fit test
bmin2 = the estimate of the parameter beta that minimizes the test statistic in the goodness-of-fit test
e1min3[i] = the expected value in the grouped cell i using amin2 and bmin2
minchisq2 = the minimum value of the test statistic
maxpvalue2 = the corresponding p-value

A useful table showing details in the grouped cells used in the goodness-of fit test can be obtained with the following command:

```
Table[{i,o3[i],e3[i],e1min3[i],beg3[i],end3[i]},{i,k2}]
```

This table shows the cell number, the total number of sample proportions in the grouped cell, the expected cell value using a and b, the expected cell value using amin2 and bmin2, and the cell boundaries in the grouped cell.

B.2: beta-bin-setupfile.mat

```
ClearAll[f,n];

m=Read[temp,Number] (*Now m = the number of data pairs.*);

Do[{n[i],f[i]}=Read[temp,{Number,Number}],{i,m}];

f1=Array[f,m];

n1=Array[n,m];

ClearAll[temp];
```

B.3: bin-gof.mat

```
(
(*This program takes data from a file called "temp" that has been processed
using a file called "beta-bin-setupfile.mat". See "beta-bin-help.mat" for more
information and detailed instructions.*);

(*This program performs an ordinary 2xm contingency table analysis to see if the
probability of failure is the same for m data sources.*);

x1=n1-f1;
```

```
totaln=Sum[n1[[i]],{i,m}];
totalf=Sum[f1[[i]],{i,m}];
totalx=Sum[x1[[i]],{i,m}];

chisq2xm=(totaln^2*Sum[x1[[i]]^2/n1[[i]],{i,m}]/(totalx*totalf)
      -totaln*totalx/totalf)//N;

chisqpvalue=1-CDF[ChiSquareDistribution[m-1],chisq2xm])
```

B.4: ord-bin-gof.mat

```
(
(*This program takes data from a file called "temp" that has been processed
using a file called "beta-bin-setupfile.mat". See "beta-bin-help.mat" for more
information and detailed instructions.*);

(*The first part of this program uses the usual estimator of p to fit a binomial
distribution to the observations. The minimum expected cell size is 0.5. That
number can be changed by changing 0.5 where it appears in two places in this
program, to any new desired minimum expected value.

The second part of this program finds the minimum chi-square test statistic
using a Mathematica® FindMinimum search procedure. The cells are the same as the
ones used in the first part of this program.*);

ClearAll[ebin,e0bin,e1bin,obin,o0bin,o1bin] (*This gets rid of hidden
definitions.*);

binp=Sum[f[i],{i,m}]/Sum[n[i],{i,m}]//N;
ebin[i_] := Sum[Binomial[n[j],i]*binp^i*(1-binp)^(n[j]-i),{j,m}];

bink=0;
sumbin=ebin[0];

While[sumbin<m-0.5, (bink=bink+1)&&(sumbin=sumbin+ebin[bink])];
Do[e0bin[i]=ebin[i],{i,0,bink-1}];
  e0bin[bink]=m-Sum[e0bin[i],{i,0,bink-1}];
Do[e2bin[i]=e0bin[i],{i,0,bink}];

Do[obin[i]=Count[f1,i],{i,0,bink-1}];
  obin[bink]=m-Sum[obin[i],{i,0,bink-1}];
Do[o0bin[i]=obin[i],{i,0,bink}];

i1=1; i3=1;j0=0; j1=0;
beg[1]=0;
end[1]=0;

Do[If[e0bin[i2]<0.5,
(j1=j1+1)&&(end[i3]=j1)&&(e0bin[j1]=e0bin[i2]+e0bin[j1]),
((e1bin[i1]=e0bin[i2])
&&(o1bin[i1]= Sum[o0bin[i3],{i3,beg[i3],end[i3]}])
&&(j0=j1+1)&&(i1=i1+1)&&(i3=i1)&&(beg[i3]=j0)&&(j1=j1+1)
&&(end[i3]=j1))],{i2,0,bink}];
```

```
kbin=i1-1;

binchisq=N[Sum[(o1bin[i])^2/e1bin[i],{i,kbin}]-m];

bindf=kbin-2;

binpvalue=1-CDF[ChiSquareDistribution[bindf],binchisq];

ClearAll[minbinp,minbinchisq,maxbinpvalue];

Do[eminbin[i1,i3_,p_] :=Sum[Sum[
        Binomial[n[j],i]*p^i*(1-p)^(n[j]-i),{j,m}],
        {i,beg[i3],end[i3]}],
        {i1,kbin-1}];

{minbinchisq,{pbin1}}=FindMinimum[N[Sum[o1bin[i]^2/eminbin[i,i,pdum],{i,kbin-1}]
+(o1bin[kbin])^2/(m-Sum[eminbin[i,i,pdum],{i,kbin-1}])-m],
{pdum,binp,0,1}];

pmin=pdum/.pbin1;

Do[e1minbin[i]=eminbin[i,i,pmin],{i,kbin-1}];
e1minbin[kbin]=m-Sum[e1minbin[i],{i,kbin-1}];

maxbinpvalue=1-CDF[ChiSquareDistribution[bindf],minbinchisq];
)
```

B.5: beta-bin-mle-gof.mat

```
(
(*This program compares the observed number of times t failures occur to the
expected values from the beta-binomial probability distribution.*);

(*This program takes data from a file called "temp" that has been processed
using a file called "beta-bin-setupfile.mat". See "beta-bin-help.mat" for more
information and detailed instructions.*);

(*This program uses maximum likelihood estimators of alpha and beta in the
beta-binimial distribution, found using the Mathematica® FindMinimum command, to
fit a beta-binomial distribution to the observations. The minimum expected cell
size is 0.5. That number can be changed by changing 0.5 where it appears in two
places in this program, to any new desired minimum expected value.*);

ClearAll[phat,log,minval,apoint,bpoint,m1,m2,a0,a1];

phat=N[(f1+1)/(n1+2)];

m1=Sum[phat[[i]],{i,m}]/m//N;

m2=Sum[(phat[[i]])^2,{i,m}]/m//N;

a0=m1*(m1-m2)/(m2-m1^2);
```

```
b0=(1-m1)*(m1-m2)/(m2-m1^2);

log[a1_,b1_] :=Sum[(Sum[Log[a1+j],{j,0,f1[[i]]-1}]
+Sum[Log[b1+j],{j,0,n1[[i]]-f1[[i]]-1}]
-Sum[Log[a1+b1+j],{j,0,n1[[i]]-1}]),{i,m}];

{minval,{apoint,bpoint}}=FindMinimum[N[-log[adum,bdum]],
    {adum,a0,0.01,Max[m,20*a0]},{bdum,b0,0.01,Max[m,20*b0]}];

a=adum/.apoint;

b=bdum/.bpoint;

ClearAll[betabinomial];
betabinomial[f_,n_] := Binomial[n,f]*Gamma[a+f]*Gamma[b+n-f]*Gamma[a+b]/
   (Gamma[a]*Gamma[b]*Gamma[a+b+n]);

sumbb=0;
i1=0;
ClearAll[e2,e3,o2,o3];
e2[0]=Sum[betabinomial[0,n[i]],{i,m}]//N;
sumbb=e2[0]//N;

mx=Max[n1];
ClearAll[frac,am,bm];
Do[Do[frac[i,j]=j/n[i],{j,n[i]}],{i,m}];

Do[am[i,j]=0,{i,m},{j,0,mx}];

Do[Do[Do[If[(frac[i,j1]<=j/(mx+1))&&(frac[i,j1]>(j-1)/(mx+1)),
am[i,j]=am[i,j]+j1],
  {j1,n[i]}],{j,mx}],{i,m}];

j1=1;
While[sumbb<m-0.5,
    (e2[j1]=Sum[If[am[i,j1]>0,N[betabinomial[am[i,j1],n[i]]],0],{i,m}])
    &&(sumbb=sumbb+e2[j1])
    &&(j1=j1+1)];
k2max=j1-1;
e2[k2max]=m-Sum[e2[i],{i,0,k2max-1}];
ClearAll[j1,sumbb,beg2,beg3,end2,end3];

o2[0]=Count[f1,0];
Do[bm[i,j]=0,{i,m},{j,k2max-1}];
Do[If[(f[i]/n[i]<=j/(mx+1))&&(f[i]/n[i]>(j-1)/(mx+1)),bm[i,j]=bm[i,j]+1],
 {i,m},{j,k2max-1}];

Do[o2[j]=Sum[bm[i,j],{i,m}],{j,k2max-1}];
o2[k2max]=m-Sum[o2[i],{i,0,k2max-1}];

i1=1;i3=1;j1=0;
beg2[1]=0;
end2[1]=0;
```

```
Do[If[e2[i2]<0.5,
(j1=j1+1)&&(end2[i3]=j1)&&(e2[j1]=e2[i2]+e2[j1]),
((e3[i1]=e2[i2])
&&(o3[i1]=Sum[o2[i3],{i3,beg2[i3],end2[i3]}])
&&(beg3[i3]=N[(beg2[i3]-1)/(mx+1)])&&(end3[i3]=N[end2[i3]/(mx+1)])
&&(i1=i1+1)&&(i3=i1)&&(j1=j1+1)&&(beg2[i3]=j1)&&(end2[i3]=j1))],
{i2,0,k2max}];
beg3[1]=0;

k2=i1-1;
end3[k2]=1;

ClearAll[i1,j0,j1,i3];

ClearAll[chisq2,pvalue2];

chisq2=(Sum[(o3[i])^2/e3[i],{i,k2}]-m)//N;

df2=k2-3 (* This subtracts two degrees of freedom for the two estimated
parameters, a and b*);

pvalue2=1-CDF[ChiSquareDistribution[df2],chisq2] (*If the p-value is > .05
then the data could have come from the distribution producing the e[]'s,
but if the p-value is less than or equal to .05 the fit is poor.*);
)
```

B.6: beta-bin-mle-gof2.mat

```
(
(*This program uses t/n as the estimator of the failure rate, and compares the
observed values of this estimator from several data sources to the expected
values from the beta-binomial probability distribution.*);

(*This program takes data from a file called "temp" that has been processed
using a file called "beta-bin-setupfile.mat". See "beta-bin-help.mat" for more
information and detailed instructions.*);

(*This program uses maximum likelihood estimators of alpha and beta in the
beta-binimial distribution, found using the Mathematica® FindMinimum command, to
fit a beta-binomial distribution to the observations. The minimum expected cell
size is 0.5. That number can be changed by changing 0.5 where it appears in two
places in this program, to any new desired minimum expected value.*);

ClearAll[phat,log,minval,apoint,bpoint,m1,m2,a0,a1];

phat=N[(f1+1)/(n1+2)];

m1=Sum[phat[[i]],{i,m}]/m//N;

m2=Sum[(phat[[i]])^2,{i,m}]/m//N;

a0=m1*(m1-m2)/(m2-m1^2);

b0=(1-m1)*(m1-m2)/(m2-m1^2);
```

```
log[a1_,b1_] :=Sum[(Sum[Log[a1+j],{j,0,f1[[i]]-1}]
+Sum[Log[b1+j],{j,0,n1[[i]]-f1[[i]]-1}]
-Sum[Log[a1+b1+j],{j,0,n1[[i]]-1}]),{i,m}];

{minval,{apoint,bpoint}}=FindMinimum[N[-log[adum,bdum]],
    {adum,a0,0.01,Max[m,20*a0]},{bdum,b0,0.01,Max[m,20*b0]}];

a=adum/.apoint;

b=bdum/.bpoint;
ClearAll[betabinomial];
betabinomial[f_,n_] := Binomial[n,f]*Gamma[a+f]*Gamma[b+n-f]*Gamma[a+b]/
   (Gamma[a]*Gamma[b]*Gamma[a+b+n]);

i1=0;
ClearAll[e2,e3,o2,o3];
e2[0]=Sum[betabinomial[0,n[i]],{i,m}]//N;
sumbb=e2[0]//N;
mx=Max[n1];

j2=1;
While[sumbb<m-0.5,(e2[j2]=Sum[
 If[(N[j1/n[i]]<=N[j2/(mx+1)])&&(N[j1/n[i]]>N[(j2-1)/(mx+1)]),
  N[betabinomial[j1,n[i]]],0],{i,m},{j1,j2}])
    &&(sumbb=sumbb+e2[j2])&&(j2=j2+1)];

k2max=j2-1;
e2[k2max]=m-Sum[e2[i],{i,0,k2max-1}];
ClearAll[j2,sumbb,beg2,beg3,end2,end3];

o2[0]=Count[f1,0];

Do[o2[j]=Sum[If[(N[f[i]/n[i]]<=N[j/(mx+1)])&&
        (N[f[i]/n[i]]>N[(j-1)/(mx+1)]),1,0],
        {i,m}],{j,k2max-1}];
o2[k2max]=m-Sum[o2[i],{i,0,k2max-1}];

i1=1;i3=1;j1=0;
beg2[1]=0;
end2[1]=0;

Do[If[e2[i2]<0.5,
(j1=j1+1)&&(end2[i3]=j1)&&(e2[j1]=e2[i2]+e2[j1]),
((e3[i1]=e2[i2])
&&(o3[i1]=Sum[o2[i3],{i3,beg2[i3],end2[i3]}])
&&(beg3[i3]=N[(beg2[i3]-1)/(mx+1)])&&(end3[i3]=N[end2[i3]/(mx+1)])
&&(i1=i1+1)&&(i3=i1)&&(j1=j1+1)&&(beg2[i3]=j1)&&(end2[i3]=j1))],
{i2,0,k2max}];
beg3[1]=0;

k2=i1-1;
end3[k2]=1;
```

```
ClearAll[i1,j0,j1,i3];

ClearAll[chisq2,pvalue2];

chisq2=(Sum[(o3[i])^2/e3[i],{i,k2}]-m)//N;

df2=k2-3 (* This subtracts two degrees of freedom for the two estimated
parameters, a and b*);

pvalue2=1-CDF[ChiSquareDistribution[df2],chisq2] (*If the p-value is > .05
then the data could have come from the distribution producing the e[]'s,
but if the p-value is less than or equal to .05 the fit is poor.*))
```

B.7: beta-bin-min-gof.mat

```
(

(*This program compares the observed number of times t failures occur, to the
expected values from the beta-binomial probability distribution.*);

(*This program takes data from a file called "temp" that has been processed
using a file called "beta-bin-setupfile.mat". See "beta-bin-help.mat" for more
information and detailed instructions.*);

(*The first part of this program uses maximum likelihood estimators of alpha and
beta in the beta-binimial distribution, found using the Mathematica® FindMinimum
command, to fit a beta-binomial distribution to the observations. The minimum
expected cell size is 0.5. That number can be changed by changing 0.5 where it
appears in two places in this program, to any new desired minimum expected
value.

The second part of this program finds the minimum chi-square test statistic
using a Mathematica® FindMinimum search procedure. The cells are the same as the
ones used in the first part of this program.*);

ClearAll[phat,log,minval,apoint,bpoint,m1,m2,a0,a1];

phat=N[(f1+1)/(n1+2)];

m1=Sum[phat[[i]],{i,m}]/m//N;

m2=Sum[(phat[[i]])^2,{i,m}]/m//N;

a0=m1*(m1-m2)/(m2-m1^2);

b0=(1-m1)*(m1-m2)/(m2-m1^2);

log[a1_,b1_] :=Sum[(Sum[Log[a1+j],{j,0,f1[[i]]-1}]
+Sum[Log[b1+j],{j,0,n1[[i]]-f1[[i]]-1}]
-Sum[Log[a1+b1+j],{j,0,n1[[i]]-1}]),{i,m}];

{minval,{apoint,bpoint}}=FindMinimum[N[-log[adum,bdum]],
    {adum,a0,0.01,Max[m,20*a0]},{bdum,b0,0.01,Max[m,20*b0]}];
```

```
a=adum/.apoint;

b=bdum/.bpoint;

ClearAll[betabinomial];
betabinomial[f_,n_] := Binomial[n,f]*Gamma[a+f]*Gamma[b+n-f]*Gamma[a+b]/
   (Gamma[a]*Gamma[b]*Gamma[a+b+n]);

sumbb=0;
i1=0;
ClearAll[e,e0,e1,o,o0,o1];

While[sumbb<m-0.5,
    (e[i1]=Sum[N[betabinomial[i1,n1[[j]]]],{j,1,m}])
     &&(sumbb=sumbb+e[i1])
     &&(i1=i1+1)];
kmax=i1-1;
e[kmax]=m-Sum[e[i],{i,0,kmax-1}];
ClearAll[i1,sumbb,beg,end];

Do[o[i]=Count[f1,i],{i,0,kmax-1}];
  o[kmax]=m-Sum[o[i],{i,0,kmax-1}];
Do[o0[i]=o[i],{i,0,kmax}];
Do[e0[i]=e[i],{i,0,kmax}];

i1=1;i3=1;j0=0;j1=0;
beg[1]=0;
end[1]=0;

Do[If[e0[i2]<0.5,
(j1=j1+1)&&(end[i3]=j1)&&(e0[j1]=e0[i2]+e0[j1]),
((e1[i1]=e0[i2])
&&(o1[i1]=Sum[o0[i3],{i3,beg[i3],end[i3]}])
&&(j0=j1+1)&&(i1=i1+1)&&(i3=i1)&&(beg[i3]=j0)&&(j1=j1+1)&&(end[i3]=j1))],
{i2,0,kmax}];

k=i1-1;
ClearAll[i1,j0,j1,i3];

ClearAll[chisq,pvalue];

chisq=(Sum[(o1[i])^2/e1[i],{i,k}]-m)//N;

df=k-3;
pvalue=1-CDF[ChiSquareDistribution[df],chisq];

ClearAll[i1,j0,j1,i3];

ClearAll[amin1,bmin1,emin];

Do[emin[i1,i3_,a2_,b2_] :=Sum[Sum[(Binomial[n1[[j]],i]*Gamma[a2+i]*
Gamma[b2+n1[[j]]-
i]*Gamma[a2+b2]/(Gamma[a2]*Gamma[b2]*Gamma[a2+b2+n1[[j]]])),
{j,1,m}],{i,beg[i3],end[i3]}],{i1,k-1}];
```

```
{minchisq,{amin1,bmin1}}=FindMinimum[N[Sum[(o1[i])^2/emin[i,i,adum,bdum],{i,k-
1}
]+
 (o1[k])^2/(m-Sum[emin[i,i,adum,bdum],{i,k-1}])-m],
 {adum,a,0.01,Max[m,20*a]},{bdum,b,0.01,Max[m,20*b]}];

amin=adum/.amin1;

bmin=bdum/.bmin1;

Do[e1min[i]=emin[i,i,amin,bmin],{i,k-1}];
e1min[k]=m-Sum[e1min[i],{i,k-1}];

maxpvalue=1-CDF[ChiSquareDistribution[df],minchisq])
```

B.8: beta-bin-min-gof2.mat

```
(

(*This program uses t/n as the estimator of the failure rate, and compares the
observed values of this estimator from several data sources to the expected
values from the beta-binomial probability distribution.*);

(*This program takes data from a file called "temp" that has been processed
using a file called "beta-bin-setupfile.mat". See "beta-bin-help.mat" for more
information and detailed instructions.*);

(*The first part of this program uses maximum likelihood estimators of alpha and
beta in the beta-binimial distribution, found using the Mathematica® FindMinimum
command, to fit a beta-binomial distribution to the observations. The minimum
expected cell size is 0.5. That number can be changed by changing 0.5 where it
appears in two places in this program, to any new desired minimum expected
value.

The second part of this program finds the minimum chi-square test statistic
using a Mathematica® FindMinimum search procedure. The cells are the same as the
ones used in the first part of this program.*);

ClearAll[phat,log,minval,apoint,bpoint,m1,m2,a0,a1];

phat=N[(f1+1)/(n1+2)];

m1=Sum[phat[[i]],{i,m}]/m//N;

m2=Sum[(phat[[i]])^2,{i,m}]/m//N;

a0=m1*(m1-m2)/(m2-m1^2);

b0=(1-m1)*(m1-m2)/(m2-m1^2);

log[a1_,b1_] :=Sum[(Sum[Log[a1+j],{j,0,f1[[i]]-1}]
+Sum[Log[b1+j],{j,0,n1[[i]]-f1[[i]]-1}]
```

```
-Sum[Log[a1+b1+j],{j,0,n1[[i]]-1}]),{i,m}];

{minval,{apoint,bpoint}}=FindMinimum[N[-log[adum,bdum]],
     {adum,a0,0.01,Max[m,20*a0]},{bdum,b0,0.01,Max[m,20*b0]}];

a=adum/.apoint;

b=bdum/.bpoint;
ClearAll[betabinomial];
betabinomial[f_,n_] := Binomial[n,f]*Gamma[a+f]*Gamma[b+n-f]*Gamma[a+b]/
   (Gamma[a]*Gamma[b]*Gamma[a+b+n]);

i1=0;
ClearAll[e2,e3,o2,o3];
e2[0]=Sum[betabinomial[0,n[i]],{i,m}]//N;
sumbb=e2[0]//N;
mx=Max[n1];

j2=1;
While[sumbb<m-0.5,(e2[j2]=Sum[
 If[(N[j1/n[i]]<=N[j2/(mx+1)])&&(N[j1/n[i]]>N[(j2-1)/(mx+1)]),
  N[betabinomial[j1,n[i]]],0],{i,m},{j1,j2}])
     &&(sumbb=sumbb+e2[j2])&&(j2=j2+1)];

k2max=j2-1;
e2[k2max]=m-Sum[e2[i],{i,0,k2max-1}];
ClearAll[j2,sumbb,beg2,beg3,end2,end3];

o2[0]=Count[f1,0];

Do[o2[j]=Sum[If[(N[f[i]/n[i]]<=N[j/(mx+1)])&&
         (N[f[i]/n[i]]>N[(j-1)/(mx+1)]),1,0],
         {i,m}],{j,k2max-1}];
o2[k2max]=m-Sum[o2[i],{i,0,k2max-1}];

i1=1;i3=1;j1=0;
beg2[1]=0;
end2[1]=0;

Do[If[e2[i2]<0.5,
(j1=j1+1)&&(end2[i3]=j1)&&(e2[j1]=e2[i2]+e2[j1]),
((e3[i1]=e2[i2])
&&(o3[i1]=Sum[o2[i3],{i3,beg2[i3],end2[i3]}])
&&(beg3[i3]=N[(beg2[i3]-1)/(mx+1)])&&(end3[i3]=N[end2[i3]/(mx+1)])
&&(i1=i1+1)&&(i3=i1)&&(j1=j1+1)&&(beg2[i3]=j1)&&(end2[i3]=j1))],
{i2,0,k2max}];
beg3[1]=0;

k2=i1-1;
end3[k2]=1;

ClearAll[i1,j0,j1,i3];

ClearAll[chisq2,pvalue2];
```

```
chisq2=(Sum[(o3[i])^2/e3[i],{i,k2}]-m)//N;

df2=k2-3 (* This subtracts two degrees of freedom for the two estimated
parameters, a and b*);

pvalue2=1-CDF[ChiSquareDistribution[df2],chisq2] (*If the p-value is > .05
then the data could have come from the distribution producing the e[]'s,
but if the p-value is less than or equal to .05 the fit is poor.*);

ClearAll[minchisq2,maxpvalue2,amin2,bmin2,e1min3,emin3];

Do[emin3[i1,i3_,a2_,b2_] :=Sum[Sum[
If[(N[j1/n[i]]<=N[j2/(mx+1)])&&(N[j1/n[i]]>N[(j2-1)/(mx+1)]),
(Binomial[n[i],j1]*Gamma[a2+j1]*
Gamma[b2+n[i]-j1]*Gamma[a2+b2]/(Gamma[a2]*Gamma[b2]*Gamma[a2+b2+n[i]])),0],
{j1,0,j2},{i,m}],{j2,beg2[i3],end2[i3]}],{i1,k2-1}];

{minchisq2,{amin1,bmin1}}=FindMinimum[N[Sum[(o3[i])^2/emin3[i,i,adum,bdum],
{i,k2-1}]+(o3[k2])^2/(m-Sum[emin3[i,i,adum,bdum],{i,k2-1}])-m],
 {adum,a,0.01,Max[m,20*a]},{bdum,b,0.01,Max[m,20*b]}];

amin2=adum/.amin1;

bmin2=bdum/.bmin1;

Do[e1min3[i]=emin3[i,i,amin2,bmin2],{i,k2-1}];
e1min3[k2]=m-Sum[e1min3[i],{i,k2-1}];

maxpvalue2=1-CDF[ChiSquareDistribution[df2],minchisq2])
```

B.9: beta-bin-mle-gof3.mat

```
(
(*This program uses (t+1)/(n+2) as the estimator of the failure rate, and
compares the observed values of this estimator from several data sources to the
expected values from the beta-binomial probability distribution.*);

(*This program takes data from a file called "temp" that has been processed
using a file called "beta-bin-setupfile.mat". See "beta-bin-help.mat" for more
information and detailed instructions.*);

(*This program uses maximum likelihood estimators of alpha and beta in the
beta-binimial distribution, found using the Mathematica® FindMinimum command, to
fit a beta-binomial distribution to the observations. The minimum expected cell
size is 0.5. That number can be changed by changing 0.5 where it appears in two
places in this program, to any new desired minimum expected value.*);

ClearAll[phat,log,minval,apoint,bpoint,m1,m2,a0,a1];

phat=N[(f1+1)/(n1+2)];

m1=Sum[phat[[i]],{i,m}]/m//N;
```

```
m2=Sum[(phat[[i]])^2,{i,m}]/m//N;

a0=m1*(m1-m2)/(m2-m1^2);

b0=(1-m1)*(m1-m2)/(m2-m1^2);

log[a1_,b1_] :=Sum[(Sum[Log[a1+j],{j,0,f1[[i]]-1}]
+Sum[Log[b1+j],{j,0,n1[[i]]-f1[[i]]-1}]
-Sum[Log[a1+b1+j],{j,0,n1[[i]]-1}]),{i,m}];

{minval,{apoint,bpoint}}=FindMinimum[N[-log[adum,bdum]],
    {adum,a0,0.01,Max[m,20*a0]},{bdum,b0,0.01,Max[m,20*b0]}];

a=adum/.apoint;

b=bdum/.bpoint;

ClearAll[betabinomial];
betabinomial[f_,n_] := Binomial[n,f]*Gamma[a+f]*Gamma[b+n-f]*Gamma[a+b]/
   (Gamma[a]*Gamma[b]*Gamma[a+b+n]);

sumbb=0;
i1=0;
ClearAll[e4,e5,o4,o5];
sumbb=0;

mx=Max[n1];
ClearAll[frac2,am2,bm2];
Do[Do[frac2[i,j]=(j+1)/(n[i]+2),{j,0,n[i]}],{i,m}];

Do[am2[i,j]=0,{i,m},{j,0,mx}];

Do[Do[Do[If[(frac2[i,j1]<=(j+1)/(mx+3))&&(frac2[i,j1]>j/(mx+3)),
 am2[i,j]=am2[i,j]+j1],
  {j1,0,n[i]}],{j,0,mx}],{i,m}];

j1=1;
While[sumbb<m-0.5,
    (e4[j1]=Sum[If[am2[i,j1]>0,N[betabinomial[am2[i,j1],n[i]]],0],{i,m}])
    &&(sumbb=sumbb+e4[j1])
    &&(j1=j1+1)];
k3max=j1-1;
e4[k3max]=m-Sum[e4[i],{i,0,k3max-1}];
ClearAll[j1,sumbb,beg4,beg5,end4,end5];

Do[bm2[i,j]=0,{i,m},{j,0,k3max-1}];
Do[If[((f[i]+1)/(n[i]+2)<=(j+1)/(mx+3))&&((f[i]+1)/(n[i]+2)>j/(mx+3)),
 bm2[i,j]=bm2[i,j]+1],
 {i,m},{j,0,k3max-1}];

Do[o4[j]=Sum[bm2[i,j],{i,m}],{j,0,k3max-1}];
o4[k3max]=m-Sum[o4[i],{i,0,k3max-1}];

i1=1;i3=1;j1=0;
```

Report Number (14) LA-UR-97-4825

Publ. Date (11) 199712

Sponsor Code (18) NRC, XF

UC Category (19) UC-000, DOE/ER

DOE