

Program Manual for FIT, DFIT, CR and Related Programs

Stuart A. Klugman

November 13, 1997

Contents

1	Introduction	1
1.1	Guarantees, Disclaimers, and Help	1
1.2	Installing the programs	2
1.3	Running the programs	2
1.4	Using the menus	3
2	FIT	5
2.1	Data format	5
2.2	Grouped data	5
2.3	Individual data	8
2.4	Distributions	8
2.5	Fitting the model	9
2.6	Results	10
2.7	Plots	11
3	DFIT	13
3.1	Data format	13
3.2	Distributions	14
3.3	Fitting the model	14
3.4	Results	14
3.5	Plots	15
4	CR...	17
4.1	Distributions	17
4.2	Computing and printing the probabilities	19
4.3	FREQADJ	21
4.4	Ruin	22
5	Comprehensive example	23
5.1	Models for the three data sets	24
5.1.1	Model for aggregate small losses	24
5.1.2	Model for the amount of individual large losses	25
5.1.3	Model for the number of large losses per family	26

5.2	A model for aggregate losses	26
5.2.1	Aggregate small losses	26
5.2.2	Aggregate large losses	27
5.2.3	Total aggregate losses	27
6	Miscellaneous	29
6.1	DOS Command	29
6.2	Menu Options	30
6.2.1	FIT and DFIT	30
6.2.2	CR	30
6.2.3	RUIN	31
6.3	Printing	31
A	An Inventory of Continuous Distributions	33
A.1	Introduction	33
A.2	Transformed beta family	34
A.2.1	Four parameter distribution	34
A.2.2	Three parameter distributions	34
A.2.3	Two parameter distributions	35
A.3	Transformed gamma family	36
A.3.1	Three parameter distributions	36
A.3.2	Two parameter distributions	36
A.3.3	One parameter distributions	37
A.4	Other distributions	38
B	An Inventory of Discrete Distributions	39
B.1	Introduction	39
B.2	The $(a, b, 0)$ class	39
B.3	The $(a, b, 1)$ class	40
B.3.1	The zero-truncated sub-class	41
B.3.2	The zero-modified sub-class	42
B.4	Compound class	42
B.4.1	Some compound distributions	43
B.5	A hierarchy of discrete distributions	44

Chapter 1

Introduction

These programs are intended for use with text *Loss Models: From Data to Decisions*. The programs FIT and DFIT find maximum likelihood estimates for 20 continuous and 16 discrete distributions respectively. They also produce estimates of the covariance matrix of the estimators as well as graphs that illustrate the quality of the fit. The chi-square goodness-of-fit test is also performed. The program CR evaluates probabilities for the collective risk model. Frequencies used can be any of the models used in DFIT while severities can be any of the models from FIT or a model entered by the user. In all cases the probabilities are obtained using recursive formulas. Prior to the recursion, individual or aggregate deductibles, limits, or coinsurances may be applied.

1.1 Guarantees, Disclaimers, and Help

The software and its documentation are provided “as is” and Stuart Klugman disclaims all warranties, either express or implied, including, but not limited to, implied warranties of fitness for a particular purpose. Stuart Klugman does not warrant that the software will meet the user’s requirements or that the operation of the software will be uninterrupted or error free.

Stuart Klugman will not be liable for any indirect, special, or consequential damages in connection with or arising from the performance of or use of this software and its documentation. Stuart Klugman shall in no event be liable for any damages whatsoever (including without limitation, damages for loss of profits or other pecuniary loss) arising out of the use or of the inability to use the software and its documentation, even if Stuart Klugman has been advised of the possibility of such damages.

There are a variety of reasons for the above disclaimer. First of all, no program can be guaranteed to be free of bugs. The best we can hope for is that every time it fails to work properly, I will correct the problem. It is not possible to anticipate every data set which may be used and so every conceivable situation cannot be tested. Second, it is impossible for the program to verify that the data you have entered means what you think it does. Third, I can only assure you that the programs run on my computer. Although there is a great deal

of compatibility between systems, yours is not exactly like mine.

Should you have comments or questions about the use of these programs, you may contact the author electronically at Stuart.Klugman@Drake.edu. Please understand that I am under no obligation to respond as your acquisition of these programs confers no rights with regard to service or help. On the other hand, I would like to make them as useful as possible for readers of the book and will make every effort to respond to inquiries. A commercial version of these programs, which provides unlimited help, notices of upgrades, and an anticipated conversion to Windows 95 in late 1998 is available for purchase from the author.

1.2 Installing the programs

The programs are in the files FIT.EXE, DFIT.EXE, CR.EXE, FREQADJ.EXE, CRC.EXE, CRC2.EXE, and RUIN.EXE and should work with any IBM compatible computer. If the 80x87 coprocessor is installed or is built in to your main processor, the program will detect and use it. If the coprocessor is not present, it is emulated, so all users enjoy the additional accuracy provided, but the coprocessor itself is needed to enjoy the extra speed. In order for the graphics to work one of the following files should be in the same directory as the program.

File	Graphics adapter
CGA.BGI	IBM CGA and MCGA
EGAVGA.BGI	IBM EGA and VGA
HERC.BGI	Hercules monochrome
PC3720.BGI	IBM 3270PC

I believe that if you have higher resolution capabilities (such as super-VGA or IBM 8514) the program will read the EGAVGA.BGI file and produce VGA resolution output.

To install the program, copy the program files and the appropriate graphics driver to a directory on your hard disk (or you can run from a floppy, but it will take longer to load).

1.3 Running the programs

Any program can be run from any directory either by typing the complete path to the program or having the program's directory in your path statement. To run, just type the program name, e.g., FIT, DFIT, etc., at the DOS prompt. If all is working well you should see a display in the middle of your screen that reads (depending on the program)

```
Severity Fit – Version LM1
  by Stuart Klugman
Copyright, December 1997
For use with Loss Models:
  From Data to Decisions
Press any key to continue
```

or

```
Discrete Fit – Version LM1
  by Stuart Klugman
Copyright, December 1997
For use with Loss Models:
  From Data to Decisions
Press any key to continue
```

or

```
Collective Risk – Version LM1
  by Stuart Klugman
Copyright, December 1997
For use with Loss Models:
  From Data to Decisions
Press any key to continue
```

The programs CRC and CRC2 and FREQADJ do not use windows, the information appears directly on the screen. All are at Version LM1, December 1997. RUIN does use windows and is at Version LM1, December 1997. Pressing any key at this point will start the program. The distribution disk contains five data sets: FIT1.DAT and FIT2.DAT contain grouped data and can be used with FIT while FIT3.DAT contains individual data for FIT. DFIT1.DAT and DFIT2.DAT can be used to try out DFIT.

1.4 Using the menus

The neutral state for each program (other than CRC, CRC2, and FREQADJ) has a window in the upper left hand corner. For FIT and DFIT it displays the current model, parameter values, and negative loglikelihood value. For CR the neutral state window displays the current frequency and severity models, parameter values, and coverage modifications. For RUIN the input parameters are displayed. The message **Select an option** is at the bottom

of the window. Any letter (upper or lower case) entered at this time will cause a command to be executed. The letter M will display the menu options. Command letters can also be entered when the menu is being displayed. The option X terminates the program. A list of the options for each program appears in at the end of the manual.

Chapter 2

FIT

FIT obtains maximum likelihood estimates for severity data that may be grouped or individual (ungrouped). The data may be truncated from below and censored or truncated from above. In all cases the model that is found is for losses starting at zero (ground-up). For each fitted model, values of the distribution function and the limited expected value can be computed.

2.1 Data format

Data may be entered with the program or using any word processor that produces ASCII output with no special characters. The DOS (5.0 or higher) editor (EDIT) is good for this purpose. The MS-Windows editor Notepad will also work. To see what a data file looks like, view any of the sample files using one of these editors, or have them listed on the screen by typing `type fit1.dat` or have one printed out by typing `copy prn fit1.dat`.

Use option C to create a data set. You will be prompted for the name of the file that will hold the data. If the file is to be saved on a directory other than the current one, the complete path must be specified.

2.2 Grouped data

For grouped data, at each subsequent prompt enter the lower limit for the interval, a space, the number of observations in that interval, a space, and the average of the observations in that interval (if you do not have the averages, you can so indicate and the program will insert the midpoints). The averages are used only for computing the empirical limited expected values; they have no influence on the maximum likelihood estimates. The lower limit entered on any line is both the lower limit for that interval and the upper limit for the previous interval. For the purpose of maximum likelihood estimation it does not matter which endpoints are included in an interval, but you should be consistent. The last line

should begin with the upper limit of the final interval and then -10 for the other two entries. If the final interval has infinity as its upper limit, enter -1 for the boundary.

If values above the final grouping were possible, but just not observed in this sample, an interval with no observations and infinite upper limit should be included. The data set in FIT1.DAT is an example. To repeat, if the final interval entered has a finite upper limit, the program will take this to mean that values above that limit were not possible (truncation from above) as opposed to just did not occur in the sample.

If the data set is truncated from below (that is, there was a deductible), reflect this by simply having the first interval begin with the deductible. Note that the program is expecting the amount of the loss, not the amount paid after the deductible was subtracted.¹

Once again, if the interval from 0 to x had no observations, but observations were possible, that interval must be included. Starting with an interval from $x > 0$ to y indicates that values below x could not be observed.

Example 2.2.1 *The following examples indicate some possibilities:*

0	200	50
100	100	200
300	50	400
500	0	700
-1	-10	-10

Observations above 500 were possible, there just were not any observed.

0	200	6
12	100	18
24	50	30
36	-10	-10

Observations above 36 were impossible (truncation from above). This might be data on the time in months from policy issue to claim report. There were no policies issued more than 36 months ago.

100	100	150
200	50	250
500	10	750
-1	-10	-10

Values below 100 were impossible (truncation from below, as in a deductible). There may have been a policy limit of 500 (censoring from above) with 10 observations at the limit. Had there been a limit, the average of 750 for losses above 500 would not have been observable.

¹The amounts paid after subtracting the deductible could be entered, but then the model found would be for amounts paid, not for ground-up losses.

For maximum likelihood estimation, the values in the third column are not used. Their only purpose is to provide empirical limited expected values for comparison to the model values. Had there been a limit, the value of 750 would be used to calculate the the empirical limited expected value at infinity (the sample mean) to compare to the model mean. With a policy limit this comparison is not valid and so the value can be ignored.

You can also input data that is from various groupings. Terminate each grouping with the entries $\mathbf{xx} \ -5 \ -1$ where \mathbf{xx} is the upper limit of the last group in that set (use -1 in place of \mathbf{xx} for infinity). The data set FIT2.DAT is an example. There is one more feature that can be incorporated into the data set. When there are multiple limits (for example) the data can be easily entered if the original frequency tables were kept separate.

Example 2.2.2 *Consider data collected from policies with two limits, 1,000 and 5,000:*

0	100	50
100	50	300
500	60	750
1000	30	2000
-1	-5	-1
0	30	50
100	20	300
500	25	750
1000	10	3000
5000	5	13000
-1	-10	-10

There were 240 losses on policies with the 1,000 limit and 90 losses with the 5,000 limit.

Example 2.2.3 *Now suppose the two groups had not been kept separate. The data set would then be entered as:*

0	130	50
100	70	300
500	85	750
1000	10	3000
5000	5	13000
-1	-1	-1
1000	30	2000
-1	-10	-10

There is now just one group with 320 losses. The -1 terminator in the second entry in line 6 indicates that the group is unchanged, but it contains some overlapping intervals. If there is truncation from above, this must be revealed by having the entry at the first -1 be the truncation point. It should be noted that when there are overlapping intervals, the empirical probabilities and LEV's are adjusted so that they can be comparable.

For option **P** the current model is utilized to provide the needed interpolation. When graphing, the current model is also used, so if several models are to be placed on one graph you should have what you think is the best model as the current model. The best results are obtained if the maximum possible number of entries are presented prior to the first **-1** as was done in the above example.

The maximum number of lines is 100.

2.3 Individual data

For individual data, values can be entered via option **C** although this may be painful. The data format is as follows (all numbers must be in ASCII code, which is easy to generate with your spreadsheet or word processor).

Line 1 – The number **-1**, to indicate individual data follows.

Line 2 – **0** if the loss entries are not sorted, **1** if they are increasing.

Line 3 – The number of observations in the data set.

Subsequent lines – The loss (total, not the payment after subtracting the deductible), followed by the deductible (**0**, if no deductible), followed by the limit (**0**, if no limit). The three numbers may be separated by spaces or by a comma. If the limit is a censoring limit (observations above the limit are included, but their value is unknown) enter the limit as is. If the limit is a truncating limit (observations above the limit are excluded) enter the negative of the limit. Any entries in the data set that are below the deductible or above a truncating limit will be ignored.

The sample file FIT3.DAT is an example. The maximum number of data points is 500.

To view the data set use option **V**. Errors can be corrected with option **E**. This is not a complete editor, but does allow you to change, delete, or add an entry for grouped data. For individual data any observation can be changed or an observation added. If major work needs to be done it is better to use an ASCII word processor such as DOS-Edit or Windows-Notepad. This can be done without leaving FIT via option **O** (DOS command, discussed later).

2.4 Distributions

Option **S** allows you to select the distribution. Brief information about the parameterization of the distributions appears in Appendix A. More detailed information is available in *Loss Models*.

2.5 Fitting the model

You must begin by specifying starting values for the parameters. Possible values for some one and two parameter distributions can be found via option B. There is no guarantee that the program will find suitable (or even legitimate) values. If the algorithm (method of moments or percentile matching, depending on the distribution) produces illegal parameter values, no starting values will be provided. Sometimes the parameters are legitimate, but strongly inconsistent with the data, causing the likelihood to be impossible to compute. This will be indicated by a negative loglikelihood of 1E100 (overflow), 1E200 (underflow) or 1E300 (both). In either case, alternative starting values must be created by you, using option N. For three and four parameter models, fit a two or three parameter member of the same family that is a special case. Option B will then select the parameters from your earlier fit and then set the additional parameter equal to 1. The Appendix can help you identify the special cases.

To enter your own starting values, use option N. You must have already selected a distribution with option S. This option can be used at any time to change the current iterate. If you are desperate for starting values try setting θ equal to the sample mean and all the other parameters equal to 2.

Two methods of doing the minimization (of the negative loglikelihood) are available. The first is the method of scoring. It is a modification of the Newton-Raphson method for finding the roots of an equation. Option I will perform one iteration of this method. The new and old loglikelihood values will be displayed so you can see if the procedure is converging. This method requires good starting values, but when it converges, it does so quickly. Continue iterating until you are satisfied with the results. If the method yields illegal parameter values, the new iterate will not be accepted and you must either get new starting values or switch to the simplex method. The scoring method is only approximate when models whose cumulative distribution function involves the incomplete gamma or beta functions are used. The final iterations must be done by the simplex method. Also, the variances and correlation matrices that are produced are approximate because they are based on approximate derivatives. The distributions for which this is the case are the gamma (5), inverse gamma (6), generalized Pareto (11), transformed gamma (12), inverse transformed gamma (13), and transformed beta (14). For the inverse Gaussian distribution (1), exact derivatives can be written in terms of the standard normal distribution function. Because this must be approximated, the scoring method will converge quickly, but to a slightly incorrect result. As with the previous distributions, it is best to finish with the simplex method. *The scoring method is not available for individual data.*

The other method is the simplex method. Use option L to initiate it. Follow the instructions. An initial simplex with values increased by 5 percent seems reasonable, but 10 to 20 percent may be more appropriate if your starting values are not good. The better you believe your starting values to be the smaller the percent that can be used. At each iteration the new value is better than at least one of the previous values. If the scoring method is available, once the simplex method appears to have gotten close to the solution, it may be

a good idea to switch.

As a final check, option **G** will compute the negative loglikelihood at a grid of values surrounding the current parameter value. You specify the percentage adjustment to make to each value and the output places a (*) next to each value that is better than the one at the current point.

You should be aware that maximum likelihood estimates do not exist for all models for some data sets. It may be that the maximum occurs when two or more parameters go to zero and/or infinity. So it is a good idea to watch the parameter values as you iterate to see if this is happening. This usually means the model is a poor choice anyway or that a model with fewer parameters will do just as well. A word of warning – when the non- θ parameters in the three and four parameter members of the transformed beta family get large (in the hundreds or more) it is likely that you have not converged (and won't) and the program will spend a long time computing under option **P** as it tries to compute various quantities via nearly infinite sums.

2.6 Results

Option **R** lists all models fitted during the session along with the most recent loglikelihood values. If the chi-square goodness-of-fit test has been done (automatic with option **P** with grouped data) this command will also present the chi-square test statistic, the degrees of freedom and the p-value. For grouped data option **P** lists more details about the current model, including variances, a correlation matrix, limited expected values, and a chi-square goodness-of-fit test. For the latter, the individual contribution of each cell to the chi-square test statistic is presented. In the presentation the number is signed to indicate if the expected count was larger (–) or smaller (+). The program groups cells together to ensure that expected counts always exceed 5. Cells which have been grouped have an entry of -1000 displayed as the “contribution” to the test statistic. The number prior to the string of -1000's is the contribution for the grouped cells. For individual data, option **P** provides only the parameter values, negative loglikelihood, variances, and the correlation matrix. To obtain the other items, use option **U** to group the data. You can then use the parameters as previously set and option **P** to display the results for the grouped version. When option **U** is used, the grouped data are placed in a new file. All previous parameters and fits are retained, but to save them as connected to the grouped data file, option **K** must be re-run with a new file name for the saved results.

As you change from one model to another (via option **S**) the program remembers the latest parameter values from all previous models. The same output that is given by **R** plus the parameter values can be saved to a file. This file can also be read in to allow a fitting session to be continued at a later time. Use option **K** to save or restore a session. **It is a good idea to save your session often.** FIT has very little in the way of error trapping, so if you or the program bomb, all your previous fits will be lost. The file **FIT1.SAV** contains the results of fitting several distributions to the data in **FIT1.DAT**.

Option **F** gives $f(x)$, $F(x)$, and $E[X;x]$ at a specified value of x . It also gives the mean

and standard deviation if they exist. This procedure ignores any truncation or censoring points that were in the original data set and computes values as if there was no deductible or limit. This option also allows you to create a table of values for $F(x)$ and $E[X;x]$ provided the entries are arithmetically spaced.

2.7 Plots

Option T creates graphs. This can only be done for grouped data. (Option U can do the grouping for individual data.) The five options are:

1. Histogram and fitted pdf for the current model.
2. Histogram and fitted pdf's for the current model and up to three other models.
3. Empirical and fitted cdf for the current model.
4. Empirical and fitted LEV for the current model.
5. Empirical and fitted LEV's for the current model and up to three other models.

You will also be given the opportunity to limit the presentation to a specified range of loss values. This lets you “blow-up” the plot to view its most important features. The lower and upper loss limits must be group boundaries. When asked if you want a two-color plot a “Y” response will produce a plot with white lines on a black background. This is best for notebooks, gray-scale projectors, or if you will be sending the graph to the printer. The quality of the plot will depend on the graphics capabilities of your machine. After viewing, hit any key to continue. Warning - this procedure will not work if the range of loss values contains more than 150 groups.

See the end of the manual for information on printing the graphs.

Chapter 3

DFIT

DFIT finds maximum likelihood estimates for discrete data. All of the models assume that probabilities are associated with non-negative integers. Some of the distributions assume zero is not a possible value.

3.1 Data format

DFIT requires a file containing the sample data. Data may be entered with the program or using any word processor that produces ASCII output with no special characters. Try using your word processor to read the file DFIT1.DAT (or just type `type DFIT1.DAT`) to see what a data set looks like.

Use option C to create a data set. You will be prompted for the name of the file that will hold the data. If the file is to be on another directory or disk you must give the complete path. At each subsequent prompt, enter the number of times that value occurred in the sample. All of the distributions used in this program assume that the lowest possible value for the random variable is either 0 or 1. If 0 is not a possible value, you still must enter 0 as the number of times 0 was observed. When 0 is impossible, only models in the AB1 class “without zeros” should be used. To indicate that no more data will be entered, enter -10. The program assumes that the last entry prior to the -10 represents the count for that value and all larger values. The only exception is for the binomial distributions where the value of m is input by the user. It is assumed that the number of observations for values larger than the last entry up through m are all zero. If all of the values were recorded as seen, then the last entry prior to the -10 should be zero. This is the case for the two sample sets.

It is also possible for data to be grouped. In that case enter the frequency for the lower limit for the group and enter -1 for the frequency for other values in that group. However, the program requires that the first group contain only 0 (that is, the first group cannot be 0-3, for example). Other groups can be any size.

Example 3.1.1

```

3
6
-1
-1
7
-1
-1
5
-1
3
-10
```

The actual observations would be 3 at 0, 6 in the interval 1-3, 7 in the interval 4-6, 5 in the interval 7-8, and 3 in the interval 9-∞.

To view the data set, use option V. Errors can be corrected with option E. This is not a complete editor, but does allow you to change, delete, or add an entry. If major work needs to be done it is better to use a word processor. You can run another program without leaving DFIT via option O.

To use a data set, select it with option D. You will be prompted for the file name. The complete path needs to be specified. If you are unsure of the file's location, use option O to view your directories.

The data set can have a maximum of 100 entries.

3.2 Distributions

Option S allows you to select the distribution. Brief information about the parameterization of the distributions appears in Appendix B. More detailed information is available in *Loss Models*.

3.3 Fitting the model

Fitting is done exactly as with FIT. The scoring method is always available and is a good choice unless the starting values are poor. If useful starting values are hard to obtain, consider setting any λ or β parameter to the sample mean and all other parameters to 1.

3.4 Results

Option R lists all models fitted during the session along with the most recent loglikelihood values as well as results from the chi-square goodness-of-fit test for those models that have been examined with option P. Option P lists more details about the current model, including a correlation matrix. As you change from one model to another (via option S) the program

remembers the latest parameter values from all previous models. The same output that is given by R along with the parameter values can be saved to a file. This file can also be read in to allow a fitting session to be continued at a later time. Use option K to save or restore a session. **It is a good idea to save your session often.** DFIT has very little in the way of error trapping, so if you or the program bomb, all your previous fits will be lost. DFIT1.SAV contains the results of some fits using DFIT1.DAT. Option F will allow you to compute the exact and cumulative probabilities for any value. It also gives the mean and standard deviation for the model.

3.5 Plots

Option T allows you to view the results of the estimation. There are two options. They are

1. Bar charts for the data and the current model.
2. Bar charts for the data, for the current model, and for up to three other models.

Except for the smaller number of options, all else is the same as for FIT.

Chapter 4

CR/FREQADJ/CRC/CRC2/RUIN

The program CR computes probabilities for the model $S = X_1 + \dots + X_N$ where N is a random variable called the frequency and the X_i are identically distributed random variables called the severity. It is assumed that all the random variables are independent. The frequency model can be selected from any of those available in the program DFIT and the same parameterizations are used. The severity model can be any of those used in the program FIT or it can be an arbitrary distribution created by the user. The severity must be a discrete distribution with probabilities at equally spaced values. Because the distributions in FIT are all continuous, the program creates a discrete distribution that is a close approximation. The user can decide what the spacing should be.

The severity distribution may be modified by imposing a deductible, limit, and or co-insurance. The same is true for the aggregate distribution. In addition to the probabilities, the program also produces limited expected values and limited expected squared values for the aggregate distribution.

The program FREQADJ is used to change the parameters of the frequency distribution in order to reflect modifications due to truncation.

The programs CRC and CRC2 are general purpose convolution programs. CRC forms the convolution of two identical distributions while CRC2 forms the convolution of two different distributions.

The program RUIN computes finite time, discrete process ruin probabilities.

4.1 Distributions

Option F allows you to select the frequency distribution. A menu appears offering the various choices. After your choice is made you will be prompted for the parameter values. The final option is called “Exactly One Claim” and provides a frequency distribution for which the probability of one claim is 1. This makes the output aggregate distribution the same as the (modified) severity distribution.

Option S allows you to select the severity distribution. A menu appears offering the various choices. The last choice is for a user specified distribution. If this one is selected,

Now suppose an ordinary deductible of 4 is applied prior to the discretization. This gives the probability of a particular amount being paid given that a payment is to be made. The results are

$$\begin{array}{ll} g_0 & \Pr(4 < X \leq 5) / \Pr(X > 4) = .186963 \\ g_1 & \Pr(5 < X \leq 7) / \Pr(X > 4) = .254519 \\ g_2 & \Pr(7 < X \leq 9) / \Pr(X > 4) = .158460 \\ g_3 & \Pr(9 < X \leq 11) / \Pr(X > 4) = .103762 \\ \vdots & \vdots \end{array}$$

If a limit of 10 (6 above the deductible) and coinsurance of .75 are to be applied the only change is to concentrate all of the excess probability at g_3 . That is, $g_3 = .400058$.

For the last two cases, the frequency distribution must be modified to reflect the per payment basis. Using FREQADJ, the new zero-modified geometric parameters are $p_0 = .747049$, $\beta = .728863$.

4.2 Computing and printing the probabilities

Aggregate probabilities are computed using the recursive formula as discussed in Chapter 4 of *Loss Models*. Aggregate modifications can be imposed by selecting option A. This will not affect the recursions as the modifications are imposed only when option P is selected for viewing the results. Recursions will be carried out to the earliest of 3,000 values, the finding of all but about one hundred-millionth of the probability, or a value preset when option C was invoked. So the aggregate modifications can be changed and their effects noted without redoing the recursions (option C). However, if you know that there will be an aggregate limit that will not be raised, you can save time by only doing recursions to just a little bit past this limit.

Option C asks the program to compute the probabilities for the collective risk model. You will be prompted for the starting value for recursions (see the discussion below on modifying the recursions). If you enter zero, ordinary recursions giving true probabilities will be done. You can also enter the aggregate amount at which aggregate loss calculations are to stop. If you enter -1 the program will terminate on its own, when either 0.99999999 of the probability has been found or 3,000 values have been set. You should always select -1 when recursions start at a value above zero.

Output is always computed at the same spacing as used for the severity distribution. This can cause some problems when the expected frequency is high. For example, a spacing of 5 may work well for a severity distribution, creating 500 intervals (so nearly all the probability is assumed to be below 2,500). If the expected frequency is 100 claims, this would require perhaps 25,000 intervals of width 5 for the aggregate distribution, well beyond the memory of the program. To solve this problem there are two possible approaches.

Begin the first method by running CR with the expected frequency divided by a power of 2 (2, 4, 8, etc.) to make it reasonable. For the Poisson or any Poisson-x distribution,

divide the Poisson parameter λ by this number. For the negative binomial, or any negative binomial- x distributions, divide the value of r by this number (recall that the geometric distribution is just another negative binomial distribution, with $r = 1$). For the binomial distribution, divide m by this number (the result must be a whole number). Obtain the aggregate distribution and when complete, write the results to disk using the option that allows for output suitable for input to the convolution program. Then run CRC using this file as input. The output will be the convolution of the aggregate distribution with itself. The output can be written to disk in a format suitable for re-input to CRC or in a form that matches the output from CR. This process must be repeated as many times as the power of 2 that was used (for example, if division was by 8, run CRC 3 times).

The program CRC2 is a general purpose convolution program that produces the convolution of two different discretized distributions. An example would be two separate aggregate distributions for two lines of business, each produced by program CR. Input is in the same format as for CRC.

The second method is to start the recursive calculations at a value higher than zero. This would be appropriate if the high frequency has pushed the probabilities well above zero. A reasonable place to start is four to six standard deviations below the mean. The program temporarily assigns a probability of one to this value and then proceeds recursively. Recursions continue until the allotted 3,000 spaces are used or the probabilities become relatively small. Once the recursions are completed the probabilities are rescaled to add to one. The major drawback to this approach is that you cannot be sure that all of the probability has been accounted for when the recursions stop (when recursions start at zero the true probabilities are obtained and so the remaining probability can be placed at the last value). The program will provide a warning when it appears that there is a problem.

Option P will print the cumulative probabilities, the limited expected values, and the limited expected squared values. You have the option of printing the results on the screen, on the printer, or to a file. As noted above, you can also write the output to disk in a form suitable for input into CRC, CRC2, or RUIN. This form is also suitable for input as a user specified severity distribution for CR itself. Finally, you also have the option of printing a table of percentiles.

Example 4.2.1 *Use program CR to obtain the aggregate distribution of losses from the model described in the previous example. Impose the individual ordinary deductible of 4, censoring limit of 10 and coinsurance of .75. The output shown below is from option P with output written to disk and then imported to this word processor.*

Deductible applied prior to discretization:

Frequency: Zero-modified Geometric

Beta = 7.2886300E-0001

p0 = 7.4704900E-0001

Severity: Pareto

Alpha = 3.0000000E+0000

Theta = 1.0000000E+0001

Discretization interval = 1.5000000E+0000
 Discretization method: Rounding
 Individual ordinary deductible: 4.0000000E+0000
 Individual censoring limit: 1.0000000E+0001
 Individual coinsurance: 0.750000
 No aggregate deductible
 No aggregate limit
 Aggregate coinsurance: 1.000000

AggLoss	PDF	CDF	LEV	LEV2
0.00	0.747049000	0.747049000		
0.00	0.031593986	0.747049000	0.0000	0.0000
5.00	0.025808282	0.916343849	0.8694	3.7214
10.00	0.006790942	0.979890827	1.0978	6.9186
15.00	0.001668970	0.994857012	1.1548	8.2800
20.00	0.000397917	0.998736040	1.1688	8.7539
25.00	0.000096620	0.999691031	1.1722	8.9035
30.00	0.000023071	0.999923851	1.1730	8.9483
35.00	0.000005578	0.999981614	1.1732	8.9612
40.00	0.000001385	0.999995536	1.1733	8.9648
45.00	0.000000334	0.999998902	1.1733	8.9658
50.00	0.000000081	0.999999735	1.1733	8.9661
55.00	0.000000020	0.999999936	1.1733	8.9662
60.00	0.000000032	0.999999984	1.1733	8.9662
65.00	0.000000000	1.000000000	1.1733	8.9662
70.00	0.000000000	1.000000000	1.1733	8.9662

Mean = 1.1733

Std. Dev. = 2.7549

4.3 **FREQADJ**

This program is a small utility that can be used to adjust the frequency distribution. This usually happens when the original frequency distribution is based on data from one deductible and/or truncating limit and the discretized distribution is appropriate for a different deductible/limit combination. The program asks for the current frequency distribution, the current severity distribution, the current deductible/limit combination and the new deductible/limit combination. The output is simply the new parameters for the frequency distribution. If needed in the middle of running program CR, option O can be used to run this program and then return to CR.

4.4 Ruin

This program uses convolutions to compute finite time ruin probabilities for a discrete time model. The input is the initial surplus, the annual premium (assumed paid at the beginning of the year), the interest rate, and the file containing the aggregate loss distribution. The latter should be a file produced by CR, CRC, or CRC2 using the output for convolutions option. It is assumed that all claims are paid at the end of the year. The program will then generate the distribution of surplus after a given number of years (option R allows you to select the number of additional years to run the model). The output (view with option P) is the probability of being ruined within that many years as well as the mean and standard deviation of surplus, given that ruin has not occurred.

Chapter 5

Comprehensive example

The format of this example is probably unlike any practical setting. It has been selected in order to illustrate many of the capabilities of these programs. The numbers themselves are made-up and are not an attempt to mimic any real data.

The problem is to learn about aggregate losses on family dental coverage. Three data sets are available. The first represents losses on small, individual claims. For each family the annual aggregate total of losses below 250 was recorded. That is, individual losses above 250 were excluded from the total.

Amount of loss	Number of losses
0-250	71
250-350	327
350-400	167
400-450	123
450-500	97
500-600	128
600-700	103
700-800	67
800-1,000	68
1,000-1,500	25
1500-	1

Individual losses above 250 were recorded individually. There were 384 such losses which were grouped as follows

Amount of loss	Number of losses
250–500	251
500–750	80
750–1,000	25
1,000–1,500	15
1,500–2,000	8
2,000–2,500	3
2,500–	3

The final piece of information is data on the number of large (excess of 250) claims by one family in one year:

Number of losses	Number of families
0	95
1	130
2	100
3	38
4	16
5	3
6	1
7	1

The three data files that correspond to this information are available as MAN1.DAT, MAN2.DAT, and MAN3.DAT respectively.

5.1 Models for the three data sets

5.1.1 Model for aggregate small losses

The first step in the fitting process is to use FIT to obtain a model for the amount of aggregate small losses in one year. After fitting all possible models (a few did not converge), the best models with each number of parameters were selected (the four parameter model did not converge). In each case the model with smallest negative loglikelihood was also the model with the smallest χ^2 -value for the goodness-of-fit test.

No. param.	Model	-loglikelihood	chi-square	df	p-value
1	Exponential	3100.94	1173.14	9	near 0
2	Inverse gamma	2506.30	32.98	7	near 0
3	Inverse trans. gamma	2505.32	31.99	7	near 0

The likelihood ratio test indicates that the inverse gamma model is an improvement over the exponential (twice the difference of the negative loglikelihoods is greater than 3.84, the 5% critical value with one degree of freedom), while the inverse transformed gamma model is not an improvement over the inverse gamma. Figure 5.1.1 shows that the inverse gamma model is satisfactory, even though it does not pass the goodness-of-fit test.

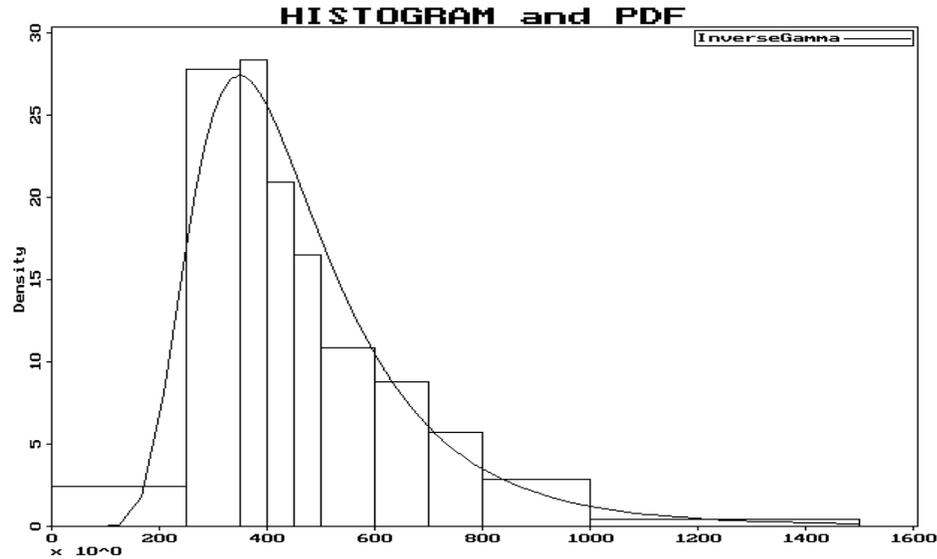


Figure 5.1.1: Inverse gamma distribution

The parameter estimates are $\alpha = 6.73021$ and $\theta = 2692.53$. Further information about this model, including standard deviations of the estimators can be found with option P.

5.1.2 Model for the amount of individual large losses

Using program FIT, a number of models were successfully fit to the data. The best fitting models with each number of parameters are summarized below.

No. param.	Model	-loglikelihood	chi-square	df	p-value
1	Exponential	427.08	9.82	2	.0074
2	Inverse Weibull	411.15	1.75	3	.6270
3	Burr	411.01	1.37	2	.5032

Identical reasoning to that from the previous subsection indicate that the inverse Weibull is the distribution of choice. This one also passes the goodness-of-fit test and had the highest p-value of any model. The quality of the fit is evident from Figure 5.1.2

The parameters for the selected inverse Weibull model are $\tau = 2.38148$ and $\theta = 317.676$.

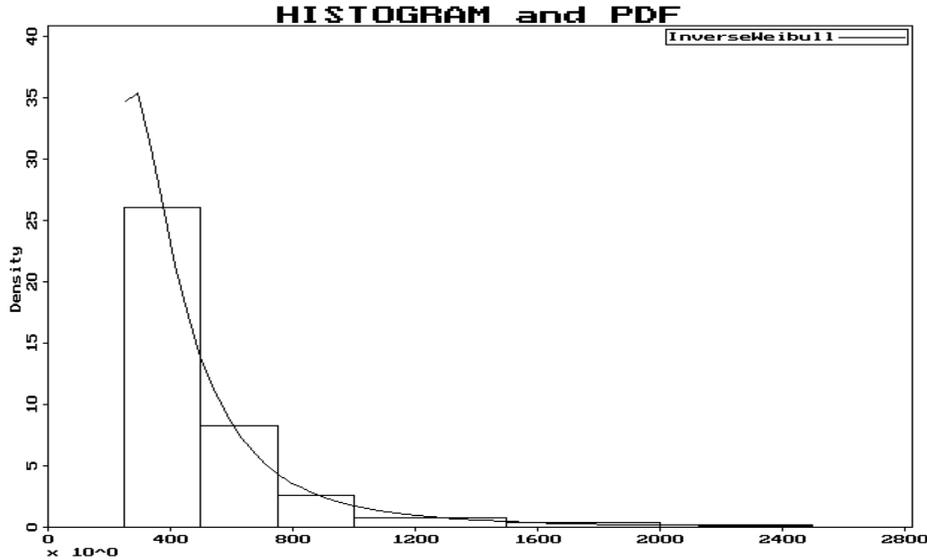


Figure 5.1.2: Inverse Weibull distribution

5.1.3 Model for the number of large losses per family

The data were entered into DFIT. Of the models which converged, the best were

No. param.	Model	-loglikelihood	chi-square	df	p-value
1	Poisson	574.97	1.34	4	.8549
2	Poisson-Inv. G.	574.96	1.40	3	.7066
2	Polya-Aeppli	574.96	1.39	3	.7077
3	Poisson-ETNB	574.88	1.53	2	.4646

Two two-parameter models are presented as one had the best loglikelihood number (Poisson-inverse Gaussian, in the third decimal place) and the other had the best chi-square value. However, the likelihood ratio test indicates that the Poisson model cannot be rejected in favor of one with more parameters and as well, it has the highest p-value for the goodness-of-fit test.

The parameter is $\lambda = 1.39583$.

5.2 A model for aggregate losses

5.2.1 Aggregate small losses

These have already been modeled with the inverse gamma distribution. To use this distribution we need only discretize it. This can be done with program CR, entering the inverse

gamma as the severity distribution, discretizing with no modifications (a span of 5 was used with the method that matches the first moment), selecting a frequency of exactly one claim, hitting C to make the transfer, and then using option P to write the distribution to disk in a format that can be used for later input.

5.2.2 Aggregate large losses

The severity for individual large losses is inverse Weibull. However, this is a ground-up distribution and we want to include only those losses in excess of 250. This can be done with a franchise deductible of 250. Because the model for the number of claims was based on those in excess of 250 there is no need to adjust the frequency distribution. The results are then saved to disk in a format that can be used for later input. A span of 25 was used with the method that matches the first moment.

5.2.3 Total aggregate losses

This is just the convolution of the two distributions found previously. This is done with program CRC2. The results are then saved to disk and then re-read into program CR with a frequency distribution of “exactly one claim.” Then output items of interest can be obtained. For this problem the following moments and percentiles were requested:

Mean	1,230
Std.Dev.	969
50th pctile	1,032
75th pctile	1,574
90th pctile	2,241
95th pctile	2,765
99th pctile	4,269

Modifications can then be entered as aggregate, not individual, modifications. This was done here to evaluate the layer from 2,500 to 4,000. The mean turns out to be 48 and the standard deviation 224.

Chapter 6

Miscellaneous

6.1 DOS Command

Option O lets you execute a DOS command without leaving the program. After selecting this option you can either enter the DOS command (e.g. "DIR") and then be returned to the program as soon as execution is completed, or hit -return- and be placed back in DOS (for example, to do word processing) and then be returned to the program only when you type EXIT in response to a DOS prompt (>).

6.2 Menu Options

6.2.1 FIT and DFIT

Option	Action
B	Starting values
C	Create a new data set
D	Read in a data set
E	Edit a data set
F	Functional values from the current model
G	Likelihood values from a grid around the current values
I	One iteration of the scoring method
K	Save or restore the results of a fitting session
L	Simplex minimization
M	Display this menu
N	New starting values
O	DOS command
P	Print results for the current model
R	Summarize results for all models used
S	Select a model
T	Plot PDF, CDF, or LEV
U	Group individual data (FIT only)
V	View the data
X	Terminate the program

6.2.2 CR

Option	Action
A	Input aggregate modifications
C	Compute aggregate probabilities
F	Select the frequency distribution
M	View this menu
O	DOS command
P	Print results
S	Select the severity distribution
V	View the severity probabilities
X	Terminate the program

6.2.3 RUIN

Option	Action
A	Get aggregate distribution
M	View this menu
O	DOS interface
P	Print results
R	Compute ruin probabilities
V	Select initial values
X	Terminate the program

6.3 Printing

There are four ways to print the graphs.

(a) The DOS 5.0 and above graphics command will work with various printers. If you have one, run the graphics command prior to running FIT or DFIT. Hitting the PrintScreen key (or perhaps Shift or Alt-PrintScreen) will cause the graph to be printed.

(b) There are a number of commercial screen capture routines. These can be used to write your graph to disk in some format that can be read by your word processor. You can then use your word processor to edit and print the graph.

(c) The distribution disk contains two programs, LJPS.EXE and LJPSRUN.EXE. Both are designed to enable you to print the contents of a graphics screen on an HP Laser Jet or Desk Jet printer. They will not work with a Hercules graphics card, nor will they work with Super-VGA. They will work with CGA, EGA, and regular VGA cards (Even if you have Super-VGA, I believe FIT and DFIT only use the resolution of regular VGA, so these programs will probably work anyway.) When running, all you need to do is hit PrintScreen (or Shift or Alt-PrintScreen). These programs will work with any screen, not just those produced by FIT and DFIT.

The program LJPS is a memory resident program. Run it by typing LJPS at the DOS prompt. Once run, it will always be the program by which screens dumps are handled. There is no way to unload this program.

LJPSRUN is not memory resident, but instead invokes a secondary DOS session. Run it by typing LJPSRUN at the DOS prompt. Later on you may type EXIT at the DOS prompt to cancel the program and the secondary DOS session. A more useful way to use this program is to type LJPSRUN FIT or LJPSRUN DFIT at the DOS prompt. This will create the secondary session, run the print screen handler, and then run the requested program. When the requested program is terminated, the secondary session will be ended and the print screen handler unloaded.

Both programs assume your printer is connected to LPT1, the first parallel printer port. If your printer is connected to LPT2 or LPT3, type LJPS -p2 or LJPS -p3 when running the first program or LJPSRUN -p2 or LJPSRUN -p3 when running the second one. You may still append a program name afterward when running LJPSRUN.

(d) If you are using Windows 3.0 or higher in enhanced mode, the following is available. First, run (D)FIT from within Windows. The default PIF works just fine. So just call the program using the RUN command from the FILE menu in either program manager or file manager. Create the desired graph (for best results, use the “two colors” option when making the graph and then hit Alt-PrtScr to copy the screen to the Windows clipboard. Then exit (D)FIT or at least switch to the Paintbrush program. Be sure the window is maximized. Choose Image Attributes from the Options menu and change the size of the picture to the resolution of your screen, using pels as the units. For example, with EGA use 640 and 350. From the View menu choose Zoom Out. From the Edit menu choose Paste. Cross hatches should appear. Place the cross-hair cursor in the upper left corner of the drawing area and click the left mouse button. The picture should now appear. Choose Zoom In from the View menu.

The colors must now be changed. The general routine is to select the old color on the color palette and click the left mouse button. Select the new color on the color palette and click the right mouse button. Then double click the color erase tool (the eraser on the left) to change every occurrence of the old color to the new color. Because the graph takes up more room than can be displayed, you will have to click the horizontal scroll bar to reveal the rest of the picture and double click the eraser tool to change the rest of it. Keep changing colors until you have black lines on a white background (unless, of course, you have a color printer, in which case you would have been printing already).

At this point you can print the graph directly from Paintbrush, or save it in PCX or BMP format for later importation to any program that recognizes these formats. Or, you can copy the graph back to the clipboard for pasting into other Windows applications (like a word processor) by first choosing Zoom Out from the View menu, selecting the pick tool (the scissors with a rectangle attached), placing the cross-hair cursor on the upper left corner of the picture, pressing the left mouse button, dragging the cursor to the lower right corner to outline the entire graph, releasing the left mouse button, and then choosing Copy from the Edit menu.

Appendix A

An Inventory of Continuous Distributions

A.1 Introduction

Descriptions of the models are given below. The distributions are presented in the following way. First the name is given along with the parameters. Next the density function ($f(x)$) and distribution function ($F(x)$) are given. The incomplete gamma function is given by

$$\Gamma(\alpha; x) = \frac{1}{\Gamma(\alpha)} \int_0^x t^{\alpha-1} e^{-t} dt, \quad \alpha > 0, \quad x > 0$$

while the incomplete beta function is

$$\beta(a, b; x) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} \int_0^x t^{a-1} (1-t)^{b-1} dt, \quad a > 0, \quad b > 0, \quad 0 < x < 1.$$

Within each family the distributions are presented in decreasing order with regard to the number of parameters. The Greek letters used are selected to be consistent. Any Greek letter that is not used in the distribution means that that distribution is a special case of one with more parameters, but with the missing parameters set equal to 1. Unless specifically indicated, all parameters must be positive.

For several of the distributions, starting values are suggested. They are not necessarily good estimators, just places from which to start an iterative procedure to maximize the likelihood or other objective function. These are found by either the methods of moments or percentile matching. The quantities used are:

$$\text{Moments: } m = \frac{1}{n} \sum_{i=1}^n x_i \quad t = \frac{1}{n} \sum_{i=1}^n x_i^2$$

Percentile matching: $p =$ 25th percentile, $q =$ 75th percentile

For grouped data or data that have been truncated or censored, these quantities may have to be approximated. Because the purpose is to obtain starting values and not a useful estimate, it is often sufficient to just ignore modifications.. For three and four parameter distributions starting values can be obtained by using estimates from a special case, then making the new parameters equal to 1. An all-purpose starting value rule (for when all else fails) is to set the scale parameter (θ) equal to the mean and all other parameters equal to 2.

A.2 Transformed beta family

A.2.1 Four parameter distribution

A.2.1.1 Transformed beta - $\alpha, \theta, \gamma, \tau$

$$f(x) = \frac{\Gamma(\alpha + \tau)}{\Gamma(\alpha)\Gamma(\tau)} \frac{\gamma(x/\theta)^{\gamma\tau}}{x[1 + (x/\theta)^\gamma]^{\alpha+\tau}}$$

$$F(x) = \beta(\tau, \alpha; u), \quad u = \frac{(x/\theta)^\gamma}{1 + (x/\theta)^\gamma}$$

A.2.2 Three parameter distributions

A.2.2.1 Generalized Pareto - α, θ, τ

$$f(x) = \frac{\Gamma(\alpha + \tau)}{\Gamma(\alpha)\Gamma(\tau)} \frac{\theta^\alpha x^{\tau-1}}{(x + \theta)^{\alpha+\tau}}$$

$$F(x) = \beta(\tau, \alpha; u), \quad u = \frac{x}{x + \theta}$$

A.2.2.2 Burr - α, θ, γ

$$f(x) = \frac{\alpha\gamma(x/\theta)^\gamma}{x[1 + (x/\theta)^\gamma]^{\alpha+1}}$$

$$F(x) = 1 - u^\alpha, \quad u = \frac{1}{1 + (x/\theta)^\gamma}$$

A.2.2.3 Inverse Burr - τ, θ, γ

$$f(x) = \frac{\tau\gamma(x/\theta)^{\gamma\tau}}{x[1 + (x/\theta)^\gamma]^{\tau+1}}$$

$$F(x) = u^\tau, \quad u = \frac{(x/\theta)^\gamma}{1 + (x/\theta)^\gamma}$$

A.2.3 Two parameter distributions**A.2.3.1 Pareto - α, θ**

$$f(x) = \frac{\alpha\theta^\alpha}{(x+\theta)^{\alpha+1}}$$

$$F(x) = 1 - \left(\frac{\theta}{x+\theta}\right)^\alpha$$

$$\hat{\alpha} = 2\frac{t - m^2}{t - 2m^2}, \quad \hat{\theta} = \frac{mt}{t - 2m^2}$$

A.2.3.2 Inverse Pareto - τ, θ

$$f(x) = \frac{\tau\theta x^{\tau-1}}{(x+\theta)^{\tau+1}}$$

$$F(x) = \left(\frac{x}{x+\theta}\right)^\tau$$

A.2.3.3 Loglogistic - γ, θ

$$f(x) = \frac{\gamma(x/\theta)^\gamma}{x[1 + (x/\theta)^\gamma]^2}$$

$$F(x) = u, \quad u = \frac{(x/\theta)^\gamma}{1 + (x/\theta)^\gamma}$$

$$\hat{\gamma} = \frac{2 \log(3)}{\log(q) - \log(p)}, \quad \hat{\theta} = \exp\left(\frac{\log(q) + \log(p)}{2}\right)$$

A.2.3.4 Paralogistic - α, θ

This is a Burr distribution with $\gamma = \alpha$.

$$f(x) = \frac{\alpha^2(x/\theta)^\alpha}{x[1 + (x/\theta)^\alpha]^{\alpha+1}}$$

$$F(x) = 1 - u^\alpha, \quad u = \frac{1}{1 + (x/\theta)^\alpha}$$

Starting values can use estimates from the loglogistic (use γ for α) or Pareto (use α) distributions.

A.2.3.5 Inverse paralogistic - τ, θ

This is an inverse Burr distribution with $\gamma = \tau$.

$$f(x) = \frac{\tau^2(x/\theta)^{\tau^2}}{x[1 + (x/\theta)^\tau]^{\tau+1}}$$

$$F(x) = u^\tau, \quad u = \frac{(x/\theta)^\tau}{1 + (x/\theta)^\tau}$$

Starting values can use estimates from the loglogistic (use γ for τ) or inverse Pareto (use τ) distributions.

A.3 Transformed gamma family**A.3.1 Three parameter distributions****A.3.1.1 Transformed gamma - α, θ, τ**

$$f(x) = \frac{\tau u^\alpha e^{-u}}{x\Gamma(\alpha)}, \quad u = (x/\theta)^\tau$$

$$F(x) = \Gamma(\alpha; u)$$

A.3.1.2 Inverse transformed gamma - α, θ, τ

$$f(x) = \frac{\tau u^\alpha e^{-u}}{x\Gamma(\alpha)}, \quad u = (\theta/x)^\tau$$

$$F(x) = 1 - \Gamma(\alpha; u)$$

A.3.2 Two parameter distributions**A.3.2.1 Gamma - α, θ**

$$f(x) = \frac{(x/\theta)^\alpha e^{-x/\theta}}{x\Gamma(\alpha)}$$

$$F(x) = \Gamma(\alpha; x/\theta)$$

$$\hat{\alpha} = \frac{m^2}{t - m^2}, \quad \hat{\theta} = \frac{t - m^2}{m}$$

A.3.2.2 Inverse gamma - α, θ

$$f(x) = \frac{(\theta/x)^\alpha e^{-\theta/x}}{x\Gamma(\alpha)}$$

$$F(x) = 1 - \Gamma(\alpha; \theta/x)$$

$$\hat{\alpha} = \frac{2t - m^2}{t - m^2}, \quad \hat{\theta} = \frac{mt}{t - m^2}$$

A.3.2.3 Weibull - θ, τ

$$f(x) = \frac{\tau(x/\theta)^\tau e^{-(x/\theta)^\tau}}{x}$$

$$F(x) = 1 - e^{-(x/\theta)^\tau}$$

$$\hat{\theta} = \exp\left(\frac{g \log(p) - \log(q)}{g - 1}\right), \quad g = \frac{\log(\log(4))}{\log(\log(4/3))}, \quad \hat{\tau} = \frac{\log(\log(4))}{\log(q) - \log(\hat{\theta})}$$

A.3.2.4 Inverse Weibull - θ, τ

$$f(x) = \frac{\tau(\theta/x)^\tau e^{-(\theta/x)^\tau}}{x}$$

$$F(x) = e^{-(\theta/x)^\tau}$$

$$\hat{\theta} = \exp\left(\frac{g \log(q) - \log(p)}{g - 1}\right), \quad g = \frac{\log(\log(4))}{\log(\log(4/3))}, \quad \hat{\tau} = \frac{\log(\log(4))}{\log(\hat{\theta}) - \log(p)}$$

A.3.3 One parameter distributions**A.3.3.1 Exponential - θ**

$$f(x) = \frac{e^{-x/\theta}}{\theta}$$

$$F(x) = 1 - e^{-x/\theta}$$

$$\hat{\theta} = m$$

A.3.3.2 Inverse exponential - θ

$$f(x) = \frac{\theta e^{-\theta/x}}{x^2}$$

$$F(x) = e^{-\theta/x}$$

$$\hat{\theta} = -q \log(3/4)$$

A.4 Other distributions**A.4.1.1 Lognormal - μ, σ (μ can be negative)**

$$f(x) = \frac{1}{x\sigma\sqrt{2\pi}} \exp(-z^2/2) = \phi(z)/(\sigma x), \quad z = \frac{\log x - \mu}{\sigma}$$

$$F(x) = \Phi(z)$$

$$\hat{\sigma} = \sqrt{\log(t) - 2 \log(m)}, \quad \hat{\mu} = \log(m) - \hat{\sigma}^2/2$$

A.4.1.2 Inverse Gaussian - μ, θ

$$f(x) = \left(\frac{\theta}{2\pi x^3} \right)^{1/2} \exp \left\{ -\frac{\theta z^2}{2x} \right\}, \quad z = \frac{x - \mu}{\mu}$$

$$F(x) = \Phi \left[z \left(\frac{\theta}{x} \right)^{1/2} \right] + \exp(2\theta/\mu) \Phi \left[-y \left(\frac{\theta}{x} \right)^{1/2} \right], \quad y = \frac{x + \mu}{\mu}$$

$$\hat{\mu} = m, \quad \hat{\theta} = \frac{m^3}{t - m^2}$$

A.4.1.3 Single parameter Pareto - α, θ

$$f(x) = \frac{\alpha \theta^\alpha}{x^{\alpha+1}}, \quad x > \theta$$

$$F(x) = 1 - (\theta/x)^\alpha, \quad x > \theta$$

$$\hat{\alpha} = \frac{m}{m - \theta}$$

Note - Although there appears to be two parameters, only α is a true parameter. The value of θ must be set in advance.

Appendix B

An Inventory of Discrete Distributions

B.1 Introduction

The sixteen models fall into three classes. The divisions are based on the algorithm by which the probabilities are computed. For some of the more familiar distributions these formulas will look different from the ones you may have learned, but they produce the same probabilities. After each name, the parameters are given. All parameters are positive unless otherwise indicated. In all cases p_k is the probability of observing k losses.

The estimators which are presented are not intended to be useful estimators but rather for providing starting values for maximizing the likelihood (or other) function. For determining starting values, the following quantities are used (where n_k is the observed frequency at k (if for the last entry n_k represents the number of observations at k or more, assume it was at exactly k) and n is the sample size):

$$\hat{\mu} = \frac{1}{n} \sum_{k=1}^{\infty} kn_k, \quad \hat{\sigma}^2 = \frac{1}{n} \sum_{k=1}^{\infty} k^2 n_k - \hat{\mu}^2$$

When the method of moments is used to determine the starting value, a circumflex (e.g., $\hat{\lambda}$) is used. For any other method, a tilde (e.g., $\tilde{\lambda}$) is used. When the starting value formulas do not provide admissible parameter values, a truly crude guess is to set the product of all λ and β parameters equal to the sample mean and all other parameters equal to 1. If there are two λ and/or β parameters, an easy choice is to set each to the square root of the sample mean.

B.2 The $(a, b, 0)$ class

The distributions in this class have support on $0, 1, \dots$. For this class, a particular distribution is specified by setting p_0 and then using $p_k = (a + b/k)p_{k-1}$. Specific members are created

by setting p_0 , a , and b .

B.2.1.1 Poisson - λ

$$\begin{aligned} p_0 &= e^{-\lambda}, \quad a = 0, \quad b = \lambda \\ p_k &= \frac{e^{-\lambda} \lambda^k}{k!} \\ \hat{\lambda} &= \hat{\mu} \end{aligned}$$

B.2.1.2 Geometric - β

$$\begin{aligned} p_0 &= 1/(1 + \beta), \quad a = \beta/(1 + \beta), \quad b = 0 \\ p_k &= \frac{\beta^k}{(1 + \beta)^{k+1}} \\ \hat{\beta} &= \hat{\mu} \end{aligned}$$

This is a special case of the negative binomial with $r = 1$.

B.2.1.3 Binomial - q, m , ($0 < q < 1$, m an integer)

$$\begin{aligned} p_0 &= (1 - q)^m, \quad a = -q/(1 - q), \quad b = (m + 1)q/(1 - q) \\ p_k &= \binom{m}{k} q^k (1 - q)^{m-k}, \quad k = 0, 1, \dots, m \\ \hat{q} &= \hat{\mu}/m \end{aligned}$$

B.2.1.4 Negative binomial - β, r

$$\begin{aligned} p_0 &= (1 + \beta)^{-r}, \quad a = \beta/(1 + \beta), \quad b = (r - 1)\beta/(1 + \beta) \\ p_k &= \frac{r(r + 1) \cdots (r + k - 1)\beta^k}{k!(1 + \beta)^{r+k}} \\ \hat{\beta} &= \frac{\hat{\sigma}^2}{\hat{\mu}} - 1, \quad \hat{r} = \frac{\hat{\mu}^2}{\hat{\sigma}^2 - \hat{\mu}} \end{aligned}$$

B.3 The $(a, b, 1)$ class

To distinguish this class from the $(a, b, 0)$ class, the probabilities are denoted $\Pr(N = k) = p_k^M$ or $\Pr(N = k) = p_k^T$ depending on which sub-class is being represented. For this class p_0^M is arbitrary (that is, it is a parameter) and then p_1^M or p_1^T is a specified function of the

parameters a and b . Subsequent probabilities are obtained recursively as in the $(a, b, 0)$ class: $p_k^M = (a + b/k)p_{k-1}^M$, $k = 2, 3, \dots$ with the same recursion for p_k^T . There are two sub-classes of this class. When discussing their members, we often refer to the “corresponding” member of the $(a, b, 0)$ class. This refers to the member of that class with the same values for a and b . The notation p_k will continue to be used for probabilities for the corresponding $(a, b, 0)$ distribution.

B.3.1 The zero-truncated sub-class

The members of this class have $p_0^T = 0$ and therefore it need not be estimated. These distributions should only be used when a value of zero is impossible. For those members of the sub-class which have corresponding $(a, b, 0)$ distributions, $p_k^T = p_k/(1 - p_0)$.

B.3.1.1 Zero-truncated Poisson - λ

$$\begin{aligned} p_1^T &= \lambda/(e^\lambda - 1), \quad a = 0, \quad b = \lambda \\ p_k^T &= \frac{\lambda^k}{k!(e^\lambda - 1)} \\ \tilde{\lambda} &= \log(n\hat{\mu}/n_1) \end{aligned}$$

B.3.1.2 Zero-truncated geometric - β

$$\begin{aligned} p_1^T &= 1/(1 + \beta), \quad a = \beta/(1 + \beta), \quad b = 0 \\ p_k^T &= \frac{\beta^{k-1}}{(1 + \beta)^k} \\ \hat{\beta} &= \hat{\mu} - 1 \end{aligned}$$

This is a special case of the zero-truncated negative binomial with $r = 1$.

B.3.1.3 Logarithmic - β

$$\begin{aligned} p_1^T &= \beta/[(1 + \beta) \log(1 + \beta)], \quad a = \beta/(1 + \beta), \quad b = -\beta/(1 + \beta) \\ p_k^T &= \frac{\beta^k}{k(1 + \beta)^k \log(1 + \beta)} \\ \tilde{\beta} &= \frac{n\hat{\mu}}{n_1} - 1 \text{ or } \frac{2(\hat{\mu} - 1)}{\hat{\mu}} \end{aligned}$$

This is a limiting case of the zero-truncated negative binomial as $r \rightarrow 0$.

B.3.1.4 Zero-truncated binomial - $q, m, (0 < q < 1, m \text{ an integer})$

$$\begin{aligned}
p_1^T &= m(1-q)^{m-1}q/[1-(1-q)^m], \\
a &= -q/(1-q), \quad b = (m+1)q/(1-q) \\
p_k^T &= \frac{\binom{m}{k}q^k(1-q)^{m-k}}{1-(1-q)^m}, \quad k = 1, 2, \dots, m \\
\tilde{q} &= \hat{\mu}/m
\end{aligned}$$

B.3.1.5 Zero-truncated negative binomial - $\beta, r, (r > -1)$

$$\begin{aligned}
p_1^T &= r\beta/[(1+\beta)^{r+1} - (1+\beta)], \\
a &= \beta/(1+\beta), \quad b = (r-1)\beta/(1+\beta) \\
p_k^T &= \frac{r(r+1)\cdots(r+k-1)}{k![(1+\beta)^r - 1]} \left(\frac{\beta}{1+\beta}\right)^k \\
\tilde{\beta} &= \frac{\hat{\sigma}^2}{\hat{\mu}} - 1, \quad \tilde{r} = \frac{\hat{\mu}^2}{\hat{\sigma}^2 - \hat{\mu}}
\end{aligned}$$

This distribution is sometimes called the extended truncated negative binomial distribution because the parameter r can extend below 0.

B.3.2 The zero-modified sub-class

A zero-modified distribution is created by starting with a truncated distribution and then placing an arbitrary amount of probability at zero. This probability, p_0^M , is a parameter. The remaining probabilities are adjusted accordingly. Values of p_k^M can be determined from the corresponding zero-truncated distribution as $p_k^M = (1-p_0^M)p_k^T$ or from the corresponding $(a, b, 0)$ distribution as $p_k^M = (1-p_0^M)p_k/(1-p_0)$. The same recursion used for the zero-truncated sub-class applies.

The maximum likelihood estimator of p_0^M is always the sample relative frequency at 0.

B.4 Compound class

Members of this class are obtained by compounding one distribution with another. That is, let N be a discrete distribution, called the primary distribution and let M_1, M_2, \dots be identically and independently distributed with another discrete distribution called the secondary distribution. The compound distribution is $S = M_1 + \dots + M_N$. The probabilities for the compound distributions are found from

$$p_k = \sum_{y=1}^k (a + by/k) f_y p_{k-y} / (1 - af_0)$$

for $n = 1, 2, \dots$ where a and b are the usual values for the primary distribution (which must be a member of the $(a, b, 0)$ class) and f_y is p_y for the secondary distribution. The only two primary distributions used here are Poisson (for which $p_0 = \exp[-\lambda(1 - f_0)]$) and geometric (for which $p_0 = 1/[1 + \beta - \beta f_0]$). As this information completely describes these distributions, only the names and starting values are given below.

In the following list the primary distribution is always named first. For the first, second, and fourth distributions, the secondary distribution is the $(a, b, 0)$ class member with that name. For the third and the last three distributions (the Poisson-ETNB and its two special cases) the secondary distribution is the zero-truncated version.

B.4.1 Some compound distributions

B.4.1.1 Poisson-binomial - λ, q, m ($0 < q < 1$, m an integer)

$$\hat{q} = \frac{\hat{\sigma}^2/\hat{\mu} - 1}{m - 1}, \quad \hat{\lambda} = \hat{\mu}/(m\hat{q}) \text{ or } \tilde{q} = 0.5, \quad \tilde{\lambda} = 2\hat{\mu}/m$$

B.4.1.2 Poisson-Poisson - λ_1, λ_2

The parameter λ_1 is for the primary Poisson distribution and λ_2 is for the secondary Poisson distribution. This distribution is also called the Neyman Type A.

$$\tilde{\lambda}_1 = \tilde{\lambda}_2 = \sqrt{\hat{\mu}}$$

B.4.1.3 Geometric-extended truncated negative binomial- β_1, β_2, r ($r > -1$)

The parameter β_1 is for the primary geometric distribution. The last two parameters are for the secondary distribution. The truncated version is used so that the extension of r is available.

$$\tilde{\beta}_1 = \tilde{\beta}_2 = \sqrt{\hat{\mu}}$$

B.4.1.4 Geometric-Poisson - β, λ

$$\tilde{\beta} = \tilde{\lambda} = \sqrt{\hat{\mu}}$$

B.4.1.5 Poisson-extended truncated negative binomial- λ, β, r

$$\tilde{r} = \frac{\hat{\mu}(K - 3\hat{\sigma}^2 + 2\hat{\mu}) - 2(\hat{\sigma}^2 - \hat{\mu})^2}{\hat{\mu}(K - 3\hat{\sigma}^2 + 2\hat{\mu}) - (\hat{\sigma}^2 - \hat{\mu})^2}, \quad \tilde{\beta} = \frac{\hat{\sigma}^2 - \hat{\mu}}{\hat{\mu}(1 + \hat{r})}, \quad \tilde{\lambda} = \frac{\hat{\mu}}{\hat{r}\tilde{\beta}}, \text{ or}$$

$$\tilde{r} = \frac{\hat{\sigma}^2 n_1/n - \hat{\mu}^2 n_0/n}{(\hat{\sigma}^2 - \hat{\mu}^2)(n_0/n) \log(n_0/n) - \hat{\mu}(\hat{\mu} n_0/n - n_1/n)},$$

$$\tilde{\beta} = \frac{\hat{\sigma}^2 - \hat{\mu}}{\hat{\mu}(1 + \hat{r})}, \quad \tilde{\lambda} = \frac{\hat{\mu}}{\hat{r}\hat{\beta}}$$

$$\text{where } K = \frac{1}{n} \sum_{k=0}^{\infty} k^3 n_k - 3\hat{\mu} \frac{1}{n} \sum_{k=0}^{\infty} k^2 n_k + 2\hat{\mu}^3$$

This distribution is also called the generalized Poisson-Pascal.

B.4.1.6 Polya-Aeppli - λ, β

$$\hat{\beta} = \frac{\hat{\sigma}^2 - \hat{\mu}}{2\hat{\mu}}, \quad \hat{\lambda} = \frac{\hat{\mu}}{1 + \hat{\beta}}$$

This is a special case of the Poisson-extended truncated negative binomial with $r = 1$. It is actually a Poisson-truncated geometric.

B.4.1.7 Poisson-inverse Gaussian - λ, β

$$\tilde{\lambda} = -\log(n_0/n), \quad \tilde{\beta} = 4(\hat{\mu} - \hat{\lambda})/\hat{\mu}$$

This is a special case of the Poisson-extended truncated negative binomial with $r = -.5$.

B.5 A hierarchy of discrete distributions

The following table indicates which distributions are special or limiting cases of others. For the special cases, one parameter is set equal to a constant to create the special case. For the limiting cases, two parameters go to infinity or zero in some special way.

Distribution	is a special case of	is a limiting case of
Poisson	ZM Poisson	Neg. bin., P-bin. P-inv. Gaussian Polya-Aeppli Neyman A
ZT Poisson	ZM Poisson	ZT Neg. bin.
ZM Poisson		ZM Neg. bin
Geometric	Neg. bin. ZM geom	Geometric-Poisson
ZT Geometric	ZT Neg. bin.	
ZM Geometric	ZM Neg. bin.	
Logarithmic		ZT Neg. bin.
ZM Logarithmic		ZM Neg. bin.
Binomial	ZM binomial 0	
Neg. bin.	ZM Neg. bin.	P-ETNB
P-inv. Gaussian	P-ETNB	
Polya-Aeppli	P-ETNB	
Neyman A		P-ETNB