COLLABORATORS			
	TITLE :		*REFERENCE* :
ACTION	*NAME*	*DATE*	*SIGNATURE*
WRITTEN BY		August 15, 2008	

REVISION HISTORY			
NUMBER	DATE	DESCRIPTION	NAME

Contents

List of Figures

1 Simulation Analysis

1.1 Introduction to Simulation Analysis

A simulation is the imitation of the operation of a real-world process or system. The behavior of a system is studied by generating an artificial history of the system through the use of random numbers. These numbers are used in the context of a simulation model, which is the mathematical, logical and symbolic representation of the relationships between the objects of interest of the system. After the model has been validated, the effects of changes in the environment on the system, or the effects of changes in the system on system performance can be predicted using the simulation model. [1]

Gnumeric includes a facility for performing Monte Carlo Simulation. Monte Carlo simulation involves the sampling of random numbers to solve a problem where the passage of time plays no substantive role. [2] In other words, each sample is not effected by prior samples. This is in contrast to discrete event simulation or continuous simulation where the results from earlier in the simulation can effect successive samples within a simulation experiment. The Monte Carlo simulation will be enabled through the use of the Random Number functions (Section A.14) and the results presented along with variances for use in statistical analysis.

1.2 Setting up Simulation Model

The remainder of this chapter will illustrate use of the simulation tool using an example from Banks et. al. A classic inventory problem is the newspaper seller's problem. A newspaper seller buys papers for 33 cents each and sells for 50 cents. Newspapers not sold are sold as scrap (recycled) for 5 cents. Newspapers are purchased by the paper seller in bundles of 10. Demand for newspapers can be categorized as “good,” “fair,” or “poor” with probability 0.35, 0.45 and 0.20 respectively, with each day's demand being independent of prior days. The problem for the paper seller is to determine the optimal number of papers to purchase when the day's demand is not yet known.

The daily profit equation is:

Profit = [(Sale revenue) - (Cost) - (Scrap value)]

EQUATION 1.1: Newsvendor Profit

To set up the model, this example will use two tabs in Gnumeric, a tab labeled 'Profit' to calculate profit, and a tab labeled 'Demand Tables' to store the various tables needed to calculate the demand for any given sampling.

[1] Adapted from Banks, Carson, Nelson and Nicol (2001), Discrete-Event System Simulation, 3rd ed.
[2] Definition from Law and Kelton (1991), Simulation Modeling & Analysis, 2nd ed, pp. 113.

For the Profit tab, set up the profit tab as in Figure 1.

At the top of the Profit' tab, the Profit table will be entered . There are three variables: Sale revenue, Cost and Scrap value, and they take the per unit coefficients of 0.5, 0.33 and 0.05 respectively. Enter the coefficients in cells B13 through D13. In cells B12 through D12, enter the equations for sale revenue, cost and Scrap value. In cell E12, enter the equation for Profit

Next, we add the values for the decision variable, which is the amount purchased, and the amount sold. For now, the amount purchased will be a fixed value, in section 1.1.5 Using SIMTABLE this will be changed to allow the testing of numerous potential values. For now, set the cell C16 to 60.

1. B13 = 0.5
2. C13 = 0.33
3. D13 = 0.05
4. B12 = B13*min(B16,B20)
5. C12 = C13*B16
6. D12 = D13*max(0,B16-B20)
7. E12 = B12-C12+D12
8. B16 = 50

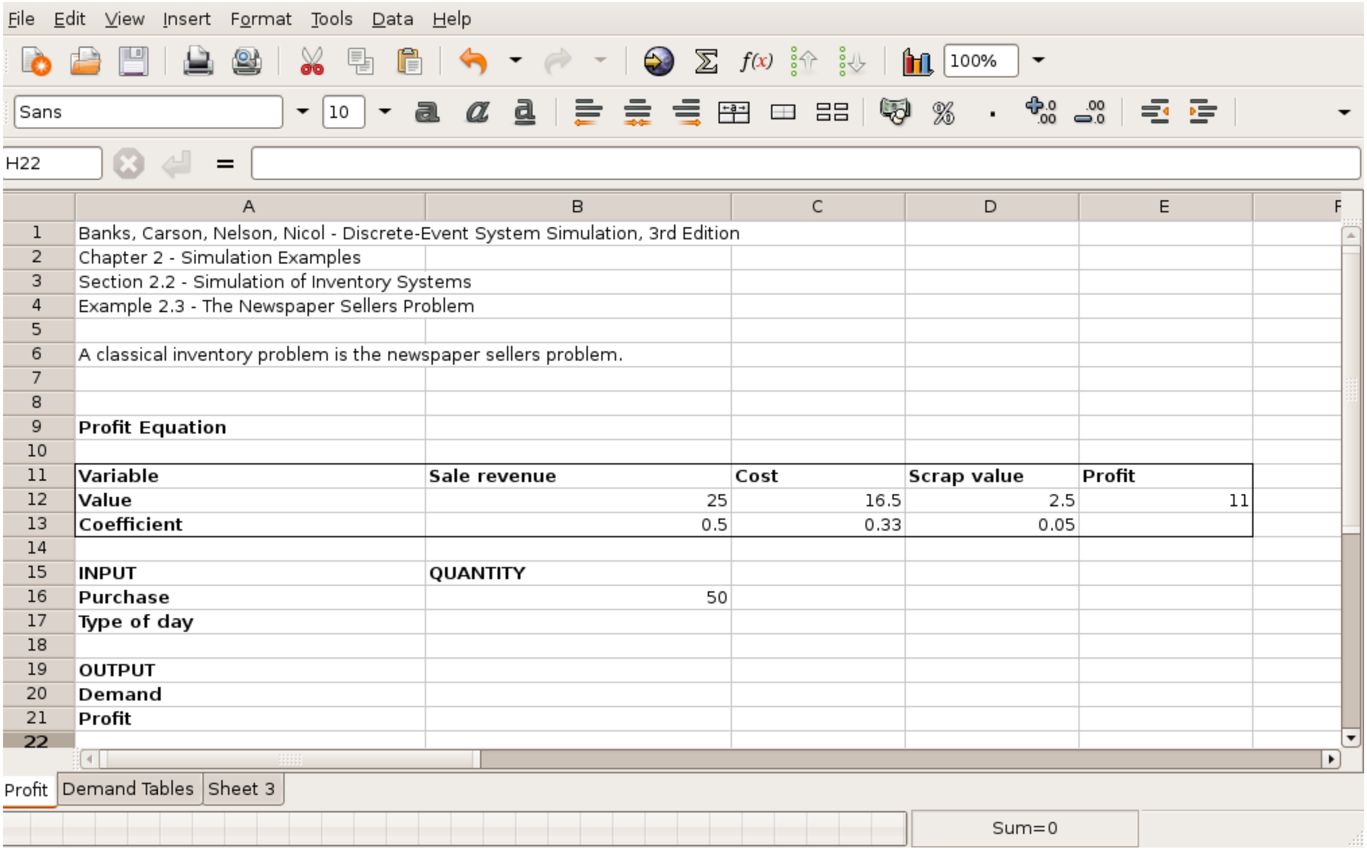

	A	B	C	D	E
1	Banks, Carson, Nelson, Nicol - Discrete-Event System Simulation, 3rd Edition				
2	Chapter 2 - Simulation Examples				
3	Section 2.2 - Simulation of Inventory Systems				
4	Example 2.3 - The Newspaper Sellers Problem				
5					
6	A classical inventory problem is the newspaper sellers problem.				
7					
8					
9	**Profit Equation**				
10					
11	**Variable**	**Sale revenue**	**Cost**	**Scrap value**	**Profit**
12	**Value**	25	16.5	2.5	11
13	**Coefficient**	0.5	0.33	0.05	
14					
15	**INPUT**	**QUANTITY**			
16	**Purchase**	50			
17	**Type of day**				
18					
19	**OUTPUT**				
20	**Demand**				
21	**Profit**				
22					

Figure 1: Profit table for newsvendor example

Next, create the demand tables from which the demand will be generated. In the tab 'Demand Tables' enter the values of the probability in cells B4 through B6. In cells C4, C5 and C6 enter the cumulative probability values (C4: 0.35; C5: 0.8; C6: 1) as shown in Figure 2

	A	B	C	D	E
1	Banks, Carson, Nelson, Nicol - Discrete-Event System Simulation, 3rd Edition				
2	Chapter 2 - Simulation Examples				
3	Section 2.2 - Simulation of Inventory Systems				
4	Example 2.3 - The Newspaper Sellers Problem				
5					
6	A classical inventory problem is the newspaper sellers problem.				
7					
8					
9	**Profit Equation**				
10					
11	**Variable**	**Sale revenue**	**Cost**	**Scrap value**	**Profit**
12	**Value**	25	16.5	2.5	11
13	**Coefficient**	0.5	0.33	0.05	
14					
15	**INPUT**	**QUANTITY**			
16	**Purchase**	50			
17	**Type of day**				
18					
19	**OUTPUT**				
20	**Demand**				
21	**Profit**				

Figure 2: Type of newsday for newspaper demand

The next table is the daily demand for newspapers based on the type of news day. In Figure 3, the table Distribution of Newspapers Demanded is in cells A11 through D18 and contains the raw demand distribution values. The cumulative distribution tables in cells A21 through G29 are derived values from the Distribution of Newspapers Demanded using values in the top Distribution of Newspapers demanded table.

	A	B	C	D	E	F	G
9	**Distribution of Newspapers Demanded**						
10							
11	**Demand**	**Good**	**Fair**	**Poor**			
12	40	0.03	0.1	0.44			
13	50	0.05	0.18	0.22			
14	60	0.15	0.4	0.16			
15	70	0.2	0.2	0.12			
16	80	0.35	0.08	0.06			
17	90	0.15	0.04	0			
18	100	0.07	0	0			
19	Total	1	1	1			
20							
21		**Cumulative Distribution**			**Values**		
22	**Demand**	**Good**	**Fair**	**Poor**	**Good**	**Fair**	**Poor**
23	**40**	0.03	0.1	0.44	0	0	0
24	**50**	0.08	0.28	0.66	0.03	0.1	0.44
25	**60**	0.23	0.68	0.82	0.08	0.28	0.66
26	**70**	0.43	0.88	0.94	0.23	0.68	0.82
27	**80**	0.78	0.96	1	0.43	0.88	0.94
28	**90**	0.93	1		0.78	0.96	1
29	**100**	1			0.93	1	0

Figure 3: Distribution of newspapers demanded tables

Finally, back in the Profit tab, the demand data will be filled in through the use of references to the Demand Tables tab as shown in Figure 4. Fill in the values that are in bold by typing in the labels.

In the following cells, enter the equations below:

1. B17: =rand()
2. C17: =if(B17<'Demand Tables'!C4,"Good",if(C19<'Demand Tables'!C5,"Fair","Poor"))
3. B18: =rand()
4. B20: =lookup($C17,$B23:$D23,$B24:$D24)
5. B21: =E12
6. B24: =lookup(Profit!$B18,'Demand Tables'!E23:E29,'Demand Tables'!$A23:$A29)
7. C24: =lookup(Profit!$B18,'Demand Tables'!F23:F29,'Demand Tables'!$A23:$A29)
8. D24: =lookup(Profit!$B18,'Demand Tables'!G23:G29,'Demand Tables'!$A23:$A29)

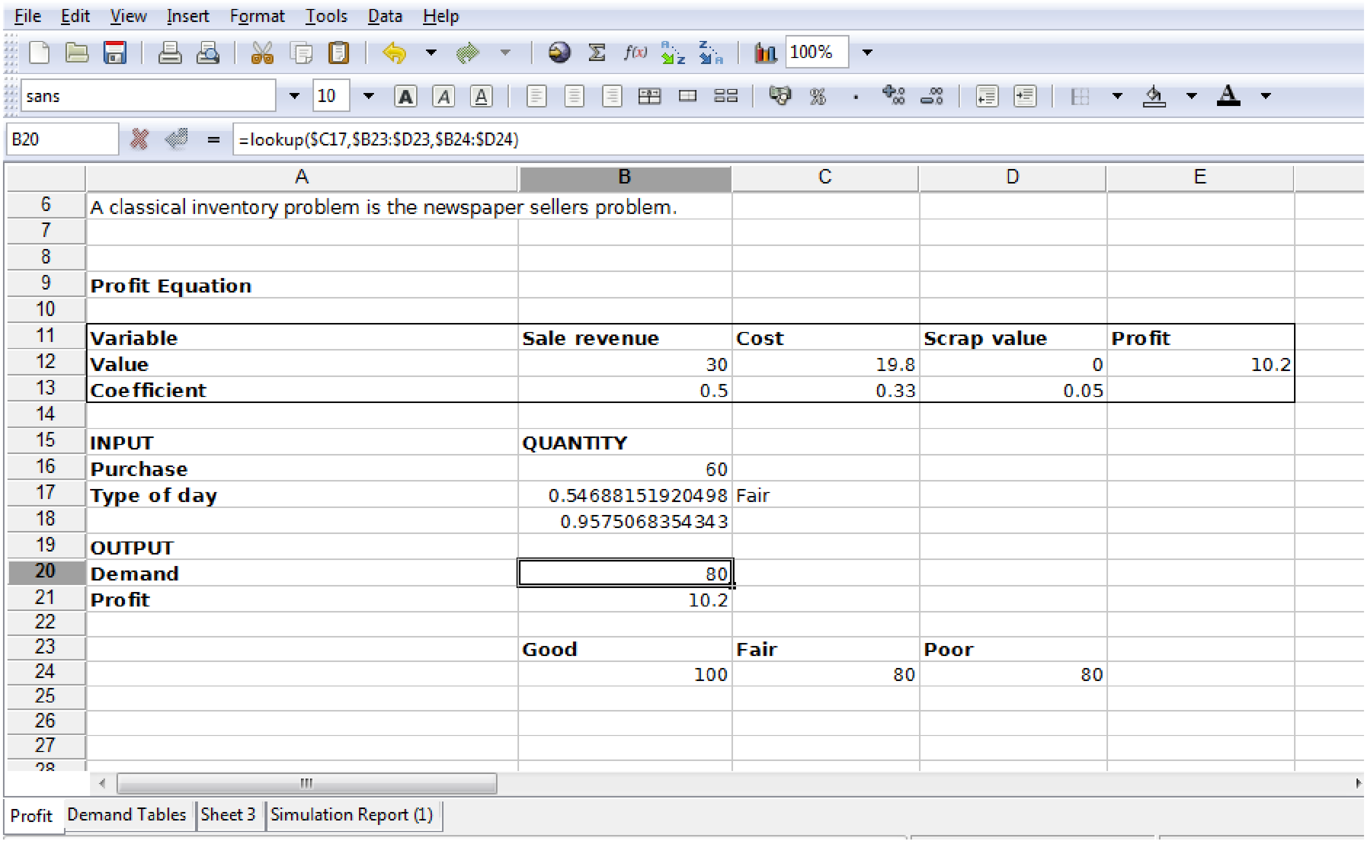

	A	B	C	D	E
6	A classical inventory problem is the newspaper sellers problem.				
7					
8					
9	**Profit Equation**				
10					
11	**Variable**	**Sale revenue**	**Cost**	**Scrap value**	**Profit**
12	**Value**	30	19.8	0	10.2
13	**Coefficient**	0.5	0.33	0.05	
14					
15	**INPUT**	**QUANTITY**			
16	**Purchase**	60			
17	**Type of day**	0.54688151920498	Fair		
18		0.9575068354343			
19	**OUTPUT**				
20	**Demand**	80			
21	**Profit**	10.2			
22					
23		**Good**	**Fair**	**Poor**	
24		100	80	80	
25					
26					
27					

Figure 4: Profit table for newsvendor example

When done, the Profit spreadsheet will be setup with a profit equation, decision variables, and random events as shown in Figure 4. The rand() functions in cells C17 and C18 return a random value between 0 and 1, which are used by the lookup() functions in cells B20, B24, C24 and D24 to calculate a randomly determined daily demand. Next, this sheet will be used for analysis through the use of simulation.

1.3 Running the Simulation

To run the simulation, from the Gnumeric toolbar, select Tools → Simulation. In the Risk Simulation window that appears, the first tab is the Variables tab. There are two entries in the Variables tab: Input variables and Output variables (Figure 5).

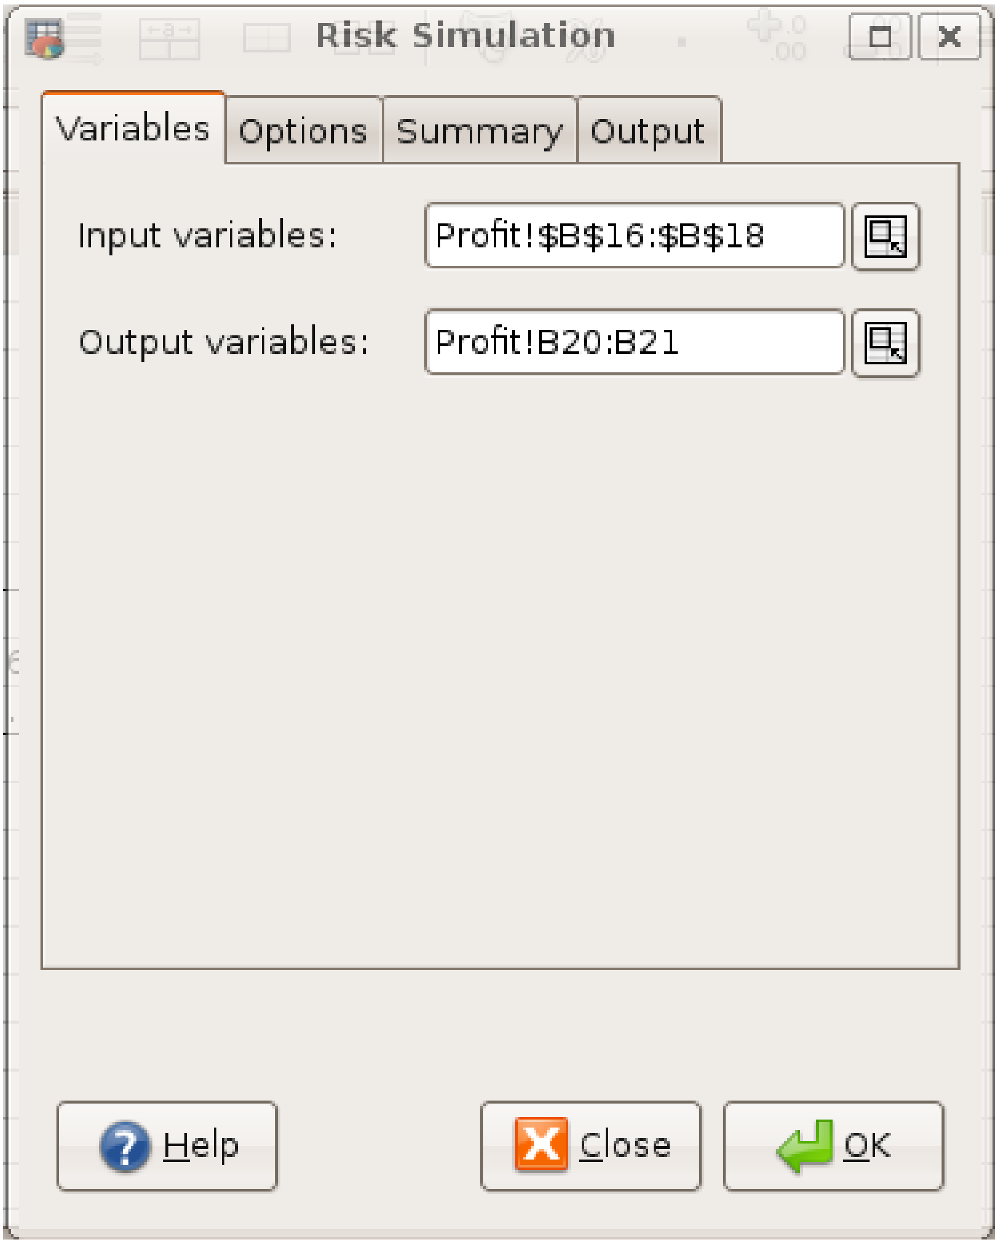

Figure 5: Profit table for newsvendor example

Input variables are the cells which hold the functions based on random numbers of the type described in Section A.14. In this case, they are the cells B17 and B18 in the Profit worksheet, which hold the rand() function. Later, when the quantity purchased is a parameter set by the `SIMTABLE` function, cell B16 which holds the purchase quantity will be added to the range of input variables.

Output variables are the results of interest, or the dependent variable. In this case, the dependent variables are the demand and the profit, which are in cells B20 and B21.

The next tab is the Options tab . There are four settings in the options as shown in Figure 6.

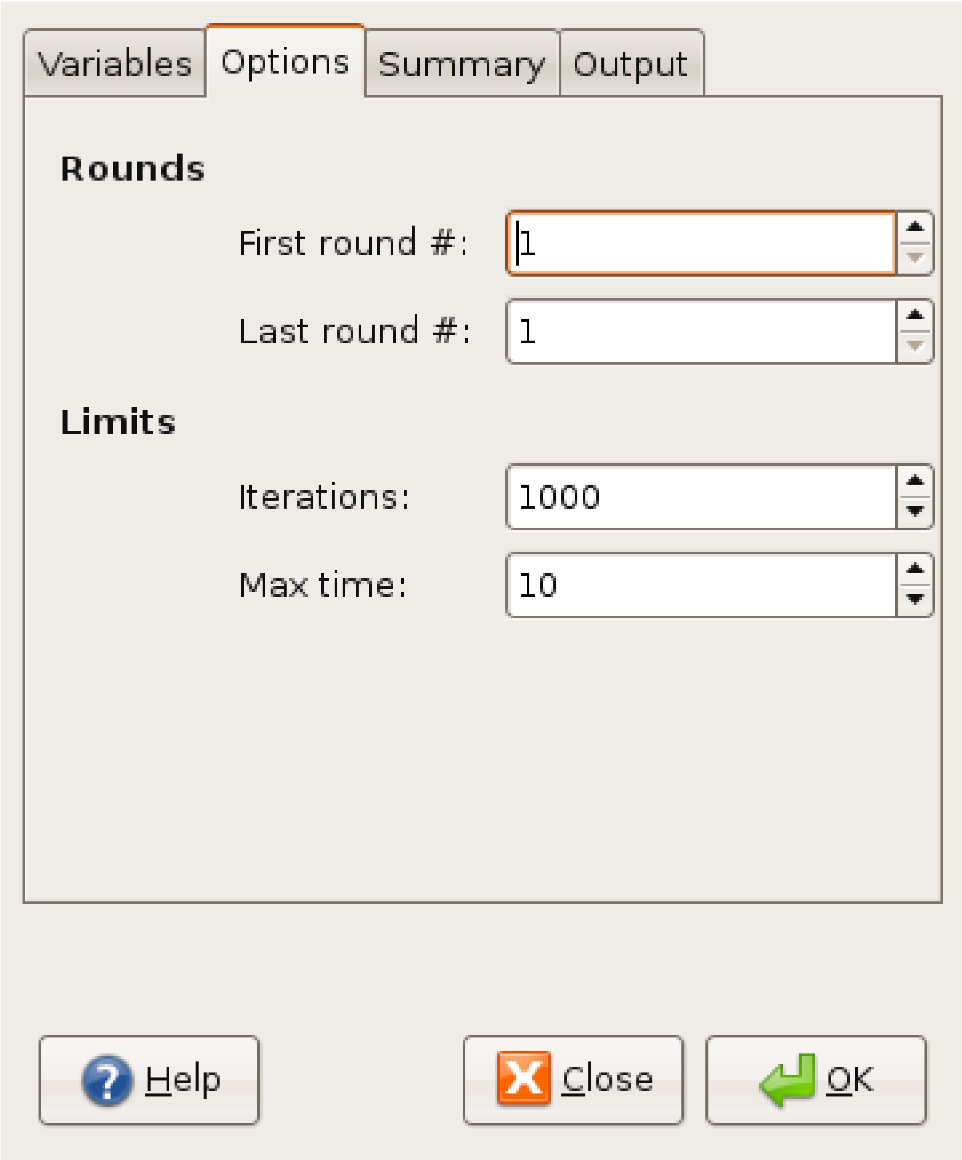

Figure 6: Profit table for newsvendor example

The second pair of options are the number of iterations and the Max time. In a simulation, each iteration is the equivalent of a sample. A sample from a random distribution is taken for each of the input values (as specified in the Variables tab) and the resulting output value(s). The more iterations, the better the estimate of the output value. However, this also takes more time to run. A Max time value is specified in seconds where the simulation will end without output if an individual simulation takes longer then the Max time alloted. If this occurs (see Figure 7), the options are to either increase the Max time value, or decrease the number of iterations. A more drastic option is to change the model so that fewer calculations or samples of random numbers need to be made.

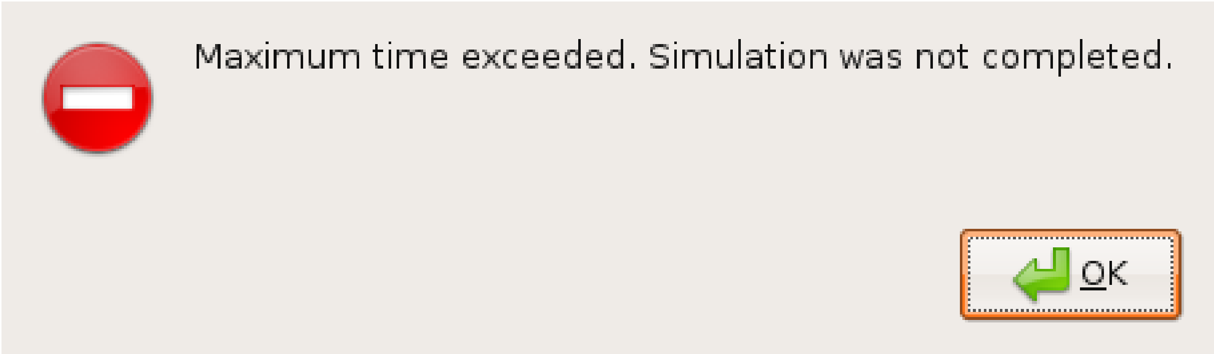

Figure 7: Maximum time for simulation exceeded message box.s

The next tab is the Summary. There are two windows, Simulation Summary and Summary of results (see Figure 8). In simulation summary, there is a description of the simulation parameters.

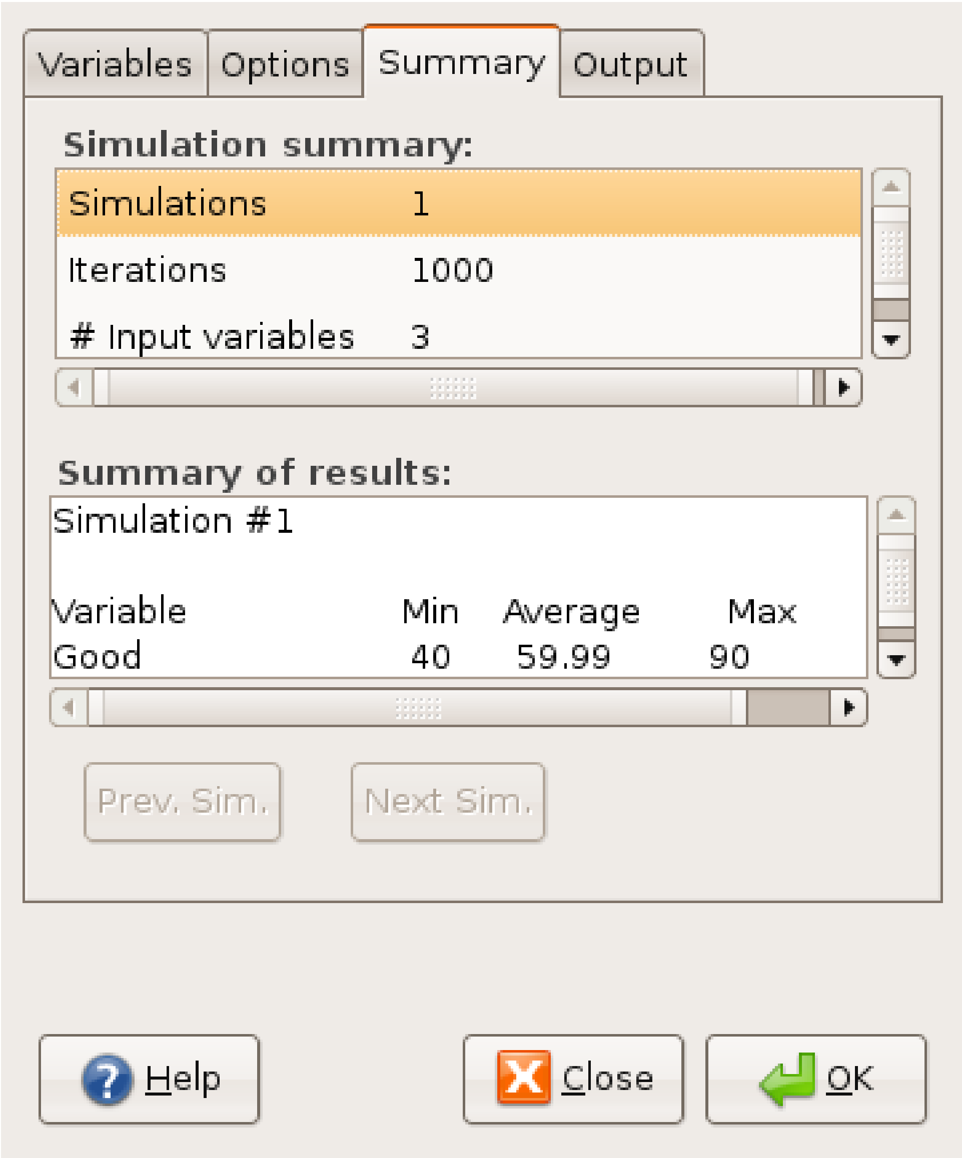

Figure 8: Summary tab for simulation tool.

1. Simulations. Number of rounds as determined in the Simulation Options box.
2. Iterations. Number of iterations in a single simulation round.
3. # input variables. Number of random numbers sampled for each iteration.

There is one set for each round of the simulation performed.

In the summary of results window, there are summary statistics for a round of the simulation. If multiple rounds were done, the results of each round can be browsed by using the 'Prev. Sim.' and 'Next Sim.' buttons below the Summary of results window. For each output and input variable, the summary shows the Min, Average and the Max value. Note that for the input variables, this shows the random number that is the average, max and min. If the statistics on intermediate values, such as a cost distribution, was desired, these intermediate values should be added to the list of output variables.

The last tab is labeled 'Output'. This tab identifies the location where the output table will be generated. There are two sets of options, first the Output Placement then Output Formatting as shown in Figure 9.

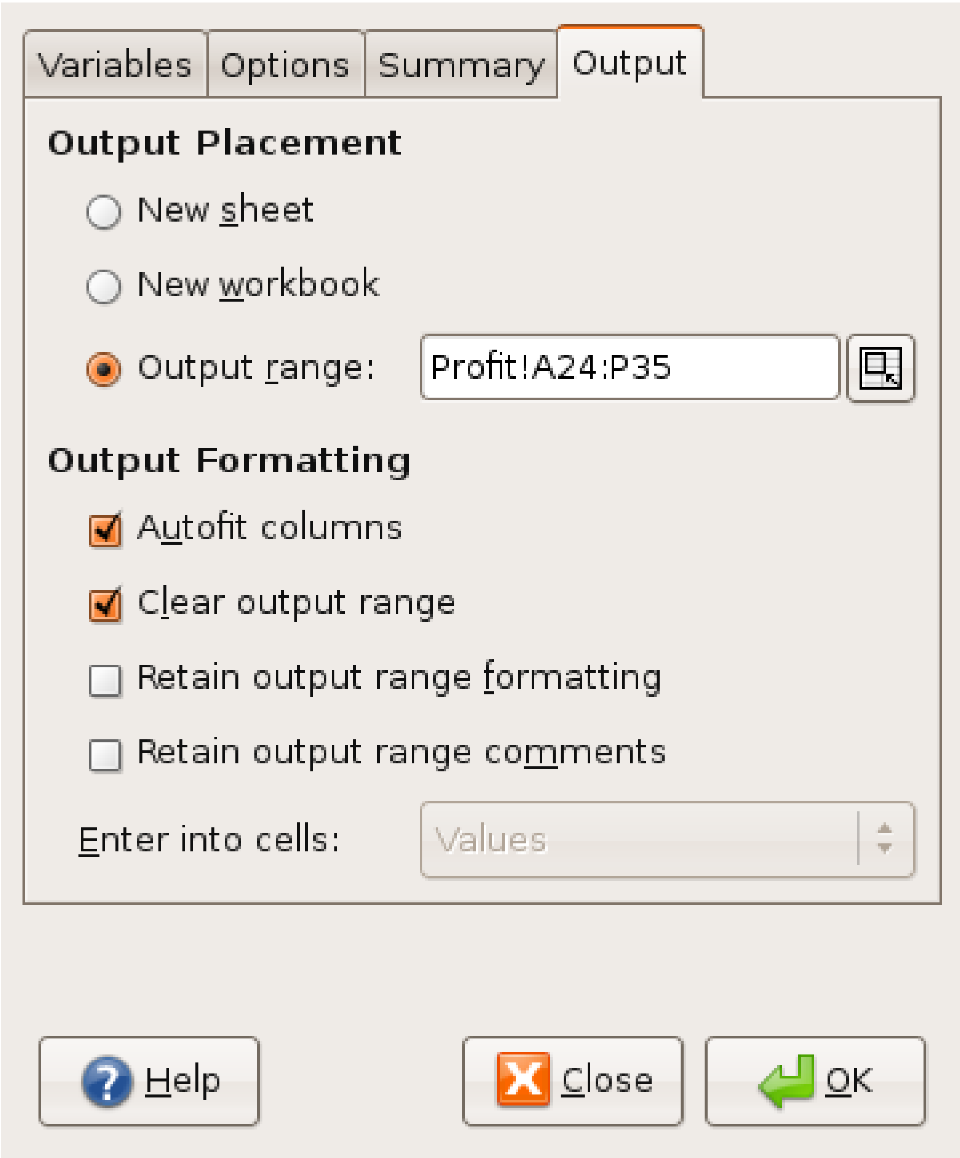

Figure 9: Output options tab for simulation.

The default output placement is 'New sheet'. This will create a new sheet in the Gnumeric workbook labeled 'Simulation Report

(1)', where '1' can be replaced with another number if a tab labeled 'Simulation Report (1)' already exists. The option 'New workbook' creates a Gnumeric workbook named 'Book2.gnumeric' with a tab labeled 'Simulation Report.'

The third option is to embed the output table into an existing worksheet. This is done by specifying the 'Output range'. Note that the output range must be large enough to include the entire table, including heading information. For a single round this requires 11 rows and 16 columns. For example, the range Profit!A24:P35 would contain the statistics for one round with the three input variables and two output variables. As input and output variables change, or the number of rounds of the simulation change, the number of rows required will change.

1.4 Simulation output

The simulation output provides statistics on the output and input variables for each round. The statistics calculated over the iterations in a single round of the simulation. These statistics are described in :

1. Varable type and name (input variables are labeled as "(Input)")
2. Min – Minimum value of variable among iterations of round
3. Mean – Arithmetic mean of variable among iterations of round
4. Max – Maximum value of variable among all iterations of round
5. Median – Median of variable among iterations of round
6. Mode – Mode value among iterations of round. For the input variable, this will be "#N/A"
7. Std. Dev. - Standard deviation of the variable
8. Variance – Second moment of variable
9. Skewness - Third moment of variable
10. Kurtosis – Fourth moment of variable
11. Range – Difference between min and max of variable among iterations of the round
12. Count – Number of iterations in round
13. Confidence (95%) - 95% confidence interval of value, centered on mean
14. Lower Limit (95%) - Lower limit of 95% confidence interval of the value, centered on the mean
15. Upper Limit (95%) - Upper limit of 95% confidence interval of the value, centered on the mean

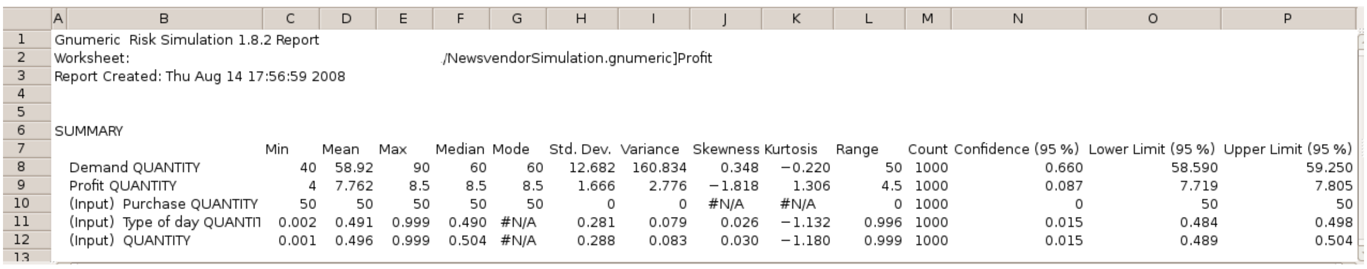

Gnumeric Risk Simulation 1.8.2 Report
Worksheet: ./NewsvendorSimulation.gnumeric]Profit
Report Created: Thu Aug 14 17:56:59 2008

SUMMARY

	Min	Mean	Max	Median	Mode	Std. Dev.	Variance	Skewness	Kurtosis	Range	Count	Confidence (95 %)	Lower Limit (95 %)	Upper Limit (95 %)
Demand QUANTITY	40	58.92	90	60	60	12.682	160.834	0.348	−0.220	50	1000	0.660	58.590	59.250
Profit QUANTITY	4	7.762	8.5	8.5	8.5	1.666	2.776	−1.818	1.306	4.5	1000	0.087	7.719	7.805
(Input) Purchase QUANTITY	50	50	50	50	50	0	0	#N/A	#N/A	0	1000	0	50	50
(Input) Type of day QUANTI	0.002	0.491	0.999	0.490	#N/A	0.281	0.079	0.026	−1.132	0.996	1000	0.015	0.484	0.498
(Input) QUANTITY	0.001	0.496	0.999	0.504	#N/A	0.288	0.083	0.030	−1.180	0.999	1000	0.015	0.489	0.504

Figure 10: Simulation output example.

The output will include a heading, then a table for each round of the simulation. Judicious choice of output variables will also include any intermediate values of interest in the simulation in this table. Each row of the output table has statistics of the values of a variable over the interations of the simulation as shown in Figure 10.

The output will be of the input variables and the output variables that were variables tab of the Simulation window . For the input variables, the output will be the statistics of the random variable used in modeling the input variables. For the output variables, the statistics will be of the output variable. These statistics, in particular the confidence interval, should be examined to ensure the simulation was at a precision adequate for the purpose. Some notes on how to use these statistics for refining the simulation design can be found in Section 1.6 .

1.5 Using SIMTABLE

The `SIMTABLE` function is intended to change a variable in the simulation so that each round of the simulation can be used to evaluate a different scenario. This automates the use of simulation for what-if questions or to build a curve of possible outcomes.

In this example, we will use the `SIMTABLE` function to find the optimal quantity of newspapers to buy. For the purchase quantity in our spreadsheet, we will replace '50' with the following formula in Profit!B16:

```
Profit!B16 = SIMTABLE(50,60,70,80,90)
```

Each entry in the list of the `SIMTABLE` arguments is a value that will be used for the purchased quantity. Each entry corresponds to one round of simulation, as used in Figure 6. In this example there are 5 entries to the SIMTABLE list, so '5' will be entered into the 'Last Round #' option in the Options tab of the Simulation dialog.

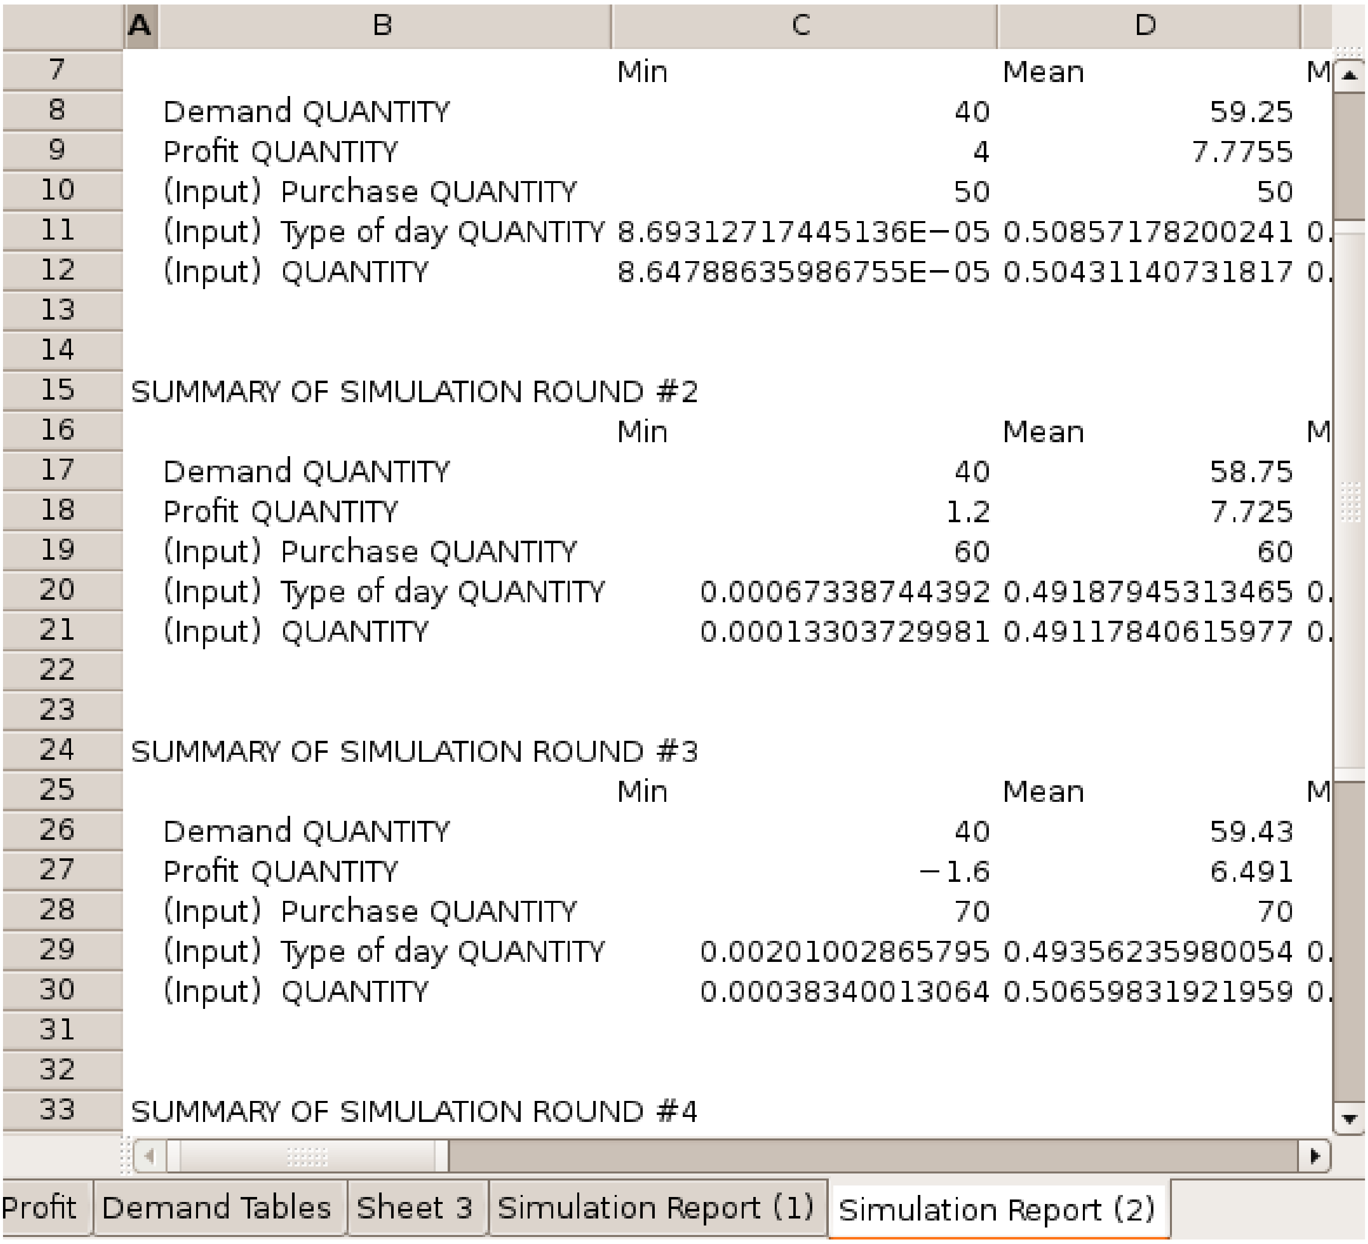

	A	B	C	D
7			Min	Mean
8		Demand QUANTITY	40	59.25
9		Profit QUANTITY	4	7.7755
10		(Input) Purchase QUANTITY	50	50
11		(Input) Type of day QUANTITY	8.69312717445136E−05	0.50857178200241
12		(Input) QUANTITY	8.64788635986755E−05	0.50431140731817
13				
14				
15	SUMMARY OF SIMULATION ROUND #2			
16			Min	Mean
17		Demand QUANTITY	40	58.75
18		Profit QUANTITY	1.2	7.725
19		(Input) Purchase QUANTITY	60	60
20		(Input) Type of day QUANTITY	0.00067338744392	0.49187945313465
21		(Input) QUANTITY	0.00013303729981	0.49117840615977
22				
23				
24	SUMMARY OF SIMULATION ROUND #3			
25			Min	Mean
26		Demand QUANTITY	40	59.43
27		Profit QUANTITY	−1.6	6.491
28		(Input) Purchase QUANTITY	70	70
29		(Input) Type of day QUANTITY	0.00201002865795	0.49356235980054
30		(Input) QUANTITY	0.00038340013064	0.50659831921959
31				
32				
33	SUMMARY OF SIMULATION ROUND #4			

Figure 11: Simulation output example using SIMTABLE and several rounds.

When this simulation is run with 5 rounds, the summary of results will have one entry for each round (value of purchase quantity) that can be previewed using the 'Prev. Sim.' and 'Next Sim.' buttons. The output also has one table for each round of the simulation.

As seen in Figure 11, each value in the original SIMTABLE statement corresponds to a simulation round, with the Purchase Quantity taking on the value from the SIMTABLE list. The analyst can then record the Profit statistics (mean, variance, skewness, kurtosis, 95% confidence intervals) and determine if the simulation results are of sufficient resolution for the analysts purposes.

The use of SIMTABLE to change parameters within the simulation provides a convenient method to do what-if analysis, and analyze the results as a whole.

1.6 Determining the number of iterations

In simulation, one major question is how many iterations are needed to reach a chosen level of precision in the results. Simulation as a tool provides an approximation of the actual relationship between the input and output variables. The precision of the approximation is based on the number of iterations of the simulation done. More iterations in the sample lead to greater precision. But the relationship between iterations and precision depends on the relationship between the variables in the precision. In addition, the analyst must decide which output variable is the variable of interest, and what degree of precision is required. The next step is to determine a sufficiently large number of iterations R be used to satisfy:

$$P\left(\left|\left(\widehat{\Theta}-\Theta\right)\right| \leq \varepsilon\right) \geq 1-\alpha$$

Where $\widehat{\Theta}$ is the estimate of the mean, Θ is the actual mean, ε is the specified error, and $(1-\alpha)$ is the probability that the estimate is within ε of the actual value (i.e. the $(1-\alpha)$ confidence interval). Common values of $(1-\alpha)$ are 95% and 99%. The Simulation Report from Gnumeric includes values for the 95% confidence interval as shown in Figure 10.

The general procedure is as follows:[3]

1. Run simulation for a sample of R_0 iterations. The default value in Gnumeric is 1000, set in the options tab of the Simulation menu, Figure 6.
2. Take the sample variance S_0^2 from the simulation output spreadsheet and determine the sample standard deviation S_0 (see Figure 10)
3. Using $z_{\alpha/2}$ as the z-value of the $(1-(\alpha/2))$ percentile of the standard normal distribution, set the initial estimate of the number of iterations required as the smallest integer R such that $R \geq \left(\frac{z_{\alpha/2}S_0}{\varepsilon}\right)^2$. Note that if R_0 is small, it would be more appropriate to use the student's t-distribution of $t_{\alpha/2,R_0}$ instead of $z_{\alpha/2}$.

In this example, to estimate the profit to within $\varepsilon = 0.05$, first run the simulation with 1000 iterations and a purchase quantity of 50 results in the following (Note that due to the random nature of the simulation, the exact results may vary between simulation runs):

	Mean	Variance	Confidence (95%)
Demand QUANTITY	59.19	152.4	0.64
Profit QUANTITY	7.85	2.51	0.08

Taking the variance of the table, and setting $\varepsilon = 0.05$ and $\alpha = 0.05$, lookup $z_{\alpha/2}$ from a standard normal table. $z_{\alpha/2} = 1.96$ so we have $R \geq \left(\frac{1.96\times\sqrt{2.51}}{0.05}\right)^2 = 3856.8$

Therefore, the minimum number of iterations is 3857. The simulation can then be re-run with 3857 iterations to create a 95% c.i for profit where $\varepsilon \leq 0.05$ In this example with 3857 iterations, we get the following Simulation Report:

	Mean	Variance	Confidence (95%)
Demand QUANTITY	59.11	163.9	0.34
Profit QUANTITY	7.72	2.88	0.04

And the 95% Confidence interval for Profit is less than 0.05.

[3] Adapted from Banks et. al. Discrete-Event System Simulation, 3rd Edition, pp. 414-416.