Chapter 8
Modeling and Simulation: Part 2

Contents

8.1 Introduction

Chapter 8 continues with our discussion of modeling. In particular, we will discuss modeling in the context of *simulation*, a term that we will soon discuss in detail. The terms **model** and **simulation** can be a bit confusing because they are often used interchangeably; that is, simulation *as* modeling and vice versa. We will make a distinction between the two terms, and we will see that in order to simulate a process, we must first create a model of the process. Thus, modeling precedes simulation, and simulation is an activity that depends on *exercising* a model. This may sound like a great deal of concern about the minor distinctions between the two terms, but as we

H. Guerrero, *Excel Data Analysis*, https://doi.org/10.1007/978-3-030-01279-3_8

discussed in Chap. 7, being systematic and rigorous in our approach to modeling helps insure that we don't overlook critical aspects of a problem. Over many years of teaching and consulting, I have observed very capable people make serious modeling errors, simply because they felt that they could approach modeling in a casual manner, thereby abandoning a systematic approach.

So why do we make the distinction between modeling and simulation? In Chap. 7 we developed deterministic models and then exercised the model to generate outcomes based on simple *what-if* changes. We did so with the understanding that not all models require sophisticated simulation. For example, Fr. Efia's problem was a very simple form of simulation. We exercised the model by *imposing* a number of conditions: weather, an expected return on bets, and an expected number of attendees. Similarly, we imposed requirements (rate, term, principal) in the modeling of mortgage payments. But models are often not this simple, and can require considerable care in conducting simulations; for example, modeling the process of patients visiting a hospital emergency department. The arrival of many types of injury and illness, the staffing required to deal with the cases, and the updating of current bed and room capacity based on the state of conditions in the emergency department make this a complex model to simulate.

The difference between the mortgage payment and a hospital emergency department simulation, aside from the model complexity, is how we deal with uncertainty. For the mortgage payment model, we used a manual approach to managing uncertainty by changing values and asking *what-if* questions individually: what if the interest rate is 7% rather than 6%, what if I change the principal amount I borrow, etc. In the OLPS model, we reduced uncertainty to **point estimates** (specific values), and then we used a manual approach to exercise a specific model configuration; for example, we set the number of attendees for *Cloudy* weather to exactly 2500 people and we considered a what-if change to *Entry Fee* from $10 to $50. This approach was sufficient for our simple what-if analysis, but with models containing more elements of uncertainty and even greater complexity due to the interaction of uncertainty elements, we will have to devise complex simulation approaches for managing uncertainty.

The focus of this chapter will be on a form of simulation that is often used in modeling of complex problems—a methodology called **Monte Carlo Simulation**. Monte Carlo simulation has the capability of handling the more complex models that we will encounter in this chapter. This does not suggest that all problems are destined to be modeled as Monte Carlo simulations, but many can and should. In the next section, I will briefly discuss several types of simulation. Emphasis will be placed on the differences between approaches and on the appropriate use of techniques. Though there are many commercially available simulation software packages for a variety of applications, remarkably, Excel is a very capable tool that can be useful with many simulation techniques. In cases where a commercially available package is necessary, Excel can still have a critical role to play in the early or **rapid prototyping** of problems. Rapid prototyping is a technique for quickly creating a model that need not contain the level of detail and complexity that an end-use model requires. It can save many, many hours of later programming and modeling effort by determining the feasibility and direction an end-use model should take.

Before we proceed, I must caution you about an important concern. Building a Monte Carlo simulation model must be done with great care. It is very easy to build faulty models due to careless consideration of processes. As such, the chapter will move methodically toward the goal of constructing a useful and thoroughly conceived simulation model. At critical points in the modeling process, we will discuss the options that are available and why some may be better than others. There will be numerous tables and figures that build upon one another, so I urge you to read all material with great care. At times, the reading may seem a bit tedious and pedantic, but such is the nature of producing a high-quality model—these things cannot be rushed. Try to avoid the need to get to the *punch-line* too soon.

8.2 Types of Simulation and Uncertainty

The world of simulation is generally divided into two categories—**continuous event simulation** and **discrete event simulation**. The difference in these terms is related to how the process of simulation evolves—how results change and develop over some dimension, usually time. For example, consider the simulation of patient arrivals to the local hospital emergency room. The patient arrivals, which we can consider to be **events**, occur sporadically and trigger other events in a *discrete* fashion. For example, if a cardiac emergency occurs at 1:23 pm on a Saturday morning, this might lead to the need of a defibrillator to restore a patient's heartbeat, specialized personnel to operate the device, as well as a call to a physician to attend to the patient. These circumstances require a simulation that triggers random events at *discrete* points in time, and we need not be concerned with tracking model behavior when events are not *occurring*. The arrival of patients at the hospital is not continuous over time, as might be the case for the flow of daytime traffic over a busy freeway in Southern California. It is not unusual to have modeling phenomenon that involves both discrete and continuous events. The importance of making a distinction relates to the techniques that must be used to create suitable simulation models. Also, commercial simulation packages are usually categorized as having either continuous, discrete, or both modeling capabilities.

8.2.1 Incorporating Uncertain Processes in Models

Now, let us reconsider some of the issues we discussed in the Chap. 7, particularly those in Fr. Efia's problem of planning the events of *Vegas Night at OLPS*, and let us focus on the issue of uncertainty. The problem contained several elements of uncertainty—the weather, number of attendees, and the outcome of games of chance. We simplified the problem analysis by assuming **deterministic values** (specific and unchanging) for these uncertainties. In particular, we considered only a single result for each of the uncertain values, for example *rainy* weather as the weather condition. We also reduced uncertainty to a single value determined as an average, for example the winning odds for the game of chance, WOD. In doing so, we fashioned the analysis to focus on various scenarios we *expected* to occur. On the face of it, this is not a bad approach for analysis. We have scenarios in which we can be relatively secure that the deterministic values represent what Fr. Efia will experience conditional on the specific weather condition being investigated. This provides a simplified picture of the event, and it can be quite useful in decision making, but in doing so, we may miss the richness of all the possible outcomes, due to the *condensation* of uncertainty that we have imposed.

What if we have a problem in which we desire a greater degree of accuracy and a more complete view of possible outcomes? How can we create a model to allow simulation of such a problem, and how do we conceptualize such a form of analysis? To answer these questions, let me remind you of something we discussed earlier in Chap. 6—sampling. As you recall, we use sampling when it is difficult, or impossible, to investigate every possible outcome in a population. If Fr. Efia had 12 uncertain elements in his problem, and if each element had 10 possible outcomes, how many distinct outcomes are possible; that is, if we want to consider *all* combinations of the uncertain outcomes, how many will Fr. Efia face? The answer is 1012 ($10 \times 10 \times 10$ . . . etc.) possible outcomes, which is a whopping 1 trillion (1,000,000,000,000). For complex problems, 12 elements that are uncertain with 10 or more possible outcome values each are not at all unusual. In fact, this is a relatively small problem. Determining 1 trillion distinct combinations of possible outcome values is a daunting task. Further, I suggest that it may be impossible.

This is where sampling comes to our rescue. If we can perform carefully planned sampling, we can arrive at a reasonably good estimate of the variety of outcomes we face: not a complete view, but one that is useful and manageable. By this, I mean that we can determine enough outcomes to produce a reasonably complete profile of the entire set of outcomes. This profile will become one of our most important tools for analysis and decision making. We call it a **risk profile** of the problem outcomes, but more on this later. Now, how can we organize our efforts to accomplish efficient and accurate sampling?

8.3 The Monte Carlo Sampling Methodology

In the 1940s, Stanislaw Ulam, working with famed mathematician John von Neumann and other scientists, formalized a methodology for arriving at approximate solutions to difficult quantitative problems. This method came to be called Monte Carlo methods. Monte Carlo methods are based on *stochastic processes*, or the study of mathematical probability and the resolution of uncertainty. The reference to Monte Carlo is due to the games of chance that are common in the gambling establishments of Monte Carlo, in the Principality of Monaco. Ulam and his colleagues determined that by using repeated statistical sampling they were able to arrive at solutions to problems that would be impossible, or very difficult, by standard analytical methods. For the types of complex simulation we are interested in performing in Chap. 8, this approach will be extremely useful. It will require knowledge of a number of well-known probability distributions, and an understanding of the use of random numbers. The probability distributions will be used to describe the behavior of the uncertain events, and the random numbers will become input for the functions generating sampling outcomes for the distributions.

8.3.1 Implementing Monte Carlo Simulation Methods

Now, let us consider the basics of Monte Carlo simulation (MCS) and how we will implement them in Excel. MCS relies on sampling through the generation of random events. The cell function which is absolutely essential to our discussion of MCS, **RAND()**, is contained in the *Math* and *Trig* functions of Excel. In the following, I present six observations and questions, steps if you will, that utilize the RAND() function to implement MCS models:

1. *Uncertain events are modeled by sampling from the distribution of the possible outcomes of each uncertain event*. A sample is the random selection of a value (s) from a distribution of outcomes, where the distribution specifies the outcomes that are possible and their related probabilities of occurrence. For example, the **random sampling** of a fair coin toss is an experiment where a coin is tossed a number of times (the sample size, *n*) and the distribution of the individual toss outcomes is heads with a 50% probability and tails with a 50% probability. Figure 8.1 shows the probability distribution for the fair (50% chance of head or tail) coin toss. If I toss the coin and record the *outcome*, this value is referred to as the **resolution of an uncertain event**, the coin toss. In a model where there are

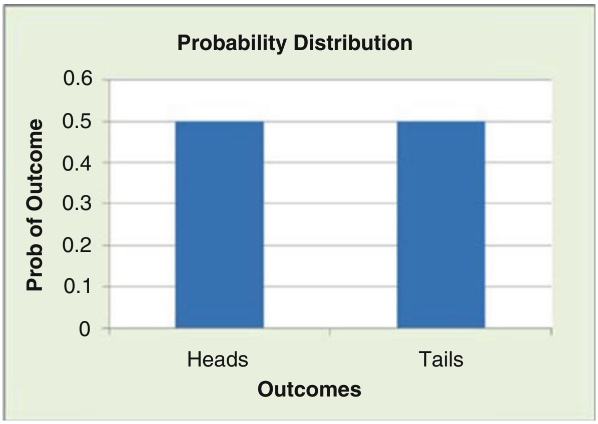

Fig. 8.1 Probability distribution of a fair coin toss

many uncertain events and many uncertain outcomes are possible for each event, the process is repeated for all relevant uncertain events. The resulting values are then used as a fair (random) representation of the model's behavior. Thus, these resolved uncertain events tell us what the condition of the model is at a point in time. Here is a simple example of how we can use sampling to provide information about a distribution of unknown outcomes. Imagine a very large bowl containing a distribution of one million colorful stones: 300,000 white, 300,000 blue, and 400,000 red. You, as an observer, do not know the number of each colorful stone type contained in the bowl. Your task is to try to determine what the true distribution of colorful stones is in the very large bowl. Of course, we can use the colors to represent other outcomes. For example, each color could represent one of the three weather conditions of the *OLPS* problem in the previous chapter. We can physically perform the selection of a colorful stone by randomly reaching into the bowl and selecting a stone, or we can use a convenient analogy. The analogy will produce random outcomes of the uncertain events, and it can be extended to all the uncertain elements in the model.

2. *The RAND() function in Excel is the tool we will use to perform sampling in MCS.* By using RAND() we can create a *virtual* bowl from which to sample. The output of the RAND() function is a Continuous *Uniform* distribution, with output greater than or equal to 0 and less than 1; thus, numbers from 0.000000 to 0.999999 are possible values. The RAND() function results in up to 16 digits to the right of the decimal point. A **Uniform distribution** is one where every outcome in the distribution has the same probability of being randomly selected. Therefore, the sample outcome 0.831342 has exactly the same probability of being selected as the sample outcome of 0.212754. Another example of a Uniform distribution is our fair coin toss example, but in this case the outcomes are discrete (only heads or tails) and not continuous. Is the distribution of colorful stones a Uniform distribution? The answer is *no*. The blue stones have a higher probability of being selected in a random sampling than white or red stones.

   We now turn to a spreadsheet model of our sampling of colorful stones. This model will allow us to discuss some of the basic tenants of sampling. Figure 8.2 shows a table of 100 RAND() functions in cell range B2:K11. We will discuss these functions in greater detail in the next section, but for now, note that cell K3 contains a RAND() function that results in a randomly selected value of 0.7682. Likewise, every cell in the range B2:K11 is the RAND() function, and importantly, every cell has a different outcome. This a key characteristic of the RAND (): each time it is used in a cell, it is independent of other cells containing the RAND() function.

3. *How do we use the RAND() function to sample from the distribution of 30% Red, 30% White, and 40% Blue?* We've already stated that the RAND() function returns Uniformly distributed values from 0 up to, but not including, 1. So, how do we use RAND() to model our bowl of colorful stones? In Fig. 8.2 you can see two tables entitled *Random Numbers Table* and *Translation of Random Numbers to Outcomes*. Each cell location in *Random Numbers Table* has an equivalent location in the *Translation of Random Numbers to Outcomes*; for example, K15 is

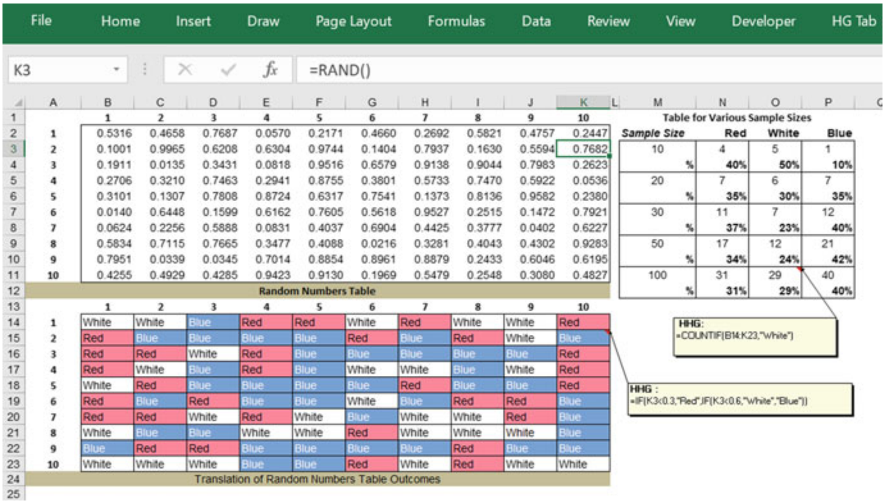

Random Numbers Table

	1	2	3	4	5	6	7	8	9	10
1	0.5316	0.4658	0.7687	0.0570	0.2171	0.4660	0.2692	0.5821	0.4757	0.2447
2	0.1001	0.9965	0.6208	0.6304	0.9744	0.1404	0.7937	0.1630	0.5594	0.7682
3	0.1911	0.0135	0.3431	0.0818	0.9516	0.6579	0.9138	0.9044	0.7983	0.2623
4	0.2706	0.3210	0.7463	0.2941	0.8755	0.3801	0.5733	0.7470	0.5922	0.0536
5	0.3101	0.1307	0.7808	0.8724	0.6317	0.7541	0.1373	0.8136	0.9582	0.2380
6	0.0140	0.6448	0.1599	0.6162	0.7605	0.5618	0.9527	0.2515	0.1472	0.7921
7	0.0624	0.2256	0.5888	0.0831	0.4037	0.6904	0.4425	0.3777	0.0402	0.6227
8	0.5834	0.7115	0.7665	0.3477	0.4088	0.0216	0.3281	0.4043	0.4302	0.9283
9	0.7951	0.0339	0.0345	0.7014	0.8854	0.8961	0.8879	0.2433	0.6046	0.6195
10	0.4255	0.4929	0.4285	0.9423	0.9130	0.1969	0.5479	0.2548	0.3080	0.4827

Translation of Random Numbers Table Outcomes

	1	2	3	4	5	6	7	8	9	10
1	White	White	Blue	Red	Red	White	Red	White	White	Red
2	Red	Blue	Blue	Blue	Blue	Red	Blue	Red	White	Blue
3	Red	Red	White	Red	Blue	Blue	Blue	Blue	Blue	Red
4	Red	White	Blue	Red	Blue	White	White	Blue	White	Red
5	White	Red	Blue	Blue	Blue	Blue	Red	Blue	Blue	Red
6	Red	Blue	Red	Blue	Blue	White	Blue	Red	Red	Blue
7	Red	Red	White	Red	White	Blue	White	White	Red	Blue
8	White	Blue	Blue	White	White	Red	White	White	White	Blue
9	Blue	Red	Red	Blue	Blue	Blue	Blue	Red	Blue	Blue
10	White	White	White	Blue	Blue	Red	White	Red	White	White

Table for Various Sample Sizes

Sample Size		Red	White	Blue
10		4	5	1
	%	40%	50%	10%
20		7	6	7
	%	35%	30%	35%
30		11	7	12
	%	37%	23%	40%
50		17	12	21
	%	34%	24%	42%
100		31	29	40
	%	31%	29%	40%

Fig. 8.2 RAND() function example

the equivalent of K3. Every cell in the translation table references the random numbers produced by RAND() in the random number table. An IF statement is used to compare the RAND() value in K3 with a set of values, and based on the comparison, the IF() assigns a color to a sample. The formula in cell K15 is—(=IF (K3 < 0.3, “Red”, IF (K3 < 0.6,“White”,“Blue”)). Thus, if the K3 value is less than 0.3 Red is returned. If the value in K3 is greater than 0.3, but less than 0.6, then White is returned. If neither of these conditions is satisfied, then Blue is returned. This logical function insures that 30% (0.3–0) of the randomly selected values are red; 30% (0.6–0.3 = 0.3) are white, and the remain- der (1.0–[0.3 + 0.3] = 0.4) are blue. Since K3 is 0.7682, the last condition is met and the value returned is the color Blue. In the case of K2, the random value is 0.2447, and the value in the translation table is the last condition, Red. Thus, values of 0 to 0.2999. . . will cause the return of Red. The value 0.2447 meets this criterion. Incidentally, the cell colors in the Translation table are produced with conditional cell formatting.

Thus, if the distribution we want to model is Discrete, as the colorful stones example, we can partition the range of the RAND() proportionally, and then we can use a logical IF to determine the outcome. For example, if the proportion of Red, White, and Blue changes to 15%, 37%, and 48%, respectively, then the cell functions in the *Translation* table can easily be changed to reflect the new distribution—(=IF (K3 < 0.15, “Red”, IF (K3 < 0.52, “White”, “Blue”)). Note that the second condition (K3 < 0.52) is the cumulative value of the first two probabilities (0.15 + 0.37 = 0.52). If there are four possible discrete outcomes, then there will be a third cumulative value in a third nested IF in the translation table.

4. *We can use larger sample sizes, to achieve greater accuracy in outcomes*. In Fig. 8.2, we see a small table entitled *Table for Various Sample Sizes*. This table collects samples of various sizes (10, 20, 30, 50, and 100 observations) to show how the accuracy of an estimate of the population (the entire bowl) proportions generally increases as sample size increases. For example, for sample size 10, cells B14:K14 form the sample. This represents the top row of the translation table. As you can see, there are 4 red, 5 white, and 1 blue randomly selected colors. If we use this sample of 10 observations to make a statement about our belief about the distribution of colors, then we conclude that red is 40% (4/10), white is 50% (5/10), and blue is 10% (1/10). This is not close to the true population color distribution.

   What if we want a sample that will provide more accuracy; that is, a sample that is larger? In the table, a sample of 20 is made up of observations in B14:K14 and B15:K15. Of course, for any one sample, there is no guarantee that a larger sample will lead to greater precision, but if the samples are repeated and we average the outcomes, it is generally true that the averages for larger samples will converge to the population proportions of colors more quickly than smaller samples. It should also be intuitively evident that more data (larger sample sizes) leads to more information regarding our population proportions of colors. At one extreme, consider a sample that includes the *entire population* of one million colorful stones. The sample estimates of such a sample would estimate population proportions exactly—30% red, 30% white, and 40% blue. Under these extreme circumstances, we no longer have a sample; we now have a **census** of the entire population.

   Note how the sample proportions in our example generally improve as sample size increases. The sample size of 50 yields proportion estimates that are about as accurate as the sample size of 30. They are clearly better estimates than sample sizes of 10 and 20. The sample size of 100 results in almost exact values of the color proportions. Another sample of 100 might not lead to such results, but you can be assured that a sample size of 100 is usually better than a sample of 10, 20, 30, or 50.

   To generate new values of RAND() we recalculate the spreadsheet by pressing the F9 function key on your keyboard. This procedure generates new RAND() values each time F9 is depressed. Alternatively, you can use Calculation Group in the Formulas Ribbon. A tab entitled Calculation Now permits you to recalculate. Also in the tab, you can control the automatic recalculation of the spreadsheet. See Fig. 8.3. You will find that each time a value or formula is placed in a cell location, all RAND() cell formulas will be recalculated if the Automatic button is selected in the Calculation Options subgroup. As you are developing models, it is generally wise to set Calculation to Manual. This eliminates the repeated, and often annoying, recalculation of your worksheet that occurs as you are entering each cell value.

5. *Why do we need to perform many replications of an experiment?* In simulation, we refer to the repeated sampling of uncertain events as **replications**. The term refers to the replication, or repetition, of an experiment, with each experiment

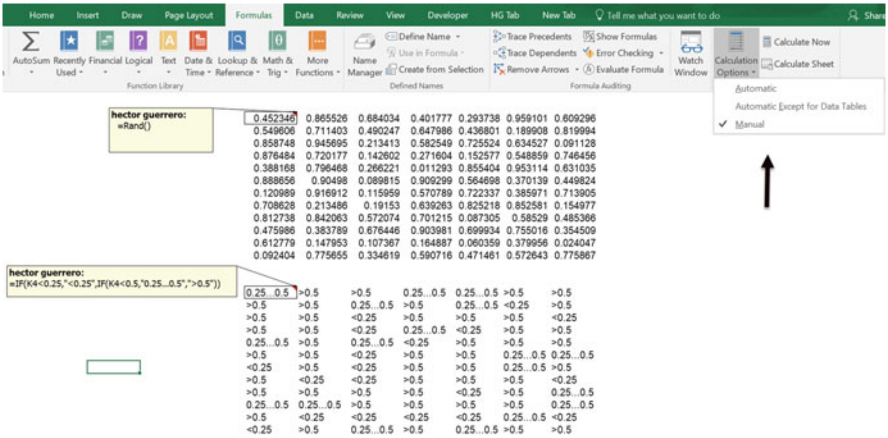

Fig. 8.3 Control of worksheet recalculation

representing a single resolution of the uncertainties of the problem. For Fr. Efia, an experiment represents a single day of operation of his event, and each replication of the experiment results in observed daily, total revenue. In a complex problem, many individual uncertain elements will be combined to produce a complex distribution of a composite outcome value. The more experiments we conduct, the closer we approach the true behavior of complex models. In fact, in many cases it is impossible to understand what the combined distribution of results might be without resorting to large numbers of replications. The resulting distribution is the risk profile that the decision maker faces, and it becomes a tool for decision making. If we perform too few replications the risk profile is likely to be inaccurate. For example, in Fig. 8.4 you can see the graphical results which we produced in Fig. 8.2. The outcomes are attached to the graph in a data table below the graph. Although risk profiles are often associated with graphs that indicate the probability of some monetary result, the results in Fig. 8.4 represents observed distribution of colors. Thus, the risk profile for the sample size of 10, the first of five, is 30% red, 20% white, and 50% blue. There are 5 risk profiles in the figure (sample size 10, 20, 30, 50, and 100), as well as the *Actual* distribution, which is provided for comparison. This exercise provides two important take-always: (a) an initial introduction to a *risk profile*, and (b) a demonstration of the value of larger sample sizes in estimating true population parameters, like proportions.

6. *In summary, the Monte Carlo methodology for simulation, as we will implement it, requires the following*: (a) develop a complete definition of the problem, (b) determine the elements of the model that are uncertain, and the nature of the uncertainty in terms of a probability distribution that represents its behavior, (c) implement the uncertain elements by using the RAND() or other Excel functions, (d) replicate a number of experiments sufficient in size to capture accurate behavior, (e) collect data from experiments, (f) present the risk profiles

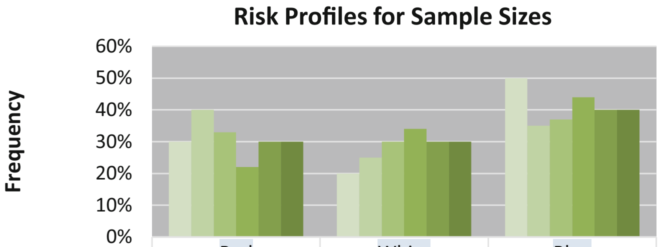

	Red	White	Blue
n= 10	30%	20%	50%
n= 20	40%	25%	35%
n= 30	33%	30%	37%
n= 50	22%	34%	44%
n= 100	30%	30%	40%
n= 1,000,000	30%	30%	40%

Fig. 8.4 Risk profiles for various sample sizes

resulting from the experiments, (g) perform sensitivity analysis on results, and (h) make the appropriate decisions based on results and the decision maker's attitude toward risk.

These steps represent a systematic approach for modeling processes and conducting simulation experiments. Now, let us turn to a discussion of probability distributions. Since it will be of utmost importance that we incorporate uncertainty into our MCS, we will need a basic understanding of how we specify uncertain events. We will introduce the basics of a Poisson Arrival process in the next section, but this is just one way to deal with arrival uncertainty. There are many other ways to describe uncertain arrivals. In the discussion that follows, we will consider some commonly used probability distributions and density functions.

8.3.2 A Word About Probability Distributions

Obviously, we could devote an entire book to this topic, but in lieu of a detailed discussion, there are a number of issues that are essential to understand. First, there are three basic ways we can model uncertain behavior with probability distributions: theoretical distributions, expert opinion, or empirically determined distributions. Here are a few important characteristics to consider about distributions:

1. *Discrete Distributions*—Recall that distributions can be classified as either Discrete or Continuous. **Discrete distributions** permit outcomes for a discrete, or countable, number of values. Thus, there will be gaps in the outcomes. For example, the outcomes of arrival of patients at an emergency room hospital during an hour of operation are discrete; they are a *countable* number of individuals. We

can ask questions about the probability of a single outcome value, four patients for example. We can also ask about the probability of a range of values, between four and eight patients. The Poisson is a very important Discrete distribution that we will discuss later in one of our examples. It is restricted to having integer values.

2. *Continuous Probability Distributions or Probability Density functions*—**Continuous distributions** permit outcomes that are continuous over some range. *Probability density functions* allow us to describe the probability of events occurring, in terms of ranges of numerical outcomes. For example, the probability of an outcome having values from 4.3 to 6.5 is a legitimate question to ask of a Continuous distribution. But, it is not possible to find the probability of a point value in a Continuous distribution. Thus, we *cannot* ask—what is the probability of the outcome five in a Continuous distribution? Probabilities of individual outcomes are undefined for Continuous distributions.
3. *Discrete and Continuous Uniform Distribution*—Figs. 8.5 and 8.6 show Discrete and Continuous Uniform distributions, respectively. As you can see in Fig. 8.5, the probability of the outcome 7 is 0.2, and all outcomes have equal probability.

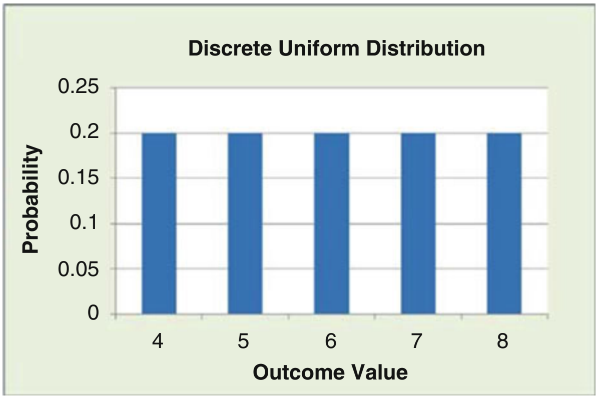

Fig. 8.5 Discrete uniform distribution example

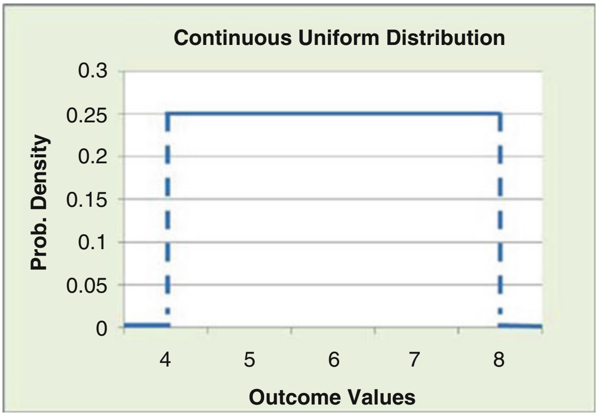

Fig. 8.6 Continuous uniform distribution example

In the case of the Continuous Uniform in Fig. 8.6, we have a distribution from the outcome range of values 4–8, and the distribution is expressed as a *probability density function*. This relates to our discussion in 2, above. The total area under the Continuous distribution curve is equal to 1. Thus, to find the probability of a range of values, say the range 4–6, we find the proportion of the area under the distribution that is implied by the range. In this case, the range 4–6 covers 2 units of a total interval of 4 (8–4).

The area between each successive integer value (4–5, 5–6, etc.) represents 25% of the area of the entire distribution. Thus, we have 50% of the area under the curve ([6–4] $*$ 0.25 = 0.5). The density is calculated as the inverse of the difference between the low and high values of the outcome range (1/[high range value—low range value]), in our case 0.25 (1/[8–4]). The probability of similarly sized intervals is equal. Although the Uniform describes many phenomena, there is one particular use of the Uniform that is very interesting and useful. It is often the case that a decision maker simply has no idea of the relative frequency of one outcome versus another. In this case, the decision maker may say the following—I just don't know how outcomes will behave relative to each other. This is when we resort to the Uniform to deal with our *total* lack of specific knowledge, and we attempt to be *fair* about our statement of relative frequency. For example, consider a neighborhood party to which I invite 50 neighbors. I receive regrets (definitely will not attend—they don't like me) from 25, but I have *no* idea about the attendance of the 25 remaining neighbors. A Discrete Uniform is a good choice for modeling the 0–25 neighbors that might attend; any number of attendees, 0–25, is equally likely, since we do not have evidence to the contrary.

4. *Specification of a distribution*—Distributions are specified by one or more parameters; that is, we provide a parameter or set of parameters to describe the specific form the distribution takes on. Some distributions, like the Normal, are described by a location parameter, the mean (Greek letter μ), and a dispersion parameter, the standard deviation (Greek letter σ). In the case of the attendees to our neighborhood party, the Uniform distribution is specified by a lower and upper value for the range, 0–25. The Poisson distribution has a single parameter, the average arrival rate, and the rate is denoted by the Greek letter λ. Figure 8.7 shows a Poisson distribution with a $\lambda = 5$. Note that the probability of obtaining an outcome of 7 arrivals is slightly greater than 0.1, and the probabilities of either 4 or 5 are equal, approximately 0.175 each. This must be taken in the context of an average arrival rate, $\lambda = 5$. It makes sense that if the average arrival rate is 5, the values near 5 will have a higher probability than those that are distant, for example 11, which has a probability of less than 0.01. Thus, the further away an outcome is from the average arrival rate, the smaller the probability of its occurrence for the Poisson.
5. *Similarity in distributions*—Many distributions often have another distribution for which they possess some similarity. Other distributions may often represent a *family* of distributions. The Beta distribution family is one such Continuous distribution. In fact, the Uniform is a member of the Beta family. The Poisson

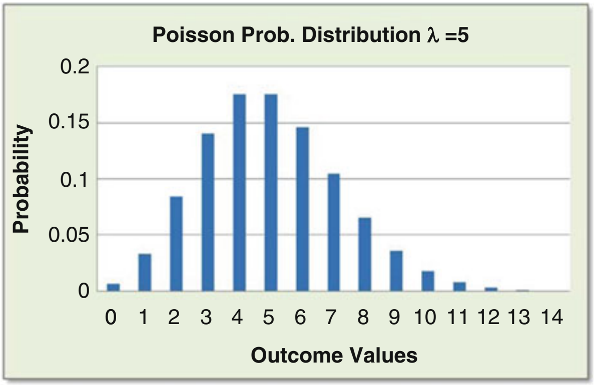

Fig. 8.7 Poisson probability distribution with $\lambda = 5$

Fig. 8.8 Binomial probability distribution example

is closely related to the Binomial distribution. Like the Poisson, the Binomial is also a Discrete distribution and provides the probability of a specific number of successes, k, in n trials of an experiment that has only two outcomes, and where the probability of success for an individual trial (experiment) is p. Thus, we could ask the probability question—what is the probability that I will get exactly 9 heads (k) in 20 tosses (n) of a coin, where the probability of a head is 0.5 (p). In Fig. 8.8 we see that the probability is approximately 0.12. In situations where the results of a trial can only be one of two outcomes, the Binomial is a very useful discrete distribution.

6. *Other important characteristics of distributions*—

   (a) *Shape*. Most of the distributions we have discussed, thus far, are *unimodal*, that is, they possess a single maximum value. Only the Uniform is *not* unimodal. It is also possible to have bimodal, trimodal, and multimodal distributions, where the modes need not be equal, but are merely localized maximum values. For example, in Fig. 8.9 we see a bimodal distribution where one local maximum is the outcome 4, and the other outcome is 12. This terminology uses the definition of the mode (the maximum occurring value) loosely, but it is generally accepted that a bimodal distribution need not have equal probability for local modes. Additionally, some distributions are

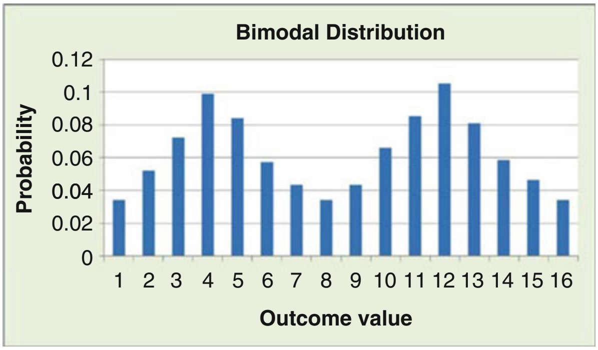

Fig. 8.9 Bimodal distribution

symmetrical in shape (Normal or Uniform), while others need not be (Poisson, Beta, or Binomial).

(b) *Behavior of uncertain events*. Often, we must assume particular behavior in order to justify using certain distributions. For example, in a Poisson arrival process, the arrivals are defined to be independent of one another, they occur at random, and the average number of events over a unit of time (or space) is constant. (More on the Poisson in the next section.) Even when we *bend* the strict nature of these conditions a bit, the approximation of a Poisson arrival process can be quite good.

(c) *Empirically based distributions*. To this point our discussion has been about theoretically based distributions. In theoretical distributions, we assume a theoretical model of uncertainty, such as a Poisson or Normal distribution. But, decision makers can often collect and record empirical data, and develop distributions based on this observed behavior. Our distribution of colorful stones is such a case. We do not assume a theoretical model of the distribution of colors; we have collected empirical data, in our case, through sampling, which leads us to a particular discrete distribution.

In the following section, we will concentrate on an introduction to the Poisson distribution. It will be somewhat complex to create an approach that will allow Poisson arrivals.

8.3.3 Modeling Arrivals with the Poisson Distribution

Earlier, I mentioned that *arrivals* in simulations are often modeled with a Poisson distribution. When we employ the Poisson to describe arrivals, we refer to this as a **Poisson Arrival Process**. The arrivals can be any physical entity, for example, autos seeking service at an auto repair facility, bank clients at an ATM, etc. Arrivals can also be more abstract, like *failures* or *flaws* in a carpet manufacturing process, or questions at a customer help desk. The arrivals can occur in time (auto arrivals during the day) or in space (location of manufacturing flaws on a large area of household carpet). So, how do we sample from a Poisson distribution to produce arrivals?

Figure 8.10 serves as an example of how we will manage the arrival data for the Autohaus simulation we will perform later in the chapter: an advanced modeling example that is essentially a discrete event simulation. Describing a process as having Poisson arrivals requires that a number of assumptions should be maintained:

1. Probability of observing a single event over a small interval of time or within a small space is proportional to the length of time or area of the interval
2. Probability of simultaneous events occurring in time and space is virtually zero
3. Probability of an event is the same for all intervals of time, and in all areas of space
4. Events are independent of one another, so the occurrence of an event does not affect another

Although it may be difficult to adhere to all these assumptions for a particular system, the Poisson's usefulness is apparent by its widespread use. For our purposes, the distribution will work quite well.

Now, let us examine in detail how we simulate Poisson arrivals. Recall how we introduced the RAND() function in Fig. 8.2. We generated uniformly distributed random numbers and then assigned an outcome value (a color) to the random

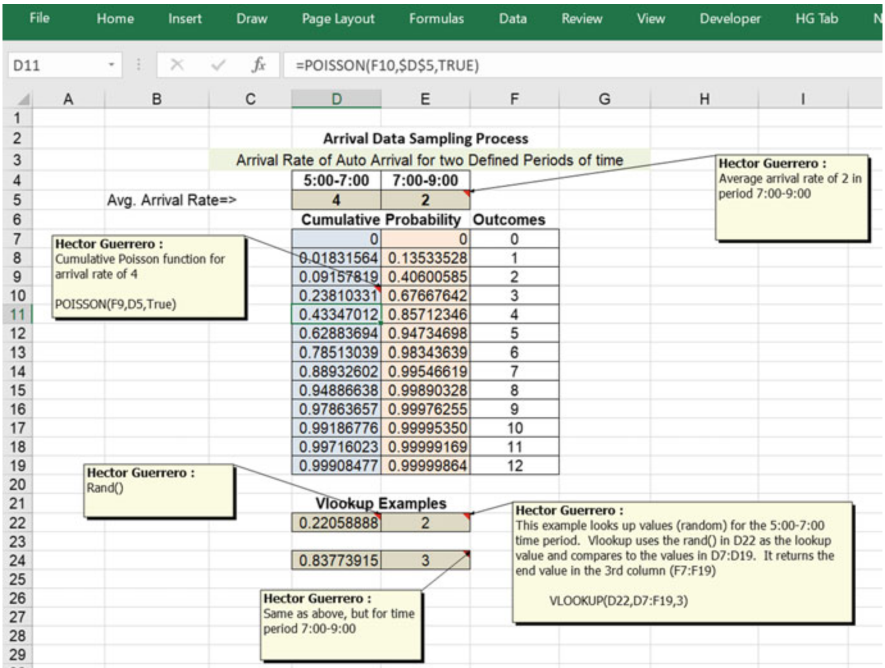

Fig. 8.10 Auto arrival sampling area of brain

number depending on its value. If the random number value was between 0 and 0.3, a Red was returned; if between 0.3 and 0.6, a White was returned; if between 0.6 and 1.0 a Blue was returned. We will use the table of *cumulative* Poisson probability values in a similar fashion. Although we did not mention the cumulative nature of the comparison values before, the IF() used cumulative values to determine the colors for our sampling. To build a cumulative Poisson probability table, we will use the internal Excel cell function, **POISSON(x, mean, cumulative)**. Then we will use the table as the basis for sampling the number of arrivals for a unit of time. The arguments of the function, *x* and *mean*, are values that the user provides. By placing the term *true* in the third argument of the function, *cumulative*, a cumulative value will be returned; that is, the value x = 3 will return the probability of 0, 1, 2, and 3 arrivals for the Poisson distribution.

Now, consider the table in Fig. 8.10 associated with the average arrival rate, λ, of 4 (in column D). In this figure we consider the arrival of autos at a repair facility in two distinct time periods: 5:00–7:00 and 7:00–9:00. The arrival rate of the later time period is 2; thus, on average more cars arrive earlier rather than later. Beginning in cell D7 and continuing to D19, the table represents the successive cumulative probabilities for a Poisson distribution with an average arrival rate of 4 (cell D5 value). Thus, the arithmetic *difference* between successive probability values represents the probability of obtaining a specific value. For example, the difference between 0.0183156 and 0.0915782 is 0.0732626, the probability of exactly 1 arrival for the Poisson distribution with an average arrival rate of 4 per unit of time. Similarly, the difference between 0 and 0.0183156 is 0.0183156, which is the probability of an outcome of exactly 0 arrivals. The numbers in cell F7 to F19 are the number of arrivals that will be returned by a lookup process of random sampling. Next, we will discuss the details of the process of sampling through the use of lookup functions.

8.3.4 VLOOKUP and HLOOKUP Functions

To demonstrate how we can use the table to sample the number of random arrivals in a time period, let us turn our attention to cells E22 and E24 in Fig. 8.10. These two cells form the heart of the sampling process, and rely on the *vertical* lookup function, *VLOOKUP(value_lookup, table_array, col_index_num)*. To understand the use of the **VLOOKUP** and its closely related partner, **HLOOKUP** (horizontal lookup), we introduce Fig. 8.11.

The VLOOKUP and HLOOKUP are a convenient way to return a value from a table based on a comparison to another value. Except for the obvious difference in vertical and horizontal orientation, there is no difference in the cell functions. Consider the utility of such a function. In Chap. 4 we used the IF() to convert a dollar value into a *payment category*. In that example, we noted that depending on the number of categories, we might have more categories than the maximum allowed

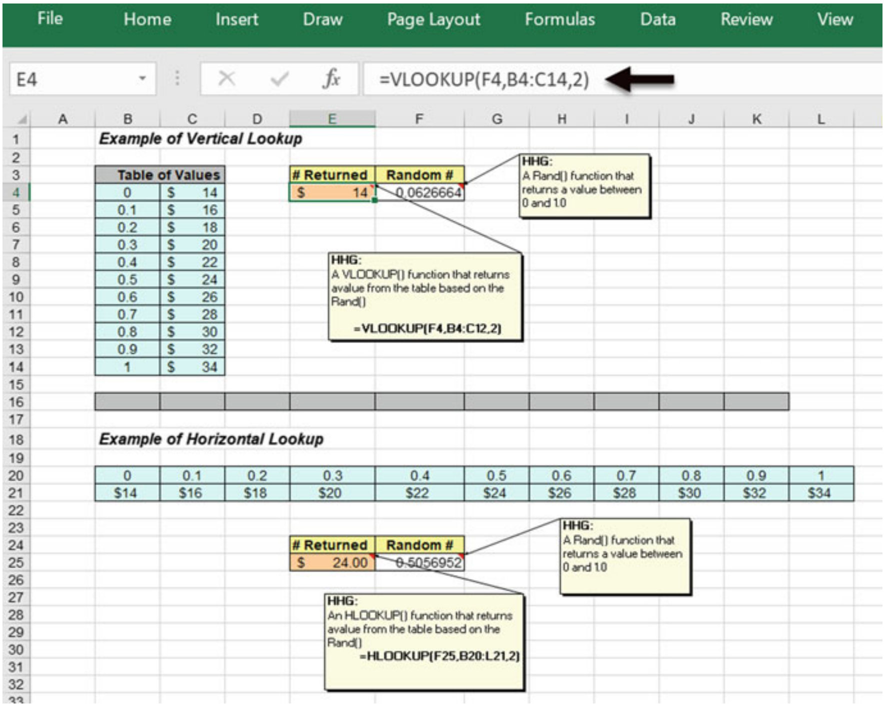

Fig. 8.11 VLOOKUP and HLOOKUP example

number of *nested* IF()s; thus, under these circumstances a lookup function should be used.

Now, let us take a look at the examples in Fig. 8.11 of lookups that convert a fractional value from 0 to 1 into a numerical dollar value. We will concentrate on the VLOOKUP due to the vertical nature of the Poisson cumulative probability table we used earlier. The function has three important arguments: *value_lookup*, *table_array*, and *col_index_num*. The lookup requires that a table be constructed with two types of values—a lookup value and a table value. The lookup value used in the example is a RAND(), and it represents a random sampling process, and this value is to be converted by the table to some associated table value.

We begin by generating a RAND(). It is compared to the leftmost column of a vertical table or the top row of a horizontal table. In Fig. 8.11 a lookup value is in cell F4 and the table is located in B4:C12. The random number generated in the vertical example is 0.0626664, cell F4. The last argument in the lookup function is the column index number which is to be returned, in this case 2, which will return a value in the second column of the table. In simple two column tables, this number is always 2. If a table has more than two columns, to permit return of other associated values or to combine several tables into one, then a column number must be chosen to represent the table value of interest. Finally, the function *takes* the value lookup (0.0626664), and it finds the region in the table that contains the value. For our example, 0.0626664 is found between 0 (B4) and 0.1 (B5). The convention used for returning a value from the second column is to return the value associated with the topmost value (0) in the range. In this case, the return is $14 which is adjacent to 0. If the lookup value from the RAND() function had been exactly 0.1, the value returned would have been $16.

Now, let's return to our Poisson table example in Fig. 8.10 to describe how we perform random sampling of a Poisson distribution. The VLOOKUP in cell E22 compares the value of a random number, D22, with the table of cumulative probability numbers. It then returns a value from 0–12 depending of the value of the RAND(). For example, the random number 0.22058888 is generated in cell D22, and the VLOOKUP searches values in D7 through D19. When it encounters a value higher than the random number, in this case 0.23810331, it returns the value in column F, in the row *above* 0.23810331. The number returned, 2, is in the third column of the table. If we repeat the process many, many times, the recorded values will have a frequency distribution that is approximately Poisson distributed with average arrival rate 4, which is what we are attempting to achieve. Note that we can use the same approach for any Discrete distribution: (1) create a cumulative probability distribution and, (2) sample using a RAND() function to determine a randomly selected value from the distribution. This simple mechanism is a fundamental tool for generating uncertain events from a distribution.

8.4 A Financial Example–Income Statement

Not all simulations are related to discrete events. Let us consider a simple financial example that simulates the behavior of an income statement. I emphasize simple because income statements can be quite complex. Our purpose in this exercise is to demonstrate the variety of simulation that is possible with the MCS method.

Table 8.1 shows the components of a typical income (profit or loss) statement and commentary on the assumptions that will be used to model each component. I have selected a variety of distributions to use in this example, including the **Normal distribution**, often referred to as the **Bell Curve** due to its shape. It is a very commonly used and convenient distribution that has a remarkable range of applications. It is distinguished by its symmetry, a central tendency ("peaked-ness"), and probabilities for lower and higher values that extend to infinity, but with very low probability of occurrence. The symmetry and central tendency of the Normal probability distribution (density function) is particularly useful to modelers. Such observable variables as individual's weight, shoe size, and many others, are often modeled with Normal distributions, although the infinitely extending low and high values could lead to a foot size that is miniscule or one that fills a soccer stadium. Fortunately, these extreme values have very, very low probability of occurrence.

A Normal distribution is described by two parameters—mean and standard deviation (or variance). The formula for generating a value from a Normal distribution in Excel is **NORMINV**(RAND(), mean, standard deviation). Note that the

Table 8.1 Income statement example data

Component	Assumptions
Sales revenue	Sales revenue = Units ∗ Unit Price Distributions: Discrete Units: 30%–75,000 units; 70%–10,000 units discrete unit price: 25%–$1.50; 50%–$2.00; 25%–$2.50
Cost of goods sold expense (COGS)	COGS = Percentage ∗ Sales revenue Distribution: Percentage = Normal Dist.—Mean–30; Std Dev–5
Gross margin	= Sales revenue-COGS
Variable operating expense (VOE)	VOE = Sales revenue ∗ Percentage Distribution: Percentage = Continuous uniform Dist.—10%–20%
Contribution margin	= Gross margin – VOE
Fixed expenses (FE)	A constant value = $6000
Operating earnings (EBIT)	= Contribution margin – Fixed expenses
Interest expense (IE)	A constant value = $3000
Earnings before income tax	= EBIT-interest expense
Income tax expense	Conditional percentage of [EBIT-interest expense] 35% < $20,000; 55% > =$20,000
Net income (Profit)	= Earnings before income tax – Income tax

function uses the familiar RAND() function as its first argument. The RAND() is the element function that samples from the distribution. In our financial example, the percentage of revenue used to calculate the cost of goods sold, COGS, is determined by a Normal distribution with mean 30 and standard deviation 5. See the COGS Expense in Table 8.1 for detail. It would be extremely rare that a value of 10 or 50 would be returned by the function, since this represents values that are four standard deviations below and above the mean. This is noteworthy since it is sometimes possible to return negative values for the Normal if the standard deviation is large relative to the mean; for example, a mean of 30 and standard deviation of 15 could easily return a negative randomly sampled value, since values 2 standard deviations below the mean are not difficult to obtain. If a Normal distribution is used with these types of parameter values, the modeler must introduce a function to eliminate negative values if they are nonsensical, set to 0. For example, a negative weight or height would be such a circumstance.

Now, let us examine the results of our simulation. I have placed the *Brain* on the same worksheet as the *Calculations*, and Fig. 8.12 shows the general structure of the combined *Brain* and *Calculation* worksheets. The *Brain* contains the important values for sampling and also the constant values used in the calculations. In the lower part of Fig. 8.12 you can see that the structure of the Profit or Loss statement is translated horizontally, with each row representing an experimental observation of the statement. Of the 13 observations visible for the model, only two, Obs. 8 and 10, results in a loss (−$1478.45, −$303.09), and the others result in a profit. The advantages of this structure are: (1) there are many more rows visible on a work-sheet than columns, and (2) since there are many more rows on a worksheet than

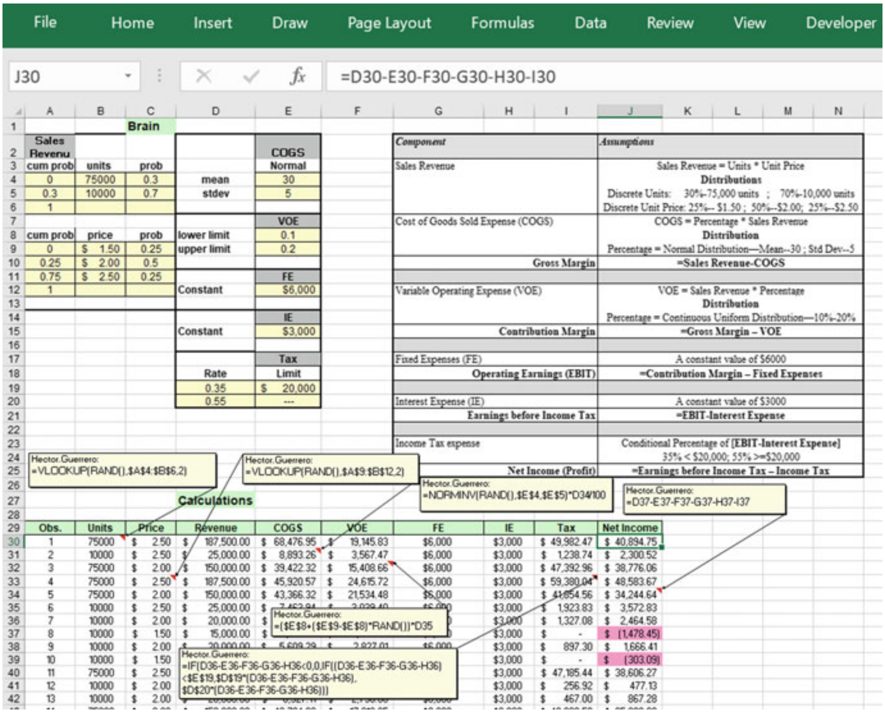

Fig. 8.12 Brain and calculation for financial example

columns, many more observations are possible. In this model, I generate 500 experimental observations of the Profit or Loss statement.

The calculations for this model are relatively straightforward. The calculations for *Units* and *Price*, which determine *Revenue*, are VLOOKUPs that sample from the *Sales Revenue* (cell A2) section of the *Brain* and are shown in Fig. 8.12. The calculation of the *COGS* uses the NORMINV function discussed above, to determine a randomly sampled percentage. Variable Operating Expense, *VOE*, uses the values for lower and upper limits (10% and 20%) through a linear transformation to return continuous uniformly distributed values:

$$\text{lower limit} + (\text{upper limit} - \text{lower limit}) * \text{RAND}().$$

The formula leads to continuous values between the lower and upper limit, since RAND() takes on values from 0 to 1. For example, if RAND() is equal to the extreme upper value 1, then the value of the expression is simply the *upper limit*. Conversely, if RAND() is equal to the extreme lower value 0, then the expression is equal to the *lower limit*.

In Fig. 8.13, we see summary statistics (cell range I530:J535) for the simulation of 500 observations. The average profit or loss is positive, leading to a profit of

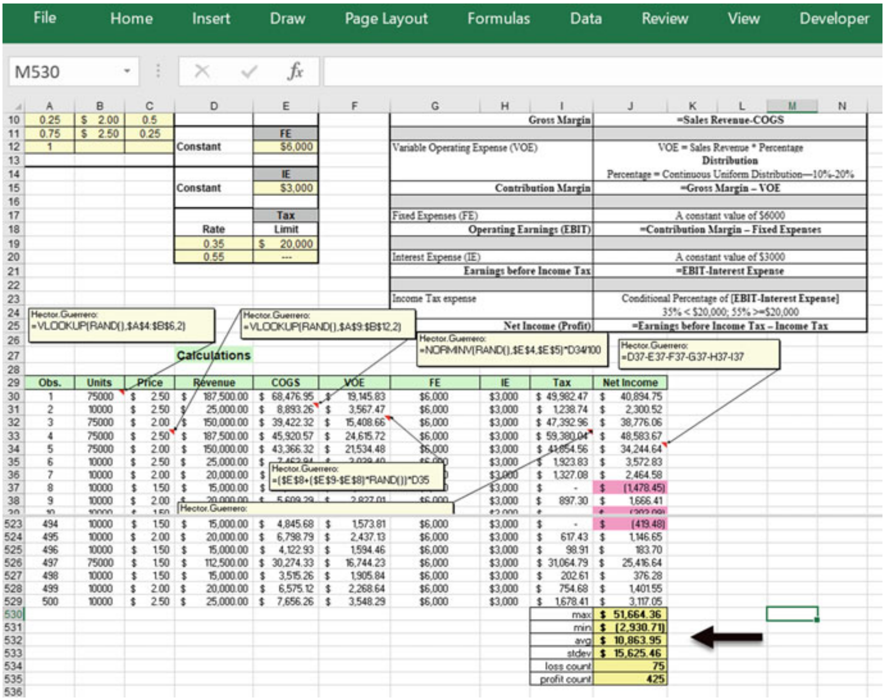

Fig. 8.13 Summary statistics for financial model

$10,863.95 and a very substantial standard deviation of $15,625.46. The coefficient of variation for the 500 observations is greater than one (15,625.46/10,863.95). Thus, our model shows a very wide range of possible values with a max of $51,957.41 and a min of −$2930.71. Also, note that 75 (cell J534) of 500 observations are losses, or 15.0% of the experimental total outcomes. This is valuable information, but the risk associated with this model is even more clearly represented by the risk profile in Fig. 8.14. In this figure we see clearly a picture of the possible outcomes of the model. This is a classic risk profile, in that it presents the range of possible monetary outcomes and their relative frequency of occurrence. Given a large sample of observations (500—more on sample size later), we can assume the frequency distribution of outcomes is representative of the probability distribution of outcomes.

Consider the calculations associated with row 30 of the worksheet in Fig. 8.13, the first observation. The number of units produced is 75,000 and the price per unit is $2.50, both resulting from VLOOKUPs involving a random sampling of Discrete distributions. Cost of Goods Sold (COGS), $68,476.95, is a percentage of revenue determined by a Normal distribution outcome value divided by 100. The variable operating expense (VOE) is determined by sampling a Uniform distribution ($19,145.83). The fixed and interest expense (FE and IE) are constant values of $6000 and $3000, respectively. Finally, the tax is determined by logic—if a loss

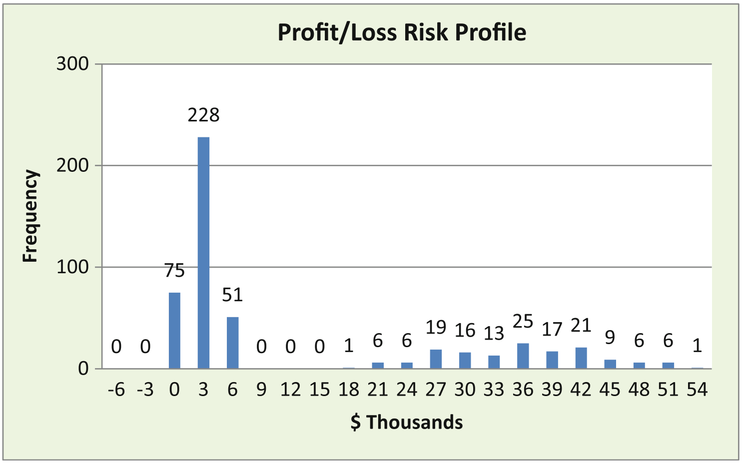

Fig. 8.14 Risk profile for profit/loss statement

occurs, tax rate is zero; if a profit occurs, determine the level and apply the appropriate tax rate. The calculation of net income is the summation of all expenses subtracted from revenue, and results in a profit of $40,894.75.

Some of the results of the model are a bit unexpected. For example, the risk profile suggests a bimodal distribution, one that is related to low or negative profits and has a relatively tight dispersion. The second mode is associated with higher profits and is more widely dispersed. This phenomenon is probably due to the distribution of demand: 70% probability of 10,000 units of demand and 30% of 75,000. In fact, the observations associated with the lower mode represent about 71% of the observations ([75 + 228 + 51]/500 = 70.8%).

8.5 An Operations Example–Autohaus

Let us now put into practice what we have learned about MCS methods on a considerably more complex problem. We begin with item *6a* and *6b* from the *Implementing Monte Carlo Simulation Methods* subsection: develop a complete model definition of the problem (6a) and determine the uncertainty associated with model behavior (6b). Consider an example that models an automobile diagnostic service. Your sister, Inez, has decided to leave her very lucrative management consulting practice in business process design for a nobler calling—providing auto repair services for the automobile driving masses. She has purchased Autohaus, a well-respected auto repair shop located in a small industrial park in her community. She wants to expand the Autohaus service options to operators of large auto fleets and light-duty truck fleets—e.g. the local police department, the US Mail service, and corporate clients with auto fleets. The advantage that she foresees in this business model is a steady and predictable flow of contract demand. Additionally, Inez wants to limit the new services Autohaus provides to *engine/electrical* diagnostics, *mechanical* diagnostics, and *oil change* service. Diagnostic service does not perform repairs, but it does suggest that repair may be necessary, or that maintenance is needed. It provides fleet operators with a preventative maintenance program that can avoid costly failures. Inez wants to perform an analysis for the new services she is contemplating.

The service facility she is planning has three service bays, each staffed by a single mechanic that works exclusively in each bay (see Fig. 8.15). There is a large parking lot associated with the facility that can accommodate a very large number of parked vehicles. It acts as a staging area. Inez has decided that the fleet operators will be required to deliver their autos to the lot every morning before the shop opens at 9:00 am. The autos are parked in the next available slot for the type of service required, and then service is performed on a first-come-first-served basis. Thus, the vehicles will form waiting lines for a particular type of service. Although the hours of operation are 9:00 am to 7:00 pm, if at the end of the day an automobile has service started before 7:00 pm, service will be completed on that auto. Also, the autos need to be removed from the lot by fleet owners by the close of business; Inez wants to avoid the legal liability of insuring that the autos are safe overnight. As we have already stated, there are three types of service that are handled at Autohaus—engine/electrical diagnostics, mechanical diagnostics, and oil change service.

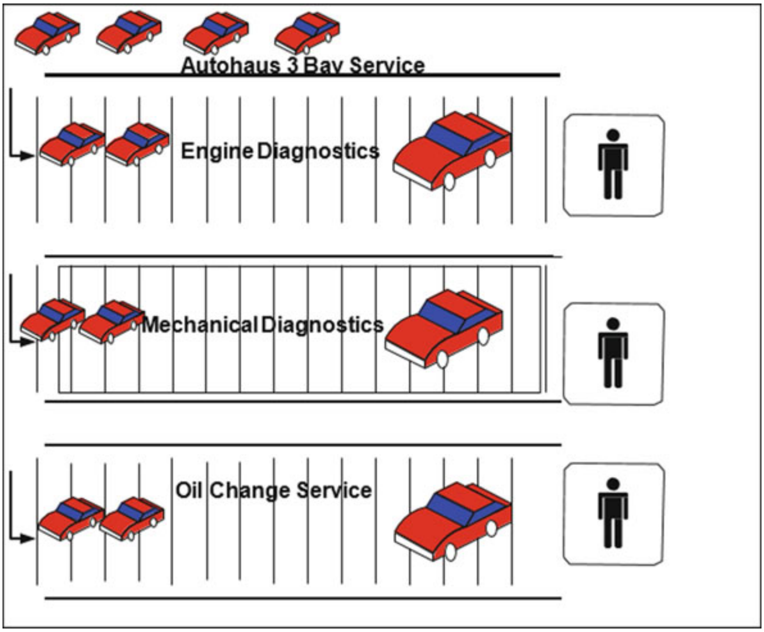

Fig. 8.15 Autohaus model

Wolfgang (Wolfie), her trusted head mechanic, has many years of experience, and he has kept mental records regarding the time that each service demand type requires. As a former modeling and simulation professional, Inez will find her expertise to be quite valuable for constructing a model of Autohaus. It is Inez's goal to understand the general behavior of this new business model in terms of the various types of service demand Autohaus will experience and the operations design that she is contemplating. This is a very reasonable goal if she is to organize the business for the highest possible return on investment. Inez decides that the first step in the assessment of the business model is to simulate the operation. She decides to have a conversation with Wolfgang to better understand the current business and to get his input on important simulation design issues. The following are questions she wants to discuss with Wolfgang to assist in step 6a and 6b of the simulation process:

1. Are the general arrival rates (autos/hour) for the arrival of customers seeking service at Autohaus different for different hours? Wolfgang has noted that the arrival rate for the early time period, 5:00–7:00 am, is different from the later time period, 7:00–9:00 am.
2. She wants to understand the uncertainty associated with the business model.
   (a) The first type of uncertainty is associated with the arrival of autos—1) when and how many autos arrive at the facility, and 2) what type of services do arriving autos request?
   (b) The second type of uncertainty is associated with service provision—1) will the requested service type be available, and 2) what is the service time required for each auto arrival?

After considerable discussion with Wolfie, she arrives at a flow diagram of the process, which is shown in Fig. 8.16. The process begins with autos arriving *prior* to the facility's start of operation at 9:00 am. (This assumption greatly simplifies the operation of the simulation model as we will see in Table 8.2). The process flow elements in diamonds represent uncertainty. The *Autos Arrive with...* process element indicates that autos arrive in unknown quantities and within one of two contiguous time periods, 5:00–7:00 am or 7:00–9:00 am. We distinguish between these two-time periods because they a have different arrival rate behavior. In general, more demand arrives in the early period than the later. This is due to traffic congestion issues becoming more acute later in the morning. Next, the type of service demand that is requested by an auto must also be determined. It will be restricted to one type of service per auto, which is another simplifying assumption.

Once we have information on the number of autos that have arrived and the service type they have requested, the model will direct the auto to one of three **queues** (waiting lines) for each type of service—engine/electrical diagnostics, mechanical diagnostics, and oil change. This occurs in the process element *Assigned to Particular Service Slot*. At this point, we exit the *Arrival Process* boundary and enter the *Service Process*. The next step, although it might appear to be simple, is complex. The bays that are performing various types of service begin to operate on the available queues of autos. This requires that autos requesting a particular type of

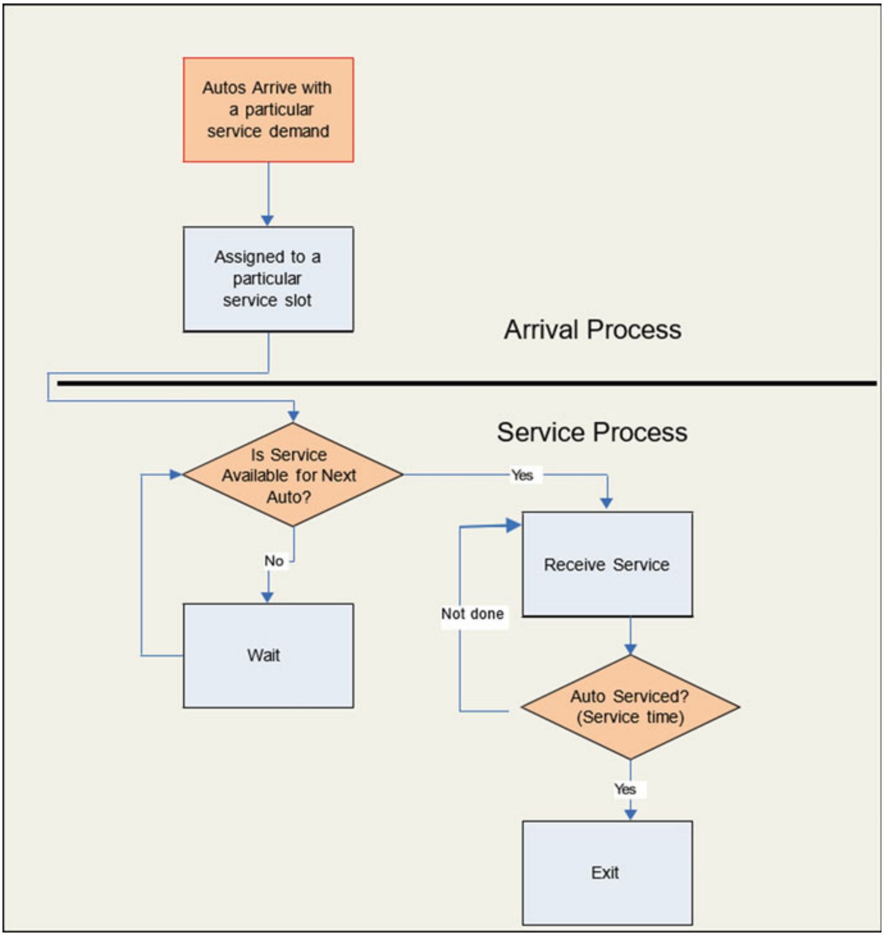

Fig. 8.16 Simple process flow model

service be available (demand) and that service capacity (supply) for the service also be available, as noted in the diamond *Is Service Available for Next Auto?* As the mechanic in the bay administers service, the uncertain service time is eventually resolved for each auto, as shown in the process element *Auto Serviced?* When the simulation starts, the model will be operated for a predetermined period of time, simulating a number of days of operation, during which it will collect data on the servicing of autos.

How will we apply the MCS method to our problem; that is, how do we execute the steps *6c-6f*? Although we have not yet specified every detail of Inez's problem, we now have a general idea of how we will structure the simulation. We can take the two processes described in Fig. 8.16, simulate each, and use the results of the *Arrival Process* as input for the *Service Process*. The steps for determining the remaining details of a model can be arduous, and there are numerous details that still need to be determined. Through interviews with Wolfie and through her own research, Inez must identify, to the best of her ability, the remaining details.

Table 8.2 is a compilation of some detail issues that remain for the model. This table takes us one step closer to the creation of an Excel simulation model by

Table 8.2 Details of Autohaus diagnosis model

Arrival process issue	Resulting model structure
When do autos arrive? 1. Strictly before the official start of service (9:00 am)? 2. All day long (5:00 am-7:00 pm)?	Review of the choices available: 1. This is the simplest choice to deal with as a simulation modeler, but maybe not very customer friendly 2. This is a much more complex modeling situation. *Our choice—1) This assures that the days demand is available prior to starting service*
How will the randomly arriving autos be assigned to service queues?	As an auto arrives, its demand for service must be determined, and only one type of service will be assigned per auto. Autos will be placed in one of 3 queues, each with a particular service demand type (engine/electrical diagnostics, etc.). The autos will be served according to a first-come-first-served service discipline. The distribution of arriving autos will be based on a *Poisson* (more on this later) distribution *Our choice—Three service queues with first-come-first- served service discipline; random arrivals of autos; random service types*
What happens to autos not served on a particular day?	This is also a relatively simple question to answer. They must leave the facility by close of business since they are there for diagnostic service only. This also eliminates the need to track vehicles that already in queues before the morning arrivals *Our choice—Cars are cleared from queues at the end of the day*
Service process issues	**Resulting model structure**
How will service be initialized each day?	A queue discipline of first-come-first-served suggests that this is how service will be administered for the 3 types of service. Thus, we must keep track of vehicles in queues. These queues will also represent the 3 bays where a single mechanic is stationed *Our choice—In accordance to first-come-first served, the auto at the head of the queue will receive service next*
How will service times be determined: Empirical data or subjective opinion?	Empirical data is data that is recorded and retained over time. Subjective opinion comes from experts that supply their opinion. Wolfie has a very good sense of the time required to perform the three service types *Our choice—In the absence of empirical data, we will use Wolfie's subjective (expert) opinion*

detailing the arrival and service processes. In the *Arrival* process we decide that arrival of demand is complete by 9:00 am. Permitting arrivals all day will make the model much more complex. Certainly, model complexity should not suggest how she designs her service, but given that this business model will have *contract* demand as opposed to *drop-in* demand, she has incentives to manage demand to suit her capacity planning. Additionally, this procedure may suit the customer well by providing all customers with a structured approach for requesting service.

In the *Service* portion of the model, the arriving autos will receive a random assignment of a service type. Although the assignment is random, it will be based on an anticipated distribution of service type. Then, service will be administered in a first-come-first-served manner at the corresponding bay, as long as capacity is available. Finally, we will use Wolfie's subjective opinion on the service time distributions, for the three services to consume service capacity at the bays.

We have resolved most of the issues that are necessary to begin constructing an Excel model. As usual, we should consider the basic layout of the workbook. I suggest that we provide for the following predictable worksheets: (1) introduction to the problem or *table of contents*, (2) a *brain* containing important data and parameters, (3) a *calculation* page for the generation of basic uncertain events of the simulation, (4) a *data collection* page for collecting data generated by the logic of the simulation, and (5) a *graph* page to display results.

8.5.1 Status of Autohaus Model

Now, let us take stock of where we are and add the final detail to our model of operation for Autohaus:

1. We will develop a sampling mechanism to determine the random Poisson arrivals of autos. Additionally, we have a similar discrete sampling mechanism for assigning the service type that is requested by each arrival.
2. The relative proportion of the type of service that is assigned to autos will remain constant over time, and it is applied to autos arriving as follows: 40% of autos arrivals will request service type *Engine/electrical Diagnosis*; 25% will request *Mechanical Diagnosis*; and 35% will request *Oil Change*. These service times sum to one—all possible service assignments.
3. We will assume that the portion of the parking lot that is allocated to the new business is sufficiently large to handle any day's demand without causing autos to **balk** (arrive and then leave). By not restricting the capacity of the parking lot, we reduce the complexity of the model, but we also eliminate the possibility of using the model to answer various questions. For example, we may want to consider how we would utilize a limited and costly lot capacity among the various services Inez provides. A *variable* lot capacity is clearly a more sophisticated modeling condition than we are currently considering.

4. Balancing the demand for service with the supply of service will be an important issue for Inez to consider. Customer demand that is greatly in excess of service supply can lead to customers spending inordinate amounts of time waiting, and subsequently lead to the potential loss of customers. Conversely, if demand is far less than supply, then the costs of operation can easily exceed the revenue generated leading to low profits or losses.

As we consider these issues, we can construct more flexible and sophisticated models, but at the cost of greater modeling complexity. The decision of how much complexity is needed should be made in light of Inez's goals for the simulation analysis.

8.5.2 Building the Brain Worksheet

Figure 8.17, the Brain worksheet, shows the results of our discussion of the arrival process. The *Brain* contains all the pertinent parameters and data to supply our *Calculation* and *Data Collection* worksheets. Note that the worksheet has five major categories of information: *Arrival Data-Cumulative Probability Distribution*, *Selection of Arrival Order*, *Type of Service*, *Service Times Distributions*, and *Worker Assumptions*. As in Figs. 8.10, 8.17 contains a table of calculations based on the cumulative Poisson distribution for the three customer categories: Corporate Client,

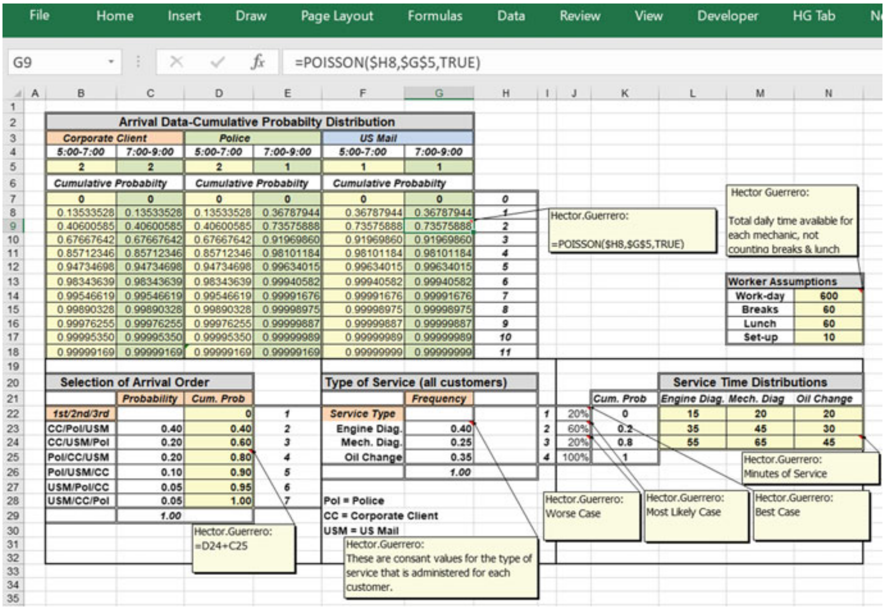

Fig. 8.17 Brain worksheet for model parameters and assumptions

Police, and US Mail. The arrival period is divided into two distinct periods of arrival—5:00–7:00 a.m. and 7:00–9:00 a.m. I have assumed Poisson average *hourly* arrival rates in the two time periods of [2, 2], [2, 1], and [1, 1] for Corporate Clients, Police, and US Mail, respectively. For example, cells B5 and C5 show the average arrival rates for *Corporate Clients* in the 5:00–7:00 period, 2 per hour, and in the 7:00–9:00 period, 2 per hour. The values in the table for the range B7:G18 are the corresponding cumulative Poisson probabilities for the various average arrival rates for all types of clients and times periods. The cell comment for G9 provides the detail for the formula used to determine the cumulative Poisson probability.

The order of customer arrival is also important in our model, and each customer arrives as a group in a period; that is, a caravan of Police autos might arrive at Autohaus in the 5:00–7:00 time period. The *Selection of Arrival Order* table provides the six possible combinations of arrival orders for the three clients, and is the result of some subjective opinion by Wolfgang. For example, there is a 40% chance that the precise order of arrival in a period will be *Corporate Client* first, *Police* second, and *US Mail* third. As you can see, the *Corporate Client* is very likely to be the first to arrive in the morning, with a 60% (40 + 20%) chance of being first overall. Order will be important when Inez begins to examine which of the various clients does not receive service should a day's demand be greater than service supply.

The table entitled *Type of Service* provides the mix or service types for the arrivals. Notice it is deterministic: a fixed or non-probabilistic value. Thus, it is not necessary to resolve uncertainty for this model factor. If 20 autos arrive in a period of time, 8 (20 $*$ 0.4) will be assigned to *Engine/electrical Diagnosis*, 5 (20 $*$ 0.25) to *Mechanical Diagnosis*, and 7 (20 $*$ 0.35) to *Oil Change*. If Inez anticipated a great deal of variation in the short-term service types, then it might be wise to determine a distribution for the service types that can be sampled, as we did for arrival order. Service could also be seasonal, with certain types of service, for example engine/electrical diagnosis, occurring more frequently in a particular season. Our model handles service type quite simply and, certainly, more sophisticated models could be constructed.

Service Time Distributions are shown next to the *Selection of Arrival Order* table. The table suggests that service times are distributed with three discrete outcomes—20% best case (shortest service time), 60% most likely case, and 20% worse case (longest service time). For the *Oil Change* service, the times are 20, 30, and 45 min, respectively, for the best, most likely, and worse case. This information is also the type of subjective information that could be derived from an interview with Wolfie. The information gathering could be as simple as asking Wolfie to make the following determination: "If we define *Best Case* occurring 20% of the time, what value of oil change service time is appropriate for this case". Similarly, times for worse case and most likely can be determined. These estimates can be quite accurate for a knowledgeable individual, but great care must be taken in the process of eliciting these subjective values. The interviewer must make sure that the interviewee is fully aware of what is meant by each question asked. There are many excellent structured techniques that provide interviewers a process for arriving at subjective probability values.

Finally, the *Worker Assumptions* relate to the employee policies that Inez will set for the work force. As stated in Table 8.2, three mechanics, each in a single bay, provide service, and it is assumed that they are available from 9:00 am to 7:00 pm (600 total minutes), with a 1 h lunch break, four 15 min breaks, and a changeover from one auto to another requiring 10 min of set-up time. The first three times are obviously a matter of workforce policy, but the set-up could be uncertain. This could include, placing the auto into the bay, selection of appropriate tools, and any number of other activities associated with the preparation for service of the next auto available. Set-up might also be dependent on the worker, the type of service, or other factors. It is possible that more experienced workers might require less set-up time than less experienced workers. Given the variety for types of automobile service, the total daily set-up time could be quite substantial, and this is time that Inez might want to reduce through special tools and standardization of work.

We have assumed a deterministic time for set-up to simplify our model. Also note that these numbers can be manually changed by the modeler to perform *what-if* or *sensitivity analysis*. For example, what if we could introduce training or equipment that might reduce set-up time significantly? We would want to know if the investment in the time reduction is worth the effort, or analyze competing technological changes to determine their cost/benefit. Having all these parameters on a single worksheet, the *Brain*, is beneficial to the modeler performing the analysis.

8.5.3 Building the Calculation Worksheet

Figure 8.18 provides a view of the calculation worksheet which simulates 250 days (sample size n = 250) of auto arrivals. We will use the 250 days of randomly selected arrivals, along with their arrival order, to determine what autos the mechanics can service each day and which services they will provide. The Calculation worksheet will be used to determine some of the fundamental calculations necessary for the model. Of particular interest are the daily totals in column N, which represent demand. Summary statistics for these calculations are shown at the bottom of the column. As you can see in the cell comment for B254, two VLOOKUP()s with a lookup value argument of RAND() determines the hourly demand in *each* of the 2 h. The cell formula uses the *Arrival Data* table in the *Brain*, B7:H18, to randomly select *two* 1-h arrivals. (Recall that each period is 2 h in length and the arrival rate is hourly, thus the need for two lookups.) It is possible to calculate one hour's arrivals and multiply by two; unfortunately, this will exaggerate a single hour's behavior and reduce the variation of that naturally occurs in each hour.

There are three categories of demand (Corporate Client, Police, and US Mail) and hourly arrivals for each: columns B, D, F, H, J, and L. The arrival order is calculated for each client in a time period in the adjacent columns: C, E, G, I, K, and M; thus, in the Day 1 row (A5:P5), we find that in the 5:00–7:00 time period the number of arrivals for the Corporate Client is 6 (B5) and they are the first (C5) to arrive. For the US Mail arrivals in the 7:00–9:00 time horizon, there are 4 arrivals (L5) and they are

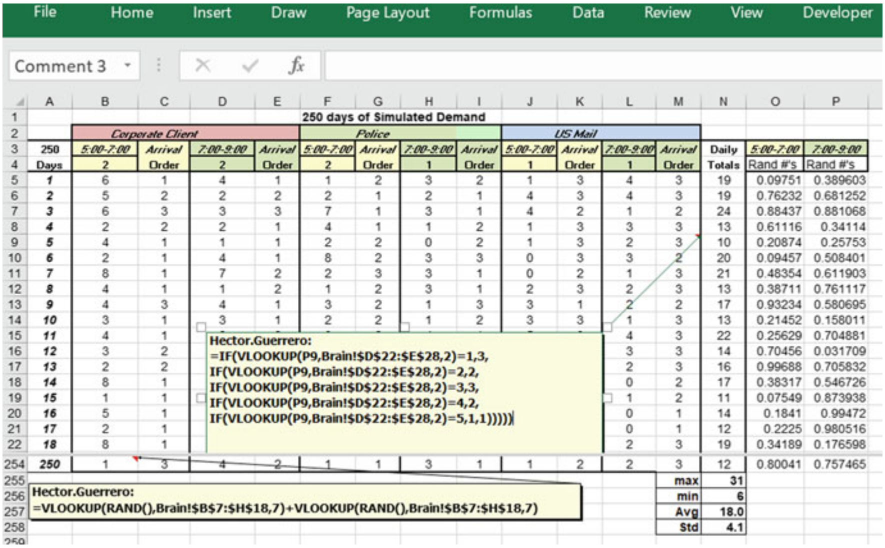

	A	B	C	D	E	F	G	H	I	J	K	L	M	N	O	P
1						250 days of Simulated Demand										
2		Corporate Client				Police				US Mail						
3	250	5:00-7:00	Arrival	7:00-9:00	Arrival	5:00-7:00	Arrival	7:00-9:00	Arrival	5:00-7:00	Arrival	7:00-9:00	Arrival	Daily	5:00-7:00	7:00-9:00
4	Days	2	Order	2	Order	2	Order	1	Order	1	Order	1	Order	Totals	Rand #'s	Rand #'s
5	1	6	1	4	1	1	2	3	2	1	3	4	3	19	0.09751	0.389603
6	2	5	2	2	2	2	1	2	1	4	3	4	3	19	0.76232	0.681252
7	3	6	3	3	3	7	1	3	1	4	2	1	2	24	0.88437	0.881068
8	4	2	2	2	1	4	1	1	2	1	3	3	3	13	0.61116	0.34114
9	5	4	1	1	1	2	2	0	2	1	3	2	3	10	0.20874	0.25753
10	6	2	1	4	1	8	2	3	3	0	3	3	2	20	0.09457	0.508401
11	7	8	1	7	2	2	3	3	1	0	2	1	3	21	0.48354	0.611903
12	8	4	1	1	2	1	2	3	1	2	3	2	3	13	0.38711	0.761117
13	9	4	3	4	1	3	2	1	3	3	1	2	2	17	0.93234	0.580695
14	10	3	1	3	1	2	2	1	2	3	3	1	3	13	0.21452	0.158011
15	11	4	1									4	3	22	0.25629	0.704881
16	12	3	2									3	3	14	0.70456	0.031709
17	13	2	2									2	3	16	0.99688	0.705832
18	14	8	1									0	2	17	0.38317	0.546726
19	15	1	1									1	2	11	0.07549	0.873938
20	16	5	1									0	1	14	0.1841	0.99472
21	17	2	1									0	1	12	0.2225	0.980516
22	18	8	1									2	3	19	0.34189	0.176598
254	250	1	3	4	2	1	1	3	1	1	2	2	3	12	0.80041	0.757465
255													max	31		
256													min	6		
257													Avg	18.0		
258													Std	4.1		

Fig. 8.18 Calculation worksheet for 250 days of operation

the third (M5) arrival. Thus, the number and sequence of arrivals for the first simulated day is:

(a) 5:00–7:00: Corporate Client = 6. . ..Police = 1. . .. US Mail = 1 . . ..subtotal 8
(b) 7:00–9:00: Corporate Client = 4 . . ..Police = 3. . .. US Mail = 4. . ..subtotal 11
(c) Thus, Total Day 1 demand =19

The logic used to determine the sequence of arrivals is shown in the cell M9 comment of Fig. 8.18. The formula consists of an IF, with four additional nested IFs, for a total of five. In the first condition, a random number is referenced in cell P9, as well as all of the nested IFs. This is done to make the comparisons consistent. Remember that every RAND() placed in any cell is independent of all other RAND ()s, and different values will be returned for each. The single random number in P9 insures that a single sequence (e.g. CC/Pol/USM) is selected, and then the IFs sort out the arrival sequence for each client type for that sequence. In this particular case, the P9 value is 0.25753. This number will be used in a VLOOKUP() in the *Brain*. See the *Selection of Arrival Order* table in the *Brain* (Fig. 8.17). The random number 0.25753 is compared to the values in D22:D28. Since the value falls between 0.0 (D22) and 0.4 (D23), the value 1 is returned in accordance to the lookup procedure. The value 1 indicates a sequence of CC/Pol/UMC; thus, the condition identifies the Corporate Client (CC) as the first position in the sequence. Cell I9 returns a 2 for the Police client (Pol) since it is second in the USM/Pol/CC sequence, etc. Although this may appear to be very complex logic, if you consider the overall structure of the cell formula, you can see that the logic is consistent for each of the six possible sequences. Finally, the calculation of *Daily Totals* is performed in column N by summing all arrival values. For Day 1, the sum of cells B5, D5, F5, H5, J5, and L5 is 19 arrivals in (N5).

8.5.4 Variation in Approaches to Poisson Arrivals: Consideration of Modeling Accuracy

Let us consider for a moment other possible options available to generate arrivals from the Poisson distributions of hourly arrivals. For the 7:00–9:00 time period, I have chosen a rather direct approach by selecting two randomly sampled values, one for the hour spanned in 7:00–8:00 and the other for 8:00–9:00. Another approach, which we briefly mentioned above, is to select a single hourly value and use it for both hours. This is equivalent to multiplying the single sample value by two. Does it really matter which approach you select? The answer is yes, it certainly does matter. The latter approach will have several important results. First, the totals will all be multiples of 2, due to the multiplication; thus, odd values for arrivals, for examples 17 or 23, will not occur. Secondly, and related to the first, the standard deviation of the arrivals in the latter approach will be greater than that of the former approach.

What are the implications of these results? For the case of no odd values, this may not be a serious matter if numbers are relatively large, but this may also be a departure from reality that a modeler may not want to accommodate. In fact, the average for the arrivals sampled for several days will be similar for both approaches, as long as the sample size of days is large, for example 250 days. The second outcome is more problematic. If there is a need for the model analysis to study the variability of arrival outcomes, the latter approach has introduced variability that may not be acceptable or representative of *real* behavior. By accentuating (multiplying by two) extreme values, it is possible that larger than realistic extreme values will be recorded. In our case, this *is* an important issue, since we want to study the possible failure to meet extreme demand for daily arrivals.

To demonstrate the differences discussed above, consider the graph in Fig. 8.19. In this simple example, the difference between the approaches becomes clear. The worksheet in Fig. 8.19 contains two areas (different colors) with different approaches for simulating the arrivals in a 2-h period. The range (A1:J25) has two separate VLOOKUP functions sampling a Poisson with an average arrival rate of one unit per hour. Why two VLOOKUPs? We have two because one is needed for each hour in the 2-h period. This is the approach we used in Fig. 8.18. The approach in range K1:T25 is a single sample which is then multiplied by two. Why a single sample? In this case we assume that the same value can be used for each of the 2 h. Both approaches collect 250 samples of arrivals for a 2-h period.

The summary statistics in show very similar behavior in terms of means and total arrivals for the 250 samples, a mean of approximately 1.9, and total arrivals of 464 and 473, respectively. Thus, both approaches appear to be similar in some summary statistics. The difference is seen in the standard deviation for the two

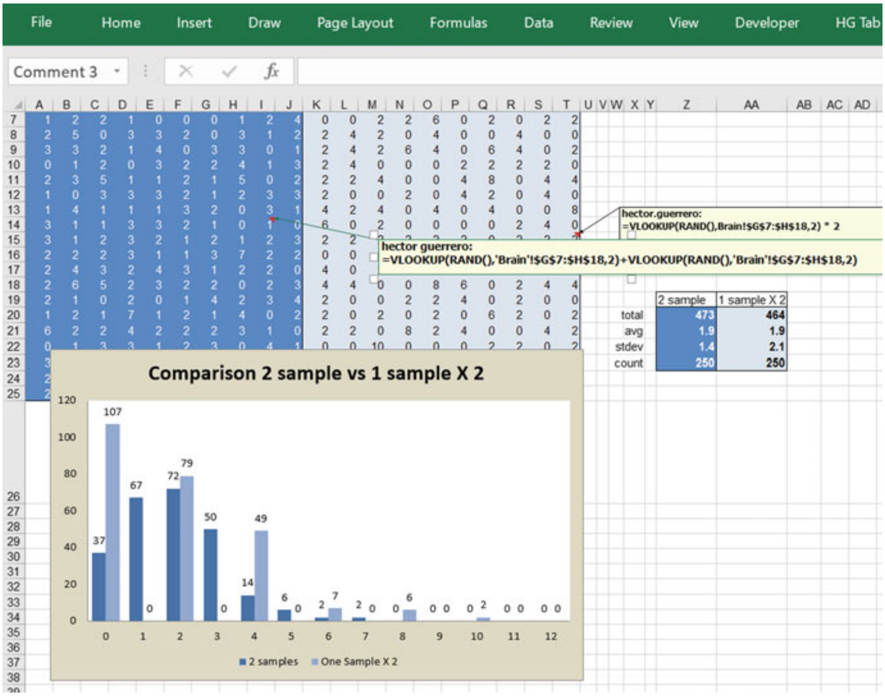

Fig. 8.19 Two approaches to arrivals

VLOOKUP approaches. The approach with two randomly sampled hours has a lower value, 1.4, than the approach which used a single VLOOKUP multiplied by two, 2.1. That is a 50% increase for the single sample times two approach. The graph in Fig. 8.19 also shows greater dispersion of results. It is far more likely that this approach will suggest a greater service capacity stock-out than the former, as evidenced by how the one sample graph extends far beyond the two sample. Additionally, the one sample graph has many more incidences of 0 arrivals. Thus, the distortion occurs for both extremes, high and low arrival values. Thus, the modeler must carefully consider the *reality* of the model before committing a sampling approach to a worksheet. The one sample approach may be fine, and certainly involves fewer excel functions to calculate, but it may also lead to exaggerated results.

8.5.5 *Sufficient Sample Size*

How many simulated days are sufficient? Obviously, simulating 1 day and basing all our analysis on that single sample would be foolish, especially when we can see the high degree of variation of *Daily Totals* that is possible in Fig. 8.18—a maximum of 28 and minimum of 5. In order to examine the model behavior carefully and accurately, I have simulated a substantial number of representative[1] days. Even with a sample size of 250 days, another 250 days can be recalculated by depressing the F9 key. By recalculating and observing the changes in summary statistics, you can determine if values are relatively stable or if variation is too great for comfort. Also, there are more formal techniques for calculating an appropriate sample size for specific **confidence intervals** for summary statistics, like the mean. Confidence intervals provide some level of assurance that a sample statistic is indicative of the true value of the population parameter that the statistic attempts to forecast.

Without going any further, notice the substantial utility offered by our simple determination of demand over the 250-workday year in Fig. 8.18. The worksheet provides excellent high-level summary statistics that Inez can use for planning capacity. There is also a substantial difference between the minimum and maximum Daily Totals, 6 and 31. Thus, planning for peak loads will not be a simple matter given that the costs of servicing such demand will certainly include human capital and capital equipment investments. Also of great interest are the average number of daily arrivals, 18.0, and the standard deviation of arrivals, 4.1. An average of 18.0 is a stable value for the model, even after recalculating the 250 days many times. But what is the variation of daily arrivals about that average? There are several ways we can answer this question. First, we can use the average and the standard deviation to calculate a coefficient of variation of approximately 22.8% (standard deviation/mean = (4.1/18.0). It is difficult to make a general statement related to the variation of demand, but we can be relatively certain that demand varying one standard deviation above and below the mean, 13.9–22.1, will include the majority of the 250 daily arrivals generated. In fact, for this particular sample, we can see in Fig. 8.20 that the number of simulated days of arrivals between 14 and 22 is 192 (24 + 44 + 36 + 51 + 37) of the 250 days of arrivals, or approximately 77%.

8.5.6 Building the Data Collection Worksheet

There is still an important element missing in our understanding of daily demand: we are unaware of the *types* of service that each of the arrivals will request and the related times for service. Although the analysis, thus far, has been very useful, the types of service requested could have a substantial impact on the demand for service time. What remains is the assignment of specific types of service to arrivals, and the

[1]We want to select a sample of days large enough to produce the diverse behavior that is possible from model operation. I have used 250 because it is a good approximation of the number of work days available in a year if weekends are not counted. After simulating many 250-day periods, I determine that the changes in the summary statistics (mean, max, min, and standard deviation) do not appear to vary significantly; thus, I feel confident that I am capturing a diversity of model behavior. If you are in doubt, increase the simulated days until you feel secure in the results.

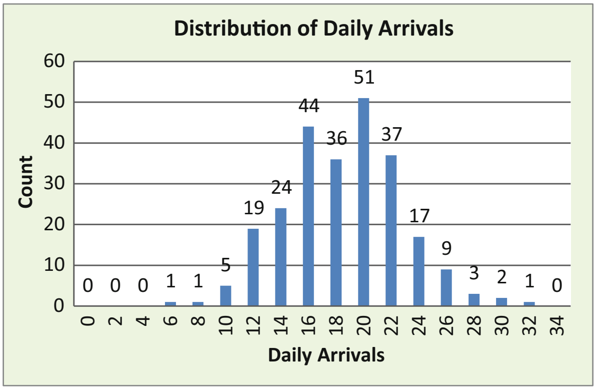

Fig. 8.20 Distribution of daily arrivals for 250 simulated days

subsequent determination of service times. Both of these issues need to be resolved to determine the amount of daily demand that is serviced.

Should the calculations take place in the *Calculation* worksheet or the *Data Collection* worksheet? I suggest that when you are in doubt you should consider two factors: (a) is one worksheet already crowded and, (b) is the calculation important to collaboration, such that you would want it to appear on the same worksheet as other calculations. The location of some model processes is often a judgment call. In our case, the *Calculation* worksheet has been used to generate daily arrivals and the sequence of arrival. Thus, we will use the *Data Collection* worksheet to deal with service. This division of calculations will make the workbook more manageable to those using the workbook for analysis.

Now, let us consider the delivery of service for the model. There is an important issue to be aware of when considering service: simply because there is demand for service, this does not imply that the Autohaus service system can process all the demand that is requested. The simulation has resulted in approximately 4500 (18.0 $*$ 250) autos seeking service during the 250-day period. Our goal now is to determine whether the operations configuration that Inez is contemplating can, in fact, manage this service demand. Depending on the availability of service capacity, the Autohaus service system may handle the demand easily, or it may be considerably short of the requested demand.

In our model, demand represents a *reservoir* of autos requesting service. To this point, we have generated demand without regard to Autohaus' capability to supply service. The question that Inez needs to answer is: given the demand that we have generated, how much of this demand can our service supply accommodate on a daily basis? Our simulation should be useful in a number of ways. Determining the amount of demand serviced will also depend on the contractual agreement that she has with customers. If she is contractually guaranteeing service over a period of time, then she may be forced to increase capacity to satisfy all demand or run the risk of a penalty for not fulfilling the contract. This should certainly be a concern for Inez as she begins to conceptualize her business model and as she finalizes plans for operation. The use of the simulation model to study service capacity will also help her understand the effects of various types of contractual agreements.

So how do we proceed? Consider Fig. 8.21, the *Data Collection* worksheet. This is the worksheet where we perform the service analysis we just discussed. It is wise to first understand the broad layout of Fig. 8.21, especially given the complex nature of the worksheet. This worksheet examines demand for each of the 250 days simulated in the *Calculation* worksheet (Fig. 8.18) and returns the service that needs to be provided. The *Data Collection* worksheet contains three major sections which we will detail in several Figures, due to their rather large size.

The first section, columns A:E, determines the percentage of each day's total demand that will be allocated to each of the three service types. In this example, day 1 demand is 17 arrivals (E4), and of this total, 7 (B4) are allocated to *Engine/electrical Diagnostics*, 4 (C4) to *Mechanical Diagnostics*, and 6 (D4) to *Oil Change*. Columns F to AX display the specific service arrivals and the service time required for each arrival. You can see in cell range F4:L4 the 7 *Engine/electrical Diagnostics* arrivals and their corresponding service times, 35, 55, 35, 35, 35, 35 and 35, respectively. Finally, Columns AY:BF determine the availability of service, providing the modeler with an opportunity to detect a **service stock-out**[2]; that is, a day for which some demand is not satisfied by available service capacity. If a stock-out condition is observed, it is because the mechanic stationed in a bay is unable to service all daily demand for that particular service type. In summary, the *Data Collection Worksheet* has: (1) an initial allocation of overall demand to specific service types (*Engine/electrical Diagnosis*, etc.), (2) a determination of actual times in minutes for the service types, and lastly, (3) a comparison of the daily service time requested versus the daily service capacity available from a mechanic.

Now, let us take a closer look at the *Data Collection* worksheet. Figure 8.22 focuses on areas (1) and (2) in the summary above. As mentioned, columns B through D allocate the total daily demand over the three service categories. Cell

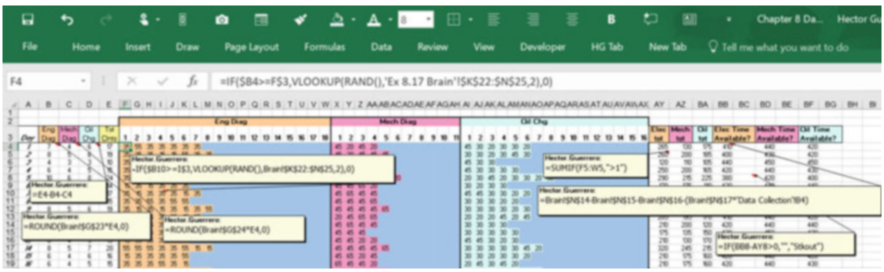

Fig. 8.21 Data collection worksheet

[2]A service stock-out is defined as a daily period for which there is demand for service that is not met. It does not suggest the size of the stock-out. It is simply a stock-out occasion, without regard to the number of autos that do not receive service.

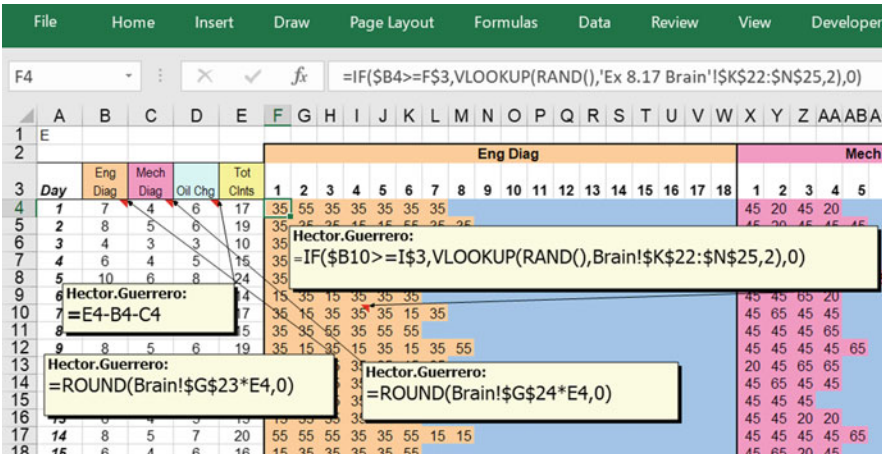

Fig. 8.22 Allocation of service and service time determination

E4 references the *Calculation* worksheet cell N5, the *Daily Total* for day 1. Cell range B4:D4 allocates the total, 11, on the basis of the deterministic percentages located in the *Brain* in the *Type of Service* table. Recall that *Engine/electrical Diagnosis* is constant at 40% of the arrivals. Therefore, 11 arrivals (E4) multiplied by 40% (G23 in the *Brain*) produces approximately 4 *Engine/electrical Diagnosis* service requests. This number is rounded to account for integer values for arrivals. Similarly, 3 *Mechanical Diagnosis* service requests are produced. Then *Oil Changes* are calculated as the difference of the *Total Clients* and the rounded *Engine/electrical* and *Mechanical Diagnostics* service requests (17-7-4 = 6) to ensure that the sum of all service types requested is equal to the total, 17.

Next, column range F4:W4 generates 4 service times. Logic is needed to select only as many service times as there are arrivals. This is done by comparing the number of service requests to the index numbers in range F3:W3 (1–18). Note that 18 columns are provided. This insures that a very large number of random arrivals are possible. It is very unlikely to ever experience 18 arrivals for any service type. For *Mechanical Diagnostics*, a maximum of 11 arrivals are permitted, and for *Oil Change*, a maximum of 16 are permitted. The cell function used to determine the service time value of *Engine/electrical Diagnostics* is a logical IF() with an embedded VLOOKUP. An example is shown in cell I10 of Fig. 8.22. The IF() tests if B10 is greater than or equal to the index number in its column, 4. If this condition is true, which it is, then the cell formula in cell I10 will randomly sample the *Service Time Distributions* table in the *Brain* via a VLOOKUP and return a value.

If the index exceeds the total number of service requests, a 0 is returned. Note that cell M10 appears blank, although it is 0. This is accomplished by using conditional formatting for the cells. A logical test of the cell content—equal to 0—is made, and if the answer is to the test is true, then the cell *and* font colors are set to similar colors, thus resulting in a blank cell appearance. This is done for clarity, and zeros could easily be allowed to appear without any loss of accuracy. The same is repeated in all other ranges for the various service types. By depressing the recalculate[3] key, F9, the contents of the three service areas change as new service arrivals are calculated in the *Calculation* worksheet.

Now, let us consider the analysis that is done in the *Data Collection* worksheet. In column range AY:BF of Fig. 8.23, we perform the calculations necessary to determine if daily service demand is satisfied by available service capacity. The time available for service is the difference between the total time available from 9:00 am to 7:00 pm, 600 min, and breaks, lunch, and set-ups. Each row determines the sum of service time requested in a particular day for each service type (e.g. AY4:BA4) then compares the request to available service time. For example, for day 1 the available *Engine/electrical Diagnosis* time is 420 min, which is calculated by subtracting

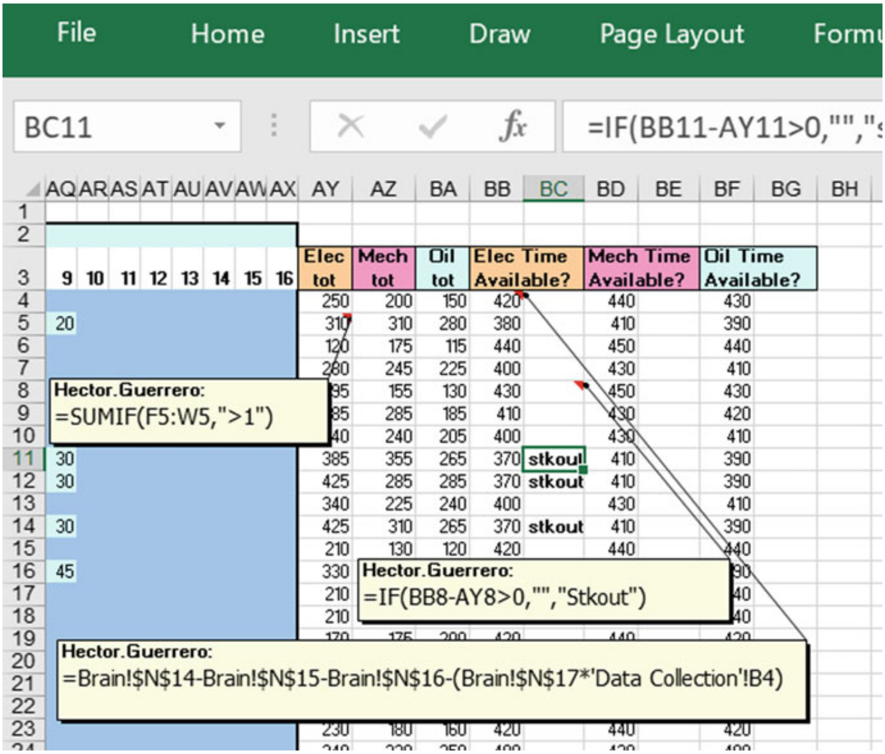

Fig. 8.23 Analysis area of data collection

---

[3]I would suggest setting the recalculation of formulas to manual in all worksheets. This gives you control and eliminates annoying and inopportune recalculations. The control can be set by using the Tools then Options menus. One of the available tabs is Recalculation and this is where you select manual.

breaks (60), lunch (60) and set-up times for each service request ($6 * 10 = 60$) from a total of 600 min. These times are found in the *Worker Assumptions* table in the *Brain*. Thus, daily breaks and lunch always account for 120 min, while the set-ups sum will vary according to the number of service requests, 6 in this case requiring 10 min each. The total is 180 min, and this results in 420 min of available service capacity ($600-180 = 420$). Since 440 min of capacity is greater than 160 min of *Engine/electrical Diagnosis* requested, there is sufficient capacity to deal with demand, thus no stock-out occurs for that service type. In row 11 of Fig. 8.23, which corresponds to day 8 of the simulation, we see that a service stock-out occurs for *Engine/electrical Diagnosis*. The comment associated with the cell BC11 shows an IF() cell formula that compares available capacity of 370 min to a capacity demand of 385 min. In this case, service demand is higher that available service time. Obviously, there is a service stock-out of 15 min, in which case the test results in a negative value; thus, the cell value s*tkout* is returned. The same is also true for row 12 and 14. There may be more detailed analysis that is possible for each stock-out. For example, is it possible to determine who will and will not receive service?

In Fig. 8.24, we see that cell BC254 counts the number of *Engine/electrical Diagnosis* stock-outs, 11. This is done by using a simple COUNTIF() function: COUNTIF(BC4:BC253,"Stkout"). We could easily count the number of stock-out minutes in column BC rather than simply determining that a stock-out occurred. This would allow us to quantify the number of minutes that we are *under* demand capacity, in each service area for 250 days.

8.5.7 Results

Inez has completed her analysis, and we are now prepared to discuss some of the results. The major focus of the analysis is on the use of capacity and the level of service that is provided to clients. So, what are the questions that Inez might ask about the simulation results? Here are a number that she might consider important:

1. How does the demand for the various services differ?
2. Will Autohaus be capable of meeting the anticipated demand for services?
3. Are there obvious surpluses or shortages of capacity? Given the revenue generation capability of each service, do the model results suggest a reconfiguration of client services?
4. Can we be sure that the simulation model has sufficient sample size (250 days) to provide confidence in the results?

The first question does not have an obvious answer from the raw data available in the *Brain*. Although the *Engine/electrical Diagnosis* represents the highest percentage of service (40%), it has shorter service times for worst, best, and most likely cases than *Mechanical Diagnosis*. The comparison to *Oil Change* also is not clear. Luckily, the summary statistics near the bottom of Fig. 8.24 show very clearly that *Engine/electrical Diagnosis* dominates demand by a substantial margin. The

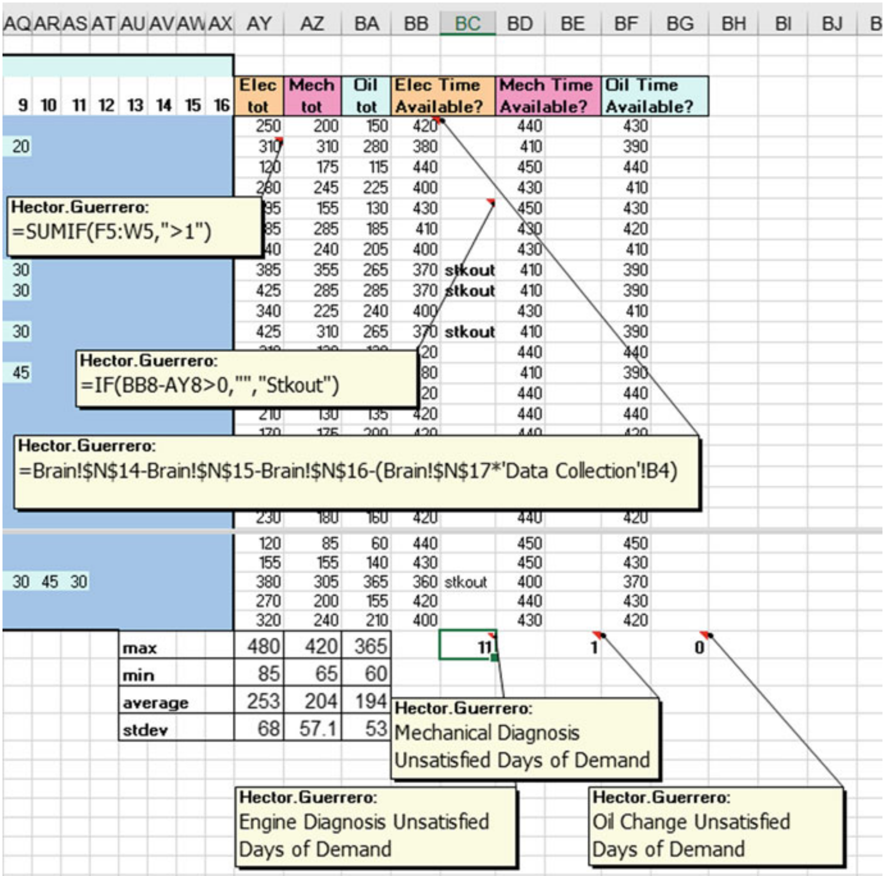

Fig. 8.24 Determination of clients not receiving service

averages for demand are 253, 204, and 194 min, respectively, for *Engine/electrical*, *Mechanical*, and *Oil*. Thus, the average for *Engine/electrical* is about 24% ([253–204]/204) greater than the average for *Mechanical*, and 30% ([253–194]/194) greater than *Oil*.

What about the variation of the service demand time? Figure 8.25 shows that the distribution for Engine/electrical Diagnosis is more widely dispersed than that of Mechanical Diagnosis or Oil Change. This is verified by the summary statistics in Fig. 8.24, where the annual (250 days) standard deviations range from 68.0 for Engine/electrical, to 57.1 for Mechanical, and 53.0 for Oil. Additionally, the range of values, max–min, for Engine/electrical (480–85 = 395) is substantially greater that Mechanical (420–65 = 355) or Oil (365–60 = 305). All this evidence indicates that Engine/electrical has much more volatile demand that the other services.

The answer to the second question has already been discussed. Again, we see in Fig. 8.24 that the only area where there appears to be any significant demand that is

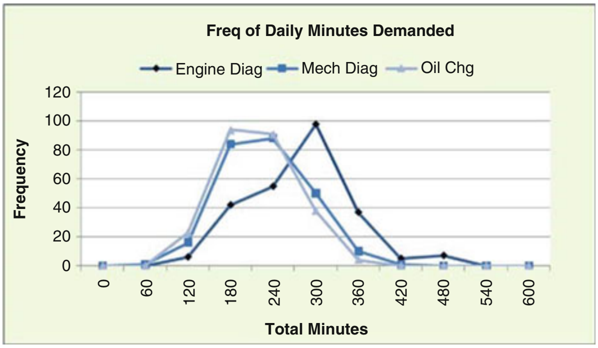

Fig. 8.25 Graph of frequency of daily service minutes demanded

not being met is for *Engine/electrical Diagnosis*. The simulation shows that there were 11 days of unsatisfied service. We need to be quite careful to understand what this suggests for capacity planning at Autohaus. It does not mean that demand was not met for *all* the demand that occurred for a service type for 11 days; it does mean that some *portion* of the *Engine/electrical Diagnosis* demand for 11 days was not met. Although the model did not indicate the amount of service time that was not met, as we mentioned earlier, it is a simple matter to change the cell formulas in columns BC, BE, and BG to return the quantity rather than the indicator term *Stkout*. The *Stkout* in Fig. 8.23 cell BC11 would then be −15 min (370–385 = −15). By comparing this time (−15) to the time requested by the last demand in the queue, we can determine if the auto receives service because it starts before the end of the day, or if the auto is not serviced because it has not yet started. Recall that if an arrival is being serviced at the end of the day, it will remain in service until it is finished. This would then not represent a *Stkout*, but would represent overtime to accommodate finishing the arrival.

Question three considers the appropriate use of capacity. Inez must consider the service stock-out for *Engine/electrical Diagnosis*. We have already suggested that some indications of a service stock-out are resolved by the *end-of-day* policy. She must also consider the cost associated with attempting to meet demand every day for every service type. A policy of 100% service could be a very costly goal. It may be far better for her to use overtime to handle excess demand, and also to accept an occasional service stock-out. Much will depend on alternative uses of the capacity she establishes for this new service. She may be able to accommodate other clients to handle the unused available capacity, in particular for *Mechanical* and *Oil* services, since they appear to have substantial excess capacity. Additionally, she could consider making the three bays more flexible by permitting multiple services in a bay. This, of course, would depend of the capital investment necessary to make the flexibility available. These questions will require further analysis and the potential adaptation of the model to perform different types of modeling. It is not unusual for the results of a model to suggest the need to change the model to capture other types of operational behavior.

Finally, question 4, regarding the sample size, is always relevant in Monte Carlo simulation. It asks whether our simulation has a sufficiently large sample of model operation (250 days in our case) to accurately describe behavior. As a modeler gains experience, it becomes clear that relatively few sources of variation are sufficient to greatly affect results, even when data from a relatively large sample is collected. So, how do we deal with this question? The obvious answer to this question is to repeat the experiments to produce results multiple times. In other words, simulate 250 days, simulate another 250 days, and repeat the process numerous times. Then examine the results of each simulation relative to the combined results of the simulations. This is easily done in our model by recalculating the spreadsheet and noting the summary statistics changes with each recalculation.

In Fig. 8.26 you can see *Stkout* data collected in 20 replications of the 250-day experiment. Notice the range of *Stkout* values for *Engine/electrical Diagnosis* in the first column. Values range from 2 to 15 and average 9.9, while the range for *Mechanical* and *Oil* are far less variable. The standard deviation for *Engine/electrical Diagnosis* is approximately five times greater than *Mechanical.* A value of 2 for *Engine/electrical Diagnosis* appears to be a rare occurrence, but this example serves to suggest that some data elements will require substantial sample size to insure a representative results, and it may be necessary to increase sample size to deal with variation in *Engine/electrical Diagnosis*.

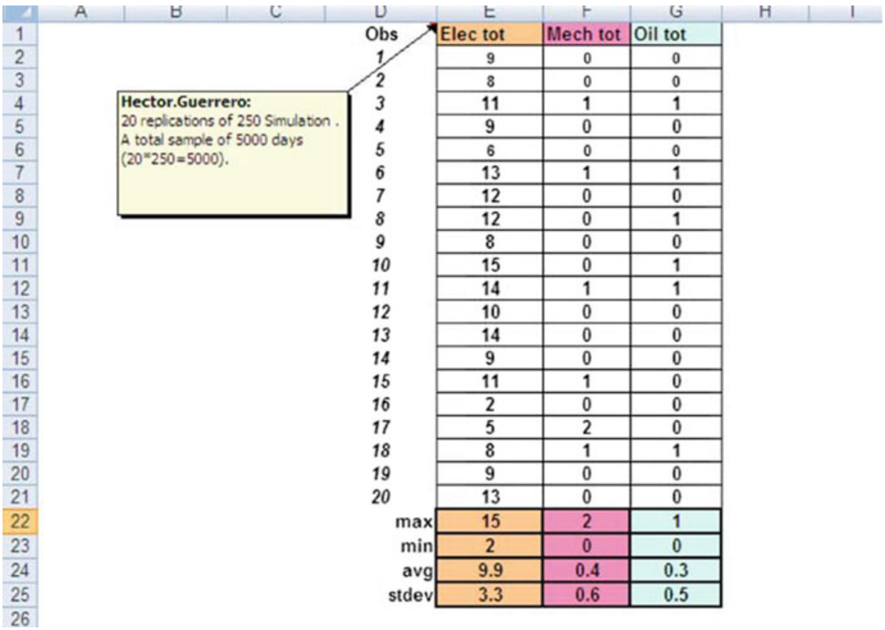

Obs	Elec tot	Mech tot	Oil tot
1	9	0	0
2	8	0	0
3	11	1	1
4	9	0	0
5	6	0	0
6	13	1	1
7	12	0	0
8	12	0	1
9	8	0	0
10	15	0	1
11	14	1	1
12	10	0	0
13	14	0	0
14	9	0	0
15	11	1	0
16	2	0	0
17	5	2	0
18	8	1	1
19	9	0	0
20	13	0	0
max	15	2	1
min	2	0	0
avg	9.9	0.4	0.3
stdev	3.3	0.6	0.5

Fig. 8.26 Multiple replications of 250 days simulation stock-out data

8.6 Summary

As we have seen in both examples, a relatively few sources of uncertainty are sufficient to produce results that are not easily predicted. Thus, the interactions of the uncertainties of a model are often unexpected and difficult to forecast. This is precisely why simulation is such a valuable business tool. It provides a systematic approach to understanding and revealing the complex interactions of uncertainty in models.

The value of a carefully conceptualized and implemented simulation can be great. Beyond the obvious benefits of providing insight into the risks associated with the problem, the process of creating a model can lead to far greater understanding of the problem. Additionally, the process of creating a simulation requires a rigorous approach and commitment to very specific steps, not least of which is problem definition. Even the cleverest simulation and most sophisticated analysis is worthless, if you are solving the "wrong" problem.

It is important to keep in mind that the goal of simulation is to determine the risk associated with a problem or decision; that is, what is the range of possible outcomes under the model conditions? One of the steps of simulation that we did not perform was *sensitivity analysis*. By changing the parameters of the problem, we can note the reaction of the model to changes. For example, suppose you find that a relatively small decrease in the set-up times used in Autohaus can lead to significant improvements in service stock-outs. You would be wise to carefully investigate the nature of step-ups and how you might be able to reduce them. This could result in great operational improvement in service.

In the next chapter, we cover a number of tools that are extremely powerful and useful. They are available in the Data Ribbon—Solver, Scenarios, and Goal Seek. Some of these tools can be used in conjunction with simulation. The modeling and simulation we performed in Chaps. 7 and 8 has been **descriptive modeling**—it has provided a method by which we can describe behavior. In Chap. 9 we will introduce **prescriptive modeling**, where we are interested in prescribing what decisions *should be made* to achieve some stated goal. Both are important, and often work together in decision making.

Key Terms

Model	Risk profile
Simulation	RAND()
Point estimates	Random sampling
Monte Carlo simulation	Resolution of uncertain events
Rapid prototyping	Uniform distribution
Continuous event simulation	Census
Discrete event simulation	Replications

(continued)

Model	Risk profile
Events	Discrete distributions
Deterministic values	Continuous distributions
NORMINV()	Average arrival rate
Empirical data	Poisson arrival process
Balk	POISSON(x, mean, cumulative)
Confidence intervals	VLOOKUP
Coefficient of variation	HLOOKUP
Service stock-out	Normal distribution
Descriptive modeling	Bell curve
Prescriptive modeling	Queues

Problems and Exercises

1. Name three categories of models and give an example of each.
2. Which of the following are best modeled by Discrete event or Continuous event simulation:

   (a) The flow of water through a large city's utility system
   (b) The arrival of Blue-birds at my backyard birdfeeder
   (c) The number of customers making deposits in their checking accounts at a drive-up bank window on Thursdays
   (d) The flow of Euros into and out of Germany's treasury
   (e) The change in cholesterol level in a person's body over a 24 h period
   (f) The cubic meter loss of polar ice cap over a 10 year time horizon.

3. Monte Carlo simulation is an appropriate modeling choice when point estimates are not sufficient to determine system behavior-T or F?
4. Give three reasons why rapid-prototyping may be useful in modeling complex business problems.
5. Risk profiles always have monetary value on the horizontal axis-T or F?
6. Create a risk profile for the following uncertain situations:

   (a) A $1 investment in a lottery ticket that may return $1,000,000
   (b) A restaurateur's estimate of daily patron traffic through a restaurant where she believes there is 30% chance of 25 patrons, 50% chance of 40 patrons, and 20% chance of 75 patrons
   (c) A skydiver's estimate of success or failure under particularly treacherous weather conditions, where the skydiver has *no* idea of the outcome (success or failure).

7. Create a simple simulation that models the toss of a fair coin. Test the results (% Heads/% Tails) for sample sizes of 5, 10, 30, and 100. Hint-Use the RAND() function.

8. Two uncertain events are related. The first event occurs and effects the second. The first event has a 35% chance of an outcome we will call *small*, and 65% chance of a *large* outcome. If the first outcome is *small* then the second event will result in equal chances of 3, 4, 5, and 6, as outcomes; if the first event is *large* then the second event has equal chances of 11, 13, 14, and 15, as outcomes. Create a simulation that provides a risk profile of outcomes. The simulation should replicate the experiment a minimum of 300 times.
9. Create a VLOOKUP that:

   (a) Allows a user to enter a percent (0–100%) and returns a categorical value based on the following data:

0–30%	31–63%	64–79%	80–92%	93–100%
A	B	C	D	E

   (b) For the same data above, create a VLOOKUP that returns a categorical value for a *randomly* generated %. Hint-Use the RAND() function.
   (c) Expand the table so that the category A and B is defined as *Good*, C as *OK*, and D and E as *Terrible*. With this new, three row table, return the new outcomes (*Good*, etc.) for exercise (a) and (b) above.

10. Create a simulation of a simple income statement for the data elements shown below—produce a risk profile and determine statistics summary data for profit (average, max, min, standard deviation, etc.):

    (a) Revenue Normally distributed mean of $140 k and standard deviation of $26 k
    (b) COGS are a % of Revenue, with outcomes of 24, 35, and 43%, all equally likely
    (c) Variable costs % of Revenue, with outcome 35% one half as likely as outcome 45%
    (d) Fixed costs of $20 k (constant).

11. The arrival of email at your home email address can be categorized as *family*, *junk*, and *friends*. Each has an arrival rate: family—5/day; junk—4/day; friends—3/day.

    (a) Create a simulation that provides you with a profile of total, daily email arrival. What is the probability that you will receive 0 emails, 3 emails or less, 15 or more emails, and between 5 and 15 emails. Hint-Use the Poisson () function and replicate for 100 days.
    (b) If you read every email you receive, and the time to read the categories is that shown below, what is the distribution of minutes spent per day reading email.

        (i) Family email—Normal distribution, mean 5 min and standard deviation 2
        (ii) Junk email—Discrete 0.5 min 80% and 3 min 20%
        (iii) Friends—Continuous Uniform from 3.5 to 6.5 min.

12. ***An advanced problem***—Father Guido Aschbach, pastor of Our Lady of Perpetual Sorrow Church, is planning a weekly (Sunday after mass) charity event, complete with food and games of chance. He has heard of Fr. Efia's success, and he would like to try his hand at raising money through games of chance. Each parishioner is charged $10 to enter the event and is given three tokens for each of the three games of chance: Wheel of Destiny, Bowl of Treachery, and the Omnipotent 2-Sided Die. See the table below for the odds and payouts of each game. The parishioners enjoy these events, but they are affected by the weather that occurs each Sunday—Rain produces the lowest turnout; Fair weather the largest turnout; Glorious sunshine results in the next largest turnout (see details below). Help Father Aschbach estimate the yearly (52 Sundays) returns for the events. Hint-You will have to decide whether, or not, you want to simulate the play of each individual parishioner, or if you will simply use the expected value of a player as we did before. If you can, and it will be much more difficult, attempt to simulate the play of each parishioner in the simulation.

Wheel of Destiny		Bowl of Treachery		Omnipotent 2-sided Die	
Type—Discrete		Type—Discrete		Type—Discrete	
Return	Prob.	Return	Prob.	Return	Prob.
$10*	0.4	$15	0.3	$100	0.35
–$10	0.6	–$20	0.7	–$40	0.65

*40% chance a parishioner will win $10

Weather Attendance/Probability		
Rain	20	0.2
Fair	50**	0.55
Glorious	35	0.25

**55% chance 50 parishioners will attend