
CHAPTER 1

Case Study 1—Uses of the Standard Deviation

In this and all the other case studies in this book, you will come across statistical terms that will be unfamiliar to you. Rather than explain each term in the body of the case study, or populate each case study with explanatory foot- or endnotes, we have provided a glossary at the end of the book to facilitate your understanding of the text.

The Steps of Data Analysis

An important step in data analysis, and hence forecasting, is to determine the shape, center, and spread of your data set. That is, how are the data distributed (normally, uniformly, or left- or right-skewed), what is its mean and median, and to what degree is it dispersed about the mean? The answers to these questions are important, as they will determine what statistical tools we can use as we attempt to determine lost profits. The questions can be answered graphically with histograms, stem and leaf plots, and box plots as well as computationally through the use of various Excel descriptive statistical functions such as AVERAGE, STDEV, SKEW, and KURT. Knowing a data set's distribution (normal, or at least near-bell shaped), its center (mean), and its spread (standard deviation), we can begin to draw conclusions about it that will allow us to meaningfully compare it with other contemporaneous data points.

Our first case, while not strictly on the subject of lost profits, is about an issue our readers see all the time in their practices. The case concerns the value of a business on June 30, 2010, for shareholder buyout purposes. The subject company's year-to-date (YTD) gross margin of 14.4 percent at the June 30 date of loss was substantially below its nine-year historical average. We want to know if this difference is statistically significant and if so, what the probability of its occurrence is. In order to run the tests to help answer these questions, we need to assure ourselves that the distribution of the nine years'¹ gross margins is, if not normal in shape, at least near-bell shaped. We begin this process with the following calculations and charts.

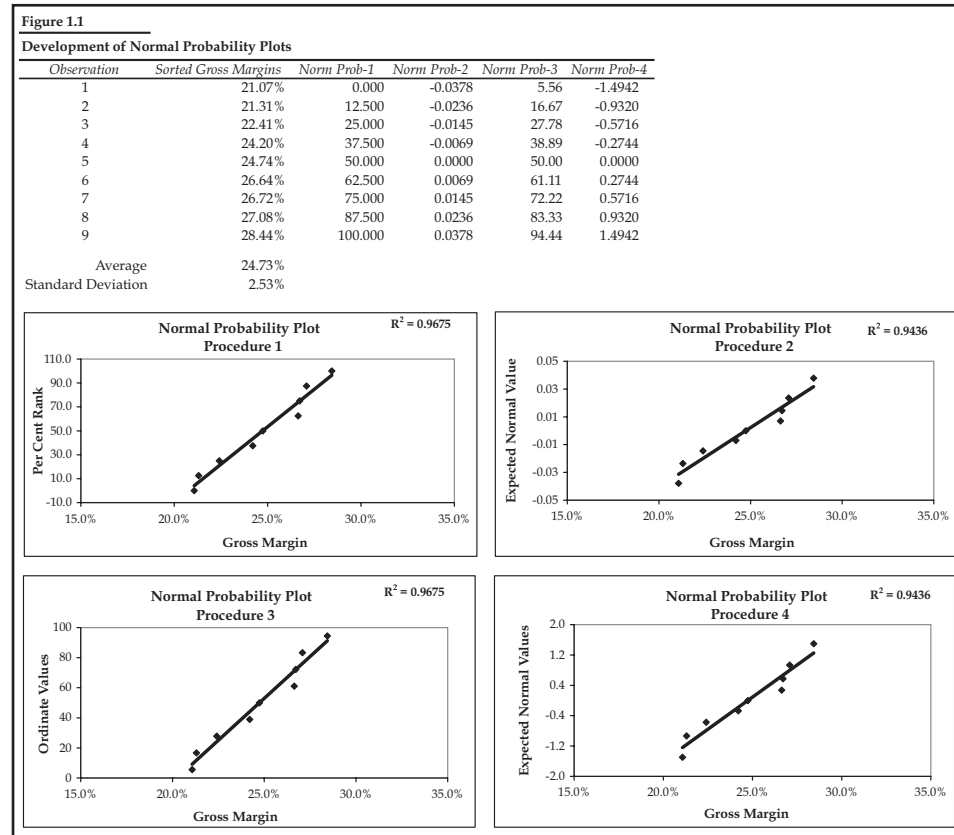
	B	C	D	E
2	Table 1.1			
3		Year	Gross Margin	
4		2001	26.6%	
5		2002	21.3%	
6		2003	26.7%	
7		2004	27.1%	
8		2005	28.4%	
9		2006	22.4%	
10		2007	24.7%	
11		2008	21.1%	
12		2009	24.2%	
13		June 30, 2010	14.4%	
14				
15	Mean		24.7%	
16	Population Standard Deviation		2.5%	
17	Kurtosis		-1.5	
18	Skewness		-0.2	
19	Median		24.7%	
20	Test statistic (14.4% - 24.7%/2.5)		-4.1	
21	p-value		0.35%	

Shape

As shown in Table 1.1, the nine-year average gross margin is 24.7 percent; its standard deviation is 2.5 percent (we are treating the nine years as the population, not a sample, so the Excel function is STDEV.P); it is symmetrical,² as its median of 24.7 percent equals its mean; and it is near-bell shaped, as both KURT and SKEW return results that are between -1.5 and $+1.5$. While there are too few data points to construct a meaningful histogram to informally test for normality, we can create a normal P-plot, or probability plot, where we match up our nine observations with the normal scores that we would expect to see if the data came from a standard normal distribution. If our nine observations follow a general straight line on the normal probability plot, especially at both ends, we can feel assured that the data are near-bell shaped.

Figure 1.1 demonstrates four methods for producing the expected values for a P-plot, with no one method being superior to the others. The reader can choose any one method he or she prefers. Since two of the methods require that the data be ranked, we first sorted the gross margins from smallest to largest, and then applied the particular formula for each method and charted the results.

For ease of understanding, we have also shown the formulas for each method in Figure 1.2. But a nagging question remains: If all the observations do not lie directly on the trendline, how do we know that the data are still near-bellshaped? The answer to that question is shown in Figure 1.3, where we have substituted nine numbers created by using Excel's random number generator, which is found in the Analysis ToolPak. Selecting normal distribution from the analysis tool's drop-down menu, and then filling in the rest of the dialogue box with the mean and standard

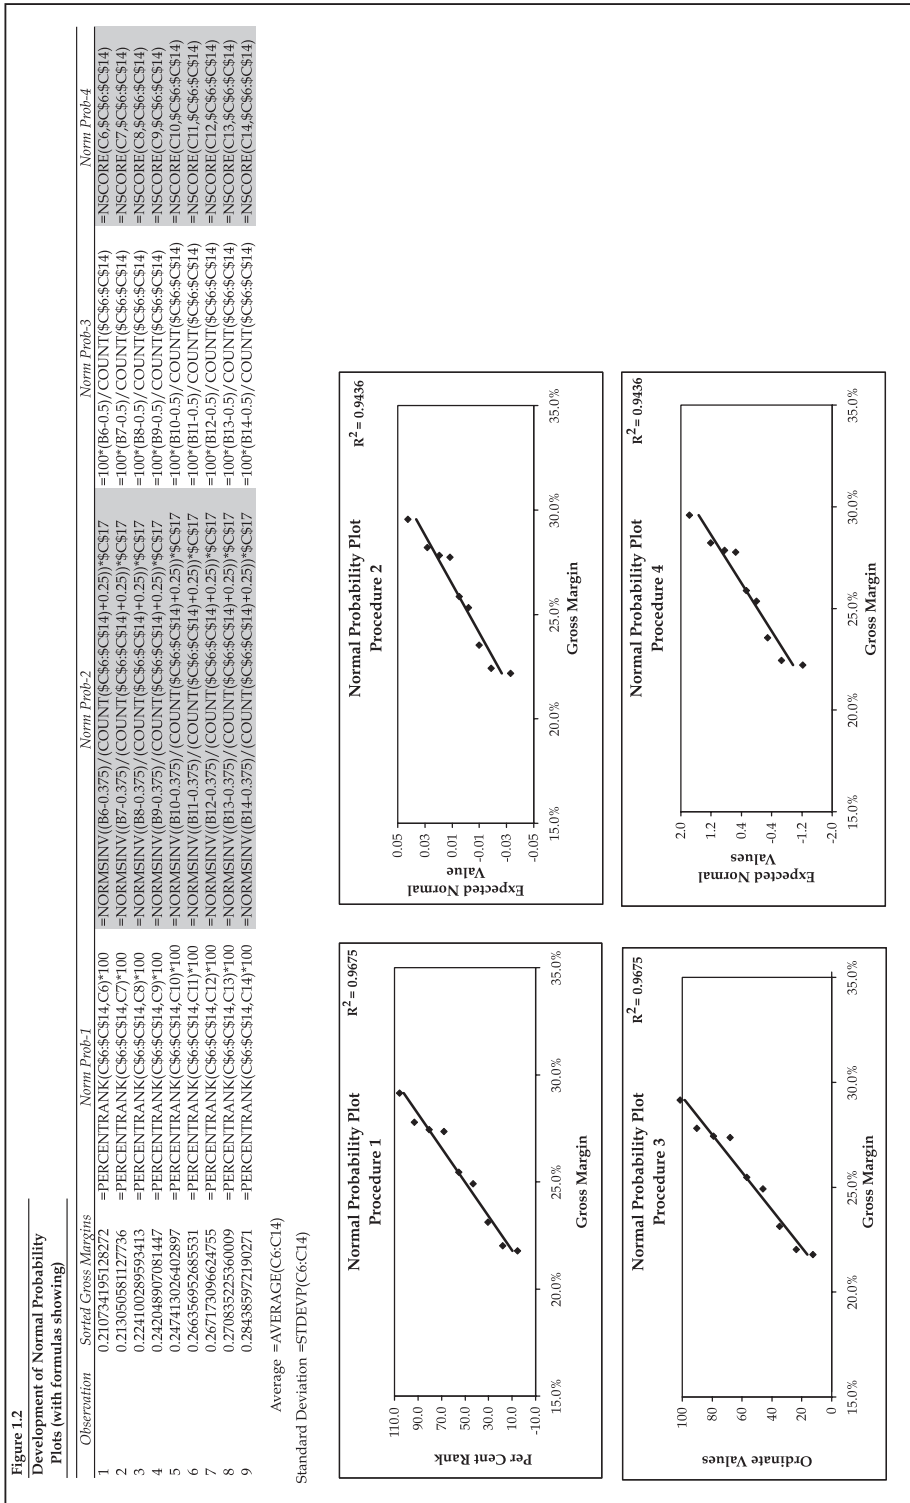


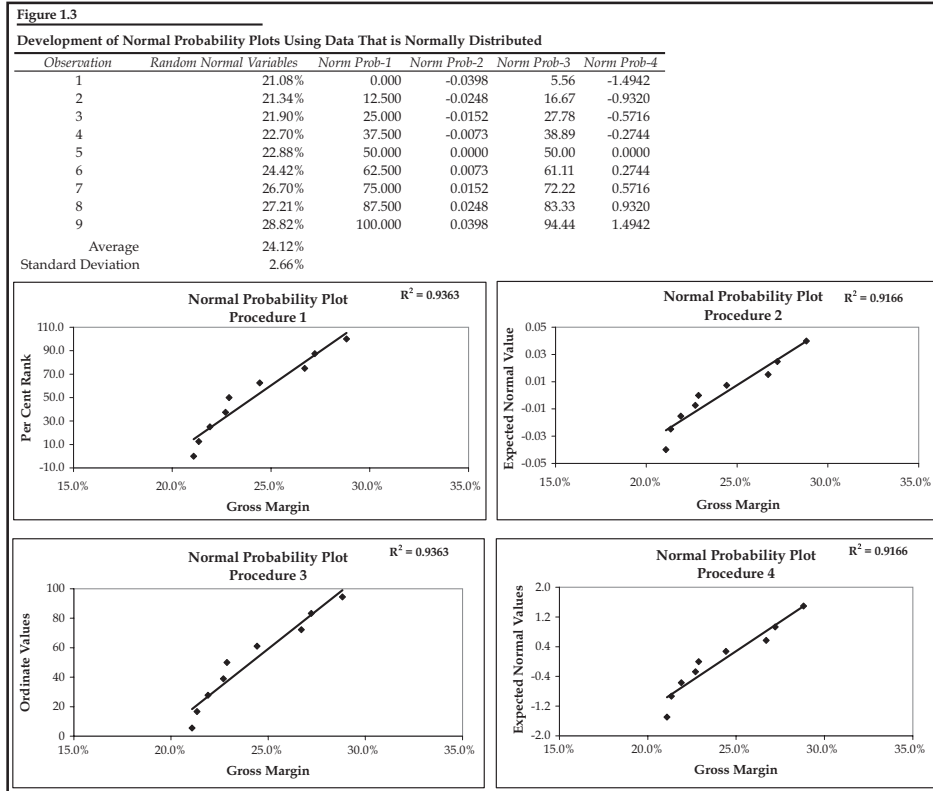
deviation of the original nine observations, produces nine random numbers that are normally distributed with approximately the same mean and standard deviation as the original nine gross margin numbers.

Comparing the two sets of charts in Figure 1.1 and Figure 1.3, we can see that both data sets have the same degree of deviation about the trendline, indicating that the original data set can be considered near-bell shaped. This near-bell shape of the distribution allows us to use parametric tests that involve the use of the two parameters of the distribution, its mean and standard deviation. Absent this shape, we would have to turn to nonparametric methods, a subject we leave to another book.

Spread

The next step is to create a test statistic that measures how far June 30's 14.4 percent is from the nine-year average gross margin. The test statistic is: $(X - \text{mean})/\text{standard deviation}$, or $(14.4 - 24.7)/2.5$. The resulting statistic, called a z-score, of -4.1 indicates that 14.4 percent is 4.1 standard deviations below the nine-year average. Is this too far from the average for our purposes? We know from the empirical rule that if a data set is near-bell shaped, then approximately 95 percent of its observations will fall within ± 2 standard deviations from the mean, and 99.7 percent of its observations





will fall within ± 3 standard deviations from the mean. At 4.1 standard deviations below average, June 30's gross margin is literally off the charts. In fact, using Excel's NORMDIST function, we can show that there is only a 0.35 percent chance that the gross margin on June 30 of 14.4 percent was drawn from the same pool as the nine year-end gross margins.

Once this, and its implication that they were deliberately cooking the books to lower the value of the company, was pointed out to the other side, they immediately increased their buyout offer by 50 percent.

Our second situation involves a car dealership that experienced a small fire in its parts department in early August 2010, and then claimed lost sales for weeks afterwards. Are the decreases in parts and service departmental sales during the month of August 2010 greater than 1.645 or more standard deviations of the historical average, and therefore statistically significant, or are the decreases in the normal course of business? Our procedure is essentially the same as that just described in the first situation. We begin by analyzing the previous six and succeeding two months of sales and find that both data sets are near-bell shaped—the median approximates the mean, and kurtosis and skewness measures are between -1.5 and $+1.5$ as shown in Table 1.2. Unlike situation 1, where the question posed was whether June 30's sales were different in either direction, that is, greater or lesser, from the nine-year average, in situation 2 we only want to know if August's sales

	A	B	C	D	E	F	G
1	Table 1.2						
2	TEST OF ABNORMALLY BELOW AVERAGE SALES						
3	PARTS AND SERVICE DEPARTMENTS						
4				Sales			
5				Parts		Service	
6		Feb-10		84,195		84,010	
7		Mar-10		96,404		94,483	
8		Apr-10		85,947		101,376	
9		May-10		83,634		92,268	
10		Jun-10		78,590		91,981	
11		Jul-10		84,336		97,762	
12		Aug-10		57,845		75,511	
13		Sep-10		92,091		87,610	
14		Oct-10		103,750		100,907	
15							
16		Monthly Average, Feb-Jul, Sept-Oct		88,618		93,800	
17		Population Standard Deviation		7,682		5,743	
18		Median		85,142		93,376	
19		Kurtosis		0.21		(0.86)	
20		Skewness		0.92		(0.27)	
21		Critical value of z - 95%, one-tail test		1.645		1.645	
22		+ - range		12,636		9,447	
23							
24		95% Confidence Level - no less than		75,983		84,353	
25							
26		August sales fall into the 5% rejection range					
27							
28		Test Statistic:					
29		Average Sales		88,618		93,800	
30		August Sales		57,845		75,511	
31		Delta		30,773		18,289	
32		Critical value of z - 95%, one-tail test		1.645		1.645	
33		No. of Std Devs from Average		4.006		3.184	
34		Greater than 1.645?		Yes		Yes	
35		Unusual Event?		Yes		Yes	
36							

are less than the average—statistically speaking, situation 1 called for a two-tailed test and situation 2 calls for a one-tailed test.

The cutoff point for a one-tailed test is 1.645 standard deviations if we are willing to be wrong about our conclusions 5 percent of the time. That is, we enter the 5 percent rejection area at 1.645 standard deviations, rather than the 1.96 standard deviations of a two-tailed test. Multiplying the population standard deviation by 1.645 and subtracting the product from the eight-month average sales for both departments gives us the lower limit of our 95 percent confidence level. For example, in the parts department, average sales equaled \$88,618, and the standard deviation was \$7,682. Multiplying \$7,682 by 1.645 gives us \$12,636. Subtracting this from average sales

gives us \$75,983, an amount greater than August's sales of \$57,845. Since actual sales for August are less than \$75,983, we can say with 95 percent confidence that those sales did not decrease in the ordinary course of business—that there was some intervention that caused them to be this far below average. As such, after subtracting avoided costs, the insured's claim for lost profits was honored.

Conclusion

In this chapter we demonstrated how the standard deviation can be applied in situations where the data set you are working with is the population, and not a sample. How does one know whether the data set is a sample or a population? The answer lies in whether or not you make inferences outside the data set. For example, a classroom of 30 students is a population if all your statistical tests and inferences are about the 30 students. But if you are going to use the statistical test results of those 30 students to make inferences about other students in other classrooms, then you are working with a sample.

In the next chapter we introduce you to various data analysis techniques that should precede any selection of a sales forecasting methodology.

Notes

1. Generally speaking, in statistics, the larger the sample size, the better the results. However, in damages cases, the analyst has to work with what is given. Therefore, while for academic research a sample size of 9 would probably be considered too small, for litigation purposes it will have to do.
2. If folded over at the midpoint, its left side would be a mirror image of its right side.

