CHAPTER 1

Profit from Accurate Forecasting

1.1 THE IMPORTANCE OF DEMAND FORECASTING

Forecasts of demand for products and services can be crucial to the operations of most companies. Inventory planning, logistics planning, production scheduling, cash flow planning, decisions on staffing levels, and purchasing decisions can all depend on forecasts. Making these forecasts perform as well as possible will lead to improved customer service levels and so foster customer goodwill and retention. It will also lower costs. There will be less need for expensive emergency production runs, and there should be a reduction in the waste associated with excessive stock levels and unsold products.

Figures for the cost reductions or increased profits that companies achieve through improved forecasting can be hard to come by – most organizations don't publish them. However, one forecasting software company (www.catchbull.com) estimates that avoidable forecast errors can add between 2% and 4% to costs of production. They quote the case of one $15 billion firm where executives estimated that "we can drive up to $200 m of avoidable costs out of the business." A survey carried out by a Triple Point Technology in 2013 indicated that reductions in inventory levels resulting from improved forecast accuracy meant that a company with a $1 billion turnover could expect savings of between $5 million and $10 million. However, the same survey found that 40% of respondents admitted that they were "not currently leveraging the advantages of statistical modeling in their demand planning operations." Of course, it's in the interests of software companies to advertise these huge benefits, but common sense suggests that better-performing forecasts will significantly benefit a company's bottom line.

1.2 WHEN IS A FORECAST NOT A FORECAST?

A forecast is an honest statement of what we expect to happen at a future data, based on the information available at the time when we make the forecast. It is not necessarily what we hope will happen, so it is not the same as a target. In fact, in some circumstances we may be doubtful that a target will be achieved – we simply created it to motivate people to try to get as close to it as possible. Nor is it the same as a plan.

A plan is what we intend to happen, assuming the future is under our control. As we shall see, ideally a forecast will acknowledge that we are uncertain about the future and provide a measure of that uncertainty.

It is also important to distinguish a forecast from a decision. A decision is what we choose to do in the light of a forecast. We may have a demand *forecast* for next week of 2,000 units, so we *decide* to hold 2,200 units in stock at the start of the week in case demand exceeds the forecast. The 2,200 units is not a forecast – it's a decision.

Sometimes people are tempted to look at the possible demand levels that may occur at a future period and decide which one will best suit their interests. For example, I might forecast a demand for next month of 3,500 units, knowing that it will please senior managers and gain me some kudos, even though I don't truly expect this demand level to be achieved. Although I might call this a forecast, in reality I'm making a decision.

1.3 WAYS OF PRESENTING FORECASTS

1.3.1 Forecasts as Probability Distributions

We can present forecasts in several different ways. A probability distribution indicates the possible levels of demand and their associated probabilities. Table 1.1 is an example. Figure 1.1 displays the distribution. It shows that relatively low levels of demand are more probable than very high demands, so the distribution is skewed. Forecasts in this form are useful because they show the risks of

Table 1.1 A Probability Distribution of Demand

Next Month's Demand (Units)	Probability (%)
20 to 29	5
30 to 39	30
40 to 49	41
50 to 59	10
60 to 69	8
70 to 79	5
80 to 89	1

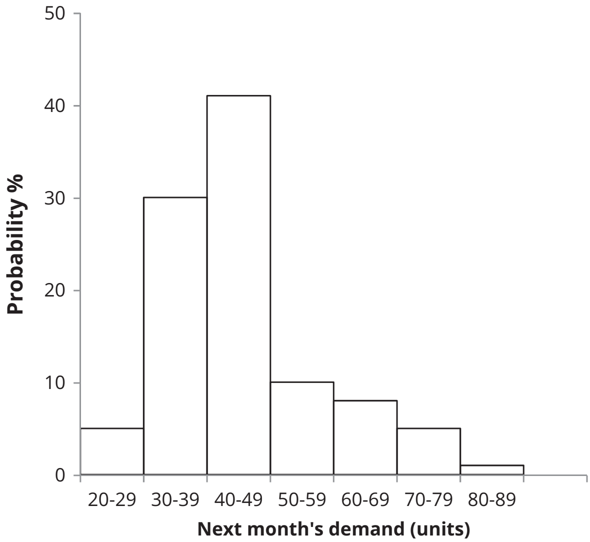

Figure 1.1 A graphical display of the probability distribution

particular decisions we may make. For example, if we decide to hold 69 units of stock at the start of the month, there will be a $5\% + 1\% = 6\%$ probability that we will be unable to meet demand and will disappoint customers.

Accurately estimating probability distributions can be difficult, particularly if we have limited past data. Usually, it is assumed that a particular distribution applies and, most commonly, this is the bell-shaped normal or Gaussian distribution. Figure 1.2 shows an example. Notice that the distribution is symmetrical about its highest point. While there are theoretical reasons why a normal distribution will apply, in many circumstances it can at best only offer a rough approximation to the probabilities of future demand. When the “true” distribution is highly skewed, the approximation will be very poor.

1.3.2 Point Forecasts

Most software products don’t currently display full probability distributions (sometimes these are called density forecasts). Instead, they produce point forecasts and prediction intervals. A point forecast is a

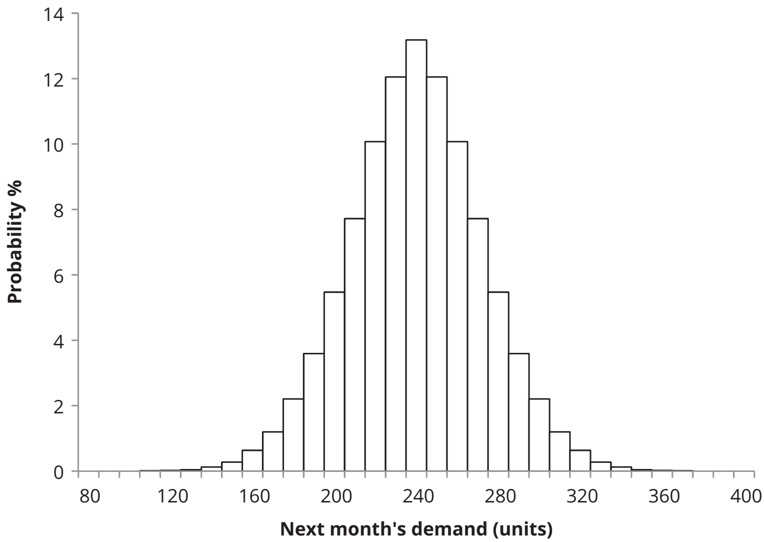

Figure 1.2 A normal distribution of demand

forecast expressed as a single number. It usually represents the mean (or average) of the probability distribution. Imagine if the month referred to in Figure 1.2 was repeated many times. On some occasions, we see demand greater than 300 units. On very rare occasions, it would exceed 350 units. In other months, demand might be well below 200 units. More often than not, it would be between 200 and 280. If we averaged all of the demands, we observed we would find that the mean demand was 240 units. This would be our point forecast. If we did the same for the month referred to in Figure 1.1, we would find that the mean demand was 45 units – slightly to the left of the range of possible demands because of the skewness in the distribution. Therefore, a point forecast produced by software is simply an average of all the possible levels of demand – taking into account their probabilities. It is not a statement that we think that that specific level of demand *will* occur – we know that the actual demand is likely to stray from its value as shown by the probability distributions. You might tell me the mean height of American males aged over 21, but I don't expect every American male I meet in this age group to be that height. In fact, meeting someone who conforms exactly to the mean would be rare.

This point is worth emphasising. I have heard of cases of senior managers who expect point forecasts always to be "100% accurate" and criticize forecasters who are not achieving this. Their attitude shows a fundamental misunderstanding of what a point forecast is. In most forecasting situations, there are bound to be random or unpredictable factors that cause the actual demand to stray from the average represented by the point forecast. In particularly unpredictable situations, such as when we forecast a long way ahead, we should not be surprised if the demand strays a significant distance from the point forecasts. However, as we will see, if we try to anticipate these random factors, we will be wasting our time and probably damaging the forecasts to boot.

1.3.3 Prediction Intervals

Point forecasts don't tell us anything about the level of uncertainty associated with a forecast. We can't tell how far actual demand might stray from the forecast, and you generally need this information to plan inventory levels. However, some idea of the level of uncertainty can be obtained from a prediction interval. A prediction interval is a range that has a stated probability of capturing the actual demand. For example, if we have the distribution shown in Figure 1.2, our software would produce a 95% prediction interval for next month's demand of 177 to 303 units. This means that there's a 95% chance that the actual demand will be captured within this range, and therefore, a 5% chance that the demand will be outside it.

The 95% is sometimes known as the coverage probability. Higher coverage probabilities and more uncertainty both lead to wider prediction intervals. Because of the greater uncertainty, prediction intervals therefore tend to be wider the further ahead you are forecasting. In Chapter 7, we will see how prediction intervals can be used to determine safety stocks and reorder levels. Note that sometimes prediction intervals are referred to as *confidence intervals*, though many statisticians prefer not to use that term in this context.

1.4 THE ADVANTAGES OF USING DEDICATED DEMAND FORECASTING SOFTWARE

Research into how companies make their sales forecasts indicates that spreadsheets are the most common of type of software employed. In one survey of US corporations, 48% of respondents used spreadsheets, while only 11% used specialized forecasting software. While spreadsheets may be accessible and allow plenty of flexibility, good-quality dedicated demand forecasting software offers many advantages.

First, they usually offer a wider range of forecasting methods than nondedicated software, and they automatically provide metrics allowing you to compare the accuracy of different methods (see Chapter 3). This increases the chances that you will find the most accurate forecasting method for a particular product. In addition, they are designed to support processes that are specific to demand forecasting such as bottom-up or top-down forecasting when you have a product hierarchy (see Chapter 7). Some software packages will also directly link your demand forecasts to inventory control, advising you on what your reorder level should be and how much safety stock you need to carry. Good-quality forecasting packages also have tried and tested algorithms to implement the different methods. It is known that statistical algorithms in some spreadsheet products contain errors that can result in serious inaccuracy in forecasting. Then there's the danger that you will introduce errors if you are setting up forecasting formulae in a spreadsheet yourself. Mistakes in cutting and pasting and references to the wrong cell ranges can compound the problem. One study found that companies who employed forecasting packages achieved average forecast errors that were almost 7% lower than those of spreadsheet users.

My experience when visiting some companies suggests that if people use a spreadsheet, there is also a danger they will set up an idiosyncratic forecasting system that no one else understands. If they leave the company, it's impossible for their successor to take over the system. Moreover, unlike methods available in good forecasting packages, these self-designed methods usually lack a theoretical underpinning, and they haven't been tested on lots of data sets or compared with other methods.

The second major reason for choosing dedicated software is that many aspects of the forecasting process can be automated and hence performed speedily. This can be particularly important in saving effort if you have to make a large number of forecasts on a frequent basis, and where timely forecasts are crucial to an organization's effectiveness. Even where fewer forecasts are needed, automation can take the form of automatic selection of a forecasting method, ideally with an associated explanation of why the method has been selected. Though automation may have some downsides (see Chapter 8), you will usually have the option of overruling the software's choice in cases where this seems necessary.

Finally, effective demand forecasting is often a team effort, involving forecasters, accounts managers, sales, and operations staff. Dedicated software is likely to be better at supporting collaboration than individuals' spreadsheets. In too many companies, there are *islands of analysis* – people producing their forecasts separately on non-interfacing, and sometimes homemade, systems. Often these systems have no direct connection with those used by production planners or other departments. Forecasters in a major retail company I visited had to copy data from one system using pen and paper, before manually reentering it into another system.

Some companies don't use computers at all to make their forecasts – they rely solely on the gut-feel of managers. As we'll see in Chapter 9, managers' judgments can bring benefits to forecasts if they are used where they are most appropriate. But we'll also see that, more often than not, judgmental forecasts suffer from both political and psychological biases that are inimical to accuracy. Replacing these with the forecasts from a dedicated forecast package is not only likely to reduce costs but also will allow managers to make more effective use of their time.

1.5 GETTING YOUR DATA READY FOR FORECASTING

Garbage in, garbage out has been a well-known phrase in computing for years and, of course, it applies to demand forecasting using computers as well. If you are supplying the software with erroneous data, then it will produce erroneous forecasts. Data cleansing is the process of removing errors and anomalies from the available data. It can be a time consuming and tedious task, but as we'll see, these days – with the availability of modern software – it should not be carried out too far. Moreover, if you use appropriate sources of data, it may be largely unnecessary.

First, we should note that our objective is to forecast future *demand* for a product or service, but historical data usually relates to *sales*. The two are not necessarily the same. Last month there may have been a demand for 2,500 units of your product, but because you only had 2,000 units available, you had sales of 2,000. In many situations, it's difficult to know how much demand was unfulfilled. Customers who saw an empty shelf where your product is usually displayed are unlikely to let you know that they were disappointed. However, some types of historical data are likely to be closer to demand than others.

Data on the quantities of a product shipped on particular dates may not reflect demand at that time because they might simply be delayed deliveries. Analysts agree that either point-of-sale (POS) or syndicated scanner data generally provide a better guide to demand. Syndicated scanner data can be obtained from specialist providers who use individually scanned purchases from large numbers of locations to build and then manipulate a database. This allows sales of products to be viewed at different levels of granularity – nationally, regionally, or at individual retail locations, or annually, quarterly, monthly weekly, or even daily. POS or syndicated scanner data is less likely to contain errors than data on past shipments that might have involved manual data entry, obviating the need to spend many joyless hours attempting to correct mistakes and chase down missing figures.

Even when historical data is correct, people can spend a lot of time cleansing it trying to remove unusual observations. These may be outliers such as exceptionally low sales figures caused by shortages of supplies or exceptionally high sales figures resulting from successful sales promotions. The worry is that these may distort forecasts, so attempts are often made to obtain *baseline* sales figures – that is, sales figures with the effects of shortages or promotions removed. The problem is that it's often difficult to know the size of the adjustment that should be applied to the anomalous sales figure. For example, how do you know how much of the sales uplift was due to a promotion and how much was simply caused by good weather that coincided with the promotion? When a promotion takes place at particular times of the year, how much of the spike in sales was simply a seasonal increase in demand that takes place every year at this time? The problem is particularly acute when, as is usual, judgment is being used to make the adjustment. As we'll see in Chapter 9, judgment can be subject to a range of cognitive biases.

The good news is that, when you have good-quality software, cleansing data for unusual events is likely to be generally unnecessary. You simply have to flag the observation as a special case, and the software will take care of the rest. In other cases, you may be using a model that employs data on demand drivers, such as advertising expenditure, to explain the sales patterns. This means that it can take into account unusual observations and even forecast when they will occur.

1.6 TRADING-DAY ADJUSTMENTS

One data preparation task that is usually well worth carrying out is adjusting past sales data for month lengths or the number of trading days. The demand for a product is likely to vary between months because they contain different numbers of days when the product is available for purchase. Weekends, national holidays, and the different lengths of months can all contribute to this variation and, if it is not taken into account, the software may falsely ascribe the variation to seasonality or randomness. Because we know the future calendar, this is an aspect of future demand that we can predict, so adjusting for the number of trading days tends to improve forecast accuracy.

Many forecasting software products will carry out the adjustments for you if you supply them with details of the relevant past and future numbers of trading days. If this facility is not available, one approach is to use the following steps:

1. Calculate the average number of trading days for the month in question (e.g., on average, how many trading days were there in October during the period for which you have data?). Call this A.

2. Determine the number of trading days for the specific month you are adjusting. Call this B.
3. Divide A by B to obtain the trading day weight.
4. Multiply the demand in the month in question by the trading day weight.

For example, suppose that the demand for a product in March of last year was 2,100 units and that month had 23 trading days. During our sales history, March on average had 22 trading days. Thus, $A = 22$ days and $B = 23$, so last March's trading day weight is $22/23 = 0.957$. This means its adjusted demand is $2{,}100 \times 0.957 = 2{,}010$ units. This puts last March's demand an equal footing with March's demand in other years. Because it had more than the average number of trading days, we have reduced its demand figure.

Once we have used the adjusted data to make the forecasts, we will need to readjust these to take into account the number of trading days in the future month. If we have a forecast of 3,000 units for March next year and it has only 20 trading days, then that March has a trading day weight of $22/20 = 1.1$. This time we *divide* the forecasted demand by the weight to give a revised forecast of $3000/1.1 = 2{,}727$ units.

This method assumes that all trading days are equally important. For example, if Tuesday is usually a busier trading day than Monday, then, ideally, we need to take into account the number of Mondays and Tuesdays in each month. Methods like regression analysis (see Chapter 6) can be used if this is the case.

1.7 OVERVIEW OF THE REST OF THE BOOK

Once your data is ready, you are all set to start making forecasts using your software. In the next chapter, we will look at how software searches for patterns in data so that it can project these ahead to produce forecasts. Chapter 3 looks at how the accuracy of forecasts can be measured and how biases, such as a systematic tendency to forecast too high, can be detected. The next three chapters explain how the most common methods for producing demand forecasts work and how to interpret the output from your software when you are using these methods.

Efficient inventory control is a key reason for demand forecasting, and Chapter 7 shows how you can use forecasts to determine when to place orders for new supplies and how to decide on safety stock levels. Often, we can arrange products into hierarchies to reflect their characteristics or markets. This chapter goes on to compare different methods for handling hierarchies that are available in some software products. It also discusses a method called *temporal aggregation* that can improve forecasts that guide inventory decisions.

Some software products will automatically suggest the best method for making forecasts in different situations and Chapter 8 explores when it is best to automate the forecasting process. It also considers whether complex forecasting methods are always more likely to give more accurate forecasts than simpler ones.

In contrast to automation, research shows that managers frequently use their judgment to override the forecasts generated by their software. Chapter 9 warns of the dangers of overzealous interventions and indicates when an override is likely to be justified.

Predicting the demand for new products is one of the most challenging tasks facing forecasters because of the lack of available historical demand data. However, historical data is often available for similar products launched in earlier periods. Chapter 10 shows how forecasting software can be used to exploit this data. It also looks at the potential pitfalls of using management judgment when forecasting the demand for new products.

The final chapter identifies the key attributes of a good-quality forecasting software product and provides a summary of the key messages in the book.

1.8 SUMMARY OF KEY TERMS

Sales and demand. Demand is the amount of a product that could be sold if there were always sufficient supplies available to meet it. The amount actually sold will be less than demand when stock-outs have occurred, so some demand is unfulfilled.

Probability distribution.	A forecast that indicates the possible levels of demand together with the probability that each level will occur.
Normal or Gaussian distribution.	A probability distribution that is often used to approximate the true probability distribution of demand. It assumes that the chances of demand being above or below average are the same and that the probability of particular levels of demand declines the further they are from the average. When plotted as a histogram it resembles a bell shape.
Point forecast.	A forecast presented as a single number.
Prediction interval.	A range that has a specified probability of capturing the actual level of demand between its lower and upper limits.
Trading-day adjustments.	Changes made to historic sales data to take into account variations in the number of days a product was traded in different months.

1.9 REFERENCES

Chase, Jr., C. W. (2015). Cleanse your historical shipment data? Why? *Journal of Business Forecasting*, **34**, 29–33.

Field, P. J. (2006). On the use and abuse of Microsoft Excel. *Foresight*, **3**, 46–47.

Kamal, J. (2013). Best practice demand planning meets unprecedented demand volatility. *Supply and Demand Chain Executive* (December 5). http://www.sdcexec.com/article/11267412/research-shows-how-best-practice-demand-planning-contributes-to-the-challenge-of-accurate-forecasting-in-an-era-of-global-demand-uncertainty

McCullough. B. D. (2006). The unreliability of Excel's statistical procedures. *Foresight*, **3**, 44–45.

Moon, M. A., Mentzer, J. T., & Smith, C. D. (2003). Conducting a sales forecasting audit. *International Journal of Forecasting*, **19**(1), 5–25.

Sanders, N. R., and Marodt, K. B. (2003). Forecasting software in practice: use satisfaction and performance. *Interfaces*, **33**, 90–93.