

Chapter

Betting on Uncertain Demand: The Newsvendor Model¹

Matching supply and demand is particularly challenging when supply must be chosen before observing demand and demand is stochastic (uncertain). To illustrate this point, suppose you are the owner of a simple business: selling newspapers. Each morning you purchase a stack of papers with the intention of selling them at your newsstand at the corner of a busy street. Even though you have some idea regarding how many newspapers you can sell on any given day, you never can predict demand for sure. Some days you sell all of your papers, while other days end with unsold newspapers to be recycled. As the newsvendor, you must decide how many papers to buy at the start of each day. Because you must decide how many newspapers to buy before demand occurs, unless you are very lucky, you will not be able to match supply to demand. A decision tool is needed to make the best out of this difficult situation. The *newsvendor model* is such a tool.

You will be happy to learn that the newsvendor model applies in many more settings than just the newsstand business. The essential issue is that you must take a firm bet (how much inventory to order) before some random event occurs (demand) and then you learn that you either bet too much (demand was less than your order) or you bet too little (demand exceeded your order). This trade-off between “doing too much” and “doing too little” occurs in other settings. Consider a technology product with a long lead time to source components and only a short life before better technology becomes available. Purchase too many components and you risk having to sell off obsolete technology. Purchase too few and you may forgo sizable profits. Cisco is a company that can relate to these issues: In 2000 they estimated that they were losing 10 percent of their potential orders to rivals due to long lead times created by shortages of parts; but by early 2001, the technology bubble had burst and they had to write off \$2.5 billion in inventory.

This chapter begins with a description of the production challenge faced by O’Neill Inc., a sports apparel manufacturer. O’Neill’s decision also closely resembles the newsvendor’s task. We then describe the newsvendor model in detail and apply it to O’Neill’s problem. We also show how to use the newsvendor model to forecast a number of performance measures relevant to O’Neill.

¹ Data in this chapter have been disguised to protect confidential information.

11.1 O’Neill Inc.

O’Neill Inc. is a designer and manufacturer of apparel, wetsuits, and accessories for water sports: surf, dive, waterski, wake-board, triathlon, and wind surf. Their product line ranges from entry-level products for recreational users, to wetsuits for competitive surfers, to sophisticated dry suits for professional cold-water divers (e.g., divers that work on oil platforms in the North Sea). O’Neill divides the year into two selling seasons: Spring (February through July) and Fall (August through January). Some products are sold in both seasons, but the majority of their products sell primarily in a single season. For example, waterski is active in the Spring season whereas recreational surf products sell well in the Fall season. Some products are not considered fashionable (i.e., they have little cosmetic variety and they sell from year to year), for example, standard neoprene black booties. With product names like “Animal,” “Epic,” “Hammer,” “Inferno,” and “Zen,” O’Neill clearly also has products that are subject to the whims of fashion. For example, color patterns on surf suits often change from season to season to adjust to the tastes of the primary user (15–30-year-old California males).

O’Neill operates its own manufacturing facility in Mexico, but it does not produce all of its products there. Some items are produced by the TEC Group, O’Neill’s contract manufacturer in Asia. While TEC provides many benefits to O’Neill (low cost, sourcing expertise, flexible capacity, etc.), they do require a three-month lead time on all orders. For example, if O’Neill orders an item on November 1, then O’Neill can expect to have that item at its distribution center in San Diego, California, ready for shipment to customers, only on January 31.


To better understand O’Neill’s production challenge, let’s consider a particular wetsuit used by surfers and newly redesigned for the upcoming spring season, the Hammer 3/2. (The “3/2” signifies the thickness of the neoprene on the suit: 3 mm thick on the chest and 2 mm everywhere else.) Figure 11.1 displays the Hammer 3/2 and O’Neill’s logo. O’Neill has decided to let TEC manufacture the Hammer 3/2. Due to TEC’s three-month lead time, O’Neill needs to submit an order to TEC in November before the start of the spring season. Using past sales data for similar products and the judgment of its designers and sales representatives, O’Neill developed a forecast of 3,200 units for total demand during the spring season for the Hammer 3/2. Unfortunately, there is considerable uncertainty in that forecast despite the care and attention placed on the formation of the forecast. For example, it is O’Neill’s experience that 50 percent of the time the actual demand deviates from their initial forecast by more than 25 percent of the forecast. In other words, only 50 percent of the time is the actual demand between 75 percent and 125 percent of their forecast.

Although O’Neill’s forecast in November is unreliable, O’Neill will have a much better forecast for total season demand after observing the first month or two of sales. At that time, O’Neill can predict whether the Hammer 3/2 is selling slower than forecast, in which case O’Neill is likely to have excess inventory at the end of the season, or whether the Hammer 3/2 is more popular than predicted, in which case O’Neill is likely to stock out. In the latter case, O’Neill would love to order more Hammers, but the long lead time from Asia prevents O’Neill from receiving those additional Hammers in time to be useful. Therefore, O’Neill essentially must “live or die” with its single order placed in November.

Fortunately for O’Neill, the economics on the Hammer are pretty good. O’Neill sells the Hammer to retailers for \$180 while it pays TEC \$110 per suit. If O’Neill has leftover inventory at the end of the season, it is O’Neill’s experience that they are able to sell that inventory with a 50 percent discount. Figure 11.2 summarizes the time line of events and the economics for the Hammer 3/2.

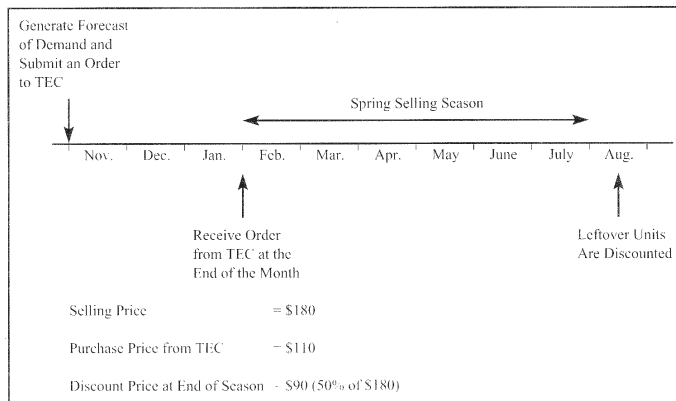
So how many units should O’Neill order from TEC? You might argue that O’Neill should order the forecast for total demand, 3,200, because 3,200 is the most likely outcome. The

FIGURE 11.1
O'Neill's Hammer
3/2 Wetsuit and Logo
for the Surf Market



forecast is also the value that minimizes the expected absolute difference between the actual demand and the production quantity; that is, it is likely to be close to the actual demand. Alternatively, you may be concerned that forecasts are always biased and therefore suggest an order quantity less than 3,200 would be more prudent. Finally, you might argue that because the

FIGURE 11.2
Time Line of Events
and Economics for
O'Neill's Hammer
3/2 Wetsuit



gross margin on the Hammer is nearly 40 percent $((180 - 110)/180 = 0.39)$, O'Neill should order more than 3,200 in case the Hammer is a hit. We next define the newsvendor model and then discuss what the newsvendor model recommends for an order quantity.

1.2 An Introduction to the Newsvendor Model

The newsvendor model considers a setting in which you have only one production or procurement opportunity. Because that opportunity occurs well in advance of a single selling season, you receive your entire order just before the selling season starts. Stochastic demand occurs during the selling season. If demand exceeds your order quantity, then you sell your entire order. But if demand is less than your order quantity, then you have leftover inventory at the end of the season.

There is a fixed cost per unit ordered: for the Hammer 3/2, $\text{Cost} = 110$. It is important that Cost includes only costs that depend on the number of units ordered; amortized fixed costs should not be included, because they are unaffected by our order quantity decision. In other words, this cost figure should include all costs that vary with the order quantity and no costs that do not vary with the order quantity. There is a fixed price for each unit you sell; in this case, $\text{Price} = 180$.

If there is leftover inventory at the end of the season, then there is some value associated with that inventory. To be specific, there is a fixed *Salvage value* that you earn on each unit of leftover inventory: with the Hammer, the Salvage value = $50\% \times 180 = 90$. It is possible that leftover inventory has no salvage value whatsoever, that is, Salvage value = 0. It is also possible leftover inventory is costly to dispose, in which case the salvage value may actually be a salvage cost. For example, if the product is a hazardous chemical, then there is a cost for disposing of leftover inventory; that is, Salvage value < 0 is possible.

To guide your production decision, you need a forecast for demand. O'Neill's initial forecast for the Hammer is 3,200 units for the season. But it turns out (for reasons that are explained later) you need much more than just a number for a forecast. You need to have a sense of how accurate your forecast is; you need a forecast on your forecast! For example, in an ideal world, there would be absolutely no error in your forecast: if the forecast is 3,200 units, then 3,200 units is surely the demand for the season. In reality, there will be error in the forecast, but forecast error can vary in size. For example, it is better to be 90 percent sure demand will be between 3,100 and 3,300 units than it is to be 90 percent sure demand will be between 2,400 and 4,000 units. Intuition should suggest that you might want to order a different amount in those two situations.

To summarize, the newsvendor model represents a situation in which a decision maker must make a single bet (e.g., the order quantity) before some random event occurs (e.g., demand). There are costs if the bet turns out to be too high (e.g., leftover inventory that is salvaged for a loss on each unit). There are costs if the bet turns out to be too low (the opportunity cost of lost sales). The newsvendor model's objective is to bet an amount that correctly balances those opposing forces. To implement the model, we need to identify our costs and how much demand uncertainty we face. We have already identified our costs, so the next section focuses on the task of identifying the uncertainty in Hammer 3/2 demand.

11.3 Constructing a Demand Forecast

The newsvendor model balances the cost of ordering too much against the cost of ordering too little. To do this, we need to understand how much demand uncertainty there is for the Hammer 3/2, which essentially means we need to be able to answer the following question:

What is the probability demand will be less than or equal to Q units?

for whatever Q value we desire. In short, we need a *distribution function*. Recall from statistics, every random variable is defined by its distribution function, $F(Q)$, which is the probability the outcome of the random variable is Q or lower. In this case the random variable is demand for the Hammer 3/2 and the distribution function is

$$F(Q) = \text{Prob}\{\text{Demand is less than or equal to } Q\}$$

For convenience, we refer to the distribution function, $F(Q)$, as our demand forecast because it gives us a complete picture of the demand uncertainty we face. The objective of this section is to explain how we can use a combination of intuition and data analysis to construct our demand forecast.

Distribution functions come in two forms. *Discrete distribution functions* can be defined in the form of a table: There is a set of possible outcomes and each possible outcome has a probability associated with it. The following is an example of a simple discrete distribution function with three possible outcomes:

Q	$F(Q)$
2,200	0.25
3,200	0.75
4,200	1.00

The Poisson distribution is an example of a discrete distribution function that we will use extensively. With *continuous distribution functions* there are an unlimited number of possible outcomes. Both the exponential and the normal are continuous distribution functions. They are defined with one or two parameters. For example, the normal distribution is defined by two parameters: its mean and its standard deviation. We use μ to represent the mean of the distribution and σ to represent the standard deviation. (μ is the Greek letter mu and σ is the Greek letter sigma.) This notation for the mean and the standard deviation is quite common, so we adopt it here.

In some situations, a discrete distribution function provides the best representation of demand, whereas in other situations a continuous distribution function works best. Hence, we work with both types of distribution functions.

Now that we know two ways to express our demand forecast, let's turn to the complex task of actually creating the forecast. As mentioned in Section 11.1, the Hammer 3/2 has been redesigned for the upcoming spring season. As a result, actual sales in the previous season might not be a good guide for expected demand in the upcoming season. In addition to the product redesign, factors that could influence expected demand include the pricing and marketing strategy for the upcoming season, changes in fashion, changes in the economy (e.g., is demand moving toward higher or lower price points), changes in technology, and overall trends for the sport. To account for all of these factors, O'Neill surveyed the opinion of a number of individuals in the organization on their personal demand forecast for the Hammer 3/2. The survey's results were averaged to obtain the initial 3,200-unit forecast. This represents the "intuition" portion of our demand forecast. Now we need to analyze O'Neill's available data to further develop the demand forecast.

Table 11.1 presents data from O'Neill's previous spring season with wetsuits in the surf category. Notice that the data include both the original forecasts for each product as well as its actual demand. The original forecast was developed in a process that was comparable to the one that led to the 3,200-unit forecast for the Hammer 3/2 for this season. For example, the forecast for the Hammer 3/2 in the previous season was 1,300 units, but actual demand was 1,696 units.

TABLE 11.1
Forecasts and Actual
Demand Data for
Selected Wetsuits from
the Previous Spring
Season


Product Description	Forecast	Actual Demand	Error*	A/F Ratio**
JR ZEN FL 3/2	90	140	50	1.56
EPIC 5/3 W/HD	120	83	37	0.69
JR ZEN 3/2	140	143	3	1.02
WMS ZEN-ZIP 4/3	170	163	7	0.96
HEATWAVL 3/2	170	212	42	1.25
JR EPIC 3/2	180	175	5	0.97
WMS ZEN 3/2	180	195	15	1.08
ZEN-ZIP 5/4/3 W/HOOD	270	317	47	1.17
WMS EPIC 5/3 W/HD	320	369	49	1.15
EVO 3/2	380	587	207	1.54
JR EPIC 4/3	380	571	191	1.50
WMS EPIC 2MM FULL	390	311	79	0.80
HEATWAVE 4/3	430	274	156	0.64
ZEN 4/3	430	239	191	0.56
EVO 4/3	440	623	183	1.42
ZEN FL 3/2	450	365	85	0.81
HEAT 4/3	460	450	10	0.98
ZEN-ZIP 2MM FULL	470	116	354	0.25
HEAT 3/2	500	635	135	1.27
WMS EPIC 3/2	610	830	220	1.36
WMS ELITE 3/2	650	364	286	0.56
ZEN-ZIP 3/2	660	788	128	1.19
ZEN 2MM 5/S FULL	680	453	227	0.67
EPIC 2MM 5/S FULL	740	607	133	0.82
EPIC 4/3	1,020	732	288	0.72
WMS EPIC 4/3	1,060	1,552	492	1.46
JR HAMMER 3/2	1,220	721	499	0.59
HAMMER 3/2	1,300	1,696	396	1.30
HAMMER 5/S FULL	1,490	1,832	342	1.23
EPIC 3/2	2,190	3,504	1,314	1.60
ZEN 3/2	3,190	1,995	1,195	0.37
ZEN-ZIP 4/3	3,810	3,289	521	0.86
WMS HAMMER 3/2 FULL	6,490	3,673	2,817	0.57

*Error = Forecast - Actual demand

**A/F ratio = Actual demand divided by Forecast

So how does O'Neill know actual demand for a product that stocks out? For example, how does O'Neill know that actual demand was 1,696 for last year's Hammer 3/2 if they only ordered 1,500 units? Because retailers order via phone or fax, O'Neill can keep track of each retailer's initial order, that is, the retailer's demand before the retailer knows a product is unavailable. (However, life is not perfect: O'Neill's phone representatives do not always record a customer's initial order into the computer system, so there is even some uncertainty with that figure. We'll assume this is a minor issue and not address it in our analysis.) In other settings, a firm may not be able to know actual demand with that level of precision. For example, a retailer of O'Neill's products probably does not get to observe what demand could be for the Hammer 3/2 once the Hammer is out of stock at the retailer. However, that retailer would know when during the season the Hammer stocked out and, hence, could use that information to forecast how many additional units could have been sold during the remainder of the season. Therefore, even if a firm cannot directly observe lost sales, a firm should be able to obtain a reasonable estimate for what demand could have been.

FIGURE 11.3
Forecasts and Actual
Demand for Surf
Wetsuits from the
Previous Season



As can be seen from the data, the forecasts ranged from a low of 90 units to a high of 6,490 units. There was also considerable forecast error: O'Neill goofed with the Women's Hammer 3/2 Full suit with a forecast nearly 3,000 units above actual demand, while the forecast for the Epic 3/2 suit was about 1,300 units too low. Figure 11.3 gives a scatter plot of forecasts and actual demand. If forecasts were perfect, then all of the observations would lie along the diagonal line.

While the absolute errors for some of the bigger products are dramatic, the forecast errors for some of the smaller products are also significant. For example, the actual demand for the Juniors Zen Flat Lock 3/2 suit was more than 150 percent greater than forecast. This suggests that we should concentrate on the relative forecast errors instead of the absolute forecast errors.

Relative forecast errors can be measured with the *A/F ratio*:

$$A/F \text{ ratio} = \frac{\text{Actual demand}}{\text{Forecast}}$$

An accurate forecast has an *A/F ratio* = 1, while an *A/F ratio* above 1 indicates the forecast was too low and an *A/F ratio* below 1 indicates the forecast was too high. Table 11.1 displays the *A/F ratios* for our data in the last column.

Those *A/F ratios* provide a measure of the forecast accuracy from the previous season. To illustrate this point, Table 11.2 sorts the data in ascending *A/F* order. Also included in the table is each product's *A/F rank* in the order and each product's percentile, the fraction of products that have that *A/F rank* or lower. (For example, the product with the fifth *A/F ratio* has a percentile of $5/33 = 15.2$ percent because it is the fifth product out of 33 products in the data.) We see from the data that actual demand is less than 80 percent of the forecast for one-third of the products (the *A/F ratio* 0.8 has a percentile of 33.3) and actual demand is greater than 125 percent of the forecast for 27.3 percent of the products (the *A/F ratio* 1.25 has a percentile of 72.7).

TABLE 11.2
Ascending *A/F Ratios* for
Wetsuits from
the Previous Spring
Season

Product Description	Forecast	Actual Demand	A/F Ratio*	Rank	Percentile**
ZEN-ZIP 2MM FULL	470	116	0.25	1	3.0
ZEN 3/2	3,190	1,195	0.37	2	6.1
ZEN 4/3	430	239	0.56	3	9.1
WMS ELITE 3/2	650	364	0.56	4	12.1
WMS HAMMER 3/2 FULL	6,490	3,673	0.57	5	15.2
JR HAMMER 3/2	1,220	721	0.59	6	18.2
HEATWAVE 4/3	430	274	0.64	7	21.2
ZEN 2MM S/S FULL	680	453	0.67	8	24.2
EPIC S/3 W/HD	120	83	0.69	9	27.3
EPIC 4/3	1,020	732	0.72	10	30.3
WMS EPIC 2MM FULL	390	311	0.80	11	33.3
ZEN FL 3/2	450	365	0.81	12	36.4
EPIC 2MM S/S FULL	740	607	0.82	13	39.4
ZEN-ZIP 4/3	3,810	3,289	0.86	14	42.4
WMS ZEN-ZIP 4/3	170	163	0.96	15	45.5
JR EPIC 3/2	180	175	0.97	16	48.5
HEAT 4/3	460	450	0.98	17	51.5
JR ZEN 3/2	140	143	1.02	18	54.5
WMS ZEN 3/2	180	195	1.08	19	57.6
WMS EPIC S/3 W/HD	320	369	1.15	20	60.6
ZEN-ZIP 5/4/3 W/HOOD	270	317	1.17	21	63.6
ZEN-ZIP 3/2	660	788	1.19	22	66.7
HAMMER S/S FULL	1,490	1,832	1.23	23	69.7
HEATWAVE 3/2	170	212	1.25	24	72.7
HEAT 3/2	500	635	1.27	25	75.8
HAMMER 3/2	1,300	1,696	1.30	26	78.8
WMS EPIC 3/2	610	830	1.36	27	81.8
EVO 4/3	440	623	1.42	28	84.8
WMS EPIC 4/3	1,060	1,552	1.46	29	87.9
JR EPIC 4/3	380	571	1.50	30	90.9
EVO 3/2	380	587	1.54	31	93.9
JR ZEN FL 3/2	90	140	1.56	32	97.0
EPIC 3/2	2,190	3,504	1.60	33	100.0

*A/F ratio = Actual demand divided by Forecast

**Percentile = Rank divided by total number of wetsuits (33)

Given that the *A/F ratios* from the previous season reflect forecast accuracy in the previous season, maybe the current season's forecast accuracy will be comparable. For example, according to Table 11.2, there is a 3.0 percent chance demand is 25 percent of the forecast. Hence, if this season's forecast error is similar to last season's forecast error, then there is a 3.0 percent chance demand will be 800 units or fewer ($0.25 \times 3,200 = 800$). Similarly, from Table 11.2 we see there is a 90.9 percent chance demand is 150 percent of the forecast or lower. Translating that to the Hammer 3/2 means there is a 90.9 percent chance demand will be 4,800 units or lower ($1.5 \times 3,200 = 4,800$). For each *A/F ratio* we observed in our historical data set, we are able to derive the corresponding demand level for the Hammer 3/2 and the probability we should observe that demand or lower. These calculations are provided in Table 11.3.

The demand forecast represented by Table 11.3 is a discrete distribution function, but we will also refer to it as an *empirical distribution function* (because it is the distribution function constructed with the empirically observed data). Exhibit 11.1 summarizes the process of constructing an empirical distribution function.

A PROCESS FOR USING HISTORICAL A/F RATIOS TO CONSTRUCT AN EMPIRICAL DISTRIBUTION FUNCTION

- Step 1. Assemble a data set of products for which the forecasting task is comparable to the product of interest. In other words, the data set should include products that you expect would have similar forecast error to the product of interest. (They may or may not be similar products.) The data should include an initial forecast of demand and the actual demand for each item. We also need an initial forecast for the product for the upcoming season. Your data table should be similar in form to Table 11.1.
- Step 2. Evaluate the A/F ratio for each product in the data set.
- Step 3. Sort the data in ascending A/F ratio order and rank the items from 1 to N , where N is the number of items in the data set. (See Table 11.2.)
- Step 4. There are N items in the empirical distribution function. The quantity of the i th item equals the initial forecast times the A/F ratio of the i th item in the sorted data set. The probability of the i th item equals i/N . (See Table 11.3.)

An attractive feature of the empirical distribution function in Table 11.3 is that it reflects our historical forecasting capability. A disadvantage of this demand forecast is that it predicts only a limited number of possible outcomes. For example, according to the table, it is possible that demand for the Hammer 3/2 will be 800 or 1,184 units, but 1,000 units is not possible. Hence, it can be desirable to fit a continuous distribution function to the data so that a broader range of outcomes are possible.

As an alternative to the empirical distribution function, let's use the well-known normal distribution. Given that there are an infinite number of potential normal distributions (essentially any mean and standard deviation combination), our challenge is to find a normal distribution that fits our data in Table 11.2.

Take the definition of the A/F ratio and rearrange terms to get

$$\text{Actual demand} = \text{A/F ratio} \times \text{Forecast}$$

For the Hammer 3/2, the forecast is 3,200 units. Note that the forecast is not random, but the A/F ratio is random. Hence, the randomness in actual demand is directly

related to the randomness in the A/F ratio. Using standard results from statistics and the above equation, we get the following results:

$$\text{Expected actual demand} = \text{Expected A/F ratio} \times \text{Forecast}$$

and

$$\text{Standard deviation of demand} = \text{Standard deviation of A/F ratios} \times \text{Forecast}$$

Expected actual demand, or *expected demand* for short, is what we should choose for the mean for our normal distribution, μ . The average A/F ratio in Table 11.2 is 0.9976. Therefore, expected demand for the Hammer 3/2 in the upcoming season is $0.9976 \times 3,200 = 3,192$ units. In other words, if the initial forecast is 3,200 units and the future A/F ratios are comparable to the past A/F ratios, then the mean of actual demand is 3,192 units. So let's choose 3,192 units as our mean of the normal distribution.

This decision may raise some eyebrows: If our initial forecast is 3,200 units, why do we not instead choose 3,200 as the mean of the normal distribution? Because 3,192 is so close to 3,200, assigning 3,200 as the mean probably would lead to a good order quantity as well. However, suppose the average A/F ratio were 0.90, that is, on average, actual demand is 90 percent of the forecast. It is quite common for people to have overly optimistic forecasts, so an average A/F ratio of 0.90 is possible. In that case, expected actual demand would only be $0.90 \times 3,200 = 2,880$. Because we want to choose a normal distribution that represents actual demand, in that situation it would be better to choose a mean of 2,880 even though our initial forecast is 3,200. (Novice golfers sometimes adopt an analogous strategy. If a golfer consistently hooks the ball to the right on her drives, then she should aim to the left of the flag. In an ideal world, there would be no hook to her shot nor a bias in the forecast. But if the data say there is a hook, then it should not be ignored. Of course, the golfer and the forecaster also should work on eliminating the bias.)

Now that we have a mean for our normal distribution, we need a standard deviation. The second equation above tells us that the standard deviation of actual demand equals the standard deviation of the A/F ratios times the forecast. The standard deviation of the A/F ratios in Table 11.2 is 0.369. (Use the "stdev()" function in Excel.) So the standard deviation of actual demand is the standard deviation of the A/F ratios times the initial forecast: $0.369 \times 3,200 = 1,181$. Hence, for the second way to express our demand forecast for the Hammer 3/2, we can use a normal distribution with a mean of 3,192 and a standard deviation of 1,181. See Exhibit 11.2 for a summary of the process of choosing a mean and a standard deviation for a normal distribution forecast.

With a discrete distribution function, it is easy to find $F(Q)$ because we need only look it up in a table. But now we need to find $F(Q)$ with a normal distribution demand forecast. There are two ways this can be done. The first way is to use spreadsheet software. For example, in Excel use the function Normdist(Q , 3192, 1181, 1). The second way, which does not require a computer, is to use the Standard Normal Distribution Function Table in Appendix B.

The *standard normal* is a particular normal distribution: its mean is 0 and its standard deviation is 1. To introduce another piece of common Greek notation, let $\Phi(z)$ be the distribution function of the standard normal. Even though the standard normal is a continuous distribution, it can be "chopped up" into pieces to make it into a discrete distribution. The Standard Normal Distribution Function Table is exactly that; that is, it is the discrete version of the standard normal distribution. The full table is in Appendix B, but Table 11.4 reproduces a portion of the table.

TABLE 11.3
Discrete/Empirical Distribution Function for the Hammer 3/2 Using Historical A/F Ratios

A/F Ratio	Q	F(Q)	A/F Ratio	Q	F(Q)	A/F Ratio	Q	F(Q)
0.25	800	0.0303	0.81	2,592	0.3636	1.23	3,936	0.6970
0.37	1,184	0.0606	0.82	2,624	0.3939	1.25	4,000	0.7273
0.56	1,792	0.0909	0.86	2,752	0.4242	1.27	4,064	0.7576
0.56	1,792	0.1212	0.96	3,072	0.4545	1.30	4,160	0.7879
0.57	1,824	0.1515	0.97	3,104	0.4848	1.36	4,352	0.8182
0.59	1,888	0.1818	0.98	3,136	0.5152	1.42	4,544	0.8485
0.64	2,048	0.2121	1.02	3,264	0.5455	1.46	4,672	0.8788
0.67	2,144	0.2424	1.08	3,456	0.5758	1.50	4,800	0.9091
0.69	2,208	0.2727	1.15	3,680	0.6061	1.54	4,928	0.9394
0.72	2,304	0.3030	1.17	3,744	0.6364	1.56	4,992	0.9697
0.80	2,560	0.3333	1.19	3,808	0.6667	1.60	5,120	1.0000

Q = A/F ratio times the initial sales forecast, 3,200 units
 $F(Q)$ = Probability demand is less than or equal to the quantity Q

A PROCESS FOR USING HISTORICAL A/F RATIOS TO CHOOSE A MEAN AND STANDARD DEVIATION FOR A NORMAL DISTRIBUTION FORECAST

- Step 1. Assemble a data set of products for which the forecasting task is comparable to the product of interest. In other words, the data set should include products that you expect would have similar forecast error to the product of interest. (They may or may not be similar products.) The data should include an initial forecast of demand and the actual demand. We also need an initial forecast for the item for the upcoming season.
- Step 2. Evaluate the A/F ratio for each product in the data set. Evaluate the average of the A/F ratios (that is, the expected A/F ratio) and the standard deviation of the A/F ratios. (In Excel use the average() and stdev() functions.)
- Step 3. The mean and standard deviation of the normal distribution that we will use as the forecast can now be evaluated with the following two equations:

$$\text{Expected demand} = \text{Expected A/F ratio} \times \text{Forecast}$$

$$\text{Standard deviation of demand} = \text{Standard deviation of A/F ratios} \times \text{Forecast}$$

where the forecast in the above equations is the initial forecast.

Although the Standard Normal Distribution Function Table is surely a table, its format makes it somewhat tricky to read. For example, suppose you wanted to know the probability that the outcome of a standard normal is 0.51 or lower. We are looking for the value of $\Phi(z)$ with $z = 0.51$. To find that value, pick the row and column in the table such that the first number in the row and the first number in the column add up to the z value you seek. With $z = 0.51$, we are looking for the row that begins with 0.50 and the column that begins with 0.01, because the sum of those two values equals 0.51. The intersection of that row with that column gives $\Phi(z)$; from Table 11.4 we see that $\Phi(0.51) = 0.6950$. Therefore, there is a 69.5 percent probability the outcome of a standard normal is 0.51 or lower.

But it is unlikely that our demand forecast will be a standard normal distribution. So how can we use the standard normal to find $F(Q)$; that is, the probability demand will be Q or lower given that our demand forecast is some other normal distribution? The answer is that we convert the quantity we are interested in, Q , into an equivalent quantity for the standard normal. In other words, we find a z such that $F(Q) = \Phi(z)$; that is, the probability demand is less than or equal to Q is the same as the probability the outcome of a standard normal is z or lower. That z is called the z -statistic. Once we have the appropriate z -statistics, we then just look up $\Phi(z)$ in the Standard Normal Distribution Function Table to get our answer.

To convert Q into the equivalent z -statistic, use the following equation:

$$z = \frac{Q - \mu}{\sigma}$$

TABLE 11.4
A Portion of the
Standard Normal
Distribution Function
Table, $\Phi(z)$

z	0	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09
0	0.5000	0.5040	0.5080	0.5120	0.5160	0.5199	0.5239	0.5279	0.5319	0.5359
0.1	0.5398	0.5438	0.5478	0.5517	0.5557	0.5596	0.5636	0.5675	0.5714	0.5753
0.2	0.5793	0.5832	0.5871	0.5910	0.5948	0.5987	0.6026	0.6064	0.6103	0.6141
0.3	0.6179	0.6217	0.6255	0.6293	0.6331	0.6368	0.6406	0.6443	0.6480	0.6517
0.4	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772	0.6808	0.6844	0.6879
0.5	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123	0.7157	0.7190	0.7224
0.6	0.7257	0.7291	0.7324	0.7357	0.7389	0.7422	0.7454	0.7486	0.7517	0.7549
0.7	0.7580	0.7611	0.7642	0.7673	0.7704	0.7734	0.7764	0.7794	0.7823	0.7852

A PROCESS FOR EVALUATING THE PROBABILITY DEMAND IS EITHER LESS THAN OR EQUAL TO Q (WHICH IS $F(Q)$) OR MORE THAN Q (WHICH IS $1 - F(Q)$)

If the demand forecast is a normal distribution with mean μ and standard deviation σ , then follow steps A and B:

A. Evaluate the z -statistic that corresponds to Q :

$$z = \frac{Q - \mu}{\sigma}$$

B. The probability demand is less than or equal to Q is $\Phi(z)$. With Excel $\Phi(z)$ can be evaluated with the function Normsdist(z); otherwise, look up $\Phi(z)$ in the Standard Normal Distribution Function Table in Appendix B. If you want the probability demand is greater than Q , then your answer is $1 - \Phi(z)$.

If the demand forecast is a discrete distribution function table (as with an empirical distribution function), then look up $F(Q)$, which is the probability demand is less than or equal to Q . If you want the probability demand is greater than Q , then the answer is $1 - F(Q)$.

For example, suppose we are interested in the probability that demand for the Hammer 3/2 will be 4,000 units or lower, that is, $Q = 4,000$. With a normal distribution that has mean 3,192 and standard deviation 1,181, the quantity $Q = 4,000$ has a z -statistic of


$$z = \frac{4,000 - 3,192}{1,181} = 0.68$$

Therefore, the probability demand for the Hammer 3/2 is 4,000 units or lower is $\Phi(0.68)$; that is, it is the same as the probability the outcome of a standard normal is 0.68 or lower. According to the Standard Normal Distribution Function Table (see Table 11.4 for convenience), $\Phi(0.68) = 0.7517$. In other words, there is just over a 75 percent probability that demand for the Hammer 3/2 will be 4,000 or fewer units. Exhibit 11.3 summarizes the process of finding the probability demand will be less than or equal to some Q (or more than Q).

In the remainder of this chapter, we will work with both demand forecasts, the discrete/empirical distribution, and the normal distribution. We study the empirical distribution function because the techniques we apply to it extend exactly to any other discrete distribution function, such as the Poisson or the negative binomial. We study the normal because it is a widely applied distribution (rightfully so). Furthermore, the techniques used with the normal distribution are actually quite similar to those used with the empirical distribution.

You may recall that it has been O'Neill's experience that demand deviated by more than 25 percent from their initial forecast for 50 percent of their products. We can now check whether that experience is consistent with our normal distribution forecast for the Hammer 3/2. Our initial forecast is 3,200 units. So a deviation of 25 percent or more implies demand is either less than 2,400 units or more than 4,000 units. The z -statistic for $Q = 2,400$ is $z = (2400 - 3192)/1181 = -0.67$, and from the Standard Normal Distribution Function Table, $\Phi(-0.67) = 0.2514$. (Find the row with -0.60 and the column with -0.07 .) If there is a 25.14 percent probability demand is less than 2,400 units and a 75.17 percent probability that demand is less than 4,000 units, then there is a $75.17 - 25.14 = 50.03$ percent probability that demand is between 2,400 and 4,000 units. Hence, O'Neill's initial assertion regarding forecast accuracy is consistent with our normal distribution forecast of demand.

FIGURE 11.4
Discrete/Empirical
Distribution Function
(diamonds) and
Normal Distribution
Function with Mean
3,192 and Standard
Deviation 1,181 (solid
line)



To summarize, the objective in this section is to develop a detailed demand forecast. A single “point forecast” (e.g., 3,200 units) is not sufficient. We need to quantify the amount of variability that may occur about our forecast; that is, we need a distribution function. We first constructed a discrete distribution function table with the empirical observed historical demand (Table 11.3) and then we fit a normal distribution to our historical forecast performance data.

Now that we can evaluate the distribution function of Hammer 3/2 demand, we can see how well our chosen normal distribution fits our data. Figure 11.4 combines the empirical distribution function in Table 11.3 with the normal distribution function. The figure shows that they match closely, which suggests that they are roughly equivalent. Furthermore, it suggests that the normal distribution is a good representation of our actual demand.

11.4 The Expected Profit-Maximizing Order Quantity

The next step after assembling all of our inputs (selling price, cost, salvage value, and demand forecast) is to choose an order quantity. The first part in that process is to decide what is our objective. A natural objective is to choose our production/procurement quantity to maximize our expected profit. This section explains how to do this. Section 11.6 considers other possible objectives.

Before revealing the actual procedure for choosing an order quantity to maximize expected profit, it is helpful to explore the intuition behind the solution. Consider again O’Neill’s Hammer 3/2 ordering decision. Should we order one unit? If we do, then there is a very good chance we will sell the unit: With a forecast of 3,192 units, it is likely we sell at least one unit. If we sell the unit, then the gain from that unit equals $\$180 - \$110 = \$70$ (the selling price minus the purchase cost). The *expected* gain from the first unit, which equals the probability of selling the first unit times the gain from the first unit, is then very close to \$70. However, there is also a slight chance that we do not sell the first unit, in which case we incur a loss of $\$110 - \$90 = \$20$. (The loss equals the difference between the purchase cost and the discount price.) But since the probability we do not sell that unit

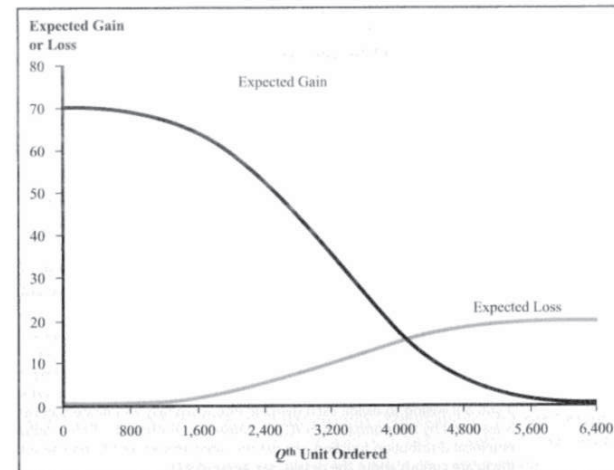
is quite small, the *expected* loss on the first unit is nearly \$0. Given that the expected gain from the first unit clearly exceeds the expected loss, the profit from ordering that unit is positive. In this case it is a good bet to order at least one unit.

After deciding whether to order one unit, we can now consider whether we should order two units, and then three units, and so forth. Two things happen as we continue this process. First, the probability that we sell the unit we are considering decreases, thereby reducing the expected gain from that unit. Second, the probability we do not sell that unit increases, thereby increasing the expected loss from that unit. Now imagine we order the 6,400th unit. The probability of selling that unit is quite low, so the expected gain from that unit is nearly zero. In contrast, the probability of *not* selling that unit is quite high, so the expected loss is nearly \$20 on that unit. Clearly it makes no sense to order the 6,400th unit. This pattern is illustrated in Figure 11.5. We see that from some unit just above 4,000 the expected gain on that unit equals its expected loss.

Let’s formalize this intuition some more. In the newsvendor model, there is a trade-off between ordering too much (which could lead to costly leftover inventory) and ordering too little (which could lead to the opportunity cost of lost sales). To balance these forces, it is useful to think in terms of a cost for ordering too much and a cost for ordering too little. Maximizing expected profit is equivalent to minimizing those costs. To be specific, let C_o be the *overage cost*, the loss incurred when a unit is ordered but not sold. In other words, the overage cost is the per-unit cost of overordering. For the Hammer 3/2, we have $C_o = 20$.

In contrast to C_o , let C_u be the *underage cost*, the opportunity cost of not ordering a unit that could have been sold. The following is an equivalent definition for C_u : C_u is the gain from selling a unit. In other words, the underage cost is the per-unit opportunity cost of underordering. For the Hammer 3/2, $C_u = 70$. Note that the overage and underage costs are defined for a *single unit*. In other words, C_o is not the total cost of all leftover inventory; instead, C_o is the cost per unit of leftover inventory. The reason for defining C_o and C_u for a single unit is simple: We don’t know how many units will be left over in inventory, or how many units of demand will be lost, but we do know the cost of each unit left in inventory and the opportunity cost of each lost sale.

FIGURE 11.5
The Expected Gain
and Expected Loss
from the Q^{th}
Hammer 3/2 Ordered
by O’Neill



Now that we have defined the overage and underage costs, we need to choose Q to strike the balance between them that results in the maximum expected profit. Based on our previous reasoning, we should keep ordering additional units until the expected loss equals the expected gain.

The expected loss on a unit is the cost of having the unit in inventory (the overage cost) times the probability it is left in inventory. For the Q th unit, that probability is $F(Q)$: It is left in inventory if demand is less than Q .² Therefore, the expected loss is $C_o \times F(Q)$. The expected gain on a unit is the benefit of selling a unit (the underage cost) times the probability the unit is sold, which in this case occurs if demand is greater than Q . The probability demand is greater than Q is $(1 - F(Q))$. Therefore, the expected gain is $C_u \times (1 - F(Q))$.

It remains to find the order quantity Q that sets the expected loss on the Q th unit equal to the expected gain on the Q th unit:

$$C_o \times F(Q) = C_u \times (1 - F(Q))$$

If we rearrange terms in the above equation, we get

$$F(Q) = \frac{C_u}{C_o + C_u} \quad (11.1)$$

The profit-maximizing order quantity is the order quantity that satisfies the above equation. If you are familiar with calculus and would like to see a more mathematically rigorous derivation of the optimal order quantity, see Appendix D.

So how can we use equation (11.1) to actually find Q ? Let's begin by just reading it. It says that the order quantity that maximizes expected profit is the order quantity Q such that demand is less than or equal to Q with probability $C_u/(C_o + C_u)$. That ratio with the underage and overage costs is called the *critical ratio*. We now have an explanation for why our forecast must be a distribution function. To choose the profit-maximizing order quantity, we need to find the quantity such that demand will be less than that quantity with a particular probability (the critical ratio). The mean alone (i.e., just a sales forecast) is insufficient to do that task.

Let's begin with the easy part. We know for the Hammer 3/2 that $C_u = 70$ and $C_o = 20$, so the critical ratio is

$$\frac{C_u}{C_o + C_u} = \frac{70}{20 + 70} = 0.7778$$

We are making progress, but now comes the tricky part: We need to find the order quantity Q such that there is a 77.78 percent probability that demand is Q or lower.

Suppose we use the discrete distribution function from Table 11.3 as our demand forecast. Reading down the table, we see that $F(4,064) = 0.7576$ and $F(4,160) = 0.7879$. Therefore, the probability we are seeking, 0.7778, falls between two entries in the table, one that

²That statement might bother you. You might recall that $F(Q)$ is the probability demand is Q or lower. If demand is exactly Q , then the Q th unit will not be left in inventory. Hence, you might argue that it is more precise to say that $F(Q - 1)$ is the probability the Q th unit is left in inventory. However, the normal distribution assumes demand can be any value, including values that are not integers. If you are willing to divide each demand into essentially an infinite number of fractional pieces, as is assumed by the normal, then $F(Q)$ is indeed the probability there is leftover inventory. With the empirical distribution function, we get the same answer, but it requires a slightly more complex logic. If you are curious about the details, see Appendix D.

corresponds to an order quantity of 4,064 units and the other that corresponds to 4,160 units. What should we do? The rule is simple, which we will call the *round-up rule*:

Round-up rule. Whenever you are looking up a target value in a table and the target value falls between two entries, choose the entry that leads to the larger order quantity.

In this case, the larger quantity is 4,160 units. So if we use the discrete distribution function as our demand forecast, then the profit-maximizing order quantity is 4,160 units.

At this point you might not be entirely comfortable with the round-up rule. What if the critical ratio were 0.7577, which is just a hair above 0.7576? Your intuition might suggest ordering the lower entry value, 4,064 units. At the very least, maybe we should extrapolate. For example, suppose the critical ratio were 0.7728, which is half-way between 0.7576 and 0.7879. Extrapolation suggests ordering the average of the two quantities, 4,112 units. Nevertheless, there are two reasons why we should stick with the round-up rule. First, if the two quantities are reasonably close together (as they are here), then the difference in expected profit among these different order quantities is generally quite small. Second, in some cases extrapolation is not even feasible. For example, suppose the two quantities were 4 units and 5 units. It makes no sense to order 4.5 units. In those cases it can be shown that the higher order quantity always generates a higher expected profit. Thus, stick with the round-up rule and you will be fine. (If you really need to know more why the round-up rule is correct, see Appendix D.)

So we have identified the optimal order quantity if the discrete distribution function is our demand forecast. Let's next consider what the optimal order is if the normal distribution is our demand forecast. This process with the normal distribution is quite similar to the process with a discrete distribution function. The only difference is that it has one additional step.

First, let's find the optimal order quantity if the standard normal is our demand forecast. Because the Standard Normal Distribution Function Table is a table like any other discrete distribution function, we already know how to do this. Just as before, we need to find the z in the table that satisfies the following equation: $\Phi(z) = C_u/(C_o + C_u)$, which is $\Phi(z) = 0.7778$ for the Hammer 3/2. In the table we see that $\Phi(0.76) = 0.7764$ and $\Phi(0.77) = 0.7794$, so we use the round-up rule and pick the larger z value, 0.77. So the optimal order quantity is 0.77 if the demand forecast is a standard normal.

Now that we have an answer for the standard normal distribution, we need to convert that answer into an order quantity for the actual normal distribution. (This is the additional step.)

That conversion is done with the following equation:

$$Q = \mu + z \times \sigma$$

where

μ = Mean of the normal distribution

σ = Standard deviation of the normal distribution

With the Hammer 3/2 we use the above to obtain $Q = 3,192 + 0.77 \times 1,181 = 4,101$. Exhibit 11.4 summarizes how to find the optimal order quantity.

An alternative to the process we just described is to use Excel to find the optimal order quantity. Just as before, the first step is to find the optimal order quantity for a standard normal via Excel's Normsinv function:

$$z = \text{Normsinv}(\text{Critical ratio})$$

where the critical ratio is $C_u/(C_o + C_u)$. The second step is the same as before: Convert z into Q with the equation $Q = \mu + z \times \sigma$. Due to rounding issues, the Excel process might lead to a slightly different order quantity. For example, in Excel we get Normsinv(7/9) = 0.7647 and then $Q = 3,192 + 0.7647 \times 1,181 = 4,095.12$. The small difference between an order quantity of 4,101 and 4,095 should not matter much.

A PROCEDURE TO FIND THE ORDER QUANTITY THAT MAXIMIZES EXPECTED PROFIT IN THE NEWSVENDOR MODEL

Step 1: Evaluate the critical ratio: $\frac{C_u}{C_o + C_u}$. In the case of the Hammer 3/2, the underage cost is $C_u = \text{Price} - \text{Cost}$ and the overage cost is $C_o = \text{Cost} - \text{Salvage value}$.

Step 2: If the demand forecast is a normal distribution with mean μ and standard deviation σ , then follow steps A and B:

A. Find the optimal order quantity if demand had a standard normal distribution. One method to achieve this is to find the z value in the Standard Normal Distribution Function Table such that

$$\Phi(z) = \frac{C_u}{C_o + C_u}$$

(If the critical ratio value does not exist in the table, then find the two z values that it falls between. For example, the critical ratio 0.7778 falls between $z = 0.76$ and $z = 0.77$. Then choose the larger of those two z values.) A second method is to use the Excel function Normsinv: $z = \text{Normsinv}(\text{Critical ratio})$.

B. Convert z into the order quantity that maximizes expected profit, Q : $Q = \mu + z \times \sigma$

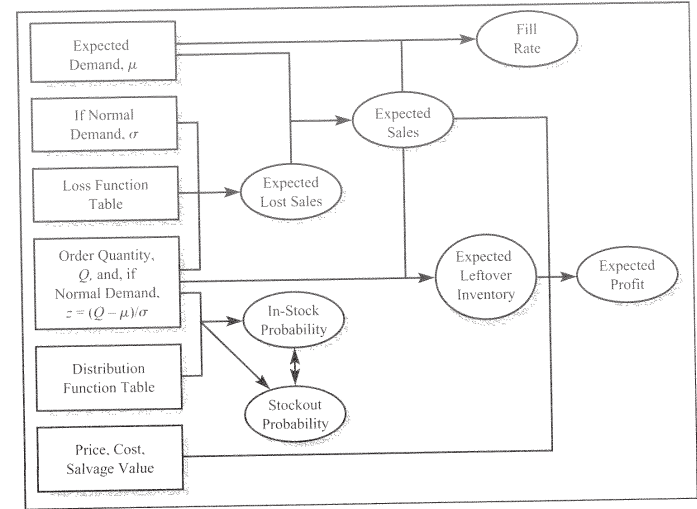
If the demand forecast is a discrete distribution function table (as with an empirical distribution function), then find the quantity in the table such that $F(Q) = \text{Critical ratio}$. If the critical ratio falls between two entries in the table, then choose the entry with the larger quantity.

To summarize this section, we determined the key equation (11.1) for finding the order quantity that maximizes expected profit: With the profit-maximizing order quantity, the critical ratio is the probability demand will be that quantity or lower. Using the discrete distribution function table (Table 11.3), we found our target probability (the critical ratio) fell between two entries in the table. We used the round-up rule to choose the larger quantity. We also evaluated the optimal order quantity if our demand forecast is a normal distribution: Look up in the Standard Normal Distribution Function Table the optimal order quantity z as if demand follows a standard normal distribution and then convert z into the order quantity for the actual normal distribution. The first step in that process is identical to the process we use with the discrete distribution function, that is, look up z in the Standard Normal Distribution Function Table just as you would look up Q in the discrete distribution function table. The second step (converting z into Q) is specific to the normal distribution.

11.5 Performance Measures

The previous section showed us how to find the order quantity that maximizes our expected profit. This section shows us how to evaluate a number of relevant performance measures. As Figure 11.6 indicates, these performance measures are closely related. For example, to evaluate expected leftover inventory, you first evaluate expected lost sales (which has up to three inputs: the order quantity, the loss function table, and the standard deviation of demand), then expected sales (which has two inputs: expected lost sales and expected demand), and then expected leftover inventory (which has two inputs: expected sales and the order quantity).

FIGURE 11.6
The Relationships between Initial Input Parameters (boxes) and Performance Measures (ovals)
Note: Some performance measures require other performance measures as inputs; for example, expected sales requires expected demand and expected lost sales as inputs.



These performance measures can be evaluated for any order quantity, not just the expected profit-maximizing order quantity. To emphasize this point, this section evaluates these performance measures assuming 3,500 Hammer 3/2s are ordered. See Table 12.1 for the evaluation of these measures with the optimal order quantity, 4,101 units.

Expected Lost Sales

Let's begin with *expected lost sales*, which is the expected number of units demand (a random variable) exceeds the order quantity (a fixed threshold). (Because order quantities are measured in physical units, sales and lost sales are measured in physical units as well, not in monetary units.) For example, if we order 3,500 units of the Hammer but demand could have been high enough to sell 3,821 units, then we would lose $3,821 - 3,500 = 321$ units of demand. Expected lost sales is the amount of demand that is not satisfied, which should be of interest to a manager even though the opportunity cost of lost sales does not show up explicitly on any standard accounting document.

Note that we are interested in the *expected lost sales*. Demand can be less than our order quantity, in which case lost sales is zero, or demand can exceed our order quantity, in which case lost sales is positive. Expected lost sales is the average of all of those events (the cases with no lost sales and all cases with positive lost sales).

How do we find expected lost sales for any given order quantity? Let's first suppose the discrete distribution function in Table 11.3 is our demand forecast. Table 11.5 replicates Table 11.3 and adds a column with the *loss function*, $L(Q)$, for each order quantity. The loss function of a random variable is the expected amount by which the outcome of the random variable exceeds a fixed threshold value, such as Q . For example, if O'Neill orders 2,592 Hammer 3/2s, then, according to Table 11.5, demand exceeds the order quantity by 841 units on average. In other words, O'Neill's expected lost sales is 841 units. To summarize, expected lost sales with an order quantity Q is exactly equal to the loss function of the demand distribution evaluated at Q . See the Statistics Tutorial in Appendix A for a more detailed description of the loss function.

TABLE 11.5
Hammer 3/2's
Empirical
Distribution and Loss
Function

Q	F(Q)	L(Q)	Q	F(Q)	L(Q)	Q	F(Q)	L(Q)
800	0.0303	2392	2,592	0.3636	841	3,936	0.6970	190
1,184	0.0606	2020	2,624	0.3939	821	4,000	0.7273	170
1,792	0.0909	1448	2,752	0.4242	743	4,064	0.7576	153
1,792	0.1212	1448	3,072	0.4545	559	4,160	0.7879	130
1,824	0.1515	1420	3,104	0.4848	542	4,352	0.8182	89
1,888	0.1818	1366	3,136	0.5152	525	4,544	0.8485	54
2,048	0.2121	1235	3,264	0.5455	463	4,672	0.8788	35
2,144	0.2424	1160	3,456	0.5758	376	4,800	0.9091	19
2,208	0.2727	1111	3,680	0.6061	281	4,928	0.9394	8
2,304	0.3030	1041	3,744	0.6364	256	4,992	0.9697	4
2,560	0.3333	863	3,808	0.6667	232	5,120	1.0000	0

Q = Order quantity
F(Q) = Probability demand is less than or equal to the order quantity
L(Q) = Loss function (the expected amount demand exceeds Q)

Table 11.5 is quite convenient because it provides us with the expected lost sales with little effort. In fact, with any discrete distribution function, we need a loss function table similar to Table 11.5. Appendix B provides the loss function for the Poisson distribution with different means. Appendix C provides a procedure to evaluate the loss function for any discrete distribution function. We relegate this procedure to the appendix because it is computationally burdensome; that is, it is the kind of calculation you want to do on a spreadsheet rather than by hand.

Now let's turn to the evaluation of the expected lost sales if our demand forecast for the Hammer 3/2 is the normal distribution with mean 3,192 and standard deviation 1,181. That process is actually not much different than what we have already done: First we look up in the Standard Normal Loss Function Table the expected lost sales for the z -statistic that corresponds to our quantity and then we convert the standard normal expected lost sales into the expected lost sales for our actual normal distribution.

Assume 3,500 Hammer 3/2s are ordered. First find the z -statistic that corresponds to $Q = 3,500$:

$$z = \frac{Q - \mu}{\sigma} = \frac{3,500 - 3,192}{1,181} = 0.26$$

Now use the Standard Normal Loss Function Table in Appendix B to look up the expected lost sales if the order quantity is $z = 0.26$; $L(0.26) = 0.2824$. To convert the standard normal expected lost sales into the expected lost sales with our actual normal distribution, use the following equation:

$$\text{Expected lost sales} = \sigma \times L(z)$$

where σ = Standard deviation of the normal distribution representing demand

$L(z)$ = Loss function with the standard normal distribution

Therefore, with 3,500 Hammer 3/2s we can expect to lose $1,181 \times 0.2824 = 334$ units of demand.

If you wish to avoid the Standard Normal Loss Function Table, then you can evaluate the loss function for the normal distribution in Excel with the following equation:

$$L(z) = \text{Normdist}(z, 0, 1, 0) - z * (1 - \text{Normdist}(z))$$

Exhibit 11.5

EXPECTED LOST SALES EVALUATION PROCEDURE

If the demand forecast is a normal distribution with mean μ and standard deviation σ , then follow steps A through D:

- Evaluate the z -statistic for the order quantity Q : $z = \frac{Q - \mu}{\sigma}$.
- Use the z -statistic to look up in the Standard Normal Loss Function Table the expected lost sales, $L(z)$, with the standard normal distribution.
- Expected lost sales = $\sigma \times L(z)$.
- With Excel, expected lost sales can be evaluated with the following equation:

$$\text{Expected lost sales} = \sigma * (\text{Normdist}(z, 0, 1, 0) - z * (1 - \text{Normdist}(z)))$$

If the demand forecast is a discrete distribution function table (as with an empirical distribution function), then expected lost sales equals the loss function for the chosen order quantity, $L(Q)$. If the table does not include the loss function, then see Appendix C for how to evaluate it.

(If you are curious about the derivation of the above function, see Appendix D.) The process of evaluating expected lost sales for both a discrete distribution and the normal distribution is summarized in Exhibit 11.5.

Expected Sales

Each unit of demand results in either a sale or a lost sale, so

$$\text{Expected sales} + \text{Expected lost sales} = \text{Expected demand}$$

We already know expected demand: It is the mean of the demand distribution, μ . Rearrange terms in the above equation and we get

$$\text{Expected sales} = \mu - \text{Expected lost sales}$$

Therefore, the procedure to evaluate expected sales begins by evaluating expected lost sales. See Exhibit 11.6 for a summary of this procedure.

Let's evaluate expected sales if 3,500 Hammers are ordered and the normal distribution is our demand forecast. We already evaluated expected lost sales to be 334 units. Therefore, $\text{Expected sales} = 3,192 - 334 = 2,858$ units.

Notice that expected sales is always less than expected demand (because expected lost sales is never negative). In other words, you can never expect to sell your demand forecast: While you might get lucky and sell more than the mean demand, on average you cannot sell more than the mean demand.

Expected Leftover Inventory

Expected leftover inventory is the average amount that demand (a random variable) is less than the order quantity (a fixed threshold). (In contrast, expected lost sales is the average amount by which demand exceeds the order quantity.)

The following equation is true because every unit purchased is either sold or left over in inventory at the end of the season:

$$\text{Expected sales} + \text{Expected leftover inventory} = Q$$

EXPECTED SALES, EXPECTED LEFTOVER INVENTORY, EXPECTED PROFIT, AND FILL RATE EVALUATION PROCEDURES

Step 1. Evaluate expected lost sales (see Exhibit 11.5). All of these performance measures can be evaluated directly in terms of expected lost sales and several known parameters: μ = Expected demand; Q = Order quantity; Price; Cost; and Salvage value.

Step 2. Use the following equations to evaluate the performance measure of interest.

$$\text{Expected sales} = \mu - \text{Expected lost sales}$$

$$\text{Expected leftover inventory} = Q - \text{Expected sales}$$

$$= Q - \mu + \text{Expected lost sales}$$

$$\text{Expected profit} = [(\text{Price} - \text{Cost}) \times \text{Expected sales}] \\ - [(\text{Cost} - \text{Salvage value}) \times \text{Expected leftover inventory}]$$

$$\text{Fill rate} = \text{Expected sales}/\mu = 1 - (\text{Expected lost sales}/\mu)$$

Rearrange the above equation to obtain

$$\text{Expected leftover inventory} = Q - \text{Expected sales}$$

We know the quantity purchased, Q . Therefore, we can easily evaluate expected leftover inventory once we have evaluated expected sales. See Exhibit 11.6 for a summary of this procedure.

If the demand forecast is a normal distribution and 3,500 Hammers are ordered, then expected leftover inventory is $3,500 - 2,858 = 642$ units because we evaluated expected sales to be 2,858 units.

It may seem surprising that expected leftover inventory and expected lost sales can both be positive. While in any particular season there is either leftover inventory or lost sales, but not both, we are interested in the expectation of those measures over all possible outcomes. Therefore, each *expectation* can be positive.

Expected Profit

We earn Price - Cost on each unit sold and we lose Cost - Salvage value on each unit we do not sell, so our expected profit is

$$\text{Expected profit} = [(\text{Price} - \text{Cost}) \times \text{Expected sales}] \\ - [(\text{Cost} - \text{Salvage value}) \times \text{Expected leftover inventory}]$$

Therefore, we can evaluate expected profit after we have evaluated expected sales and leftover inventory. See Exhibit 11.6 for a summary of this procedure.

With an order quantity of 3,500 units and a normal distribution demand forecast, the expected profit for the Hammer 3/2 is

$$\text{Expected profit} = (\$70 \times 2,858) - (\$20 \times 642) = \$187,220$$

Fill Rate

The fill rate is the percentage of demand that is satisfied:

$$\text{Fill rate} = \frac{\text{Expected sales}}{\text{Expected demand}} = \frac{\text{Expected sales}}{\mu}$$

IN-STOCK PROBABILITY AND STOCKOUT PROBABILITY EVALUATION

If the demand forecast is a normal distribution with mean μ and standard deviation σ , then follow steps A through D:

A. Evaluate the z-statistic for the order quantity: $z = \frac{Q - \mu}{\sigma}$.

B. Use the z-statistic to look up in the Standard Normal Distribution Function Table the probability the standard normal demand is z or lower, $\Phi(z)$.

C. In-stock probability = $\Phi(z)$ and Stockout probability = $1 - \Phi(z)$.

D. In Excel, In-stock probability = Normsdist(z) and Stockout probability = $1 - \text{Normsdist}(z)$.

If the demand forecast is a discrete distribution function table (as with an empirical distribution function), then In-stock probability = $F(Q)$ and Stockout probability = $1 - F(Q)$, where $F(Q)$ is the probability demand is Q or lower.

The fill rate is a measure of customer service: The higher the fill rate, the more likely a customer will find a unit available to purchase. See Exhibit 11.6 for the procedure to evaluate the fill rate. With an order of 3,500 Hammer 3/2s and the normal distribution demand forecast, O'Neill's fill rate is $2,858/3,192 = 89.5$ percent.

In-Stock Probability and Stockout Probability

The fill rate is one measure of customer service, but it is not the only measure of customer service. The in-stock probability is the probability the firm ends the season having satisfied all demand. (Equivalently, the in-stock probability is the probability the firm has stock available for every customer.) That occurs if demand is less than the order quantity.

$$\text{In-stock probability} = F(Q)$$

The stockout probability is the probability the firm stocks out for some customer during the selling season (i.e., a lost sale occurs). Because the firm stocks out if demand exceeds the order quantity,

$$\text{Stockout probability} = 1 - F(Q)$$

(The firm either stocks out or it does not, so the stockout probability equals 1 minus the probability demand is Q or lower.) We also can see that the stockout probability and the in-stock probability are closely related:

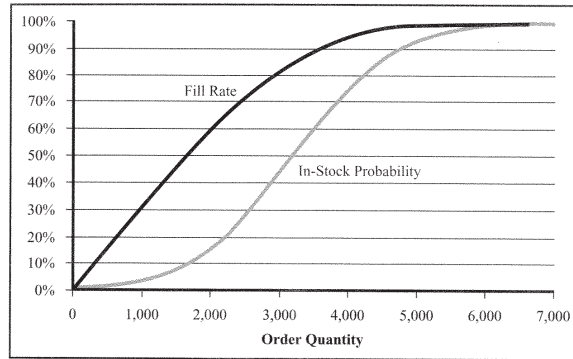
$$\text{Stockout probability} = 1 - \text{In-stock probability}$$

See Exhibit 11.7 for a summary of the procedure to evaluate these probabilities. With an order quantity of 3,500 Hammers, the z-statistic is $z = (3,500 - 3,192)/1,181 = 0.26$. From the Standard Normal Distribution Function Table, we find $\Phi(0.26) = 0.6026$, so the in-stock probability is 60.26 percent. The stockout probability is $1 - 0.6026 = 39.74$ percent.

Notice that O'Neill satisfies nearly 90 percent of its demand (the fill rate) with the order quantity 3,500, but stocks out with nearly a 40 percent probability and only has an in-stock probability of 60 percent. Figure 11.7 displays the in-stock probability and fill rate for a wide range of order quantities for the Hammer 3/2, assuming the normal distribution is the demand forecast.

How is it possible for the in-stock probability to be so low while the fill rate is so high? The in-stock probability is approximately the fill rate at the very end of the season: It is

FIGURE 11.7
In-Stock Probability
and Fill Rate for the
Hammer 3/2 with a
Normal Distribution
Demand Forecast



the probability that a customer's demand can be satisfied at the end of the season. Most of the customers during the season can be satisfied (resulting in a high fill rate) even if the poor customer at the end of the season experiences a stockout (resulting in a low in-stock probability).

This discussion raises the following question: Which is the better measure of customer service, the fill rate or the in-stock probability? The fill rate is a good measure of average customer service since it treats each customer as equally important. The in-stock probability is a conservative measure of customer service because it is the fill rate when the fill rate is expected to be at its worst, that is, at the end of the season. Therefore, neither measure dominates the other in all situations. A catalog retailer's customers probably do not expect a high fill rate late in the Christmas season, so a stockout the week before Christmas might not have a significant long-run impact on customer goodwill. However, a grocery retailer probably does not want to be out of stock on most items at the end of the day because certain customers are only able to do their grocery shopping at that time. Therefore, the fill rate may be more appropriate for the catalog retailer, whereas the in-stock probability may be more appropriate for the grocery retailer.

11.6 Other Objectives for Choosing an Order Quantity

Maximizing expected profit is surely a reasonable objective for choosing an order quantity, but it is not the only objective. As we saw in the previous section, the expected profit-maximizing order quantity may generate an unacceptable fill rate or in-stock probability from the firm's customer service perspective. This section explains how to determine an order quantity that satisfies a customer service objective, in terms of either a minimum fill rate or in-stock probability.

Let's begin with the objective of choosing an order quantity to achieve a target fill rate. For example, suppose O'Neill insists that all products have a 99 percent fill rate. In that case, we need to find the order quantity for the Hammer 3/2 that generates a 99 percent fill rate.

In the previous section, we started with an order quantity and evaluated the fill rate. Now we need to reverse that process, that is, start with a fill rate and end up with an order quantity. This requires a couple of algebraic gymnastics. Let's start with the fill rate equation,

$$\text{Fill rate} = \frac{\text{Expected sales}}{\mu}$$

Recall that

$$\text{Expected sales} = \mu - \text{Expected lost sales}$$

If we combine the above two equations and rearrange terms, we get

$$\text{Expected lost sales} = \mu \times (1 - \text{Fill rate}) \quad (11.2)$$

Suppose the discrete distribution is our demand forecast (Table 11.5). We can evaluate the right-hand side of the above equation: $(1 - 0.99) \times 3,192 = 31.92$. Therefore, we want to find a Q such that the expected lost sales with that quantity is 31.92 units. Looking at Table 11.5, we see that $L(4,672) = 36$ and $L(4,800) = 20$. Therefore, an order quantity of 4,672 would yield an expected lost sales slightly above our target (36 versus 31.92), while the larger order quantity of 4,800 would yield an expected lost sales slightly below our target (20 versus 31.92). We again shall follow our round-up rule: When the target value falls between two values in a table, choose the entry with the larger quantity. In this case, we round up to the order quantity 4,800. Therefore, if we order 4,800 Hammer 3/2s and the discrete distribution is our demand forecast, then our fill rate will be at least 99 percent.

Just to double-check our decision, we can evaluate the fill rate with the order quantity 4,800 and discover that we indeed make our target:

$$\text{Fill rate} = \frac{\text{Expected sales}}{\mu} = \frac{\mu - \text{Expected lost sales}}{\mu} = \frac{3,192 - 20}{3,192} = 99.37 \text{ percent}$$

How do we find an order quantity if our demand forecast is a normal distribution? Again, the process is quite similar, but we need to work with the standard normal. Remember that with a normal distribution

$$\text{Expected lost sales} = \sigma \times L(z)$$

If we combine the above equation with equation (11.2) and rearrange terms, we get

$$L(z) = \left(\frac{\mu}{\sigma} \right) \times (1 - \text{Fill rate})$$

Therefore, we first evaluate the right-hand side of the above equation, then look up $L(z)$ in the Standard Normal Loss Function Table, and then convert our z back into Q . The right-hand side of the above equation is $(3,192/1,181) \times (1 - 0.99) = 0.0270$. We now need to find a z -statistic such that $L(z) = 0.0270$. Looking in the Standard Normal Loss Function Table, we see that $L(1.53) = 0.0274$ and $L(1.54) = 0.0267$. Use our round-up rule again and choose $z = 1.54$, which generates $L(1.54) = 0.0267$. (Unfortunately, Excel does not provide a built-in function to perform this look up.) To finish our process, convert our z -statistic into an order quantity with the equation $Q = \mu + z \times \sigma$, so our order quantity is $3,192 + 1.54 \times 1,181 = 5,011$. In other words, if our demand forecast is the normal distribution, then we need to order 5,011 Hammer 3/2s to achieve a 99 percent fill rate. Exhibit 11.8 summarizes the process for finding an order quantity to satisfy a target fill rate.

As we have already discovered, the fill rate and the in-stock probability are not the same for any order quantity. Therefore, a firm may prefer the more conservative in-stock probability requirement. For example, let's now suppose O'Neill wants to find the order quantity that generates a 99 percent in-stock probability with the Hammer 3/2. It turns out that this objective is easier to work with than the target fill rate objective.

A PROCEDURE TO DETERMINE AN ORDER QUANTITY THAT SATISFIES A TARGET FILL RATE

If the demand forecast is a normal distribution with mean μ and standard deviation σ , then follow steps A through C:

A. Evaluate the target $L(z)$,

$$L(z) = \left(\frac{\mu}{\sigma}\right) \times (1 - \text{Fill rate})$$

B. Find the z -statistic in the Standard Normal Loss Function Table that corresponds to the target $L(z)$. If $L(z)$ falls between two z values in the table, choose the higher z value.

C. Convert the chosen z -statistic into the order quantity that satisfies our target fill rate, $Q = \mu + z \times \sigma$.

If the demand forecast is a discrete distribution function table (as with an empirical distribution function), then follow steps D and E:

D. Evaluate the target expected lost sales:

$$\text{Expected lost sales} = \mu \times (1 - \text{Fill rate})$$

E. Find the order quantity in the expected lost sales table that generates the target expected lost sales evaluated in step D. If the target expected lost sales fall between two entries in the table, choose the entry with the larger order quantity. If your table does not include expected lost sales, use the procedure in Appendix C to evaluate them.

The in-stock probability is $F(Q)$. So we need to find an order quantity such that there is a 99 percent probability that demand is that order quantity or lower. If the discrete distribution function is our demand forecast, then we look in Table 11.3 (or Table 11.5) and see $F(4,992) = 0.9697$ and $F(5,120) = 1.0000$. The round-up rule suggests that we need to order 5,120 Hammer 3/2s to achieve a 99 percent in-stock probability.

If our demand forecast is normally distributed, then we first find the z -statistic that achieves our objective with the standard normal distribution. In the Standard Normal Distribution Function Table, we see that $\Phi(2.32) = 0.9898$ and $\Phi(2.33) = 0.9901$. Again, we choose the higher z -statistic, so our desired order quantity is now $Q = \mu + z \times \sigma = 3,192 + 2.33 \times 1,181 = 5,944$. You can use Excel to avoid looking up a probability in the Standard Normal Distribution Function Table to find z :

$$z = \text{Normsinv}(\text{In-stock probability})$$

Notice that a substantially higher order quantity is needed to generate a 99 percent in-stock probability than a 99 percent fill rate (5,944 vs. 5,011) and both order quantities are substantially higher than the one that maximizes expected profit (4,101). Exhibit 11.9 summarizes the process for finding an order quantity to satisfy a target in-stock probability.

11.7 Managerial Lessons

Now that we have detailed the process of implementing the newsvendor model, it is worthwhile to step back and consider the managerial lessons it implies.

With respect to the forecasting process, there are three key lessons.

A PROCEDURE TO DETERMINE AN ORDER QUANTITY THAT SATISFIES A TARGET IN-STOCK PROBABILITY

If the demand forecast is a normal distribution with mean μ and standard deviation σ , then follow steps A and B:

A. Find the z -statistic in the Standard Normal Distribution Function Table that satisfies the in-stock probability, that is,

$$\Phi(z) = \text{In-stock probability}$$

If the in-stock probability falls between two z values in the table, choose the higher z . In Excel, z can be found with the following formula:

$$z = \text{Normsinv}(\text{In-stock probability}).$$

B. Convert the chosen z -statistic into the order quantity that satisfies our target in-stock probability,

$$Q = \mu + z \times \sigma$$

If the demand forecast is a discrete distribution function table (as with an empirical distribution function), then find the order quantity in the table such that $F(Q) = \text{In-stock probability}$. If the in-stock probability falls between two entries in the table, choose the entry with the larger order quantity.

- For each product, it is insufficient to have just a forecast of expected demand. We also need a forecast for how variable demand will be about the forecast. That uncertainty in the forecast is captured by the standard deviation of demand.

- It is important to track actual demand. Two common mistakes are made with respect to this issue. First, do not forget that actual demand may be greater than actual sales due to an inventory shortage. If it is not possible to track actual demand after a stockout occurs, then you should attempt a reasonable estimate of actual demand. Second, actual demand includes potential sales only at the regular price. If you sold 1,000 units in the previous season, but 600 of them were at the discounted price at the end of the season, then actual demand is closer to 400 than 1,000.

- You need to keep track of past forecasts and forecast errors in order to assess the standard deviation of demand. Without past data on forecasts and forecast errors, it is very difficult to choose reasonable standard deviations; it is hard enough to forecast the mean of a distribution, but forecasting the standard deviation of a distribution is nearly impossible with just a "gut feel." Unfortunately, many firms fail to maintain the data they need to implement the newsvendor model correctly. They might not record the data because it is an inherently undesirable task to keep track of past errors: Who wants to have a permanent record of the big forecasting goofs? Alternatively, firms may not realize the importance of such data and therefore do not go through the effort to record and maintain it.

There are also a number of important lessons from the order quantity choice process.

- The profit-maximizing order quantity generally does not equal expected demand. If the underage cost is greater than the overage cost (i.e., it is more expensive to lose a sale than it is to have leftover inventory), then the profit-maximizing order quantity is larger than expected demand. (Because then the critical ratio is greater than 0.50.) On the other

hand, some products may have an overage cost that is larger than the underage cost. For such products, it is actually best to order less than the expected demand.

- The order quantity decision should be separated from the forecasting process. The goal of the forecasting process is to develop the best forecast for a product's demand and therefore should proceed without regard to the order quantity decision. This can be frustrating for some firms. Imagine the marketing department dedicates considerable effort to develop a forecast and then the operations department decides to produce a quantity above the forecast. The marketing department may feel that their efforts are being ignored or their expertise is being second-guessed. In addition, they may be concerned that they would be responsible for ensuring that all of the production is sold even though their forecast was more conservative. The separation between the forecasting and the order quantity decision also implies that two products with the same mean forecast may have different expected profit-maximizing order quantities, either because they have different critical ratios or because they have different standard deviations.

- Explicit costs should not be overemphasized relative to opportunity costs. Inventory at the end of the season is the explicit cost of a demand-supply mismatch, while lost sales are the opportunity cost. Overemphasizing the former relative to the latter will cause you to order less than the profit-maximizing order quantity.

- It is important to recognize that choosing an order quantity to maximize expected profit is only one possible objective. It is also a very reasonable objective, but there can be situations in which a manager may wish to consider an alternative objective. For example, maximizing expected profit is wise if you are not particularly concerned with the variability of profit. If you are managing many different products so that the realized profit from any one product cannot cause undue hardship on the firm, then maximizing expected profit is a good objective to adopt. But if you are a startup firm with a single product and limited capital, then you might not be able to absorb a significant profit loss. In situations in which the variability of profit matters, it is prudent to order less than the profit-maximizing order quantity. The expected profit objective also does not consider customer service explicitly in its objective. With the expected profit-maximizing order quantity for the Hammer 3/2, the fill rate is about 95 percent and the in-stock probability is about 78 percent. Some managers may feel this is an unacceptable level of customer service, fearing that unsatisfied customers will switch to a competitor.

- Finally, while it is impossible to perfectly match supply and demand when supply must be chosen before random demand, it is possible to make a smart choice that balances the cost of ordering too much with the cost of ordering too little. In other words, uncertainty should not invite ad hoc decision making.

11.8 Summary

The newsvendor model is a tool for making a decision when there is a "too much-too little" challenge: Bet too much and there is a cost (e.g., leftover inventory), but bet too little and there is a different cost (e.g., the opportunity cost of lost sales). (See Table 11.6 for a summary of the key notation and equations.) To make this trade-off effectively, it is necessary to have a complete forecast of demand. It is not enough to just have a single sales forecast; we need to know the potential variation about that sales forecast.


In the case of O'Neill's Hammer 3/2 wetsuit, we discovered that there exists an order quantity that maximizes expected profit, but that quantity might not lead to a desirable level of customer service. In other words, there is a trade-off between profit and service. This trade-off is illustrated in Figure 11.8. While the newsvendor optimal quantity leads to an outcome in the upper-left-hand side of the trade-off curve, we also developed methods for moving down and to the right along the curve, that is, to a higher service level, albeit at a lower profit

TABLE 11.6
Summary of Key
Notation and
Equations in
Chapter 11

Q = Order quantity	C_u = Underage cost	C_o = Overage cost	Critical ratio = $\frac{C_u}{C_u + C_o}$
μ = Expected demand	σ = Standard deviation of demand		
$F(Q)$: Distribution function	$\Phi(Q)$: Distribution function of the standard normal		
Expected actual demand = Expected A/F ratio \times Forecast			
Standard deviation of actual demand = Standard deviation of A/F ratios \times Forecast			
Expected profit-maximizing order quantity: $F(Q) = \frac{C_u}{C_u + C_o}$			
z-statistic or normalized order quantity: $z = \frac{Q - \mu}{\sigma}$			
$Q = \mu + z \times \sigma$			
$L(z)$ = Expected lost sales with the standard normal distribution			
Expected lost sales = $\sigma \times L(z)$ Expected sales = $\mu -$ Expected lost sales			
Excel: Expected lost sales = $\sigma * (\text{Normdist}(z, 0, 1, 0) - z * (1 - \text{Normdist}(z)))$			
Expected leftover inventory = $Q -$ Expected sales			
Expected profit = [(Price - Cost) \times Expected sales] - [(Cost - Salvage value) \times Expected leftover inventory]			
Fill rate = Expected sales/ μ			
In-stock probability = $F(Q)$		Stockout probability = $1 -$ In-stock probability	
Excel: $z = \text{Normsinv}(\text{Target in-stock probability})$			
Excel: In-stock probability = $\text{Normsdist}(z)$			
To achieve a fill rate with normally distributed demand, $L(z) = (\mu/\sigma) \times (1 - \text{Fill rate})$; otherwise, target expected lost sales = $\mu \times (1 - \text{Fill rate})$			

level. We see that achieving a very high in-stock probability can be quite expensive, but because the curve is relatively flat at its top, it is possible to increase service without much of a profit sacrifice. This curve allows a manager to determine the correct position for the firm.

FIGURE 11.8
The Trade-off
between Profit and
Service with the
Hammer 3/2
The circle indicates the
in-stock probability
and the expected profit
of the optimal order
quantity, 4,101 units.



11.9 Further Reading

The newsvendor model is one of the most extensively studied models in operations management. It has been extended theoretically along numerous dimensions (e.g., multiple periods have been studied, the pricing decision has been included, the salvage values could depend on the quantity salvaged, the decision maker's tolerance for risk can be incorporated into the objective function, etc.) Several textbooks provide more technical treatments of the newsvendor model than this chapter. See Nahmias (2000), Porteus (2002), or Silver, Pyke, and Peterson (1998).

For a review of the theoretical literature on the newsvendor model, with an emphasis on the pricing decision in a newsvendor setting, see Petruzzi and Dada (1999).

11.10 Practice Problems

- Q11.1*** (McClure Books) Dan McClure owns a thriving independent bookstore in artsy New Hope, Pennsylvania. He must decide how many copies to order of a new book, *Power and Self-Destruction*, an exposé on a famous politician's lurid affairs. Interest in the book will be intense at first and then fizzle quickly as attention turns to other celebrities. The book's retail price is \$20 and the wholesale price is \$12. The publisher will buy back the retailer's leftover copies at a full refund, but McClure Books incurs \$4 in shipping and handling costs for each book returned to the publisher. Dan believes his demand forecast can be represented by a normal distribution with mean 200 and standard deviation 80.
- Dan will consider this book to be a blockbuster for him if it sells more than 400 units. What is the probability *Power and Self-Destruction* will be a blockbuster?
 - Dan considers a book a "dog" if it sells less than 50 percent of his mean forecast. What is the probability this exposé is a "dog"?
 - What is the probability demand for this book will be within 20 percent of the mean forecast?
 - What order quantity maximizes Dan's expected profit?
 - Dan prides himself on good customer service. In fact, his motto is "McClure's got what you want to read." What order quantity should he choose to satisfy a 95 percent fill rate?
 - How many books should Dan order if he wants to achieve a 95 percent in-stock probability?
 - If Dan orders the quantity chosen in part f to achieve a 95 percent in-stock probability, then what is the probability that "Dan won't have what some customer wants to read" (i.e., what is the probability some customer won't be able to purchase a copy of the book)?
 - Suppose Dan orders 300 copies of the book. What would Dan's expected profit be in this case?
 - Suppose Dan orders 150 copies of the book. What would Dan's fill rate be in this case?
 - Being an introspective chap, Dan decided to evaluate his forecasting skills in his spare time. He collected the following data on recent books he felt matched the characteristics of *Power and Self-Destruction*. If Dan used these data to construct an empirical distribution function, then what would be his optimal order quantity? (Assume Dan's initial forecast is 200 units.)

Book	Actual Demand	Forecast	A/F Ratio	Book	Actual Demand	Forecast	A/F Ratio
1	27	100	0.27	9	88	100	0.88
2	209	170	1.23	10	57	70	0.81
3	83	160	0.52	11	188	140	1.34
4	77	100	0.77	12	157	130	1.21
5	205	150	1.37	13	65	110	0.59
6	228	190	1.20	14	135	170	0.79
7	12	60	0.20	15	155	160	0.97
8	88	60	1.47	16	82	90	0.91

(* indicates that the solution is in Appendix E)

- Q11.2*** (EcoTable Tea) EcoTable is a retailer of specialty organic and ecologically friendly foods. In one of their Cambridge, Massachusetts, stores, they plan to offer a gift basket of Tanzanian teas for the holiday season. They plan on placing one order and any leftover inventory will be discounted at the end of the season. Expected demand for this store is 4.5 units and demand should be Poisson distributed. The gift basket sells for \$55, the purchase cost to EcoTable is \$32, and leftover baskets will be sold for \$20.

- If they purchase only 3 baskets, what is the probability that some demand will not be satisfied?
- If they purchase 10 baskets, what is the probability that they will have to mark down at least 3 baskets?
- How many baskets should EcoTable purchase to maximize its expected profit?
- Suppose they purchase 4 baskets. How many baskets can they expect to sell?
- Suppose they purchase 6 baskets. How many baskets should they expect to have to mark down at the end of the season?
- Suppose EcoTable wants to minimize its inventory while satisfying all demand with at least a 90 percent probability. How many baskets should they order?
- Suppose they wish to minimize their inventory investment while achieving at least a 90 percent fill rate. How many baskets should they order?
- Suppose EcoTable orders 8 baskets. What is its expected profit?

- Q11.3*** (Pony Express Creations) Pony Express Creations Inc. (www.pony-ex.com) is a manufacturer of party hats, primarily for the Halloween season. (80 percent of their yearly sales occur over a six-week period.) One of their popular products is the Elvis wig, complete with sideburns and metallic glasses. The Elvis wig is produced in China, so Pony Express must make a single order well in advance of the upcoming season. Ryan, the owner of Pony Express, expects demand to be 25,000 and the following is his entire demand forecast:

Q	Prob(D = Q)	F(Q)	L(Q)
5,000	0.0183	0.0183	20,000
10,000	0.0733	0.0916	15,092
15,000	0.1465	0.2381	10,550
20,000	0.1954	0.4335	6,740
25,000	0.1954	0.6289	3,908
30,000	0.1563	0.7852	2,052
35,000	0.1042	0.8894	978
40,000	0.0595	0.9489	425
45,000	0.0298	0.9787	170
50,000	0.0132	0.9919	63
55,000	0.0053	0.9972	22
60,000	0.0019	0.9991	8
65,000	0.0006	0.9997	4
70,000	0.0002	0.9999	2
75,000	0.0001	1.0000	2

Prob(D = Q) = Probability demand D equals Q
 F(Q) = Probability demand is Q or lower
 L(Q) = Expected lost sales if Q units are ordered

The Elvis wig retails for \$25, but Pony Express's wholesale price is \$12. Their production cost is \$6. Leftover inventory can be sold to discounters for \$2.50.

- Suppose Pony Express orders 40,000 Elvis wigs. What is the chance they have to liquidate 10,000 or more wigs with a discounter?

(* indicates that the solution is in Appendix E)

- What order quantity maximizes Pony Express's expected profit?
- If Pony Express wants to have a 90 percent fill rate, then how many Elvis wigs should be ordered?
- If Pony Express orders the quantity chosen in part c, what is Pony Express's actual fill rate?
- If Pony Express wants to have a 90 percent in-stock probability, then how many Elvis wigs should be ordered?
- If Pony Express orders 50,000 units, then how many wigs can they expect to have to liquidate with discounters?
- If Pony Express insists on a 100 percent in-stock probability for its customers, then what is its expected profit?

Q11.4* **(Flextrola)** Flextrola, Inc., an electronics systems integrator, is planning to design a key component for their next-generation product with Solectrics. Flextrola will integrate the component with some software and then sell it to consumers. Given the short life cycles of such products and the long lead times quoted by Solectrics, Flextrola only has one opportunity to place an order with Solectrics prior to the beginning of its selling season. Flextrola's demand during the season is normally distributed with a mean of 1,000 and a standard deviation of 600.

Solectrics' production cost for the component is \$52 per unit and it plans to sell the component for \$72 per unit to Flextrola. Flextrola incurs essentially no cost associated with the software integration and handling of each unit. Flextrola sells these units to consumers for \$121 each. Flextrola can sell unsold inventory at the end of the season in a secondary electronics market for \$50 each. The existing contract specifies that once Flextrola places the order, no changes are allowed to it. Also, Solectrics does not accept any returns of unsold inventory, so Flextrola must dispose of excess inventory in the secondary market.

- What is the probability that Flextrola's demand will be within 25 percent of its forecast?
 - What is the probability that Flextrola's demand will be more than 40 percent greater than its forecast?
 - Under this contract, how many units should Flextrola order to maximize its expected profit?
- For parts d through i, assume Flextrola orders 1,200 units.
- What are Flextrola's expected sales?
 - How many units of inventory can Flextrola expect to sell in the secondary electronics market?
 - What is Flextrola's expected gross margin percentage, which is $(\text{Revenue} - \text{Cost}) / \text{Revenue}$?
 - What is Flextrola's expected profit?
 - What is Solectrics' expected profit?
 - What is the probability that Flextrola has lost sales of 400 units or more?

j. A sharp manager at Flextrola noticed the demand forecast and became wary of assuming that demand is normally distributed. She plotted a histogram of demands from previous seasons for similar products and concluded that demand is better represented by the log normal distribution. Figure 11.9 plots the density function for both the log normal and the normal distribution, each with mean of 1,000 and standard deviation of 600. Figure 11.10 plots the distribution function for both the log normal and the normal. Using the more accurate forecast (i.e., the log normal distribution), approximately how many units should Flextrola order to maximize its expected profit?

Q11.5* **(Fashionables)** Fashionables is a franchisee of The Limited, the well-known retailer of fashionable clothing. Prior to the winter season, The Limited offers Fashionables the choice

(* indicates that the solution is in Appendix E)

FIGURE 11.9
Density Function


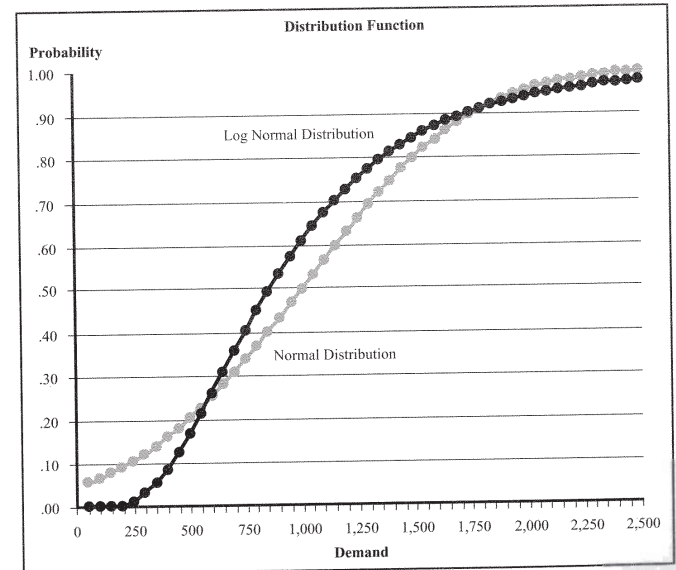


FIGURE 11.10
Distribution Function



of five different colors of a particular sweater design. The sweaters are knit overseas by hand, and because of the lead times involved, Fashionables will need to order its assortment in advance of the selling season. As per the contracting terms offered by The Limited, Fashionables also will not be able to cancel, modify, or reorder sweaters during the selling season. Demand for each color during the season is normally distributed with a mean of 500 and a standard deviation of 200. Further, you may assume that the demands for each sweater are independent of those for a different color.

The Limited offers the sweaters to Fashionables at the wholesale price of \$40 per sweater and Fashionables plans to sell each sweater at the retail price of \$70 per unit. The Limited delivers orders placed by Fashionables in truckloads at a cost of \$2,000 per truckload. The transportation cost of \$2,000 is borne by Fashionables. Assume unless otherwise specified that all the sweaters ordered by Fashionables will fit into one truckload. Also assume that all other associated costs, such as unpacking and handling, are negligible.

The Limited does not accept any returns of unsold inventory. However, Fashionables can sell all of the unsold sweaters at the end of the season at the fire-sale price of \$20 each.

- How many units of each sweater type should Fashionables order to maximize its expected profit?
- If Fashionables wishes to ensure a 97.5 percent in-stock probability, what should its order quantity be for each type of sweater?
- If Fashionables wishes to ensure a 97.5 percent fill rate, what should its order quantity be for each type of sweater?

For parts d through f, assume Fashionables orders 725 of each sweater.

- What is Fashionables' expected profit?
- What is Fashionables' expected fill rate for each sweater?
- What is the stockout probability for each sweater?
- Now suppose that The Limited announces that the unit of truckload capacity is 2,500 total units of sweaters. If Fashionables orders more than 2,500 units in total (actually, from 2,501 to 5,000 units in total), it will have to pay for two truckloads. What now is Fashionables' optimal order quantity for each sweater?

Q11.6 (Teddy Bower Parkas) Teddy Bower is an outdoor clothing and accessories chain that purchases a line of parkas at \$10 each from its Asian supplier, TeddySports. Unfortunately, at the time of order placement, demand is still uncertain. Teddy Bower forecasts that its demand is normally distributed with mean of 2,100 and standard deviation of 1,200. Teddy Bower sells these parkas at \$22 each. Unsold parkas have little salvage value; Teddy Bower simply gives them away to a charity.

- What is the probability this parka turns out to be a "dog," defined as a product that sells less than half of the forecast?
- How many parkas should Teddy Bower buy from TeddySports to maximize expected profit?
- If Teddy Bower wishes to ensure a 98.5 percent fill rate, how many parkas should it order?
- If Teddy Bower wishes to ensure a 98.5 percent in-stock probability, how many parkas should it order?

For parts e through g, assume Teddy Bower orders 3,000 parkas.

- Evaluate Teddy Bower's expected profit.
- Evaluate Teddy Bower's fill rate.
- Evaluate Teddy Bower's stockout probability.

Q11.7 (Teddy Bower Boots) To ensure a full line of outdoor clothing and accessories, the marketing department at Teddy Bower insists that they also sell waterproof hunting boots. Unfortunately, neither Teddy Bower nor TeddySports has expertise in manufacturing those kinds of boots. Therefore, Teddy Bower contacted several Taiwanese suppliers to request quotes. Due to competition, Teddy Bower knows that it cannot sell these boots for

more than \$54. However, \$40 per boot was the best quote from the suppliers. In addition, Teddy Bower anticipates excess inventory will need to be sold off at a 50 percent discount at the end of the season. Given the \$54 price, Teddy Bower's demand forecast is for 400 boots, with a standard deviation of 300.

- If Teddy Bower decides to include these boots in its assortment, how many boots should it order from its supplier?
- Suppose Teddy Bower orders 380 boots. What would its fill rate be?
- Suppose Teddy Bower orders 380 boots. What would its expected profit be?
- The marketing department will not be happy with the planned order quantity (from part a). They are likely to argue that Teddy Bower is a service-oriented company that requires a high fill rate. In particular, they insist that Teddy Bower order enough boots to have at least a 98 percent fill rate. What order quantity yields a 98 percent fill rate for Teddy Bower?
- What would Teddy Bower's expected profit be with the order quantity from part d?
- John Briggs, a buyer in the procurement department, overheard at lunch a discussion of the "boot problem." He suggested that Teddy Bower ask for a quantity discount from the supplier. After following up on his suggestion, the supplier responded that Teddy Bower could get a 10 percent discount if they were willing to order at least 800 boots. If the objective is to maximize expected profit, how many boots should it order given this new offer?
- After getting involved with the "boot problem," John Briggs became curious about using A/F ratios to forecast. He directed his curiosity to another product, Teddy Bower's standard hunting boot, which has a demand forecast for 1,000 units. This boot sells for \$55, and because of Teddy Bower's volume, the supplier of this boot only charges \$30. The standard hunting boot never goes out of style (therefore, all leftover boots will be sold next year, but it is a seasonal product). It costs Teddy Bower \$2.50 to hold a boot over from one season to the season in the following year. Furthermore, Teddy Bower anticipates that the selling price and procurement cost of this boot will be the same next year (i.e., \$55 and \$30, respectively). He collected the following data on 20 items that he felt were similar in nature to hunting boots. Using these data collected by John Briggs, what is Teddy Bower's profit-maximizing order quantity?

Actual				Actual			
Item	Demand	Forecast	A/F Ratio	Item	Demand	Forecast	A/F Ratio
1	2,512	2,041	1.23	11	1,317	1,667	0.79
2	1,003	916	1.09	12	366	1,216	0.30
3	32	264	0.12	13	1,009	1,266	0.80
4	829	1,471	0.56	14	1,501	778	1.93
5	95	1,946	0.05	15	1,918	1,599	1.20
6	2,122	1,184	1.79	16	2,306	2,042	1.13
7	165	418	0.39	17	2,058	1,170	1.76
8	769	1,514	0.51	18	794	1,607	0.49
9	1,120	595	1.88	19	552	323	1.71
10	762	872	0.87	20	638	801	0.80

Q 11.8 (Land's End) Geoff Gullo owns a small firm that manufactures "Gullo Sunglasses." He has the opportunity to sell a particular seasonal model to Land's End. Geoff offers Land's End two purchasing options:

- Option 1. Geoff offers to set his price at \$65 and agrees to credit Land's End \$53 for each unit Land's End returns to Geoff at the end of the season (because those units did not sell). Since styles change each year, there is essentially no value in the returned merchandise.
- Option 2. Geoff offers a price of \$55 for each unit, but returns are no longer accepted. In this case, Land's End throws out unsold units at the end of the season.

This season's demand for this model will be normally distributed with mean of 200 and standard deviation of 125. Land's End will sell those sunglasses for \$100 each. Geoff's production cost is \$25.

- How much would Land's End buy if they chose option 1?
- How much would Land's End buy if they chose option 2?
- Which option will Land's End choose?
- Suppose Land's End chooses option 1 and orders 275 units. What is Geoff Gullo's expected profit?

Q11.9 (Three Kings) Erica Zhang is the chief buyer for housewares at a large department store. Unfortunately, her store faces stiff competition from specialty retailers that carry imported cooking and dining articles. To meet this competitive challenge, Erica has reorganized her housewares department to create a "store within a store" that has the same ambiance as her competitors'. To kick off her concept, she plans a one-month promotion that features a sale on several special items, including an imported Three Kings baking dish. (The Three Kings baking dish leaves a "crown" image on the top of the cake.) These baking dishes must be ordered six months in advance. Any leftover inventory at the end of the promotion will be sold to a discount chain at a reduced price.

Erica has collected some pricing and cost data, listed in the following table, to help her decide how many of the Three Kings baking dishes to order. Erica predicts total demand for the imported baking dish to be normally distributed with mean of 980 and standard deviation of 354. Leftover dishes will be sold to the discount chain for \$15.

Three Kings baking dish:	
Selling price	\$40.00
Purchase price	\$16.00
Shipping cost	\$ 3.00
Handling cost*	\$ 0.80
Warehouse surcharge**	\$ 1.10
Total cost	\$20.90

*Estimate of variable cost to uncrate, clean, and transport a dish to the housewares department.

**Allocation of fixed overhead expenses in the shipping and receiving department.

For parts a through c, suppose 1,200 imported Three Kings baking dishes are ordered.

- What is the fill rate?
- What is the in-stock probability?
- How many Three Kings cake dishes should she purchase?
- Erica is concerned about customer service. Suppose she feels there is a loss of goodwill of \$10 for every customer that wants to purchase the Three Kings dish but is unable to do so due to a stockout. Now how many Three Kings cake dishes should she purchase?

Q 11.10 (CPG Bagels) CPG Bagels starts the day with a large production run of bagels. Throughout the morning, additional bagels are produced as needed. The last bake is completed at 3 p.m. and the store closes at 8 p.m. It costs approximately \$0.20 in materials and labor to make a bagel. The price of a fresh bagel is \$0.60. Bagels not sold by the end of the day are sold the next day as "day old" bagels in bags of six, for \$0.99 a bag. About two-thirds of the day-old bagels are sold; the remainder are just thrown away. There are many bagel flavors, but for simplicity, concentrate just on the plain bagels. The store manager predicts that demand for plain bagels from 3 p.m. until closing is normally distributed with mean of 54 and standard deviation of 21.

- How many bagels should the store have at 3 p.m. to maximize the store's expected profit (from sales between 3 p.m. until closing)? (*Hint:* Assume day-old bagels are sold for $\$0.99/6 = \0.165 each; i.e., don't worry about the fact that day-old bagels are sold in bags of six.)

- Suppose the manager would like to ensure at least a 99 percent fill rate on demand that occurs after 3 p.m. How many bagels should the store have at 3 p.m. to ensure that fill rate?
- Suppose that the store manager is concerned that stockouts might cause a loss of future business. To explore this idea, the store manager feels that it is appropriate to assign a stockout cost of \$5 per bagel that is demanded but not filled. (Customers frequently purchase more than one bagel at a time. This cost is per bagel demanded that is not satisfied rather than per customer that does not receive a complete order.) Given the additional stockout cost, how many bagels should the store have at 3 p.m. to maximize the store's expected profit?
- Suppose the store manager has 101 bagels at 3 p.m. How many bagels should the store manager expect to have at the end of the day?

Q 11.11 (The Kiosk) Weekday lunch demand for spicy black bean burritos at the Kiosk, a local snack bar, is approximately Poisson with a mean of 22. The Kiosk charges \$4.00 for each burrito, which are all made before the lunch crowd arrives. Virtually all burrito customers also buy a soda that is sold for 60¢. The burritos cost the Kiosk \$2.00, while sodas cost the Kiosk 5¢. Kiosk management is very sensitive about the quality of food they serve. Thus, they maintain a strict "No Old Burrito" policy, so any burrito left at the end of the day is disposed of. The distribution function of a Poisson with mean 22 is as follows:

Q	F(Q)	Q	F(Q)	Q	F(Q)	Q	F(Q)
1	0.0000	11	0.0076	21	0.4716	31	0.9735
2	0.0000	12	0.0151	22	0.5564	32	0.9831
3	0.0000	13	0.0278	23	0.6374	33	0.9895
4	0.0000	14	0.0477	24	0.7117	34	0.9936
5	0.0000	15	0.0769	25	0.7771	35	0.9962
6	0.0001	16	0.1170	26	0.8324	36	0.9978
7	0.0002	17	0.1690	27	0.8775	37	0.9988
8	0.0006	18	0.2325	28	0.9129	38	0.9993
9	0.0015	19	0.3060	29	0.9398	39	0.9996
10	0.0035	20	0.3869	30	0.9595	40	0.9998

- Suppose burrito customers buy their snack somewhere else if the Kiosk is out of stock. How many burritos should the Kiosk make for the lunch crowd?
- Suppose that any customer unable to purchase a burrito settles for a lunch of Pop-Tarts and a soda. Pop-Tarts sell for 75¢ and cost the Kiosk 25¢. (As Pop-Tarts and soda are easily stored, the Kiosk never runs out of these essentials.) Assuming that the Kiosk management is interested in maximizing profits, how many burritos should they prepare?