CHAPTER 5

Batching and Other Flow Interruptions: Setup Times and the Economic Order Quantity Model

So far we have considered processes in which one flow unit consistently enters the process and one flow unit consistently exits the process at fixed intervals of time, called the process cycle time. For example, in the scooter example of **Chapter 4**, we establish a cycle time of 3 minutes, which allows for a production of 700 scooters per week.

In an ideal process, a cycle time of 3 minutes would imply that every resource receives one flow unit as an input each 3-minute interval and creates one flow unit of output each 3-minute interval. Such a smooth and constant flow of units is the dream of any operations manager, yet it is rarely feasible in practice. There are several reasons for why the smooth process flow is interrupted, the most important ones being setups and variability in processing times or quality levels. The focus of this chapter is on setups, which are an important characteristic of batch-flow operations.

To discuss setups, we return to the Xootr production process. In particular, we consider the computer numerically controlled (CNC) milling machine that is responsible for making two types of parts on each Xootr—the steer support and 2 ribs (see **Figure 5.1**). The steer support attaches the Xootr's deck to the steering column, and the ribs help the deck support the weight of the rider. Once the milling machine starts producing one of these parts, it can produce them reasonably quickly. However, a considerable setup time, or changeover time, is needed before the production of each part type can begin. Our primary objective is to understand how setups like these influence the three basic performance measures of a process: inventory, flow rate, and flow time.

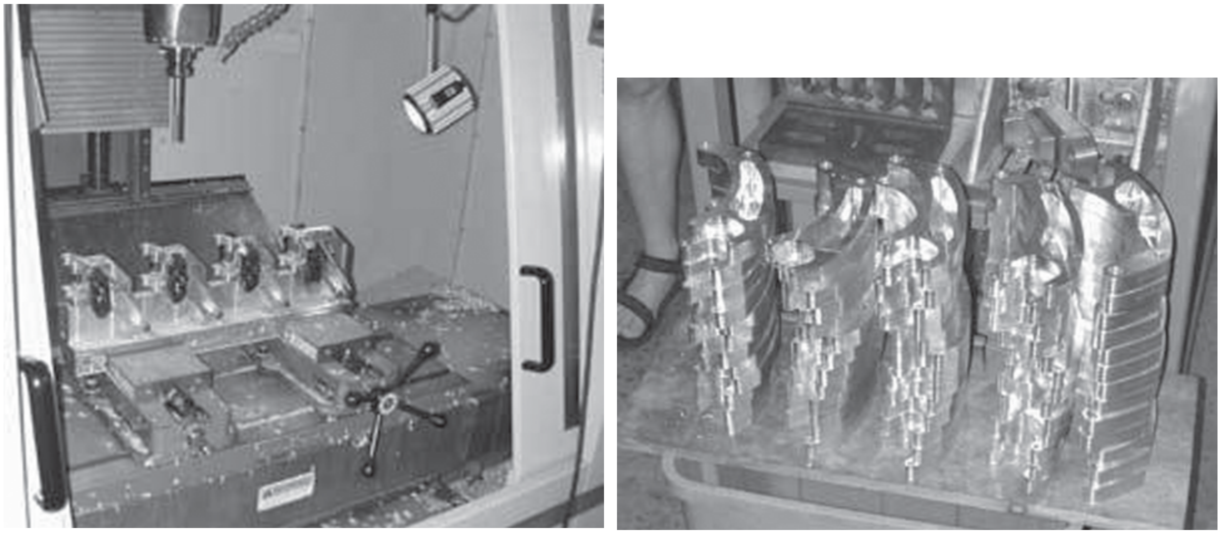

FIGURE 5.1 **Milling Machine (left) and Steer Support Parts (right)**

Karl Ulrich, Xootr LLC.

5.1 The Impact of Setups on Capacity

LO 5-1

Explain what batch production is and why increasing the batch size increases capacity.

To evaluate the capacity of the milling machine, we need some more information. Specifically, once set up to produce a part, the milling machine can produce steer supports at the rate of one per minute and can produce ribs at the rate of 2 per minute. Recall, each Xootr needs one steer support and 2 ribs. Furthermore, 1 hour is needed to set up the milling machine to start producing steer supports, and 1 hour is also needed to begin producing ribs. Although no parts are produced during those setup times, it is not quite correct to say that nothing is happening during those times either. The milling machine operator is busy calibrating the milling machine so that it can produce the desired part.

It makes intuitive sense that the following production process should be used with these two parts: set up the machine to make steer supports, make some steer supports, set up the machine to make ribs, make some ribs, and finally, repeat this sequence of setups and production runs. We call this repeating sequence a ***production cycle***: one production cycle occurs immediately after another, and all productions cycles "look the same" in the sense that they have the same setups and production runs.

We call this a batch production process because parts are made in batches. Although it may be apparent by what is meant by a "batch," it is useful to provide a precise definition:

A ***batch*** is a collection of flow units.

Throughout our analysis, we assume that batches are produced in succession. That is, once the production of one batch is completed, the production of the next batch begins and all batches contain the same number and type of flow unit.

Given that a batch is a collection of flow units, we need to define our flow unit in the case of the Xootr. Each Xootr needs 1 steer support and 2 ribs, so let's say the flow unit is a “component set” and each component set is composed of those three parts. Hence, each production cycle produces a batch of component sets.

One might ask why we did not define the flow unit to be one of the two types of parts. For example, we could call the steering supports made in a production run a batch of steering supports. However, our interest is not specifically on the capacity to make steering supports or ribs in isolation. We care about the capacity for component sets because one component set is needed for each Xootr. Thus, for the purpose of this analysis, it makes more sense to define the flow unit as a component set and to think in terms of a batch of component sets.

Because no output is produced while the resource is in setup mode, it is fairly intuitive that frequent setups lead to lower capacity. To understand how setups reduce the capacity of a process, consider **Figure 5.2**. As nothing is produced at a resource during setup, the more frequently a resource is set up, the lower its capacity. As discussed above, the milling machine underlying the example in **Figure 5.2** has the following processing times/setup times:

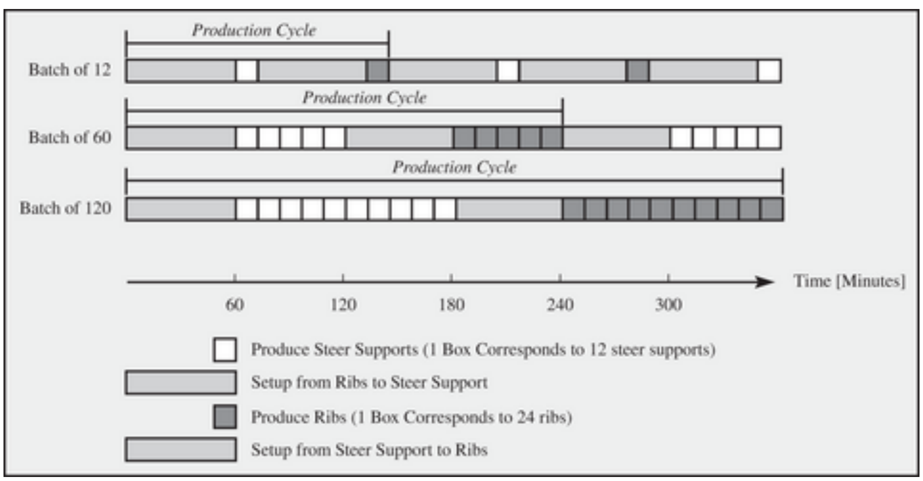

FIGURE 5.2 The Impact of Setup Times on Capacity

- It takes 1 minute to produce one steer support unit (of which there is one per Xootr).
- It takes 60 minutes to change over the milling machine from producing steer supports to producing ribs (setup time).
- It takes 0.5 minute to produce 1 rib; because there are 2 ribs in a Xootr, this translates to 1 minute/per pair of ribs.
- Finally, it takes another 60 minutes to change over the milling machine back to producing steer supports.

Now consider the impact that varying the batch size has on capacity. Recall that we defined capacity as the maximum flow rate at which a process can operate. If we produce in small batches of 12 component sets per batch, we spend a total of 2 hours of setup time (1 hour to set up the production for steer supports and 1 hour to set up the production of ribs) for every 12 component sets we produce. These 2 hours of setup time are lost from regular production.

The capacity of the resource can be increased by increasing the batch size. If the machine is set up every 60 units, the capacity-reducing impact of setup can be spread out over 60 units. This results in a higher capacity for the milling machine. Specifically, for a batch size of 60, the milling machine could produce at 0.25 component set per minute. **Table 5.1** summarizes the capacity calculations for batch sizes of 12, 60, and 120.

TABLE 5.1 **The Impact of Setups on Capacity**

Batch Size	Time to Complete One Batch [minutes]	Capacity [units/minute]
12	60 minutes (set up steering support) + 12 minutes (produce steering supports) + 60 minutes (set up ribs) + 12 minutes (produce ribs) 144 minutes	12/144 = 0.0833
60	60 minutes (set up steering support) + 60 minutes (produce steering supports) + 60 minutes (set up ribs) + 60 minutes (produce ribs) 240 minutes	60/240 = 0.25
120	60 minutes (set up steering support) + 120 minutes (produce steering supports) + 60 minutes (set up ribs) + 120 minutes (produce ribs) 360 minutes	120/360 = 0.333

Generalizing the computations in **Table 5.1**, we can compute the capacity of a resource with setups as a function of the batch size:

$$\text{Capacity given batch size} = \frac{\text{Batch size}}{\text{Setup time} + \text{Batch size} \times \text{Processing time}}$$

Basically, the above equation is spreading the "unproductive" setup time over the members of a batch. To use the equation, we need to be precise about what we mean by batch size, the setup time, and processing time:

- The batch size is the number of flow units that are produced in one production cycle (i.e., before the process repeats itself, see **Figure 5.2**).
- The setup time includes all setups in the production of the batch (i.e., all of the setups in the production cycle). In this case, this includes S = 60 minutes + 60 minutes = 120 minutes. It can also include any other nonproducing time associated with the production of the batch. For example, if the production of each batch requires a 10-minute worker break, then that would be included. Other "setup times" can include scheduled maintenance or forced idle time (time in which literally nothing is happening with the machine—it is neither producing nor being prepped to produce).
- The processing time includes all production time that is needed to produce one complete flow unit of output at the milling machine. In this case, this includes 1 minute/unit for the steer support as well as 2 times 0.5 minute/unit for the ribs. The processing time is thus p = 1 minute/unit + 2 × 0.5 minute/unit = 2 minutes/unit. Notice that the processing time is 2 minutes even though no single component set is actually produced over a single period of 2 minutes of length. Due to setups, the processing time for a component set is divided over two periods of 1 minute each, and those two periods can be separated by a considerable amount of time. Nevertheless, from the perspective of calculating the capacity of the milling machine when operated with a given batch size, it does not matter whether each component set is produced over a continuous period of time or disjointed periods of time. All that matters is that a total of 2 minutes is needed for each component set.

Given these definitions, say we operate with a batch size of 100 units. Our capacity in this case would be

$$\begin{aligned}
\text{Capacity (for } B = 100) &= \frac{\text{Batch size}}{\text{Setup time} + \text{Batch size} \times \text{Processing time}} \\
&= \frac{100 \text{ units}}{120 \text{ minutes} + 100 \text{ units} \times 2 \text{ minutes/unit}} \\
&= 0.3125 \text{ unit/minute}
\end{aligned}$$

No matter how large a batch size we choose, we never are able to produce faster than 1 unit every p units of time. Thus, $1/p$ can be thought of as the maximum capacity the process can achieve. This is illustrated in **Figure 5.3**.

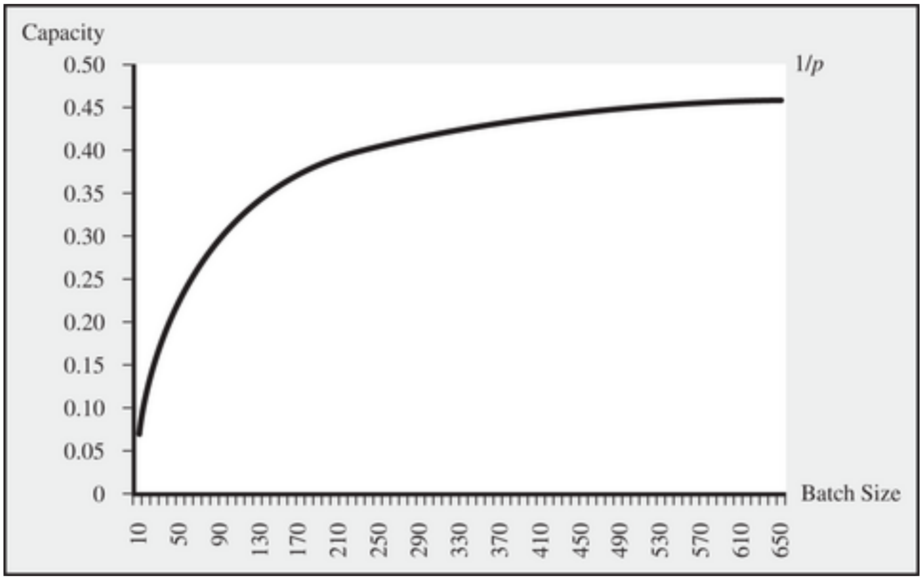

FIGURE 5.3 **Capacity as a Function of the Batch Size**

5.2 Interaction between Batching and Inventory

LO 5-2

Describe the impact of batch size on inventory.

Given the desirable effect that large batch sizes increase capacity, why not choose the largest possible batch size to maximize capacity? While large batch sizes are desirable from a capacity perspective, they typically require a higher level of inventory, either within the process or at the finished goods level. Holding the flow rate constant, we can infer from Little's Law that such a higher inventory level also leads to longer flow times. This is why batch-flow operations generally are not very fast in responding to customer orders (remember the last time you bought custom furniture?).

The interaction between batching and inventory is illustrated by the following example. Consider an auto manufacturer producing a sedan and a station wagon on the same assembly line. For simplicity, assume both models have the same demand rate, 400 cars per day each. The metal stamping steps in the process preceding final assembly are characterized by especially long setup times. Thus, to achieve a high level of capacity, the plant runs large production batches and produces sedans for 8 weeks, then station wagons for 8 weeks, and so on.

The production schedule results in lumpy output of sedans and station wagons, but customers demand sedans and station wagons at a constant rate (say). Hence, producing in large batches leads to a mismatch between the rate of supply and the rate of demand.

To make this schedule work, in addition to producing enough to cover demand over the 8 weeks of production, the company needs to also produce enough cars to satisfy demand in the subsequent 8 weeks while it is producing the other type of car. Assuming 5 days per week, that means that when sedan production finishes, there needs to be 400 cars per day × 5 days per week × 8 weeks = 16,000 sedans in inventory. Those 16,000 sedans are sold off at a constant rate of 2,000 cars per week for the 8 weeks while station wagons are made. On average there are 8,000 sedans in inventory. The same applies to station wagons—when production of station wagons ends there needs to be 16,000 of them in inventory and they are then depleted over the subsequent 8 weeks, leaving an average of 8,000 station wagons in inventory. This pattern of inventory rising and falling is illustrated in the left side of **Figure 5.4**.

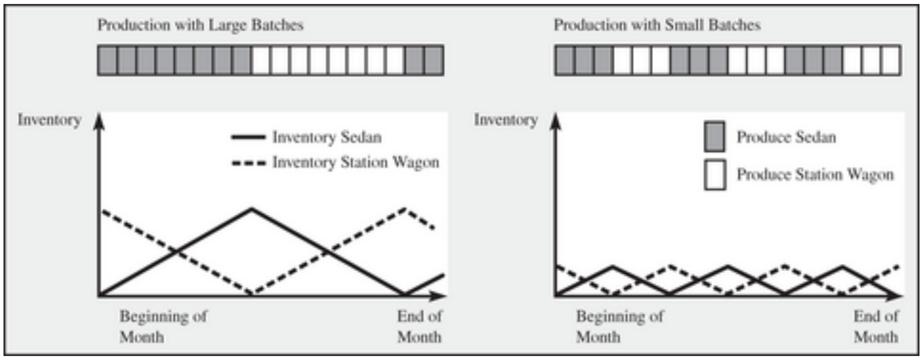

FIGURE 5.4 The Impact of Batch Sizes on Inventory

Instead of producing one type of car for 8 weeks before switching, the company may find that it is feasible to produce cars for 3 weeks before switching. Now at the end of each production run the company needs 400 cars per day × 5 days per week × 3 weeks = 6,000 cars to cover demand over the 3 weeks the other product is produced. That means that the average inventory of each type is only 3,000 cars, which is dramatically lower than the inventory needed with the 8-week schedule. This is illustrated in the right side of **Figure 5.4**. Thus, smaller batches translate to lower inventory levels!

In the ideal case, which has been propagated by the Toyota Production Systems (see **Chapter 8**) under the word ***heijunka*** or *mixed-model* production, the company would alternate between producing one sedan and producing one station wagon, thereby producing in batch sizes of one. This way, a much better synchronization of the demand flow and the production flow is achieved, and the inventory is basically eliminated.

Now let's turn our attention back to the milling machine at Nova Cruz. Similar to **Figure 5.4**, we can draw the inventory of components (ribs and steer supports) over the course of a production cycle. Remember that the assembly process following the milling machine requires a supply of 1 unit every 3 minutes. This 1 unit consists, from the view of the milling machine, of 2 ribs and a steer support unit. If we want to ensure a sufficient supply to keep the assembly process operating, we have to produce a sufficient number of ribs such that during the time we do not produce ribs (e.g., setup time and production of steer support) we do not run out of ribs. Say the milling machine operates with a batch size of 200 units, $B = 200$. In that case, the inventory of ribs changes as follows:

- During the production of ribs, inventory accumulates. As we produce 1 rib pair per minute, but assembly takes only 1 rib pair every 3 minutes, rib inventory accumulates at the rate of 2 rib pairs every 3 minutes, or 2/3 rib pairs per minute.
- The production of 200 rib pairs requires 200 minutes. Hence, the inventory of rib pairs at the end of the production run is 200 minutes × 2/3 rib pairs per minute = 133.3 rib pairs (i.e., 266 ribs).

The resulting production plan as well as the corresponding inventory levels are summarized in **Figure 5.5**. Notice that each production cycle takes 200 scooters × 3 minutes per scooter = 600 minutes, and this includes 80 minutes of idle time. Why is there idle time in the milling machine's production schedule? The answer is that without the idle time, the milling machine would produce too quickly given our batch size of 200 units. To explain, assembly takes 600 minutes to produce a batch of 200 scooters but the milling machine only needs 520 minutes to produce that batch of 600 scooters (120 minutes of setup and 400 minutes of production). Hence, if the milling machine produced one batch after another (without any idle time between them), it would produce 200 component sets every 520 minutes (or 200/520 = 0.3846 component set per minute), which is faster than assembly can use them (which is 1/3 component sets per minute). This analysis suggests that maybe we want to choose a different batch size, as we see in the next section.

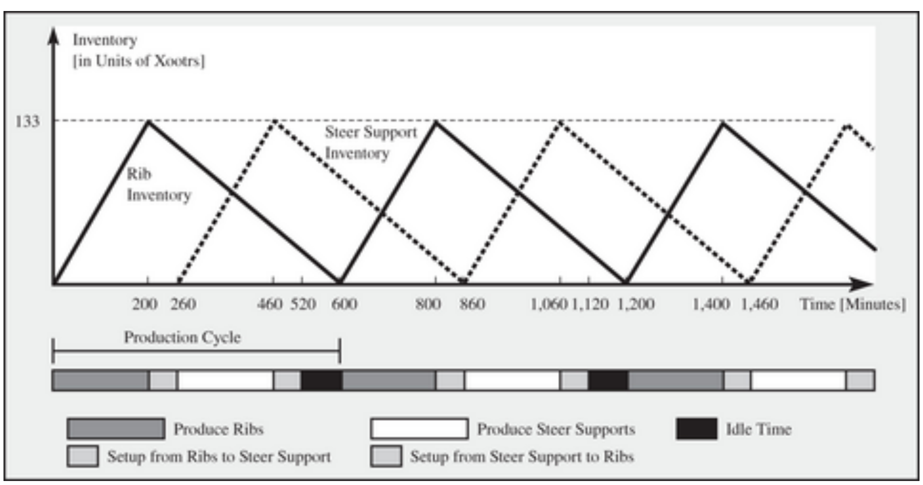

FIGURE 5.5 The Impact of Setup Times on Capacity

Figure 5.5 helps us to visualize the pattern of inventory for both rib pairs and steer supports. We see that the inventory of rib pairs makes a “sawtooth” pattern over time, with a minimum of 0 and a maximum of 133.3. If we were to average over all of the inventory levels, we would discover the average inventory to be 133.3/2 = 66.7. (The average across a triangle is half of its height.) But using a graph is not an efficient way to evaluate average inventory for each item in the production schedule. A better approach is to use the equation

$$\text{Average inventory} = \frac{1}{2}\text{Batch size }(1 - \text{Flow rate} \times \text{Processing time})$$

Page 88

In our case, the batch size is 200 rib pairs, the flow rate is 1/3 rib pairs per minute, and the processing time is 1 minute per rib pair. Hence,

$$\begin{aligned}\text{Average inventory} &= \tfrac{1}{2}200 \text{ rib pairs} \times (1 - (1/3 \text{ rib pairs per min} \times 1 \text{ min per rib pair})\\ &= 66.7 \text{ rib pairs}\end{aligned}$$

We see that the equation’s answer matches what we found from the graph, as it should.

It is essential to emphasize that when using the inventory equation we must be consistent with units. In particular, if we want to evaluate the average inventory of rib pairs, then batch size, flow rate, and processing time must all be given in terms of rib pairs. It makes no sense to define the batch size and flow rate in rib pairs but the processing time in component sets. Furthermore, we can’t use the above equation to evaluate the inventory of a set of parts, such as a component set, because the sum of saw-toothed inventory patterns is no longer saw-toothed. To evaluate inventory we must consider each possible part individually. For example, we can evaluate the average inventory of ribs and the average inventory of steer supports, and then we can add those two averages together. But the shortcut of trying to the evaluate inventory of all parts, all at once doesn’t work. Finally, the above inventory equation only applies if the batch size is sufficiently large that we never run out of inventory for an extended period of time, that is, there are no flat zones in the graph in **Figure 5.5**. As we see in the next section, we will generally want to operate with such a sufficiently large batch size.

We can end this section by repeating the key observation that larger batches lead to more inventory, which is readily apparent in our average inventory equation. Thus, if we want to reduce inventory, we need to operate with smaller batches.

5.3 Choosing a Batch Size in the Presence of Setup Times

LO 5-3

Determine how to choose an appropriate batch size for a process flow.

When choosing an appropriate batch size for a process flow, it is important to balance the conflicting objectives: capacity and inventory. Large batches lead to large inventory but more capacity; small batches lead to losses in capacity but less inventory.

To balance the conflict between our desire for more capacity and less inventory, we benefit from the following two observations:

- Capacity at the bottleneck step is extremely valuable (as long as the process is capacity-constrained; i.e., there is more demand than capacity) as it constrains the flow rate of the entire process.
- Capacity at a nonbottleneck step is free, as it does not provide a constraint on the current flow rate.

This has direct implications for choosing an appropriate batch size at a process step with setups.

- If the setup occurs at the bottleneck step (and the process is capacity-constrained), it is desirable to increase the batch size, as this results in a larger process capacity and, therefore, a higher flow rate.
- If the setup occurs at a nonbottleneck step (or the process is demand-constrained), it is desirable to decrease the batch size, as this decreases inventory as well as flow time.

The scooter example summarized in **Figure 5.6** illustrates these two observations and how they help us in choosing a good batch size. Remember that *B* denotes the batch size, *S* the setup time, and *p* the per unit processing time.

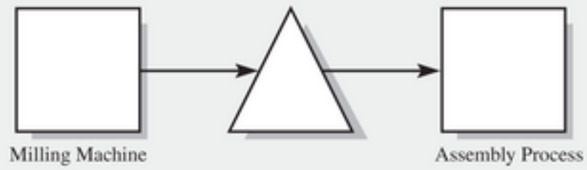

	Milling Machine	Assembly Process
Setup Time, *S*	120 Minutes	—
Per-Unit Time, *p*	2 Minutes/Unit	3 Minutes/Unit
Capacity (B = 12)	0.0833 Unit/Minute	0.33 Unit/Minute
Capacity (B = 300)	0.4166 Unit/Minute	0.33 Unit/Minute

FIGURE 5.6 Data from the Scooter Case about Setup Times and Batching

The process flow diagram in **Figure 5.6** consists of only two activities: the milling machine and the assembly operations. We can combine the assembly operations into one activity, as we know that its slowest step (bottleneck of assembly) can create 1 Xootr every 3 minutes, which is therefore assembly's *processing time*. The capacity of the assembly operation is 1/*Processing time*, so its capacity is $\frac{1}{3}$ unit per minute.

Let's evaluate this process with two different batch sizes. First, say B = 12. The capacity of the milling machine can be evaluated with the formula

$$\begin{aligned}\text{Capacity}\,(B) &= \frac{\text{Batch size}}{\text{Setup time} + \text{Batch size} \times \text{Processing time}} \\ &= \frac{B}{S + B \times p} = \frac{12}{120 + 12 \times 2} = 0.0833\,\text{unit/minute}\end{aligned}$$

With B = 12, the milling machine is the bottleneck because its capacity (0.0833 unit/minute) is lower than the capacity of assembly (0.3333 unit/minute).

Next, consider what happens to the same calculations if we increase the batch size from 12 to 300. While this does not affect the capacity of the assembly operations, the capacity of the milling machine now becomes

$$\text{Capacity}\,(B) = \frac{B}{S\ +\ B\ \times\ p} = \frac{300}{120\ +\ 300\ \times\ 2} = 0.4166\,\text{unit/minute}$$

Thus, we observe that the location of the bottleneck has shifted from the milling machine to the assembly operation; with B = 300 the milling machine's capacity (0.4166 unit/minute) now exceeds assembly's capacity (0.3333 unit/minute). Just by modifying the batch size, we can change which activity is the bottleneck! Now, which of the two batch sizes is the "better" one, 12 or 300?

- The batch size of 300 is too large. A smaller batch size would reduce inventory but as long as assembly remains the bottleneck, the smaller batch size does not lower the process's flow rate.
- The batch size of 12 is probably too small. As long as demand is greater than 0.0833 unit per minute (the milling machine's capacity with $B = 12$), a larger batch size can increase the flow rate of the process. It would also increase inventory, but the higher flow rate almost surely justifies a bit more inventory.

As a batch size of 12 is too small and a batch size of 300 is too large, a good batch size is "somewhere in between." Specifically, we are interested in the smallest batch size that does not adversely affect process capacity.

To find this number, we equate the capacity of the step with the setup (in this case, the milling machine) with the capacity of the step from the remaining process that has the smallest capacity (in this case, the assembly operations):

$$\frac{B}{120 + B \times 2} = \frac{1}{3}$$

and solve this equation for *B:*

$$\begin{aligned} \tfrac{B}{120 + B \times 2} &= \tfrac{1}{3} \\ 3 \times B &= 120 + 2 \times B \\ B &= 120 \end{aligned}$$

which gives us, in this case, $B = 120$. This algebraic approach is illustrated in **Figure 5.7**. If you feel uncomfortable with the algebra outlined above (i.e., solving the equation for the batch size B), or you want to program the method directly into Excel or another software package, you can use the following equation:

$$\text{Recommended batch size} = \frac{\text{Flow rate} \times \text{Setup time}}{1 - \text{Flow rate} \times \text{Processing time}}$$

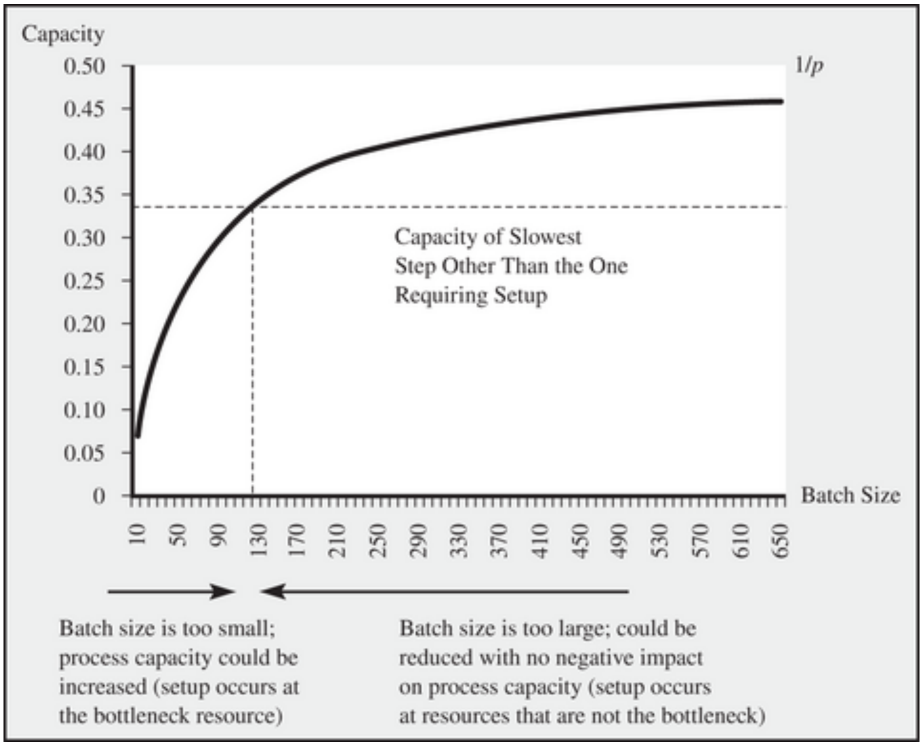

FIGURE 5.7 Choosing a "Good" Batch Size

which is equivalent to the analysis performed above. To see this, simply substitute Setup time = 120 minutes, Flow rate = 0.333 unit per minute, and Processing time = 2 minutes per unit and obtain

$$\text{Recommended batch size} = \frac{\text{Flow rate} \times \text{Setup time}}{1 - \text{Flow rate} \times \text{Processing time}} = \frac{0.333 \times 120}{1 - 0.333 \times 2} = 120$$

Figure 5.7 shows the capacity of the process step with the setup (the milling machine), which increases with the batch size *B,* and for very high values of batch size *B* approaches $1/p$ (similar to the graph in **Figure 5.3**). As the capacity of the assembly operation does not depend on the batch size, it corresponds to a constant (flat line).

The overall process capacity is—in the spirit of the bottleneck idea—the minimum of the two graphs. Thus, before the graphs intersect, the capacity is too low and the flow rate is potentially given up. After the intersection point, the assembly operation is the bottleneck and any further increases in batch size yield no return. **Exhibit 5.1** provides a summary of the computations leading to the recommended batch size in the presence of setup times.

Exhibit 5.1

FINDING A GOOD BATCH SIZE IN THE PRESENCE OF SETUP TIMES

1. Compute Flow rate = Minimum {Available input, Demand, Process capacity}.
2. Define the production cycle, which includes the processing and setups of all flow units in a batch. Let B be the number of units produced in the production cycle.
3. Compute the total time in a production cycle that the resource is in setup; setup times are those times that are independent of the batch size. Call this total the *Setup time*.
4. Compute the total time in a production cycle to process a single unit. If a single unit has multiple parts, then sum the times to process each of the parts. Call this total the *Processing time*.
5. Compute the capacity of the resource with setup for a given batch size:

$$\text{Capacity}\,(B) = \frac{B}{\text{Setup time} + B \times \text{Processing time}}$$

6. We are looking for the batch size that leads to the lowest level of inventory without affecting the flow rate; we find this by solving the equation

$$\text{Capacity}\,(B) = \text{Flow rate}$$

for the batch size B. This also can be done directly using the following formula:

$$\text{Recommended batch size} = \frac{\text{Flow rate} \times \text{Setup time}}{1 - \text{Flow rate} \times \text{Processing time}}$$

5.4 Setup Times and Product Variety

LO 5-4

Discuss the impact of product variety on a process with setup times.

As we have seen in the case of the Xootr production process, setup times often occur due to the need to change over production from one product to another. This raises the following question: What is the impact of product variety on a process with setup times? To explore this question, let's consider a simple process that makes two kinds of soup: chicken noodle and tomato.

Demand for chicken soup is 100 gallons per hour, while demand for tomato soup is 75 gallons per hour. Switching from one type of soup to another requires 30 minutes to clean the production equipment so that one flavor does not disrupt the flavor of the next soup. Once production begins, the process can make 300 gallons per hour of either type of soup. Given these parameters, let's evaluate a production cycle that minimizes inventory while satisfying demand.

We first need to define our flow unit. In this case, it is natural to let our flow unit be 1 gallon of soup. Hence, a production cycle of soup contains a certain number of gallons, some chicken and some tomato. In this case, a "batch" is the set of gallons produced in a production cycle. While the plant manager is likely to refer to batches of tomato soup and batches of chicken soup individually, and unlikely to refer to the batch that combines both flavors, we cannot analyze the production process of tomato soup in isolation from the production process of chicken soup. (For example, if we dedicate more time to tomato production, then we will have less time for chicken noodle production.) Because we are ultimately interested in our capacity to make soup, we focus our analysis at the level of the production cycle and refer to the entire production within that cycle as a "batch."

Our desired flow rate is 175 gallons per hour (the sum of demand for chicken and tomato), the setup time is 1 hour (30 minutes per soup and two types of soup) and the processing time is 1/300 hour per gallon. The batch size that minimizes inventory while still meeting our demand is then

$$\begin{aligned}\text{Recommended batch size} &= \frac{\text{Flow rate}\times\text{Setup time}}{1-\text{Flow rate}\times\text{Processing time}} = \frac{175\times(2\times1/2)}{1-175\times(1/300)}\\ &= 420\,\text{gallons}\end{aligned}$$

We should produce in proportion to demand (otherwise at least one of the flavors will have too much production and at least one will have too little), so of the 420 gallons, 420 × 100/(100 + 75) = 240 gallons should be chicken soup and the remainder, 420 − 240 = 180 gallons, should be tomato.

To evaluate the average inventory of chicken noodle soup, let's use the equation

$$\text{Average inventory} = \frac{1}{2}\text{Batch Size} \times (1 - \text{Flow rate} \times \text{Processing time})$$

The flow unit is 1 gallon of chicken noodle soup, the batch size is 240 gallons, the flow rate is 100 gallons per hour, and the processing time is 1/300 hours per gallon. Thus, the average inventory of chicken noodle soup is 1/2 × 240 gallons × (1 – 100 gallons per hour × 1/300 hours per gallon) = 80 gallons.

To understand the impact of variety on this process, suppose we were to add a third kind of soup to our product offering, onion soup. Furthermore, with onion soup added to the mix, demand for chicken remains 100 gallons per hour, and demand for tomato continues to be 75 gallons per hour, while onion now generates 30 gallons of demand on its own. In some sense, this is an ideal case for adding variety—the new variant adds incrementally to demand without stealing any demand from the existing varieties.

The desired flow rate is now 100 + 75 + 30 = 205, the setup time is 1.5 hours (three setups per batch), and the inventory minimizing quantity for the production cycle is

$$\begin{aligned}\text{Recommended batch size} &= \frac{\text{Flow rate}\times\text{Setup time}}{1-\text{Flow rate}\times\text{Processing time}} = \frac{205\times(3\times1/2)}{1-205\times(1/300)}\\ &= 971\,\text{gallons}\end{aligned}$$

Again, we should produce in proportion to demand:

971 × (100/205) = 474 gallons of chicken,

971 × (75/205) = 355 gallons of tomato, and

971 × (30/205) = 142 gallons of onion.

What happened when we added to variety? In short, we need more inventory. With the batch size of 474 gallons, the average inventory of chicken noodle soup becomes 1/2 × 474 gallons × (1 – 100 gallons per hour × 1/300 hours per gallon) = 158 gallons. Because the batch size of chicken noodle soup nearly doubles (474/240 = 1.98), the average inventory of chicken noodle soup also nearly doubles.

Page 93

Why did inventory of chicken soup increase when onion soup was added to the mix? Setup times are to blame. With more varieties in the production mix, the production process has to set up more often per production cycle. This reduces the capacity of the production cycle (no soup is made during a setup). To increase the capacity back to the desired flow rate (which is even higher now), we need to operate with larger batches (longer production cycles), and they lead to more inventory.

One may argue that the previous analysis is too optimistic—adding onion soup to the mix should steal some demand away from the other flavors. It turns out that our result is not sensitive to this assumption. To demonstrate, let's consider the opposite extreme—adding onion soup does not expand overall demand, it only steals demand from the other flavors. Specifically, the overall flow rate remains 175 gallons per hour, with or without onion soup. Furthermore, with onion soup, the demand rate for chicken, tomato, and onion are 80, 65, and 30 gallons per hour, respectively. The processing time is still 1/300 gallons per hour, and the setup time per batch is now 1.5 hours (three changeovers due to three types of soup). The batch size that minimizes our inventory while meeting our demand is

$$\begin{aligned}\text{Recommended batch size} &= \frac{\text{Flow rate}\times\text{Setup time}}{1-\text{Flow rate}\times\text{Processing time}} = \frac{175\times(3\times1/2)}{1-175\times(1/300)}\\ &= 630\,\text{gallons}\end{aligned}$$

The chicken noodle batch size is (80 gallons/175 gallons) × 630 gallons = 240 gallons. Average inventory is 1/2 × 240 gallons × (1 – 100 gallons per hour × 1/300 hours per gallon) = 96 gallons. Recall, with just the two flavors and a flow rate of 175 gallons of soup per hour, there are only 80 gallons of chicken noodle soup. So inventory does not increase as much in this case, but it still increases.

The conclusion from this investigation is that setup times and product variety do not mix very well. Consequently, there are two possible solutions to this challenge. The first is to offer only a limited amount of variety. That was Henry Ford's approach when he famously declared that "You can have any color Model-T you want, as long as it is black." It was also the approach recently taken by many firms during the COVID-19 pandemic. For example, if you were making consumer toilet paper, you found yourself with much higher demand than normal—even beyond panic buying, people were spending more time at home, less at work, and so the need for toilet paper at home naturally rose considerably. Equipment to make toilet paper is very large and expensive. While before the pandemic there was sufficient capacity to switch between different variants of the product, during the pandemic more capacity was needed. Adding more equipment to increase capacity was not an option because it can take years to bring a new manufacturing line into service. So the only option was to cut back on the variety of toilet paper. This reduced the number of setups which increased the amount of time in actual production. Customers might not have found their preferred type of toilet paper, but at the time, they were happy to purchase any variant.

While the "any color ... as long as it is black" strategy is a convenient solution for a production manager, it is not necessarily the best strategy for satisfying demand in a competitive environment. The other approach to the incompatibility of setups and variety is to work to eliminate setup times. This is the approach advocated by Shigeo Shingo, one of the most influential thought leaders in manufacturing. When he witnessed changeover times of more than an hour in an automobile plant, he responded with the quote, “The flow must go on,” meaning that every effort must be made to ensure a smooth flow of production. One way to ensure a smooth flow is to eliminate or reduce setup times. Shigeo Shingo developed a powerful technique for doing exactly that, which we will revisit later in the chapter.

5.5 Setup Time Reduction

LO 5-5

Explain the SMED method and how it addresses setup time reduction.

Despite improvement potential from the use of "good" batch sizes and smaller transfer batches, setups remain a source of disruption for a smooth process flow. For this reason, rather than taking setups as "God-given" constraints and finding ways to accommodate them, we should find ways that directly address the root cause of the disruption.

This is the basic idea underlying the single-minute exchange of die (SMED) method. The SMED method states that any setup of 10 or more minutes is an unacceptable source of process flow disruption and thus, should be reduced to a setup that is a single-digit number of minutes (i.e., 0–9 minutes). The 10-minute rule is not necessarily meant to be taken literally: the method was developed in the automotive industry, where setup times used to take as much as 4 hours. The practical implementation of the SMED method defines an aggressive, yet realistic, setup time goal and proceeds to identify potential opportunities for setup time reduction.

The basic underlying idea of SMED is to carefully analyze all tasks that are part of the setup time and then divide those tasks into two groups, ***internal setup*** tasks and ***external setup*** tasks.

- Internal setup tasks are those tasks that can only be executed while the machine is stopped.
- External setup tasks are those tasks that can be done while the machine is still operating, meaning they can be done *before* the actual changeover occurs.

Experience shows that companies are biased toward using internal setups and that, even without making large investments, internal setups can be translated into external setups.

Similar to our discussion about choosing a good batch size, the biggest obstacles to overcome are ineffective cost accounting procedures. Consider, for example, the case of a simple heat treatment procedure in which flow units are moved on a tray and put into an oven. Loading and unloading of the tray is part of the setup time. The acquisition of an additional tray that can be loaded (or unloaded) while the other tray is still in process (before the setup) allows the company to convert internal setup tasks to external ones. Is this a worthwhile investment?

The answer is, as usual, it depends. SMED applied to nonbottleneck steps is not creating any process improvement at all. As discussed previously, nonbottleneck steps have excessive capacity and therefore setups are entirely free (except for the resulting increase in inventory). Thus, investing in any resource, technical or human, is not only wasteful, but it also takes scarce improvement capacity/funds away from more urgent projects. However, if the oven in the previous example were the bottleneck step, almost any investment in the acquisition of additional trays suddenly becomes a highly profitable investment.

The idea of internal and external setups as well as potential conversion from internal to external setups is best visible in car racing. Any pit stop is a significant disruption of the race car's flow toward the finish line. At any point and any moment in the race, an entire crew is prepared to take in the car, having prepared for any technical problem from tire changes to refueling. While the technical crew might appear idle and underutilized throughout most of the race, it is clear that any second they can reduce from the time the car is in the pit (internal setups) to a moment when the car is on the race track is a major gain (e.g., no race team would consider mounting tires on wheels during the race; they just put on entire wheels).

5.6 Balancing Setup Costs with Inventory Costs: The EOQ Model

LO 5-6

Explain the importance of balancing setup costs with inventory costs to achieve an ideal economic order quantity.

Up to now, our focus has been on the role of setup times, as opposed to setup costs. Specifically, we have seen that setup time at the bottleneck leads to an overall reduction in process capacity. Assuming that the process is currently capacity-constrained, setup times thereby carry an opportunity cost reflecting the overall lower flow rate (sales).

Independent of such opportunity costs, setups frequently are associated with direct (out-of-pocket) costs. In these cases, we speak of setup costs (as opposed to setup times). Consider, for example, the following settings:

- The setup of a machine to process a certain part might require scrapping the first 10 parts that are produced after the setup. Thus, the material costs of these 10 parts constitute a setup cost.
- When receiving shipments from a supplier, there frequently exists a fixed shipment cost as part of the procurement cost, which is independent of the purchased quantity. This is similar to the shipping charges that a consumer pays at a catalog or online retailer. Shipping costs are a form of setup costs.

All of those settings reflect ***economies of scale***: the more we order or produce as part of a batch, the more units there are in a batch over which we can spread out the setup costs, thereby lowering the per-unit cost.

If we can reduce per-unit costs by increasing the batch size, what keeps us from using infinite (or at least very large) batches? Similar to the case of setup times, we again need to balance our desire for large batches (fewer setups) with the cost of carrying a large amount of inventory.

In the following analysis, we need to distinguish between two cases:

- If the quantity we order is produced or delivered by an outside supplier, all units of a batch are likely to arrive at the same time.
- In other settings, the units of a batch might not all arrive at the same time. This is especially the case when we produce the batch internally.

Figure 5.8 illustrates the inventory levels for the two cases described above. The lower part of **Figure 5.8** shows the case of the outside supplier and all units of a batch arriving at the same moment in time. The moment a shipment is received, the inventory level jumps up by the size of the shipment. It then falls up to the time of the next shipment.

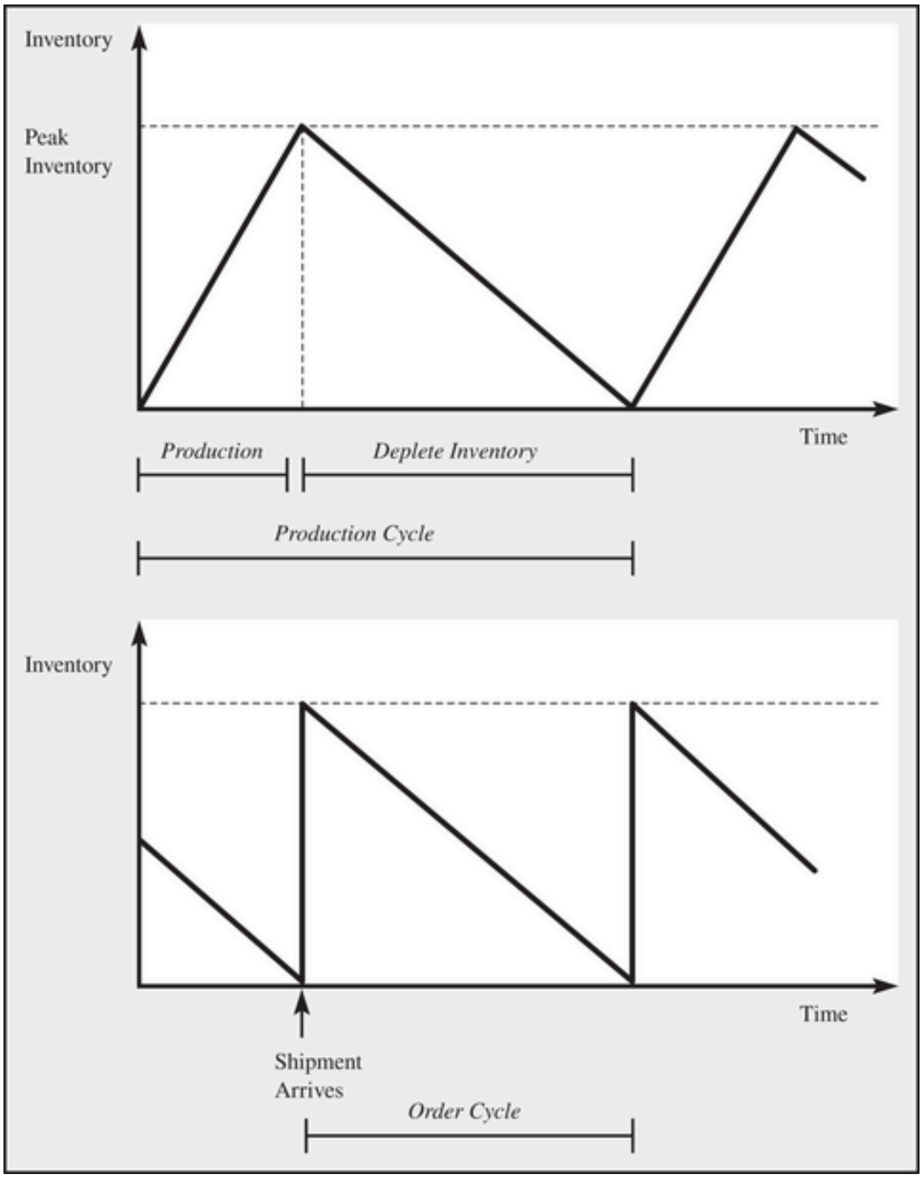

FIGURE 5.8 Different Patterns of Inventory Levels

The upper part of **Figure 5.8** shows the case of units created by a resource with (finite) capacity. Thus, while we are producing, the inventory level increases. Once we stop production, the inventory level falls. Let us consider the case of an outside supplier first (lower part of **Figure 5.8**). Specifically, consider the case of the Xootr handle caps that Nova Cruz sources from a supplier in Taiwan for $0.85 per unit. Note that the maximum inventory of handle caps occurs at the time we receive a shipment from Taiwan. The inventory is then depleted at the rate of the assembly operations, that is, at a flow rate, *R,* of 700 units (pairs of handle caps) per week, which is equal to 1 unit every 3 minutes.

For the following computations, we make a set of assumptions. We later show that these assumptions do not substantially alter the optimal decisions.

- We assume that production of Xootrs occurs at a constant rate of 1 unit every 3 minutes. We also assume our orders arrive on time from Taiwan. Under these two assumptions, we can deplete our inventory all the way to zero before receiving the next shipment.
- There is a fixed setup cost per order that is independent of the amount ordered. In the Xootr case, this largely consists of a $300 customs fee.
- The purchase price is independent of the number of units we order, that is, there are no quantity discounts. We talk about quantity discounts in the next section.

Page 96

The objective of our calculations is to minimize the cost of inventory and ordering with the constraint that we must never run out of inventory (i.e., we can keep the assembly operation running).

We have three costs to consider: purchase costs, delivery fees, and holding costs. We use 700 units of handle caps each week no matter how much or how frequently we order. Thus, we have no excuse for running out of inventory and there is nothing we can do about our purchase costs of

$$\$0.85/\text{unit} \times 700\ \text{units/week} = \$595\ \text{per week}$$

So when choosing our ordering policy (when and how much to order), we focus on minimizing the sum of the other two costs, delivery fees and inventory costs.

Page 97

The cost of inventory depends on how much it costs us to hold 1 unit in inventory for a given period of time, say 1 week. We can obtain the number by looking at the annual inventory costs and dividing that amount by 52. The annual inventory costs need to account for financing the inventory (cost of capital, especially high for a start-up like Nova Cruz), costs of storage, and costs of obsolescence. Nova Cruz uses an annual inventory cost of 40 percent. Thus, it costs Nova Cruz 0.7692 percent to hold a piece of inventory for 1 week. Given that a handle cap costs \$0.85 per unit, this translates to an inventory cost of h = 0.007692 × \$0.85/unit = \$0.006538 per unit per week. Note that the annual holding cost needs to include the cost of capital as well as any other cost of inventory (e.g., storage, theft).

How many handle caps will there be, on average, in Nova Cruz's inventory? As we can see in **Figure 5.8**, the average inventory level is simply

$$\text{Average inventory} = \frac{\text{Order quantity}}{2}$$

If you are not convinced, refer to **Figure 5.8** the "triangle" formed by one order cycle. The average inventory during the cycle is half of the height of the triangle, which is half the order quantity, $Q/2$. Thus, for a given inventory cost, h, we can compute the inventory cost per unit of time (e.g., inventory costs per week):

$$\text{Inventory costs [per unit of time]} = \frac{1}{2}\text{Order quantity} \times h = \frac{1}{2} Q \times h$$

Before we turn to the question of how many handle caps to order at once, let's first ask ourselves how frequently we have to place an order. Say at time 0 we have I units in inventory and say we plan our next order to be Q units. The I units of inventory will satisfy demand until time I/R (in other words, we have I/R weeks of supply in inventory). At this time, our inventory will be zero if we don't order before then. We would then again receive an order of Q units (if there is a lead time in receiving this order, we simply would have to place this order earlier).

Do we gain anything by receiving the Q handle caps earlier than at the time when we have zero units in inventory? Not in this model: demand is satisfied whether we order earlier or not and the delivery fee is the same too. But we do lose something by ordering earlier: we incur holding costs per unit of time the Q units are held.

Given that we cannot save costs by choosing the order time intelligently, we must now work on the question of how much to order (the order quantity). Let's again assume that we order

Q units with every order and let's consider just one order cycle. The order cycle begins when we order Q units and ends when the last unit is sold, Q/R time units later. For example, with Q = 1,000, an order cycle lasts 1,000 units/700 units per week = 1.43 weeks. We incur one ordering fee (setup costs), K, in that order cycle, so our setup costs per week are

$$\begin{aligned}\text{Setup costs [per unit of time]} &= \frac{\text{Setup cost}}{\text{Length of order cycle}} \\ &= \frac{K}{Q/R} = \frac{K \times R}{Q}\end{aligned}$$

Let $C(Q)$ be the sum of our average delivery cost per unit time and our average holding cost per unit time (per week):

$$\begin{aligned}\text{Per unit of time cost } C(Q) &= \text{Setup costs} + \text{Inventory costs} \\ &= \frac{K \times R}{Q} + \frac{1}{2} \times h \times Q\end{aligned}$$

Note that purchase costs are not included in $C(Q)$ for the reasons discussed earlier. Page 98 From the above, we see that the delivery fee per unit time decreases as Q increases: we amortize the delivery fee over more units. But as Q increases, we increase our holding costs.

Figure 5.9 graphs the weekly costs of delivery, the average weekly holding cost, and the total weekly cost, $C(Q)$. As we can see, there is a single order quantity Q that minimizes the total cost $C(Q)$. We call this quantity Q^*, the economic order quantity, or ***EOQ*** for short. Hence the name of the model.

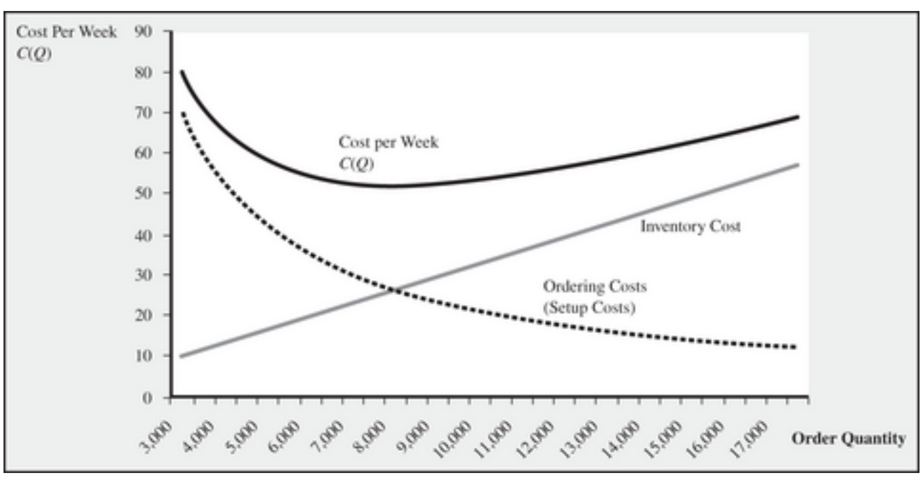

FIGURE 5.9 Inventory and Ordering Costs for Different Order Sizes

In Figure 5.9, it appears that Q^* is the quantity at which the weekly delivery fee equals the weekly holding cost. In fact, that is true, as can be shown algebraically. Further, using calculus it is possible to show that

$$\begin{aligned} \text{Economic order quantity} &= \sqrt{\frac{2\times\text{Setup cost}\times\text{Flow rate}}{\text{Holding cost}}} \\ Q^* &= \sqrt{\frac{2\times K\times R}{h}} \end{aligned}$$

As our intuition suggests, as the setup costs K increase, we should make larger orders, but as holding costs h increase, we should make smaller orders.

We can use the above formula to establish the economic order quantity for handle caps:

$$\begin{aligned} Q^* &= \sqrt{\frac{2\times\text{ Setup cost }\times\text{ Flow rate}}{\text{Holding cost}}} \\ &= \sqrt{\frac{2\times300\times700}{0.006538}} = 8{,}014.69 \end{aligned}$$

The steps required to find the economic order quantity are summarized in Exhibit 5.2.

Exhibit 5.2

FINDING THE ECONOMIC ORDER QUANTITY

1. Verify the basic assumptions of the EOQ model:
   - Replenishment occurs instantaneously.
   - Demand is constant and not stochastic.
   - There is a fixed setup cost K independent of the order quantity.
2. Collect information on
   - Setup cost, K (only include out-of-pocket cost, not opportunity cost).
   - Flow rate, R.
   - Holding cost, h (not necessarily the yearly holding cost; needs to have the same time unit as the flow rate).
3. For a given order quantity Q, compute

$$\begin{aligned} \text{Inventory costs [per unit of time]} &= \tfrac{1}{2} Q \times h \\ \text{Setup costs [per unit of time]} &= \tfrac{K \times R}{Q} \end{aligned}$$

4. The economic order quantity minimizes the sum of the inventory and the setup costs and is

$$Q^* = \sqrt{\frac{2 \times K \times R}{h}}$$

The resulting costs are

$$C(Q^*) = \sqrt{2 \times K \times R \times h}$$

5.7 Observations Related to the Economic Order Quantity

LO 5-7

Understand the economies of scale that are present when managing inventory with setup costs.

If we always order the economic order quantity, our cost per unit of time, $C(Q^*)$, can be computed as

$$C\left(Q^*\right) = \frac{K \times R}{Q^*} + \frac{1}{2} \times h \times Q^* = \sqrt{2 \times K \times R \times h}$$

While we have done this analysis to minimize our average cost per unit of time, it should be clear that Q^* would minimize our average cost per unit (given that the rate of purchasing handle caps is fixed). The cost per unit can be computed as

$$\text{Cost per unit} = \frac{C\left(Q^*\right)}{R} = \sqrt{\frac{2 \times K \times h}{R}}$$

As we would expect, the per-unit cost is increasing with the ordering fee K as well as with our inventory costs. Interestingly, the per-unit cost is decreasing with the flow rate R. Thus, if we doubled our flow rate, our ordering costs increase by less than a factor of 2. In other words, there are economies of scale in the ordering process: the per-unit ordering cost is decreasing with the flow rate R. Put yet another way, an operation with setup and inventory holding costs becomes more efficient as the demand rate increases.

While we have focused our analysis on the time period when Nova Cruz experienced a demand of 700 units per week, the demand pattern changed drastically over the product life cycle of the Xootr. As discussed in **Chapter 4**, Nova Cruz experienced a substantial demand growth from 200 units per week to over 1,000 units per week. **Table 5.2** shows how increases in demand rate impact the order quantity as well as the per-unit cost of the handle caps. We observe that, due to scale economies, ordering and inventory costs are decreasing with the flow rate *R*.

TABLE 5.2 **Scale Economies in the EOQ Formula**

Flow Rate, *R*	Economic Order Quantity, *Q**	Per-Unit Ordering and Inventory Cost, *C*(*Q**)/*R*	Ordering and Inventory Costs as a Percentage of Total Procurement Costs
200	4,284	0.14 [$/unit]	14.1%
400	6,058	0.10	10.4%
600	7,420	0.08	8.7%
800	8,568	0.07	7.6%
1,000	9,579	0.06	6.8%

A nice property of the economic order quantity is that the cost function, $C(Q)$, is relatively flat around its minimum Q^* (see graph in **Figure 5.9**). This suggests that if we were to order Q units instead of Q^*, the resulting cost penalty would not be substantial as long as Q is reasonably close to Q^*. Suppose we order only half of the optimal order quantity, that is, we order $Q^*/2$. In that case, we have

$$C\left(Q^*/2\right) = \frac{K \times R}{Q^*/2} + \frac{1}{2} \times h \times Q^*/2 = \frac{5}{4} \times \sqrt{2 \times K \times R \times h} = \frac{5}{4} \times C\left(Q^*\right)$$

Thus, if we order only half as much as optimal (i.e., we order twice as frequently as optimal), then our costs increase only by 25 percent. The same holds if we order double the economic order quantity (i.e., we order half as frequently as optimal).

This property has several important implications:

- Consider the optimal order quantity Q^* = 8,014 established above. However, now also assume that our supplier is only willing to deliver in predefined quantities (e.g., in multiples of 5,000). The robustness established above suggests that an order of 10,000 will only lead to a slight cost increase (increased costs can be computed as $C(Q = 10{,}000)$ = $53.69, which is only 2.5 percent higher than the optimal costs).
- Sometimes, it can be difficult to obtain exact numbers for the various ingredients in the EOQ formula. Consider, for example, the ordering fee in the Nova Cruz case. While this fee of $300 was primarily driven by the $300 for customs, it also did include a shipping fee. The exact shipping fee in turn depends on the quantity shipped and we would need a more refined model to find the order quantity that accounts for this

effect. Given the robustness of the EOQ model, however, we know that the model is "forgiving" with respect to small misspecifications of parameters.

A particularly useful application of the EOQ model relates to ***quantity discounts***. When procuring inventory in a logistics or retail setting, we frequently are given the opportunity to benefit from quantity discounts. For example,

- We might be offered a discount for ordering a full truckload of supply.
- We might receive a free unit for every 5 units we order (just as in consumer retailing settings of "buy one, get one free").
- We might receive a discount for all units ordered over 100 units.
- We might receive a discount for the entire order if the order volume exceeds 50 units (or say $2,000).

We can think of the extra procurement costs that we would incur from not taking advantage of the quantity discount—that is, that would result from ordering in smaller quantities—as a setup cost. Evaluating an order discount therefore boils down to a comparison between inventory costs and setup costs (savings in procurement costs), which we can do using the EOQ model.

If the order quantity we obtain from the EOQ model is sufficiently large to obtain the largest discount (the lowest per-unit procurement cost), then the discount has no impact on our order size. We go ahead and order the economic order quantity. The more interesting case occurs when the EOQ is less than the discount threshold. Then we must decide if we wish to order more than the economic order quantity to take advantage of the discount offered to us.

Let's consider one example to illustrate how to think about this issue. Suppose our supplier of handle caps gives us a discount of 5 percent off the entire order if the order exceeds 10,000 units. Recall that our economic order quantity was only 8,014. Thus, the question is "should we increase the order size to 10,000 units in order to get the 5 percent discount, yet incur higher inventory costs, or should we simply order 8,014 units?"

We surely will not order more than 10,000; any larger order does not generate additional purchase cost savings but does increase inventory costs. So we have two choices: either stick with the EOQ or increase our order to 10,000. If we order Q^* = 8,014 units, our total cost per unit time is

$$
\begin{aligned}
&700\ \text{units/week}\ \times\ \$0.85/\text{unit} + C\left(Q^*\right)\\
&\quad = \ \$595/\text{week}\ +\ \$52.40/\text{week}\\
&\quad = \ \$647.40/\text{week}
\end{aligned}
$$

Notice that we now include our purchase cost per unit time of 700 units/week × \$0.85/unit. The reason for this is that with the possibility of a quantity discount, our purchase cost now depends on the order quantity.

If we increase our order quantity to 10,000 units, our total cost per unit time would be

$$
\begin{aligned}
&700\ \text{units/week} \times \$0.85/\text{unit}\ \times\ 0.95 + C\,(10,000)\\
&\quad = \ \$565.25/\text{week} + \$52.06/\text{week}\\
&\quad = \ \$617.31/\text{week}
\end{aligned}
$$

where we have reduced the procurement cost by 5 percent (multiplied by 0.95) to reflect the quantity discount. (*Note:* The 5 percent discount also reduces the holding cost h in C.) Given that the cost per week is lower in the case of the increased order quantity, we want to take advantage of the quantity discount.

After analyzing the case of all flow units of one order (batch) arriving simultaneously, we now turn to the case of producing the corresponding units internally (upper part of **Figure 5.8**).

All computations we performed above can be easily transformed to this more general case (see, e.g., Nahmias 2005). Moreover, given the robustness of the economic order quantity, the EOQ model leads to reasonably good recommendations even if applied to production settings with setup costs. Hence, we will not discuss the analytical aspects of this. Instead, we want to step back for a moment and reflect on how the EOQ model relates to our discussion of setup times at the beginning of the chapter.

A common mistake is to rely too much on setup *costs* as opposed to setup *times.* For example, consider the case in **Figure 5.6** and assume that the monthly capital cost for the milling machine is \$9,000, which corresponds to \$64 per hour (assuming 4 weeks of 35 hours each). Thus, when choosing the batch size, and focusing primarily on costs, Nova Cruz might shy away from frequent setups. Management might even consider using the economic order quantity established above and thereby quantify the impact of larger batches on inventory holding costs.

There are two major mistakes in this approach:

- This approach to choosing batch sizes ignores the fact that the investment in the machine is already sunk.
- Choosing the batch size based on cost ignores the effect setups have on process capacity. As long as setup costs are a reflection of the cost of capacity—as opposed to direct financial setup costs—they should be ignored when choosing the batch size. It is the overall process flow that matters, not an artificial local performance measure! From a capacity perspective, setups at nonbottleneck resources are free. And if the setups do occur at the bottleneck, the corresponding setup costs not only reflect the capacity costs of the local resource, but of the entire process!

Thus, when choosing batch sizes, it is important to distinguish between setup costs and setup times. If the motivation behind batching results from setup times (or opportunity costs of capacity), we should focus on optimizing the process flow. **Section 5.3** provides the appropriate way to find a good batch size. If we face "true" setup costs (in the sense of out-of-pocket costs) and we only look at a single resource (as opposed to an entire process flow), the EOQ model can be used to find the optimal order quantity.

Finally, if we encounter a combination of setup times and (out-of-pocket) setup costs, we should use both approaches and compare the recommended batch sizes. If the batch size from the EOQ is sufficiently large so that the resource with the setup is not the bottleneck, minimizing costs is appropriate. If the batch size from the EOQ, however, makes the resource with the setups the bottleneck, we need to consider increasing the batch size beyond the EOQ recommendation.

5.8 Summary

Setups are interruptions of the supply process. These interruptions on the supply side lead to mismatches between supply and demand, visible in the form of inventory and—where this is not possible—lost throughput.

While in this chapter we have focused on inventory of components (handle caps), work-in-process (steer support parts), or finished goods (station wagons vs. sedans, **Figure 5.4**), the supply–demand mismatch also can materialize in an inventory of waiting customer orders. For example, if the product we deliver is customized and built to the specifications of the customer, holding an inventory of finished goods is not possible. Similarly, if we are providing a substantial variety of products to the market, the risk of holding completed variants in finished goods inventory is large. Independent of the form of inventory, a large inventory corresponds to long flow times (Little's Law). For this reason, batch processes are typically associated with very long customer lead times.

In this chapter, we discussed tools to choose a batch size. We distinguished between setup times and setup costs. To the extent that a process faces setup times, we need to extend our process analysis to capture the negative impact that setups have on capacity. We then want to look for a batch size that is large enough to not make the process step with the setup the bottleneck, while being small enough to avoid excessive inventory.

To the extent that a process faces (out-of-pocket) setup costs, we need to balance these costs against the cost of inventory. We discussed the EOQ model for the case of supply arriving in one single quantity (sourcing from a supplier), as well as the case of internal production. **Figure 5.10** provides a summary of the major steps you should take when analyzing processes with flow interruptions, including setup times, setup costs, or machine downtimes. There are countless extensions to the EOQ model to capture, among other things, quantity discounts, perishability, learning effects, inflation, and quality problems.

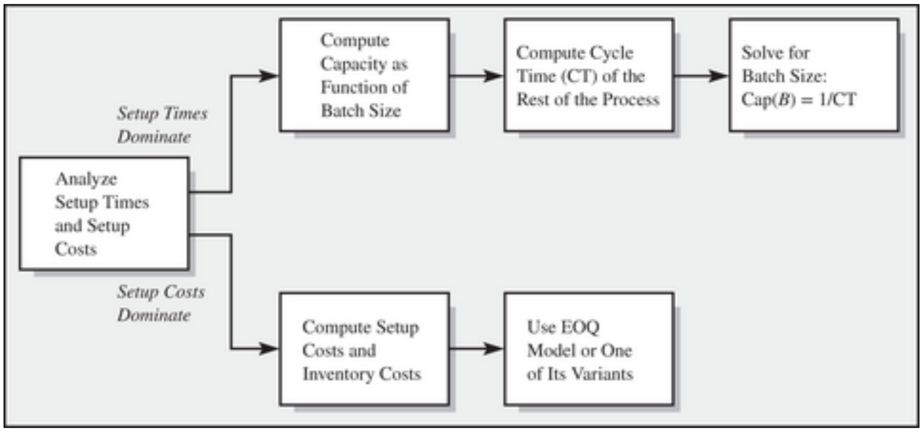

FIGURE 5.10 Summary of Batching

Our ability to choose a "good" batch size provides another example of process improvement. Consider a process with significant setup times at one resource. As a manager of this process, we need to balance the conflicting objectives of

- Fast response to customers (short flow times, which correspond, because of Little's Law, to low inventory levels), which results from using small batch sizes.
- Cost benefits that result from using large batch sizes. The reason for this is that large batch sizes enable a high throughput, which in turn allows the firm to spread out its fixed costs over a maximum number of flow units.

This tension is illustrated in **Figure 5.11**. Similar to the case of line balancing, we observe that adjustments in the batch size are not trading in one performance measure against the other, but allow us to improve by reducing current inefficiencies in the process.

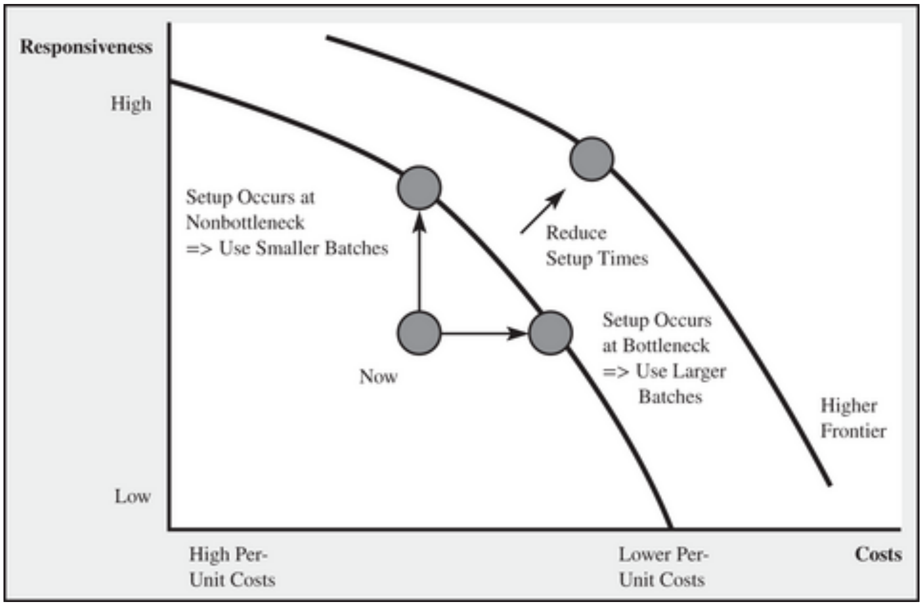

FIGURE 5.11 Choosing a Batch Size

Despite our ability to choose batch sizes that mitigate the tension between inventory (responsiveness) and costs, there ultimately is only one way to handle setups: eliminate them wherever possible or at least shorten them. Setups do not add value and are therefore wasteful.

Methods such as SMED are powerful tools that can reduce setup times substantially. Similarly, the need for transfer batches can be reduced by locating the process resources according to the flow of the process.

5.9 Practice Problems and Selected Solutions

The following questions will help in testing your understanding of this chapter. After each question, we show the relevant section in brackets [Section *x*].

Solutions to problems marked with an "*" appear at the end of this section. Video solutions to select problems are available in Connect.

Q5.1* **(Window Boxes)** Metal window boxes are manufactured in two process steps: stamping and assembly. Each window box is made up of three pieces: a base (one part A) and two sides (two part Bs).

The parts are fabricated by a single stamping machine that requires a setup time of 120 minutes whenever switching between the two part types. Once the machine is set up, the processing time for each part A is 1 minute while the processing time for each part B is only 30 seconds.

Currently, the stamping machine rotates its production between one batch of 360 for part A and one batch of 720 for part B.

In assembly, parts are assembled manually to form the finished product. One base (part A) and two sides (two part Bs), as well as a number of small purchased components, are required for each unit of final product. Each product requires 27 minutes of labor time to assemble. There are currently 12 workers in assembly. There is sufficient demand to sell every box the system can make.

a. What is the capacity of the stamping machine? [5.1]

b. What batch size would you recommend for the process? [5.6]

c. Suppose they operate with a production cycle of 1,260 part As and 2,520 part Bs. What would be the average inventory of part A? [5.6]

Q5.2 **(Two-step)** Consider the following two-step process:

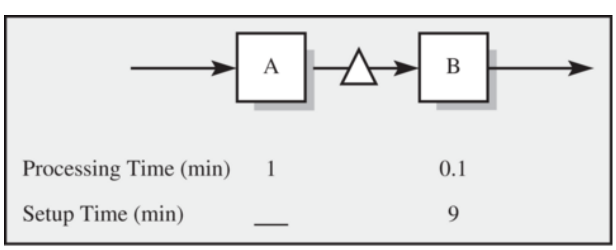

Step A has a processing time of 1 minute per unit, but no setup is required. Step B has a processing time of 0.1 minute per unit, but a setup time of 9 minutes is required per batch.

a. Suppose units are produced in batches of five (i.e., after each set of 5 units is produced, step B must incur a setup of 9 minutes). What is the capacity of the process (in units per minute)? [**5.1**]

b. Suppose they operate with a batch size of 15 and with this batch size step A is the bottleneck. What would be the average inventory after step B? [**5.6**]

c. What is the batch size that maximizes the flow rate of this process with minimal inventory? Assume there is ample demand. [**5.6**]

Q5.3 **(Simple Setup)** Consider the following batch-flow process consisting of three process steps performed by three machines:

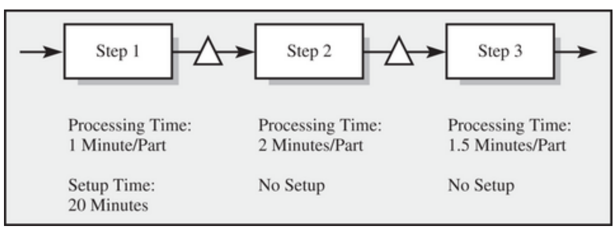

Work is processed in batches at each step. Before a batch is processed at step 1, the machine has to be set up. During a setup, the machine is unable to process any product.

a. Assume that the batch size is 50 parts. What is the capacity of the process? [**5.1**]

b. For a batch size of 10 parts, which step is the bottleneck for the process? [**5.1**]

c. What batch size would you choose? [**5.6**]

d. Suppose the batch size is 40 parts. What would be the average inventory after step 1? [**5.6**]

Q5.4 **(Setup Everywhere)** Consider the following batch-flow process consisting of three process steps performed by three machines:

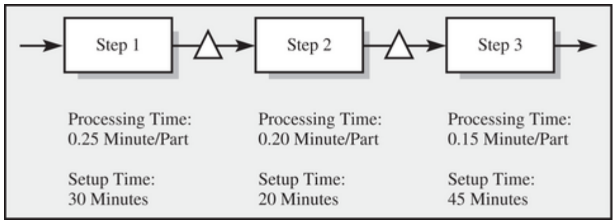

Work is processed in batches at each step. Before a batch is processed at a step, the machine at that step must be set up. (During a setup, the machine is unable to process any product.) Assume that there is a dedicated setup operator for each machine (i.e., there is always someone available to perform a setup at each machine.)

a. What is the capacity of step 1 if the batch size is 35 parts? [**5.6**]

b. For what batch sizes is step 1 (2, 3) the bottleneck? [**5.6**]

Q5.5 **(JCL Inc.)** JCL Inc. is a major chip manufacturing firm that sells its products to computer manufacturers such as Dell, HP, and others. In simplified terms, chip making at JCL Inc. involves three basic operations: depositing, patterning, and etching.

- **Depositing:** Using chemical vapor deposition (CVD) technology, an insulating material is deposited on the wafer surface, forming a thin layer of solid material on the chip.
- **Patterning:** Photolithography projects a microscopic circuit pattern on the wafer surface, which has a light-sensitive chemical like the emulsion on photographic film. It is repeated many times as each layer of the chip is built.
- **Etching:** Etching removes selected material from the chip surface to create the device structures.

The following table lists the required processing times and setup times at each of the steps. Assume that the unit of production is a wafer, from which individual chips are cut at a later stage.

Note: A setup can only begin once the batch has arrived at the machine.

Process Step	1 Depositing	2 Patterning	3 Etching
Setup time	45 min.	30 min.	20 min.
Processing time	0.15 min./unit	0.25 min./unit	0.20 min./unit

a. What is the process capacity in units per hour with a batch size of 100 wafers? [**5.1**]

b. For what batch size is step 3 (etching) the bottleneck? [5.6]

c. Suppose JCL Inc. came up with a new technology that eliminated the setup time for step 1 (deposition) but increased the processing time to 0.45 minute/unit. What would be the batch size you would choose so as to maximize the overall capacity of the process? [5.6]

Q5.6 **(Kinga Doll Company)** Kinga Doll Company manufactures eight versions of its popular girl doll, Shari. The company operates on a 40-hour workweek. The eight versions differ in doll skin, hair, and eye color, enabling most children to have a doll with a similar appearance to them. It currently sells an average of 4,000 dolls (spread equally among its eight versions) per week to boutique toy retailers. In simplified terms, doll making at Kinga involves three basic operations: molding the body and hair, painting the face, and dressing the doll. Changing over between versions requires setup time at the molding and painting stations due to the different colors of plastic pellets, hair, and eye color paint required. The table below lists the setup times for a batch and the processing times for each unit at each step. Unlimited space for buffer inventory exists between these steps.

Assume that (i) setups need to be completed first, (ii) a setup can only start once the batch has arrived at the resource, and (iii) all flow units of a batch need to be processed at a resource before any of the units of the batch can be moved to the next resource.

Process Step	1 Molding	2 Painting	3 Dressing
Setup time	15 min.	30 min.	No setup
Processing time	0.25 min./unit	0.15 min./unit	0.30 min./unit

a. What is the process capacity in units per hour with a batch size of 500 dolls? [5.1]

b. Which batch size would minimize inventory without decreasing the process capacity? [5.6]

c. Which batch size would minimize inventory without decreasing the current flow rate? [5.6]

Q5.7 **(PTests)** Precision Testing (PTests) does fluid testing for several local hospitals. Consider their urine testing process. Each sample requires 12 seconds to test, but after 300 samples, the equipment must be recalibrated. No samples can be tested during the recalibration process and that process takes 30 minutes.

a. What is PTests's maximum capacity to test urine samples (in samples per hour)? [5.1]

b. Suppose 2.5 urine samples need to be tested per minute. What is the smallest batch size (in samples) that ensures that the process is not supply constrained? (Note; A batch is the number of tests between calibrations.) [5.6]

c. PTests also needs to test blood samples. There are two kinds of tests that can be done—a "basic" test and a "complete" test. Basic tests require 15 seconds per sample, whereas "complete" tests require 1.5 minutes per sample. After 100 tests, the equipment needs to be cleaned and recalibrated, which takes 20 minutes. Suppose PTests runs the following cyclic schedule: 70 basic tests, 30 complete tests, recalibrate, and then repeat. With this schedule, how many *basic* tests can they complete per minute on average? [5.1]

Q5.8 **(Gelato)** Bruno Fruscalzo decided to set up a small production facility in Sydney to sell to local restaurants that want to offer gelato on their dessert menu. To start simple, he would offer only three flavors of gelato: fragola (strawberry), chocolato (chocolate), and bacio (chocolate with hazelnut). After a short time, he found his demand and setup times to be

	Fragola	Chocolato	Bacio
Demand (kg/hour)	10	15	5
Setup time (hours)	3/4	1/2	1/6

Bruno first produces a batch of fragola, then a batch of chocolato, then a batch of bacio, and then he repeats that sequence. For example, after producing bacio and before producing fragola, he needs 45 minutes to set up the ice cream machine, but he needs only 10 minutes to switch from chocolato to bacio. When running, his ice cream machine produces at the rate of 50 kg per hour no matter which flavor it is producing (and, of course, it can produce only one flavor at a time).

a. Suppose Bruno wants to minimize the amount of each flavor produced at one time while still satisfying the demand for each of the flavors. (He can choose a different quantity for each flavor.) If we define a batch to be the quantity produced in a single run of each flavor, how many kilograms should he produce in each batch? [5.6]

b. Given your answer in part (a), how many kilograms of fragola should he make with each batch? [5.6]

c. Given your answer in part (a), what is the average inventory of chocolato? (Assume production and demand occur at constant rates.) [5.6]

Q5.9 (**Carpets**) A carpet manufacturer makes four kinds of carpet on a machine. For simplicity, call these four types of carpet, A, B, C, and D. It takes 3 hours to switch production from one type of carpet to another. The demand rates (yards/hr) for the four types of carpet are as follows: 100, 80, 70, and 50. When producing, the machine can produce at the rate of 350 yards/hr. Batch sizes are chosen to minimize inventory while also satisfying the demand requirements. The manufacturer produces with a schedule that cycles through each of the four types, (e.g., A, B, C, D, A, B, . . .).

a. What batch size (yards) is chosen for carpet type A? [**5.6**]

b. Suppose they produce 16,800 yards of carpet A in each production cycle (and 50,400 yards of carpet in total within the production cycle). What would be the average inventory of carpet A? [**5.6**]

Q5.10* **(Cat Food)** Cat Lovers Inc. (CLI) is the distributor of a very popular blend of cat food that sells for $1.25 per can. CLI experiences demand of 500 cans per week on average. They order the cans of cat food from the Nutritious & Delicious Co. (N&D). N&D sells cans to CLI at $0.50 per can and charges a flat fee of $7 per order for shipping and handling. CLI uses the economic order quantity as their fixed order size. Assume that the opportunity cost of capital and all other inventory cost is 15 percent annually and that there are 50 weeks in a year.

a. How many cans of cat food should CLI order at a time? [**5.6**]

b. What is CLI's total order cost for 1 year? [**5.6**]

c. What is CLI's total holding cost for 1 year? [**5.6**]

d. What is CLI's weekly inventory turns? [**5.6**]

Q5.11* **(Beer Distributor)** A beer distributor finds that it sells on average 100 cases a week of regular 12-oz. Budweiser. For this problem assume that demand occurs at a constant rate over a 50-week year. The distributor currently purchases beer every 2 weeks at a cost of $8 per case. The inventory-related holding cost (capital, insurance) for the distributor equals 25 percent of the dollar value of inventory per year. Each order placed with the supplier costs the distributor $10. This cost includes labor, forms, postage, and so forth.

a. Assume the distributor can choose any order quantity it wishes. What order quantity minimizes the distributor's total inventory-related costs (holding and ordering)? [**5.6**]

For the next three parts, assume the distributor selects the order quantity specified in part (a).

b. What are the distributor's inventory turns per year? [**5.6**]

c. What is the inventory-related cost per case of beer sold? [**5.6**]

d. Assume the brewer is willing to give a 5 percent quantity discount if the distributor orders 600 cases or more at a time. If the distributor is interested in minimizing its total cost (i.e., purchase and inventory-related costs), should the distributor begin ordering 600 or more cases at a time? [**5.6**]

Q5.12 **(Millennium Liquors)** Millennium Liquors is a wholesaler of sparkling wines. Their most popular product is the French Bete Noire. Weekly demand is for 45 cases. Assume demand occurs over 50 weeks per year. The wine is shipped directly from France. Millennium's annual cost of capital is 15 percent, which also includes all other inventory-related costs. Below are relevant data on the costs of shipping, placing orders, and refrigeration.

- Cost per case: $120
- Shipping cost (for any size shipment): $290
- Cost of labor to place and process an order: $10
- Fixed cost for refrigeration: $75/week

a. Calculate the weekly holding cost for one case of wine. [**5.6**]

b. Use the EOQ model to find the number of cases per order and the average number of orders per year. [**5.6**]

c. Currently orders are placed by calling France and then following up with a letter. Millennium and its supplier may switch to a simple ordering system using the Internet. The new system will require much less labor. What would be the impact of this system on the ordering pattern? [**5.6**]

Q5.13 **(Powered by Koffee)** Powered by Koffee (PBK) is a new campus coffee store. PBK uses 50 bags of whole bean coffee every month, and you may assume that demand is perfectly steady throughout the year.

PBK has signed a year-long contract to purchase its coffee from a local supplier, Phish Roasters, for a price of $25 per bag and an $85 fixed cost for every delivery independent of the order size. The holding cost due to storage is $1 per bag per month. PBK managers figure their cost of capital is approximately 2 percent per month.

a. What is the optimal order size, in bags? [**5.6**]

b. Given your answer in (a), how many times a year does PBK place orders? [**5.6**]

c. Given your answer in (a), how many months of supply of coffee does PBK have on average? [**5.6**]

d. On average, how many dollars per month does PBK spend to hold coffee (including cost of capital)? [**5.6**]

Suppose that a South American import/export company has offered PBK a deal for the next year. PBK can buy a year’s worth of coffee directly from South America for $20 per bag and a fixed cost for delivery of $500. Assume the estimated cost for inspection and storage is $1 per bag per month and the cost of capital is approximately 2 percent per month.

e. Should PBK order from Phish Roasters or the South American import/export company? Quantitatively justify your answer. [**5.6**]

Selected Solutions

Q5.1 (Window Boxes)

The following computations are based on **Exhibit 5.1**.

Part a

Step 1. Since there is sufficient demand, the step (other than the stamping machine) that determines the flow rate is assembly. Capacity at assembly is $\frac{12}{27}$ unit/minute.

Step 2. The production cycle consists of the following parts:

- Setup for A (120 minutes).
- Produce parts A (360 × 1 minute).
- Setup for B (120 minutes).
- Produce parts B (720 × 0.5 minute).

Step 3. There are two setups in the production cycle, so the setup time is 240 minutes.

Step 4. Every completed window box requires one part A (1 minute per unit) and two parts B (2 × 0.5 minute per unit). Thus, the per-unit activity time is 2 minutes per unit.

Step 5. Use formula

$$\begin{aligned}\text{Capacity given batch size} &= \frac{360 \text{ units}}{240 \text{ minutes} + 360 \text{ units} \times 2 \text{ minutes/unit}} \\ &= 0.375 \text{ unit/minute}\end{aligned}$$

Step 6. Capacity at stamping for a general batch size is

$$\frac{\text{Batch size}}{240 \text{ minutes} + \text{Batch size} \times 2 \text{ minutes/unit}}$$

We need to solve the equation

$$\frac{\text{Batch size}}{240 \text{ minutes} + \text{Batch size} \times 2 \text{ minutes/unit}} = \frac{12}{27}$$

for the batch size. The batch size solving this equation is Batch size = 960. We can obtain the same number directly by using

$$\text{Recommended batch size} = \frac{\text{Flow rate} \times \text{Setup time}}{1 - \text{Flow rate} \times \text{Time per unit}} = \frac{\frac{12}{27} \times 240}{1 - \frac{12}{27} \times 2} = 960$$

Q5.10 (Cat Food)

$$\frac{7 \times 500}{\text{EOQ}} = 1.62$$

Part a

Holding costs are \$0.50 × 15%/50 = 0.0015 per can per week. Note, each can is purchased for \$0.50, so that is the value tied up in inventory and therefore determines the holding cost.

The EOQ is then $\sqrt{\dfrac{2 \times 7 \times 500}{0.0015}} = 2160$

Part b

The ordering cost is \$7 per order. The number of orders per year is 500/EOQ. Thus, order cost $= \frac{7 \times 500}{EOQ} = 1.62$ \$/week = 81\$/year.

Part c

The average inventory level is EOQ/2. Inventory costs per week are thus 0.5 × EOQ × 0.0015 = \$1.62. Given 50 weeks per year, the inventory cost per year is \$81.

Part d

Inventory turns = Flow rate/Inventory

Flow Rate = 500 cans per week

Inventory = 0.5 × EOQ

Thus, Inventory Turns = R/(0.5 × EOQ) = 0.462 turns per week = 23.14 turns per year

Q5.11 (Beer Distributor)

The holding costs are 25% per year = 0.5% per week = 8*0.005 = \$0.04 per week

(a) EOQ $= \sqrt{\frac{2 \times 100 \times 10}{0.04}} = 223.6$

(b) Inventory turns = Flow Rate/Inventory = 100 × 50/(0.5 × EOQ) = 5,000/EOQ = 44.7 turns per year

(c) Per unit inventory cost $= \sqrt{\frac{2 \times 0.04 \times 10}{100}} = 0.089\$/unit$

(d) You would never order more than Q = 600.

For Q = 600, we would get the following costs: 0.5 × 600 × 0.04 × 0.95 + 10 × 100/600 = 13.1.

The cost per unit would be 13.1/100 = \$0.131.

The quantity discount would save us 5%, which is \$0.40 per case. However, our operating costs increase by \$0.131 − 0.089 = \$0.042. Hence, the savings outweigh the cost increase and it is better to order 600 units at a time.