Chapter 6

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

Up to this point, we were working under the assumption that during every X units of time, one flow unit would enter the process and one flow unit would leave the process. We defined X as the process cycle time. In the scooter example of Chapter 4, we established a cycle time of three minutes in conjunction with Table 4.3, allowing us to fulfill demand of 700 scooters per week.

In an ideal process, a cycle time of three minutes would imply that every resource receives one flow unit as an input each three-minute interval and creates one flow unit of output each three-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. Problems related to variability are discussed in Chapters 7 and 8. And quality problems are discussed in Chapter 9.

Unlike mass production systems with their highly specialized tools, batch operations typically use general-purpose technology to produce a larger variety of products in production runs. Given the general nature of the production technology and the high level of product variety, the production resources in a batch-flow operation commonly have to be set up before beginning work on a specific product.

We define a production batch as a collection of flow units that are processed before the resource (usually the equipment being used at that step) needs to go through another setup. Such a setup might involve changing the equipment configuration from producing product A to producing product B (in which case, we also speak of a changeover time), an example

FIGURE 6.1 Milling Machine (left) and Steer Support Parts (right)

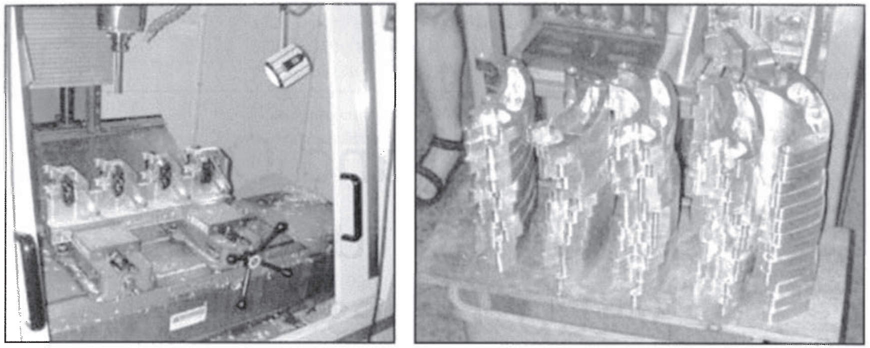

common in low-volume, high-variety manufacturing. A setup also might be the result of some other, recurring flow interruption, such as breaks for workers or downtime for machines.

In this chapter, we continue our discussion of the Xootr production process. While the assembly operations discussed in Chapter 4 do not include any setup times, the CNC milling machine that produces the steer support as well as the ribs required for final assembly of the Xootr follow a batch operation (see Figure 6.1). Specifically, the milling machine must undergo a one-hour setup whenever it is switched from producing steer support parts to ribs and, similarly, a one-hour setup whenever it is switched back to producing steer support parts. Every Xootr includes one steer support unit and two ribs.

6.1 The Impact of Setups on Capacity

The objective of our analysis remains to predict the three basic performance measures of a process: inventory, flow rate, and flow time. Toward that end, we need to be able to identify the bottleneck of the process, which requires the computation of the capacity for each of the various process steps.

However, in determining the capacity of each process step, it is important to not only count the per-unit activity times, but also to include the effect of setup times on capacity. As no output is produced while the resource is in setup mode, it is fairly intuitive that frequent setups lead to lower capacity.

Shigeo Shingo, one of the most influential thought leaders in manufacturing, is quoted as saying, "The flow must go on," when he witnessed changeover times of more than an hour in an automotive plant he studied. The idea of "the flow must go on" is helpful to us for two reasons. First, it illustrates that setups are interruptions of the process flow. They thereby "steal" capacity. Second, while we can do many things to choose batch sizes intelligently—which will be explained in the following pages—there fundamentally exists only one response to setups: eliminate them, or at least try to reduce the time it takes to perform the setup. Shigeo Shingo developed a powerful technique toward that end, which we revisit at the end of this chapter.

To understand how setups reduce the capacity of a process, consider Figure 6.2. The impact of setups on capacity is fairly intuitive. As nothing is produced at a resource during setup, the

FIGURE 6.2 The Impact of Setup Times on Capacity

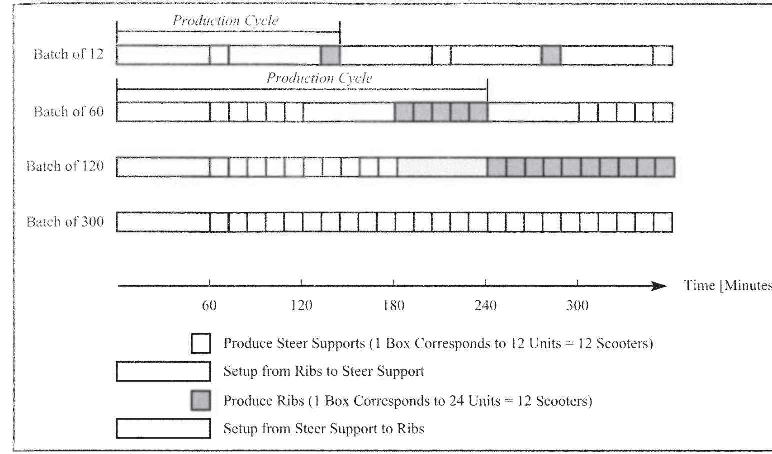

more frequently a resource is set up, the lower its capacity. As discussed above, the milling machine underlying the example of Figure 6.2 has the following activity times/setup times:

- It takes one 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 one rib; since there are two ribs in a Xootr, this translates to one minute/unit.
- 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 scooters per batch, we spend a total of two hours of setup time (one hour to set up the production for steer supports and one hour to set up the production of ribs) for every 12 scooters we produce. These two hours of setup time are lost for 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 scooter per minute. Table 6.1 summarizes the capacity calculations for batch sizes of 12, 60, 120, and 300.

Generalizing the computations in Table 6.1, we can compute the capacity of a resource with setups as a function of the batch size:

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

TABLE 6.1
The Impact of Setups on Capacity

Batch Size	Time to Complete One Batch [minutes]	Capacity [units/minute]
12	60 minutes (set up steering support)	12/144 = 0.0833
	+ 12 minutes (produce steering supports)	
	+ 60 minutes (set up ribs)	
	+ 12 minutes (produce ribs)	
	144 minutes	
60	60 minutes (set up steering support)	60/240 = 0.25
	+ 60 minutes (produce steering supports)	
	+ 60 minutes (set up ribs)	
	+ 60 minutes (produce ribs)	
	240 minutes	
120	60 minutes (set up steering support)	120/360 = 0.333
	+ 120 minutes (produce steering supports)	
	+ 60 minutes (set up ribs)	
	+ 120 minutes (produce ribs)	
	360 minutes	
300	60 minutes (set up steering support)	300/720 = 0.4166
	+ 300 minutes (produce steering supports)	
	+ 60 minutes (set up ribs)	
	+ 300 minutes (produce ribs)	
	720 minutes	

Basically, the above equation is spreading the "unproductive" setup time over the members of a batch. To use the equation, we need to be careful how exactly we define the batch size, the setup time, and the time per unit:

- The batch size is the number of scooters that are produced in one "cycle" (i.e., before the process repeats itself, see Figure 6.2). Let's say the batch size would be $B = 100$ scooters.
- The setup time includes all setups within a production cycle. In this case, this includes $S = 60$ minutes + 60 minutes = 120 minutes.
- The time per unit includes all production time that is needed to produce one complete unit of output at the milling machine. In this case, this includes 1 minute/unit for the steer support as well as two times 0.5 minute/unit for the two ribs. The total time per unit is thus $p = 1$ minute/unit + 2 × 0.5 minute/unit = 2 minutes/unit.

With these more careful definitions, we can now use the above equation to compute the capacity of the milling machine as

$$\text{Capacity (for } B = 100) = \frac{\text{Batch size}}{\text{Setup time} + \text{Batch size} \times \text{Time per unit}}$$
$$= \frac{100 \text{ units}}{120 \text{ minutes} + 100 \text{ units} \times 2 \text{ minutes/unit}}$$
$$= 0.3125 \text{ unit/minute}$$

No matter how large a batch size we choose, we will never be able to produce faster than one 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 6.3.

FIGURE 6.3
Capacity as a Function of the Batch Size

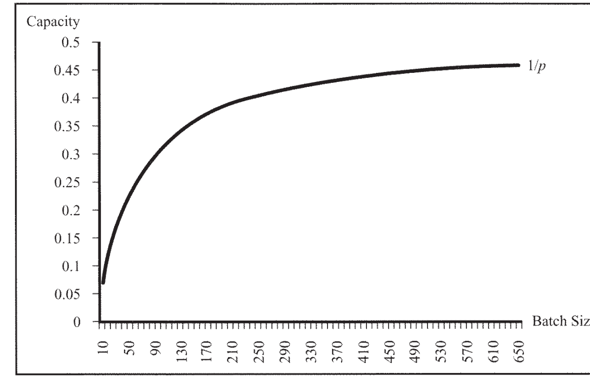

6.2 Interaction between Batching and Inventory

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 will lead 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 furniture?).

The interaction between batching and inventory is illustrated by the following two examples. First, 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 from the first of a month to the 15th and station wagons from the 16th to the end of the month.

However, it seems fairly unrealistic to assume that customers only demand sedans at the beginning of the month and station wagons at the end of the month. In other words, producing in large batches leads to a mismatch between the rate of supply and the rate of demand.

Thus, in addition to producing enough to cover demand in the first half of the month, to satisfy demand for sedans the company needs to produce 15 days of demand to inventory, which then fulfills demand while the line produces station wagons. This is illustrated by the left side of Figure 6.4. Observe that the average level of inventory is 3,000 cars for each of the two models. Now, ignoring setup times for a moment, consider the case in which the firm produces 400 station wagons and 400 sedans a day. In this setting, one would only need to carry 0.5 day of cycle inventory, a dramatic reduction in inventory. This is illustrated by the right side of Figure 6.4. Thus, smaller batches translate to lower inventory levels!

In the ideal case, which has been propagated by Toyota Production Systems 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.

FIGURE 6.4 The Impact of Batch Sizes on Inventory

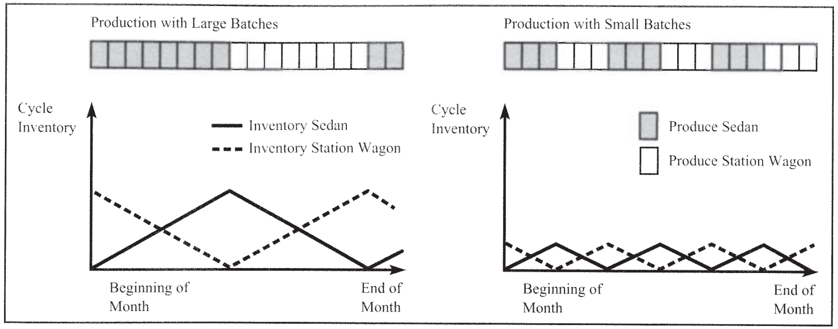

This way, a much better synchronization of the demand flow and the production flow is achieved and cycle inventory is eliminated entirely.

Second, consider a furniture maker producing chairs in batch sizes of 100. Starting with the wood-cutting step and all the way through the finishing process, the batch of 100 chairs would stay together as one entity.

Now, take the position of one chair in the batch. What is the most dominant activity throughout the process? Waiting! The larger the batch size, the longer the time the flow unit waits for the other "members" of the same batch—a situation comparable with going to the barber with an entire class of children. Given Little's Law, this increase in wait time (and thereby flow time) leads to a proportional increase in inventory.

With these observations, we can turn our attention back to the milling machine at Nova Cruz. Similar to Figure 6.4, 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 is requiring a supply of one unit every three minutes. This one unit consists, from the view of the milling machine, of two 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. If we assume, for the moment, a batch size of $B = 200$, the inventory of ribs changes as follows:

- During the production of ribs, inventory accumulates. As we produce ribs for one scooter per minute, but only supply ribs to the assembly line at a rate of one scooter every three minutes, rib inventory accumulates. It does so at a rate of enough ribs for two scooters every three minutes.
- Because we produce for 200 minutes, the inventory of ribs at the end of the production run will be enough to produce 200 minutes $\times$ 2/3 scooters per minute = 133 scooters (i.e. 266 ribs).
- How long does the rib inventory for 133 scooters last? The inventory ensures supply to the assembly for 400 minutes (cycle time of assembly operations was three minutes). After these 400 minutes, we need to start producing ribs again. During these 400 minutes, we have to accommodate two setups (together 120 minutes) and 200 minutes for producing the steer supports.

The resulting production plan as well as the corresponding inventory levels are summarized by Figure 6.5.

FIGURE 6.5 The Impact of Setup Times on Capacity

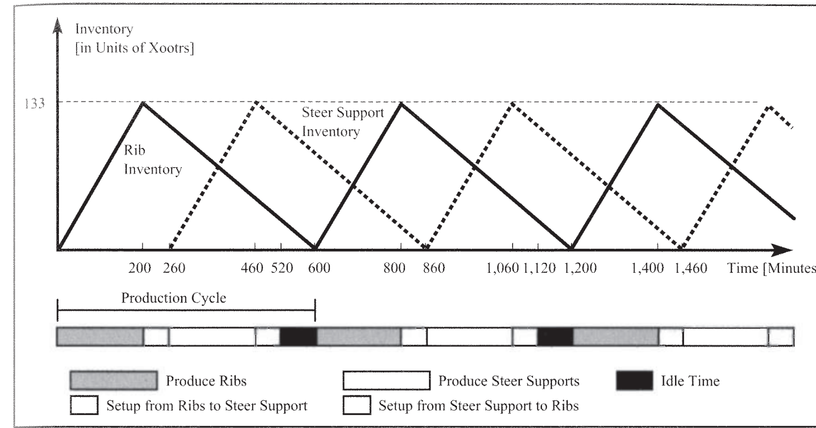

6.3 Choosing a Batch Size in the Presence of Setup Times

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; small batches lead to losses in capacity.

In balancing these two conflicting objectives, 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 by Figure 6.6 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 activity time.

The process flow diagram in Figure 6.6 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 one Xootr every three minutes.

FIGURE 6.6
Data from the Scooter Case about Setup Times and Batching

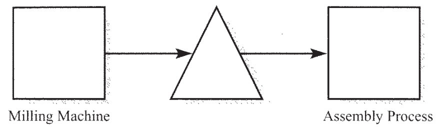

	Milling Machine	Assembly Process
Setup Time, S	120 Minutes	—
Per-Unit Time, p	2 Minutes/Unit	3 Minutes/Unit
Capacity (B = 12)	0.0833 Unit/Minute	0.33 Unit/Minute
Capacity (B = 300)	0.4166 Unit/Minute	0.33 Unit/Minute

To determine the capacity of the milling machine for a batch size of 12, we apply the formula

$$\text{Capacity } (B) = \frac{\text{Batch size}}{\text{Setup time} + \text{Batch size} \times \text{Time per unit}}$$
$$= \frac{B}{S + B \times p} = \frac{12}{120 + 12 \times 2} = 0.0833 \text{ unit/minute}$$

The capacity of the assembly operation is easily computed based on its bottleneck capacity of $\frac{1}{3}$ unit per minute. Note that for $B = 12$, the milling machine is the bottleneck.

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, just by modifying the batch size. Now which of the two batch sizes is the "better" one, 12 or 300?

- The batch size of 300 is clearly too large. The milling machine incurs idle time as the overall process is constrained by the (substantially) smaller capacity of the assembly operations (note, based on Figure 6.5, we know that even for the smaller batch size of $B = 200$, there exists idle time at the milling machine). This large batch size is likely to create unnecessary inventory problems as described above.
- The batch size of 12 is likely to be more attractive in terms of inventory. Yet, the process capacity has been reduced to 0.0833 unit per minute, leaving the assembly operation starved for work.

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 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:

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

which gives us, in this case, $B = 120$. This algebraic approach is illustrated by Figure 6.7. If you feel uncomfortable with the calculus 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:

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

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 Time per unit = 2 minutes per unit and obtain

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

Figure 6.7 shows the capacity of the process step with setup, 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 6.3). 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 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 6.1 provides a summary of the computations leading to the recommended batch size in the presence of setup times.

FIGURE 6.7
Choosing a "Good" Batch Size

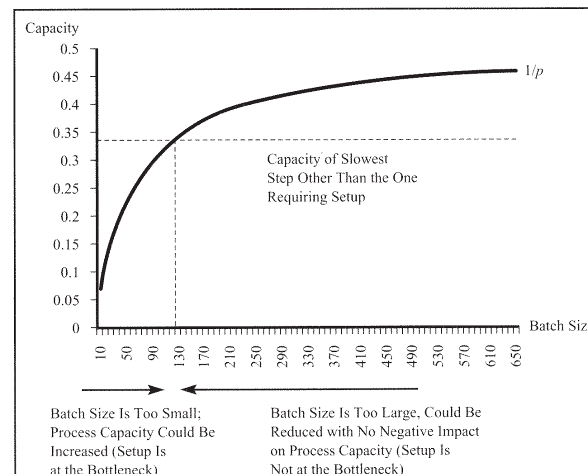

Exhibit 6.1

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 before the resource starts processing the first type of flow units again.
3. Compute the time in a production cycle that the resource is in setup; setup times are those times that are independent of the batch size.
4. Compute the time in a production cycle that the resource is processing; this includes all the activity times that are incurred per unit (i.e., are repeated for every member of the batch).
5. Compute the capacity of the resource with setup for a given batch size:

$$\text{Capacity } (B) = \frac{B}{\text{Setup time} + B \times \text{Time per unit}}$$

6. We are looking for the batch size that leads to the lowest level of inventory without affecting 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:

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

6.4 Balancing Setup Costs with Inventory Costs: The EOQ Model

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.
- Assume that we are charged a per-time-unit usage fee for a particular resource (e.g., for the milling machine discussed above). Thus, every minute we use the resource, independent of whether we use it for setup or for real production, we have to pay for the resource. In this case, "time is money" and the setup time thereby translates directly into setup costs. However, as we will discuss below, one needs to be very careful when making the conversion from setup times to setup costs.
- 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 three 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.

If we can reduce per-unit costs by increasing the batch size, what keeps us from using infinitely (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 6.8 illustrates the inventory levels for the two cases described above. The lower part of Figure 6.8 shows the case of the outside supplier and all units of a batch arriving at

FIGURE 6.8
Different Patterns of Inventory Levels

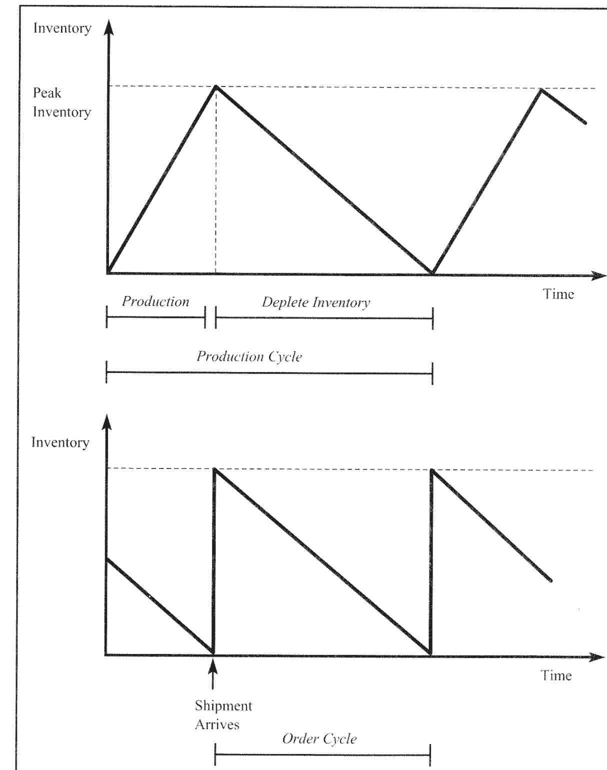

the same moment in time. The moment a shipment is received, the inventory level jumps up by the size of the shipment. It then falls up to the time of the next shipment.

The upper part of Figure 6.8 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 6.8). 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 one unit every three 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 one unit every three 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.

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.

The cost of inventory depends on how much it costs us to hold one unit in inventory for a given period of time, say one 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 one week. Given that a handle cap costs \$0.85 per unit, this translates to an inventory cost of $h = 0.007692 \times \$0.85/\text{unit} = \$0.006538$ per unit and 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, etc).

How many handle caps will there be, on average, in Nova Cruz's inventory? As we can see in Figure 6.8, the average inventory level is simply

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

If you are not convinced, refer in Figure 6.8 to 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

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

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

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

Note that purchase costs are not included in $C(Q)$ for the reasons discussed earlier. 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 6.9 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.

From Figure 6.9 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

$$\text{Economic order quantity} = \sqrt{\frac{2 \times \text{Setup cost} \times \text{Flow rate}}{\text{Holding cost}}}$$

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

FIGURE 6.9 Inventory and Ordering Costs for Different Order Sizes

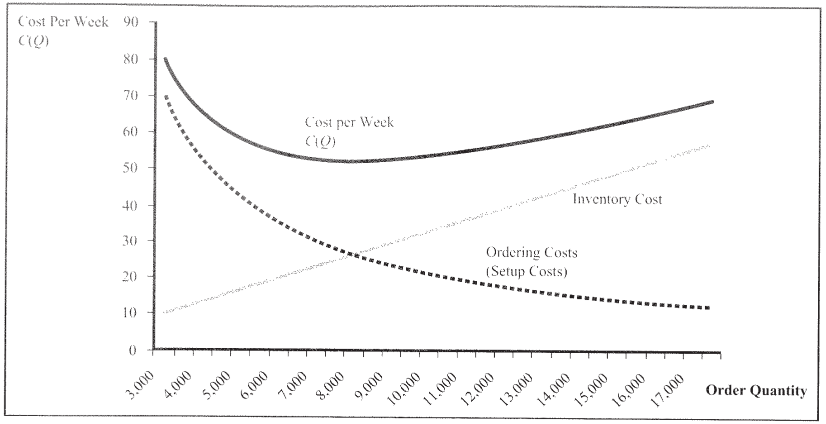

As our intuition suggests, as the setup costs K increase, we should make larger orders, but as holding costs h increase, we should make smaller orders.

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

$$Q^* = \sqrt{\frac{2 \times \text{Setup cost} \times \text{Flow rate}}{\text{Holding cost}}}$$

$$= \sqrt{\frac{2 \times 300 \times 700}{0.006538}} = 8{,}014.69$$

The steps required to find the economic order quantity are summarized by Exhibit 6.2.

6.5 Observations Related to the Economic Order Quantity

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

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

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

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

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.

Exhibit 6.2

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

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

$$\text{Setup costs [per unit of time]} = \frac{K \times R}{Q}$$

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

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

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 6.2 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.

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 6.9). This suggests that if we were

TABLE 6.2 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%

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

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

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 retailing 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 five 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

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

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

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

where we have reduced the procurement cost by 5 percent (multiplied by 0.95) to reflect the quantity discount. (Note: the 5% 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 6.8).

All computations we performed above can be easily transformed to this more general case (see, e.g., Nahmias 2000). 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 of Figure 6.6 and assume that the monthly capital cost for milling machine 1 is $9,000, which corresponds to $64 per hour (assuming four 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 6.3 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.

6.6 Transfer Batches

Up to this point, we have assumed that a batch would stay together as a collection of flow units throughout the process. However, in many settings, processed units are forwarded to the next process step, although some of the flow units in the same batch are still in process at the previous step. The following definitions are useful distinctions.

- For process steps involving setups, we can define a *production batch* as a collection of flow units that is produced between two setups.
- We can define a *transfer batch* as a collection of flow units that is transferred as an entity or group from one process step to another.

Both production and transfer batches reflect setups/economies of scale. Just as we have seen from Table 6.1 that it is not economical to produce in batch sizes of one when there are setups, similarly, material handling and transportation aspects of the process become increasingly complex and costly with smaller batch sizes. For example, a forklift transferring five cases of goods takes as long for the transportation time as a forklift with 50 cases.

If the nature of the transportation steps permits, it might be possible to reduce the transfer batch size significantly below the production batch size. In this case, flow units of the same production batch would be processed at a downstream step, while some of their "peers" are still being produced at the step with setup. Working with smaller transfer batches has the direct benefit of shortened flow time and thereby—because of Little's Law—of lower overall inventory.

It is also important to keep in mind that internal transportation processes are not adding any value to the customer. Hence, one should look for opportunities to eliminate the need for internal transportation by grouping resources with flows between them physically close to each other. The extreme case of this is an assembly-line layout.

However, even if not organized as an assembly line (e.g., the milling machine in the Xootr example), a firm should always attempt to transfer flow units one by one. This is the only way to have "the flow go on" (or, in the spirit of the book, have a flow rate matching the demand rate). The Toyota Production System also advocates the idea of piece-by-piece transfer of parts under the word *ikko-nagashi*.

6.7 Setup Time Reduction

Despite improvement potential from the use of "good" batch sizes and smaller transfer batches, setups remain a source of disruption of 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 creators of the SMED method referred to any setup exceeding 10 minutes as an unacceptable source of process flow disruption. 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 four hours. The SMED method helps to define an aggressive, yet realistic setup time goal and to identify potential opportunities of 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).

6.8 Other Flow Interruptions: Buffer or Suffer

In addition to illustrating the SMED method, the race car example also helps to illustrate how the concept of batching can be applied to *continuous process flows,* as opposed to discrete manufacturing environments. First of all, we observe that the calculation of the average speed of the race car is nothing but a direct application of the batching formula introduced at the beginning of this chapter:

$$\text{Average speed (number of miles between stops)} = \frac{\text{Number of miles between stops}}{\text{Duration of the stop} + \text{Time to cover one mile} \times \text{Number of miles between stops}}$$

In continuous flow processes, the quantity between two flow interruptions is frequently referred to as a production run.

Consider the production of orange juice, which is produced in a continuous flow process. At an abstract level, orange juice is produced in a three-step process: extraction, filtering, and bottling. Given that the filter at the second process step has to be changed regularly, the process needs to be stopped for 30 minutes following every four hours of production. While operating, the step can produce up to 100 barrels per hour.

To determine the capacity of the filtering step, we use

$$\text{Capacity }(B) = \frac{B}{S + B \times p}$$

$$= \frac{\text{Amount processed between two stops}}{\text{Duration of stop} + \text{Time to produce one barrel} \times \text{Amount processed between two stops}}$$

$$= \frac{400 \text{ barrels}}{30 \text{ minutes} + 60/100 \text{ minutes per barrel} \times 400 \text{ barrels}}$$

$$= \frac{400 \text{ barrels}}{270 \text{ minutes}}$$

$$= 1.48 \; barrels/minute = 88.88 \; barrels/hour$$

While in the case of batch flow operations we have allowed for substantial buffer sizes between process steps, the process as described in Figure 6.10 is currently operating without buffers. This has substantial implications for the overall flow rate.

Analyzing each step in isolation would suggest that the extraction step is the bottleneck, which would give us a process capacity of 80 barrels per hour. However, in the absence of buffers, the extraction step needs to stop producing the moment the filtering step is shut down. Thus, while running, the process is constrained by the extraction step, producing an output of 80 barrels per hour, and while being shut down, the process step is constrained by the filtering step (at 0 barrel per hour).

Previously, we considered the setup step in isolation from the rest of the process. That is a valid analysis if the setup step indeed works in isolation from the rest of the process, that is, if there is sufficient inventory (buffers) between steps. That assumption is violated here: The filtering step cannot operate at 88 barrels per hour because it is constrained by the extraction step of 80 barrels per hour.

For this reason, when we use our equation

$$\text{Capacity} = \frac{\text{Amount processed between two stops}}{\text{Duration of stop} + \text{Time to produce one barrel} \times \text{Amount processed between two stops}}$$

FIGURE 6.10
Data for the Production of Orange Juice

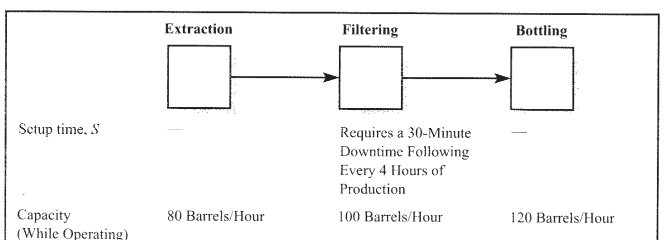

it is important that we acknowledge that we are producing at a rate of 80 barrels per hour (i.e., $\frac{1}{80}$ hour per barrel) while we are at the filtering step. This leads to the following computation of process capacity:

$$\text{Capacity} = \frac{320 \text{ barrels}}{0.5 \text{ hour} + 320 \text{ barrels} \times \frac{1}{80} \text{ hour per barrel}}$$

$$= 320 \text{ barrels}/4.5 \text{ hours}$$

$$= 71.11 \text{ barrels}/\text{hour}$$

This prompts the following interesting observation: In the presence of flow interruptions, buffers can increase process capacity. Practitioners refer to this phenomenon as "buffer or suffer," indicating that flow interruptions can be smoothed out by introducing buffer inventories. In the case of Figure 6.10, the buffer would need to absorb the outflow of the extraction step during the downtime of the reduction step. Thus, adding a buffer between these two steps would indeed increase process capacity up to the level where, with 80 barrels per hour, the extraction step becomes the bottleneck.

6.9 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 (see orange juice example)—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 versus sedans, Figure 6.4), 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 (this will be further discussed in Chapter 14). 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 6.11 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.

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.

FIGURE 6.11 **Summary of Batching**

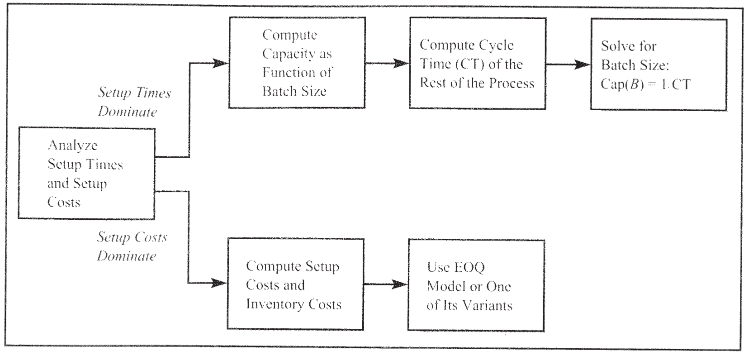

This tension is illustrated by Figure 6.12. Similar to the case of line balancing, we observe that adjustments in the batch size are not trading in one performance measure against the other, but allow us to improve by reducing current inefficiencies in the process.

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.

FIGURE 6.12
Choosing a Batch Size

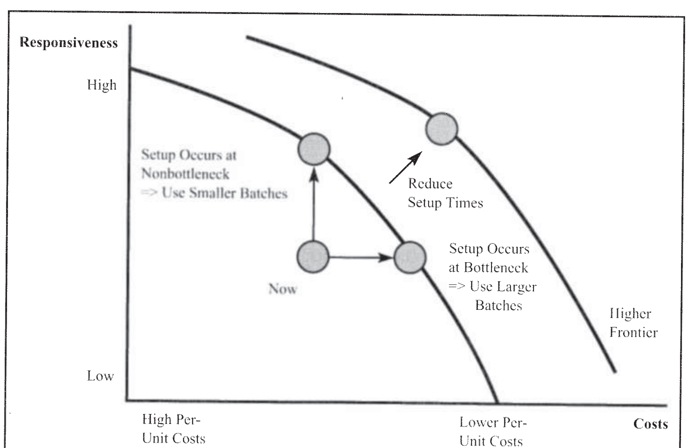

6.10 Further Reading

Nahmias (2000) is a widely used textbook in operations management that discusses, among other things, many variants of the EOQ model.

6.11 Practice Problems

Q6.1* **(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 activity time for each part A is one minute while the activity 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. Completed parts move from the stamping machine to the assembly only after the entire batch is complete.

At 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?
b. What batch size would you recommend for the process?

Q6.2 **(Simple Setup)** Consider the following batch flow process consisting of three process steps performed by three machines:

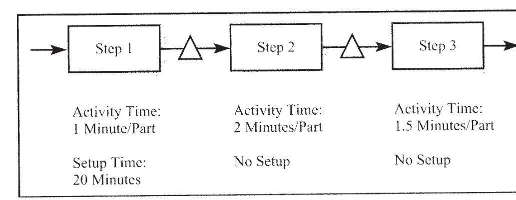

Work is processed in batches at each step. Before a batch is processed at step 1, the machine has to be set up. During a setup, the machine is unable to process any product.

a. Assume that the batch size is 50 parts. What is the capacity of the process?
b. For a batch size of 10 parts, which step is the bottleneck for the process?
c. Using the current production batch size of 50 parts, how long would it take to produce 20 parts starting with an empty system? Assume that the units in the batch have to stay together (no smaller transfer batches allowed) when transferred to step 2 and to step 3. A unit can leave the system the moment it is completed at step 3. Assume step 1 needs to be set up before the beginning of production.
d. Using the current production batch size of 50 parts, how long would it take to produce 20 parts starting with an empty system? Assume that the units in the batch do *not* have to stay together; specifically, units are transferred to the next step the moment they are completed at any step. Assume step 1 needs to be set up before the beginning of production.
e. What batch size would you choose, assuming that all units of a batch stay together for the entire process?

(* indicates that the solution is at the end of the book)

Q6.3 **(Setup Everywhere)** Consider the following batch-flow process consisting of three process steps performed by three machines:

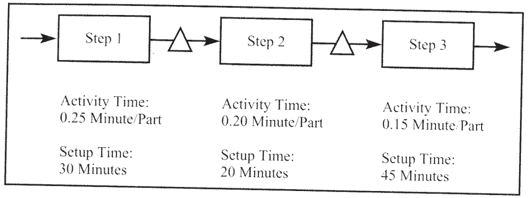

Work is processed in batches at each step. Before a batch is processed at a step, the machine at that step must be set up. (During a setup, the machine is unable to process any product.) Assume that there is a dedicated setup operator for each machine (i.e., there is always someone available to perform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

a. What is the capacity of step 1 if the batch size is 35 parts?

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

Q6.4 **(JCL Inc.)** JCL Inc. is a major chip manufacturing firm that sells its products to computer manufacturers like Dell, Gateway, 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. There is unlimited space for buffer inventory between these 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.
Activity 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?

b. For the current batch size of 100 wafers, how long would it take to produce 50 wafers? Assume that the batch needs to stay together during deposition and patterning (i.e., the firm does not work with transfer batches that are less than the production batch). However, the 50 wafers can leave the process the moment all 50 wafers have passed through the etching stage. Recall that a setup can only be started upon the arrival of the batch at the machine.

c. For what batch size is step 3 (etching) the bottleneck?

d. Suppose JCL Inc. came up with a new technology that eliminated the setup time for step 1 (deposition), but increased the activity time to 0.45 min./unit. What would be the batch size you would choose so as to maximize the overall capacity of the process, assuming all units of a batch stay together for the entire process?

Q6.5 **(Kinga Doll Company)** Kinga Doll Company manufactures eight versions of its popular girl doll, Shari. The company operates on a 40-hour work week. 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 activity 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
Activity 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?

b. What is the time it takes for the first batch of 500 dolls to go through an empty process?

c. Which batch size would minimize inventory without decreasing the process capacity?

d. Which batch size would minimize inventory without decreasing the current flow rate?

Q6.6 **(Bubba Chump Shrimp)** The Bubba Chump Shrimp Company processes and packages shrimp for sale to wholesale seafood distributors. The shrimp are transported to the main plant by trucks that carry 1,000 pounds (lbs.) of shrimp. Once the continuous flow processing of the shrimp begins, *no* inventory is allowed in buffers due to spoilage and all of the shrimp must be processed within 12 hours to prevent spoilage. The processing begins at the sorter, where the trucks dump the shrimp onto a conveyor belt that feeds into the sorter, which can sort up to 500 lbs. per hour. The shrimp then proceed to the desheller, which can process shrimp at the rate of 400 lbs. per hour. However, after 3 hours and 45 minutes of processing, the desheller must be stopped for 15 minutes to clean out empty shrimp shells that have accumulated. The veins of the shrimp are then removed in the deveining area at a maximum rate of 360 lbs. per hour. The shrimp proceed to the washing area, where they are processed at 750 lbs. per hour. Finally, the shrimp are packaged and frozen.

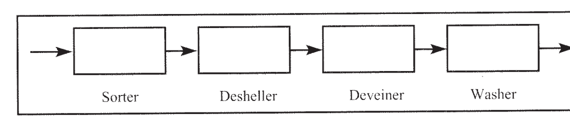

Note: All unit weights given are in "final processed shrimp." You do *not* need to account for the weight of the waste in the deshelling area. The plant operates continuously for 12 hours per day beginning at 8:00 a.m. Finally, there is negligible time to fill the system in the morning.

a. What is the daily process capacity of the desheller (in isolation from the other processes)?

b. What is the daily process capacity of the deveiner (in isolation from the other processes)?

c. What is the daily process capacity of the processing plant (excluding the packaging and freezing)?

d. If five trucks arrive one morning at 8:00 a.m., what is the total number of pounds of shrimp that must be wasted?

Q6.7 **(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?
b. What is CLI's total order cost for one year?
c. What is CLI's total holding cost for one year?
d. What is CLI's weekly inventory turns?

Q6.8 **(Millennium Liquors)** Millennium Liquors is a wholesaler of sparkling wines. Their most popular product is the French Bete Noire 1989. 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 and handling. These costs include the usual ordering and handling costs, plus the cost of refrigeration, which includes a fixed component (mainly depreciation of the cooling equipment) and a variable component that depends on the number of cases in inventory.

- Cost per case: $120
- Shipping cost (for any size shipment): $290
- Cost of labor to place and process an order: $10
- Cost of labor to place cases into warehouse: $2/case
- Cost of labor to pick case when sold: $2/case
- Fixed cost for refrigeration: $75/week
- Variable cost for refrigeration: $3/case/week

a. Calculate the weekly holding cost for one case of wine.
b. Use the EOQ model to find the number of cases per order and the average number of orders per year.
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?

Q6.9 **(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?
b. Given your answer in (a), how many times a year does PBK place orders?
c. Given your answer in (a), how many months of supply of coffee does PBK have on average?
d. On average, how many dollars per month does PBK spend to hold coffee (including cost of capital)?

Suppose that a South American import/export company has offered PBK a deal for the next year. PBK can buy a years' worth of coffee directly from South America for $20 per bag and a fixed cost for delivery of $2,000. 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.

Q6.10* **(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 two weeks at a cost of $8 per case. The inventory-related holding cost (capital, insurance, etc.) 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)?
b. For the next three parts, assume the distributor selects the order quantity specified in part (a).
c. What are the distributor's inventory turns per year?
d. What is the inventory-related cost per case of beer sold?
e. 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?

(* indicates that the solution is at the end of the book)