
Introduction

A central premise in economics is that prices adjust to match supply with demand: if there is excess demand, prices rise; if there is excess supply, prices fall. But while an economist may find comfort with this theory, managers in practice often do not. To them excess demand means lost revenue and excess supply means wasted resources. They fully understand that matching supply with demand is extremely difficult and requires more tools than just price adjustments.

Consider the following examples:

- In 2006, Nintendo launched the Wii game console with much success—so much success that the company could not make enough units to keep up with demand. Some entrepreneurs would wait in long lines to purchase scarce units only to turn around and sell them online for several hundred dollars over the retail price.
- In 2007, Dell lost its worldwide market share leadership to HP. Trying to regain momentum, Dell offered laptop computers to consumers in various colors. Unfortunately, problems with dust contamination in the painting process prevented Dell from ramping up production, causing long delays, which in turn caused some customers to cancel their order.
- At 4 p.m. on weekdays, it is hard to find a taxi in Manhattan because that is when taxis tend to change between shifts. Consequently, customers wait longer for a cab.
- In March 2011, a massive earthquake hit Japan, followed by devastating tsunamis. Supplies for some key automobile and electronic components were unavailable or scarce for months, disrupting production around the globe.
- In 2008, Boeing was unable to deliver on time its new 777s to Emirates Airlines because growth in demand caught its supplier of kitchen galleys off guard and short on capacity.
- In early 2002, a victim of a car crash in Germany died in a rescue helicopter after the medical team together with their dispatcher had unsuccessfully attempted to find a slot in an operating room at eight different hospitals. In the United States, every day there are thousands of patients requiring emergency care who cannot be transported to the nearest emergency room and/or have to wait considerable time before receiving care.
- The average customer to Disney World experiences only nine rides per day, in part because of long queues. To give customers a better experience (read, “more rides”), Disney developed several mechanisms to encourage customers to find rides with short or no queues.

All of these cases have in common that they suffer from a mismatch between demand and supply, with respect either to their timing or to their quantities.

This book is about how firms can design their operations to better match supply with demand. Our motivation is simply stated: By better matching supply with demand, a firm gains a significant competitive advantage over its rivals. A firm can achieve this better

match through the implementation of the rigorous models and the operational strategies we outline in this book.

To somewhat soften our challenge to economic theory, we do acknowledge it is possible to mitigate demand–supply mismatches by adjusting prices. For example, the effective market price of the Wii game console did rise due to the strong demand. But this price adjustment was neither under Nintendo’s control, nor did Nintendo (or its retailers) collect the extra surplus. In other words, we view that price adjustment as a symptom of a problem, rather than evidence of a healthy system. Moreover, in many other cases, price adjustments are impossible. The time period between the initiation of demand and the fulfillment through supply is too short or there are too few buyers and sellers in the market. There simply is no market for emergency care in operating rooms, waiting times in call centers, or piston rings immediately after an earthquake.

Why is matching supply with demand difficult? The short answer is that demand can vary, in either predictable or unpredictable ways, and supply is inflexible. On average, an organization might have the correct amount of resources (people, product, and/or equipment), but most organizations find themselves frequently in situations with resources in the wrong place, at the wrong time, and/or in the wrong quantity. Furthermore, shifting resources across locations or time is costly, hence the inflexibility in supply. For example, physicians are not willing to rush back and forth to the hospital as they are needed and retailers cannot afford to immediately move product from one location to another. While it is essentially impossible to always achieve a perfect match between supply and demand, successful firms continually strive for that goal.

Table 1.1 provides a sample of industries that we will discuss in this book and describes their challenge to match supply with demand. Take the airline industry (last column in Table 1.1.). For fiscal year 2007, British Airways achieved a 76.1 percent utilization; that is, a 160-seat aircraft (the average size in their fleet) had, on average, 122 seats occupied with a paying passenger and 38 seats flying empty. If British Airways could have had four more (paying) passengers on each flight, that is, increase its utilization by about 2.5 percent, its corporate profits would have increased by close to £242 million, which is about 44 percent of its operating profit for 2007. This illustrates a critical lesson: Even a seemingly small

TABLE 1.1 Examples of Supply–Demand Mismatches

	Retailing	Iron Ore Plant	Emergency Room	Pacemakers	Air Travel
Supply	Consumer electronics	Iron ore	Medical service	Medical equipment	Seats on specific flight
Demand	Consumers buying a new video system	Steel mills	Urgent need for medical service	Heart surgeon requiring pace-maker at exact time and location	Travel for specific time and destination
Supply exceeds demand	High inventory costs; few inventory turns	Prices fall	Doctors, nurses, and infrastructure are underutilized	Pacemaker sits in inventory	Empty seat
Demand exceeds supply	Forgone profit opportunity; consumer dissatisfaction	Prices rise	Crowding and delays in the ER; potential diversion of ambulances	Forgone profit (typically not associated with medical risk)	Overbooking; customer has to take different flight (profit loss)
Actions to match supply and demand	Forecasting; quick response	If prices fall too low, production facility is shut down	Staffing to predicted demand; priorities	Distribution system holding pacemakers at various locations	Dynamic pricing; booking policies

(continued)

TABLE 1.1 Concluded

Managerial importance	Per-unit inventory costs for consumer electronics retailing all too often exceed net profits	Prices are so competitive that the primary emphasis is on reducing the cost of supply	Delays in treatment or transfer have been linked to death	Most products (valued \$20k) spend 4–5 months waiting in a trunk of a salesperson before being used	About 30% of all seats fly empty; a 1–2% increase in seat utilization makes the difference between profits and losses
Reference	Chapter 2, The Process View of the Organization; Chapter 12, Betting on Uncertain Demand: The Newsvendor Model; Chapter 13, Assemble-to-Order, Make-to-Order, and Quick Response with Reactive Capacity	Chapter 3, Understanding the Supply Process: Evaluating Process Capacity; Chapter 4, Estimating and Reducing Labor Costs	Chapter 8, Variability and Its Impact on Process Performance: Waiting Time Problems; Chapter 9, The Impact of Variability on Process Performance: Throughput Losses	Chapter 14, Service Levels and Lead Times in Supply Chains: The Order-up-to Inventory Model	Chapter 16, Revenue Management with Capacity Controls

improvement in operations, for example, a utilization increase of 2.5 percent, can have a significant effect on a firm's profitability precisely because, for most firms, their profit (if they have a profit) is a relatively small percentage of their revenue. Hence, improving the match between supply and demand is a critically important responsibility for a firm's management.

The other examples in Table 1.1 are drawn from a wide range of settings: health care delivery and devices, retailing, and heavy industry. Each suffers significant consequences due to demand–supply mismatches, and each requires specialized tools to improve and manage its operations.

To conclude our introduction, we strongly believe that effective operations management is about effectively matching supply with demand. Organizations that take the design of their operations seriously and aggressively implement the tools of operations management will enjoy a significant performance advantage over their competitors. This lesson is especially relevant for senior management given the razor-thin profit margins firms must deal with in modern competitive industries.

1.1 Learning Objectives and Framework

In this book, we look at organizations as entities that must match the supply of what they produce with the demand for their product. In this process, we will introduce a number of quantitative models and qualitative strategies, which we collectively refer to as the “tools of operations management.” By “quantitative model” we mean some mathematical procedure or equation that takes inputs (such as a demand forecast, a processing rate, etc.) and outputs a number that either instructs a manager on what to do (how much inventory to buy, how many nurses to have on call, etc.) or informs a manager about a relevant performance measure (e.g., the average time a customer waits for service, the average number of patients in the emergency room, etc.). By “qualitative strategy” we mean a guiding principle: for example, increase the flexibility of your production facilities, decrease the variety of products offered, serve customers in priority order, and so forth. The next section gives a brief description of the key models and strategies we cover. Our learning objective for

this book, put as succinctly as we can, is to teach students how and when to implement the tools of operations management.

Just as the tools of operations management come in different forms, they can be applied in different ways:

1. Operations management tools can be applied to ensure that resources are used as efficiently as possible; that is, the most is achieved with what we have.
2. Operations management tools can be used to make desirable trade-offs between competing objectives.
3. Operations management tools can be used to redesign or restructure our operations so that we can improve performance along multiple dimensions simultaneously.

We view our diverse set of tools as complementary to each other. In other words, our focus is neither exclusively on the quantitative models nor exclusively on the qualitative strategies. Without analytical models, it is difficult to move beyond the “blah-blah” of strategies and without strategies, it is easy to get lost in the minutia of tactical models. Put another way, we have designed this book to provide a rigorous operations management education for a strategic, high-level manager or consultant.

We will apply operations tools to firms that produce services and goods in a variety of environments—from apparel to health care, from call centers to pacemakers, and from kick scooters to iron ore fines. We present many diverse settings precisely because there does not exist a “standard” operational environment. Hence, there does not exist a single tool that applies to all firms. By presenting a variety of tools and explaining their pros and cons, students will gain the capability to apply this knowledge no matter what operational setting they encounter.


Consider how operations tools can be applied to a call center. A common problem in this industry is to find an appropriate number of customer service representatives to answer incoming calls. The more representatives we hire, the less likely incoming calls will have to wait; thus, the higher will be the level of service we provide. However, labor is the single largest driver of costs in a call center, so, obviously, having more representatives on duty also will increase the costs we incur per call.

The first use of operations management tools is to ensure that resources are used as effectively as possible. Assume we engage in a benchmarking initiative with three other call centers and find that the performance of our competitors behaves according to Figure 1.1: Competitor A is providing faster response times but also has higher costs. Competitor B has longer response times but has lower costs. Surprisingly, we find that competitor C outperforms us on both cost and service level. How can this be?

It must be that there is something that competitor C does in the operation of the call center that is smarter than what we do. Or, in other words, there is something that we do in our operations that is inefficient or wasteful. In this setting, we need to use our tools to move the firm toward the frontier illustrated in Figure 1.1. The frontier is the line that includes all benchmarks to the lower left; that is, no firm is outside the current frontier. For example, a premium service might be an important element of our business strategy, so we may choose not to compromise on service. And we could have a target that at least 90 percent of the incoming calls will be served within 10 seconds or less. But given that target, we should use our quantitative tools to ensure that our labor costs are as low as possible, that is, that we are at least on the efficiency frontier.

The second use of operations management tools is to find the right balance between our competing objectives, high service and low cost. This is similar to what is shown in Figure 1.2. In such a situation, we need to quantify the costs of waiting as well as the costs of labor and then recommend the most profitable compromise between these two objectives.

FIGURE 1.1
Local Improvement
of Operations by
Eliminating
Inefficiencies



Moving to the frontier of efficiency and finding the right spot on the frontier are surely important. But outstanding companies do not stop there. The third use for our operations management tools is to fundamentally question the design of the current system itself. For example, a call center might consider merging with or acquiring another call center to gain scale economies. Alternatively, a call center might consider an investment in the development of a new technology leading to shorter call durations.

FIGURE 1.2
Trade-off between
Labor Productivity
and Responsiveness

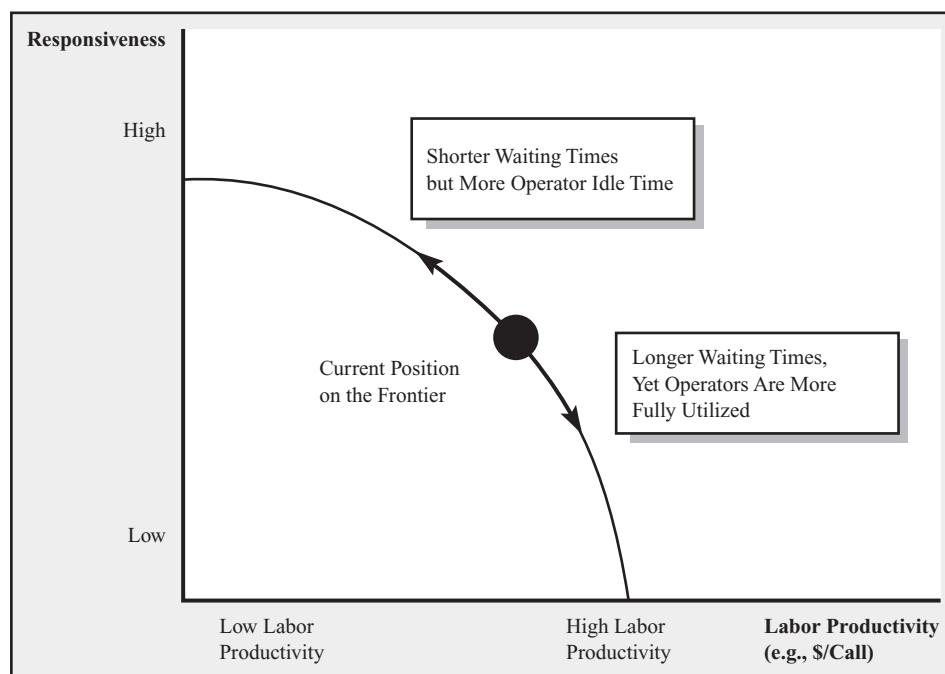
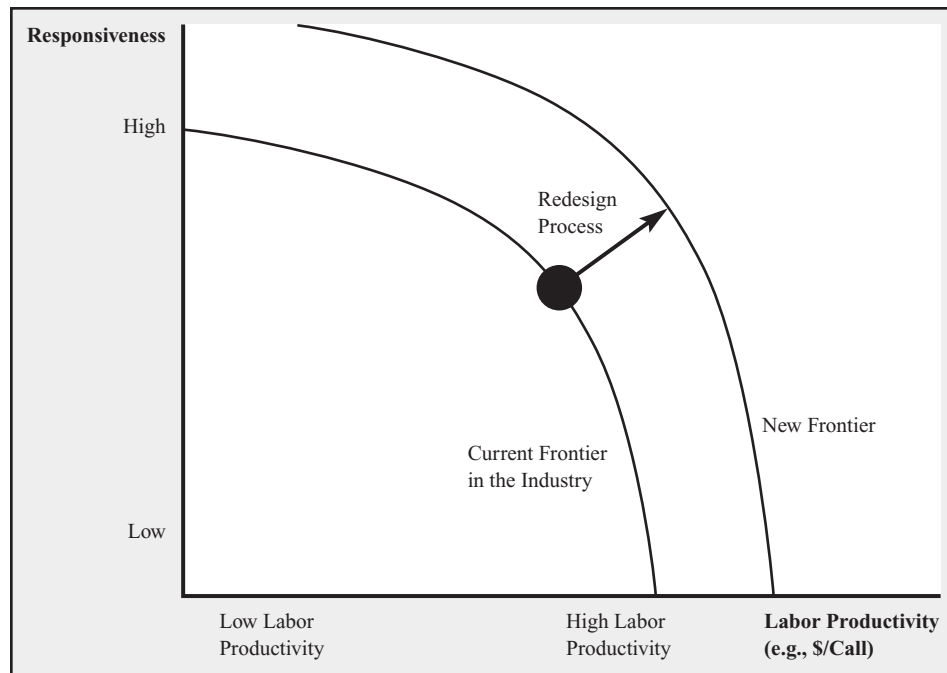


FIGURE 1.3
Redesigning the
Process to Operate at
an Improved Frontier



In such cases, a firm pushes the envelope, that is, moves the frontier of what previously was feasible (see Figure 1.3). Hence, a firm is able to achieve faster responsiveness and higher labor productivity. But, unfortunately, there are few free lunches: while we have improved both customer service and labor productivity, pushing out the frontier generally requires some investments in time and effort. Hence, we need to use our tools to quantify the improvements we can achieve so that we can decide whether the effort is justifiable. It is easy to tell a firm that investing in technology can lead to shorter call durations, faster service, and higher labor productivity, but is that investment worthwhile? Our objective is to educate managers so that they can provide “big ideas” and can back them up with rigorous analysis.

1.2 Road Map of the Book

This book can be roughly divided into five clusters of closely related chapters.

The first cluster, Chapters 2–7, analyzes business processes (the methods and procedures by which a service is completed or a good is produced). For the most part, the view taken in those chapters is one of process without variability in service times, production times, demand arrival, quality, and so forth. Hence, the objective is to organize the business process to maximize supply given the resources available to the firm. Chapters 8–11 introduce variability into business process analysis. Issues include the presence of waiting times, lost demand due to poor service, and lost output due to poor quality. This cluster concludes with an overview of the Toyota Production System.

Chapters 12–15 discuss inventory control, information management, and process flexibility. Issues include demand forecasting, stocking quantities, performance measures, product design, and production flexibility.

Chapter 16 departs from a focus on the supply process and turns attention to the demand process. In particular, the chapter covers the tools of revenue management that allow a firm to better match its demand to its fixed supply.

TABLE 1.2
A High-Level
Grouping of
Chapters

Chapters	Theme
2–7	Process analysis without variability in service times, production rates, demand arrival, quality, etc.
8–11	Process analysis with variability in service times, production rates, demand arrival, quality, etc.
12–15	Inventory control, information management, process flexibility
16	Revenue management
17–19	Strategic operations management: supply chains, sustainability, and business models

Chapters 17–19 conclude the book with several strategic topics, including the management and control of the supply chain, sustainability, and business model innovation.

Table 1.2 summarizes these clusters.

The following provides a more detailed summary of the contents of each chapter:

- Chapter 2 defines a process, introduces the basic process performance metrics, and provides a framework for characterizing processes (the product–process matrix). Little’s Law is introduced, an essential formula for understanding business processes and the link between operations management and financial accounting.
- Chapter 3 introduces process analysis tools from the perspective of a manager (as opposed to an engineer): how to determine the capacity of a process and how to compute process utilization.
- Chapter 4 looks at assembly operations with a specific focus on labor costs, an extremely important performance metric. It frequently drives location decisions (consider the current debate related to offshoring) and has—especially in service operations—a major impact on the bottom line. We define measures such as labor content, labor utilization, and idle time. We also introduce the concept of line balancing.
- Chapter 5 investigates project management, a process that is designed for a single, somewhat unique, project such as a ship, a new building, or a satellite.
- Chapter 6 connects the operational details of process analysis with key financial performance measures for a firm, such as return on invested capital. Through this chapter we discover how to make process improvement translate into enhanced financial performance for the organization.
- Chapter 7 studies production in the presence of setup times and setup costs (the EOQ model). A key issue is the impact of product variety on production performance.
- Chapter 8 explores the consequences of variability on a process. As we will discuss in the context of a call center, variability can lead to long customer waiting times and thereby is a key enemy in all service organizations. We discuss how an organization should handle the trade-off between a desire for minimizing the investment into capacity (e.g., customer service representatives) while achieving a good service experience for the customer.
- Chapter 9 continues the discussion of variability and its impact on service quality. As we will discuss in the context of emergency medicine, variability frequently can lead to situations in which demand has to be turned away because of insufficient capacity. This has substantial implications, especially in the health care environment.
- Chapter 10 details the tools of quality management, including statistical process control, six-sigma, and robust design.
- Chapter 11 describes how Toyota, via its world-famous collection of production strategies called the Toyota Production System, achieves high quality and low cost.

- Chapter 12 focuses on the management of seasonal goods with only one supply opportunity. The newsvendor model allows a manager to strike the correct balance between too much supply and too little supply.
- Chapter 13 expands upon the setting of the previous chapter by allowing additional supply to occur in the middle of the selling season. This “reactive capacity” allows a firm to better respond to early season sales information.
- Chapter 14 continues the discussion of inventory management with the introduction of lead times. The order-up-to model is used to choose replenishment quantities that achieve target availability levels (such as an in-stock probability).
- Chapter 15 highlights numerous risk-pooling strategies to improve inventory management within the supply chain: for example, location pooling, product pooling, universal design, delayed differentiation (also known as postponement), and capacity pooling.
- Chapter 16 covers revenue management. In particular, the focus is on the use of booking limits and overbooking to better match demand to supply when supply is fixed.
- Chapter 17 identifies the bullwhip effect as a key issue in the effective operation of a supply chain and offers coordination strategies for firms to improve the performance of their supply chain.
- Chapter 18 applies operations management thinking to the challenge of sustainability.
- Chapter 19 concludes the book with how operations management enables new business models. A framework is presented for understanding business model innovation that can assist in the creation of new business models.

Some of the chapters are designed to be “entry level” chapters, that is, chapters that can be read independently from the rest of the text. Other chapters are more advanced, so they at least require some working knowledge of the material in another chapter. Table 1.3 summarizes the contents of the chapters and indicates prerequisite chapters.

TABLE 1.3 Chapter Summaries and Prerequisites

Chapter	Managerial Issue	Key Qualitative Framework	Key Quantitative Tool	Prerequisite Chapters
2: The Process View of the Organization	Understanding business processes at a high level; process performance measures, inventory, flow time, and flow rate	Product–process matrix; focus on process flows	Little’s Law Inventory turns and inventory costs	None
3: Understanding the Supply Process: Evaluating Process Capacity	Understanding the details of a process	Process flow diagram; finding and removing a bottleneck	Computing process capacity and utilization	Chapter 2
4: Estimating and Reducing Labor Costs	Labor costs	Line balancing; division of labor	Computing labor costs, labor utilization Minimizing idle time	Chapters 2, 3
5: Project Management	Time to project completion	Critical path	Critical path analysis	Chapters 2, 3
6: The Link between Operations and Finance	Process improvement to enhance corporate performance	Return on Invested Capital (ROIC) tree	Computing ROIC	Chapters 2, 3
7: Batching and Other Flow Interruptions: Setup Times and the Economic Order Quantity Model	Setup time and setup costs; managing product variety	Achieving a smooth process flow; deciding about setups and ordering frequency	EOQ model Determining batch sizes	Chapters 2, 3

(continued)

TABLE 1.3 Concluded

Chapter	Managerial Issue	Key Qualitative Framework	Key Quantitative Tool	Prerequisite Chapters
8: Variability and Its Impact on Process Performance: Waiting Time Problems	Waiting times in service processes	Understanding congestion; pooling service capacity	Waiting time formula	None
9: The Impact of Variability on Process Performance: Throughput Losses	Lost demand in service processes	Role of service buffers; pooling	Erlang loss formula Probability of diverting demand	Chapter 8
10: Quality Management, Statistical Process Control, and Six-Sigma Capability	Defining and improving quality	Statistical process control; six-sigma	Computing process capability; creating a control chart	None
11: Lean Operations and the Toyota Production System	Process improvement for competitive advantage	Lean operations; Toyota Production System	—	None
12: Betting on Uncertain Demand: The Newsvendor Model	Choosing stocking levels for seasonal-style goods	Improving the forecasting process	Forecasting demand The newsvendor model for choosing stocking quantities and evaluating performance measures	None
13: Assemble-to-Order, Make-to-Order, and Quick Response with Reactive Capacity	How to use reactive capacity to reduce demand–supply mismatch costs	Value of better demand information; assemble-to-order and make-to-order strategies	Reactive capacity models	Chapter 12
14: Service Levels and Lead Times in Supply Chains: The Order-up-to Inventory Model	Inventory management with numerous replenishments	Impact of lead times on performance; how to choose an appropriate objective function	The order-up-to model for inventory management and performance-measure evaluation	Chapter 12 is highly recommended
15: Risk Pooling Strategies to Reduce and Hedge Uncertainty	How to better design the supply chain or a product or a service to better match supply with demand	Quantifying, reducing, avoiding, and hedging uncertainty	Newsvendor and order-up-to models	Chapters 12 and 14
16: Revenue Management with Capacity Controls	How to manage demand when supply is fixed	Reserving capacity for high-paying customers; accepting more reservations than available capacity	Booking limit/protection level model; overbooking model	Chapter 12
17: Supply Chain Coordination	How to manage demand variability and inventory across the supply chain	Bullwhip effect; supply chain contracts	Supply chain contract model	Chapter 12
18: Sustainability	How to employ operations management techniques to a sustainability initiative	Measuring resource use and emissions	—	None
19: Business Model Innovation	How to create a business model innovation	Customer value curve and supply process transformation	—	None