
WHY STUDY SERVICES?

1.1 WHAT ARE SERVICES?

We use services every day of our lives. When we wake up to the voice of our favorite radio broadcaster on our bedside alarm clock, we are using the services of the radio station as well as those of the electric company providing our home with power. A quick shower entails the use of the water supplied by the water company. The bowl of cereal that we eat at breakfast was purchased from a retail grocery store, or perhaps delivered to our house by a food delivery service, either of which received the box of cereal using trucking services. The pants that we put on before leaving the house were cleaned at the local dry cleaners, yet another service we use implicitly before leaving the house.

After alighting from our bus (a public transportation service), we stop at the local coffee shop for a latte before walking to the office. Chances are we work for a service provider, such as a lawyer, doctor, accountant, or consultant. During the day at the office, we use numerous communication services, including e-mail, phones, faxes, the Internet, and instant messaging. To learn more about the new software our employer installed last week, we may employ the services of an online training program or a consultant at the software vendor's call center.

Lunch with friends at the nearby burger house demands the services of this establishment, including the waiters, chefs, busboys, and cashier. As we pay with our credit card, we utilize the financial services of our bank. The afternoon call to our dentist to schedule our next checkup brings us in touch with the vast array of medical services at our disposal. After a quiet dinner at home, we drive to a nearby theater to catch the latest movie, taking advantage of the myriad of entertainment services that surrounds us.

We take many services for granted. For example, in driving to the movie theater, we may have forgotten about the public road maintenance services that patched the potholes left by last winter's snows. We surely forgot about the snow plowing services that removed the snow. Unless we recently had our car repaired, we didn't consider the auto maintenance and repair services that keep our cars going. We took for granted the police services that kept our streets and neighborhoods safe. Most of us are probably unaware of the traffic monitoring services provided by many large municipalities and counties to ensure the smooth flow of traffic on our highways and arterials. If the theater was one we frequent often, we probably did not need the GPS service in our car or on our cell phone. And throughout the day, we implicitly relied on our life, auto, health, homeowner's, and property insurance services.

In short, we rely on and utilize services for many of our daily needs.

So what are services and how can we characterize them? Table 1.1 lists some of the services that we frequently encounter. This should give the reader a sense of the broad spectrum of economic activity that is encompassed by the service sector. In further defining the service sector, it may be easier to define what it is not. Agriculture is not part of the service sector, though agriculture relies on transportation carriers (part of the service sector) to move seed, feed, and fertilizer to farms and to ship produce and animals to producers and markets. The construction industry, which builds roads, airports, manufacturing plants, office buildings, and much more, is not part of the service sector, though it too depends on many services. Manufacturers of durable goods such as automobiles, washers, driers, and refrigerators are not part of the service sector, though they too rely heavily on services for the goods they produce. Manufacturers depend on the financial industry for loans for inventory, raw materials and machinery; on accountants to help manage their finances; on the legal profession to ensure that they comply with government regulations; and consultants for a wide range of professional services. Some durable goods manufacturers are also in the service business. For example, General Motors seems to spend as much time advertising its OnStar emergency and vehicle diagnostic service as it does marketing the vehicles it produces. Often, government activities and employees are separated from the service sector, though many governmental activities are rightly considered services. For example, police, fire, and emergency medical personnel clearly provide important services to the communities in which they are located, as do teachers, many of whom are employed by their local public school districts.

Several attributes characterize most services. First, it is difficult to inventory services. Empty seats at a Sunday matinee of a musical cannot be sold for the

TABLE 1.1. Some of the services we commonly use

Financial	Savings Checking ATMs Credit card Mortgage services CDs Stocks Bonds Mutual funds	Education	Primary school Secondary school Community college Four-year college University Online training Private school Parochial school
Insurance	Life Health Dental Disability Automobile Homeowner/property Long-term care	Entertainment	Movies Live theater Concerts Sports Television Radio
Retail	Grocery store Book store Gasoline station Jewelry store Department store Specialty clothing Electronics warehouse	Transportation	Parking Taxi Bus Train Airplane Truck Rail freight
Health	Social worker Psychiatrist Personal physician Specialist Dentist Clinic Hospital Emergency room	Personal Services	Hair stylist Spa Massage Dry cleaning Shoe repair Accountant Financial consultant Lawyer
Communication Services	Post office/U.S. mail Express mail Land line telephone Cell phone E-mail Text message Internet	Public Services	Library Police Fire Emergency services Garbage collection Roads Water Electricity
Dining	Fast food Family dining School cafeteria Ethnic restaurant Fine dining	Other	Call center Lawn care Snow removal Pet care

following Saturday night when demand may be high. Students who miss a class on optimization cannot sit in on a class on French literature and make up for the material on optimization that they missed. A dry cleaner who goes on vacation for two weeks cannot make up for the lost business during that time by working overtime because her customers are likely to have taken their clothing to another cleaner during that time. Airline seats that are flown empty on one flight cannot later be used for a flight that is overbooked. Nevertheless, services often depend critically on the careful management of inventories. Without an adequate inventory of sterile equipment and materials, surgeries cannot be performed. A retail establishment must maintain an adequate inventory of goods to be able to serve its customers. Airlines think of their seats as a perishable inventory that they manage using sophisticated revenue management techniques. Videotaping and recording of lectures, coupled with online material, may allow students to “inventory” a lecture for future reference.

Second, there is an intangible quality associated with services. You cannot touch the interaction between a psychiatrist and his or her patient, just as you cannot hold the education that a student receives in high school. Third, services typically produce some value for customers, or the service provides a solution to a customer need. The value may be in the form of improved health care, entertainment, or education. The value may also provide peace of mind (as in the case of insurance, police, fire, and emergency medical services) or convenience (as provided, for example, by dry cleaners). Fourth, services typically involve an interaction between a service provider and a consumer of the service. For example, doctors interact with patients in providing health care services. Teachers and professors interact with students in the educational process. Lawyers provide advice and counsel to their clients.

Finally, the need for services varies significantly with time. This is one of the key distinguishing features of services as opposed to manufacturing operations where the number of products produced each day is likely to remain roughly constant. An automobile assembly plant can produce approximately one vehicle a minute or about 1,000 cars per day during two eight-hour shifts. The production rate remains the same throughout the day. Emergency room arrivals, on the other hand, vary significantly over the course of the day. For example, in Austin, Texas, we found that the ratio of the peak hourly demand for ambulance services to the demand during the least busy hour exceeded four to one (Eaton et al., 1979). Traffic on most urban highways also exhibits clear daily patterns with low volumes in the late evening and early morning hours and heavy traffic during the morning and evening rush hours. The demand for air travel also exhibits weekly and annual peaks. The Sunday of Thanksgiving weekend is typically one of the busiest times of the year for most airlines. Retail establishments must often sell a large fraction of their goods during the Christmas holiday season if they are to have a successful year. Accountants experience a rush of business in the first three and a half months of the year, as clients try to meet the April 15 tax filing deadline. Some services, such as summer-stock theaters or winter ski resorts,

operate only during certain times of the year. Other services, such as insurance sales and library usage, exhibit significantly lower degrees of temporal variation.

Planning for these peaks and valleys in demand is a critical task for a service manager. If possible, one would like to smooth the demand. In some cases, pricing can be used to shift demand from peak to off-peak periods. For some services, this is not possible. The demand for emergency medical services, for example, is not likely to be influenced significantly by any pricing scheme that a hospital or municipality might consider putting in place. Similarly, Thanksgiving holiday travel may decrease overall with significant price increases, but people are not likely to shift that travel to other times of the year since Thanksgiving is when most employees and most students have two (or more) days off from work or classes. Supply can also be adjusted in many cases. Additional flights can be scheduled for busy periods. Larger aircraft can be assigned to those routes that are known to experience the most significant percentage increases in demand during a busy period. Additional temporary accounting staff are hired during the tax filing period. More doctors and nurses can be scheduled to work in an emergency room during the busier hours of the day and week. Reversible traffic lanes are used in many urban areas to accommodate the daily rush hour traffic.

Many services have evolved significantly since the turn of the century. The Internet has spurred much of this evolution. For example, in booking an airline reservation, one used to have to contact the airline or a travel agent directly. Today, many people book reservations online with little to no human interaction. This function of travel agents has been largely replaced by the Internet. Toward the end of the twentieth century, it was nearly impossible to research the housing stock in a community without the aid of a realtor. Today, most sophisticated home buyers conduct a significant amount of research regarding communities and home prices before visiting their first house. Smart and sophisticated realtors encourage and facilitate this sort of research.

To illustrate the evolution of services, consider what has happened to photography during the twenty-first century. In the pre-Internet, pre-digital camera era, photographs were taken on film. The film was developed at a local photo shop and converted into pictures and negatives or slides. Alternatively, some film manufacturers provided customers with mailers to return the film for processing by the company. Today, relatively few people rely on film; most have converted to digital photography. With this conversion has come an explosion in the number of options people have for processing photographs. It is worth noting, that most digital photographs are probably never printed as they can be viewed on the camera or on a computer. Only the winners need be printed and processed. In addition to the local photo shop as a means of printing pictures, consumers can now print their photographs on an inexpensive color printer hooked to their home computer. Many printers now allow users to bypass the computer completely with the camera or memory card linked directly to the printer. Photographs can be uploaded to an online photo

processor. Pictures can then be ordered for mail delivery or pickup at a local store. Many drug stores and other establishments have stand-alone kiosks to allow photo enthusiasts to print pictures, again with little or no interaction with the store owners. Finally, digital photos are finding their way into far more than pictures. Online companies allow customers to order mugs, photo books, magnets, calendars, clothing and much more.

Many other services are also being transformed by the Internet. In his seminal book entitled, *The World is Flat*, Thomas Friedman (2005) points out that many accountants are sending simple tax returns to India and other low-labor-cost countries for processing. Accountants are now finding more sophisticated ways of meeting their clients' needs. Medical images taken at a hospital in Boston, Massachusetts, may be read by a radiologist in Boston; they could just as easily, however, be read by a radiologist in Beer Sheva, Israel or Bangalore, India.

1.2 SERVICES AS A PERCENT OF THE ECONOMY

In measuring the economy, the U.S. Bureau of Economic Analysis divides the gross domestic product—the value of all goods and services produced in the country—into these broad categories:

- a. **Durable goods**, including such items as motor vehicles and parts, furniture and household equipment
- b. **Non-durable goods**, including food, clothing, shoes, gasoline, oil, fuel, and other energy supplies
- c. **Exports**, which is the value of all exported goods and services
- d. **Imports**, which is the value of all imported goods and services and which contributes *negatively* to the gross domestic product
- e. **Gross private domestic investment**, which includes the value of all inventory and machinery in the country
- f. **Government expenditures**, including federal defense and non-defense related expenditures as well as all state and local government expenditures
- g. **Services** including expenditures on housing, transportation, household operating expenses (electricity, gas and other operating costs), medical care, recreation, and so on.

Figure 1.1 plots the value of these components of the gross domestic product beginning with the first quarter of 1947 and running through the first quarter of 2008 in current dollars. The explosive growth of the value of the service sector of the economy is readily apparent. In 2000 dollars, the GDP increased from \$1.57 trillion to \$11.7 trillion (\$0.237 to \$14.2 trillion in current dollars) during

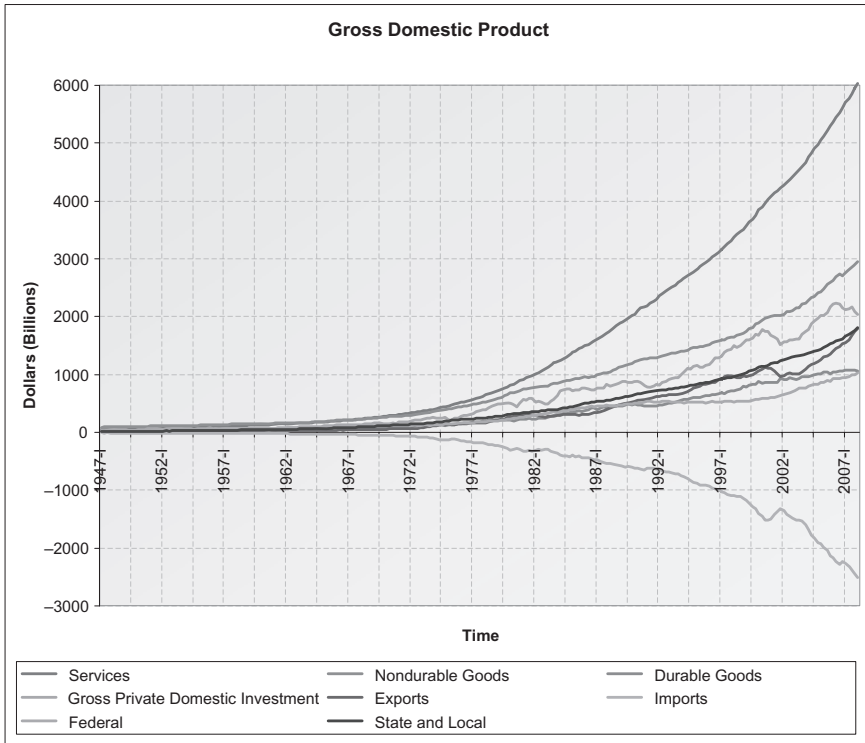


Figure 1.1. Gross domestic product over 60 years. (See color insert)

Based on the U.S. Department of Commerce, Bureau of Economic Analysis, National Income and Product Accounts Table 1.1.5 Gross Domestic Product (Seasonally adjusted at annual rates)

this interval. This is an increase of 745% in real dollars or a compounded growth rate of over 3.3 percent per year. In the 60-year period from 1940 to 2000, the U.S. population grew from 142.2 million to 291.4 million, a compound growth rate of only 1.2 percent.

Figure 1.2 is perhaps more enlightening, as it plots the percentage of the gross domestic product in each of these areas in 10-year increments beginning with the first quarter of 1948 and going through the first quarter of 2008. While the value of durable goods remained roughly constant at about 8 or 9 percent of the GDP, the value of services more than doubled from 20.5 percent in 1948 to 42.5 percent of the GDP in 2008. During this interval, the value of non-durable goods has decreased from 36.5 percent of the GDP to 20.8 percent. Up until the late 1960s, the United States was generally a net exporter; the United States has been a net importer of goods and services during the subsequent four decades (<http://www.u-s-history.com/pages/h980.html>).




Figure 1.2. Percentage breakdown of the gross domestic product

While services represent roughly 45 percent of the gross domestic product, the number of people employed in the service industry is an even greater percentage of the total employment. Table 1.2 shows the breakdown in non-farm employment in the United States in June 2009. (Note that less than 2 percent of the U.S. workforce is employed on farms, so looking only at non-farm employment is quite reasonable.) Of the 131.4 million employees, less than 15 percent are employed in goods-producing industries including natural resource extraction, construction, and manufacturing. About the same percentage (17.1 percent) are government employees (many of whom can be viewed as being in service-related jobs), while over two out of every three employees (nearly 69 percent) work in the service sector according to the U.S. Bureau of Labor Statistics.

Figure 1.3 shows the breakdown of the U.S. employment in percentage terms for June 2009. Figure 1.4 breaks down the service sector employment into the major categories shown in Table 1.2. Two out of every seven service sector employees work in the trade, transportation, or utility industries, another 19 percent work in professional and business services, while 17 percent work in health care.

TABLE 1.2. Breakdown of U.S. employment in June 2009

Non-farm Jobs in the US—June, 2009			
		Raw #s	% of previous
Total Non-farm Employment ($\times 10^3$)	131,411		
Goods Producing		18,713	14.2%
Natural Resources		715	3.8%
Construction		6,162	32.9%
Manufacturing		11,836	63.3%
Private Service Providing		90,223	68.7%
Trade, Transportation, Utilities		25,174	27.9%
Information		2,834	3.1%
Financial Activities		7,737	8.6%
Professional and Business Services		16,624	18.4%
Education		3,072	3.4%
Health care		16,190	17.9%
Leisure and Hospitality		13,177	14.6%
Other		5,415	6.0%
Government		22,475	17.1%
Federal		2,826	12.6%
State		5,149	22.9%
Local		14,500	64.5%

Based on the U.S. Bureau of Labor Statics, Table B-3. Employees on nonfarm payrolls by major industry sector and selected industry detail, seasonally adjusted (<http://www.bls.gov/ces/#tables>, <ftp://ftp.bls.gov/pub/suppl/empst/empst.ceseeb3.txt>)

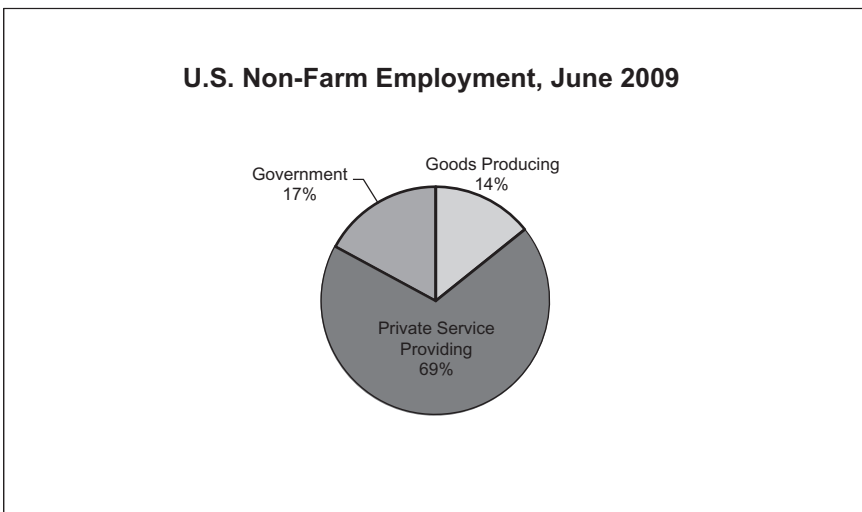


Figure 1.3. U.S. non-farm employment by major sector, June 2009

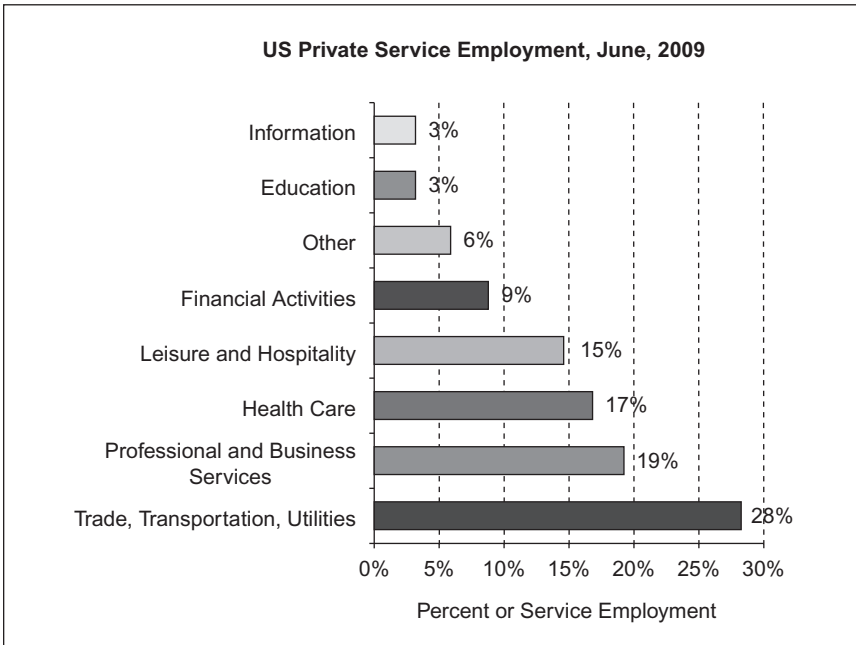


Figure 1.4. U.S. private service employment, June 2009

1.3 PUBLIC VERSUS PRIVATE SERVICE DELIVERY

Services are provided by both the public and private sectors. The public sector refers to governmental agencies, while the private sector encompasses private companies. For the purposes of this discussion, we will consider non-profit agencies, like many food pantries, to be governmental agencies since, in important ways, they behave in a manner that is similar to government agencies.

One of the fundamental differences between services provided by a public agency and those delivered by a private company has to do with the service provider's objective. Private companies, even though they are in the service provision business, often act just like other companies. Their primary objectives have to do with profit maximization or maximizing the return on their shareholders' equity. For example, the primary objective of a company providing cellular phone service is to maximize its profits. In those cases in which their (short-term) revenues are fixed, such firms attempt to minimize the cost of delivering the service. In modeling such firms, it is important to include a constraint that guarantees a minimum level of service. For example, while the long-term objective of a package delivery service company is to maximize profit, its short-term day-to-day objective might be to minimize its costs. Clearly, the firm can minimize its costs by doing nothing! However, it typically is constrained to deliver the packages scheduled for that day. With this constraint, it wants to do so at minimum total cost.

Public agencies and most non-profit agencies operate with different objectives and different constraints. Such agencies typically want to maximize some measure of service, such as the number of clients served, while operating within a fixed, and often very tight, budget. For example, an ambulance service might want to locate its bases to maximize the number of people who live within a given time of the nearest ambulance base, since time is of the essence in medical emergencies. The ambulance service will typically have sufficient funds to build and operate only a limited number of bases. In other cases, an agency may want to maximize some measure of the equity of the service provided to its clients. For example, maximizing the minimum fill rate (the ratio of the delivered food to the requested food) was the objective used in allocating collected food from food donors to food pantries in the Greater Chicago Food Depository program (Lien, Iravani, and Smilowitz, 2007). A vehicle would visit donor companies—typically grocery stores or restaurants—collecting fresh food donations in the morning and would then deliver the food to roughly half a dozen food kitchens over the course of the afternoon. Without some attempt to achieve equity between the food kitchens, the first kitchen visited might receive a large allocation, leaving little if any food for the last kitchens visited on the vehicle’s route.

Not only is the profit maximization objective typically replaced by a service-related objective when considering public sector service providers, but also we typically have to capture multiple objectives (Cohon, 1978). For example, in locating ambulances, not only is it important to maximize the number of people served adequately—those for whom an ambulance is located within a specified time standard—but it is also important to ensure equity across different socio-economic groups within the community. There may well be important tradeoffs between the number of people who can be served adequately and a measure of equity across all groups. Also, a public agency may want to minimize the disruption to the community, thereby introducing a third objective. In the case of ambulance bases, the city providing the service may want to maximize the number of bases that are located on existing city land, or equivalently to minimize the number of pieces of property that need to be seized using eminent domain laws. This will introduce additional tradeoffs and will make the analysis and modeling of the problem that much more complicated and difficult. Johnson and Smilowitz (2007) outline many important issues in community-based operations research, a growing field that is closely related to public sector service modeling.

While many of the models that we outline will be appropriate for both public and private sector service providers, it is important that we bear in mind the differences between these two contexts when determining the appropriate objective(s) and constraints for any model.

1.4 WHY MODEL SERVICES?

From the discussion above, it is clear that services constitute a significant part of the U.S. economy and increasingly of the global economy. Over two out of

every three employees work in the service industry and over 45 percent of the \$13.8 trillion 2007 GDP is attributable to services. A number of good books have been written outlining the service industry (Davis and Heineke, 2003; Fitzsimmons and Fitzsimmons, 2008; Metters, King-Metters, and Pullman, 2003). While providing an excellent introduction to the service industry and its components in general, these texts do not focus on modeling some of the key decisions that managers of service-oriented firms or agencies need to make.

Many service providers face similar decisions. Much of the focus of this book is on those decisions or classes of problems that are common across many service providers, be they in the public or private sectors. These decisions include:

- Location decisions – how many facilities to have and where should they be located
- Resource allocation decisions – how to allocate scarce resources (such as class capacity) to demands (such as students wishing to take each course)
- Short-term workforce scheduling issues – how many employees to schedule for each shift and when each shift should begin
- Long-term workforce management problems—how many employees to hire, release, and retrain, at which locations, and in which positions
- Inventory problems (arising in retail services as well as a range of other services) – how much safety stock inventory to maintain, when to place orders, and how much inventory to order
- And vehicle routing problems—how many routes are needed to serve a customer base (e.g., clients for Meals on Wheels), how to assign customers to routes, and how to sequence the stops on the route to achieve a service provider's objectives

In many of these cases, the fundamental underlying issue is that of striking the right balance between the demand for services and the supply of those services. Alternatively, this can often be viewed as striking a balance between the level of customer service that is achieved and the cost of providing that level of service. Intuitively we would expect that as more service is provided, the cost of providing the service increases but the quality of service also improves. For example, if a grocery store increases the number of cashiers on duty, its labor costs will go up, but the quality of service, as measured by the time patrons must wait to pay for their groceries, will improve. Similarly, if a popular restaurant increases its seating capacity with a concomitant increase in wait staff, chefs and other personnel, its costs will increase. At the same time, however, the waiting time for a table during the peak dining hours will decrease. As more screeners are added at an airport security checkpoint, costs increase but passenger delay times decrease.

1.5 KEY SERVICE DECISIONS

In the long-run, service providers must make decisions regarding the *number and location of facilities*. For example, a hospital system may need to decide how to allocate services to hospitals and whether additional sites are needed to accommodate a growing demand for medical services. Complicating the planning process is the need to have some services co-located at the same site. Also, any plan must evolve over time as the population ages and changes demographically (Santibáñez, Bekiou, and Yip, 2009).

A coffee company like Starbucks must decide how many stores to open and where they should be. In fact, Starbucks seems to have opened too many stores and, in the summer of 2008, announced that it was closing 600 stores in the United States (de la Merced, 2008). Just as the decision of where to open stores can be modeled, so too can the decision regarding the number of stores to close and the location of those stores. While this is a private sector problem for which we earlier argued that profit maximization should be the long-run objective, news reports, editorials and commentaries (Wong and Rose, 2008) suggest that even such corporate decisions must account for community-based objectives if the firm is to avoid negative publicity.

In some industries, the service region needs to be partitioned or divided between sales representatives. This is a districting problem that generally has multiple objectives. Typically, one wants the districts to be contiguous so that a salesperson does not service one county in New York, for example, and another in southern California. In addition, one often wants to equalize the workload (e.g., the number of potential customers) in such districts. At the same time, one would typically prefer a compact (close to circular) district with small customer-to-base travel distances to one that is highly elongated with longer travel distances. This too is a form of a location problem.

Resource allocation problems are also common in the provision of services. For example, at Northwestern University, half of the freshmen are in the Weinberg College of Arts and Sciences (WCAS). Each of these 1,000 students must take a seminar in the fall quarter and another in the winter or spring quarter. During the fall quarter, there are roughly 70 seminars, each with a capacity of about 15 students. Students must be assigned to seminars so that each student is in exactly one seminar, the seminar capacities are not violated and the assignment maximizes a measure of student satisfaction.

Similarly, each quarter or semester, at every college or university in the country, classes need to be assigned to rooms and times. Clearly one objective is to minimize the number of classes that do not have adequate seating capacity for those wishing to enroll in the course. At the same time, we do not want to assign a small class to a large auditorium as this makes for an uncomfortable teaching and learning experience for all involved. (I know this from personal experience.) In addition, we typically want to minimize both faculty and student walking times. Faculty walking times or distances are generally easier to estimate because we can generally assume that faculty leave for class from their offices,

while the location of students immediately before a class is less certain. They could be in their dormitories, at a student cafeteria, in the library, or in another class. Finally, we need to ensure that no faculty member has to teach two classes at the same time. We also would like to respect faculty requests for teaching schedules as much as possible. Some faculty members prefer mornings to afternoons; some prefer longer Tuesday/Thursday classes to shorter Monday/Wednesday/Friday classes.

Airlines operating at large hubs like Chicago, Atlanta, Los Angeles, the various New York airports, Dallas, Houston, and San Francisco must assign planes to gates. Airlines today operate hub and spoke systems with aircraft converging on the hub at about the same time and then leaving at about the same time roughly an hour or so later. This allows time for passengers to change planes and to make connections. Airlines may try to assign flights to gates to minimize the total walking distance of all connecting passengers. Typically, this problem is constrained by the capacity of specific gates to accommodate different aircraft types.

Almost all services, with the possible exception of Internet-based services, need to *schedule employees*. As noted above, the demand for services varies dramatically over time. In many cases, such as the provision of emergency services and restaurant services, the demand exhibits daily patterns with significant peaks and troughs. Key questions that can be answered with appropriate models include:

- How many employees to have on duty at each time of the day?
- When should the employees begin and end their shifts?
- How many full-time and part-time employees should be utilized?
- How should cross-trained employees (those with multiple skills) be deployed?
- What is the tradeoff between employee costs and customer service?

For some industries, the temporal variability occurs over the time span of months. For example, the demand for accounting services peaks in the first three and a half months of each year in the United States as citizens prepare their personal income tax returns. Community organized or for-profit summer camps, parks, and swimming pools must add staff during the summer months. Retail stores add staff for the pre-holiday sales season. The key decision in these cases revolves around how much staff to add and when to begin the staff in any required training program.

Many service providers must deal with *long-term personnel management*. For example, a management consulting firm must determine how many new hires to bring on board with various skills and how to assign those employees to projects. While this may seem on the surface to be a simple issue of managing an inventory of employees—a problem for which inventory management is, perhaps, well-suited—employees differ significantly from an inventory of drill bits at a

home improvement store. First, employees (unlike drill bits) have preferences for where they are deployed and which projects they work on. A failure to account for such preferences can result in disgruntled employees and higher than necessary turnover rates. Second, employees typically come with a bundle of skills whereas a $\frac{1}{4}$ -inch drill bit has only a limited number of attributes. Third, and perhaps most importantly, employees acquire and can be taught new skills while untapped and unused knowledge may be forgotten or become out-of-date. A drill bit will not “forget” what it is supposed to be able to do; at the same time, it is unlikely to acquire a new skill and suddenly know how to rip a sheet of $\frac{1}{2}$ -inch plywood. In short, while inventory theory may teach us some lessons about personnel management, we need to account for the unique features of people in managing personnel.

While most writers use the presence or absence of inventory as a key distinguishing feature between manufacturing sector and service sector, *inventory decisions* play a key role in the provision of many services. As such, it is appropriate and important that students and managers of the service industries understand fundamental issues of inventory management. For example, retail stores must manage large quantities of inventory if they are to remain in business. When the skies open up in a sudden downpour, customers running into the local drug store for shelter want to be able to buy an umbrella. They are not willing to backorder one, only to pick it up a week later on a sunny day. Auto maintenance and repair facilities must carry an inventory of the commonly used spare parts. Few, if any, drivers would be willing to wait three days while a local service station ordered a few quarts of motor oil for a \$20 oil change.

But inventory issues extend far beyond the obvious ones of managing a physical stock of items. Airlines view the seats on their flights as a perishable inventory in the sense that once the flight has departed any unsold seats have “perished” or “spoiled.” Yield and profit management systems are used to manage this inventory of seats. Hotels face a similar problem. A room that is not rented on a given night cannot be placed in inventory to be rented during an evening when the hotel is full. At the same time, the hotel does not want to rent a room several months in advance at a low family rate when they could hold the room in “inventory” and rent it to a businessperson willing to pay a far higher rate.

Finally, the newsvendor problem in inventory theory has much to teach us about services as well. In its classical statement, the newsvendor problem is that of a small convenience store determining how many of the daily newspaper to order. If it orders too many, it loses money on unsold copies. If the store orders too few, customers arriving late in the day will be disappointed at not being able to buy a paper, resulting in a loss of goodwill on the part of the store. In the long run, the store may lose such customers as they choose to buy not only newspapers, but other items as well, at a store that does have the paper in stock. Such models are applicable not only in this context but also in the context of seasonal clothing purchases for which there is likely to be a single purchasing window months in advance of the realization of demand. These models can also inform decisions about contracting for such services as snow plowing of one’s driveway.

Finally, employees who must make decisions about how many pre-tax dollars to set aside in flexible health care spending accounts are making decisions similar to those made by the newsvendor; the dollars must be committed before the need for those moneys is realized.

Finally, *vehicle routing* is important in the delivery of many services. Express mail companies such as UPS and FedEx are in the business of delivering small packages to home and business locations. For such firms, efficient vehicle routes can mean the difference between profitability and bankruptcy. In other cases, vehicle routing is a necessary part of doing business. For example, a company selling white goods (refrigerators, freezers, washing machines, and driers) must often deliver and install the items purchased by its customers. The firm must decide when to offer delivery to customers in each region that it services and then how to sequence customers on routes on the day of delivery. In a different context, many regions share an inventory of books across multiple libraries serving the area. If a book is not available at your local library, you can often request the book on an inter-library loan system. Such books are sent from one library to another on a dedicated fleet of vehicles operated by a regional library system. A library system may encompass hundreds of primary, middle and high school, university, hospital, and public libraries. Determining the frequency of pickups and dropoffs at each library in the system as well as the route of each of the vehicles in the fleet is a complex vehicle routing problem that significantly impacts the time a library patron must wait for a book to become available at his or her local library.

1.6 PHILOSOPHY ABOUT MODELS

George Box, the noted statistician, said, “Essentially, all models are wrong, but some are useful” (Box and Draper, 1987). A model is an abstraction of reality that may ignore many details of the real world context. This is true of all models. A picture of the heart in a text on biology is a model of the heart. Clearly, the model does not beat. The picture also does not pump any blood. The picture is often cut away in a manner that would be fatal to any patient with such a “wound” in his or her heart. Nevertheless, the picture, coupled with the associated text, is useful in explaining to biology and medical students the various components of the heart, their placement relative to each other, and the functions they perform.

The same is true of models of services. For example, while we know that the demand for emergency services varies by time of day, most facility location models ignore this temporal variability as it is unlikely to impact long-term decisions about where to locate fixed ambulance bases. The model may not explicitly account for the uncertain nature of ambulance demand. The model may also assume that all ambulances are always available for service. Differences between patients in the required on-scene time as well as whether or not transport to a hospital is or is not required are also typically ignored. Nevertheless, location

models are often used to assist decision makers in locating ambulances because the models provide valuable *insights* into the problem and the impact of the decisions that need to be made (Eaton et al., 1985).

Similarly, a model of the delays experienced by callers to a call center is likely to make a number of assumptions that may not be supported by empirical studies. For example, we are likely to assume that calls arrive at the center according to a particular stochastic process known as a Poisson process. While this is often supported by the data, we may also assume that the duration of calls follows an exponential distribution. This is rarely validated by empirical results. Despite these flaws, the models are likely to provide valuable insights into the operation of a call center and are likely to help us determine the appropriate number of customer service agents to employ to balance operating costs and customer waiting time.

But the value of modeling a process extends far beyond any insights that the model may provide regarding the operation of the underlying system and the decisions needed to design the service. The process of developing a model often yields benefits in and of itself. This process is shown in Figure 1.5.

Modeling generally begins with a problem. For example, the problem might be that of assigning freshmen at a university to seminars in a more efficient and more equitable manner than the current manual process. The problem might be that of determining how many parking spots to build in a municipal parking garage. The problem might be that of identifying bases for emergency medical service vehicles. It could be identifying improvements in a hospital discharge planning process.

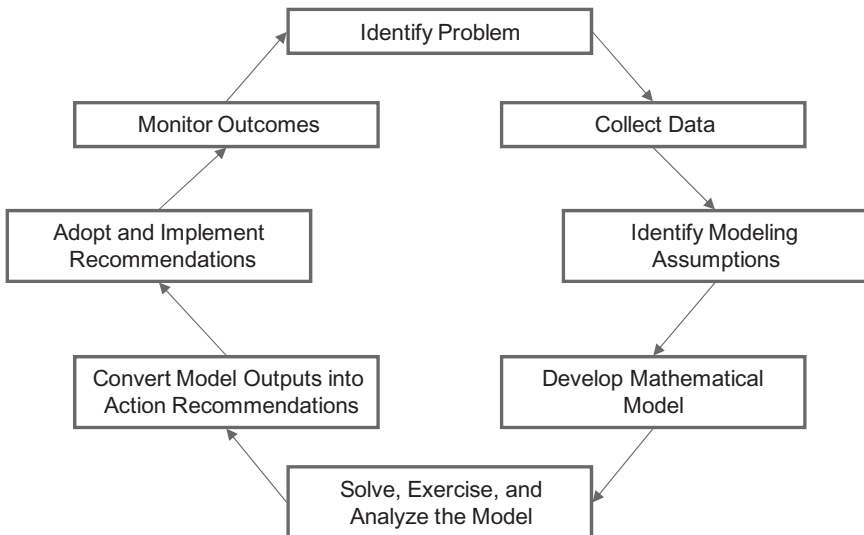


Figure 1.5. Schematic of the modeling/decision-making process

There are three key facets associated with identifying a problem for modeling and analysis. First, we need to identify the key stakeholders or groups impacted by the process. For example, in the case of hospital discharge planning, key constituents clearly include: (a) the patients and their families, (b) the physicians, (c) the nursing staff, (d) other patient-oriented staff including social workers and discharge planners, (e) the hospital administration, and (f) any institutions such as nursing facilities and rehabilitation institutes to which patients might be sent. Each of these groups may have different subgroups. For example, patients might be stratified by the cause of their hospitalization. The discharge process for a teenager with a broken leg is likely to differ significantly from that of an elderly patient who was hospitalized for a heart attack. Physicians include in-house hospitalists, specialists who treated the patient, and the patient's own internist. Each of these groups might have different perspectives on the issue.

Second, we need to identify the key objectives to be achieved. Two objectives in the case of discharge planning are (a) expediting the process to increase the flow of patients through the hospital and (b) enhancing the communication of discharge orders between the medical personnel (physicians and nursing staff) and the patient.

Third, we need to identify any constraints that may be affecting the system. In the case of discharge planning, constraints may include language barriers between patients and the medical staff. One community hospital in the Chicago area has to deal with over 40 languages for informed consent forms for surgery; a large metropolitan hospital is likely to encounter far more languages. Such language barriers are likely to be present not only at the time a patient is admitted for a surgical or medical procedure but also at the time the patient is discharged. Another significant constraint is that the medical staff has many other tasks to perform each day other than discharging patients.

Once the problem—stakeholders, objectives and key constraints—is identified, the next step involves collecting data on the current process. In the case of discharge planning, this would include data on the number of patients discharged each day, the destinations to which they are discharged, the reasons for their initial hospitalization, and so on. It would also include mapping the current discharge process so that the key steps are clearly identified.

The third key step in modeling a service operation is to identify any key assumptions that the model(s) might make. For example, if we are developing a simulation model of the discharge process, we need to structure the “arrival” process of patients ready to be discharged on any given day. What probabilistic laws does this process follow? Will we model the process as if the same number of patients are discharged each day? This is likely to be a very bad assumption. Which of the various steps in the discharge process will we model explicitly and which will we handle implicitly? How will we represent the delays in this process that are caused by physicians or nurses dealing with other patients who are not being discharged? If the model is an optimization model, explicitly enumerating the assumptions often entails identifying the nature of the relationships between the key decision variables as well as how those variables interact in the objective

function(s). For queueing models, we need to be explicit about the probabilistic form of the customer arrival and service processes as well as the queue discipline that specifies the order in which customers are to be served.

Developing the model is the next key step. In the case of a simulation model, this entails writing the computer code in the relevant simulation language to represent the modeled process. In the case of an optimization model, this involves writing down the objective function and constraints in terms of the input data and the decision variables. This mathematical model must then be translated into some form that a computer can understand. For example, many optimization models can be solved within spreadsheet systems like Microsoft Excel. Larger and more complex problems may require the use of specialized optimization languages such as AMPL (Fourer, Gay, and Kernighan, 2002) and IBM's ILOG CPLEX optimizer. Developing the model involves setting up the appropriate spreadsheets and/or writing the code needed to invoke specialized languages. Developing a queueing model may be simple if the model already exists in the literature. If it does not, the underlying equations for the queueing process need to be written down. Sometimes these can be solved in closed form, while in other cases they need to be solved numerically. In the latter case, developing the model entails writing the code or developing the spreadsheet model capable of solving these equations.

Solving, exercising, and analyzing the model is the next step. In many cases, solving the model is actually fairly straightforward once the model has been developed. For simulation models, one simply has to run the model. Similarly, many optimization models can be solved with software built into spreadsheet systems or with off-the-shelf optimization packages. If the results of a queueing model are known, there is nothing to solve! If they are not, numerical procedures in a spreadsheet can often solve a given instance of a queueing model quite adequately.

Exercising the model entails changing the key inputs to see how the model responds. This is necessary (1) to validate the model, (2) to test alternative policies, and (3) to determine the sensitivity of the results to changes in key model inputs. For example, in a simulation model we might increase the rate at which customers arrive at the system. If the performance metrics improve (i.e., if the customer waiting time decreases as the arrival rate increases) we have strong reasons to suspect that the model is not valid. Similarly, if attempts to replicate current conditions (e.g., the current discharge process at a hospital) result in times that differ significantly from those observed in practice, we again have reason to suspect that there is a problem with the model. If the model appears to be valid, we generally want to test alternative policies with the model. For example, if we simulated the discharge planning process, we might want to use the model to understand the impact of moving some steps of the process earlier. We might want to know, for example, how the average time required to discharge a patient—measured from the time a physician says a patient is ready to leave the hospital until the time the patient actually vacates the bed—changes if we begin arranging for a nursing home bed at the time critically ill patients are

admitted to the hospital, rather than waiting to do so until they are discharged. In the case of an optimization model for assigning students to seminars, we might want to see how many more students can get their first-choice seminar if we increase the number of students allowed in each seminar by one. For a queueing model, we might want to see how the average waiting time before service is impacted by changes in the average arrival rate of customers, the number of servers on duty, or the mean service time.

The model results must also be analyzed, often statistically. In the case of a simulation model, this is particularly important since the output of any simulation model is typically one realization of a random process. Running the simulation model a second time with different random numbers is likely to result in different output values. Thus, it is important to analyze the results to see if the predicted results of two different policies are indeed statistically significantly different. In the case of an optimization model, there may be multiple objectives and the tradeoff between these objectives may need to be identified as part of analyzing the results. For example, in assigning students to seminars, there is likely to be a tradeoff between maximizing the satisfaction with the assignments averaged over all students and minimizing the number of students who are particularly adversely affected by the assignment (e.g., those who get their fourth choice or worse!).

Once the model has been validated, exercised, and analyzed, we must convert the model results into action recommendations. In the case of discharge planning, this might involve recommending that additional hospitalists be employed to expedite the discharge process. It might mean recommending that physicians spend *additional* time with patients to ensure that the patients fully understand their discharge instructions. Perhaps one of the recommendations would be to begin certain processes, like planning for rehabilitation care or nursing care, as soon as it becomes apparent that the patient will need this additional level of care upon discharge, even if this means beginning the planning significantly earlier in a patient's stay.

It is worth noting that this is often the stage at which the formal modeling by an analyst ends and the work of the decision maker takes over. For example, in assigning students to seminars, it is probably not appropriate for a technical analyst to recommend one assignment plan over another. Rather, the job of the analyst may be to make explicit to the decision maker—the dean in this case—the tradeoffs that must be made in choosing one assignment over another.

The next step involves adopting and implementing the action recommendations. This is often complicated by the many stakeholders involved in the process. For example, a recommendation might be that a patient not be discharged until he or she can tell the nurse the top three discharge instructions. This sounds simple and is intended to ensure that the physician has adequately communicated the discharge instructions to the patient. However, adopting and implementing such a recommendation may be anything but simple. First, it requires that the physician identify the top three instructions and communicate to the nurse that these are, in fact, the key instructions that the patient needs to understand before

being discharged from the hospital. Second, the nurse now has added work as he or she must ask the patient what these key instructions are. Third, what should happen if the patient does not understand the instructions or cannot recite them back for the nurse? Should the physician be called back in to discuss the instructions again with the patient? Should the nurse discuss the instructions with the patient? In assigning students to seminars, the seminar assignments must be uploaded to the university or college course scheduling software and conflicts that might arise between the assigned seminars and other courses individual students have enrolled in must be resolved.

The last step involves monitoring the process to assess how well the recommendations have worked. Are discharge times reduced? Is there evidence that patients understand their discharge instructions better than they did before and that they are healthier following the revised discharge process than they had been under the old process? Are students happier with the new assignment process for freshman seminars? Have the number of complaints from parents been reduced as a result of the new assignment process?

Finally, Figure 1.5 shows a feedback link between the step of monitoring the process and the initial step of identifying the problem. In some cases, we find that the problems that initiated the study are still not resolved and the process needs to be repeated. Perhaps we need a more careful understanding of the objectives of the constituents. Maybe the initial process failed to account for some key constraints that resulted in an incomplete or failed implementation of the recommendations. Maybe the data that were collected were biased or erroneous in important ways, which resulted in serious discrepancies between the modeled results and those obtained in the real world. Perhaps the differences between the real and modeled world resulted from unrealistic assumptions that were made in the process of modeling the problem. In any event, it may be necessary to cycle through this modeling/decision making process until the monitored results are satisfactory.

It should be clear that there are likely to be many ancillary benefits associated with this modeling process aside from the direct benefits that result from the recommendations. First, the process forces us to identify the key stakeholders, their objectives, and the constraints under which they and the process operate. Identifying these people and sharing the various objectives with all stakeholders is likely to be a valuable exercise as all parties will begin to see the problem from a new perspective. Doctors will be forced to consider the discharge process from the perspective of patients, nurses, and social workers, while administrators will better understand the myriad demands placed on a doctor's limited time.

Similarly, the data collection process is likely to yield important side benefits. In the course of collecting data relevant to a problem, one often finds that important information is missing, thereby making it hard to assess quantitatively how well the current system is performing. What metrics should be used to measure a patient's understanding of the discharge instructions that he or she has been given? In many hospitals, there is limited, if any, data on such issues. Proxies,

such as the number of calls back to a physician regarding medication instructions, may be the best information available. Thus the data collection process may suggest that new metrics need to be devised and new data collected. In the course of collecting and using the data, errors in the dataset are likely to be discovered. Correcting these errors for future use is another ancillary benefit of any modeling process.

The process of identifying the key modeling assumptions further facilitates discussion among the stakeholders and between the stakeholders and the technical staff performing the analysis. Agreement must often be reached regarding which aspects of a problem are likely to impact the key decisions and are therefore worth modeling. Is there reason to believe that it is important to model patients being discharged to home separately from those being discharged to a nursing home or rehabilitation facility? Is the age of the patient an important determinant of the discharge time and, if so, should this be included in the model? If we are developing a model for surgical scheduling, is there agreement that uncertainty in the duration of surgical procedures may be contributing to delays in the starting times of subsequent surgeries and that modeling this uncertainty is worthwhile? Finally, by making the assumptions explicit, we are likely to be better able to discuss the results of the model and the recommendations that are derived from the modeling process. Both the technical analysts and the stakeholders will have a better appreciation for whether or not a particular result or recommendation is likely to have been caused by a modeling assumption or is the consequence of underlying properties of the process being analyzed.

In summary, in addition to the insights and action recommendations provided by the modeling process, there are likely to be significant ancillary benefits in the form of more reliable and useful data for future modeling, enhanced communication among stakeholders, and an improved understanding by all of the modeling being done.

1.7 OUTLINE OF THE BOOK

Thomas Davenport and Jeanne Harris argue that “analytics” is the new tool for competitive firms (Davenport and Harris, 2007). Analytics means data-driven decision making. They cite numerous examples of companies and organizations—from professional sports teams to pharmaceutical manufacturers, from banks to book sellers, from casinos to international insurance companies—that are getting ahead faster than their competitors by using quantitative data-based decision tools. They argue that firms that use analytic techniques routinely manage to out-think and out-perform their competitors in a systematic way.

Ian Ayers, in a related book, argues that our lives are increasingly being ruled by “super crunchers” (Ayers, 2007), individuals and organizations capable of managing and analyzing huge quantities of data to extract important trends. Like the Davenport and Harris’s firms that compete on analytics, Ayers’ super crunchers are found in a broad variety of industries, from medical diagnostics to

airline pricing, from studies of Supreme Court decisions to online matchmaking and dating services. Ayers argues that two methodologies—randomization and regression—are the keys to super crunching. Randomization basically means that treatments are applied randomly to elements of a population to see which treatment is most effective. Medical science has used randomization for many years in testing new drugs and procedures. Other industries are taking up the banner of randomization as well. Ayers reports, for example, that it is now possible for a firm like Amazon.com to determine in a matter of hours which of several different home page web designs generates the most viewing time and the most clicks by randomly showing the different designs to different customers over the course of the testing period. Because pages are shown randomly, preference differences that might exist based on age or gender or geographical location or ethnicity or any other individual characteristic are washed out due to the large sample that Amazon.com can amass in a short period of time. For samples that include data on individual characteristics, regression can be used to assess the impact of a treatment as a function of the individual characteristics. For example, in drug testing, individual characteristics that might impact the efficacy of a new drug include (1) how long a patient has had a disease, (2) the age of the patient, (3) the patient's gender, (4) the patient's family history, and (5) any other health conditions (e.g., obesity) that the patient may currently have or may have had in the past. Near the end of the text, Ayers introduces a third methodology, hypothesis testing, which can be used to determine whether or not the differences between two samples are statistically significant. For example, if, in a sample of 200 teenagers, we find that boys have more traffic accidents per mile driven than do their female counterparts, we can test whether the difference is statistically significant or if the difference could have been caused purely by chance.

Just as Ayers focuses on two methodologies, so too this text will focus primarily on two methodologies: optimization modeling and queueing theory. Optimization involves determining the values of key decision variables—those factors associated with a problem about which we need to make decisions—to minimize or maximize some objective. For example, in locating ambulances, we may want to determine the base locations that maximize the number of people who are within 8 minutes of the nearest ambulance base using a limited number of ambulances. Queueing models are typically closed-form mathematical equations that tell us how long a customer will have to wait for service. For example, we can use queueing theory to estimate the average time a person will spend in line and being served at a vehicle emission testing station if we know how many testing bays are operating, some characteristics of the rate at which cars arrive to be tested, as well as key attributes of the time required to test each vehicle. In some important cases, optimization and queueing will merge. In scheduling nurses in an emergency room, for example, we may want to examine two competing objectives: minimizing the cost of the nursing staff that we use and minimizing the total time that patients wait.

Chapter 2 reviews key concepts and issues in optimization modeling. The chapter outlines five key questions that need to be addressed in developing any

optimization model. It then reviews linear programming, one of the fundamental tools used in optimization. Network optimization, a special form of linear programming modeling, is then summarized, as is integer programming, another extension of linear programming, in which some variables must take on only integer values. Integer programming is particularly useful when yes/no decisions need to be made. Returning to the example of locating emergency medical service bases, there is likely to be a yes/no variable associated with each candidate location at which a base can be located. As indicated above, many problems are characterized by multiple objectives. Chapter 2 discusses two methods of finding efficient points on a tradeoff curve. Chapter 2 concludes with a list of 10 rules for formulating problems, rules that many students and practitioners seem to violate with disturbing regularity.

Chapter 3 introduces queueing theory. After introducing four key performance metrics of interest in almost any queueing problem and Little's law relating the average time in the system to the average number in the system, the chapter develops a general approach to analyzing a commonly studied class of queues. This general approach is then applied to a range of queueing problems. A more general model is introduced for the case of systems with a single server. Finally, since many service systems are characterized by time-varying conditions (arrivals of customers and service capabilities), the chapter concludes with a discussion of a numerical approach to solving queueing problems in the face of temporally changing conditions.

After reviewing these two key methodologies, the remainder of the text focuses on problems that arise frequently in the design and operation of services. Chapter 4 examines location models, which are typically used to assist in making strategic decisions regarding the number and location of service facilities. The chapter includes a discussion of districting problems that also arise in services.

Inventory theory is closely related to queueing theory and serves as the topic of Chapter 5. While services themselves can rarely be carried in inventory, most service providers require an inventory of some goods or equipment to aid in the service delivery process. As in other chapters, we begin with a simple model, the economic order quantity model, and progressively extend the model in a variety of ways. The chapter concludes with an analysis of the newsvendor problem.

Chapter 6 focuses on decisions related to the allocation of resources to customers or demands. For example, if we know where schools are located, we need to determine which students should attend each of the available schools. This can be thought of as a districting problem or as a resource allocation problem in which we are allocating the scarce resource of classroom space to the students in the district. Similarly, assigning students to seminars is also a resource allocation problem in which the resources are the courses allocated to the students. Assigning gates to aircraft at an airport is also a resource allocation problem. Many such problems can be solved using what is called an *assignment model*, perhaps with minor modifications.

Short-term workforce scheduling decisions and problems are the focus of Chapter 7. Many simple scheduling problems can be modeled as network problems. This allows us to solve large instances of such problems very quickly and easily. The basic model can be extended in important ways to include part-time workers, overtime shifts, and cross-trained employees. More complex problems, including cases in which employees must be scheduled over a 24-hour period or situations in which employees are permitted scheduled breaks, require more complex modeling techniques. The chapter concludes with a discussion of a multiple-objective employee scheduling problem in which employee costs and customer delays are traded off against each other.

Chapter 8 discusses long-term workforce management. This is an evolving area and models of this important topic are still in their nascent stage. The chapter begins with a very simple model and then evolves to increasingly complex problems including a model that accounts (approximately) for employee attrition, hiring and firing in response to anticipated workload requirements. The chapter concludes with a model that examines the tradeoff between employee costs and schedule delays for projects. Such a model could be of particular value in assessing the true cost of new projects in a consulting environment, for example.

Chapter 9 examines a variety of applications of queueing theory in the design and operation of service systems. Priority queueing problems arise in the operation of emergency response systems. Clearly, an elderly patient in active cardiac arrest will be given priority over a young child with a sore throat in an emergency room. Call center design is an important application of queueing theory in service systems. Finally, many services, including non-emergent health care systems, can schedule their customers (or patients). Chapter 9 concludes with a discussion of this problem.

Chapter 10 deals with vehicle routing decisions and problems. As in the case of long-term workforce management, vehicle routing models explicitly devoted to service problems are relatively few and far between. The chapter therefore focuses on traditional models and outlines ways in which such models might be extended to account for the unique issues that arise in the service industries.

Any text on a topic as broad as modeling the service industries can only begin to scratch the surface of the problems that we are likely to encounter. This is certainly true of this book as well. Chapter 11 outlines a number of additional directions for future modeling of the service industries that are not explicitly covered in this book.

1.8 PROBLEMS

1. List the services that you employed:
 - a) in the last 24 hours
 - b) in the last week
 - c) in the last month

- Be sure to include services that you implicitly used, such as insurance.
- Identify two services that you have used recently and briefly discuss ways in which each can be improved.

REFERENCES

- Ayers, I., 2007, *Super Crunchers: Why Thinking by Numbers Is the New Way to Be Smart*, Bantam Books, New York.
- Box, G. E. P., and N. R. Draper, 1987, *Empirical Model-Building and Response Surfaces*, Wiley, New York, NY.
- Cohon, J. L., 1978, *Multiobjective Programming and Planning*, Academic Press, New York, NY.
- Davenport, T. H., and J. G. Harris, 2007, *Competing on Analytics: The New Science of Winning*, Harvard Business Press, Boston.
- Davis, M. M., and J. Heineke, 2003, *Managing Services: Using Technology to Create Value*, McGraw Hill/Irwin, Boston, MA.
- de la Merced, M. J., 2008, "Starbucks announces it will close 600 stores," *The New York Times*, July 2, 2008, online at http://www.nytimes.com/2008/07/02/business/02sbux.html?_r=1&scp=1&sq=starbucks%20600%20stores&st=cse.
- Eaton, D. J. et al., 1979, *Location Techniques for Emergency Medical Service Vehicles: Volume I—An Analytical Framework for Austin, Texas*, Lyndon B. Johnson School of Public Affairs, The University of Texas at Austin, Austin, TX.
- Eaton, D., M. S. Daskin, D. Simmons, B. Bulloch, B., and G. Jansma, 1985, "Determining emergency medical service vehicle deployment in Austin, Texas," *Interfaces*, 15, 1, 96–108.
- Fitzsimmons, J. A., and M. J. Fitzsimmons, 2008, *Service Management: Operations, Strategy, Information Technology*, 6th ed., McGraw Hill/Irwin, Boston, MA.
- Fourer, R., D. M. Gay, and B. W. Kernighan, 2002, *AMPL: A Modeling Language for Mathematical Programming*, Duxbury Press/Brooks/Cole Publishing Company, South San Francisco, CA.
- Friedman, T. L., 2005, *The World Is Flat: A Brief History of the Twenty-first Century*, Farar, Straus and Giroux, New York, NY.
- Johnson, M. P., and Smilowitz, K., 2007, "Community-Based Operations Research," in *TutORials in Operations Research*, INFORMS, Institute for Operations Research and the Management Sciences, Chapter 6, 102–123.
- Lien, R., Iravani, S., and K. Smilowitz, 2007, *Sequential allocation problems for nonprofit agencies*, Technical report, Department of Industrial Engineering and Management Sciences, Northwestern University, Evanston, IL.
- Metters, R., K. King-Metters, and M. Pullman, 2003, *Successful Service Operations Management*, Thomson/South-Western, Mason, OH.
- Santibáñez, P., G. Bekiou, and K. Yip, 2009, "Fraser Health uses mathematical programming to plan its outpatient hospital network," *Interfaces*, 39, 3, 196–208.
- Wong, W., and B. Rose, 2008, "Starbucks' exit leaves some areas wondering," *The Chicago Tribune*, July 19, 2008, online at www.chicagotribune.com/business/chi-sat-starbucksjul19,0,7595492.story.