Chapter 1

INTRODUCTION

This chapter provides an introduction to the topic of revenue management (RM). We begin with an explanation of RM and its history and origins. We then provide a conceptual framework for understanding the objectives of RM, the types of business conditions under which it is applied, and the ways RM systems work. Finally we conclude by giving an outline of the remaining chapters of the book.

1.1 What Is "RM"?

Every seller of a product or service faces a number of fundamental decisions. A child selling lemonade outside her house has to decide on which day to have her sale, how much to ask for each cup, and when to drop the price (if at all) as the day rolls on. A homeowner selling a house must decide when to list it, what the asking price should be, which offer to accept, and when to lower the listing price—and by how much—if no offers come in. A stamp dealer selling on an Internet auction site has to select the duration of the auction, what reserve price to set (if any), and so on.

And anyone who has ever faced such decisions knows the uncertainty involved. You want to sell at a time when market conditions are most favorable, but who knows what the future might hold? You want the price to be right—not so high that you put off potential buyers and not so low that you lose out on potential profits. You would like to know how much buyers value your product, but more often than not you must just guess at this number.

Indeed, it is hard to find anyone who is entirely satisfied with their pricing and selling decisions. Even if you succeed in making a sale, you often wonder whether you should have waited for a better offer or whether you accepted a price that was too low.

Businesses face even more complex selling decisions. For example, how can a firm segment buyers by providing different conditions and terms of trade that profitably exploit their different buying behavior or willingness to pay? How can a firm design products to prevent cannibal–ization across segments and channels? Once it segments customers, what prices should it charge each segment? If the firm sells in different channels, should it use the same price in each channel? How should prices be adjusted over time based on seasonal factors and the observed demand to date for each product? If a product is in short supply, to which segments and channels should it allocate the products? How should a firm manage the pricing and allocation decisions for products that are complements (seats on two connecting airline flights) or substitutes (different car categories for rentals)?

RM is concerned with such *demand-management* decisions[1] and the methodology and systems required to make them. It involves managing the firm's "interface with the market" as it were—with the objective of *increasing revenues.* RM can be thought of as the complement of *supply-chain management* (SCM), which addresses the *supply decisions* and processes of a firm, with the objective (typically) of *lowering the cost* of production and delivery.

Other roughly synonymous names have been given to the practice over recent years—*yield management* (the traditional airline term), *pricing and revenue management, pricing and revenue optimization, revenue process optimization, demand management, demand-chain management* (favored by those who want to create a practice parallel to supply-chain management)—each with its own nuances of meaning and positioning. However, we use the more standard term revenue management to refer to the wide range of techniques, decisions, methods, processes, and technologies involved in demand management.

1.1.1 Demand-Management Decisions

RM addresses three basic categories of demand-management decisions:

[1]These can be referred to as either *sales decisions* (we are making decisions on where and when to sell and to whom and at what price) or *demand-management decisions* (we are estimating demand and its characteristics and using price and capacity control to "manage" demand). We use the latter consistently and use the shorter *demand management* whenever appropriate.

- Structural decisions: Which selling format to use (such as posted prices, negotiations or auctions); which segmentation or differentiation mechanisms to use (if any); which terms of trade to offer (including volume discounts and cancellation or refund options); how to bundle products; and so on.

- Price decisions: How to set posted prices, individual-offer prices, and reserve prices (in auctions); how to price across product categories; how to price over time; how to markdown (discount) over the product lifetime; and so on.

- Quantity decisions: Whether to accept or reject an offer to buy; how to allocate output or capacity to different segments, products or channels; when to withhold a product from the market and sale at later points in time; and so on.

Which of these decisions is most important in any given business depends on the context. The timescale of the decisions varies as well. Structural decisions about which mechanism to use for selling and how to segment and bundle products are normally strategic decisions taken relatively infrequently. Firms may also have to commit to certain price or quantity decisions, for example, by advertising prices in advance or deploying capacity in advance, which can limit their ability to adjust price or quantities on a tactical level. The ability to adjust quantities may also be a function of the technology of production—the flexibility of the supply process and the costs of reallocating capacity and inventory. For example, the use of capacity controls as a tactic in airlines stems largely from the fact that the different "products" an airline sells (different ticket types sold at different times and under different terms) are all supplied using the same, homogeneous seat capacity. This gives airlines tremendous quantity flexibility, so quantity control is a natural tactic in this industry. Retailers, in contrast, often commit to quantities (initial stocking decisions) but have more flexibility to adjust prices over time. The ability to price tactically, however, depends on how costly price changes are, which can vary depending on the channel of distribution such as online versus catalog.

Whether a firm uses quantity or price-based RM controls varies even across firms within a given industry. For instance, while most airlines commit to fixed prices and tactically allocate capacity, low-cost carriers tend to use price as the primary tactical variable.

Firms can also find innovative ways to increase their ability to make price or quantity recourse decisions. For example, retailers may hold back some stock in a centralized warehouse and then make a mid season replenishment decision rather than precommit all their stock to stores up front. Some major airlines have experimented with movable partitions that allow them to reallocate seats from coach to business cabins on a short-term basis. And other major airlines have recently experimented with a practice called *demand-driven dispatch* (D^3), in which aircraft of different sizes are dynamically assigned to each flight departure in response to fluctuations in demand, and are not precommitted to flights [50]. Car rental companies also may reallocate their fleet from one city to another. In terms of pricing, using online channels or advertising products without price ("call for our low price") provides firms with more price flexibility. All these innovations increase the opportunity for quantity and price-based RM.

Broadly speaking, RM addresses all three categories of demand-management decisions—structural, pricing, and quantity decisions. We qualify RM as being either *quantity-based RM* or *price-based RM* if it uses (inventory- or) capacity-allocation decisions or prices as the primary tactical tool respectively for managing demand. Both the theory and practice of RM differ depending on which control variable is used, and hence we use this dichotomy as necessary.

1.1.2 What's New About RM?

In one sense, RM is a very old idea. Every seller in human history has faced RM-type decisions. What price to ask? Which offers to accept? When to offer a lower price? And when to simply "pack up one's tent" as it were and try selling at a later point in time or in a different market. In terms of business practice, the problems of RM are as old as business itself.

In terms of theory, at a broad level the problems of RM are not new either. Indeed, the forces of supply and demand and the resulting process of price formation—the "invisible hand" of Adam Smith—lie at the heart of our current understanding of market economics. They are embodied in the concept of the "rational" (profit-maximizing) firm, and define the mechanisms by which market equilibria are reached. Modern economic theory addresses many advanced and subtle demand-management decisions, such as nonlinear pricing, bundling, segmentation, and optimizing in the presence of asymmetric information between buyers and sellers.

What *is* new about RM is not the demand-management decisions themselves but rather *how* these decisions are made. The true innovation of RM lies in the *method* of decision making—a technologically sophisticated, detailed, and intensely operational approach to making demand-management decisions.

This new approach is driven by two complementary forces. First, scientific advances in economics, statistics, and operations research now make it possible to model demand and economic conditions, quantify the uncertainties faced by decision makers, estimate and forecast market response, and compute optimal solutions to complex decision problems. Second, advances in information technology provide the capability to automate transactions, capture and store vast amounts of data, quickly execute complex algorithms, and then implement and manage highly detailed demand-management decisions. This combination of science and technology applied to age-old demand management is the hallmark of modern RM.

And both the science and technology used in RM are quite new. Much of the science used in RM today (demand models, forecasting methods, optimization algorithms) is less than fifty years old, most of the information technology (large databases, personal computers, Internet) is less than twenty years old, and most of the software technology (Java, object-oriented programming) is less than five years old. Prior to these scientific developments, it would have been unthinkable to accurately model real world phenomena and demand-management decisions. Without the information technology, it would be impossible to *operationalize* this science. These two capabilities *combined* make possible an entirely new approach to decision making—one that has profound consequences for demand management.

The first consequence is that science and technology now make it possible to manage demand on a *scale and complexity* that would be unthinkable through manual means (or would require a veritable army of analysts to achieve). A modern large airline, for example, can have thousands of flights a day and provide service between hundreds of thousands of origin-destination pairs, each of which is sold at dozens of prices—and this entire problem is replicated for hundreds of days into the future! A similar complexity is found at most large retail chains, which can have tens of thousand of SKUs[2] sold in hundreds of stores and over the Web with prices monitored and updated on a daily basis. The sheer scale and complexity of the decision-making task in these cases is beyond the ability of human decision makers. And if not automated, the task has to be so highly aggregated and simplified that significant opportunities for incremental gains—on particular products, at particular locations, at specific points in time—are simply lost.

[2] A *SKU* (stock-keeping unit) is the lowest level at which we identify inventory—such as men's Arrow blue Oxford shirts, long sleeves, size medium.

The second consequence of science and technology is that they make it possible to improve the *quality* of demand-management decisions. The management tasks that are involved—quantifying the risks and rewards in making demand-management decisions under uncertainty; working through the often subtle economics of pricing; accurately interpreting market conditions and trends and reacting to this information with timely, accurate, and consistent real-time decisions; optimizing a complex objective function subject to many constraints and business rules—are tasks most humans, even with many years of experience, are simply not good at. Models and systems are better at separating market signals from market noise, evaluating complex tradeoffs, and optimizing and producing consistent decisions. The application of science and technology to demand decisions often produces an improvement in the quality of the decisions, resulting in a significant increase in revenues.

Of course, even with the best science and technology, there will always be decisions that are better left to human decision makers. Models can detect only what's in the data. They cannot reason through the consequences of a demand shock, new technologies, a sudden shift in consumer preferences, or the surprise price war of a competitor. These higher-level analyses are best left to experienced, human analysts. Most RM systems recognize this fact and parse the decision-making task, with models and systems handling routine demand-management decisions on an automated basis and human analysts overseeing these decisions and intervening (based on flags or alerts from the system) when extraordinary conditions arise. Such man-machine interaction offers a firm the best of both human and automated decision making.

The process of managing demand decisions with science and technology—implemented with disciplined processes and systems, and overseen by human analysts (a sort of "industrialization" of the entire demand-management process)—defines modern RM.

1.2 The Origins of RM

Where did RM come from? In short, the airline industry. There are few business practices whose origins are so intimately connected to a single industry. Here we briefly review the history of airline RM and then discuss the implications of this history for the field.

1.2.1 Airline History

The starting point for RM was the Airline Deregulation Act of 1978. With this act, the U.S. Civil Aviation Board (CAB) loosened control of airline prices, which had been strictly regulated based on standardized price and profitability targets. Passage of the act led to rapid change and a rash of innovation in the industry. Established carriers were now free to change prices, schedules, and service without CAB approval. Large airlines accelerated their development of computerized reservation systems (CRSs) and global distribution systems (GDSs), and the CDS business became profitable in its own right. The majors developed hub-and-spoke networks, which allowed them to offer service in many more markets than was possible with point-to-point service but also made pricing and operations more complex.

At the same time, new low-cost and charter airlines entered the market. Many of these upstarts—because of their lower labor costs, simpler (point-to-point) operations, and no-frills service—were able to profitably price much lower than the major airlines. These new entrants tapped into an entirely new and vast market for discretionary travel—families on a holiday, couples getting away for the weekend, college students visiting home—many of whom might otherwise have driven their cars or not traveled at all. It turned out (quite surprisingly to some at the time) that air travel was quite price elastic; with prices sufficiently low, people switched from driving to flying, and demand from this segment surged.

The potential of this market was embodied in the rapid rise of PeopleExpress, which started in 1981 with cost-efficient operations and fares 50 to 70% lower than the major carriers. By 1984, its revenues were approaching $1 billion, and for the year 1984 it posted a profit of $60 million, its highest profit ever (Cross [137]).

While these developments resulted in a significant migration of price-sensitive discretionary travelers to the new, low-cost carriers, the major airlines had strengths that these new entrants lacked. They offered more frequent schedules, service to more city pairs and an established brand name and reputation. For many business travelers, schedule convenience and service was (and still is) more important than price, so the threat posed by low-cost airlines was less acute in the business-traveler segment of the market. Nevertheless, the cumulative losses in revenue from the shift in traffic were badly damaging the profits of major airlines.

A strategy to recapture the leisure passenger was needed. Yet, for the majors, a head-to-head, across-the-board price war with the upstarts was deemed almost suicidal; with their much lower costs, airlines like PeopleExpress could earn a profit at the new low prices, while most majors would lose money at a staggering rate.

Robert Crandall, American Airline's vice president of marketing at the time, is widely credited with the breakthrough in solving this problem. He recognized that his airline was already producing seats at a marginal cost near zero because most of the costs of a flight (capital costs, wages, fuel) are fixed. As a result, American could in fact afford to "compete on cost" with the upstarts using its surplus seats.

However, two problems had to be solved to execute this strategy. First, American had to have some way of identifying the "surplus" seats on each flight. The scheme would not be profitable if a sale of a low-price seats displaced high-paying business customers.[3] Second, they had to ensure that American's business customers did not switch and buy the new low-price products it offered to discretionary, leisure customers.

American solved these problems using a combination of *purchase restrictions* and *capacity-controlled fares.* First, they designed discounts that had significant restrictions for purchase: they had to be purchased 30 days in advance of departure, were nonrefundable, and required a seven day minimum stay. These restrictions were designed to prevent most business travelers from utilizing the new low fares. At the same time, American limited the number of discount seats sold on each flight: they *capacity-controlled* the fares. This combination provided the means to compete on price with the upstart airlines without damaging their core business-traveler revenues.

The new pricing scheme was launched in 1978 as American Super-Saver Fares. The fares were quite successful at stemming the tide of defections of discretionary travelers to the low-cost airlines.

Despite this initial success, American experienced some significant problems implementing its new strategy. Initially, American's capacity controls were based on setting aside a fixed portion of seats on each flight for the new low-fare products. But as American gained experience with its Super-Saver fares, it realized that not all flights were the same. Flights on different days and at different times had very different patterns of demand. Some had many excess seats and could profitably support a higher allocation of discount seats; others had sufficient demand for regular-priced seats and warranted very little if any allocation to the new, discounted products.

American realized that a more intelligent approach was needed to realize the full potential of capacity-controlled discounts. It therefore embarked on the development of what became known as the *Dynamic Inventory Allocation and Maintenance Optimizer* system (DINAMO). These efforts on DINAMO represent, in many ways, the first large-scale RM system development in the industry. (Though on a more modest scale, the capacity-control problem dates back to the mid-1970s, and other airlines and the Boeing Aircraft Company were experimenting with

[3]As we show in the chapters that follow, a notion of this sort of *displacement cost* is central to the theory of RM.

similar ideas at the time.) The DINAMO system was large and complex and took several years to develop and refine.

DINAMO was implemented in full in January 1985 along with a new fare program entitled Ultimate Super-Saver Fares, which matched or undercut the lowest discount fares available in every market American served.

DINAMO made all this possible. American could now be much more aggressive on price. It could announce low fares that spanned a large swath of individual flights, confident in its capability to accurately capacity-control the discounts on each individual departure. If a rival airline advertised a special fare in one of American's markets, American could immediately match the offer across the board, knowing that the DINAMO system would carefully control the availability of this fare on the thousands of departures affected by the price change. Moreover, the competition could not observe American's capacity controls unlike prices themselves, which, thanks to GDSs, instantly became public information. This feature of pricing aggressively and competitively at an aggregate, market level, while controlling capacity at a tactical, individual-departure level still characterizes the practice of RM in the airline industry today.

The effect of this new capability was dramatic. PeopleExpress was especially hard hit as American repeatedly matched or beat their prices in every market it served. PeopleExpress's annual profit fell from an all-time high in 1984 (the year prior to implementation of DINAMO) to a loss of over $160 million by 1986 (one year after DINAMO was implemented). It soon went bankrupt as a result of mounting losses, and in September 1986 the company was sold to Continental Airlines.

Donald Burr, CEO of PeopleExpress, summarized the reasons behind the company's failure [137]:

> We were a vibrant, profitable company from 1981 to 1985, and then we tipped right over into losing $50 million a month. We were still the same company. What changed was American's ability to do widespread Yield Management in every one of our markets. We had been profitable from the day we started until American came at us with Ultimate Super Savers. That was the end of our run because they were able to under-price us at will and surreptitiously.
>
> Obviously PeopleExpress failed … We did a lot of things right. But we didn't get our hands around Yield Management and automation issues. … [If I were to do it again,] the number one priority on my list every day would be to see that my people got the best information technology tools. In my view, that's what drives airline revenues today more than any other factor—more than service, more than planes, more than routes."

This story was played out in similar fashion throughout the airline industry in the decades following deregulation. And airlines that did not have similar RM capabilities scrambled to get them.

As a result of this history, the practice of RM in the airline industry today is both pervasive and mature, and RM is viewed as critical to running a modern airline profitably. For example, American Airlines' estimates that its RM practices generated $1.4 billion in additional incremental revenue over a three-year period starting around 1988 [477]. Many other carriers also attribute similar improvements in their revenue due to RM.

1.2.2 Consequences of the Airline History

The intimate connection of RM to the airline industry is both a blessing and a curse for the field of RM. The blessing is that RM can point to a major industry in which the practice of RM is pervasive, highly developed, and enormously effective. Indeed, a large, modern airline today would just not be able to operate profitably *without* RM. By most estimates, the revenue gains from the use of RM systems are roughly comparable to many airlines' total profitability in a good year (about 4 to 5% of revenues).[4] And the scale and complexity of RM at major airlines is truly mind-boggling. Therefore, the airline success story validates both the economic importance of RM and the feasibility of executing it reliably in a complex business environment. This is the good-news story for the field from the airline experience.

The bad news—the curse if you will—of the strong association of RM with airlines is that it has created a certain myopia inside the field. Many practitioners and researchers view RM solely in airline-specific terms, and this has at times tended to create biases that have hampered both research and implementation efforts in other industries.

A second problem with the airline-specific association of RM is that airline pricing has something of a bad reputation among consumers. While on the one hand customers love the very low fares made possible by RM practices, the fact that fares are complex, are available one minute and gone the next, and can be drastically different for two people sitting side by side on the same flight, has led to a certain hostility toward the way airlines price. As a result, managers outside the industry are at times, quite naturally, somewhat reluctant to try RM practices for fear of engendering a similar hostile reaction among their customers.

[4] Many skeptics point to Southwest Airlines as a counterexample, but Southwest does use RM systems. However, because its pricing structure is simpler than most other airlines the use of RM is less obvious to consumers and casual observers.

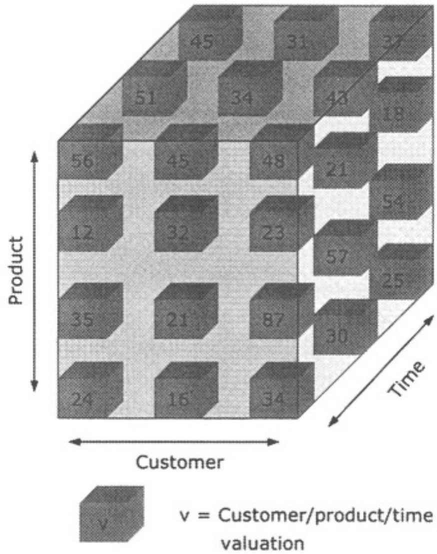

Figure 1.1. Conceptual view of a firm's demand landscape

Yet the reality is that, in most cases, applying RM does *not* involve radically changing the structure of pricing and sales practices; rather, it is a matter of making more intelligent decisions.

1.3 A Conceptual Framework for RM

So if airlinelike conditions aren't strictly necessary for RM, then exactly where *does* it apply? A short answer is: in any business where tactical demand management is important and the technology and management culture exists to implement it. But this in turn begs the question: when do *these* conditions arise? To answer this question, it helps to begin with a conceptual framework for thinking about the demand management process.

1.3.1 The Multidimensional Nature of Demand

A firm's demand has multiple dimensions, including (1) the different products it sells, (2) the types of customers it serves, their preferences for products, and their purchase behaviors, and (3) time. Other dimensions (such as, locations or channels) also affect the nature of a firm's demand, but these three dimensions—products, customers, and time—suffice to illustrate the idea.

Figure 1.1 shows these three demand dimensions. A single cell in the figure indicates a particular customer's valuation for a particular product at a particular point in time. RM addresses the structural, price, timing and quantity decisions a firm makes in trying to exploit the potential of this multidimensional demand landscape.

For example, some RM problems look at exploiting heterogeneity in valuations among customers for a single product at a single point in time: they fix the product and time dimension and try to optimize over the customer dimension. This problem is characteristic of the classical auction-design problems discussed in Chapter 6 and classical price-discrimination problems discussed in Chapter 8. Other RM problems look at dynamically pricing a single product to heterogeneous customers over time: they fix the product dimension and optimize over the customer and time dimensions. Such problems are addressed in Chapter 5. Others, such as the network problems in Chapter 3, address managing demand decisions for multiple products over multiple time periods, and the customer-behavior dimension is not explicitly considered. Of course, all three dimensions are important factors in practice. However, methodologically one often has to decompose and simplify the problem to develop implementable solutions.

1.3.2 Linkages Among Demand-Management Decisions

If the decisions affecting the demand landscape in Figure 1.1 were independent, then the decision-making problem would be considerably simpler. However, typically one or more of the following three factors link the demand across these dimensions.

First, multiple products may share production capacity or have joint production costs. In such cases, the demand-management decision for different products or for a given product in different periods of time are interrelated. For example, because of joint capacity constraints, accepting demand from a customer for a particular product at a specific point in time may mean giving up opportunities to accept demand at later points in time, or because lowering the price of one product increases its demand, this may reduce the capacity available for producing other products.

Second, even if production constraints do not link demand decisions, customer behavior often does. Customers may choose among substitute products at any given point in time, or customers may strategize over their timing in purchasing a given product. As a result, the price or quantity decisions that a firm makes about one product may affect demand for related products—or may affect the future demand for the same product.

Finally, demand decisions for different products, customers, and time-periods may also be linked in terms of the information the firm gains. The most common link is over time; observed demand to date may reveal information about future demand. Thus, a decision about price today may affect the information we gain about demand sensitivity, which will affect future pricing decisions. Also, a firm selling the same product in geographically separated markets or in different channels may gain information in one location or channel as a result of observing demand that impacts its decisions in other locations and channels. Or the observed purchase decisions of a given customer may reveal information about that customer's future purchase decisions. Such linkages complicate demand-management decisions, and managing the often subtle tradeoffs they create is a key motivation for RM.

1.3.3 Business Conditions Conducive to RM

Given this conceptualization of the demand-management problem, one can begin to gain insights into conditions in which RM is likely to be beneficial. Here, we discuss a few such conditions.

1.3.3.1 Customer Heterogeneity

If all customers value a product identically and exhibit similar purchase behavior, then the customer dimension of Figure 1.1 is essentially lost. As a result, there is less potential to exploit variations in willingness to pay, variations in preference for different products, and variations of purchase behavior over time. Therefore, the more heterogeneity in customers, the more potential there is to exploit this heterogeneity strategically and tactically to improve revenues.

Customers in the airline and hotel industries certainly exhibit this characteristic. They have widely varying patterns of usage and behavior in terms of when they purchase and how flexible their plans are, and they place very different valuations on the need to travel.

1.3.3.2 Demand Variability and Uncertainty

The more demand varies over time (due to seasonalities, shocks and so on) and the more uncertainty one has about future demand (the more variance there is along the time dimension in Figure 1.1) the more difficult the demand-management decisions become. Hence, the potential to make bad decisions rises, and it becomes important to have sophisticated tools to evaluate the resulting complex tradeoffs.

Consider the demand for air travel. It exhibits significant variations (by season, time of day, day of week, holidays) and even correcting for this predictable seasonal variation is highly uncertain for a given flight.

1.3.3.3 Production Inflexibility

As mentioned, joint production constraints and costs complicate the demand-management problem. If a firm can "absorb" variations in *demand* easily and costlessly through variations in *supply,* then the complexity of managing demand diminishes; you just supply enough to meet demand. However, the more inflexible the production—the more delays involved in producing units, the more fixed costs or economies of scale involved in production, the more the switch-over costs, the more capacity constraints—the more difficult or costly it becomes to match demand variations with supply variations. As a result, inflexibility leads to more interaction in the demand management at different points in time, between different segments of customers, across different products of a product line, and across different channels of distribution (the different cells in Figure 1.1). The complexity increases and the consequences of poor decisions become more acute. Hence, RM becomes more beneficial.

Again, the airline industry is one in which production is very inflexible. Essentially, when committing to fly a flight from A to B, an airline both *fixes* the level of its output (the number of seats) and, for all practical purposes, the total cost of that output—independent of how many customer actually fly on the flight. Its unit cost per seat sold, therefore, varies tremendously with the volume of sales, and once the capacity constraint is reached, no more production is possible. Worse yet, like all services, output cannot be inventoried, so production of air transport output in one period cannot be used to satisfy demand in later periods (an unsold seat on Monday cannot be used to supply the need of an excess passenger on Tuesday). All these factors combine to create extreme inflexibility in the technology of air transport service, and this is one of the key driving factors in the importance of RM in this industry.

1.3.3.4 Price as a Signal of Quality

The extent to which price is a signal of quality is also a factor. For example, people buy a $10,000 Patek Philippe watch partly for its aesthetics and functionality but also, to a large extent, because they want the exclusivity of a $10,000 watch. The price is a key feature of the watch, as it is with most luxury goods. They are status symbols, and to lower or manipulate the price risks damaging this status.

A more subtle case is observed in situations where it is hard to assess quality through other, objective means. For example, the hourly rate of a prominent attorney or consultant, the tuition at an Ivy League university, and the price of a bottle of wine on a dinner menu—all play important roles as signals of quality to consumers. Again, tampering with prices for tactical reasons in such settings jeopardizes the signaling value of prices. Therefore, RM is more suited to products where price is *not* a status symbol and *not* a significant signal of value—where price and quality are decoupled in the consumer's mind.

Airlines are arguably a good example. While different airlines position themselves differently with respect to price and quality (e.g., no-frills discount carriers and full-service, mainstream carriers), consumers generally do not associate the price of an airline ticket with the quality of the particular flight. We do not expect a "nicer" flight when paying $300 more because we booked our ticket at the last minute or because we booked our flight on a holiday weekend as opposed to a normal weekday. Moreover—despite what some airline marketers might like to believe—most consumer do not have strong quality preferences among airlines, at least not sufficient to outweigh even relatively small differences in schedule and price. This is one of the main motivations behind the introduction of loyalty schemes in the industry, which are really an attempt to "synthesize" a high level of brand loyalty among a group of consumers who innately have very little of it.

1.3.3.5 Data and Information Systems Infrastructure

To operationalize RM requires data to accurately characterize and model demand. It also requires systems to collect and store the data and to implement and monitor the resulting real-time decisions. In most industries it is usually feasible—in theory, at least—to collect and store demand data and automate demand decisions. However, attempting to apply RM in industries that do not have databases or transactions systems in place can be a time-consuming, expensive, and risky proposition. RM, therefore, tends to be more suited to industries where and transaction-processing systems are already employed as part of incumbent business processes.

Again, the airline industry is a perfect case in point. It is an industry whose pricing and distribution processes were largely automated with the introduction of GDSs in the 1960's and 1970's. In fact, it is one of the earliest industries to move almost entirely to electronic selling and distribution—decades before the advent of e-commerce. This long history of using information systems to automate business processes meant that it was quite natural to implement RM in the airline industry when the time came.

1.3.3.6 Management Culture

RM is a technically complex and demanding practice. There is a risk, therefore, that a firm's management may simply not have sufficient familiarity with—or confidence in—science and technology to make implementing a RM system a realistic prospect. The culture of the firm may not be receptive to innovation or may value more intuitive approaches to problem solving. This is often due to the culture of the industry and its managers: their educational backgrounds, their professional experiences and responsibilities en route to leadership positions, and the skills required to succeed in the industry.

Again, the airline industry serves as a good example. Modern airlines cannot run without information systems: systems for ticketing and reservations, scheduling crews and aircraft, handling baggage, planning meals and operational control (rerouting aircraft because of delays and breakdowns, and so on). Also, airline managers are accustomed to applying scientific methods in managing these various operations. In fact, long before RM was practiced in the industry, most large airlines had staffs of operations researchers working on complex problems of scheduling and fleet assignment. When RM came along, the management and culture in the industry were therefore well conditioned to accept it.

1.3.4 Industry Adopters Beyond the Airlines

What do these conditions imply for adopters of RM technology? Chapter 10 reviews specific industry adopters in detail, so here we only briefly mention some of them.

The production-inflexibility characteristics of airlines are shared by many other service industries, such as hotels, cruise ship lines, car rental companies, theaters and sporting venues, and radio/TV broadcasters, to name a few. Indeed, RM is strongly associated with service industries.

Retailers have recently begun to adopt RM, especially in the fashion apparel, consumer electronics, and toy sectors. Retail demand is highly volatile and uncertain, consumers' valuations change rapidly over time, and with short selling seasons and long production and distribution lead times, supply is quite inflexible. On the technology front, the introduction of bar codes and point-of-sale (POS) technology has resulted in a high degree of automation of sales transactions for most major retailers.

The energy sector has been a recent adopter of RM methods as well, principally in the area of managing the sale of pipeline capacity for gas transportation. Again, energy demands are volatile and uncertain, and the technology for generating and transmitting electricity and gas can be inflexible. Also, thanks to deregulation in the industry, there has been a lot of experimentation and innovation in the pricing practices of energy, gas, and transmission markets.

Manufacturing is potentially a vast market for RM methods, though to date relatively few instance of the practice have been documented. To a large extent this is due to the fact that supply is more flexible, and, for durable goods, customers have more flexibility in their purchase timing. This somewhat diminishes the impact of RM and creates unique challenges for the methodology as well. Still, there is immense interest in RM in manufacturing. Enterprise resource planning (ERP), supply-chain management (SCM), and customer-relationship management (CRM) systems are commonplace in the industry, and most manufacturers have huge amounts of data and heavily automated business processes, which could form the foundations for RM. Indeed, in the auto industry Ford Motor Corporation recently completed a high-profile implementation of RM technology [135].

What about future adopters of RM? Given the criteria outlined above, one can argue that many industries are potential candidates. Almost all businesses must deal with demand variability, uncertainty, and customer heterogeneity. Most are subject to some sort of supply or production inflexibility. Finally, thanks largely to the wave of enterprise software and e-commerce innovation of late, many firms have now automated their business processes. All of these factors bode well for the future of RM.

Nevertheless, as with any technological and business-practice innovation, the case for RM ultimately boils down to a cost-benefit analysis for each individual firm. For some, the potential benefit will simply never justify the costs of implementing RM systems and business processes. However, we believe that for the majority of firms, RM will eventually be justified once the technology and methodology in their industry matures. Indeed, the history of RM in industries such as airlines, hotels, and retail suggests that once the technology gains a foothold in an industry, it spreads quite rapidly. As a result, we would not be surprised to see RM systems (or systems performing RM functions under a different label) become as ubiquitous as ERP, SCM, and CRM systems are today.

1.4 An Overview of a RM System

Here, we give a brief description of the generic operations of a RM system. This introduces the key components and gives an overview of the information flows, controls, and design of a RM system. The details of the science and systems involved in each component are covered in later chapters.

RM generally follows four steps:

1. Data collection: Collect and store relevant historical data (prices, demand, causal factors).
2. Estimation and forecasting: Estimate the parameters of the demand model; forecast demand based on these parameters; forecast other relevant quantities like no-show and cancellation rates, based on transaction data.
3. Optimization: Find the optimal set of controls (allocations, prices, markdowns, discounts, overbooking limits) to apply until the next re-optimization.
4. Control: Control the sale of inventory using the optimized control. This is done either through the firm's own transaction-processing systems or through shared distribution systems (such as GDSs).

The RM process typically involves cycling through these steps at repeated intervals. The frequency with which each step is performed is a function of many factors such as the volume of data, the speed that business conditions change, the type of forecasting and optimization methods used, and the relative importance of the resulting decisions. For example, most RM systems in airline and hotel applications stagger the dates—data collection points (DCPs)—when they collect data, reforecast, and reoptimize, with the cycle occurring more frequently (at least daily) as the service time nears. This is because in these industries, a substantial portion of the reservations occurs during the last few days before the time of service.

Figure 1.2 shows the process flow in a RM system. Data is fed to the forecaster; the forecasts become input to the control optimizer; and finally the controls are uploaded to the transaction-processing system, which controls actual sales.

1.5 The State of the RM Profession

On the practice side, the profession can be divided into *users* (the firms and individuals who use RM methods to manage their business) and *vendors* (the firms and individuals who develop and supply technology and consulting services to users). Of course, this division is not always sharp. Many users of RM, especially in the airline industry, have research and development organizations that provide significant components of their firm's RM technology. Still, most users—even those with their own RM staff—rely on vendors in part or whole for their technology. Often, the role of a user's R&D staff is to serve as in-house technology advisers and consultants, helping senior management evaluate new technologies and manage the relationships with the firm's technology vendors.

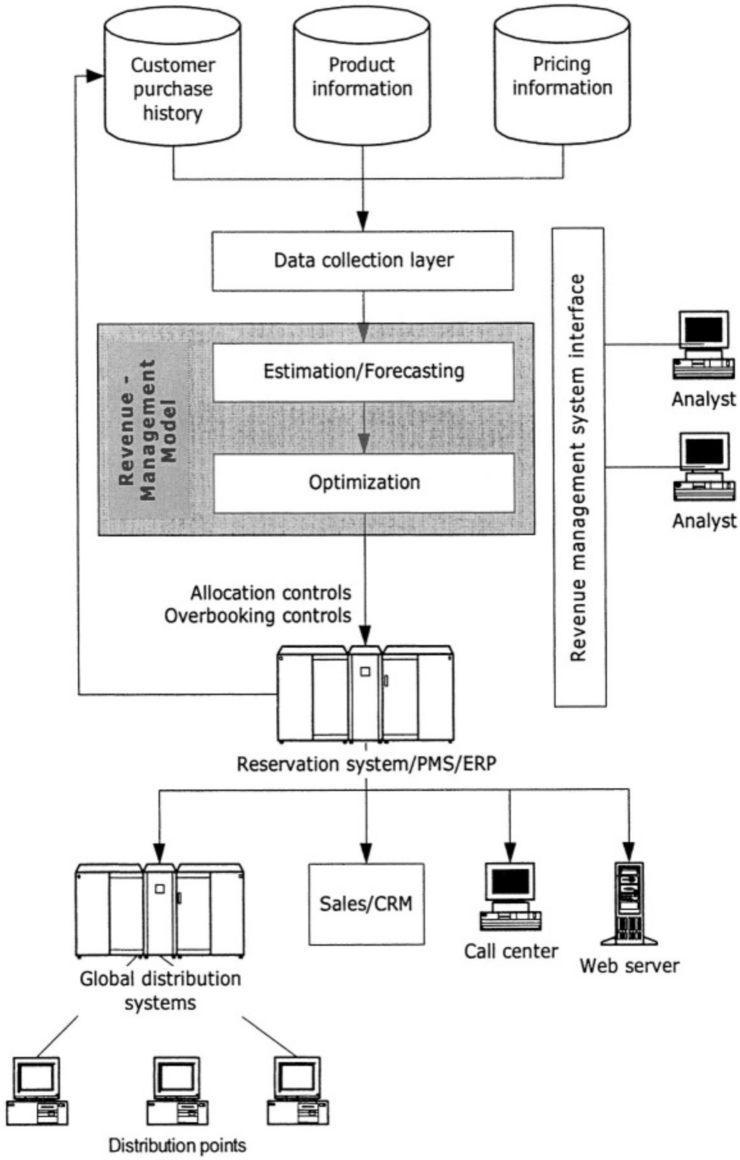

Figure 1.2. RM process flow.

In most user organizations, the vast majority of RM staff are involved in day-to-day RM operational activities: training and supervision of field staff, managing pricing and capacity controls in individual markets, overseeing automated decisions and intervening where necessary or maintaining computer systems. The typical senior management titles in such organizations are *VP of RM* or *VP of pricing and inventory control.* The organization typically has a corporate staff that is responsible for overall RM strategy, policy and systems and line management and staff responsible for RM processes in specific business units and markets.

Except for a few large airlines that develop their own systems, RM software is developed by a handful of RM vendors, many of whom specialize in a particular industry. Most of these firms have both a scientific staff to develop models and algorithms (operations researchers, marketing scientists, statisticians, economists), an IT and software staff to develop the associated software and systems, and a consulting staff to provide training and implementation services. The resulting products are usually customized for each user's particular business conditions. Vendors also provide training and consulting in the use of the systems. To this list of RM-specific vendors, one ought to add the major enterprise software and technology companies and general IT consulting and software companies that, though not specifically identified as RM vendors, nevertheless provide some RM products and services.

1.6 Chapter Organization and Reading Guide

We next describe the organization of the book and then provide our suggestions for how to approach the material.

1.6.1 Chapter Organization

The book is divided into three main parts. Part I addresses *quantity-based RM,* in which the primary demand-management decisions concern product rationing and availability control—how much to sell to whom, whether to accept or reject requests for products, and so on. These are the core set of problems behind traditional airline RM and closely related industries like hotels and rental car industries. Part I is comprised of these chapters:

- Chapter 2, Single-Resource Capacity Control: This chapter looks at capacity controls for a single resource (seats on a single flight, hotel rooms on a single night) that is sold to differentiated demand classes—the so-called single-leg problem in airline RM parlance. We provide a comprehensive treatment of the classic exact and heuristic approaches to this problem, as well as a number of more recent advances.
- Chapter 3, Network Capacity Control: This chapter looks at the same capacity-control decisions, but in a setting in which products require multiple resources—called the *network problem.* The main

motivation is controlling availability of discount classes at an origin-destination (O&D) level in an airline network. However, hotels face a similar network problem when they control capacity by length of stay. Because the network capacity-control problem is significantly more complex than the single-resource problem, most of the methods in this chapter are based on approximations.

- Chapter 4, Overbooking: This chapter looks at the practice of overbooking—accepting more reservations than physical capacity as a hedge against cancellations. The topic is somewhat specialized to reservation-based industries like airlines, hotels, and car rentals. While in a sense overbooking is a demand-management decision, it is somewhat different from the pricing and allocation decisions of the other chapters. However, overbooking is intimately connected to RM in the airline and hotel industries and is almost always implemented in conjunction with RM capacity controls. It is also extremely important economically in these industries and forms a significant and visible part of RM.

Part II of the book examines *price-based RM,* in which the primary demand decisions are prices—how to price to various customer groups or how to vary prices over time. Both posted price and auction mechanisms are considered. These price-based RM problems are more typical of retail and manufacturing RM. Part II has two chapters:

- Chapter 5, Dynamic Pricing: In this chapter we look at a problem in which the principle demand decision is how to adjust prices over time, subject to demand variability and uncertainty and various constraints or costs on re-supply. Many of the retail RM systems are based on the types of models discussed in this chapter.

- Chapter 6, Auctions: Auctions are an important and long-standing pricing mechanism in many industries and, with the rise of e-commerce, have gained popularity as a alternatives to posted pricing. The basic types of auctions are discussed along with the theory of optimal auction design. We discuss the implications of this theory for dynamic pricing in general and look at classical auctions, dynamic auctions and network auctions.

Finally, the five chapters in Part III of the book examine components of RM that are common to both quantity and price-based RM:

- Chapter 7, Customer Behavior and Market-Response Models: This chapter summarizes the core demand-modeling theory and methodology underlying RM. We discuss the basic theory of consumer behavior and develop several of the demand models used in both quantity and price-based RM. Both individual customer choice and aggregate market-demand models are covered.

- Chapter 8, The Economics of RM: Here we discuss the economic theory of RM. We briefly survey classical monopoly and oligopoly pricing theory as well as the theory of price discrimination, peak-load pricing, and pricing under demand uncertainty, all of which are particularly relevant to understanding the strategies and tactics used in RM practice.

- Chapter 9, Estimation and Forecasting: This chapter addresses the broad range of issues involved in estimating models from data and building forecasts of future demand. We survey the main estimation and forecasting methods commonly used in practice. The coverage is not intended to be as in-depth as specialized books on these topics but rather to review the basic assumptions and theory of each method and its role in RM practice.

- Chapter 10, Industry Profiles: This chapter provides detailed descriptions of several industries practicing RM, including information on consumers, products, sales practice and technology—all of which impact the real world practice of RM. For experienced industry insiders, much of this material may be well-known. However, for new employees in an industry, for academics, and for industry practitioners looking at a different industry, the chapter provides useful information on the institutional context in which RM is practiced.

- Chapter 11, Implementation: This chapter discusses issues involved in implementing a RM system, including product design, organizational and technology-management issues, all factors critical in making a RM system effective in application.

1.6.2 Reading Guide

Some readers will not want to read the book in strict sequential order. It is certainly possible to read Parts I and II independently of one another. Readers who are interested primarily in traditional quantity-based RM should begin with Part I, while those interested primarily in price-based RM problems could begin with Part II and then look at Part I afterwards. However, within Parts I and II chapters are interrelated, with later chapters building on ideas developed in earlier chapters.

Each chapter provides a comprehensive introduction as well, so readers may wish to begin by looking through each of the chapter introductions to get a sense of the scope of each one and then read individual chapters in detail according to their level of interest.

Parts I and II can also be read largely independently of Part III, though the material in Part III provides useful background. While some readers may choose to use Part III only as a reference, in our view each chapter in Part III is also of significant independent interest. Readers interested in the theory underlying RM will find Chapter 7 on demand modeling and Chapter 8 on economics of particular interest. Those interested primarily in the applied elements of RM will find Chapter 9 on forecasting methods and Chapters 10 and 11 on industry profiles and implementation (respectively) most useful.

The chapters in Part III are not strongly interrelated and may be read in any order. However, the material in Part III is best understood in the context of the topics covered in Parts I and II; hence, we recommend at least skimming the introductions of chapters in Parts I and II before reading Part III in detail.

1.7 Notes and Sources

The 1997 book by Robert Cross, *RM: Hard Core Tactics for Market Domination* [137] was influential in popularizing the story of airline RM and introducing the concept of RM to the general business community. Bob Cross was then chairman and CEO of Aeronomics, a RM consultancy and software firm. It is a nontechnical and lively book for a general audience , and is informative reading, providing nice descriptions of the early history of RM in the airline industry, many practical anecdotes, and insights into the philosophy and challenges of implementing RM. Several other books on RM have been published recently. One is an edited volume by Ingold, McMahon-Beattie, and Yeoman [263] that focuses primarily on the hotel industry. Another, Daudel and Vialle [146], focuses on air transportation. Both, however, deal more with practical and conceptual issues and do not cover the scientific methods of RM in much depth. The book by Nagel and Holden [400] provides a comprehensive overview of many managerial issues involved in pricing and is useful reading. However, it does not address tactical RM decision making in depth.

Several survey articles provide general coverage of RM. The *Handbook of Airline Economics* edited by Jenkins [268] provides several good practice-oriented articles on RM in the airline industry. Kimes [301] provides a conceptual introduction to RM with a hotel RM focus. Smith et al. [477] provide a nice description of the practice of RM at American Airlines and the DINAMO system.

As for guides to the research literature, Weatherford and Bodily [556] propose a taxonomy for classifying the sets of assumptions used in many traditional RM models, although the taxonomy itself is little used. McGill and van Ryzin [374] provide a comprehensive overview and annotated bibliography of the published academic literature in the field through 1998. Elmaghraby and Keskinocak [177] provide a survey on research in the area of dynamic pricing.