PRICE AND REVENUE MANAGEMENT

I-MBA 8TH SEMESTER

COURSE OBJECTIVES

- To offer fundamental understandings of pricing and revenue management with respect to operations management
- To analyze the impact of different types of pricing and economy on revenue management
- To provide Network and Capacity Control ideas in Revenue Management
- To practice the students by apply pricing and revenue management in various fields

Module – I: Introduction to Revenue Management [RM] :

Prices & Revenue Management Concept, Application in Air Lines, Railways, Hospitality Industries, Steps Involved in setting the price, Types of Pricing, Dynamic Pricing – Introduction and overview, Single Product Dynamic Pricing with and without Replenishment, Multi Product and Multi Source Pricing, Finite Population Models and Price Skimming, Cost revenue trade off, Relationship to List Pricing

Module – II : Economy of Revenue Management :

Introduction, Perfect Competition, Perfectly Competitive Markets, Firm level decision under perfect competition, Pre–commitment and Demand Uncertainty, Peak–load pricing under perfect competition, Identifiable peak periods competition, Monopoly Pricing, Price and capacity competition in Oligopoly and monopolistic market.

Module – III : Network and Capacity Control in Revenue Management

Promise and Challenge of Network Control, Types of Controls, Theory of Optimal Network Control–Structure of Optimal Control, Bid Price Control, Non Optimality of Bid Price Control, Evidence in support of Bid Price, Bid Prices and Opportunity Cost, Approximations based on network models– Deterministic Linear Programming, Simulation method for price revenue trade off. Pricing and Revenue Management Practice applied to Airlines, Hotels and hospitals, Revenue Opportunity Assessment and Revenue Benefits Measurement

PRICING AND REVENUE MANAGEMENT
Module-I

Revenue management is the use of pricing to increase the profit generated from a limited supply of supply chain assets

– SCs are about matching demand and capacity

– Prices affect demands

- Yield management similar to RM but deals more with quantities rather than prices
- Supply assets exist in two forms

  – Capacity: expiring

  – Inventory: often preserved
- Revenue management may also be defined as offering different prices based on customer segment, time of use and product or capacity availability to increase supply chain profits
- Most common example is probably in airline ticket pricing

  – Pricing according to customer segmentation at any time

  – Pricing according to reading days for any customer segment

  » Reading days: Number of days until departure
- The value of the product varies in different market segments

  – Airline seats: Leisure vs. Business travel

  – Films: Movie theater goers, DVD buyers, Cheap movie theater goers, TV watchers.
- The product is highly perishable or product waste occurs

  – Fashion and seasonal apparel

  – High tech products
- Demand has seasonal and other peaks

  – Products ordered at Amazon.com, peaking in December

  – Supply Chain textbook orders peaking in August and January.
- The product is sold both in bulk and on the spot market

  – Owner of warehouse who can decide whether to lease the entire warehouse through long-term contracts or save a portion of the warehouse for use in the spot market

  – Truck capacities for a transportation company

RM for Multiple Customer Segments:

- If a supplier serves multiple customer segments with a fixed asset, the supplier can improve revenues by setting different prices for each segment

– Must figure out customer segments

- Prices must be set with barriers such that the segment willing to pay more is not able to pay the lower price

  – Barriers: Time, location, prestige, inconvenience, extra service

- In the case of time barrier,

  – The amount of the asset reserved for the higher price segment is such that quantities below are equal

  » the expected marginal revenue from the higher priced segment

  » the price of the lower price segment

RM in the Service Industries

- Airline Industry uses RM the most.
- Evidence of airline revenue increases of 4 to 6 percent:

  – With effectively no increase in flight operating costs

- RM allows for tactical matching of demand vs. supply:

  – Booking limits can direct low-fare demand to empty flights

  – Protect seats for highest fare passengers on forecast full flights

- Hotel, Restaurant, Car rental, Overseas shipping, Cruise travel, Transportation capacity providers, Computation capacity providers (computer farms) and sometimes Health care industries show similarities to airline industry in using RM.

RM for Perishable Assets

- Any asset that loses value over time is perishable
- Examples: high-tech products such as computers and cell phones, high fashion apparel, underutilized capacity, fruits and vegetables
- Two basic approaches:

  – Dynamic Pricing: Vary price over time to maximize expected revenue

  – Overbooking: Overbook sales of the asset to account for cancellations

  - » Airlines use the overbooking most
  - » Passengers are “offloaded” to other routes
  - » Offloaded passengers are given flight coupons
  - » This practice is legal

  – Dynamic pricing belongs to RM while overbooking can be said to more within the domain of Yield management.

  - » But concepts are more important than the names!

- Overbooking or overselling of a supply chain asset is valuable if order cancellations occur and the asset is perishable
- The level of overbooking is based on the trade-off between the cost of wasting the asset if too many cancellations lead to unused assets (spoilage) and the cost of arranging a backup (offload) if too few cancellations lead to committed orders being larger than the available capacity
- Spoilage and offload are actually terms used in the airline industry

RM for Seasonal Demand

- Seasonal peaks of demand are common in many SCs\
  – Most retailers achieve a large portion of total annual demand in December
  » Amazon.com
- Off-peak discounting can shift demand from peak to nonpeak periods
- Charge higher price during peak periods and a lower price during off-peak periods

Types of Pricing Methods:

An organization has various options for selecting a pricing method. Prices are based on three dimensions that are cost, demand, and competition.

The organization can use any of the dimensions or combination of dimensions to set the price of a product.

Different pricing methods:

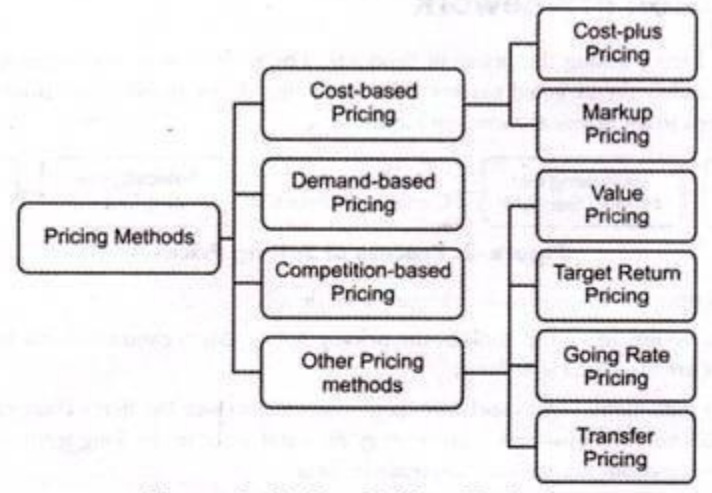

Figure-4: Various Pricing Methods

Cost-based Pricing:

Cost-based pricing refers to a pricing method in which some percentage of desired profit margins is added to the cost of the product to obtain the final price. In other words, cost-based pricing can be defined as a pricing method in which a certain percentage of the total cost of production is added to the cost of the product to determine its selling price. Cost-based pricing can be of two types, namely, cost-plus pricing and mark-up pricing.

These two types of cost-based pricing are as follows:

i. Cost-plus Pricing:

Refers to the simplest method of determining the price of a product. In cost-plus pricing method, a fixed percentage, also called mark-up percentage, of the total cost (as a profit) is added to the total cost to set the price. For example, XYZ organization bears the total cost of Rs. 100 per unit for producing a product. It adds Rs. 50 per unit to the price of product as' profit. In such a case, the final price of a product of the organization would be Rs. 150.

Cost-plus pricing is also known as average cost pricing. This is the most commonly used method in manufacturing organizations.

In economics, the general formula given for setting price in case of cost-plus pricing is as follows:

P = AVC + AVC (M)

AVC= Average Variable Cost

M = Mark-up percentage

AVC (m) = Gross profit margin

Mark-up percentage (M) is fixed in which AFC and net profit margin (NPM) are covered.

AVC (m) = AFC+ NPM

ii. For determining average variable cost, the first step is to fix prices. This is done by estimating the volume of the output for a given period of time. The planned output or normal level of production is taken into account to estimate the output.

The second step is to calculate Total Variable Cost (TVC) of the output. TVC includes direct costs, such as cost incurred in labour, electricity, and transportation. Once TVC is calculated, AVC is obtained by dividing TVC by output, Q. [AVC= TVC/Q]. The price is then fixed by adding the mark-up of some percentage of AVC to the profit [P = AVC + AVC (m)].

iii. The advantages of cost-plus pricing method are as follows:

a. Requires minimum information

b. Involves simplicity of calculation

c. Insures sellers against the unexpected changes in costs

The disadvantages of cost-plus pricing method are as follows:

a. Ignores price strategies of competitors

b. Ignores the role of customers

iv. Mark-up Pricing:

Refers to a pricing method in which the fixed amount or the percentage of cost of the product is added to product's price to get the selling price of the product. Mark-up pricing is more common in retailing in which a retailer sells the product to earn profit. For example, if a retailer has taken a product from the wholesaler for Rs. 100, then he/she might add up a mark-up of Rs. 20 to gain profit.

It is mostly expressed by the following formulae:

a. Mark-up as the percentage of cost= (Mark-up/Cost) *100

b. Mark-up as the percentage of selling price= (Mark-up/ Selling Price)*100

c. For example, the product is sold for Rs. 500 whose cost was Rs. 400. The mark up as a percentage to cost is equal to (100/400)*100 =25. The mark up as a percentage of the selling price equals (100/500)*100= 20.

Demand-based Pricing:

Demand-based pricing refers to a pricing method in which the price of a product is finalized according to its demand. If the demand of a product is more, an organization prefers to set high prices for products to gain profit; whereas, if the demand of a product is less, the low prices are charged to attract the customers.

The success of demand-based pricing depends on the ability of marketers to analyse the demand. This type of pricing can be seen in the hospitality and travel industries. For instance, airlines during the period of low demand charge less rates as compared to the period of high demand. Demand-based pricing helps the organization to earn more profit if the customers accept the product at the price more than its cost.

Competition-based Pricing:

Competition-based pricing refers to a method in which an organization considers the prices of competitors' products to set the prices of its own products. The organization may charge higher, lower, or equal prices as compared to the prices of its competitors.

The aviation industry is the best example of competition-based pricing where airlines charge the same or fewer prices for same routes as charged by their competitors. In addition, the introductory prices charged by publishing organizations for textbooks are determined according to the competitors' prices.

Other Pricing Methods:

In addition to the pricing methods, there are other methods that are discussed as follows:

i. Value Pricing:

Implies a method in which an organization tries to win loyal customers by charging low prices for their high- quality products. The organization aims to become a low cost producer without sacrificing the quality. It can deliver high- quality products at low prices by improving its research and development process. Value pricing is also called value-optimized pricing.

ii. Target Return Pricing:

Helps in achieving the required rate of return on investment done for a product. In other words, the price of a product is fixed on the basis of expected profit.

iii. Going Rate Pricing:

Implies a method in which an organization sets the price of a product according to the prevailing price trends in the market. Thus, the pricing strategy adopted by the organization can be same or similar to other organizations. However, in this type of pricing, the prices set by the market leaders are followed by all the organizations in the industry.

iv. Transfer Pricing:

Involves selling of goods and services within the departments of the organization. It is done to manage the profit and loss ratios of different departments within the organization. One department of an organization can sell its products to other departments at low prices. Sometimes, transfer pricing is used to show higher profits in the organization by showing fake sales of products within departments.

Introduction to Dynamic Pricing

Dynamic pricing, also referred to as **surge pricing**, **demand pricing**, or **time-based pricing** is a pricing strategy in which businesses set flexible prices for products or service based on current market demands. Businesses are able to change prices based on algorithms that take into account competitor pricing, supply and demand, and other external factors in the market.

Dynamic pricing is a common practice in several industries such as hospitality, travel, entertainment, retail, electricity, and public transport. Each industry takes a slightly different approach to repricing based on its needs and the demand for the product.

Dynamic pricing can be unpopular with consumers and favours the wealthy, which are less likely to be priced out of a market when there is high demand, such as for electricity during a heat wave or for food during a famine.

Overview of Dynamic Pricing

Dynamic pricing implies a technique focusing on the setting of any product/product's price in accordance with its supply and demand but in a relatively small span of time. In addition, a retailer has to consider a few other factors, like customers' perception, competitor pricing and value of a specific brand.

The psychology behind Dynamic Pricing

A majority of online retailers perceives dynamic pricing an entirely new concept and thinks that it has arrived only after the booming of e-commerce industry. However, offline retailers have remained involved in dynamic pricing for a long time. At present, dynamic pricing is prominent on varying online platforms because of the fact that the online market is relatively more agile and allows easy access to the available information.

Once online airline companies or booking sites become sure that you will book their offered services, they create a price in accordance to the number of seats left, the specific location, from where you book and the type of device you use. Each of these steps helps in determining the ability to shell out.

Techniques related to dynamic pricing are according to demand fluctuations and behavior of people. However, it is somewhat different from personalized pricing, as dynamic pricing collects the innumerable number of data and determines an appropriate price to help any vendor to earn the highest possible profits and/or revenues.

The technology behind Dynamic Pricing

Dynamic pricing today includes/adopts innovative technologies to obtain dynamic price ranges. Whenever e-commerce websites target prospective customers and realize that a person requires a specific product but until now does not become a customer, they opt to remarket their products.

For this, eCommerce marketers apply decent price reductions to convert a window shopper into a permanent shopper. Especially, they analyze cookie data of the respective individual and use other related tools, like Facebook Pixel and Google Adwords to generate a discounted price to entice the same individual. Dynamic pricing is perfect in this case to determine whether a person would purchase any product after changes made in its price.

For instance, Dealdash, eBay and other auction-based eCommerce stores use dynamic pricing to make the most from their customers. They have used and continue using special algorithms to increase the product price with each of the placed bids and sell the same to customers with a unique one.

Benefits of Dynamic Pricing

Dynamic pricing gives the following key benefits to an online retailer or eCommerce marketer.

- Provides higher control over the already existing pricing strategy
- Gives forecasts on upcoming trends in the customers' demands
- Helps in huge revenue and income growth
- Gives a strong platform to conduct customer analysis to perform personalized pricing, while market analysis to achieve price competition
- Facilitates for highly précised and SKU pricing levels
- Assures fast response to fluctuations in the demand
- As changes in price consider different factors, including the perception about the price of a product among customers, it leads to a long-term increase in profit and sales.

Different Types of Dynamic Pricing

Segmented Pricing: Segmented pricing features a pricing strategy that offers different pricing for different customers. For example, high-value customers might be offered prices that are higher because they are willing to pay more money for a quicker or higher quality service.

Peak Pricing: Peak pricing is used by different industries to charge extra money for the use of services that take place during peak hours. For example, a train ticket may cost twice the price between 8AM and 10AM but the price will drop after the peak commuting hours.

Time Based Pricing: Businesses use a time-based pricing strategy when they charge more money for faster services. Different industries will use this pricing structure to charge higher prices for same-day services. Another way to implement this structure includes charging more for orders that are processed close to the end of business hours.

Penetration Pricing: Penetration pricing is a dynamic pricing strategy that involves a business setting the initial price of a product below normal market level in order to drive demand. The low price is designed to reach a large portion of the market but can also be used to increase market share.

Changing Conditions: The changing conditions strategy involves a business using a strategy to boost profits when the conditions of the market are changing. For example, when there is a lot of uncertainty in the market or the market opportunity has only a short window of time. This strategy works by lowering the price as sales begin to fall and then raising the prices back up again.

B. Decision Variables

The following set of decision variables are the most significant ones and does not include all decision variables:

1. **The Demand:** Demand can be modelled in a simple way: assume 10 purchases per hour; or complex: random arrivals with time varying means and deviations. Demand is usually considered as exogenous to the model by most, but that is not the case, and is generally a simplifying assumption. Market always responds to product changes, even under a perfect monopoly.

2. **The Inventory:** Pricing and indexing of inventory can also be done easily: last 80 seats to be sold at 2.5x the price; or modelled statistically and probabilistically: demand responsive, time dependent and fare class conscious, with demand and purchase probabilities for each bracket of customer and fare class combination statistically estimated.

3. **The Purchase Probability:** The probability of an arriving customer (modelled in 1) who will buy a product (priced and indexed in 2) for the listed price is computed. For example, 7 in 10 people will buy 3rd AC, 2 will buy 2nd AC and 1 will buy 1st AC. It can also be done by computing the probability of purchase by applying utility theory, logit models, probabilistic regression of past purchases, etc.

4. **Nesting:** Purchases are not often straightforward. There are subclasses of products than can be nested. When a customer goes to buy a ticket, he usually has a few conditions that must be met (say arrival and departure times), some that should be met within a reasonable range (price, travel time), and some that are perceived as added bonus (on board catering, included in the ticket price or number of stops on the route). Most of these can be nested. Consider departure time is the parent super-nest, journey time will be a sub-nest, pantry service either can be a function of journey time and departure time, or can be nested within journey time, number of revenue stops can again be modelled as a nest within journey time. Thus, Indian Railways can fix price limits based on the amenities provided like, journey time, departure and arrival times, stops and pantry service.

Dynamic price with or without Replenishment

Dynamic pricing, in which prices are set to partially influence the demand may be seen as an ex ante tool to provide price-based product substitutions. In the same light, product upgrades are often implemented after demand is observed, and the practice may be viewed as an ex post tool in the form of a one-way, availability based product substitution.

The following main questions that naturally arise when pricing, replenishment, and upgrades are jointly considered:

(1) How should a firm decide on extending product upgrade offers in any given period?

(2) How should the firm set prices and replenishment levels for its products in each period?

(3) What is the impact of a firm's willingness to offer upgrades on its pricing and replenishment decisions?

The optimal upgrade, replenishment, and pricing policies:

The second-stage optimal upgrade policy is defined by a protection level on the higher quality item, where the protection level depends on the intermediate inventory positions of the products through their sum. Further, we provide monotonicity results on how this threshold changes with the total intermediate inventory.

Second, we analytically explore the impact of offering product upgrades to the firm's pricing and replenishment decisions. Particularly, we find that offering product upgrades in instances where there are inventory imbalances between the products may allow the firm to restrain its reliance on dynamic pricing as the sole mechanism to counteract the imbalance. Consequently, the optimal price difference between the products in such instances is closer to their respective list prices, enabling the firm to have a more consistent price positioning between the products

Replenishment costs influence the optimal policy

Third, we study how the quality differentiation between the products or changes in replenishment costs influence the optimal policy. An increase in the quality difference between the products leads the firm to increase the base-stock level for the higher quality product, lower the base-stock level for the lower quality product and apply list prices that are further apart. Regarding replenishment costs, we find that the firm's pick of replenishment level and prices leads it to offer fewer subsequent customer upgrades if the replenishment cost for the higher quality product increases and to offer more upgrades if the cost for the lower quality product increases. We also provide sensitivity results for a correlated cost structure with a change in an underlying cost parameter driving the replenishment costs for both products.

Multiple Product Pricing:

Generally, organizations produce more than one product in their line of production. Even a single product of an organization can differ in styles and sizes. For example, a refrigerator manufacturing organization produces refrigerators in different colors, sizes, and features. Similarly, an automobile organization manufactures vehicles in different colors, sizes, and mileage. The pricing in case of multiple products is called multiple product pricing. The demand curve for multiple products would be different. However, the MC curve of these products is same as these are produced under interchangeable production facilities. Therefore, AR and MR curves are different for each product. On the other hand, AC and MC are inseparable. Therefore, the condition of MR=MC cannot be applied directly to fix the prices of each product.

The solution of this problem was provided by E.W. Clemens who stated how the multi-product organizations fix prices of their products. Suppose there are four differentiated products. A, B, C, and D produced by an organization.

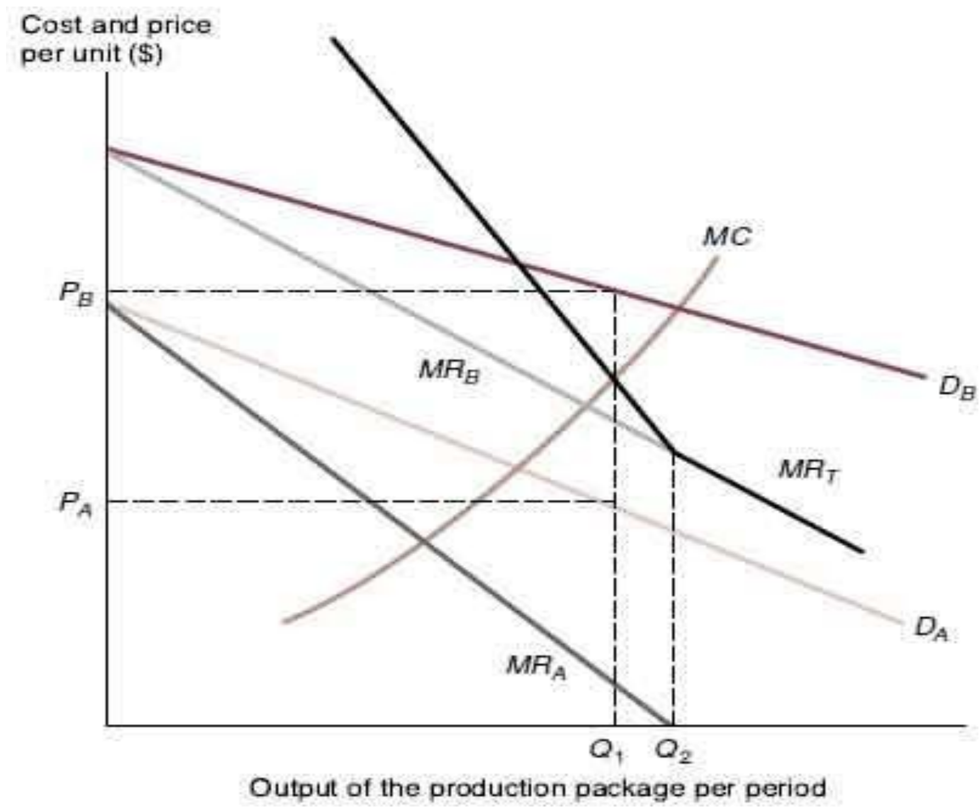

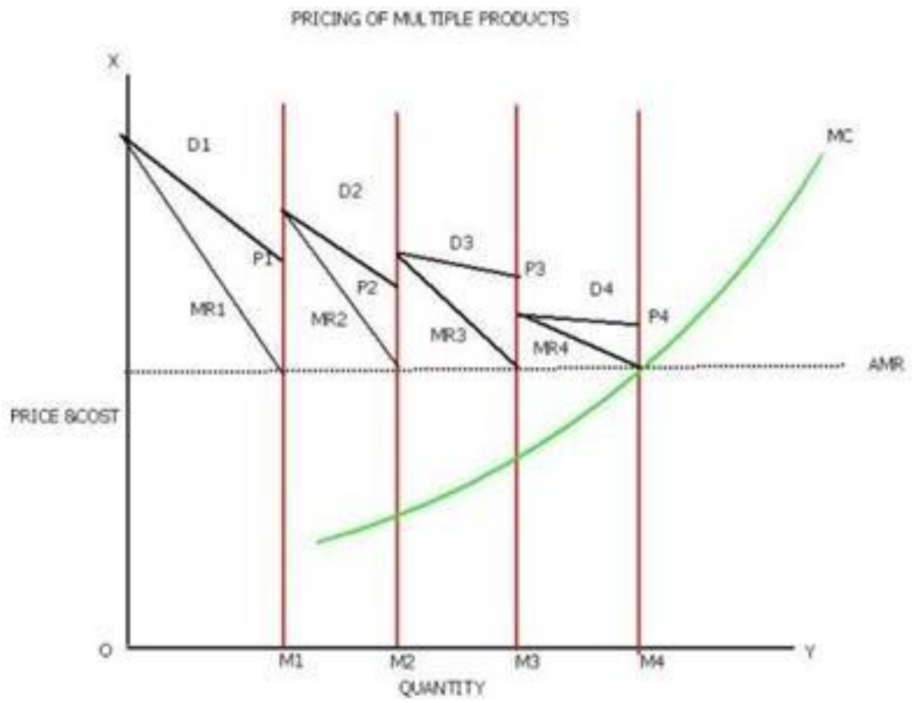

As shown in Figure, the AR (price) and MR curves for four products are shown as four different curves and MC curve is shown as the total of MC of all the products. Suppose the aggregate MR curve, which is the total of all individual MR curves, passes through point E on the MC curve. From point E, a parallel line, Average marginal revenue (AMR) is drawn towards Y- axis (parallel to X-axis). This parallel line passes through the M Rs of A, B, C, and D. The output and prices of these four products are determined at the points where their respective MC and MR curves intersect each other.

Pricing decisions - the volume / revenue trade-off:

The **adjustment in pricing** will affect the **volume change**, and affect the revenue.

Raising Price:

When you raise the price of your product, two things will happen.

1) **Get more revenue for every unit**
2) Able to **sell fewer units**

The danger here is that your price increase could **drive away too many customers**. The math is simple, the but information you need can be difficult to capture. You need to determine the following:

- What is the price increase you have in mind?
- How much will your demand decrease based on the price increase?

To illustrate the calculation, let's assume you currently sell 100 widgets at Rs.100 each, bringing in Rs.10,000 in revenues. You plan on raising prices 10% to Rs.110 per widget.

Scenario 1: You think your price increase will cause demand to drop by 10%

In this scenario, you will still sell 90 widgets.

At $110 per widget, you bring in $9,900.

Your revenue has decreased by $100

Scenario 2: You think your price increase will cause demand to drop by 5%

In this scenario, you will still sell 95 widgets

At $110 per widget, you bring in $10,450.

Your revenue has increased by $450

As you can see, **what you believe about the impact of price on demand** is a critical part of this concept.

Decreasing Price:

When you raise the price of your product, two things will happen.

1) You'll **sell more units**
2) You'll **make less money per unit**.

The danger here is that your **price cuts are too deep** and you give up too much revenue to your current customers compared to the new revenue you'll get from new customers. Again, the math is simple, the but you need that critical information on what impact your price change will have on demand.

- What is the price decrease you have in mind?
- How much additional demand will that attract?

To illustrate the calculation, let's again assume you currently sell 100 widgets at $100 each, bringing in $10,000 in revenues. You plan on lowering prices 5% to $95 per widget.

Scenario 1: You think your price decrease will attract 5% incremental demand

You attract 5 new customers, for a total of 105

But all of your customers - including the ones who used to pay $100 - now only pay $95

You make a total of $9,975

Your revenue has decreased by $25

Scenario 2: You think your price decrease will attract 10% incremental demand

You attract 10 new customers, for a total of 110

But all of your customers - including the ones who used to pay $100 - now only pay $95

You make a total of $10,450

Your revenue has increased by $450

Price Sensitivity or Price Elasticity of Demand:

Price sensitivity is just what it sounds like - it defines **how strongly your customers will react** to a change in price. **High price sensitivity** - or price elasticity of demand - means a relatively small change in price will drive a big change in demand. **Low price sensitivity** means even large changes in price will cause little change in demand.

- **Luxury goods** and pleasure travel have high elasticity or are highly price sensitive. **Essentials** like fuel, medicine, and food staples often have low elasticity or are not very price sensitive.
- **High price sensitivity means raising prices is risky** - a small price increase might result in a big drop in demand. But **lowering prices could drive a big increase in demand**.
- **Low price sensitivity means lowering prices is risky** - you might lower prices a lot and not see much increase in demand. But it might be **safe to raise prices** since your customers are less driven by price.
- If your product is a **commodity and not highly differentiated** from what your competitors sell, your customers are likely to be **highly price sensitive**.

Pricing and Revenue Management
Module-II

Perfect Competition:

The Perfect Competition is a market structure where a large number of buyers and sellers are present, and all are engaged in the buying and selling of the homogeneous products at a single price prevailing in the market.

In other words, perfect competition also referred to as a pure competition, exists when there is no direct competition between the rivals and all sell identically the same products at a single price.

Features of Perfect Competition:

Large number of buyers and sellers: In perfect competition, the buyers and sellers are large enough, that no individual can influence the price and the output of the industry. An individual customer cannot influence the price of the product, as he is too small in relation to the whole market. Similarly, a single seller cannot influence the levels of output, which is too small in relation to the gamut of sellers operating in the market.

Homogeneous Product: Each competing firm offers the homogeneous product, such that no individual has a preference for a particular seller over the others. Salt, wheat, coal, etc. are some of the homogeneous products for which customers are indifferent and buy these from the one who charges a less price. Thus, an increase in the price would let the customer go to some other supplier.

Free Entry and Exit: Under the perfect competition, the firms are free to enter or exit the industry. This implies, If a firm suffers from a huge loss due to the intense competition in the industry, then it is free to leave that industry and begin its business operations in any of the industry, it wants. Thus, there is no restriction on the mobility of sellers.

Perfect knowledge of prices and technology: This implies that both the buyers and sellers have complete knowledge of the market conditions such as the prices of products and the latest technology being used to produce it. Hence, they can buy or sell the products anywhere and anytime they want.

No transportation cost: There is an absence of transportation cost, i.e. incurred in carrying the goods from one market to another. This is an essential condition of the perfect competition since the homogeneous product should have the same price across the market and if the transportation cost is added to it, then the prices may differ.

Absence of Government and Artificial Restrictions: Under the perfect competition, both the buyers and sellers are free to buy and sell the goods and services. This means any customer can buy from any seller, and any seller can sell to any buyer. Thus, no restriction is imposed on either party. Also, the prices are liable to change freely as per the demand-supply conditions. In such a situation, no big producer and the government can intervene and control the demand, supply or price of the goods and services.

Perfect Competition Short Run Equilibrium: Supernormal Profits:

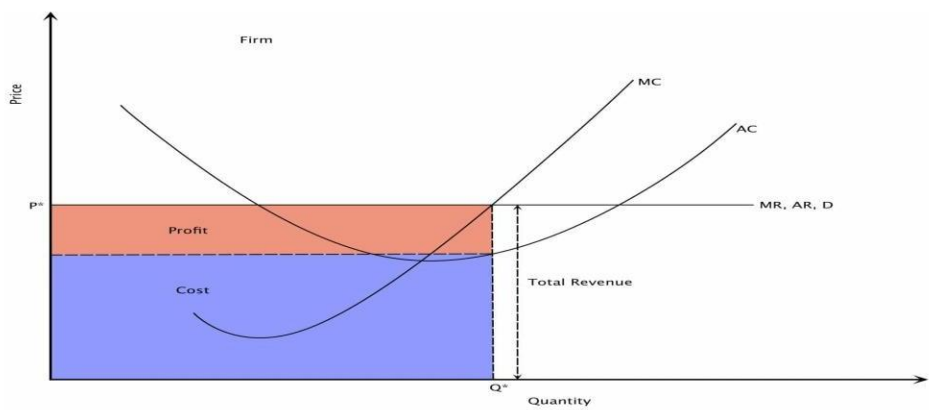

- Supernormal profit is all the excess profit a firm makes above the minimum return necessary to keep a firm in business.

- Supernormal profit is calculated by Total Revenue – Total Costs (where total cost includes all fixed and variable costs, plus minimum income necessary for the owner to be happy in that business.)
- Normal profit is defined as the minimum level of profit necessary to keep a firm in that line of business. This level of normal profit enables the firm to pay a reasonable salary to its workers and managers. The definition of normal profit occurs when AR=ATC (average revenue = average total cost)
- Supernormal profit is defined as extra profit above that level of normal profit.
- Supernormal profit is also known as abnormal profit. Abnormal profit means there is an incentive for other firms to enter the industry.

The theory of perfect competition suggests that supernormal profit can only be earned in the short term. In the long-term firms will make normal profit. Perfect competition is a market structure which involves:

- Perfect information
- Freedom of Entry and exit

Suppose there is a rise in demand, price rises and a firm can make supernormal profit in the short-term.

Perfect Competition Short Run Equilibrium Loss Making:

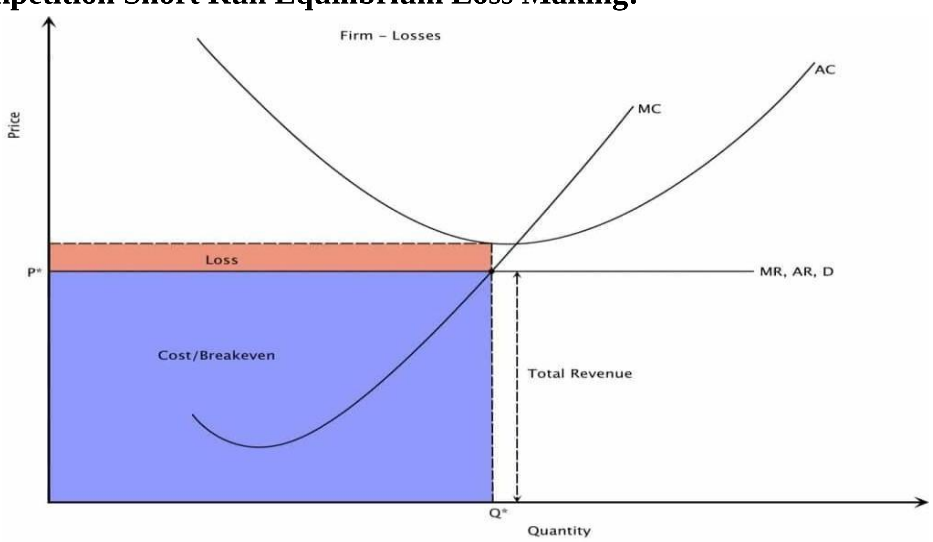

Loss:

Loss is relatively self-explanatory: if Jack's business takes in less revenue than it spends in a given period, it has experienced loss. Losses can occur for any number of reasons in any market condition. It might be the result of customer dissatisfaction with the company, even in a strong economy. The business might have rising expenses over time while simultaneously taking in less revenue.

Perfect Competition Short Run Zero Economic Profits:

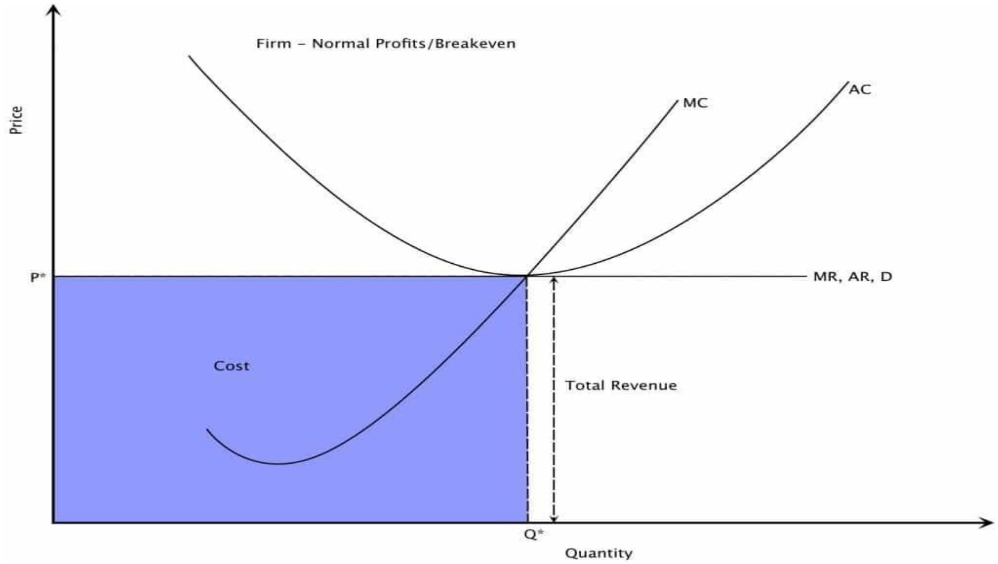

Normal profit

In markets which are perfectly competitive, the profit available to a single firm in the long run is called normal profit. This exists when total revenue, TR, equals total cost, TC. Normal profit is defined as the minimum reward that is just sufficient to keep the entrepreneur supplying their enterprise. In other words, the reward is just covering opportunity cost - that is, just better than the next best alternative.

The accounting definition of profits is rather different because the calculation of profits is based on a straightforward numerical calculation of past monetary costs and revenues, and makes no reference to the concept of opportunity cost. Accounting profit occurs when revenues are greater than costs, and not equal, as in the case of normal profit.

To the economist, normal profit is a cost and is included in the total costs of production.

Imperfect Competition:

In real life, perfect competition or even pure competition are seldom met with. On the other hand, it is imperfect competition which is the rule, and perfect competition is the exception. There are different degrees of imperfect competition, ranging from what is called- 'monopolistic competition' to 'simple monopoly'. In between these two forms of imperfect competition are 'oligopoly' and 'duopoly".

Monopoly Market

The Monopoly is a market structure characterized by a single seller, selling the unique product with the restriction for a new firm to enter the market. Simply, monopoly is a form of market where there is a single seller selling a particular commodity for which there are no close substitutes.

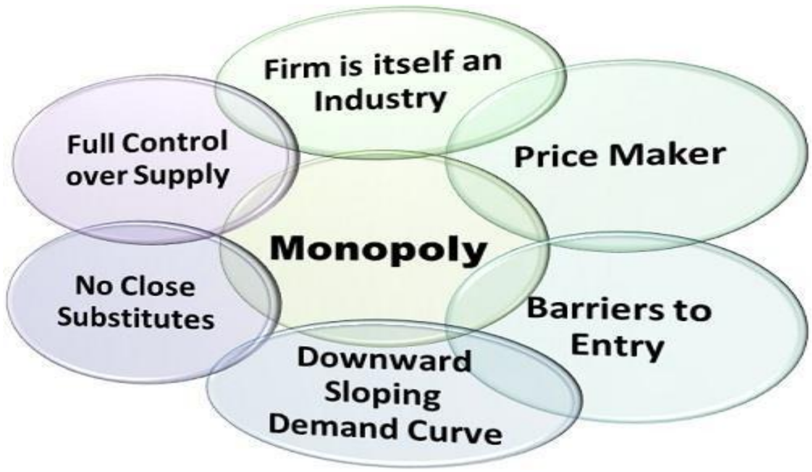

- Under monopoly, the firm has full control over the supply of a product. The elasticity of demand is zero for the products.
- There is a single seller or a producer of a particular product, and there is no difference between the firm and the industry. The firm is itself an industry.
- The firms can influence the price of a product and hence, these are price makers, not the price takers.
- There are barriers for the new entrants.
- The demand curve under monopoly market is downward sloping, which means the firm can earn more profits only by increasing the sales which are possible by decreasing the price of a product.
- There are no close substitutes for a monopolist's product.
- Under a monopoly market, new firms cannot enter the market freely due to any of the reasons such as Government license and regulations, huge capital requirement, complex technology and economies of scale. These economic barriers restrict the entry of new firms.

Price and Output Decisions in Pure Monopoly Markets

Demand in Monopoly Markets

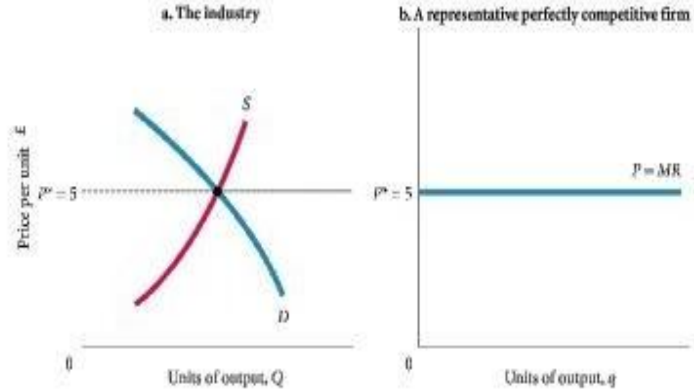

The Demand Curve Facing a Perfectly Competitive Firm Is Perfectly Elastic

Perfectly competitive firms are price-takers; they are small relative to the size of the market and thus cannot influence market price. The implication is that the demand curve facing a perfectly competitive firm is perfectly elastic. If the firm raises its price, it sells nothing and there is no reason for the firm to lower its price if it can sell all it wants at P* = £5.

Monopolistic Competition:

The monopolistic competition is also called as imperfect competition because this market structure lies between the pure monopoly and the pure competition. Under, the Monopolistic Competition, there are a large number of firms that produce differentiated products which are close substitutes for each other. Here large sellers selling the products that are similar but not identical and compete with each other on other factors besides price.

Features of Monopolistic Competition:

Product Differentiation: This is one of the major features of the firms operating under the monopolistic competition that produces the product which is not identical but is slightly different from each other. The products being slightly different from each other remain close substitutes of each other and hence cannot be priced very differently from each other.

Large number of firms: A large number of firms operate under the monopolistic competition, and there is a stiff competition between the existing firms. Unlike the perfect competition, the firms produce the differentiated products which are substitutes for each other, thus make the competition among the firms a real and a tough one.

Heavy expenditure on Advertisement and other Selling Costs: Under the monopolistic competition, the firms incur a huge cost on advertisements and other selling costs to promote the sale of their products. Since the products are different and are close substitutes for each other; the firms need to undertake the promotional activities to capture a larger market share.

Product Variation: Under the monopolistic competition, there is a variation in the products offered by several firms. To meet the needs of the customers, each firm tries to adjust its product accordingly. The changes could be in the form of new design, better quality, new packages or container, better materials, etc. Thus, the amount of product a firm is selling in the market depends on the uniqueness of its product and the extent to which it differs from the other products.

Demand curve of Monopolistic market:

The demand curve for the output produced by a monopolistically competitive firm is relatively elastic. The firm can sell a wide range of output within a relatively narrow range of prices. As a price maker, the firm has some ability (not much, but some) to control price. The demand curve is negatively sloped, but relatively elastic, because each firm produces a slightly differentiated product, but faces competition from a large number of very, very close substitutes.

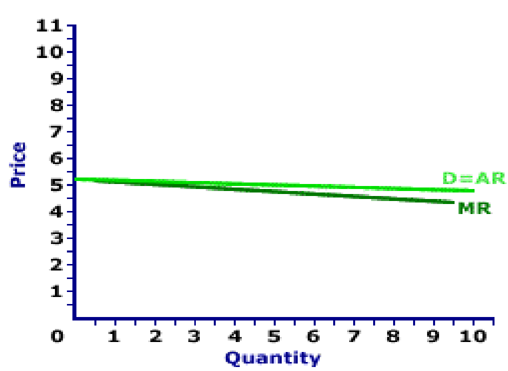

Relatively Elastic Demand

The four characteristics of monopolistic competition mean that a monopolistically competitive firm faces a relatively, but not perfectly, elastic demand curve, such as the one labeled D=AR and displayed in the exhibit to the right. Each firm in a monopolistically competitive market can sell a wide range of output within a relatively narrow range of prices.

Demand is relatively elastic in monopolistic competition because each firm faces competition from a large number of very, very close substitutes. However, demand is not perfectly elastic (as in perfect competition) because the output of each firm is slightly different from that of other firms. Monopolistically competitive goods are close substitutes, but not perfect substitutes.

Oligopoly Market

The Oligopoly Market characterized by few sellers, selling the homogeneous or differentiated products. In other words, the Oligopoly market structure lies between the pure monopoly and monopolistic competition, where few sellers dominate the market and have control over the price of the product.

Under the Oligopoly market, a firm either produces:

- **Homogeneous product:** The firms producing the homogeneous products are called as Pure or Perfect Oligopoly. It is found in the producers of industrial products such as aluminum, copper, steel, zinc, iron, etc.
- **Heterogeneous Product:** The firms producing the heterogeneous products are called as Imperfect or Differentiated Oligopoly. Such type of Oligopoly is found in the producers of consumer goods such as automobiles, soaps, detergents, television, refrigerators, etc.

There are five types of oligopoly market, for detailed description, click on the link below:

Features of Oligopoly Market

1. **Few Sellers:** Under the Oligopoly market, the sellers are few, and the customers are many. Few firms dominating the market enjoys a considerable control over the price of the product.
2. **Interdependence:** it is one of the most important features of an Oligopoly market, wherein, the seller has to be cautious with respect to any action taken by the competing firms. Since there are few sellers in the market, if any firm makes the change in the price or promotional scheme, all other firms in the industry have to comply with it, to remain in the competition.

   Thus, every firm remains alert to the actions of others and plan their counterattack beforehand, to escape the turmoil. Hence, there is a complete interdependence among the sellers with respect to their price-output policies.
3. **Advertising:** Under Oligopoly market, every firm advertises their products on a frequent basis, with the intention to reach more and more customers and increase their customer base. This is due to the advertising that makes the competition intense.

   If any firm does a lot of advertisement while the other remained silent, then he will observe that his customers are going to that firm who is continuously promoting its product. Thus, in order to be in the race, each firm spends lots of money on advertisement activities.
4. **Competition:** It is genuine that with a few players in the market, there will be an intense competition among the sellers. Any move taken by the firm will have a considerable impact on its rivals. Thus, every seller keeps an eye over its rival and be ready with the counterattack.
5. **Entry and Exit Barriers:** The firms can easily exit the industry whenever it wants, but has to face certain barriers to entering into it. These barriers could be Government license, Patent, large firm's economies of scale, high capital requirement, complex technology, etc. Also, sometimes the government regulations favor the existing large firms, thereby acting as a barrier for the new entrants.
6. **Lack of Uniformity:** There is a lack of uniformity among the firms in terms of their size, some are big, and some are small.

Types of Oligopoly Market

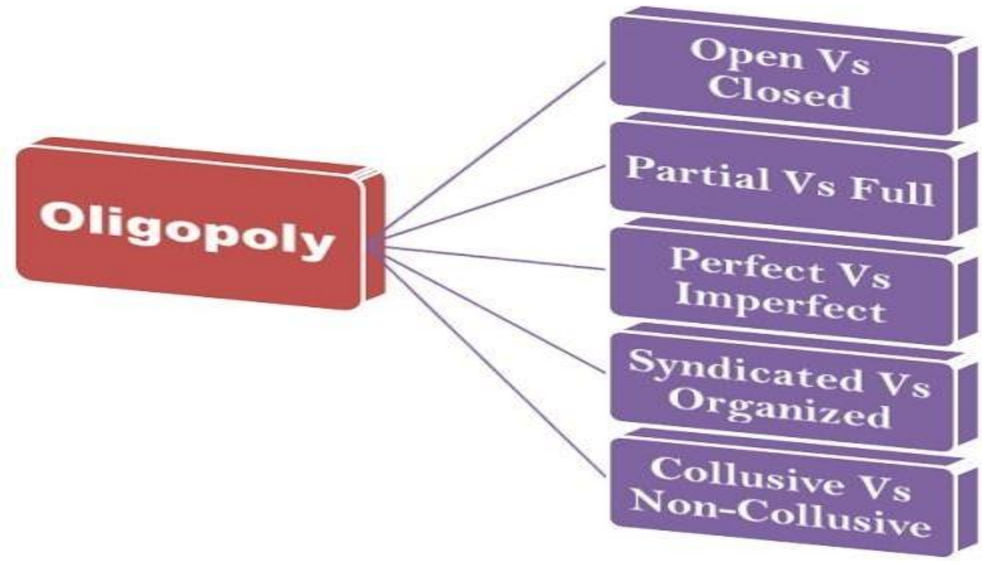

1. **Open Vs Closed Oligopoly:** This classification is made on the basis of freedom to enter into the new industry. An open Oligopoly is the market situation wherein firm can enter into the industry any time it wants, whereas, in the case of a closed Oligopoly, there are certain restrictions that act as a barrier for a new firm to enter into the industry.
2. **Partial Vs Full Oligopoly:** This classification is done on the basis of price leadership. The partial Oligopoly refers to the market situation, wherein one large firm dominates the market and is looked upon as a price leader. Whereas in full Oligopoly, the price leadership is conspicuous by its absence.
3. **Perfect (Pure) Vs Imperfect (Differential) Oligopoly:** This classification is made on the basis of product differentiation. The Oligopoly is perfect or pure when the firms deal in the homogeneous products. Whereas the Oligopoly is said to be imperfect, when the firms deal in heterogeneous products, i.e. products that are close but are not perfect substitutes.
4. **Syndicated Vs Organized Oligopoly:** This classification is done on the basis of a degree of coordination found among the firms. When the firms come together and sell their products with the common interest is called as a Syndicate Oligopoly. Whereas, in the case of an Organized Oligopoly, the firms have a central association for fixing the prices, outputs, and quotas.
5. **Collusive Vs Non-Collusive Oligopoly:** This classification is made on the basis of agreement or understanding between the firms. In Collusive Oligopoly, instead of competing with each other,

the firms come together and with the consensus of all fixes the price and the outputs. Whereas in the case of a non-collusive Oligopoly, there is a lack of understanding among the firms and they compete against each other to achieve their respective targets.

Demand Uncertainty:

Demand uncertainty occurs during times when a business or an industry is unable to accurately predict consumer demand for its products or services. This can cause a number of problems for the business, especially in managing orders and stocking levels, with effects magnifying through the supply chain.

Causes

The causes of demand uncertainty may result from inherent qualities of the business and its customer base, or from external factors. Seasonal fluctuations, for example, are a type of inherent uncertainty, although industries that experience seasonal fluctuations can often use records from past years to anticipate and estimate the current seasonal shift. Businesses with a very innovative product or service will face a great deal of demand uncertainty, simply because their uniqueness means that there is no track record from which to draw conclusions about demand.

Problems

When demand is uncertain, it’s difficult to determine the right quantity of supplies and goods to order for the next sales cycle. A business that anticipates a normal or high level of sales, only to see the demand drop, will have leftover goods that must be stored, returned or discarded. Each of these scenarios leads to extra costs. If demand increases, however, and the company does not have a sufficient supply of goods to sell, the result is dissatisfied customers, some of whom may purchase from a competitor that does have a supply of the desired goods. Some customers may never return to the original seller, resulting in loss of business for that company.

Bullwhip effect

Demand uncertainty at the retail customer level has a way of becoming magnified down the supply chain. When wholesalers notice that the retail outlets they serve have cut back or increased their orders, the wholesalers are likely to cut back or increase purchases from manufacturers by a greater percentage, in order to hedge their bets. Manufacturers, in turn, may hedge their bets even further on their orders from suppliers of raw materials. As result, raw materials suppliers may end up with a six month or greater backlog, on a supply chain that began with retailers ordering an extra supply of only one month. This is known as the bullwhip effect, because it charts in a way that resembles the action of a bullwhip when the "handle" is moved slightly while the "tail" swings in wider and wider fluctuations.

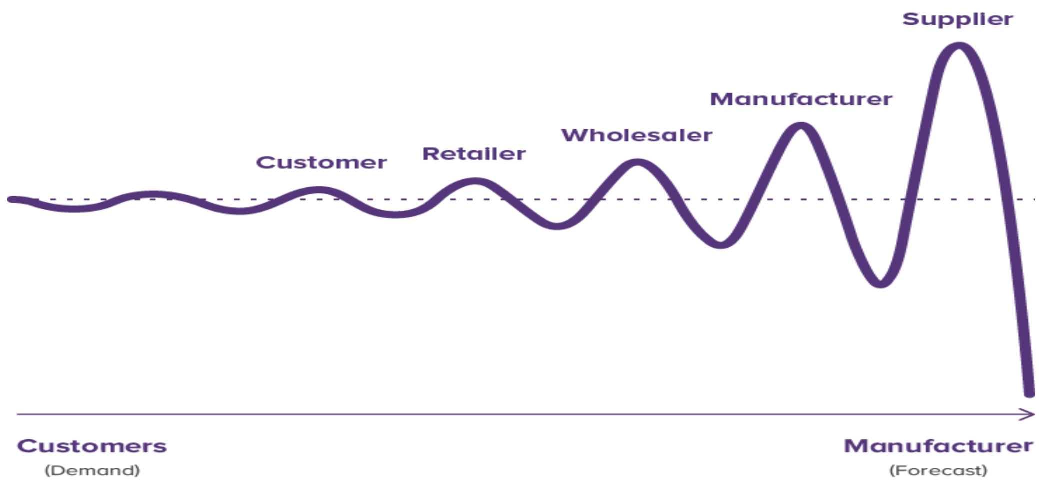

Factors said to cause or contribute to the bullwhip effect in supply chains:

- **Disorganization** between each supply chain link; with ordering larger or smaller amounts of a product than is needed due to an over or under reaction to the supply chain beforehand.
- **Lack of communication** between each link in the supply chain makes it difficult for processes to run smoothly. Managers can perceive a product demand quite differently within different links of the supply chain and therefore order different quantities.
- **Free return policies**; customers may intentionally overstate demands due to shortages and then cancel when the supply becomes adequate again, without return forfeit retailers will continue to exaggerate their needs and cancel orders; resulting in excess material.
- **Order batching**; companies may not immediately place an order with their supplier; often accumulating the demand first. Companies may order weekly or even monthly. This creates variability in the demand as there may for instance be a surge in demand at some stage followed by no demand after.
- **Price variations** – special discounts and other cost changes can upset regular buying patterns; buyers want to take advantage on discounts offered during a short time period, this can cause uneven production and distorted demand information.

- **Demand information** – relying on past demand information to estimate current demand information of a product does not take into account any fluctuations that may occur in demand over a period of time.

Industries Affected

Certain industries are more vulnerable to demand uncertainty than others. Ranked high in demand uncertainty, while low in technological uncertainty, were consumer-level industries, such as restaurants and hotels, healthcare services, and retail. However, materials suppliers such as coal, mining and steelworks ranked even higher in demand uncertainty. Furthermore, some industries face high demand uncertainty, along with high technological uncertainty. These industries include transportation, computers, software, medical equipment and pharmaceuticals.

Pricing and Revenue Management

Module-III

Revenue Management (RM) is concerned with the methodology and systems required making *demand management* decisions, which can be categorized into:

(i) **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 cancelation or refund options); how to bundle products; and so on.

(ii) **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.

(iii) **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.

SINGLE-RESOURCE CAPACITY CONTROL

In this section, we examine some basic results on the problem of quantity-based RM for a single resource; specifically, optimally allocating capacity of a resource to different classes of demand. Two prototypical examples are controlling the sale of different fare classes on a single flight leg of an airline and the sale of hotel rooms for a given date at different rate classes. This is to be contrasted with the multiple resources—or network—problems, in which customers require a bundle of different resources (such as two connecting flights or a sequence of nights at the same hotel).

We assume that the firm sells its capacity in n distinct classes (in the case of airlines, these are called *fare classes*) that require the same resource. Classes are ranked from 1 to n in decreasing order of their revenue values, with $r_1 > r_2 > \cdots > r_n$. In the airline and hotel context, these classes represent different discount levels with differentiated sale conditions and restrictions. The units of capacity are assumed homogeneous, and customers demand a single unit of capacity for the resource. The central problem of this section is how to optimally allocate the capacity of the resource to the various classes. This allocation must be done dynamically as demand materializes and with considerable uncertainty about the quantity or composition of future demand.

2.2 Static Models

The static model makes several simplifying assumptions that are worth examining in some detail. The first is that demand for the different classes arrives in non-overlapping intervals in the order of increasing prices of the classes. This could be justified by the observation that advance purchase discount demand typically arrives before full-fare coach demand in the airline case. Moreover, the optimal controls that emerge from the model can be applied—at least heuristically—even when demand comes in arbitrary order. As for the strict low-before-high assumption, this represents something of a worst-case scenario; for instance, if high-revenue demand arrives before low-revenue demand, the problem is trivial because we simply accept demand first come, first serve.

The second main assumption is that the demands for different classes are independent random variables. Largely, this assumption is made for analytical convenience because dealing with dependence in the demand structure would require introducing complex state variables on the history of observed demand.

A third assumption is that demand for a given class does not depend on the capacity controls; in particular, it does not depend on the availability of other classes. However, in practice, customers in a high revenue class may buy down to a lower class if the latter is available, and customers in a lower class may buy up to a higher class if the lower class is closed.

A fourth assumption in the static model is that it suppresses many details about the demand and control process within each of the periods. However, the form of the optimal control is not sensitive to this assumption.

A fifth assumption of the model is that either there are no groups, or if there are group bookings, they can be partially accepted. Finally, the static models assume risk-neutrality. This is a reasonable assumption in practice, since a firm implementing RM typically makes such decisions for a large number of products sold repeatedly (for example, daily flights or daily hotel room stays). Maximizing the average revenue, therefore, is what matters in the end. While we do not cover this case here, some researchers have recently analyzed the single-resource problem with risk-averse decision makers or using worst-case analysis, for example, Lan et al. (2007).

Types of Controls

In the travel industry, reservation systems provide different mechanisms for controlling availability. These mechanisms are usually deeply embedded in the software logic of the reservation system and, as a result, can be quite expensive and difficult to change. Therefore, the control mechanisms chosen for a given implementation are often dictated by the reservation system.

The first types of control are *booking limits* that ration the amount of capacity that can be sold to any particular class at a given point in time. For example, a booking limit of 18 on class 2 indicates that at most 18 units of capacity can be sold to customers in class 2. Beyond this limit, the class would be "closed" to additional class 2 customers. This limit of 18 may be less than the physical capacity, for example, when we protect capacity for future demand from class 1 customers. Booking limits are either *partitioned* or *nested*: A *partitioned booking limit* divides the available capacity into separate blocks (or *buckets*)-one for each class—that can be sold only to the designated class. With a *nested booking limit*, the capacity available to different classes overlaps in a hierarchical manner-with higher-ranked classes having access to all the capacity reserved for lower-ranked classes (and perhaps more). Let the nested booking limit for class *j* be denoted b_j . Then b_j is the maximum number of units of capacity we are willing to sell to classes j to n. So, naturally, b_1 is equal to the capacity. Effectively, nesting logic simply allows any capacity "left over" after selling to classes of lower ranks to become available for sale to classes of higher rank. Nesting booking limits in this way avoids the problem of capacity being simultaneously unavailable for a higher-ranked class yet available for lower-ranked classes. Most reservations systems that use booking-limit controls quite sensibly use nested rather than partitioned booking limits for this reason.

The second type of control are *protection levels* that specify an amount of capacity to reserve (protect) for a particular class or set of classes. Again, protection levels can be *nested* or *partitioned*. A partitioned protection level is trivially equivalent to a partitioned booking limit. In the nested case, protection levels are again defined for sets of classes—ordered in a hierarchical manner according to class order. The protection level j, denoted y_j , is defined as the amount of capacity to save for classes $j, j-1, \ldots, 1$ combined. The booking limit for class j, b_j is simply the capacity minus the protection level for classes $j-1$ and higher. That is,

$$b_j = C - y_{j-1}, \quad j = 2, \ldots, n,$$

where C is the capacity. For convenience, we define $b_1 = C$ (the highest class has a booking limit equal to the capacity) and $y_n = C$ (all classes combined have a protection level equal to capacity).

The third types of controls are *bid-price controls*. What distinguishes bid-price controls from both booking limits and protection levels is that they are revenue-based rather than class-based controls. Specifically, a bid-price control sets a threshold price (which may depend on variables such as the remaining capacity or time), such that a request is accepted if its revenue exceeds the threshold price and rejected if its revenue is less than the threshold price. Bid price controls are, in principle, simpler than booking-limit or protection-level controls because they require only storing a single threshold value at any point in time—rather than a set of capacity numbers, one for each class. But to be effective, bid prices must be updated after each sale—and possibly also with time as well—and this typically requires storing a table of bid price values indexed by the current available capacity, current time, or both.

As with single-resource problems, in network problems there are a variety of ways one can control the availability of capacity. We next look at the major categories of network controls. Most are network versions of the controls used for single-resource problems. But others, virtual nesting in particular, is somewhat unique to the network setting.

Virtual Nesting Controls

Nested booking limits, of the type we saw in Section 2 for the single-resource case, are difficult to translate directly into a network setting. However, the ability of nested controls to dynamically share the capacity of a resource - and thereby recover the pooling economies lost in partitioned controls - is an attractive feature. Thus, it is desirable to have a control that combines these features.

Virtual nesting control –a hybrid of network and single resource controls– provides one solution. This control scheme was developed by American Airlines beginning in 1983 as a strategy for incorporating some degree of network control within the single-leg nested allocation structure of American's (then leg-based) reservation systems; see Smith, Leimkuhler and Darrow (1992).

Virtual nesting uses single-resource nested booking controls at each resource in the network. However, the classes used in these nested allocations are not the fare classes themselves. Rather, they are based on a set of *virtual classes*. Products are assigned to a virtual class through a process known as *indexing*. This indexing could be updated over time as network demand patterns change, though typically indexing is not a "real time" process.

Nested booking limits (or protection levels) for each resource is then computed using these virtual classes.

Virtual nesting has proven to be quite effective and popular in practice, especially in the airline industry. It preserves the booking-class control logic of most airline computer reservation systems (CRS) yet incorporates network displacement cost information. It therefore provides a nice compromise between leg-level and full network O&D control.

Bid-price controls

While nested allocations are difficult to extend directly to networks, network bid-price controls are a simple extension of their single-resource versions described in Section 2.

In a network setting, a bid-price control sets a threshold price - or *bid-price* - for each resource in the network. Roughly, this bid-price is an estimate of the marginal cost to the network of consuming the next incremental unit of the resource's capacity. When a request for a product comes in, the revenue of the request is compared to the *sum* of the bid prices of the resources required by the product. If the revenue exceeds the sum of the bid prices, the request is accepted; if not, it is rejected.

Network Management:

- Once a seller with constrained capacity begins selling products that use multiple resources—such as an airline selling tickets on connecting flights or a hotel selling multiple-night stays—the revenue management problem becomes much more complex.
- The seller can no longer maximize total contribution by maximizing contribution from each resource independently.
- Rather, he needs to consider the interactions among the various products he sells and their effect on his ability to sell other products. This is the challenge of *network management.*
- *Network management* is applicable to any company controlling a set of constrained and perishable resources and selling products that consist of combinations of those resources.

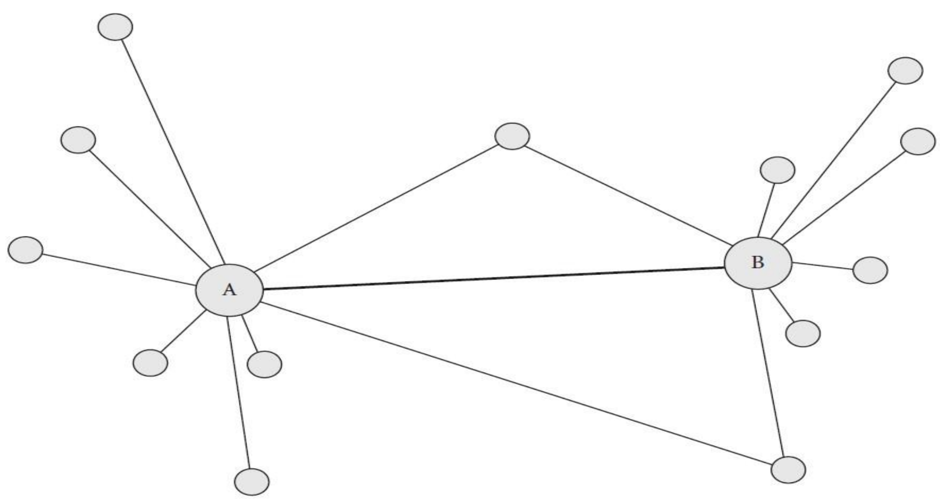

Figure 8.2 Example of a hub-and-spoke network. Nodes A and B are hubs; the remaining nodes are spoke cities.

- Each spoke city is connected to one or more of the hubs.
- The hub-and-spoke network allows an airline to offer a large number of products with a relatively small number of flights.
- With 20 eastbound flights into a hub connecting with 20 westbound flights, an airline can offer 440 products with only 40 total flights.

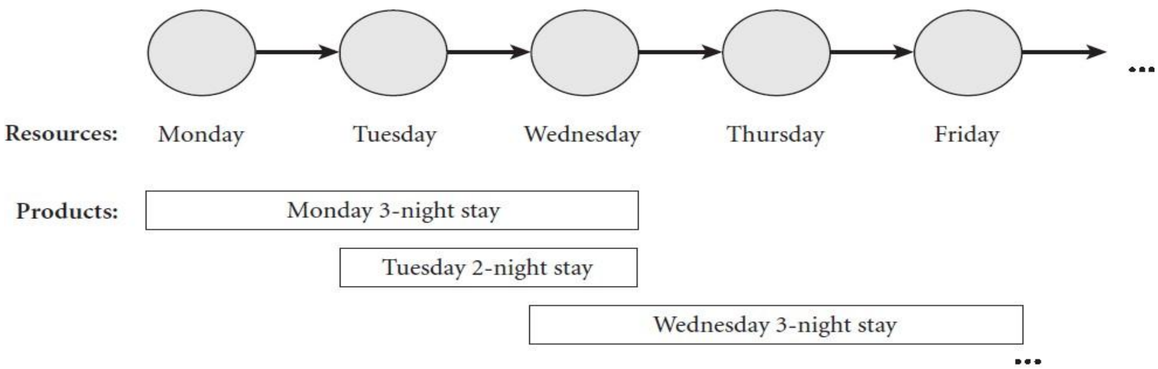

Figure 8.3 A linear network, typical of hotels and rental car companies.

- The resources for a hotel are room nights, and its products are all possible combinations of arrival date and length of stay.
- In a *linear network*, such as the one shown in Figure 8.3, each room night connects to the preceding night and to the following night.
- A passenger railway network is also linear; for example, Amtrak's *California Zephyr*, which starts at Emeryville, CA, and terminates in Chicago, makes 33 intermediate stops along the way.
- Each of the 34 legs is a resource, and each combination of origin and destination is a product.

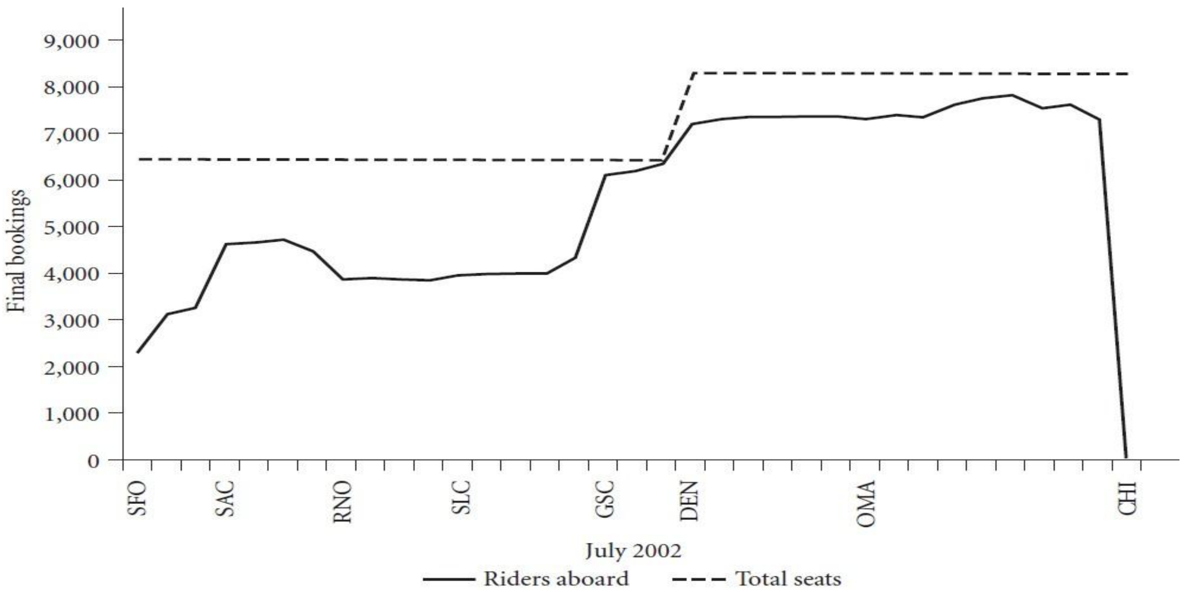

Figure 8.4 Capacity and passengers by leg for the *California Zephyr* train (July 2002). *Source*: Amtrak.

- Each product in a network has one or more associated fare classes.
- The network management problem is to determine which booking requests to accept for every possible combination of product and fare class at every time.
- We use the term *origin-destination fare class* (or ODF) for a product /fare class combination.
- For example, for a hotel, a four-night booking arriving June 13 is a different ODF than a three-night discount booking arriving June 14.
- The goal of network management is to manage and update the availability of all ODFs over time in order to maximize expected contribution.
- The number of products being offered by a hotel or rental car company at one time can be very large—much larger, in fact, than the number of rooms in the hotel.
- A hotel accepting reservations for customers arriving for the next 365 days with stays from 1 up to 15 days in length is offering 15x365=5,475 different products.
- In fact, the number of products can be even larger, since each room type creates an entirely new product dimension.
- A particular hotel might have four different types of room (for example, standard, deluxe, deluxe view, and suite), for which it can charge different rates.
- Since a customer can book any one of these room types for any arrival date and length of stay, the number of products being offered by the hotel is actually 4x15x365=21,900.
- In other words, a hotel may well be offering 100 times more products than it has rooms.

A Greedy Heuristic for Network Management

- Consider a 100-room hotel facing the *unconstrained demand* pattern for some future week shown in Figure 8.5.
- In this example, unconstrained demand exceeds capacity on Wednesday. This means that the hotel will need to reject some bookings or it will be oversold on Wednesday
- The network management problem faced by the hotel is to decide which bookings to accept and which to reject *considering both room rate and length of stay* in order to maximize expected contribution margin.

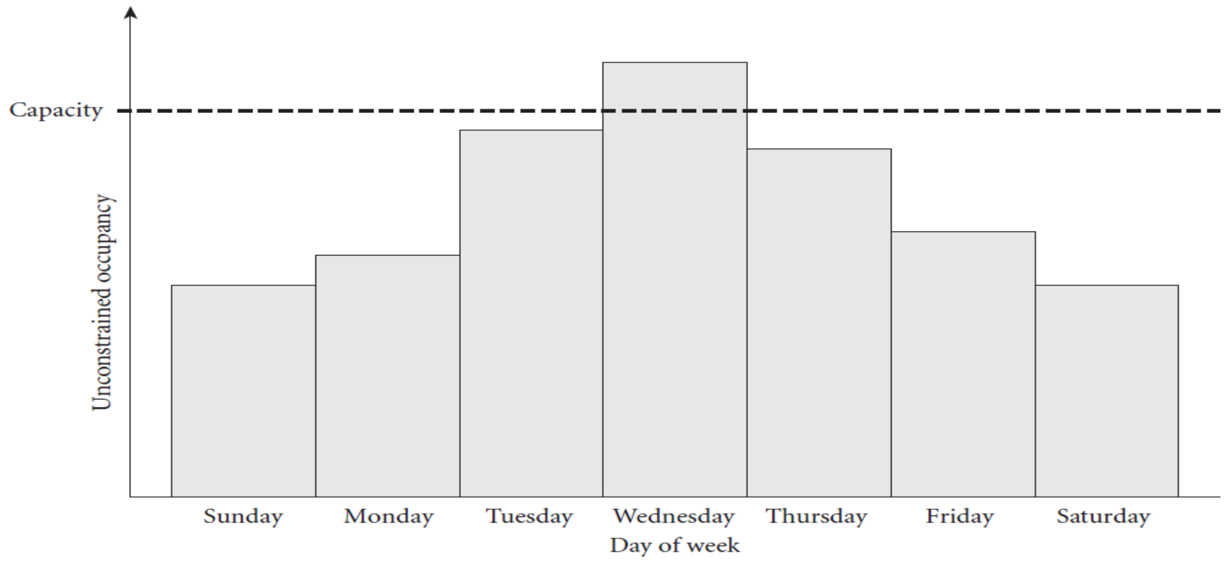

Figure 8.5 Typical unconstrained occupancy pattern for a midtown business hotel.

Unconstrained demand refers to the quantity of rooms in a hotel that could be sold if there were no constraints, no limits.

- Let us say that the hotel offers two rates: a rack-rate of $200 per night and a discount rate of $150 per night.
- The hotel manager might look at the demand pattern in Figure 8.5 and conclude (correctly) that the hotel's constrained resource is Wednesday night capacity.
- The manager might reason that the best way to maximize operating profit is to limit the number of discount-rate customers booking any product that includes Wednesday night in order to protect availability for rack-rate customers.
- This approach might improve revenue relative to a first-come, first-served policy, but it has an obvious drawback. Rack-rate customers arriving Wednesday for a one-night stay pay $200. But discount-rate customers arriving Tuesday for a three-night stay pay $450. If the hotel limits discount bookings to save room for full-fare bookings, it runs the risk of turning down a $450 customer to save room for a $200 customer
- Clearly a more sophisticated approach is required if the hotel is to maximize expected revenue.
- When there is only a single constrained resource, the manager could order all of the ODFs that use the constrained resource by total rate paid and solve the multiclass problem using Expected Marginal Seat Revenue (EMSR) or a similar heuristic.

- This would mean that the lowest class on Wednesday would be one-night discount Wednesday customers paying $150, the next-highest class would be one-night rack customers paying $200, followed by two-night discount stays for either Tuesday–Wednesday or Wednesday–Thursday paying $300, followed by two-night rack customers paying $400, and so on.
- By establishing classes in this fashion and setting the appropriate booking limits, the hotel would maximize expected operating contribution as long as unconstrained demand exceeded capacity *only on Wednesday night*.
- Sorting all the ODFs on each leg in fare order and setting availabilities based on total fare is called the *greedy heuristic*.
- The greedy heuristic is optimal if there is only a single constrained resource or bottleneck.
- However, the greedy heuristic breaks down quickly when more than one resource is constrained.
- This would be the case for the hotel if unconstrained demand exceeded 100 rooms on both Wednesday and Thursday nights.
- In this case, it is no longer optimal for the hotel to order fare classes strictly by total revenue.
- The hotel should accept a two-night Wednesday–Thursday discount booking paying $300 if it believes that Thursday is not going to sell out. But if it believes that Thursday is going to sell out, it may be better off rejecting the $300 two-night booking in favor of waiting for two $200 one-night rack-rate bookings for Wednesday night and Thursday night.
- When multiple resources are constrained, the right ordering of ODFs on a leg can depend on the fares and demands for all the other ODFs as well as the capacities of all the other resources in the network.
- There is no stable optimal ordering of ODFs when multiple resources in a network can be constrained.

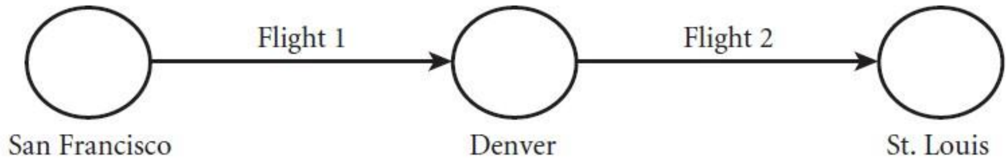

Figure 8.6 A simple hub-and-spoke network.

- The failure of the greedy heuristic can be seen with the simple network shown in Figure 8.6, where an airline offers two flights.
- For simplicity, let us assume that the airline offers a single fare for each product and that the San Francisco–to-Denver fare is $200, the Denver-to–St. Louis fare is $160, and the San Francisco–to–St. Louis fare is $300.
- Assume now that the airline has exactly one seat left on each of the two flights.
- The airline has a customer on the phone who wants to fly from San Francisco to St. Louis and is willing to pay the $300 fare for the privilege of doing so. Should the airline let her book? Or should the airline reject the booking in the hopes of booking local passengers on both legs?
- Let $p1$ be the probability that the airline will receive at least one future booking request for flight 1 and $p2$ be the probability that the airline will receive at least one future request for flight 2.
- Then the expected value of refusing the connecting customer and letting only the local customers book is $\$200p1+\$160p2$.
- If this amount is greater than $300, then the airline is better off refusing the connecting customer and relying on local traffic to fill the last seats.
- If, on the other hand, $\$200p1+\$160p2 < \$300$, then the airline should go ahead and allow the connecting passenger to book.
- Whether or not the airline should accept the through passenger depends on the forecasts of $p1$ and $p2$.
- This observation can be generalized to a full network: There is no unambiguous answer to the question of whether an airline should prefer connecting traffic to local traffic on a given leg. The right answer depends on the forecasts of future demand for all ODFs.
- Ultimately, when lots of resources are constrained, we need to optimize over the entire network in order to determine which bookings we prefer on any leg in the network.

A Linear Programming Approach:

- If we assume that future demands are known with certainty, linear programming provides an exact solution to the network management problem.
- Since uncertainty is a key element of revenue management decision making, it shouldn't be too surprising that a solution derived in this fashion is not wholly satisfactory.
- However, the linear programming formulation of the network management problem is worth studying, for three reasons.
- First, it solves the "capacity allocation" piece of network management and provides insight into the nature of the optimal solution.
- Second, the solution to the network management linear program provides a good starting point for a fully optimal solution. In many hotels and airlines, linear programming is used to determine an initial solution, which is then adjusted to account for the uncertainty in future demand.
- Finally, the linear program generates "marginal values" of capacity as a byproduct. These marginal values approximate the opportunity costs of the capacity and can serve as an estimate of network bid prices.
- Assume we know the future demand for every ODF with certainty.
- Furthermore, assume that this demand will not be influenced by which ODFs we open and close (in other words, we have independent demands both on each flight and among flight legs).
- We have m resources and n ODFs utilizing combinations of those resources.
- We assume each resource is used in at least one ODF, so $n \geq m$.
- We use the subscript i to index resources and the subscript j to index ODFs. Each resource i has a constrained capacity $c_i > 0$. Each ODF has a known demand $d_j > 0$ and a net margin $p_j > 0$.

$$a_{ij} = \begin{cases} 1 & \text{if resource } i \text{ is used in ODF } j \\ 0 & \text{otherwise} \end{cases}$$

- Let $x_j \geq 0$ be the amount of ODF j we will sell (our allocation).

- Then the *deterministic network management problem* is to find the values of x_j for j=1, 2, . . . , n that maximize total net contribution subject to the constrained capacities of the resources.

The Deterministic Network Linear Program:

$$\max_{x_j} \sum_{j=1}^{n} p_j x_j$$

subject to

$$\sum_{j=1}^{n} a_{ij} x_j \leq c_i \qquad \text{for all } i$$

$$x_j \leq d_j \qquad \text{for all } j$$

$$x_j \geq 0 \qquad \text{for all } j$$

A Linear Programming Approach:

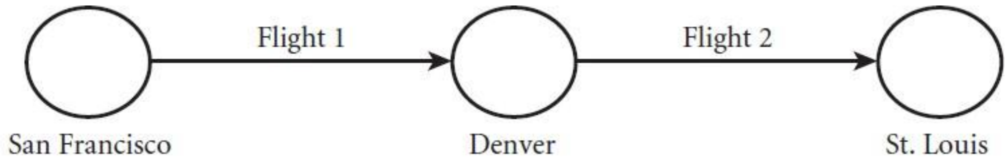

Figure 8.6 A simple hub-and-spoke network.

- In Figure 8.6., there are two resources: (1) flight 1 from San Francisco to Denver and (2) flight 2 from Denver to St. Louis.
- With three products and two fare classes, the airline is offering six ODFs in all:

1. San Francisco to Denver full fare
2. San Francisco to Denver discount
3. Denver to St. Louis full fare
4. Denver to St. Louis discount
5. San Francisco to St. Louis full fare
6. San Francisco to St. Louis discount

TABLE 8.2
Values of the incidence variables (a_{ij}) for the two-flight example

Resource (i)	ODF (j) 1	2	3	4	5	6
1	1	1	0	0	1	1
2	0	0	1	1	1	1

- A hotel that offers one-night, two-night, and three-night products only.
- The resources for the hotel are the room nights, and the products are combinations of arrival night and length of stay.
- Table 8.3 shows the values of the incidence variables (a_{ij}) for three different lengths of stay and a week of resources.
- In theory, the network management problem for a hotel stretches out indefinitely into the future and the number of products are infinite, since customers can buy any length of stay.

TABLE 8.3
Values of the incidence variables (a_{ij}) for the hotel example

	ARRIVAL DATE														
	SUNDAY			MONDAY			TUESDAY			WEDNESDAY			THURSDAY		
Resource	1	2	3	1	2	3	1	2	3	1	2	3	1	2	3
Sunday	1	1	1	0	0	0	0	0	0	0	0	0	0	0	0
Monday	0	1	1	1	1	1	0	0	0	0	0	0	0	0	0
Tuesday	0	0	1	0	1	1	1	1	1	0	0	0	0	0	0
Wednesday	0	0	0	0	0	1	0	1	1	1	1	1	0	0	0
Thursday	0	0	0	0	0	0	0	0	1	0	1	1	1	1	1
Friday	0	0	0	0	0	0	0	0	0	0	0	1	0	1	1
Saturday	0	0	0	0	0	0	0	0	0	0	0	0	0	0	1

Ex: An airline is operating the two-flight airline network in Figure 8.6 and offering a full fare and a discount fare for each of the three itineraries. There are six ODFs using two resources. The airline has assigned a 100-seat aircraft on flight 1 from San Francisco to Denver and a 120-seat aircraft on flight 2 from Denver to St. Louis; demands and fares are as shown in Table 8.4. Notice that the unconstrained demand is 160 on the San Francisco–to-Denver flight and 170 on the Denver-to–St. Louis flight. Since these demands exceed the flight capacities for the two flights, we will need to turn some passengers away.

TABLE 8.4

Two-flight network management example fares and demands

Number	ODF	Fare	Demand
1	San Francisco to Denver full fare	$150	30
2	San Francisco to Denver discount	100	60
3	Denver to St. Louis full fare	120	20
4	Denver to St. Louis discount	80	80
5	San Francisco to St. Louis full fare	250	30
6	San Francisco to St. Louis discount	170	40

The Deterministic Linear Programming:

$$\text{maximize} \quad 150x_1 + 100x_2 + 120x_3 + 80x_4 + 250x_5 + 170x_6$$

subject to

$$\begin{aligned}
x_1 + x_2 + x_5 + x_6 &\le 100\\
x_3 + x_4 + x_5 + x_6 &\le 120\\
x_1 &\le 30\\
x_2 &\le 60\\
x_3 &\le 20\\
x_4 &\le 80\\
x_5 &\le 30\\
x_6 &\le 40\\
x_1, x_2, x_3, x_4, x_5, x_6 &\ge 0
\end{aligned}$$

TABLE 8.5

Two-flight network management example solution

Number	ODF	Fare	Allocation	Revenue
1	San Francisco to Denver full fare	$150	30	$4,500
2	San Francisco to Denver discount	100	40	4,000
3	Denver to St. Louis full fare	120	20	2,400
4	Denver to St. Louis discount	80	70	5,600
5	San Francisco to St. Louis full fare	250	30	7,500
6	San Francisco to St. Louis discount	170	0	0
		TOTAL REVENUE		$24,000

Weaknesses of Linear Programming Approach:

- Working with a large linear program is often cumbersome and produces a large number of outputs (the ODF allocations) that can be difficult to interpret and implement.

- It is not immediately clear how to incorporate demand uncertainty. Planning for mean demand does not maximize expected revenue as seen in the single-leg case.
- The allotment approach breaks down when bookings do not arrive in strict fare order. For this reason, we need to extend the concept of nesting to booking requests that utilize multiple legs. The most common way to do this is *virtual nesting*.

Virtual Nesting:

- The first step in airline virtual nesting is to define a set of *buckets* on each flight leg. Each bucket represents a range of fare values.
- The second step is to map each ODF into a bucket on each of its legs based on an estimate of the ODF's value to the airline. The process of mapping ODFs into buckets on legs is called *indexing*.
- At the end of the indexing process, each bucket contains ODFs of similar value.
- We will use the convention that the lowest numbered bucket on each leg contains the highest-value ODFs.
- The buckets are nested so that bucket 1 has access to the entire capacity of the leg, bucket 2 has access to all the inventory except that protected for bucket 1, and the lowest bucket has access only to its own allocation.
- Each bucket has a booking limit and a protection level that is updated every time a booking is accepted or a cancellation occurs. This means we can apply all of the mechanics for leg-based revenue management to buckets.

Indexing:

- How should ODFs be mapped to their constituent legs—in other words, how should they be indexed?
- One immediate thought might be to map an ODF into buckets on each leg based on its total fares. That is, we would map the ODFs with the highest total fare into the highest bucket, those with slightly lower total fare into the second bucket, and so on for every leg.
- This is equivalent to the greedy heuristic for network management. However, as we saw before, the greedy heuristic fails when multiple resources in a network are constrained.
- Furthermore, it ignores the opportunity cost in determining how an ODF should be bucketed on one of its legs.

- Specifically, an ODF with a high total fare but a high opportunity cost on its other legs should potentially be bucketed below an ODF with lower total fare but no opportunity cost on its other legs.
- This type of bucketing can never be achieved by the greedy heuristic.
- The failure of the greedy heuristic suggests that we need a way of indexing ODFs that includes their opportunity costs.
- The best way to value an ODF for bucketing on a resource is on the basis of its total fare *minus the opportunity cost (if any) on other resources it uses*. This quantity, known as the *net leg fare*, is defined as follows:

Net leg fare for ODF *i* on leg *k* = Total fare for ODF *i* - Sum of opportunity costs on *all resources other than k* (if any) in ODF *i*

- The same calculation holds for hotels, where the resources are room nights, and for rental car companies, where the resources are rental days.
- Ex: A hotel has estimated that the opportunity cost for a room on Tuesday night, January 23, is $80 and the opportunity cost for Wednesday night, January 24, is $105. The hotel receives a booking request for two nights, arriving January 23 and departing January 25, with an associated total rate of $220. The net leg fare associated with that booking request is $220 - $105 = $115 for the night of January 23 and $220 - $80 = $140 for the night of January 24.
- Note that the net leg fare is calculated by subtracting the opportunity cost on *other resources* from the total fare. Only if the fare associated with the ODF exceeds the opportunity cost for all of the resources it uses should it be accepted.

Example 8.5: Consider the three-city network in Figure 8.6. The opportunity cost on the SFO-DIA leg is $250 and on the DIA-STL leg is $800. Each leg has three buckets, defined as follows:

Bucket 1: Net leg fares $\geq$ $1,000
Bucket 2: $600 $\leq$ Net leg fares $<$ $1,000
Bucket 3: Net leg fares $<$ $600

TABLE 8.7

Net local fares and index for six ODFs in the network from Figure 8.6

Product	Class	Total fare	FLIGHT 1 NLF	FLIGHT 1 Bucket	FLIGHT 2 NLF	FLIGHT 2 Bucket
SFO-DIA	M	$400	$400	3	—	—
SFO-DIA	Y	$650	$650	2	—	—
DIA-STL	M	$550	—	—	$550	3
DIA-STL	Y	$750	—	—	$750	2
SFO-STL	M	$1,300	$500	3	$1,050	1
SFO-STL	Y	$1,800	$1,000	1	$1,550	1

The opportunity cost on flight 1 is $250, and the opportunity cost on flight 2 is $800.

Note that SFO-STL M-Class passengers are mapped into the highest bucket on flight 2 but the lowest bucket on flight 1. This means that on Flight 1 we will be protecting seats for local M-Class passengers from the SFO-STL M-Class passengers; but on flight 2 the situation is reversed.

Virtual Nesting:

Ex: An airline operating the simple network shown in Figure 8.6 offers a Y-Class fare and an M-Class fare for each product. This is a total of six ODFs. The airline has established three buckets on each leg. Then one possible indexing is illustrated in Figure 8.7. In this case the airline maps SFO-STL Y-Class passengers into bucket 1 on flight 1 but into bucket 2 on flight 2. This means that the SFO-STL Y passengers will have access to the entire capacity of flight 1, but they only have access to the bucket 2 allocation of flight 2, which is less than the capacity of the plane.

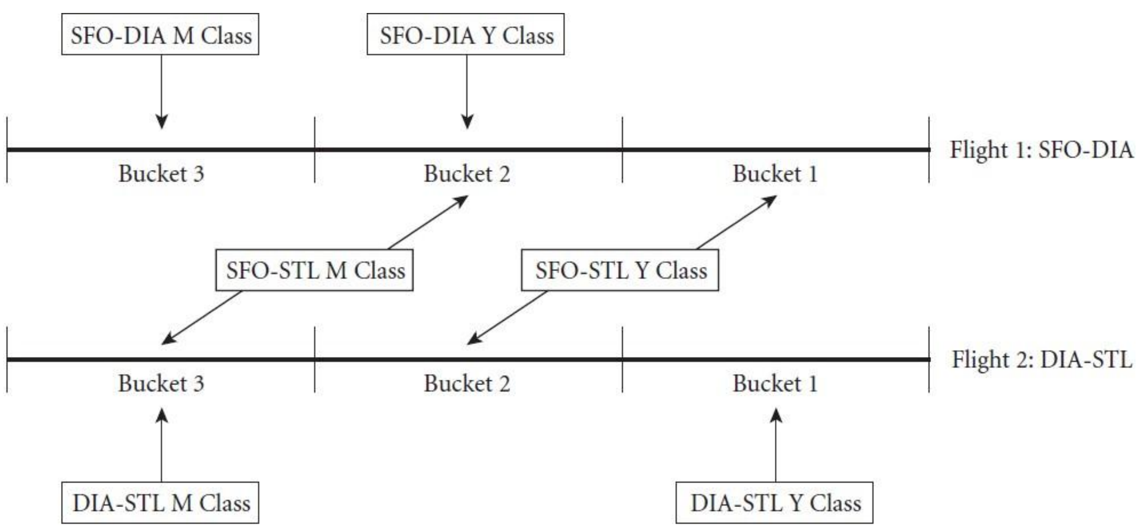

Figure 8.7 Example mappings of ODFs for a two-flight network.

Virtual Nesting Booking Control

1. Initialization

- Establish an indexing that maps every ODF into a bucket on each leg in the ODF.
- For each leg bucket, estimate the mean and standard deviation of total forecasted demand. Using the mean and standard deviation of demand by bucket, average contribution by bucket, and leg capacity; use EMSR or a similar approach to determine booking limits and protection levels for each bucket on each leg.

2. Operation

- When a booking request for an ODF is received, check each leg in the ODF. If there is sufficient capacity in the corresponding bucket on each leg, accept the request. Otherwise reject it.
- If the booking is accepted, reduce bucket availabilities on all legs in the booked ODF.
- When cancellations occur, increase bucket availabilities for each leg in the cancelled ODF.

3. Reoptimization

- Periodically reforecast demand for each ODF and update expected bucket demand by leg.
- Periodically rerun EMSR or other algorithm to recalculate nested booking limits for each bucket based on the new forecasts and capacity remaining on each leg.

Static Virtual Nesting:

- One obvious way to estimate the future opportunity cost for a product is to base it on past experience. This is, in essence, equivalent to forecasting opportunity cost.

Ex: A Tuesday 2:00 flight from San Francisco to Dallas had opportunity costs of $180, $130, $190, $205, and $95 at departure. One way to estimate the opportunity cost for next Tuesday's flight would be to average the past five opportunity costs; that is, the estimate of next Tuesday's opportunity cost would be ($180 + $130 + $190 + $205 + $95)/5 = $160.

- The method of using historic opportunity costs to map ODFs into buckets is called *static virtual nesting*.
- With static virtual nesting, an airline (or other revenue management company) estimates an opportunity cost λ_k for each resource based on the history of that resource.
- If the resource rarely sells out (e.g., as a Tuesday night in a midtown hotel in Detroit), the opportunity cost will be small. If, on the other hand, the resource sells out with regularity, its opportunity cost will be much higher.
- One approach to forecasting the future opportunity cost would be to average past opportunity costs;

- Another would be to calculate the fraction of time a resource sells out and to use the average opportunity cost of the resource when it sells out. Recall that the opportunity cost of a resource that does not sell out is $0.

Ex: Flight leg 135 has historically sold out 30% of the time, and its average opportunity cost when it sold out is $750. An estimate of the expected future opportunity cost for this leg would be 0.3 x $750 = $225.

- There are other approaches to static virtual nesting, but they all have one thing in common: Indexing is based entirely on past experience.
- This means that the indexing for a particular leg is unlikely to change during the booking period for a flight. In static virtual nesting, an airline might recalculate its indexing on a monthly or quarterly basis.
- Another shortcoming of static virtual nesting is implicit in the name: Since it is “static,” it does not respond to changes in the course of the booking process. An airline’s initial forecast of the opportunity cost for a leg might change considerably between first booking and departure.
- One approach to forecasting the future opportunity cost would be to average past opportunity costs;

- Another would be to calculate the fraction of time a resource sells out and to use the average opportunity cost of the resource when it sells out. Recall that the opportunity cost of a resource that does not sell out is $0.

  Ex: Flight leg 135 has historically sold out 30% of the time, and its average opportunity cost when it sold out is $750. An estimate of the expected future opportunity cost for this leg would be 0.3 x $750 = $225.

- There are other approaches to static virtual nesting, but they all have one thing in common: Indexing is based entirely on past experience.
- This means that the indexing for a particular leg is unlikely to change during the booking period for a flight. In static virtual nesting, an airline might recalculate its indexing on a monthly or quarterly basis.
- Another shortcoming of static virtual nesting is implicit in the name: Since it is “static,” it does not respond to changes in the course of the booking process. An airline’s initial

forecast of the opportunity cost for a leg might change considerably between first booking and departure.

- One approach to forecasting the future opportunity cost would be to average past opportunity costs;
- Another would be to calculate the fraction of time a resource sells out and to use the average opportunity cost of the resource when it sells out. Recall that the opportunity cost of a resource that does not sell out is $0.

  Ex: Flight leg 135 has historically sold out 30% of the time, and its average opportunity cost when it sold out is $750. An estimate of the expected future opportunity cost for this leg would be 0.3 x $750 = $225.
- There are other approaches to static virtual nesting, but they all have one thing in common: Indexing is based entirely on past experience.
- This means that the indexing for a particular leg is unlikely to change during the booking period for a flight. In static virtual nesting, an airline might recalculate its indexing on a monthly or quarterly basis.
- Another shortcoming of static virtual nesting is implicit in the name: Since it is “static,” it does not respond to changes in the course of the booking process. An airline’s initial forecast of the opportunity cost for a leg might change considerably between first booking and departure.

Network Bid Pricing:

- If we update opportunity costs often enough, they can almost serve as the only controls we need for network management. This is the idea behind bid pricing.
- We saw that we should accept a booking request for a single resource only if the fare associated with the request exceeded the bid price for that resource.
- There is an equivalent condition for products consisting of multiple resources: We should only accept a multiple-resource booking if the fare exceeds the sum of the bid prices for the constituent resources. This condition can be written as follows:

Accept a booking for an ODF only if its fare is greater than the sum of the bid prices on the constituent resource.

- The idea behind network bid pricing is to use this condition as the basis of the revenue management system.

The revenue management system calculates a bid price for each leg at each time before departure. When a booking request is received, its fare is compared to the sum of the bid prices for its constituent legs.

- In the extreme, this approach would allow us to dispense with the need for any sophisticated booking control structures and to manage bookings on bid price alone.
- However, we also need to allow for the fact that not all booking requests will be for a single seat.
- The easiest way to manage multiseat bookings is to keep track of the remaining unbooked seats on each leg and to use that number as a total booking limit.
- Without no-shows and cancellations, the total booking limit would be equal to the unbooked capacity.
- Managing bookings in this fashion is called *network bid price control*.

Network Bid Price Control

- 1. Calculate an initial bid price and a total booking limit for each leg.
- 2. When a booking request is received, calculate the *product bid price* as the sum of the bid prices for all the legs in the product and the *product availability* as the *minimum* of the total booking limits on all the legs used by the product.
- 3. If the fare for the booking exceeds the product bid price *and* the number of seats requested is less than the product availability, accept the request and go to step 4. Otherwise reject the request and go to step 2.
- 4. Decrease the total booking limit on each resource within the product by the number of seats in the booking and go to step 2.
- 5. When a booking cancels, increment the total booking limits for all flights in the booking.
- 6. Periodically re-optimize all flight booking limits and all bid prices.

Ex: Consider Table 8.6. If we define $75 as the bid price for the Denver-to–St. Louis leg and $125 as the bid price for the San Francisco–to-Denver leg, the bid price for the San Francisco–to–St. Louis will be $75+$125=$200. We have established a set of bid prices consistent with the optimal solution. By using the bid prices for the two legs plus a total booking limit for each leg in the dynamic bid price control structure we could achieve a result very close to the optimal network solution in Table 8.6. And we have done so by using just four total controls—two bid prices plus two booking limits—instead of the 15 allocations in Table 8.6.

TABLE 8.6
Expanded network solution

Number	Product	Class	Fare	Demand	Allocation	Revenue
1	San Francisco–Denver	A	$180	10	10	$1,800
2		B	160	20	20	3,200
3		C	140	15	15	2,100
4		D	130	18	18	2,340
5		E	100	27	0	0
6	Denver–St. Louis	A	$130	15	15	$1,950
7		B	110	20	20	2,200
8		C	90	15	15	1,350
9		D	80	20	20	1,600
10		E	75	30	13	975
11	San Francisco–St. Louis	A	$260	20	20	$5,200
12		B	240	10	10	2,400
13		C	200	10	7	1,400
14		D	190	15	0	0
15		E	170	15	0	0
				TOTAL REVENUE		$26,515

- Network bid price control is a remarkably simple approach to revenue management. It requires only two controls per leg—a bid price and a total booking limit.
- However, we have not yet described how bid prices can be calculated.
- Furthermore, bid prices need to be updated each time a booking is accepted or a cancellation occurs.
- If calculating network bid prices is time consuming or complex (as it is), the airline will require an interim control scheme to handle bookings between reoptimizations.
- In addition, network bid price control would require additional controls if multiple-seat booking requests are to be handled correctly.
- Finally, any network management approach needs to be implemented within the existing CRS infrastructure, and existing reservation systems are not designed to easily enable pure bid price management.
- Network bid price control is a remarkably simple approach to revenue management. It requires only two controls per leg—a bid price and a total booking limit.
- However, we have not yet described how bid prices can be calculated.

- Furthermore, bid prices need to be updated each time a booking is accepted or a cancellation occurs.
- If calculating network bid prices is time consuming or complex (as it is), the airline will require an interim control scheme to handle bookings between reoptimizations.
- In addition, network bid price control would require additional controls if multiple-seat booking requests are to be handled correctly.
- Finally, any network management approach needs to be implemented within the existing CRS infrastructure, and existing reservation systems are not designed to easily enable pure bid price management.
- In the single leg case, the bid price for a resource was equal to the opportunity cost for that resource and could be calculated as the fare of the lowest open fare class.
- Network bid prices are also related to the opportunity costs of the resources in the network. We can extend this observation to a network by noting that the following six quantities are (approximately) equivalent.
  - *Bid price*—the minimum price we should accept for a customer on a leg
  - *Opportunity cost*—the increased revenue we would see from an additional seat on a leg
  - *Displacement cost*—the revenue we would lose if we had one seat less on a leg
  - Marginal value of the capacity constraint in the network linear program
  - Boundary between the highest closed and lowest open fare classes in a single-leg problem
  - *Closed bucket boundary*—boundary between the highest closed and lowest open buckets on a leg in a virtual nest
- To be sure, these six quantities will not always be exactly equal. For example, the revenue gained from an additional seat on a leg will not necessarily be the same as the revenue lost from removing a seat, etc.
- Calculating bid prices can be quite complex—developing faster and more accurate bid-pricing algorithms is a very active area of research.

1. Calculating bid prices by linear programming:

- One way to calculate the bid price on a leg would be to solve the network linear program twice: the first time with the actual capacities and the second time with the capacity of the leg reduced by one seat.
- The difference in the total contribution between the two runs would be the displacement cost on the leg, which is an estimate of the bid price on the leg.
- Deterministic displacement costs for all the legs in the network can be more practically estimated using linear programming in two different ways by using the marginal values associated with the capacity constraints as bid prices.
- The marginal values associated with the capacity constraints of the deterministic network management linear program suffer from the same problem as the LP formulation itself—they ignore uncertainty.

2. Calculating bid prices by sequential estimation:

- The method of *sequential estimation* is based on the observation that the bid price on a leg should be equal to its *closed bucket boundary* in a virtual nest, where the closed bucket boundary is defined as the boundary between the lowest open bucket and the highest closed bucket on the leg.
- It should be clear that the closed bucket boundary should approximate the bid price on a leg. After all, we will accept a local booking for a single seat only if its fare exceeds the closed bucket boundary, which is the definition of bid price.
- Furthermore, we will accept a multi-leg booking request for a single seat only if its net leg fare is greater than the closed bucket boundary on each leg in the ODF.
- You can confirm that this is the same as ensuring that its total net fare exceeds the sum of the closed bucket boundaries.
- We can calculate bid prices for a network by finding a set of closed bucket boundaries that is consistent across the network.
- The challenge arises from the fact that the closed bucket boundary (bid price) on any leg depends on the closed bucket boundaries on all other legs via the net local fares of multi-leg ODFs.
- If we change the bid price on a leg we will change the net local fares for multi-leg ODFs on all connecting legs. This will in turn change the buckets into which those ODFs are mapped on the other legs.

- When we re-optimize availabilities on those legs, the new mapping may well change the bid prices on those legs, which will change the ODF mappings on the original leg.
- A consistent set of bid prices is one in which *the closed bucket boundaries on all legs are locally optimal given the closed bucket boundaries on all other legs*.
- Sequential estimation calculates bid prices by starting with an initial estimate of all bid prices and then updating the bid price on each leg until the bucket boundaries are consistent across the entire network.
- Assume we have a network consisting of two legs. We initially map all the ODFs into buckets assuming that each leg has a bid price of 0 (recall that this is equivalent to the "greedy heuristic").
- Given this indexing, we estimate the total demand by bucket on each leg.
- Then, for leg 1, we use EMSR to calculate the bucket allocations on leg 1 given the current indexing. Once we have calculated the allocations on leg 1, we re-estimate the bid price on leg 1 as its closed bucket boundary.
- Given this new estimate of the leg 1 bid price, we can recalculate the net leg fares on leg 2 and re-bucket all of the ODFs on leg 2 accordingly. We can then apply EMSR on leg 2, determine its closed bucket boundary, and use that as the new bid price for leg 2.
- Given this new leg 2 bid price, we can re-bucket the ODFs on leg 1 and continue. Under the right conditions, this procedure will converge to a situation where the mappings and allocations on each leg are consistent with each other.

Sequential Estimation Algorithm to Calculate Bid Prices:

- 1. Establish a set of buckets on all legs, and calculate probabilistic demand forecasts for all buckets. Set all leg bid prices to 0.
- 2. Loop over all legs, $k = 1, 2, \ldots, N$.
- 3. For the current leg k and every ODF i that includes leg k, calculate the net local fare as the net fare for ODF i minus the sum of the current bid prices for all other legs in the ODF.
- 4. Map all the ODFs on the current leg into buckets based on their net local fares. Once this is complete, use EMSR to determine allocations by bucket.

- 5. If all buckets are open, set the new bid price for leg *k* to 0. If all buckets are closed, set the new bid price for leg *k* to the highest bucket boundary. Otherwise, set the new bid price to the boundary between the lowest open bucket and the highest closed bucket.
- 6. Continue to the next leg until all legs have had new bid prices calculated.
- 7. When all leg bid prices have been recalculated, check the change between the new bid price and the old bid price on each leg. If the bid prices are sufficiently close for all legs, stop, the current set of bid prices is optimal.
- 8. Otherwise go to step 2.

Strengths and Weaknesses of Bid Pricing:

- As opportunity costs, bid prices specify the value of an additional unit for each resource in the network. If a flight leg has a bid price of $250, this indicates that the airline could realize (approximately) $250 in additional revenue from adding another seat to that flight.
- Since the bid prices are marginal signals, this logic cannot be extended to determine the effect of adding more than one seat: A bid price of $250 on a flight leg does not mean that an airline could expect to gain $2,500 from adding 10 seats.
- However, even with this limitation, the bid prices can provide very useful input into capacity decisions.
- A flight leg that consistently has a high opportunity cost at departure would be an excellent candidate to be considered for assignment to a larger aircraft, while a flight leg that consistently departs with a bid price of 0 is an excellent candidate to be considered for downsizing.
- A rental car company with the option to move its cars around should consider moving cars from locations with consistently low bid prices to those with consistently high bid prices.
- Bid pricing is particularly well suited to hotel applications where there is one bid price for each room type for each future night. The product bid price is the sum of the bid prices for the nights stayed.
- Hotel and rental car bid prices can be conveniently displayed using a calendar, as shown in Figure 8.10.

The calendar shows the bid prices (labeled b) and the forecast unconstrained occupancies (labeled d) for a hotel as they might be displayed by a revenue management system for a future month.

March						
Mon	Tue	Wed	Thur	Fri	Sat	Sun
28 $d = 85$ $b = \$84.34$	1 $d = 93$ $b = \$92.07$	2 $d = 112$ $b = \$153.12$	3 $d = 108$ $b = \$112.34$	4 $d = 99$ $b = \$92.57$	5 $d = 65$ $b = \$54.30$	6 $d = 80$ $b = \$62.33$
7 $d = 91$ $b = \$88.47$	8 $d = 102$ $b = \$122.00$	9 $d = 135$ $b = \$172.15$	10 $d = 120$ $b = \$142.34$	11 $d = 92$ $b = \$95.67$	12 $d = 53$ $b = \$42.34$	13 $d = 44$ $b = \$32.34$
14 $d = 67$ $b = \$54.37$	15 $d = 85$ $b = \$72.48$	16 $d = 110$ $b = \$122.47$	17 $d = 97$ $b = \$99.97$	18 $d = 93$ $b = \$92.34$	19 $d = 72$ $b = \$55.18$	20 $d = 66$ $b = \$54.54$
21 $d = 86$ $b = \$89.11$	22 $d = 104$ $b = \$130.02$	23 $d = 157$ $b = \$199.93$	24 $d = 140$ $b = \$178.25$	25 $d = 122$ $b = \$122.20$	26 $d = 95$ $b = \$100.69$	27 $d = 85$ $b = \$85.18$
28 $d = 84$ $b = \$84.33$	29 $d = 92$ $b = \$93.44$	30 $d = 114$ $b = \$155.67$	31 $d = 100$ $b = \$101.01$	1 $d = 82$ $b = \$78.77$	2 $d = 60$ $b = \$53.92$	3 $d = 75$ $b = \$62.74$

Figure 8.10 Bid price calendar.

- A customer wishing to rent a room for the nights of March 23, 24, and 25 would need to pay at least the sum of the bid prices for these nights ($199.93 + $178.25 + $122.20 = $500.38) in order to be allowed to book.
- A customer wishing to rent for March 12, 13, and 14 would be allowed to book if he were paying more than $42.34 + $32.34 +$54.37 = $129.05.
- The product bid price is the *minimum* price we should accept for a product given remaining capacity and anticipated future demand by fare class.
- It does not tell us how we should actually be *pricing* the product.
- A hotel manager looking at the bid price calendar in Figure 8.10 should not make the mistake of thinking that he should set the rack rate on Wednesday, March 16, to $122.47.
- In this sense the term *bid price* is unfortunate, since, as we have seen, the bid price is better regarded as an opportunity *cost*.
- In the extreme case, a bid price of zero for a product tells us nothing about how we should price the product. Bid price control is a mechanism for ensuring that we do not

accept any business whose margin does not exceed its opportunity cost—it does not tell us whether we are priced correctly in the market.

Dynamic Virtual Nesting:

- *Dynamic virtual nesting* is the marriage of bid pricing and virtual nesting. It combines the conceptual elegance and intuitiveness of bid pricing with the practical orientation of virtual nesting.
- Under dynamic virtual nesting, new bid prices are frequently recalculated for all legs based on current bookings and forecasts of future demands.
- The recalculation of bid prices could be purely periodic (e.g., nightly or weekly), it could be event driven (e.g., recalculated whenever a flight closes), or it could be ordered by a revenue analyst.
- Every time new resource bid prices are calculated, new net leg fares are calculated for each leg in the network. These net leg fares are used to define a new indexing. This indexing is put in place until the next recalculation of bid prices.
- The strength of dynamic virtual nesting is that it updates bid prices to reflect the current situation on each leg, while utilizing virtual nesting to take advantage of leg-based control structures.
- If bid prices are not recalculated very often, then dynamic virtual nesting is not much of an improvement over static virtual nesting. If bid prices are continually updated, then dynamic virtual nesting will approximate network bid price control for single-seat bookings.