Introduction to pricing and revenue optimization

Brian Kallehauge

42134 Advanced Topics in Operations Research

Fall 2009

Revenue Management Session 02

DTU Management Engineering
Department of Management Engineering

Outline

- Introduction to pricing and revenue optimization
  - The challenge of pricing
  - Traditional approaches to pricing
  - The scope of pricing and revenue optimization
  - The pricing and revenue optimization process

Introduction to pricing and revenue optimization

- Pricing and revenue optimization is a process for managing and updating pricing decisions in a consistent and effective way.
- The goal is to find a set of prices which maximizes total expected profit given certain constraints such as business goals and/or limitations of e.g. capacity.
- Most companies know the list prices for their products but the prices customers actually pay for the products are often unclear.
- Selling prices are influenced by e.g. discounts, adjustments, rebates, etc.
- There may be big differences between product list prices and *pocket prices* (actual selling prices).
- While list prices are generic, pocket prices may vary from customer to customer, dependent on which pricer reductions he or she gets.

Price waterfall (by McKinsey and Company)

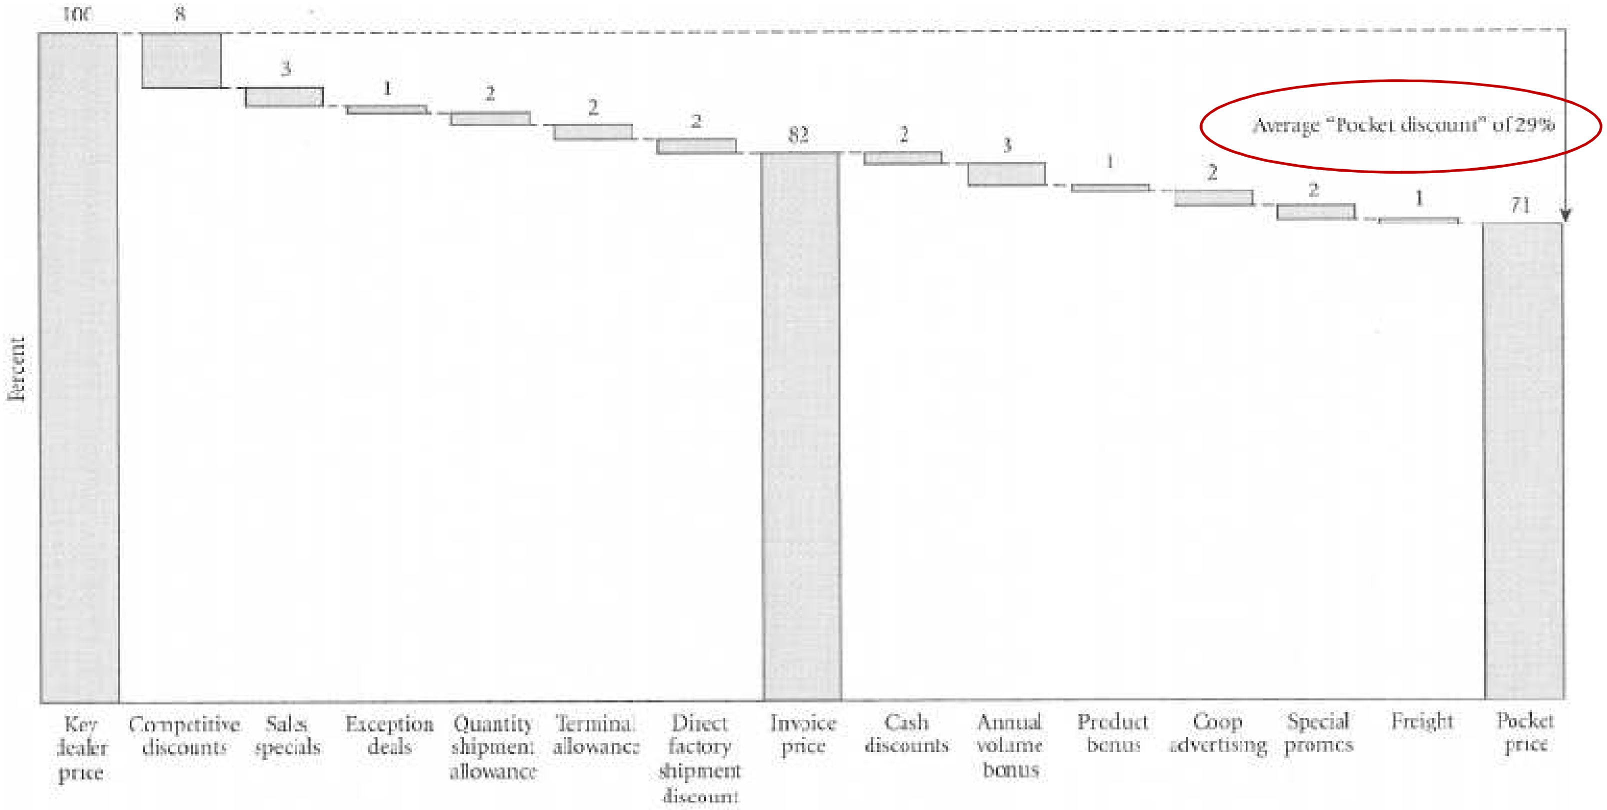

Price waterfall for a consumer package goods (CPG) company.

Source: Mike Reopel of A. T. Kearney.

Disorganized pricing

- Each pocket price is rarely the result of a single decision but rater the cumulative result of several decisions – often made by different parts of the company without measuring or tracking the individual decisions.
  - I.e. no specific person or section is in charge of the pricing and no one is responsible for the differentiation in selling prices.

- Due to disorganized pricing organizations there is a rather low probability that a particular procket price maximizes customer profitability.

- Even worse: sophisticated and experienced buyers who are aware of the price waterfall of the selling company quickly learn how to ”divide and conquer” to obtain the largest possible price reductions.
  - Several buying agents call different parts of the selling company to bargain for extra discounts.

- Even though very few percent of a company’s goods are sold without several discounts, the management focuses on the list prices and, therefore, analyze the business and make strategies on a wrong basis.

Pocket price distribution for a CPG company

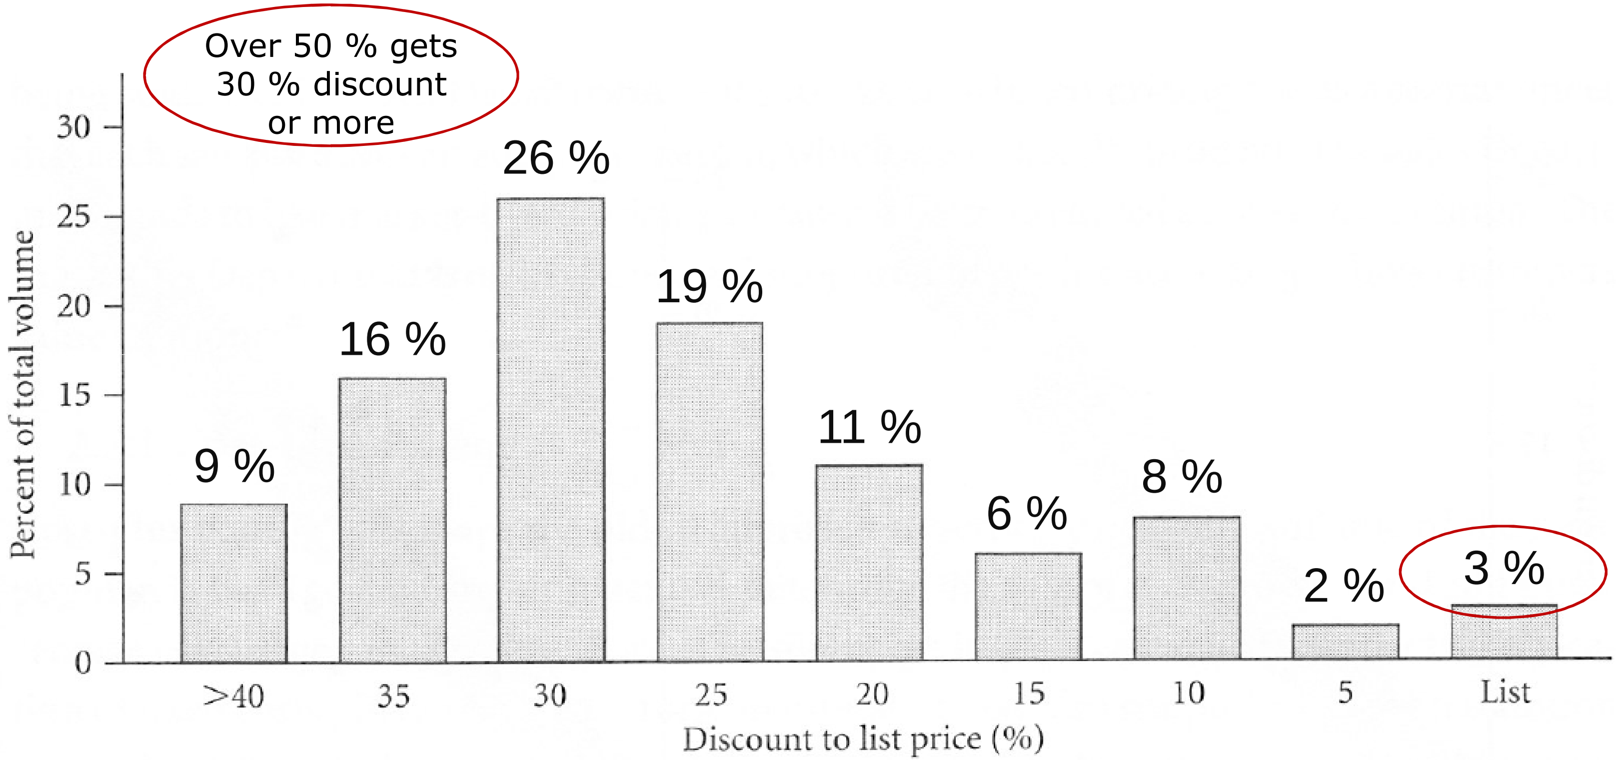

Pocket price distribution for a consumer package goods (CPG) company.
Source: Mike Reopel of A. T. Kearney.

Pricing and revenue optimization

- Since only 3 % of the buyers pay the list price, this price is not important to set optimally.

- Rather, the pricing and revenue optimization (PRO) decision is to set the list price high and use discounts to target prices to individual customers.

- The wide pocket price distribution does not tell anything about the quality of the pricing decisions – sophisticated PRO may very well result in a wide price distribution.

The important thing is to control the pocket prices in a concious way based on sound analysis rather than letting the price distribution be a result of an arbitrary process.

- The quality of the PRO decisions is measured by the extent to which the pocket prices correlate with customer characteristics which indicate price sensitivity, i.e. higher discounts to larger, more price sensitive, customers.

Correlation of discount with customer size

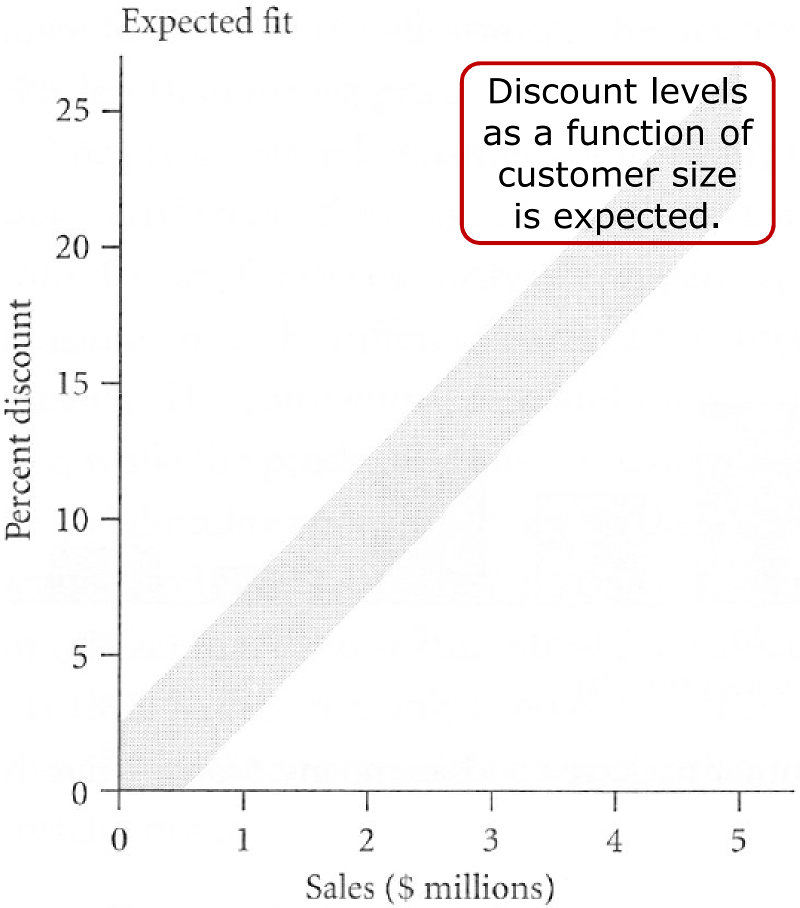

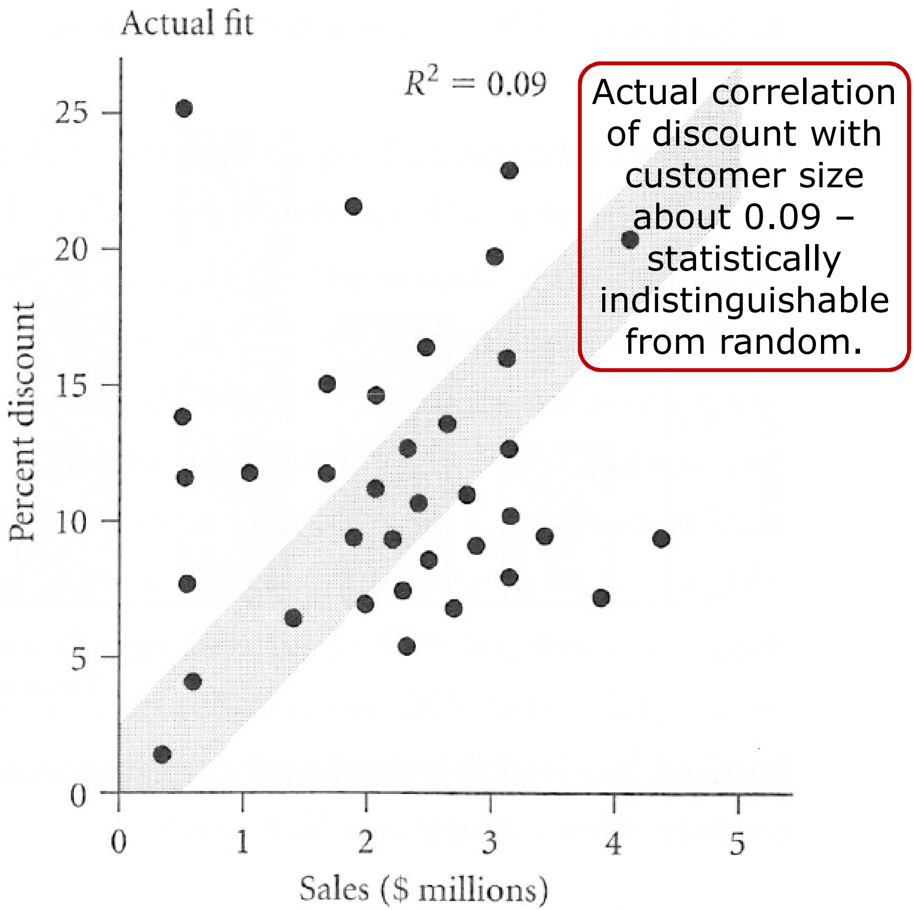

Correlation of discount with customer size for a CPG company.

Source: Mike Reopel of A. T. Kearney.

Traditional approaches to pricing

- Pricing and revenue optimization (PRO) is applied to determine prices that maximizes the expected net contribution, incorporating:
  - costs,
  - the competitive environment , and
  - customer demand (or *willingness to pay*).

Alternative approaches to pricing

Approach	Based on	Ignores	Liked by
Cost-plus	Costs	Competition, customers	Finance
Market based	Competition	Cost, customers	Sales
Value based	Customers	Cost, competition	Marketing

- *Cost-plus pricing* calculates prices based on costs plus a standard margin.
- *Market-based pricing* determines prices based on competitors' actions.
- *Value-based pricing* sets prices based on estimates of how customers value the product.

Cost-plus pricing

- Cost-plus pricing is preferred by Finance departments since it guarantees that each sale produces an adequate profit margin.

- Cost-plus pricing is perhaps the oldest approach to price control and still one of the most popular.

The technique is very simple: determine the cost of the product and add a certain percentage to determine the price.

- The surchage often reflects a fixed cost plus a required return on capital. E.g. in the restaurant industry:
  - food is marked up three times,
  - beer is marked up four times, and
  - liquor is marked up six times.

- Cost-plus pricing seems very appropriate as the company is guaranteed to gain the profit margin if all competitors have similar cost structures.

Disadvantages of cost-plus pricing

- The major drawback of cost-plus pricing is the complete ignoration of the market.

- Customers' willingness to pay is not considered when calculating prices and all buyers are charged the same price.
  - Recall that setting different prices for different customer segments is the core of pricing and revenue optimization (PRO).

- Furthermore, the assumption that the cost prices can be retrieved rarely holds, i.e. the selling prices are calculated based on a very weak foundation and only on historical data, not forecasts of future product costs.

- In spite of its obvious disadvantages, cost-plus pricing is still very widely used, perhaps due it the simplicity of the approach.

Market-based pricing

- Market-based pricing is preferred by Sales departments since it helps them sell against competition.

The approach bases prices solely on the prices being offered by competing companies.

  - Market-based pricing is often applied by smaller players in situations where there is a clear market leader to compare prices to.
  - There is no alternative to market-based pricing in markets where competitors sell completely identical products, and where transaction prices are transparent and known by the buyers.
  - Upcoming low-cost suppliers may also apply market-based pricing in order to enter a new market by ensuring lower prices than the current players.

- Though being very approapriate in the above situations, market-based pricing does not capitalize on customers who value the company's products or brand and does not make use of price differentation.

Value-based pricing

- Value-based pricing is preferred by Marketing departments since they focus on customer behaviour and how customers value a product.

The approach focuses on customer value as the key driver of the price.

- Historically, value-based pricing usually referred to the use of methodologies such as customer surveys, focus groups, and conjoint analysis to estimate how customers value a product relative to the alternatives – which is then used to determine price.
  - This type of value-based pricing is frequently employed when introducing new products in a market.

- Value-based pricing is a very suitable and promising approach *if*:
  - the company is a monopoly,
  - customer values of the products can actually be determined,
  - customers are willing to pay the calculated price, and
  - there is no threat of regulation, cannibalization, etc.

Disadvantages of value-based pricing

- Essentially, value-based pricing is impossible!
  - Determining each individual customer's value for a product at the point of sale is not possible to do.
  - There is a constant possibiliy of cannibalization, arbitrage, etc.
  - Competitive pressure means that companies almost always have to price lower that they would like to any group of customers.

- The main drawback of value-based pricing is that the approach ignores competition.
  - There is a great difference between the value that a potential buyer might place on a product in isolation and what the customer is actually willing to pay in the market.
  - A customer may value the product highly but he or she has alternatives.

Pricing in practice

- Cost-plus pricing, market-based pricing, and value-based pricing are puristic approaches, each focusing narrowly on one or a few parameters and leaving out others.

- In practice, companies apply combinations of different pricing approaches and modify their pricing strategies from time to time to achieve different goals.

- Applying different approaches rather than being strictly devoted to one specific pricing method is better for the company.
  - For example letting Marketing use value-based pricing when introducing a new product, letting Sales apply market-based pricing when seeking to increase the market share, etc.

Pricing and revenue optimization

- Pricing and revenue optimization (PRO) provides a consistent approach to pricing decisions across the organization.
- This means that a company needs to have a clear view of all the prices it is setting in the marketplace and the ways in which those prices are set.
- The goal of PRO is to provide the right price,
  - for every product,
  - to every customer segment,
  - through every channel,

  and to update those prices over time in response to changing market conditions.
- Each combination of these three dimensions – *product*, *customer type*, and *channel* – has an associated price. For example:
  - turbines sold to large customers in the north via direct sales channel,
  - replacement gears sold to small companies via online sales.

Dimensions of the pricing and revenue optimization cube

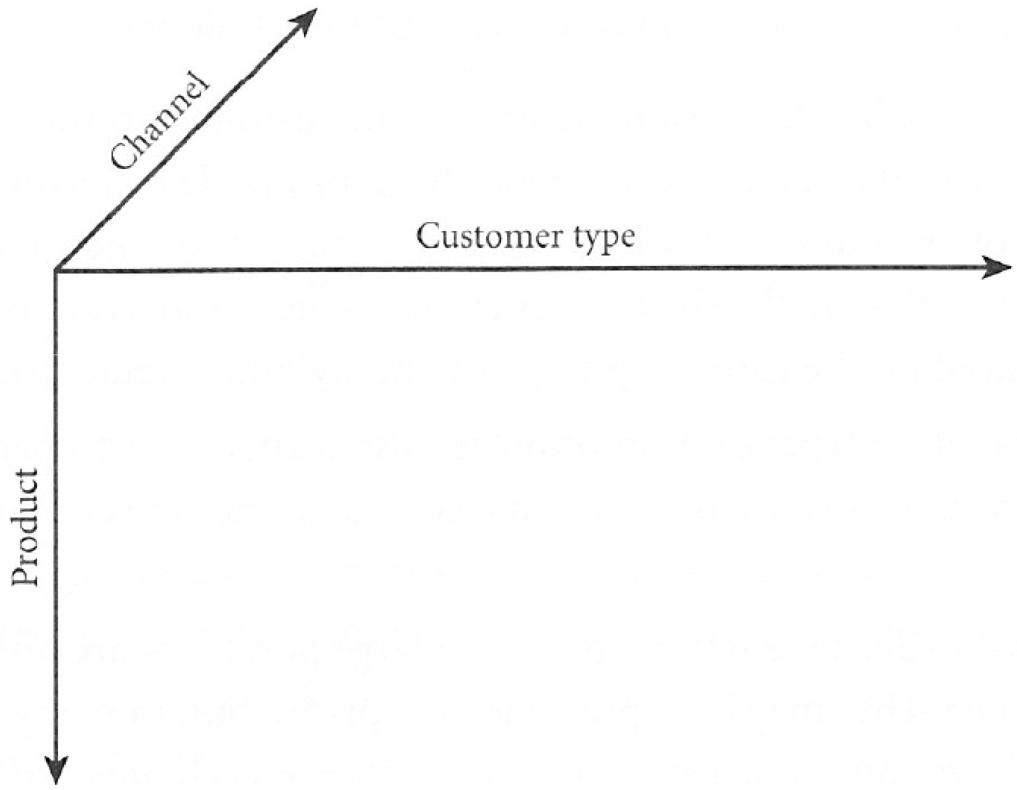

Dimensions of the PRO cube.

- In theory, each cell within the PRO cube could correspond to a different price.
- In practice, some combinations of product, channel, and customer type is not meaningful.
- Furthermore, our ability to charge different prices through different channels may be constrained by practical considerations or by strategic goals.

The pricing and revenue optimization cube in practice

- The PRO cupe is a useful starting point for a company seeking to understand the magnitude of the pricing challenge that it faces.

- Enumerating the combinations of products, market segments, and channels gives a rough estimate of the total number of prices a company needs to manage.

- In practice, a successful PRO process includes such an overview of the different prices associated with the different *product-market-channel* combinations but also strategic goals such as e.g. charging small customers a lower price in online sales to encourage these to purchase online rather than through direct sales.

Customer commitments

- A core concept in PRO is the idea of customer commitment, occurring when a seller agrees to provide a customer with products or services, now or in the future, at a price.

- The elements of a customer commitment are:
  - the products and services being offered,
  - the price,
  - the time period over which the commitment will be delivered,
  - the time for which the offered commitment is valid (how long the customer has to make up his/her mind),
  - other element of the contract or transaction (e.g. payment terms or return policy),
  - firmness of the commitment and risk sharing.

- The forms and types of commitments that sellers make (and buyers expect) vary from industry to industry.
  - E.g. airline tickets are often refundable whereas theater tickets are not.

The pricing and revenue optimization process

- Successful pricing and revenue optimization (PRO) involves two components:
  1. A consistent business process focused on pricing as a critical set of decisions.
  2. The software and analytical capabilities required to support the process.

- There has been a focus on mathematical analysis in PRO but in order to provide sustainable improvements, quantitative analysis needs to be embedded in the right process, including eight activities:
  – Four activities are part of *operational PRO*. These are executed every time the company needs to change prices.
  – Four activities are part of *supporting PRO*. These occur at longer intervals.

- PRO is a closed-loop process, i.e. feedback from the market must be incorporated into both the operational activities and the more periodic activity of updating market response curves.

The pricing and revenue optimization process

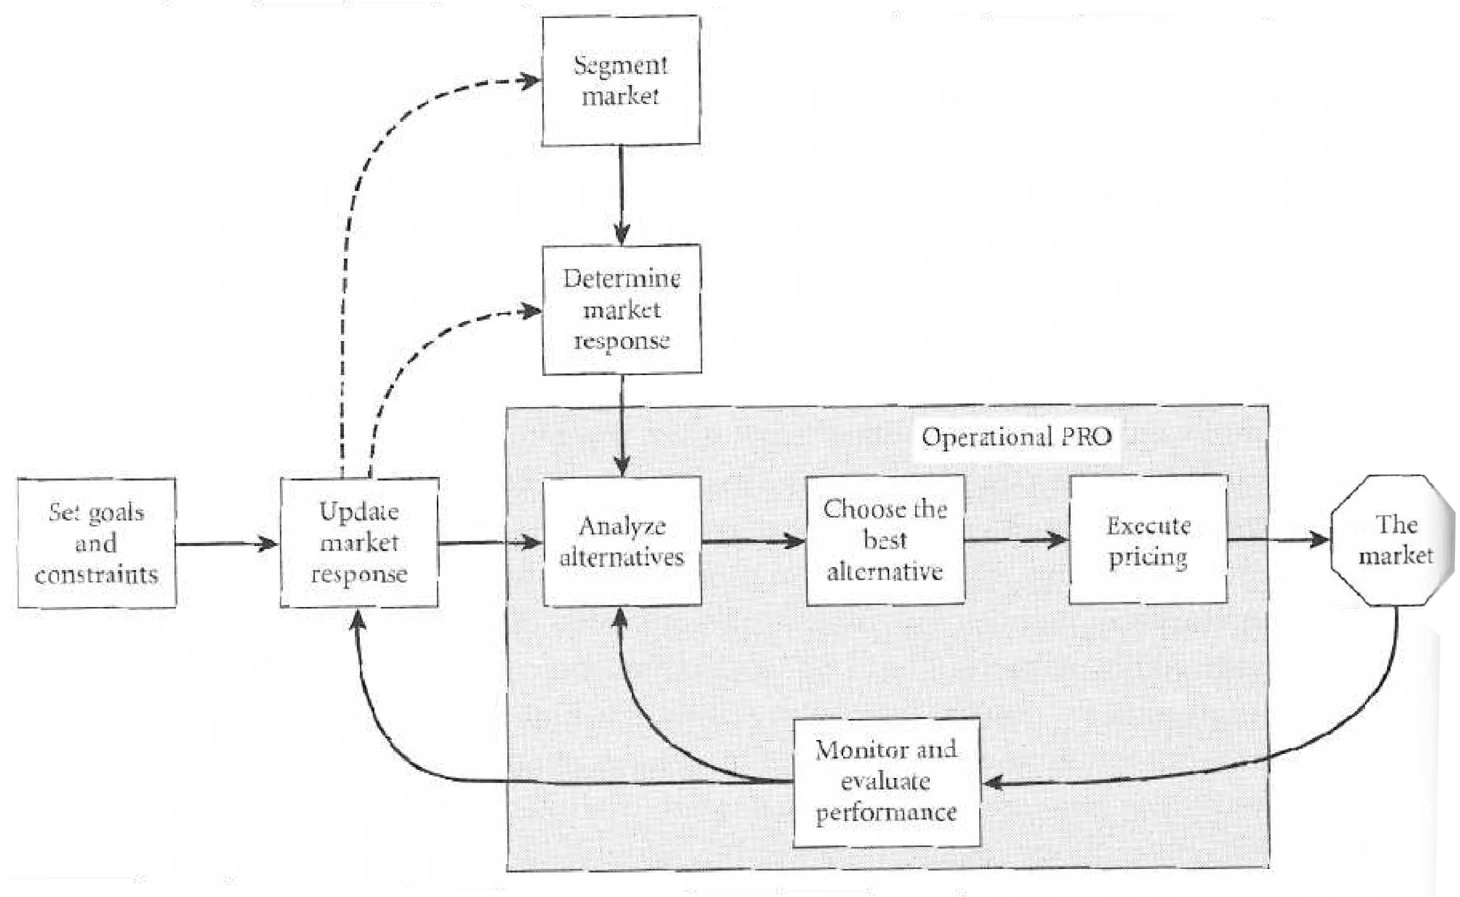

The closed-loop process of pricing and revenue optimization.

Operational PRO activities

- The operational PRO activities work continuously to set and update prices.
  - Some companies, such as airlines, change fares from moment to moment and others, e.g. production companies, update their prices weekly or monthly.

- For all types of companies, the operational PRO activities are similar:
  - **Analyze alternatives.** E.g. analysts comparing pricing alternatives under different scenarios using spreadsheets or using optimization systems to recommend new prices for each element of the PRO cube.
  - **Choose the best alternative.** Human beings evaluating the price alternatives provided by software.
  - **Execute pricing.** Communication of the calculated prices to the market. E.g. opening and closing fare classes in the airline or hotel industry, construction of a ”pricing matrix” database wich specifies which prices are available to which customers through which channels.
  - **Monitor and evaluate performance.** Comparison of expections with results from the marketplace and evaluation of the overall performance against the company goals.

Supporting PRO activities

- The supporting PRO activities provide key inputs to the operational PRO activities.

- Supporting activities occur on a much longer time frame – e.g. every quarter or year – than the operational ones. They include:
  - **Set goals and business rules.** Key initial step of specifying the overall goals of the process; typically changed no more often than quarterly. In general, the goal of PRO is to maximize expected total profit. Other temporary goals may be to increase market share or attain a volume sales goal. The business goals determine the form of the objective function when optimizing the PRO cube entries.
  - **Segment the market.** Defining customer segments in order to maximize the opportunity to extract profit; typically performed anually to reflect changes in the underlying market.
  - **Determine price response.** Calculation of the price-response function for each of the market segments; often updated weekly or monthly.
  - **Update price response.** Updating model parameters if performance differs significantly from the expectations.

Pricing and revenue optimization as a business process

- PRO should be treated like any other business process.
  - The company needs to identify clearly what it is trying to achieve, the constraints it is facing, and the alternatives available.
  - Based on an understanding of the market and the constraints, the alternative most likely to achieve the goals is chosen.
  - When this alternative is implemented, results must be monitored and measured against the expectations.
  - The understanding of the market must be updated in order to make better decisions in the future.

The time dimension

- Previously, prices changed e.g. once a quater whereas today, prices change weekly, daily, or even hour by hour.
  - More frequent price changes makes it less feasible to apply in-depth analysis to each pricing decision.

- This makes the situation more challenging for the companies and players who are able to change prices rapidly in response to changing market conditions will gain in advantage.

- Pricing has different cadence in every industry, e.g.:
  - gasoline prices fluctuate more or less randomly,
  - fashion goods are priced high at the beginning of the season and then subsequently marked down as the season progresses, whereas
  - passenger airlines have successfully segmented their customer base into early-booking leisure passengers and later-booking business customers, enabling the company to increase ticket prices as departure approaches.

The role of optimization

- To cope with the complexity of adjusting prices very frequently, most companies consider adopting computerized pricing support systems.
- Treating pricing decisions as constrained optimization problems is at the hart of PRO.
- The formulation and solution of these constrained optimization problems draws on techniques from statistics, operations research, and management science.
- However, though optimization is very important in PRO, no company actually *optimizes* its prices since determining the *optimal* prices is, in general, impossible or at least well beyond our current ability to model and solve.
- Therefore, optimization approaches make assumptions and exclude some features of the real-world problem to solve.
  ...good prices on time are far better than perfect prices late

Group arrivals, dynamic models, and customer-choice behavior

Brian Kallehauge

42134 Advanced Topics in Operations Research

Fall 2009

Revenue Management Session 02

DTU Management Engineering
Department of Management Engineering

Outline

- 2.4 Group Arrivals
- 2.5 Dynamic Models
  - 2.5.1 Formulation and Structural Properties
    - 2.5.1.1 Dynamic Program
  - 2.5.2 Optimal Policy
    - 2.5.2.1 Computation
- 2.6 Customer-Choice Behavior
  - 2.6.1 Buy-Up Factors

Source: Talluri and van Ryzin (2004), chapter 2

Group Arrivals

- A group request is a single request for multiple units of capacity, e.g. a family of four traveling together
- Group arrivals can be handled as:
  - Partially accepted demand (simple case)
  - All-or-none demand (complex case)
- Basically when a large fraction of the demand is size one or two means group arrivals can be ignored
  - This is an assumption in most airline RM systems, hence group arrivals are ignored
- In some applications it is not possible to ignore group arrivals, e.g. container shipping

Group Arrivals - References

- Partially accepted case
  - T. C. Lee and M. Hersh. A model for dynamic airline seat inventory control with multiple seat bookings. Transportation Science, 27:252-265, 1993.

- All-or-none accepted case
  - D. Walczak and S. Brumelle. Dynamic airline revenue management with multiple semi-Markov demand. Operations Research, 51:137-148, 2003.
  - A. J. Kleywegt and J. D. Papastavrou. The dynamic and stochastic knapsack problem. Operations Research, 46:17-35, 1998.
  - R. Van Slyke and Y. Young. Finite horizon stochastic knapsacks with applications to yield management. Operations Research, 48:155-172, 2000.

Dynamic Models

- Recall that a *static model* is defined by the assumption that demand for classes arrives in a strict low-to-high revenue order
- Dynamic models allow for an arbitrary order of arrival
- However, dynamic models require the following assumptions:
  - Markovian (such as Poisson) arrivals, which puts restrictions on modeling demand variability (main drawback)
  - Estimate of the pattern of arrivals of bookings over time – the *booking curve (difficult to estimate)*
- The choice of dynamic versus static models is basically a choice of:
  - Which approximations is more acceptable
  - What data is available

Formulation of Dynamic Model

- *n* classes
- prices $p_1 \geq p_2 \geq ... \geq p_n$
- $t = 1,...,T$ time periods
  - 1 is the first time period - in contrast to static model where the stages (classes) run from *n* to 1
- One arrival occurs in each time period (by sufficiently fine time discretization)
- Probability of arrival of class *j* in time period *t* is $\lambda_j(t)$, i.e.

$$\sum_{j=1}^{n} \lambda_j(t) \leq 1$$

- We do not require lower classes to arrive earlier than higher classes

Dynamic Program #1

- x denote the remaining capacity
- $V_t(x)$ denote the value function in period t
- $R(t)$ is a random variable
  - $R(t) = p_j$ if a demand arrives in period t, 0 otherwise
  - Note that $P(R(t)=p_j) = \lambda_j(t)$
  - Let u = 1 if we accept the demand, 0 otherwise
- Maximize the sum of current revenue and the revenue to go:

$$R(t)u + V_{t+1}(x-u)$$

Dynamic Program #2a

The Bellman equation is therefore:

$$V_t(x) = E\left[\max_{u\in\{0,1\}}\left\{R(t)u + V_{t+1}(x-u)\right\}\right]$$

see next slide

$$= V_{t+1}(x) + E\left[\max_{u\in\{0,1\}}\left\{\left(R(t) - \Delta V_{t+1}(x)\right)u\right\}\right]$$

where $$\Delta V_{t+1}(x) = V_{t+1}(x) - \Delta V_{t+1}(x-1)$$

is the expected marginal value of capacity in period $t+1$

The boundary conditions are:
(take-off at time *T*)

$$V_{T+1}(x) = 0, \quad x = 0,1,\dots,C,$$
$$V_t(0) = 0, \quad t = 1,\dots,T.$$

Dynamic Program #2b

$$V_t(x) = E\left[\max_{u\in\{0,1\}}\left\{R(t)u + V_{t+1}(x-u)\right\}\right]$$

- For *u=0* (rejecting the demand):

$$\begin{aligned} V_t(x) &= E\left[R(t)\cdot 0 + V_{t+1}(x-0)\right] \\ &= V_{t+1}(x) \end{aligned}$$

- For *u=1* (accepting the demand):

$$\begin{aligned} V_t(x) &= E\left[R(t)\cdot 1 + V_{t+1}(x-1)\right] \\ &= E\left[R(t) + V_{t+1}(x-1)\right] \end{aligned}$$

$$\Downarrow$$

$$V_t(x) = V_{t+1}(x) + E\left[\max_{u\in\{0,1\}}\left\{\left(R(t) - \Delta V_{t+1}(x)\right)u\right\}\right]$$

$$\Delta V_{t+1}(x) = V_{t+1}(x) - V_{t+1}(x-1)$$

- For *u=0* (rejecting the demand):

$$\begin{aligned} V_t(x) &= V_{t+1}(x) + E\left[\left(R(t) - \Delta V_{t+1}(x)\right)\cdot 0\right] \\ &= V_{t+1}(x) \end{aligned}$$

- For *u=1* (accepting the demand):

$$\begin{aligned} V_t(x) &= V_{t+1}(x) + E\left[\left(R(t) - \Delta V_{t+1}(x)\right)\cdot 1\right] \\ &= V_{t+1}(x) + E\left[R(t) - \left(V_{t+1}(x) - V_{t+1}(x-1)\right)\right] \\ &= E\left[R(t) + V_{t+1}(x-1)\right] \end{aligned}$$

Optimal Policy #1

- If a class *j* request arrives, so that $R(t)=p_j$, then it is optimal to accept the request if and only if

$$p_j \geq \Delta V_{t+1}(x) = V_{t+1}(x) - V_{t+1}(x-1)$$

- Thus, the optimal control can be implemented using a bid-price control where the bid price is equal to the marginal value:

$$\pi_t(x) = \Delta V_{t+1}(x)$$

- Revenues that exceed this threshold are accepted; those that do not are rejected

Optimal Policy #2

- For the dynamic model the optimal control can be achieved using either:
  - Time-dependent nested protection levels

$$y_j^*(t) = \max\left\{ x : p_{j+1} < \Delta V_{t+1}(x) \right\}, \quad j = 1,2,\ldots,n-1$$

  - Time-dependent nested booking limits

$$b_j^*(t) = C - y_{j-1}^*(t), \quad j = 2,\ldots,n$$

  - Bid-price tables

$$\pi_t(x) = \Delta V_{t+1}(x)$$

- In practice, because the value function is not likely to change much over short periods of time, fixing the protection levels or booking levels computed by a dynamic model and updating them periodically is usually close to optimal

Optimal Policy #3

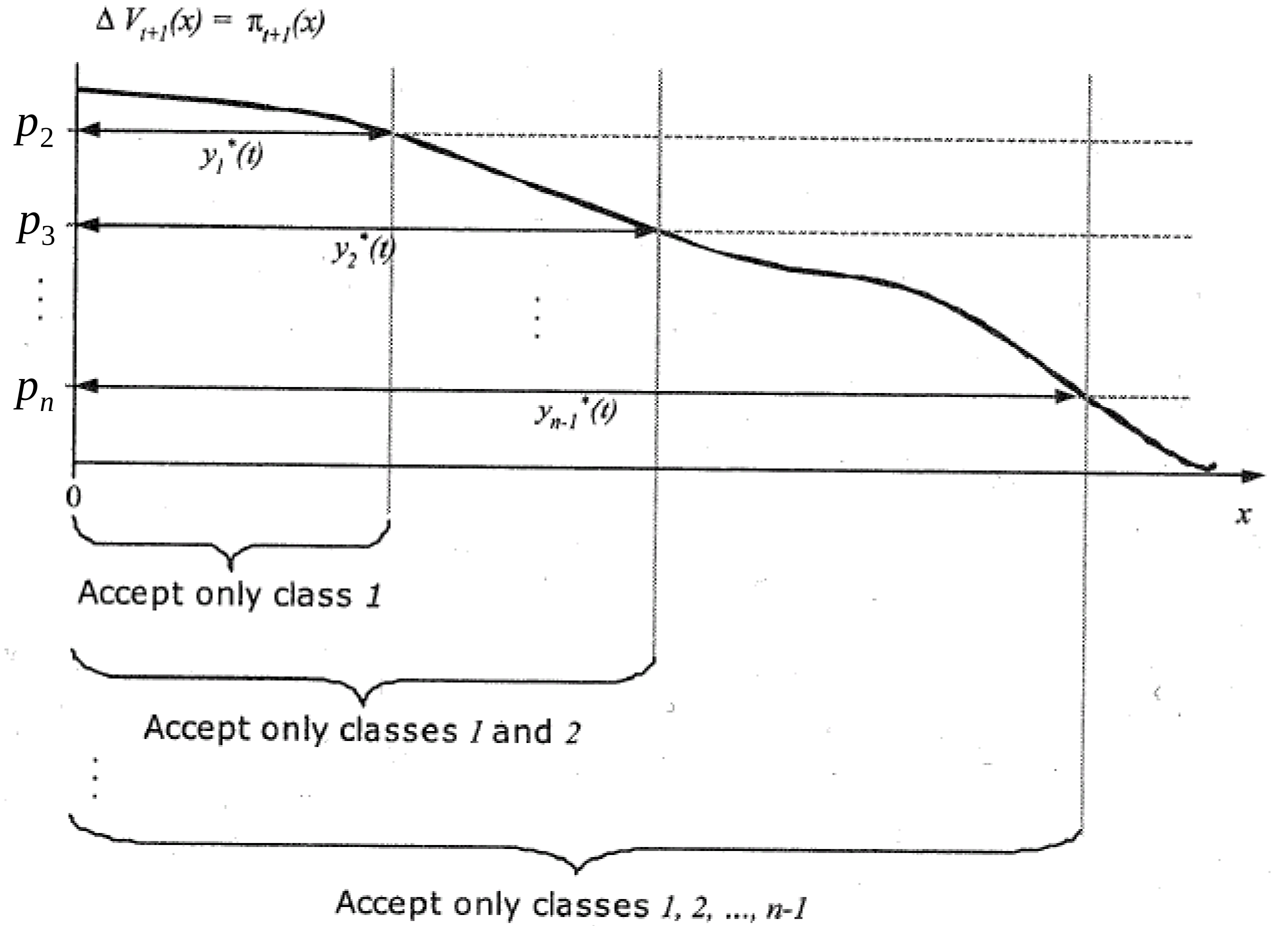

Figure 2.6. Optimal protection level $y_j^*(t)$ in the dynamic model.

Computation

- The dynamic model is solved using the recursion:

$$V_t(x) = V_{t+1}(x) + E\left[\left(R(t) - \Delta V_{t+1}(x)\right)^+\right]$$

$$= V_{t+1}(x) + \sum_{j=1}^{n} \lambda_j(j)(p_j - \Delta V_{t+1}(x))^+, \quad t = 1, \ldots, T$$

errata!

- Starting with the boundary condition:

$$V_{T+1}(x) = 0, \forall x$$

we proceed with the recursion backward in time *t*

Customer-Choice Behavior

- Key assumption in basic static model and dynamic model:

  Demand for each of the classes is completely independent of the capacity controls being applied by the seller

- That is, it is assumed that the likelihood of receiving a request for any given class does not depend on which other classes are available at the time of request

- This is an unrealistic assumption:
  - Availability of full fare tickets vs. discount tickets
  - Likelihood that customer buys may depend on the lowest available fare

- When a customer buys a higher fare when discounts are closed it is called *buy-up*

Buy-Up Factors – Two-Class Model #1

- Suppose there is a probability *q* that a class 2 customer will buy class 1 if class 2 is closed

- If a customer buy up to class 1 we earn a net benefit of:

$$p_1 - p_1 P(D_1 > x)$$

(the class 1 revenue minus the expected marginal cost)

- Thus, it is optimal to accept class 2 if

$$p_2 - p_1 P(D_1 > x) \geq q p_1 \left(1 - P(D_1 > x)\right)$$

- What happens when *q*=0, and *q*=1, respectively?

Buy-Up Factors – Two-Class Model #2

- Littlewood's rule with buy-up:

$$p_2 \geq (1-q)\,p_1 P\left(D_1 > x\right) + q p_1$$

- Note that the right-hand-side is strictly larger than the normal rule, which means that the modified rule is more likely to reject class 2 demand

- This is intuitive because with the possibility of customers upgrading we should be more eager to close class 2

Buy-Up Factors – EMSR-b

- EMSR-b protection levels with buy-up factor:

$$p_{j+1} = (1 - q_{j+1})\,\overline{p}_j P\left(S_j > y_j\right) + q_{j+1}\hat{p}_{j+1}$$

where q_{j+1} is the probability that a customer of class j+1 buys up to one of the class j, j-1,...,1.

$\hat{p}_{j+1} > p_{j+1}$ is an estimate of the average revenue received given that a class j +1 customer buys up to one of the classes j, j-1,...,1.

$\hat{p}_{j+1} = p_j$ if customers are assumed to buy up to the next class.

Buy-Up Factors – Criticism of EMSR-b approach

- EMSR-b protection levels with buy-up factor provides a simple heuristic way to incorporate choice behavior
- Ad-hoc adjustment to an already heuristic approach
- Serious difficulties in estimating the buy-up factors (often made-up, reasonably sounding numbers)
- Despite limitations the buy-up factors have proven useful in practice