MARKDOWN MANAGEMENT IN PRACTICE

Wolfgang Weidinger, Lisa Gimpl-Heersink, Jürgen Gimpl, Alfred Taudes

[1] Prolytic - Marketing Engineering, Consulting & Software GmbH, 1220 Vienna, Austria
[2] Vienna University of Economics and Business Administration, 1090 Vienna, Austria
[3] EBCONT systems & solutions, 1230 Vienna, Austria
[4] Vienna University of Economics and Business Administration, 1090 Vienna, Austria

wolfgang.weidinger@prolytic.com(Wolfgang Weidinger)

Abstract

Price is a key element for retailing because it communicates information about the brand, product, value proposition and overall strategy of the retailer. A permanent price reduction of an item over time is known by definition as 'markdown'. Markdown is a pricing strategy, found in nearly every retail store including apparel retail, consumer products, fashion style goods or just any products with a limited life cycle. It is often used to make up for buying errors resulting from demand unpredictability and to clear out stock over time. While discounts increase sales on the one hand they erode profits on the other hand. Thus optimizing the timing and magnitude of markdowns is crucial for liquidating a specified inventory quantity at the maximum profit over a set amount of time. Presently retail buyers do this balancing mainly based on intuitive experience and 'rules of thumb'.
In this work we use a mathematical model for markdown industries with the aim of maximizing total expected profits over the end of season sale. In order to provide a powerful DSS (decision support system) for retailers we have to understand the customer-demand patterns and price sensitivities and we have to forecast consumer demand. A major part of this procedure is parameter estimation for gaining valuable demand data. Once the selling season begins, there is an opportunity of revising prior demand estimates using actual sales data. In our model we implement an adaptive learning mechanism such that we can expect the estimates of demand to get better tuned as the season progresses. Based on these results we calculate an optimal markdown-pricing policy via a dynamic programming approach. Furthermore we give results of a case study which is conducted at a renowned Austrian fashion retailer.

Keywords: Retail Management, Markdown Pricing, Adaptive Learning, Dynamic Programming

Presenting Author's Biography

Wolfgang Weidinger. Master of Science in Technical Mathematics (Modelling and Simulation, Numerical Mathematics, OR) at Vienna University of Technology. Working experience in indemnity insurance and aviation industry. Currently at the R&D department of Prolytic - Marketing Engineering, Consulting & Software GmbH.
Research interests: Pricing models, Dynamic Programming, Adaptive learning methods.

1 Introduction

Today's retail environment is getting more and more complex and managers are faced with thousands of daily decisions. Some of them are small and routine but occur frequently, others are occasional but have greater financial impact. Nowadays technology can help support managers' decision-making process by ways of DSS (decision support systems). A DSS is a computerized system for helping make decisions which are a choice between alternatives based on estimates of the values of those alternatives. One way to obtain these recommendations is provided by optimization of mathematical models. The strength of mathematical optimization includes the ability to consistently evaluate far more alternatives than a human can. Especially when the same kind of decisions are repeated many times, such applications provide effective means.

The goal of almost any industry should be to maximize profits or revenues on the long-term. Here such optimization techniques come into play. Since revenues depend on prices and sales, pricing policies that are responsive to sales can be an efficient tool to enhance revenues. Customers will hardly be willing to buy a perishable product in retail stores whose price oscillates, from their point of view, randomly over the season. This explains the need for non increasing pricing strategies which are commonly used for short life cycle products. Here retailers start with a fairly high price in the hope of reaping large revenues if the item turns out to be a hot seller and knowing that they can lower it, if the item does not sell well. Such strategies where prices are continuously reduced over time are called permanent markdowns. The pricing decisions are made dynamically, taking into account the inventory on hand and time left until the end of the season. Permanent markdowns are rarely advertised, because each product is discounted at different points in time and in different amounts. Thus customers have little information about current prices before going to the store. This and the fact that they have only partial information about inventories, prevents them to some extent, from acting strategically. If however, customers would act strategically and adjust their buying behavior in response to the firm's pricing strategy, this would require a different type of model including a game theoretic formulation which is beyond the scope of our analysis. As motivated above the assumption of non-strategic customers is approximately true in 'impulse purchase' settings containing consumer products, apparel retail, fashion-style-goods or just any products with a limited life cycle like seasonal goods. Other examples that fit this framework could also be travel and leisure industry, which markets space such as cabins on vacation cruises, hotel rooms or seats on airline flights which become worthless if not sold by a specific time.

Many industries dealing with such settings face the problem of selling a given stock of items by a deadline with no resupply option during the sales season. This results from high production lead times compared to the shrinking products' life cycles and is a fairly common situation in practice, with seasonal products such as fashion clothing (e.g. production lead time of three to six months and a sales season of less than 12 weeks). Moreover in most industries capacity decisions are fixed for the sales horizon and cannot be changed in the short run. Buyers order inventories for the entire season well before the item goes on sale. Thus once the selling season is underway, there is no further opportunity of reordering. Even though products usually don't perish at the end of the season, stores usually liquidate the inventory and do not store merchandise for the next season due to rapid changes in fashion (obsolescing inventory) and high inventory carrying costs. The merchandise manager's job is to tactically adjust the price throughout the selling season in response to the realized demand, in order to maximize expected total revenues over the entire selling season.

Until recently, retailers typically based their initial pricing and subsequent markdown decisions on arbitrary, time-honored rules that they believed had worked well in the past. A common rules-based approach would be to apply a fixed percentage markup onto the cost (keystone markup) and then take a fixed percentage markdown on merchandise that has been in the store for a certain number of weeks, followed by an additional markdown a few weeks later. The magnitude of the markdown typically ranges between 25% and 50% with the exact rate depending on the sales rates and current inventories. Furthermore many retailers tend to offer the same list price and the same discounted price at different location which usually do not follow the same demand patterns. Another difficulty is the following: Since demand is price sensitive and stochastic, managers have to rely on uncertain demand forecasts. In practice retailers tend to buy stocks according to their optimistic forecasts while planning to mark down potential excess inventory to stimulate demand and sell out the excess by the end of the season. However, in order to hedge against the possibility of stronger than expected demand, they set relatively high initial prices which are then reduced over a set amount of time. If demand falls short of projections, this results in too much unsold and practically worthless inventory at the end of the season. Thus he needs to apply necessary markdowns to clear out the stocks as profitably as possible. If he marks down too less and/ or too late, this could lead to unsold stock at the end of the season causing heavy losses. On the other hand giving large discounts too early could exhaust stock at unnecessarily low profitability early in the season with massive opportunity losses. Getting the initial pricing, and timing and depth of discounts right, is a challenging job. Most retailers will agree that markdowns are one of the most inefficiently run areas of their business. Given todays small margins, the effectiveness of markdown policies can make the difference between a profitable and unprofitable season. Decision support systems, like ours, can help taking these crucial decisions.

1.1 Literature Review

Several recent papers address the problem of marking down style goods. Theoretical models have been proposed in marketing, operations research and economic

literature. [1] provides a survey of pricing models in marketing literature, and [2] provides a survey of pricing models in the economic literature.
[3] formulate a basic model for the continuous and discrete time case where demand is described by a Poisson process or is deterministic and provide structural results for the optimal price path as a function of inventory level and the remaining length of the selling period. They furthermore give an exact solution for an exponential price-dependent demand function.[4] use a similar approach to determine the optimal timing and duration of a single price (markup or markdown) change and show conditions where the optimal policy possesses a threshold property. [5] also develop a continuous-time model in which customers arrive according to a Poisson process, but they describe the price sensitivity by a reservation price function. They also compare optimal continuous pricing policies with more realistic periodic pricing policies. [6] determine an optimal clearance pricing policy which takes into account the reduced stock of inventory which influences demand.
Several authors, e.g. [7], [8], [9], [10] and [11] study retail pricing decisions working with data sets on sales and prices of specific retailers empirically. [10] furthermore explain how hypothesized models are fit to the actual data in order to obtain estimates of revenues under various pricing policies.
Furthermore [12] propose shrinkage estimation procedures to estimate separate elasticities for different chain-brand combinations by using a hierarchical Bayes model. More fundamental is [13] which deal with the linear model which is adapted to Bayesian methods. You can find an enhancement of this paper by reading [14]. It examines the Bayesian model in more detail and deals with the resulting estimators. [15] proposes a Bayesian Mixture Model for demand estimation where retail demand was parametrized as semilog and doublelog. [16] review the use of a Gibbs sampler as a method for calculating Bayesian posterior densities and illustrate it with a range of normal data models. A decent overview of Bayesian methods is provided by [17]. Recently several authors examine the nexus between retail practice and research. E.g. [18] discuss recent advances in retail pricing optimization and lists critical components which should be incorporated in order to determine optimal prices.

1.2 Outline and Overview

This paper reports the results of an empirical study of pricing policies at an Austrian fashion apparel retailer.
Till nowadays the pricing policy of this firm consists of essentially two price changes for the analyzed categories. These happen at the end of each season whereas during each season no markdowns are allowed. The amount of reduction is determined by rules of thumb, weather, remaining days of sale and so on. The same decision rules are applied to the timing of all price changes and their quantity.
Most garments are available at a limited supply and are sold over a short selling season. The inventory stock is reviewed periodically and at each review period the manager has the opportunity to implement a markdown from the current retail price. Using data provided by this company we will analyse the company's current pricing policies and analyse the loss of expected profit for different pricing policies. In order to provide a powerful DSS, we first need to understand the customer-demand patterns and the corresponding price sensitivities. Adaptive learning techniques, like bayesian updating, will be used to revise prior demand estimates using actual sales data and get better tuned demand forecasts as the season progresses. A large section will be assigned to the formulation of the underlying dynamic program which numerically solves the actual price optimization problem by taking the estimation parameters as input variables.

2 The Model

We assume our firm operates in a market with imperfect competition. Either the firm is a monopolist or it's product is new and innovative. In that case the firm holds a temporary monopoly. Furthermore the firm applies periodic reviewing, and markdown pricing decisions are to be taken over a finite selling horizon with T periods. We assume that stochastic demands are price-sensitive (their distributions depend on the product price) and independent from each other in different periods. The period's pricing decisions (selecting a specific price from a given permissible set) are made at the beginning of each period, before the period's random demand is realized, depending on how much inventory is left from the last period x_t and the previous period's price p_{t-1}. Randomness in demand does not depend on the price. We define demand additively as $D_t(p_t, \epsilon_t) = E[D_t(p_t, \epsilon_t)] + \epsilon_t$, where ϵ_t has a probability mass function $P_{\epsilon_t}[\cdot]$ with mean $E[\epsilon_t] = 0$ or multiplicatively as $D_t(p_t, \epsilon_t) = E[D_t(p_t, \epsilon_t)]\epsilon_t$, where ϵ_t has a probability mass function $P_{\epsilon_t}[\cdot]$ with mean $E[\epsilon_t] = 1$. Demand is indexed by t to denote time dependence. Revenues are collected at the end of each time period as the stock is sold. x_t denotes the inventory level at the beginning of time period t which is left over from the last time-period, respectively x_0 denotes the starting inventory. At the end of each period t a cost h_t is incurred which represents the inventory holding costs per unit of leftover salable inventory $x_t > 0$. No backlogging of demand is allowed and an excess of demand is lost, but penalty costs of l_t per unit of excess of demand arise, accounting for the loss of goodwill of the customers. Unsold items at the end of the selling season have a given salvage value s_T per unit, which could be obtained by selling the remaining stock to a deep-discount retailer. All costs related to the purchase or production of those items are considered sunk costs. The objective is to find a pricing sequence subject to the constraint that $p_t \leq p_{t-1}$ (markdown pricing policy) which maximizes expected (discounted) profit over the finite horizon:

$$\max_{p_t \leq p_{t-1}} E[\sum_{t=0}^{T-1} \gamma^t p_t D_t(p_t, \epsilon_t)],$$

Tab. 1 Variables, Relationships, Objective

Quantity	Formulation	Description
State Variables	x_t p_{t-1}	inventory on hand at the beginning of period t previous period's price
Decision Variable	$p_t(p_{t-1}, x_t)$	optimal price decision in the current period t
Transition Function	$x_{t+1} = x_t -$ $- \min(x_t, D_{a,t})$	subsequent period's inventory with current period's actual demand $D_{a,t}$
Objective Function	$\max_{p_t \le p_{t-1}} \Pi$	maximize total expected profit Π over the selling horizon

where γ is a possible discounting factor $(1 + \rho/100)^{-1}$ with ρ being the discounting rate. According to Bellman this maximization problem can be written as a dynamic program which we are going to describe during the next paragraphs.

A dynamic program essentially consists of a system of relationships among variables and time. The system is composed of a vector of state variables (which define the state of the system at any point in time), a vector of decision variables (which affect the system's evolution), a system of transition functions (which relate the state and decision variables over time), the objective (which reflects the goals of the management) and an external disturbance on the system. The decision problem is one of choosing optimal time paths for the set of decision variables from a permissible set of possible decision paths.
In case of the here considered markdown optimization problem Tab.1 specifies the variables of interest and the relationship between them. Fig. 1 gives an idea of how the dynamic program works.

Suppose, we are at the beginning of some period t, $t = 0 \dots T-1$, and suppose p_{t-1} was the price applied in the previous period, resulting in an actual (observed) demand $D_{a,t-1}$. Note that from now on we use the subscript a to denote realized values of random variables. Then the gross quantity of stock on hand at the beginning of period t is given by $x_{a,t} = x_{a,t-1} - \min(x_{a,t-1}, D_{a,t-1})$ (see Tab.1) which is the inventory on hand, $x_{a,t-1}$ at the start of the previous period less the total quantity sold during that period.
Next we consider the single period profit $\pi_t(x_t, p_t)$ is given by the revenues from all units sold less the inventory holding and inventory stockout (penalty) costs in this period. Thus the single period profit is given by the following equation:

$$\begin{aligned} \pi_t(x_t, p_t) = p_t \cdot \min(x_t, D_{a,t}) - \\ - h_t \cdot \max((x_t - D_{a,t}), 0) - \\ - l_t \cdot \max((D_{a,t} - x_t), 0) \end{aligned} \quad (1)$$

In the terminal period $t = T - 1$ there is a possibility that the leftover stock has some salvage value s_T at the end of the season, as it could perhaps be returned at

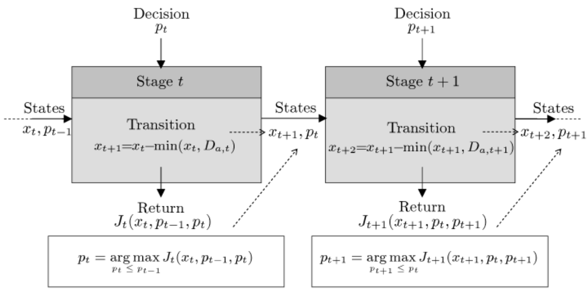

Fig. 1 Dynamic program

some buyback price to the vendor or sold at a very low price (at or even below cost) in a liquidation sale.

We are now going to formulate the buyer's dynamic pricing problem. We define the function $V_t^*(x, p_{t-1})$ as the maximum expected discounted profit from period t onwards (profit-to-go function), if the initial inventory is x and the price in the previous period was p_{t-1}. The mathematic formulation is given by

$$V_t^*(x, p_{t-1}) = \max_{p \le p_{t-1}} \{J_t(x, p_{t-1}, p)\} \quad (2)$$

$$\begin{aligned} J_t(x, p_{t-1}, p) = pE\left[D_t(p, \epsilon_t)\right] - \\ -h_t E\left[x_t - D_t(p_t, \epsilon_t) | x_t \ge D_t(p_t, \epsilon_t)\right] - \\ -l_t E\left[D_t(p_t, \epsilon_t) - x_t | x_t < D_t(p_t, \epsilon_t)\right] + \\ +\gamma E\left[V_{t+1}^*(x - D_t(p, \epsilon_t), p)\right], \end{aligned} \quad (3)$$

with the boundary conditions

$$V_t(0, p) = 0, \ \forall t, p \text{ and} \quad (4)$$
$$V_T^*(x, p) = x s_T, \ \forall x, p. \quad (5)$$

Since only integer values of demand can be realised we can reformulate Eq.(3) substituting expected values in the following way:

$$\begin{aligned} J_t(x, p_{t-1}, p) = \sum_{D_{a,t}=0}^{\infty} p \cdot \min(x, D_{a,t}) + \\ - \sum_{D_{a,t}=0}^{\infty} h_t \cdot \max((x_t - D_{a,t}), 0) - \\ - \sum_{D_{a,t}=0}^{\infty} l_t \cdot \max((D_{a,t} - x_t), 0) + \\ +\gamma \sum_{D_{a,t}=0}^{\infty} V_{t+1}^*(x - \min(x, D_{a,t}), p) P_{D_t}\left[D_t(p, \epsilon_t) = D_{a,t}\right], \end{aligned} \quad (6)$$

where $P_{D_t}[\cdot]$ is the probability mass function of the number of actual sales in period t.

ISBN 978-3-901608-32-2

3 Case study

This case study was initiated when the CEO of a big Austrian fashion retailer decided that their markdown strategy needed a fundamental redesign. After a short presentation it was clear that we would start a pilot project with the purpose to price a few categories each with 50-60 items during the following end of season sale and to evaluate the results. When we talked about the precise implementation of our software it turned out that our model of demand was a bit too sophisticated for this firm. So we did not use parameters like inventory holding and penalty costs as well as salvage values for our demand estimation. Our literature review suggested that the additive and multiplicative are two popular demand specifications for retail demand. After a lot of reading (especially [19] and [20]) and finding a lot of evidence that the linear specification is a bad choice for a decision model we decided that we will use the latter. One of the big advantages of multiplicative demand is the possibility of non-linear relations between the different variables.

$$D_{i,t}(P_{i,t}, \epsilon_{i,t}) = \beta_0 P_{i,t}^{\beta_1} \epsilon_{i,t} \tag{7}$$

In Eq.(7) t is the time index and i specifies the category, whereas β_0 and β_1 are the respective coefficients. The random disturbance terms $\epsilon_{i,t}$ are log-normally distributed due to earlier work conducted in [11].The resulting probability distribution for a certain category is depicted below.

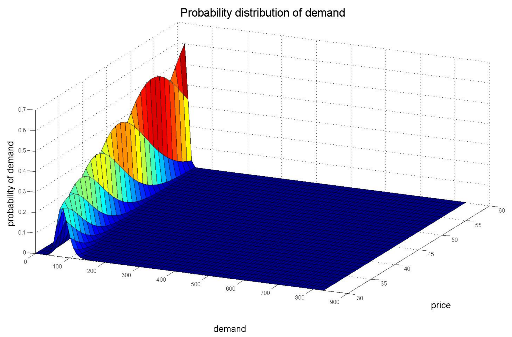

Fig. 2 Probability demand

For estimation of the coefficients we used a Bayesian approach to include current sales figures. In doing so methods like MCMC (Markov Chain Monte Carlo) and the Gibbs sampler were applied. In the process a system of m regression equations are related through correlated error terms:

$$\begin{aligned} y_i &= X_i \beta_i + \epsilon_i, \\ \epsilon_i &\sim N(0, \Sigma) \end{aligned} \tag{8}$$

MCMC is based on the idea that rather than compute a probability density $p(\theta|y)$ we would be satisfied to have a large random sample from $p(\theta|y)$ instead of knowing the precise form of the density. This gives rise to the question of how to simulate a large number of random samples from $p(\theta|y)$. One of the most used approaches for this task is the mentioned Gibbs sampling.

After the estimation of our demand function was finished Dynamic Programming methodology was used for optimization of the different price paths. So the first thing we did was testing if our optimization behaves like it should and one of the first results is depicted below.

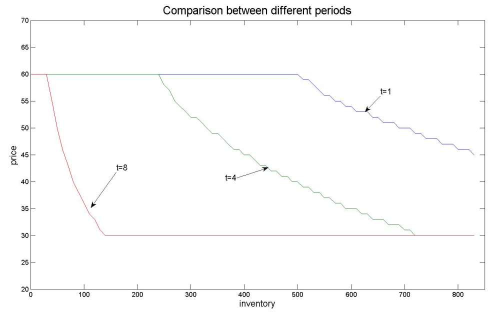

Fig. 3 Comparison between different periods

In Fig.3 these items are priced accordingly to their stock of inventory. So as stock decreases, prices for these items increase. Each line describes different behaviour at different time periods. So as time progresses our algorithm tends to drop prices more sharply as you can see by comparing $t = 1$ and $t = 8$. These characteristics are completely sound because as time progresses there's less time to clear the inventory which is the ultimate goal (while maximizing profit).

But as this paper is about markdown there has to be a third dimension in the following graphs namely the price at the previous period because $p_t \le p_{t-1}$. So let's have a short look what happens if we don't act accordingly to this rule (graph below).

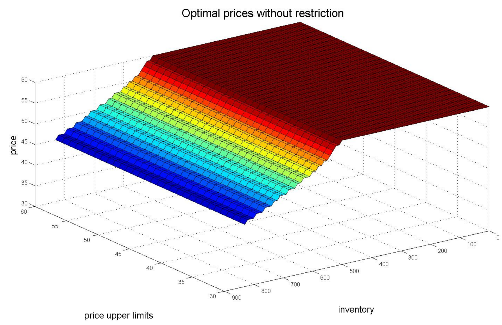

Fig. 4 Optimal prices without restriction

The new dimension is named 'price upper limits' and consists of the same numbers as our permissible prices. As is clearly visible despite the fact that eg. price 30 was applied at the previous period our program suggests that the new price should be approx. 45. That's because this new price is mathematically optimal but you don't consider prices which were set in the previous period. This was done for the creation of Fig. 5. As you can see all prices which are 'optimal' but violate the assumption $p_t \le p_{t-1}$ are 'cut'.

ISBN 978-3-901608-32-2

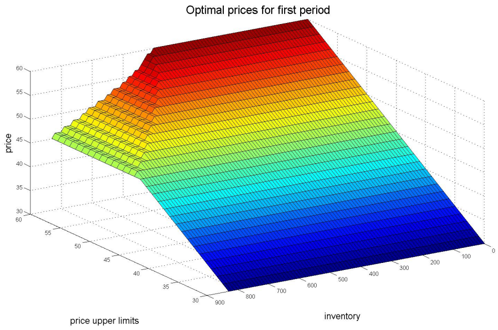

Fig. 5 Optimal prices for the first period

The figure above depicts all different price suggestions for a certain category for one period of time. So you get such a graph for each time period which behaves like Fig. 3 over time with incorporation of the rule $p_t \leq p_{t-1}$.

780	I N V E N	50,00	48,33	46,67	45,00	43,33	41,67	40,00	38,33
790		50,00	48,33	46,67	45,00	43,33	41,67	40,00	38,33
800		50,00	48,33	46,67	45,00	43,33	41,67	40,00	38,33
810		50,00	48,33	46,67	45,00	43,33	41,67	40,00	38,33
820		50,00	48,33	46,67	45,00	43,33	41,67	40,00	38,33
830		50,00	48,33	46,67	45,00	43,33	41,67	40,00	38,33
		Price reduction of the previous period%							
		50,00	**48,333**	**46,667**	**45,00**	**43,33**	**41,67**	**40,00**	**38,33**

Fig. 6 Excel output

As you can't expect a category manager to look at tons of graphs like Fig. 5 the output for this retailer consists of simple Excel-files like Fig. 6.

As an additional task we had a look at the expected loss of profit for different price politics. Fig. 7 illustrates the various effects which could happen when you allow only one price reduction (eg. -60%) during this sales period.

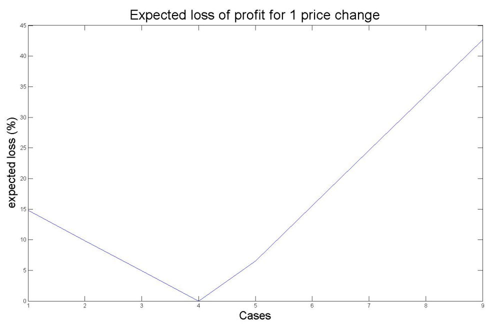

Fig. 7 Expected loss of profit for 1 price change

As one can easily see you'll have no expected loss of profit if you act accordingly to Case 4. What is meant by the different cases is demonstrated in Fig. 8 - 11.

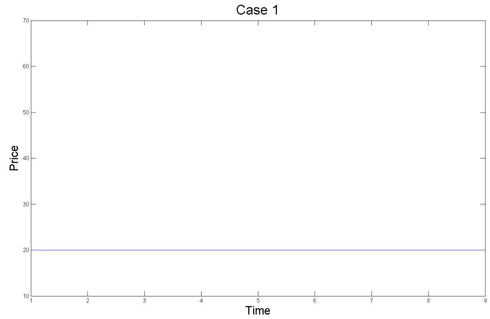

Fig. 8 Case 1 (for 1 permissible price change)

If you're acting accordingly to Case 1 you should markdown at the first time period. Looking at Fig. 7 this policy brings you an expected loss of profit of approx. 15%! This is because you start too early and don't get the full potential profit.

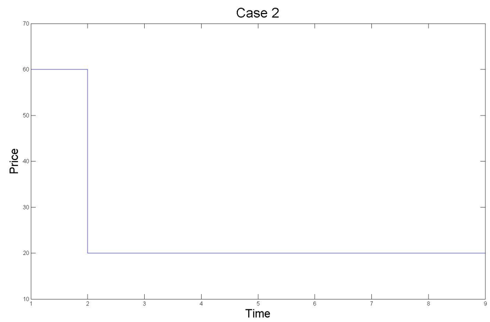

Fig. 9 Case 2 (for 1 permissible price change)

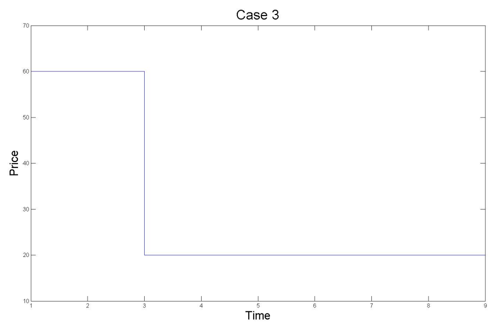

Fig. 10 Case 3 (for 1 permissible price change)

Case 9 (Fig.11) demonstrates the effect you get if you don't lower your price at all. Looking at Fig.7 you can see that your expected loss of profit would rise to approx. 42%.

ISBN 978-3-901608-32-2

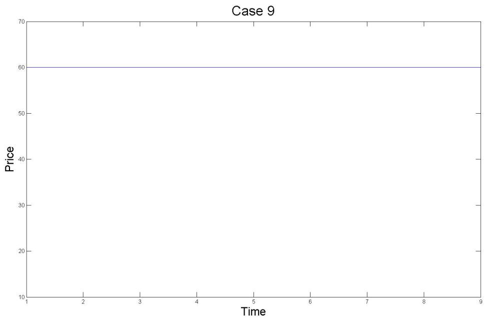

Fig. 11 Case 9 (for 1 permissible price change)

Continuing our analysis it seemed reasonable to have a look at the possible effects of a policy with two permissible price reductions. Results are depicted in Fig.12 but aren't that easy to interpret as in Fig.7.

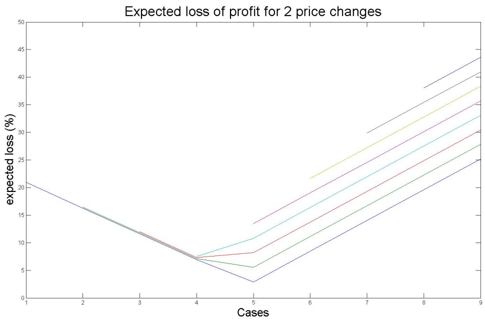

Fig. 12 Expected loss for 2 price changes

In this graph every line represents up to 9 additional cases. That's why we changed our notation in Fig.13 - 16 to Case X,X.

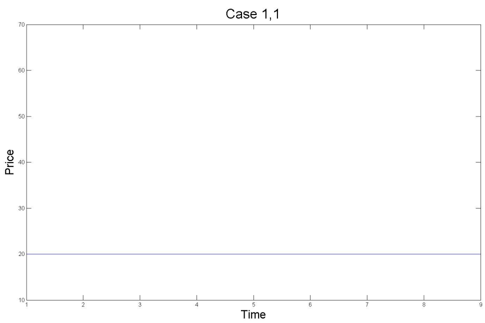

Fig. 13 Case 1,1

This policy (Fig.13) applies to the left-most point on the blue line in Fig.12. So by applying the biggest price drop in the beginning of your sales period you get an expected loss of profit of approx. 20%.

This policy (Fig.14) applies to the left-most point on the green line in Fig.12. So by applying the biggest

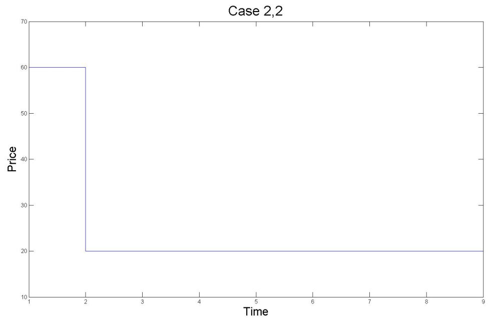

Fig. 14 Case 2,2

price drop during your second time period you get an expected loss of profit of approx. 15%. This is again because you don't get your full potential profit.

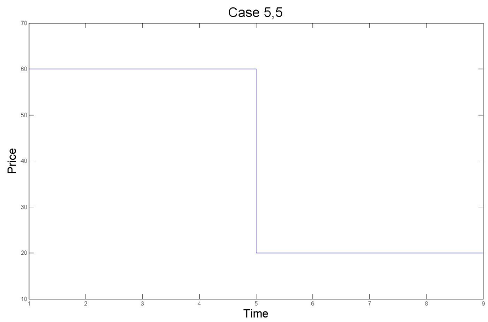

Fig. 15 Case 5,5

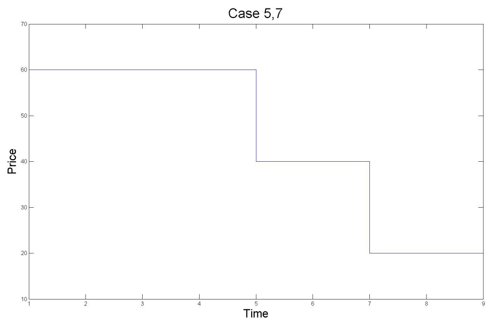

Fig. 16 Case 5,7

Fig.15 and Fig.16 correspond to the same line now (the violet one) but Fig.16 is the more interesting one. In this case you begin with full price till time period 5, then you lower your price to the first permissible price and in the beginning of time period 7 you lower again. For this policy your expected loss of profit amounts to approx. 25%.

Fig.12 still keeps two important findings ready. First the bottom line (x-axis) doesn't represent the expected optimal profit for two permissible price changes but for 100! So by looking at the graph you can easily see that the optimal policy for two price changes (the lowest point of the blue line) only accounts to additional

3% compared to a pricing policy with 100 permissible price changes! Applying this finding to the mentioned one price change policy (Fig.7) means that we only lose approx. 8% by acting this way. So by taking two price changes instead of one you gain 5% of your expected profit.

4 Issues for Future research

Several issues for future research fall in the areas of parameter estimation and model enhancements.
First we want to extend our demand model by means of Hierarchical Bayes methods. These are designed to measure differences between units using a particular prior structure because as consumer sensitivities become more and more diverse, it becomes less and less efficient to consider the market in the aggregate. Hierarchical Bayes approaches are ideal for these problems as it is possible to produce posterior distributions for a large number of unit-level parameters.
Second we want to include prices of competitors in our demand model. The reason is the dominance of this aspect in the minds of many business managers and the big influence this has on their decisions. An example provides [12].
Third is the enhancement of our optimization in terms of adding actions. Again the reason is the wish of several managers to include such actions as promotion and not only price. There exist Dynamic Programming models which combine inventory and pricing decisions so an extension in this form seems straightforward.
Fourth is the modification of our objective function in terms of risk aversion. This issue is best described by the newsvendor problem where a decision maker orders inventory before a selling season with stochastic demand. If too much is ordered, stock is left over at the end of the period, whereas too little is ordered, sales are lost. A good introduction gives [21].

5 References

[1] Vithala Rao. Pricing research in marketing: The state of the art. *Journal of Business*, 57:39–60, 1984.

[2] Thomas Nagle. Economic foundations for pricing. *Journal of Business*, 57(1):3–26, 1984.

[3] Guillermo Gallego and Garrett Van Ryzin. Optimal dynamic pricing of inventories with stochastic demand over finite horizons. *Management Science*, 40(8):999–1020, August 1994.

[4] Youyi Feng and Guillermo Gallego. Optimal starting times for end-of-season sales and optimal stopping times for promotional fares. *Management Science*, 41(8):1371–1391, August 1995.

[5] Gabriel Bitran and Susana Mondschein. Periodic pricing of seasonal products in retailing. *Management Science*, 43(1):64–79, January 1997.

[6] Stephen A. Smith and Dale Achabal. Clearance pricing and inventory policies for retail chains. *Management Science*, 44(3):285–300, Mar 1998.

[7] Arthur Allaway, Barry Mason, and Gene Brown. An optimal decision support model for department-level promotion mix planning. *Journal of retailing*, 63:215–242, 1987.

[8] Dale Achabal, Shelby McIntyre, and Stephen Smith. Maximizing profits from periodic department store promotions. *Journal of retailing*, 66(4):383–407, 1990.

[9] Phillip Carlson. Fashion retailing: the sensitivity of rate of sale to mark?down. *Journal of retailing*, 59:67–76, 1983.

[10] Aliza Heching, Guillermo Gallego, and Garrett van Ryzin. Mark-down pricing: An empirical analysis of policies and revenue potential at one apparel retailer. *Journal of revenue and pricing management [1476-6930] Heching yr:2002 vol:1 pg:139*, 1(2):139–160, 2002.

[11] Murali K. Mantrala and Surya Rao. A decision-support system that helps retailers decide order quantities and markdowns for fashion goods. *Interfaces*, 31(3):146–165, 2001.

[12] Robert C. Blattberg and Edward I. George. Shrinkage estimation of price and promotional elasticities: Seemingly unrelated equations. *Journal of the American Statistical Association*, 86(414):304–315, June 1991.

[13] D.V. Lindley and A.F.M Smith. Bayes estimates for the linear model. *Journal of the Royal Statistical Society*, 34:1–41, 1972.

[14] A.F.M. Smith. A general bayesian linear model. *Journal of the Royal Statistical Society*, 35:67–75, 1973.

[15] Kirthi Kalyanam. Pricing decisions under demand uncertainty. *Marketing Science*, 15:207–221, 1996.

[16] Alan E. Gelfand, Susan E. Hills, Amy Racine-Poon, and Adrian F.M. Smith. Illustration of bayesian inference in normal data models using gibbs sampling. *Journal of the American Statistical Association*, 85:972–985, 1990.

[17] Peter E. Rossi, Greg M. Allenby, and Robert McCulloch. *Bayesian Statistics and Marketing*. John Wiley & Sons, 2005.

[18] Michael Levy, Dhruv Grewal, Praveen Kopalle, and James Hess. Emerging trends in retail pricing practice: implications for research. *Journal of retailing*, 80(3):13–21, 2004.

[19] John D.C. Little. Models and managers: The concept of a decision calculus. *Management Science*, 16:B466–B485, 1970.

[20] Alan L. Montgomery and Eric T. Bradlow. Why analyst overconfidence about the funtional form of demand models can lead to overpricing. *Marketing Science*, 18:569–583, 1999.

[21] Maurice E. Schweitzer and Gerard P. Cachon. Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Science*, 46(3):404–420, 2000.

ISBN 978-3-901608-32-2