

---

**Barry Leventhal**

is a marketing statistician with more than 30 years experience in market research, geodemographics and mining large databases. In 2009 he founded BarryAnalytics, an independent advanced analytics consultancy. Barry sits on the board of the IDM journal and chairs the MRS Census & Geodemographics Group. He is a fellow of the IDM, MRS and RSS.

**Tom Breur**

runs consulting firm XLNT Consulting ([www.xlntconsulting.com](http://www.xlntconsulting.com)), dedicated to helping companies make more money with their data. His fields of interest span data mining, analytics, data quality, IT governance, data warehousing, Agile development methods and business models.

**Keywords:** Business analytics, markdown pricing, price optimization, yield management

Barry Leventhal  
BarryAnalytics Ltd  
9 Markham Close  
Borehamwood  
Hertfordshire  
WD6 4PQ UK  
Tel: +44 7803 231870  
E-mail: [barry@barryanalytics.com](mailto:barry@barryanalytics.com)

# Intelligent markdown pricing

Barry Leventhal and Tom Breur

Received (in revised form): 13th October, 2011

## Abstract

Many businesses, in a diverse range of industries, employ price markdown strategies in order to reduce wastage and losses incurred when products reach the end of their selling periods. This paper aims to demonstrate the commonalities in markdown pricing between different industries and explain the role that analytics can play in setting intelligent price reductions. We cover the nature of the markdown pricing problem and the type of infrastructure required to address the problem analytically. We discuss the importance of using the most granular data available and ensuring that analytics are integrated with existing business processes. We provide a generic approach to markdown pricing using advanced analytics and include a variety of illustrative case studies.

*Journal of Direct, Data and Digital Marketing Practice* (2012) **13**, 207–220.  
doi:10.1057/dddmp.2011.38

## Introduction

Consumers in an affluent society expect to be able to purchase the products that they desire, at the time of their choice — whether these are food items in a supermarket, clothes in a department store, airline tickets or hotel accommodation. Therefore, maintaining high product availability is a key goal for retailers. However, the flipside of high availability is high wastage levels on the items that do not get sold. It is important for retailers to minimize their losses incurred due to product wastage — such losses can run into hundreds of millions of pounds (or euros). Therefore, many businesses, in a diverse range of industries, employ price markdown strategies in order to reduce the wastage losses incurred when products reach the end of their selling periods.

The aims of this paper are to demonstrate the commonalities in markdown pricing between different industries and to explain the role that analytics can play in setting intelligent markdown prices.

However, intelligent markdown pricing relies on more than just analytics — we will show the importance of, and interplay between, people, process and technology. These need to be employed *together* to help the retailer manage its waste losses as part of the regular business operation process.

## Markdown pricing across industries

Markdown pricing strategies are employed in many industries. The applications we will discuss in this paper all concern products that

approach the end of their lifecycle. In order to avoid having to dispose of wasted goods, the retailer would prefer to mark down the product while it is still 'sellable'. This is done in order to fetch at least some money from a product that would otherwise need to be written off.

Markdowns benefit the consumer, in terms of providing a choice between paying the full price for a product with a long shelf life or a reduced price for a short shelf life — this choice tends to be welcomed by many consumers. At the same time, store managers generally prefer to sell those products reaching their 'sell-by' dates, rather than have to remove them from the shelves and dispose of them in the waste bin. Therefore, the store manager should support the introduction of an intelligent markdown pricing strategy.

**The analytics challenge is to find the pricing 'sweet spot'**

What different industries have in common is the analytical challenge of setting a price that is neither too high nor too low. If the chosen price point is too high, stock remains on the shelves, and does not get sold before it has to be written off. If the chosen price point is too low, the product will sell out too fast, even before the end of the trading cycle, potentially resulting in stock-outs. On top of this, excessive discounts lead to lower revenues than could be achieved from this product.

Rather diverse industries all have comparable dynamics, each with a different manifestation of 'perishable' goods. It is easy to see how bananas that are turning black need to be sold quickly, while consumers are still interested in paying for them. Similarly, an airline seat that does not get filled cannot be sold after the plane has left the gate. But also in retail and fashion, in particular with seasonal products such as clothing, there is a limited window of opportunity to sell. This year's summer clothes effectively need to be sold before June/July, or else people will not be able to wear them this season. Next year, new colours and designs will come in vogue and *nobody* wants to be seen in last year's styles. An unbooked hotel room loses all economic value as the occupation date passes by. In all these cases, there is an 'optimal' window of opportunity for selling a product. Once this window of opportunity has passed, the retailer is forced to write off on the product.

Markdown pricing strategies across disparate industries pertain to widely different products, with their own life cycle lengths. Although these markets and products are quite different, they have an overarching dynamic in common: How does one find the 'price reduction sweet spot' that will maximize product sales, yet at the same time minimize the discount the seller has to offer, in order to optimize profits.

### **The importance of people, process and technology to markdown pricing**

Intelligent markdown pricing is not solely about analytics, because it has to work within a business operation that is based on people, process and technology:

**Intelligent markdown pricing must integrate seamlessly with existing business processes**

*People* — In the traditional 'bricks and mortar' retailer, markdown pricing usually relies on store staff to identify items that are nearing the end of their selling life, check their stock levels and apply the

appropriate process to reprice those items. Furthermore, subsequent stock checks and further price reductions may need to be made, requiring additional staff time. The store manager has the authority to override the system-generated markdown prices — for example, if the manager believes that a deeper price reduction is needed in order to clear all the stock that has reached its sell-by date.

*Process* — Any new markdown pricing strategy will have to be consistent with the existing process for pricing products. And every business will have such a process, whether crude or sophisticated.

*Technology* — The technical infrastructure that supports the pricing process is another ‘given’, although we comment below on the data infrastructure that is likely to be most suitable for this application.

The interplay between these components can be complex — in a supermarket, for example, store staff may need to make multiple checks (or scans) of perishable products reaching their sell-by dates, either to capture data or reprice items:

- An initial check on the day before, in order to obtain up-to-date stock levels for the items that will need to be cleared.
- Another scan on the last day of sale, in order to generate reduced-price stickers and reprice the items.
- A further scan later in the day, in case a second markdown is required.

All three components, making up the existing pricing/markdown operation, will need to be fully understood and allowed for, as part of any approach to ‘intelligent markdown pricing’.

### The value of timely and relevant information

In order to respond *accurately* to the current stock level, one needs to confront a forecast<sup>1,2</sup> on expected sales (assuming *no* intervention by the retailer) with current stock levels. If the forecast expects (considerably) lower sales than available in stock, discounting is one of the options to consider in order to avoid waste.

Since the forecast is based on the last known level of inventory, a *timely* feed of input data is crucial in order to decide on any discount recommendation. In many cases, these data feeds are passed along into a data warehouse (DWH),<sup>3</sup> and therefore the speed with which they can be updated is crucial. The meaning of ‘timely’ depends on the volatility of sales, the particular industry and stock levels (and promotions) of complementary and substitute products. Data that arrive late can be highly accurate and properly cleansed, but are utterly useless if they come too late for making markdown recommendations.

### The value of historic information

In order to identify the ‘pricing sweet spot’ for a product, one needs historic sales and price data to calculate price elasticity (PE). At the very least, the analyst would want to have weekly data for at least one complete year, to support PE modelling.

**Timely availability of data is crucial**

### Having several years of history allows seasonal patterns to be identified

Analysts always seem to want more data, of better quality. However, in the end, they will simply need to do the best they can, with whatever data they have available. Since one needs to consider cyclical (seasonal) patterns in sales, *at least* one entire cycle is needed, but preferably several. So when looking to discount summer clothes, at least a full year is required, but if at all possible one should obtain 2–3 years or more. Analysts attempt to determine seasonal patterns; therefore, it is preferable to have *multiple* seasonal recordings in order to separate ‘signal’ from ‘noise’, that is, to discern ‘true’ seasonal variation from erratic fluctuations. Global retailers can sometimes benefit from the fact that summer is only a half year later on the southern hemisphere, but even then, having more data is always better.

### The value of granular data

Granular data, at basket level, are of great value in two ways — at the first and last stages of the analytics process.

First, they enable the retailer to examine current patterns of markdown sales and answer business questions such as:

- How do marked-down items sell across the trading day? When do the first reduced-price transactions take place?
- How many price reductions are required in order to clear the entire stock of a product that is reaching the end of its selling period? And how deep are the markdowns that are being made?
- How successful is the retailer in clearing marked-down products, and how does the success rate vary between stores, product types and price points?
- What are the characteristics of customers who purchase marked-down products at different times of the day, week, month or year (depending on the cycles that are typical of this particular industry)?

At the last stage of the analytics process, which we outline later in this paper, the granular data set can be used as a ‘test bed’ for simulating the effects of repricing the product. It allows one to predict the markdown sales and calculate the reduction in waste loss. In other words, one can estimate the return on investment that could be achieved through intelligent markdown pricing, which invariably is the most compelling way to convince the retailer’s management team to move forward with this analytic approach.

### Data quality goes beyond ‘mere’ accuracy

#### What data are needed?

Analysts always seem to complain about their data. Why is this? Their discontent seems to centre around two areas: *infrastructure* and *quality* of the data. The *infrastructure* has to do with how easy it is to get access to the data needed to do the job. Although writing complex queries is fascinating work to do, in and of itself it does not provide any added value to the business. The *quality* goes beyond ‘mere’ accuracy. Are the data timely, relevant, complete, trusted and accessible (Olson<sup>4</sup>)? In

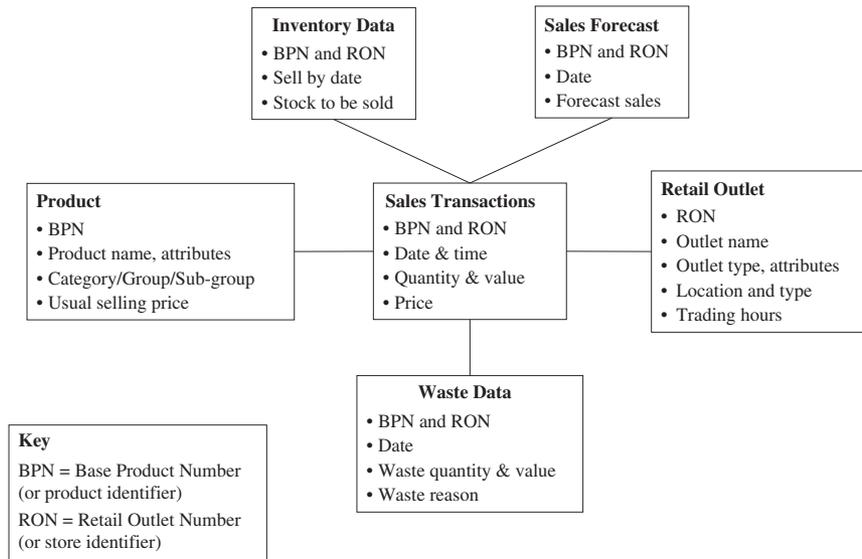


Figure 1: Data requirements for intelligent markdown pricing in retail

short, the *infrastructure* that data analysts need revolves around proper modelling of the data. But it also has to do with making access efficient and ergonomic. The *quality* of the data is about their accuracy, validity and usefulness.

The data sources that are likely to be required for markdown pricing in a typical ‘bricks and mortar’ retailer are outlined in Figure 1. As Figure 1 shows, the primary sources cover all stages in the selling process: inventory available to be sold, forecasts of expected sales, actual sales and residual wastage.

**Intelligent markdown pricing is an ideal application to gain traction from a data warehouse**

**Data infrastructure**

The analytical data infrastructure encompasses both hardware and software components. Users typically consider hardware as ‘merely’ computing power (CPUs). However, analytical applications like the ones we describe in this paper tend to push the boundaries of what is technically possible or feasible, more so than almost any other application of a DWH. No matter how powerful the machine, it is possible to dream up a query that will bring the system to its knees — the so-called ‘query that dims the light’.

For example, PEs vary between products. Some product sales hardly change unless one offers very deep discounts. This typically refers to what are considered ‘discretionary’ or ‘luxury’ goods. Because these products are not necessarily ‘needed’, customers are less likely to stock up when they have an opportunity to buy at a reduced price.

Often, the impact of markdown pricing will differ between stores, depending on the affluence levels of their customers. Therefore, the analyst will wish to work at product-store level; however, this leads to huge amounts of computations. A large supermarket chain will have several thousand fresh food products and possibly over a thousand

stores. This implies more than a million product–store combinations that simultaneously require intelligent pricing. It is easy to see how this will place enormous demands on CPUs.

For these performance reasons, DWH systems have come to the fore, which can store vast amounts of information and have the required organization of data and processing power to handle large numbers of pricing combinations. Furthermore, intelligent markdown pricing is an ideal application for a dynamic markdown process, in which near-real-time pricing recommendations are made during the selling period, based on the latest information on stock levels and rates of sale.

### **Data quality**

Insufficient accuracy and questionable validity seem to be perennial data quality issues.<sup>4,5</sup> According to Pyle,<sup>6</sup> ‘... it [data] seems intent on leading the unwary astray. Invariably, the data that a data miner has to use seems particularly well constructed to promote frustration’. The fact that better data are preferable is hard to deny. Data quality seems like the ultimate ‘motherhood and apple pie’ issue. Everybody agrees on its importance, and everybody agrees the data should be improved. But when the time comes to actually do the work, it is remarkably difficult to find volunteers who want to *do* something about it. However, despite either subtle or egregious shortcomings in the data, there may still be an excellent business case (financially) for extracting value from the available data at hand anyway. Perhaps after managing to secure a profitable business case for working with suboptimal data, one can convince (senior) management of a business case for structurally improving data quality.

### **High-level approach to intelligent markdown pricing**

We would advise the following generic steps in order to create an intelligent markdown pricing solution:

- Understand the existing markdown pricing process.
- Carry out preliminary analysis to examine levels and patterns of waste losses currently being incurred.
- Calculate PEs for products being marked down.
- Create an alternative markdown pricing strategy designed to reduce the losses.
- Simulate the impact of the pricing strategy and measure its benefits.
- Carry out a live test to see how customers react to the change in markdown pricing.

A brief explanation of each step follows below. Bear in mind that the detailed approach will depend on industry and the available data.

- *Understand the existing markdown pricing process*

For intelligent markdown pricing to add value, it needs to fit into the existing markdown process — therefore, a detailed understanding of the current process and markdown strategy is an essential first stage. For

example, as identified above, retailers selling perishable products may have to carry out daily checks to identify the items reaching the end of their selling periods, quantify the numbers of units that need to be marked down, calculate reduced prices and apply these to the items. This process can be highly complex and staff intensive.

The performance metrics that are used to measure each store manager may include a target on product wastage (due to reaching the end of the 'sell-by' date). If so, it is important to identify and understand the metric — for a new markdown pricing strategy to be successful, it will have to support (or at least not conflict with) the manager's performance target.

The existing operational process needs to be fully understood in order to introduce a seamless intelligent markdown process that avoids creating additional workload or unnecessary complications. Although running the analyses might seem difficult, changing existing business processes may prove to be at least as hard. And *that is* where the money is made.

- *Preliminary analysis to examine business questions concerning waste losses*

**Preliminary analysis will reveal strengths and weaknesses of the existing pricing strategy**

The initial analysis should quantify the levels of losses and understand how these vary by factors such as product price, day of week and outlet type. If granular data are available, then more detailed business questions can be answered, as discussed above. Where does the business think they are losing money? How does this line up with the data? Losses are not always greatest where they appear to be, and therefore the preliminary analysis can yield unexpected findings and new insights. Data analysts can drive change here by putting together compelling business cases. Providing *insight* alongside 'bare' predictions fosters adoption of intelligent markdown pricing, and may spur innovations in current (store operation) practices.

- *Calculate PEs for products to be marked down*

Historic sales and price data are used, together with other attributes (such as seasonality) to calculate the PE coefficient of each product. The PE coefficient measures the slope of the relationship between price and demand, defined and illustrated hypothetically in Figure 2.

Since demand usually decreases as price increases, the PE invariably has a negative value (see, eg Tellis<sup>7</sup>). If the PE is large in size (eg  $-2$  or less), then a small price reduction will greatly increase the demand and therefore should be sufficient to clear the stock to be sold. On the other hand, if the PE has a smaller size (eg between  $0$  and  $-1$ ), then the product is less sensitive to price, implying that a deeper price reduction will be required.

For the most accurate and valid calculations of these demand curves, historic sales as affected by price changes (markdowns) are required. This typically involves analyzing data for each day on which markdowns took place, modelling the relationship between price reduction and quantity sold. Given the central role these elasticities

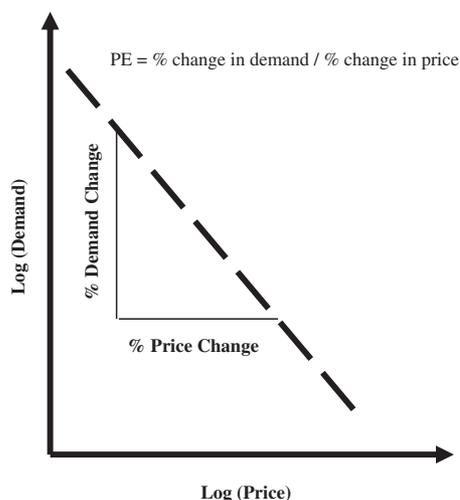


Figure 2: The PE coefficient

play, great care should be exercised to gather as much useful data as possible, and to ensure their validity.

In some cases, reducing the price of Product A will impact sales of Products B and C, which either compete with Product A or are complementary items. This situation is common in supermarket retailing — for example, reducing the price of apples will increase their sales but may have the effect of reducing purchases of oranges. In this case, it is also necessary to calculate the ‘cross-PE’, that is, the relationship between price of apples and sales of oranges, and to take account of this relationship when deciding the markdown reduction on apples.

- *Create an alternative markdown pricing strategy designed to reduce the losses*

Putting the PEs together with the stock levels and forecast sales, the analyst arrives at a markdown pricing strategy for each product. This will include both the markdown discount and the optimal time at which the reduction(s) should be made. An understanding of current systems is essential in order to ensure that the new markdown strategy can effectively be deployed within the existing business process — for example, there is no point in producing a strategy that requires different fresh food products to be marked down at different times of the trading day, because this would be impractical to put into practice. This demonstrates the intricate interplay between analytics, people and process.

- *Simulate the impact of the pricing strategy and measure its benefits*

**Simulation predicts likely returns from markdown strategies**

The analyst applies the new markdown strategy to a sample of recent trading days, using the PE models to predict each product’s sales at its reduced price. The following steps are involved:

- The new markdown pricing strategy is applied, in order to reprice each product, for each day on which markdowns were required.

- The effect on sales is predicted, by applying the PE models.
- This gives a predicted outcome for that day, that is, the quantity sold (at reduced price), the quantity unsold (for wastage) and the resulting financial loss.

Needless to say that to make this simulation valid, the sample days should *not* have been used for building the model!

The overall return for each product category can then be derived and hence it is possible to calculate the loss saving to be made by introducing the proposed intelligent markdown pricing strategy.

- *Carry out a live test to see how customers react to the change in markdown pricing*

### Performing a live test is vital for measuring actual returns

Although the simulation results will *predict* the return from intelligent markdown pricing using the data alone, the only way to measure its actual success is to carry out a live test using a sample of customers or a small number of stores. The proof of the pudding is in the eating, and management should not attempt to bypass the live test in order to save time — if their need to implement a new price reduction strategy is that great, then think how much more deeply in trouble they will be if they bypass the test and the new strategy fails.

Critically, the test should include a matched control sample, for comparison purposes. Matching is preferred over one large random sample of test and control stores, in order to increase the statistical power of this test. Matching implies finding pairs of stores that are as comparable as possible, and then randomly selecting which of each pair eventually becomes the test store. In ‘ordinary’ sampling one would randomly select test and control stores, but this might lead to large stores being included as test store, and smaller ones for the control group (by chance). Matching reduces this between-group variance. It needs to be carefully designed in order to avoid biasing the results in any way. The design merits care, because a more valid test design leads to more compelling business cases.

Two important reasons to include live tests are firstly to gather the most valid empirical evidence to convince all stakeholders of the value of intelligent markdown pricing, and secondly to experience in a realistic setting where adoption of analytics will affect the operational processes currently in place.

### Industry-specific issues

Although intelligent markdown pricing fits into a generic approach that draws on the commonalities between industries, local issues always come into play when considering a specific business. For example:

- In supermarket and food retailers, the selling period for marked down products is very short — typically just one day — and store staff may not always be able to follow the laid-down business process due to other short-term priorities.

- In fashion retailing, new product lines are introduced every season, and therefore no sales history is available for calculating their PEs. This greatly complicates the analyst's task — typically, new lines are matched with previous season merchandise, using a process known as 'stock-keeping unit (SKU) likening', in order to assign sensible PE values based on last season's products (or SKUs).
- In the airline industry, a generic system of price banding and discounting (using letters A–Z) is commonly embraced to facilitate resale of seats among carriers and intermediaries. Any intelligent markdown pricing scheme needs to fit into this model to make it feasible to implement.
- Hotel rooms are largely sold through intermediaries, with differential pricing for 'blocks' of rooms. However, these are not physical blocks; this pertains to a certain *volume* of sales that is used to predict occupancy. Therefore, it is important that resellers not only provide actual sales, but also sales dates, so that the industry can employ further discounting accordingly.

## Case study examples

**The role of analytics was to *support* decision making, not to *replace* it**

### Fashion retail case

A fashion chain faced the following challenge: although a central analytics department was quite capable of delivering reliable predictions, more often than not headquarters played no more than a superficial role in actual purchase and markdown pricing decisions for fashion collections. This chain of fashion stores had widely variable sales rates across product lines, and relatively high write-offs. Analytics staff were quite capable of providing accurate predictions. The ingredients for profitable application of analytics were all present: timely data of good quality, sufficient history and good insight into the repricing process. Why then was it so difficult to convince store managers to adopt stock predictions, and make purchasing and markdown decisions accordingly?

Because fashion is considered such a personal product, many people in the profession feel (felt) it cannot be predicted by something as impersonal as an algorithm. Because of this, store managers still aim to determine assortment and stock levels themselves, possibly because they know local trends and/or preferences best. Or arguably, because — in part due to cognitive dissonance<sup>8</sup> — they will go out of their way to sell what they have committed to carry. 'Doesn't that look absolutely fabulous on you'?

Store managers would receive recommendations on *what* to buy and *how much* and a summary explanation of *why*. Towards the end of the season, it is hugely important to start marking down in time to still be able to sell stock, but not too early as to 'hurt' sales at regular price levels. The markdowns tend to be quite deep, in particular when a large supply needs to be sold. The summary explanations proved to be quite important. They provided 'insight' rather than 'bare' predictions. Despite the fact that this was no more than an aggregate display of sales (and remaining stock) at comparable benchmark stores,

it provided a sense of context, reliability and sanity checks of all numbers provided.

What this case exemplifies, among other things, is that analytics is foremost used to *support* decision making, not *replace* it. This distinction is not always obvious.

**Understand how the pricing strategy will be deployed before planning your approach to creating it**

### Supermarket case

This case study exemplifies that analytic predictions work ‘backwards’: one first needs to consider deployment before investing in making predictions.

A supermarket chain was experiencing heavy losses on perishable products, as the ‘downside’ of maintaining high product availability levels for its customers. From the preliminary analysis of wastage losses, it was found that store managers were doing a very good job at minimizing wastage, but this was at the cost of making high reduced-to-clear markdown reductions.

This introduced the possibility that overall losses could be reduced if the markdowns were less generous to the consumer, even if this meant a small increase in the quantity of unsold products that would need to go in the waste bin.

The results of extensive analysis and modelling demonstrated the potential for reducing losses and the need for a store pilot as a ‘live test’. The retailer accepted this recommendation and piloted the new intelligent markdown rates in ten stores, and the benefits were significant. The overall losses were generally lower in the pilot stores, in comparison with a control group, and decreased over the pilot period — at a time of year when losses on fresh foods were generally increasing.

However, the pilot also demonstrated poor compliance rates with the retailer’s existing generic markdown rules — compliance was tracked by monitoring the reduced prices paid for items and the times of those transactions. Therefore, the retailer decided to focus on improving adherence to the existing process, and then return to the issue of optimizing markdown rates after satisfactory compliance levels had been achieved.

**Pricing in the airline industry can be hampered by non-availability of key inputs such as sales dates**

### Airline case

In the airline industry, seats are offered through multiple channels. When deciding how to price individual seats, companies employ a process commonly known as ‘yield management’. This requires a holistic view on sales to date, hence merging sales data across all channels. Timely data are crucial for this. Proprietary sales data (direct sales, without using resellers) may be a significant portion of all sales, but they do not provide the whole picture. Internet data tend to come in early after the consumer displays his interest; sales through ‘traditional’ booking channels tend to come in with a delay. This time lag would not be so much of a problem (at least for early, heavily discounted sales) if the *date of sale* to the consumer were always known when the reseller provided sales data to the airline carrier.

A relatively small portion of airline seats are offered early (say, three months before departure) at steeply discounted rates. Depending

on how quickly these first 10 per cent of seats are sold out, subsequent blocks of seats are sold at increasingly higher rates. Because the difference between early sales and last-minute sales (at 'full fare') is so large, the carrier has an interest in keeping some seats available for last-minute bookings. What ordinary consumers consider to be 'last-minute deals' (deep discounts) are actually a rarity, and largely due to a surplus of capacity. Sometimes it may not be possible to cancel a flight when demand is low, either because too few alternative flight schedules are available or because too many seats have been sold already to accommodate in later flights. This situation illustrates poor yield management. A better tactic is to either (almost) sell out or to raise the seat prices early and consider cancelling a flight if alternative times on the same leg (or connection) are available.

Depending on historical 'no-shows', and availability of alternative flights connecting the same destinations (with or without a stop-over), the airline will consider some degree of overbooking. This is the practice of structurally selling more seats than are available on the flight. Airlines risk disappointing 1–2 per cent of travellers if by chance (almost) everybody shows up. These customers then need to be offered some compensation, such as hotel stays, upgrades and/or additional airline loyalty points. These costs need to be weighed against the option of selling more tickets, and is essentially a business decision, although informed by analytics.

Many seats are sold through intermediaries. Because the eventual confirmation of the seat typically runs through an airline booking system, at least the carrier stays reasonably well informed on how much interest there is from consumers for these first (least expensive) seats. Thus, timeliness of data is somewhat less of an issue. What remains opaque for the airline is how much 'marketing pressure' the reseller is putting in (typically purchasing web traffic through banner advertising). Discretionary marketing spend by resellers heavily influences the sales cycle.

**Heavy reliance on intermediaries makes it harder for hotel chains to manage markdown pricing**

**Hotel case**

In the hospitality industry, it is common for the majority of rooms to be sold through intermediaries, often in conjunction with additional hotel services (eg: conference facilities and catering). Because 'blocks' of hotel rooms are offered to intermediaries, at a fixed discount rate, the hotel may not learn which portion have been sold until the entire block sells out. Therefore although the mechanics of yield management are similar to the airline industry, there are two main differences.

The first major difference is that a much larger portion of rooms are sold through intermediaries, which means that proprietary hotel room sales do not provide much accuracy in terms of committed sales by end consumers (resellers can often 'return' unsold hotel rooms, although at some penalty). Resellers may well provide 90 per cent+ of sales. Secondly, the rooms that resellers have already sold do not become 'visible' to the hotel, until the reseller's entire block is committed. This causes significant delays in predicting occupancy rates and exacerbates the problems of accurate estimation.

Because resellers are so important, the ‘balance of power’ is rather different (less favourable for hotel owners) for the hospitality industry. There is much less IT standardization than for airline booking systems, and thus hotel managers cannot be overly demanding in terms of the kind of data provisioning one might desire. Furthermore, overbooking is not really an option for hotels, and thus they are somewhat more dependent on accurate predictions to achieve close to 100 per cent occupancy rates.

Predicting occupation rates is further complicated, because events such as conferences or wedding parties can trigger additional bookings that make the process of yield management somewhat crude; at least it may appear this way to outsiders. What this case exemplifies, foremost, is that unless one obtains *timely* information on how much sales have been *committed* by consumers, it is impossible to set prices accurately to avoid the risk of having many rooms unsold.

**Analytics indicate the benefits of a strategy that’s been optimized to achieve your objectives**

### **The benefits of intelligent markdown pricing**

From our experience, existing markdown pricing strategies often either use arbitrary generic price reductions or are driven by parameters that were decided by managers’ ‘gut feel’, rather than being optimized using advanced analytics. In either case, these strategies can err towards making deep price reductions that ensure that products are fully sold off rather than wasted. However, their impact is to give too much money back to the consumer and achieve inadequate revenues for the retailer. Timing of price reductions is also never optimal if it is based on intuition, rather than analytics.

An intelligent markdown pricing approach will optimize across the losses arising from price reductions and the wastage losses from products that have not sold. It may result in smaller markdown reductions that will leave some stock unsold at the end of the period, but will achieve lower losses overall.

Since PEs and buying patterns differ between products, we advocate creating markdown strategies at the finest level of granularity allowed by the retailer’s system. Often, the system will allow price reductions to be input at a finer level than is actually being used, and thus a finer strategy may be employed cost-free within the existing infrastructure. Differentiation across products should be derived analytically.

We also recommend that, in the retail case, one considers whether markdown pricing should be carried out for individual stores or by store formats. For example, a store located in an affluent area may be able to make relatively small percentage reductions and encourage customers to purchase those products. On the other hand, a store in a poor area may need to make deep reductions in order to keep its customers satisfied. Therefore, a single strategy for all stores rarely makes sense.

### **Conclusions**

Retailers want to offer their customers a wide portfolio of products, and they never want to run out of stock. Yet many products have an inherent lifecycle and therefore have a limited window of opportunity

to sell. For perishable products this ends at the 'sell-by' date. For clothing there are seasons and styles that come into vogue, only later to go out of fashion. Hotels offer accommodation that can only be sold until the day of stay arrives.

In order to sell all remaining 'stock', retailers can reduce prices for excess inventory to avoid wasting products that have reached the end of their selling periods. There is an interesting commonality across industries that requires optimization: it is desirable to mark down products enough so as to clear all (or at least most) inventory, yet at the same time not discount too deeply so as to offer too favourable prices for consumers.

Intelligent markdown pricing is the application of sophisticated analytics in order to optimize the timing and level of discount that should be offered. It is a data-intensive process that needs to be carefully integrated into existing business practices. There is an intriguing interplay between people, process and (analytic) technology. People and process in existing (store) operations tend to take precedence. Whatever bandwidth remains for analytics in large part determines what options remain for markdown pricing.

The existing DWH technology plays a large role; however, development of such capabilities tends to require significant investments, hence presenting a 'barrier to entry' into the analytics space. At the end of the day, markdown pricing models are only as good as the data that go into them. Frustrating as this may be in and of itself, experience shows that even suboptimal data can often be leveraged to produce highly profitable models.

#### References

1. Mentzer, J.T. and Moon, M.A. (2005) *Sales Forecasting Management — A Demand Management Approach*, Sage Publications, Thousands Oaks, CA.
2. Thomas Wallace, T.F. and Stahl, R. (2002) *Sales Forecasting — A New Approach*, T.F. Wallace & Co, Cincinnati, OH.
3. Inmon, W.H. (2005) *Building the Data Warehouse*, 4th Ed. Wiley, Indianapolis, IN.
4. Olson, J.E. (2003) *Data Quality: The Accuracy Dimension*, Morgan Kaufman, San Francisco, CA.
5. Thomas Redman, T. (2008) *Data Driven: Profiting from Your Most Important Business Assets*, Harvard Business School Press, Boston, MA.
6. Pyle, D. (2003) *Business Modeling and Data Mining*, Morgan Kaufman, San Francisco, CA.
7. Tellis, G.J. (1988) 'The price elasticity of selective demand: A meta-analysis of econometric models of sales', *Journal of Marketing Research*, Vol. 25, No. 4, pp. 331–341.
8. Festinger, L. (1957) *A Theory of Cognitive Dissonance*, Stanford University Press, Stanford, CA.