Revenue management

Riding up the demand curve

This article reviews the technicalities of revenue management, a pricing tactic widely applied in a range of service industries. A case study is calculated using an Excel spreadsheet, whereupon the concept is then explained theoretically, methodologically and graphically.

Marian Mazdra, October 2018

Revenue Management[1] has its origin in the airline industry and is well known to airline customers around the globe. The earlier you buy a ticket for a certain flight, the lower the price, while late bookers often pay substantially more. Since its inception in the aviation industry in the 1980s, it has since become popular in many other sectors.

Technically, revenue management tries to understand the individual value which a company creates for each customer or segment of customers, and then it determines the price for that value. Simply put, it is about selling the right product to the right customer at the right time for the right price.

There are nine requirements for revenue management:

1. The business operates in the for-profit arena (i.e. it may be morally acceptable to hold back hotel rooms for guests that are willing to pay more but withdrawing beds in a hospital is apparently unacceptable).
2. The business entails significant freedom of action and no legal constraints are in place (e.g., in regulated businesses such as the pharmaceutical industry and in financial services, there are restrictions on how to set prices).
3. There must be no arbitrage opportunities between market participants (e.g., no one is able to buy at cheap prices and resell the products or services).
4. Prices must be set with barriers so that the segment willing to pay more cannot pay the lower price. Such barriers can be time, location, inconvenience, prestige, special customer service (e.g., business travellers usually can´t book a flight three months ahead as they usually don´t know where they need to go, nor when).
5. Customers can be clearly classified into various segments with different demand curves (e.g., business travellers are willing to pay more for the ability to book late for a given destination, while leisure travellers seek lower prices but are more flexible on destination).
6. There is a limited capacity of a *perishable* good (e.g., a hotel room for tomorrow night is perishable in that it cannot be consumed the night after tomorrow, and the hotel cannot sell it the night after tomorrow; any hotel room that is not booked on a given day is lost money for the hotel)
7. The same unit of capacity (rooms, seats, etc.) can be used to deliver all segments
8. You have a sufficiently large database and know your demand curve.

Companies that apply price discrimination typically have low or even no ability to differentiate between products, and they very high fix costs and relatively low variable costs (e.g., the costs of servicing a hotel guest upon her arrival is very low compared to the fixed costs the hotel must cover). A hotel room is a hotel room, a seat on a plane is a seat on a plane – there is little scope for making a difference.

On the other hand, airlines, for example, offer the opportunity to book at very short notice while they have to, however, bear relatively high fixed costs. If bookings do not occur, the airline has to assume all these costs themselves. In a nutshell, then, the higher prices a late-booking business traveller pays compared to the tourist who books early is justified by the convenience of obtaining a ticket at short notice, sometimes even on the day of departure.

A simple example

Suppose a hotel has a limited number of rooms to sell. The demand schedule of these rooms is defined by how much the rooms are desired at any given price level. The law of demand says that as prices increase, the demand for a product will go down. If price decreases, demand will go up.

Exhibit 1 Demand curve[2]

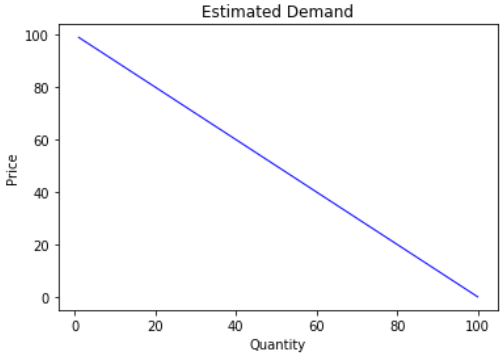

A simple demand curve is shown in exhibit 1: The demand curve suggests that for 1 currency unit increase (decrease) in price, the quantity sold decreases (increases) by 1 unit. It follows then, that the maximum revenue the hotel can generate is 2,500 currency units. This is the case if the hotel charges a price of 50, resulting in bookings of 50 rooms. Any decrease in price would lead to a higher number of rooms sold, though the overall revenue would come in lower. Similarly, any price increase would not be able to compensate for the decrease in quantity.

However, the demand curve illustrates overall average willingness to pay across all customer groups, and not all these groups are equally price sensitive. Suppose the hotel can identify three heterogeneous segments:

- *Business travellers*: They book very late, often only a day or two before arrival. Consequently, representatives of this segment are willing to pay a higher price relative to other groups.
- *Leisure travellers* usually book well in advance, have some freedom of choice regarding destination and hotel and are thus much more price sensitive than a typical business traveller.
- A third segment are *students*. Given their usually tiny budgets, their willingness (and assumingly ability) to pay is lowest.

What choices does the hotel have at its disposal? Let's start with the extreme case of selling only to students. This segment pays the lowest price, but there would be a high probability that all rooms could be sold. In that case, the maximum revenue equals 1,600 units (see the top graph in exhibit 2).

However, the hotel can do better. Suppose the hotel's data marketing department suggests that 60 rooms could be booked by leisure travellers. In that case, it is better to sell only 20 rooms to students and charge the tourists a higher price. Total revenue would then come in at 2,800 units from the combined sales to students (20*20=400) and to leisure travellers (60*40=2,400). See the second graph in exhibit 2.

Exhibit 2: Skimming off the consumer surplus

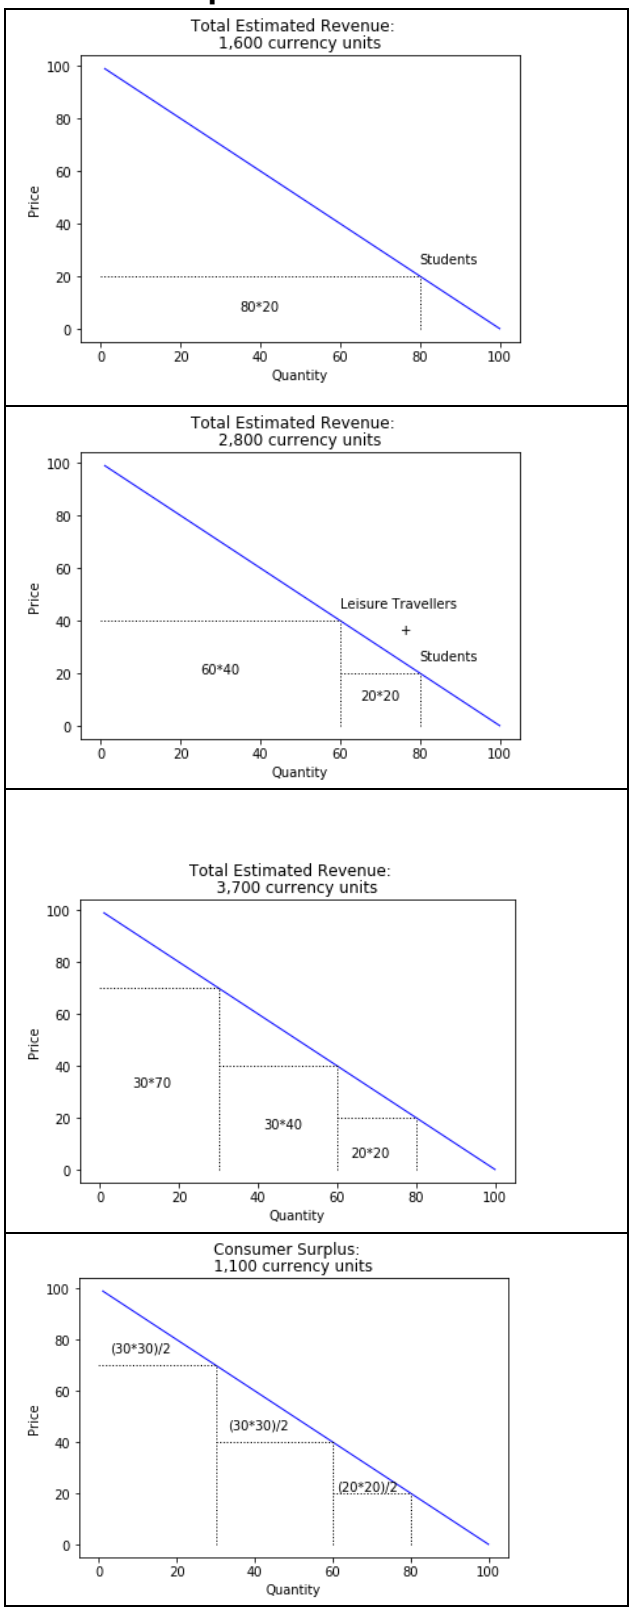

However, the hotel can do even better. Past data allows us assume that there are 30 business travellers scheduled to check in later, so the hotel holds back this quantity for them. Now the total revenue equals 3,700 units: the students' 400 plus the 30*30=900 (see the graph at the bottom in exhibit 2).

However, even more can be extracted for the hotel. The total revenue of 3,700 has been generated by using a uniform pricing strategy across the three segments. As the demand curve suggests, the hotel could realise a higher price for almost any given quantity. This is illustrated by the triangles above each revenue segment as shown in the graph at the bottom of exhibit 2.

These triangles represent the consumer surplus. This is the difference between the price a customer is willing to pay and the price she actually pays. Simply put, if I am willing to pay 3.50 Euro for a cup of coffee but you charge 3.00 Euro, then the consumer surplus is 50 cents.

Consequently, assuming all else remains constant, the consumer surplus increases as prices fall, and it decreases as prices rise.

From an entrepreneurial point of view, consumer surplus is the amount of money the company leaves on the table. In our simple example, this is 1,100 currency units, thus some 30% of the total revenue the hotel has estimated.

To skim off the consumer surplus as much as possible, the hotel could apply a technique based on statistical data form the past. This technique is revenue management.

The next section describes how this can be done by illustrating a simple example.

Basics of revenue management

Revenue management is one of three types of price discrimination[3] (see the exhibit below).

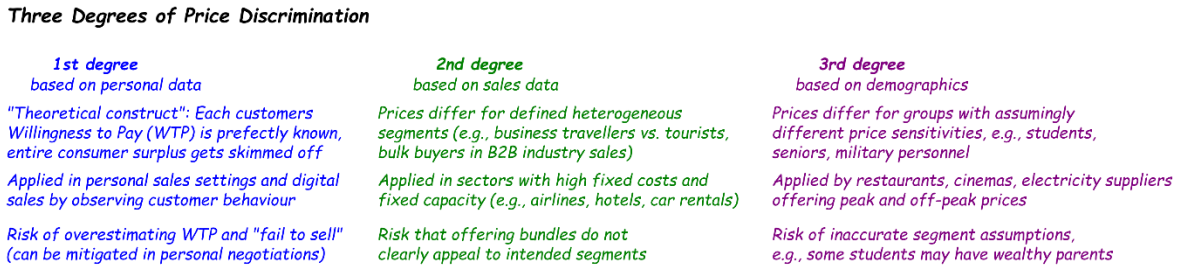

From a pricing perspective, the ideal world would be one where the willingness of each individual customer is perfectly known. The company then sets a price for each individual customer according to her and his willingness to pay, reducing the consumer surplus to zero. Such price discrimination is called *first-degree price discrimination*.

Although such an ideal world does not exist, a variation of this kind of revenue management is applied in personal sales settings. Some claim that first degree price discrimination has its origin in traditional Middle Eastern souks of medieval times, where the seller would screen each prospective customer from head to toe and set the price according to one's physical appearance: Well-dressed women may then be charged more than others [4].

In case knowing the customer in person is impossible, a business can base its pricing on broad demographics such as age or gender. Economists call this *third-degree price discrimination*. A "lady's night" at a discotheque offering free entrance to women (hoping that the men will follow for a price), and a movie theatre offering discounts to students Monday afternoon (to fill seats that would otherwise be empty), constitute two typical examples that may have their reason, although this type of pricing might not be exploiting each segment's full potential.

This is where the *second-degree price discrimination* comes into play. It creates and communicates product packages to clearly defined heterogeneous customer segments. Each of these offering bundles appeals to one segment. Lower flight ticket prices for tourists who book earlier than for business travellers who mostly book at short notice, is a good example for this tactic.

If it works, customers segment themselves and pay the corresponding price, while product bundling comes at low or even no cost to the company. The goal of revenue management is to utilise the maximum of a fixed capacity – selling all seats on an airplane, all rooms in a hotel – while at the same time maximising revenue.

When this fixed capacity is divided into the different customer segments, the selling procedure starts low. For example, airlines offer tickets at the lowest price and increase the price over time.

This is for good reason. If the airline started with the highest price and then decreased it, people would simply wait longer before they book. Hence, you'd be better off starting with a low price and riding up the demand curve.

Let us consider the example of a hotel with 100 rooms. The hotel introduces different prices for each of its two segments: business travellers who are supposed to pay more and purchase late, and leisure travellers who might pay less and purchase soon.

The hotel faces two risks:

- If it sells too many rooms early on, it forgoes revenue, as, later it will no longer have enough tickets for the higher-paying segment who are willing to pay more. This risk is called *spill*.
- If the hotel sells too few tickets, however, it bears the risk that late bookers do not purchase them all, leaving some rooms empty (therein again foregoing revenue). This risk is called *spoilage*.

To achieve the "no spill-no spoilage" balance, the hotel must know its *booking limit*. That is, it has to know the number of rooms that should be sold to leisure travellers. You subtract this ten from the 100 rooms available in total to calculate the *protection level*, the capacity reserved for the higher-priced segment.

The formula simply is:

Booking limit + protection level = 100%

The Expected Marginal Revenue: A practical example

The concept that deals with this is the Expected Marginal Seats Revenue (EMSR), developed by Peter Belobaba[5]. Belobaba developed his theory specifically for the airline industry, hence the term "seat" in the name, but you can replace it by any other item – rooms, beds, cabins, spots, whatever your fixed capacity is.

The theoretical concept behind it is the normal inverse cumulative distribution function (normal inverse CDF), also called quantile function.

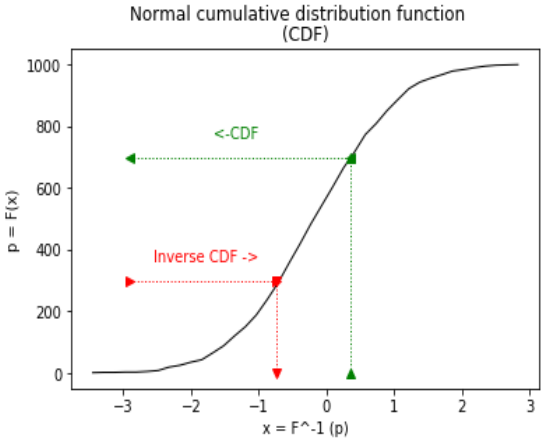

The cumulative distribution function tells us the function that returns the probability p that a certain value X is equal to or lower than a value x. We denote: P (X<=x) = F(x), or simply: p = F(x). Therefore, value x is the input and we seek the probability p.

By the inverse cumulative distribution function (which we deal here with), we denote the value x that would make the function F(x) return a probability p: $x = F^{-1}(p)$. Therefore, we have p as the input and are seeking x.

The relationship between the CDF and its inverse is illustrated in the graph to the left.

In revenue management, we have the probability as the input, because we can estimate it from past data (hence we deal with the inverse CDF). As we shall see, the probability is then used to compute the quantity x to be reserved for the higher-priced segment. Let´s turn to an example: Suppose we run a hotel with 100 rooms. It features 2 fares: rooms at €150 and at €110, and we want to sell these rooms on a particular Friday. From comparable past bookings, we have the data shown in exhibit 3.

Exhibit 3 Sales data for a hotel (example)

1	A	B	C	D
2		Rooms sold at €150	Rooms sold at €110	
3		€ 150,00	€ 110,00	
4	1	58	70	
5	2	43	80	
6	3	39	85	
7	4	39	61	
8	5	61	86	
9	6	36	78	
10	7	38	81	
11	8	29	91	
12	9	48	69	
13	10	34	71	
14	11	33	88	
15	12	47	77	

The question is: how many rooms should be sold at the lower price, and how many reserved for the higher price segment?

Let's write down what we see from the data:

Price for segment H (higher-priced segment) => p_H = *€150.00*

Price for segment L (lower-priced segment) = > p_L = *€110.00*

We calculate the mean and the standard deviation for the higher-priced segment.

Mean μ demand for segment H => D_H = *42.08*

Standard deviation σ *for segment H =>* SD_H = *9.83*

What we seek is the optimal allocation of capacity, i.e. the capacity to be reserved for the higher-priced segment, C_H. This optimum is reached when the probability for additional revenue made in the higher-priced segment equals the lower price.

To provide a simple example, let's suppose the hotel has already sold a certain number of rooms at the lower price of €110. Hence, the hotel has remaining capacity "c". Now a further customer calls in and demands a room at the same low price (let's call this inquiry x). If the hotel accepts, the hotel will generate €110 for certain.

However, there are so-called opportunity costs. That is, if the hotel accepts, it cannot generate the higher price of €150 later, as this room has already been booked at the lower price. Yet these opportunity costs only incur only if demand D_H for the higher price is equal or higher than the remaining capacity c. Mathematically, this means: $c = D_H - x$.

Theoretically, we therefore write: P (Demand from segment H > C_H) * p_H = p_L, thus:

$$P \text{ (Demand from segment H} > C_H) = p_L / p_H$$

The related formula in Excel[6] is NORMINV (1-(p_L / p_H), D_A, SD_A), hence the protection level calculates as follows:

NORMINV(1-(110/150),B18,B20)

Exhibit 4 illustrates this using a spreadsheet. Excel returns a protection level of 35 (see cell B23 in exhibit 6). It follows then, that you should reserve 35 rooms for the higher price segment and sell 65 rooms to the lower price segment. As an average of 78 people are willing to pay €110 (see cell B17), you will likely be able to sell the 65 vacancies resulting in a total revenue of

35 rooms * €150 + 65 rooms *€110 = €12,400

Exhibit 4 Sales data for a hotel (cont.)

1		A	B	C	D
2			Rooms sold at €150	Rooms sold at €110	
3			€ 150,00	€ 110,00	
4		1	58	70	
5		2	43	80	
6		3	39	85	
7		4	39	61	
8		5	61	86	
9		6	36	78	
10		7	38	81	
11		8	29	91	
12		9	48	69	
13		10	34	71	
14		11	33	88	
15		12	47	77	
16					
17		Probability	0,267		
18					
19		Mean	42,08	78,08	
20					
21		Standard Deviation	9,83	8,94	
22					
23		Protection Quantity	35,96		
24					
25					
26					
27					

= 1-(C2/B2)

=MEAN(B4:B15)

=STDEV(B4:B15)

=NORMINV(1-(C2/B2);B19;B21)

As the capacity is fixed at 100 rooms and the protection level is 35, the booking limit is 100 – 35 = 65. Hence, the hotel can sell 65 room reservations to the lower price segment before it should raise its price. As an average of 78 are expected to pay 110, the booking limit is lower, and the business can expect to generate more revenue when not pursuing a uniform pricing tactic.

Brief outlook for the future

Given the uncertainty of our contemporary economic transformation, a detailed forecast about the future of revenue management would mean little more than gazing into the crystal ball. I shall therefore limit myself to three simple points about how entrepreneurs and managers should deal with this situation:

Manage customer relationships rather than inventory.

Look at all changes with the eyes of each department and possibly each individual and listen to what they say.

Hire well and let them do their work.

The first point deals with competition. Airbnb, the world´s fastest growing hospitality company, does not own hotel buildings; Uber, the fastest growing taxi operator, has no cars – in the last 20 years we have been watching an awful lot of start-ups turning their industries upside down (like Amazon, Linkedin, Netflix, to name a few more). Today, technological innovation is pushing managers into the obligation to see their own industries from different angles. Knowing what customers want and what they are willing to pay is a key success factor more than ever before.

As a result, the three types of price discrimination discussed above will increasingly converge into a single instrument. It is a further shift from inventory management (how much do we sell at a given price?) to customer equity management (what is the right price for a given transaction?)[7].

The second point deals with implementing change. There is no shortage of providers of advanced technology (data analysis, online reservations, automated billing, etcetera) allowing companies to model, document, optimize and eventually automate processes. However, software implementation often leads to confusion among managers and employees alike. People within the entire organisation (including the owner-manager or the CEO who makes the decision) are going to have concerns, uncertainties and fears to overcome, and everyone involved may temporarily experience a drop in productivity.

Professional staff training and education on how to handle new technologies and at the same time leaving as much room as possible for everyone to work through the changes at one's own pace is therefore crucial. Yet it also involves selling not only by price but by value, too, and knowing when to gently say 'no´ to a customer inquiry.

Rule number three deals with people. The game for the lowest price goes on (and perhaps never stops). The good news from the marketplace is that it is not all about price. As we see across almost all industries, customers have different degrees of sensitivity to price and quality[8]. It is up to revenue management (or, better, the new customer equity management) to explore these degrees of sensitivity and create new offers for satisfying this very demand.

Employees with enough room to use own initiative play a central role as not everything can be read from statistical data. Frontline staff, for example, are the customers' first point of contact for inquiring about price changes, and they first instruct customers on how to use the benefits of new offers. Service staff know about the sources of poor quality within the organisation, and the finance department can tell what regular (and lost) customers mean to the company. Let's bring this knowledge on board, e.g. through frequent staff meetings to discuss how demand may evolve, the causes of recent complaints, denials, no-shows, and other data that effect the business.

The author is founder of afersu.com. He maybe contact at marianm@afersu.com.

[1] The terms Revenue Management and Yield Management (sometimes also called Dynamic Pricing) are often used interchangeably. Sometimes it is claimed there is a subtle difference in that Yield Management deals with selling the core product (e.g., the rooms in a hotel or the seats on a plane) while Revenue Management has a broader picture considering also cross-selling issues (e.g., selling laundry services in a hotel or food on a plane).
For this article, the difference does not matter as the principles described remain the same. I refer here to Revenue Management as a collective term.

[2] Technically, the graph shows the inverse demand curve. A demand curve maps the quantity as a function of price, thus plotting the price on the horizontal x-axis and quantity on the vertical y-axis. We denote: Q = a – b(P), where *Q* is the quantity of demand, *a* equals the intercept (all factors affecting the demand other than price, e.g., quality and fashion), *b* is the slope of the curve (the number of units the quantity changes for each unit of a price change) and *P* is the price. Contrary to that, an inverse demand curve maps the price as a function of quantity, plotting the quantity on the horizontal x-axis and price on the vertical y-axis. Technically, you solve the above equation for P = (a – Q) / b

[3] A fourth-degree price discrimination is often mentioned. This tactic is based on different costs occurring for the supplier. For example, an airline may charge a higher ticket price if a passenger calls in before his departure to order a vegan menu that may cost more than the regular one. For reasons of simplification, I skip this here.

[4] The practice of setting prices individually has long made its way into the digital space. Amongst other tactics, some e-commerce businesses identify a user's OS (Operating system). The idea behind it is that those with more expensive devices (such as a smartphone by Apple) are supposed to show a higher willingness to pay than those owning a no-name product.
However, this can result in a harmful customer backlash, especially if price differences appear to occur arbitrarily. Some time ago, online retailer Amazon applied such a tactic and later said it was "a mistake". See: Heffernan, Virginia: "Amazon´s Prime Suspect", New York Times Magazine, 6 August 2010, https://www.nytimes.com/2010/08/08/magazine/08FOB-medium-t.html?smid=tw-share, accessed 1 October 2018).
Due to bad experience, many outlets have turned to some sort of "reverse" revenue management. If customers search for something on a shop's website, they encounter a recommendation list. Prices for each product are the same for all people, but suppliers adjust the recommendation list according to people's purchase history –those who bought higher-priced products in the past are shown higher-priced products again as they are expected to spend more.

[5] Belobaba, Peter: "Airline Travel Demand and Airline Seat Inventory Management," Flight Transportation Laboratory Report R87-7, Massachusetts Institute of Technology (MIT), May 1987

[6] In Excel we use NORMDIST and NORMINV for calculating the normal CDF and normal inverse CDF, respectively. In the above case, for example, we are using the expected probability 0.267 in the function NORMINV (0.267; 42.08;9.83) and Excel returns x=35.96. Now, if we used this x in the function NORMDIST(35.96;42.08;9.83;TRUE) we got p=0.267. That demonstrates that the one is the reverse of the other.
The Boolean value TRUE in the NORMDIST function refers to the cumulative distribution function, FALSE returned the probability density function which is not needed here.

[7] For the hotel industry, see, for example:
Noone, Breffni M.; McGuire, Kelly A.; Rohlfs, Kristin V.: "Social Media Meets Hotel Revenue Management: Opportunities, Issues and Unanswered Questions", Journal of Revenue and Pricing Management, 10(4) 2011, pp. 293-305).
http://corsi.unibo.it/emt/Documents/Social%20Media%20Meets%20Revenue%20Mngmt%202011.pdf (accessed 1 October 2018)

[8] Rao, Akshar R.; Bergen, Mark E.; Davis, Scott: „How to Fight a Price War", Harvard Business Review, March-April 2000 (https://hbr.org/2000/03/how-to-fight-a-price-war, accessed 1 October 2018)