Dynamic pricing

Part 2

The model-driven approach

We have an e-commerce site where we offer a single product.

We are able to split the customers visiting our site in many *segments*, differentiated by profile. Segments can be based on demographics attributes (e.g. gender or age), or on behavioral attributes (what they visited or purchased in the past), or in other ways.

We are able to offer a different price to each segment.

We want to find the best *pricing policy*, i.e. the association of a different price to each segment which gives us the greatest overall revenue (in expectation).

The basic equation is

$$R = p \times d(p)$$

where

R is the total revenue, given by the price *p* multiplied by the demand *d*.

The basic equation is

$$R = \sum_i R_i = \sum_i (p_i \times d_i(p_i))$$

where

- R is the total revenue,
- R_i is the revenue from segment i, given by the price offered to the segment p_i multiplied by the demand of the segment $d_i(p_i)$.

It is very similar to the equation we used when managing the dynamic pricing problem as a MAB problem.

The same assumptions still hold.

$$R = \sum_i R_i = \sum_i (p_i \times d_i(p_i))$$

This time we use a different approach, a *model-driven* one.

First, we analyze historical data in order to build a predictive model of demand, answering to the question *Which will be the selling rate for each price offered to a customer of each segment.*

Then, we design methods to select the best pricing policy based on the predictive model.

Demand prediction

In its simplest form, a predictive model of the demand can be build applying a regression to past sales data.

A different regression can be applied to each segment.

Assume we have a dataset about sales for segment *Men* and make a linear regression on it.

Price	Demand
50	50
54	49
57	41
60	40
62	33
68	32
72	25
75	22
81	21

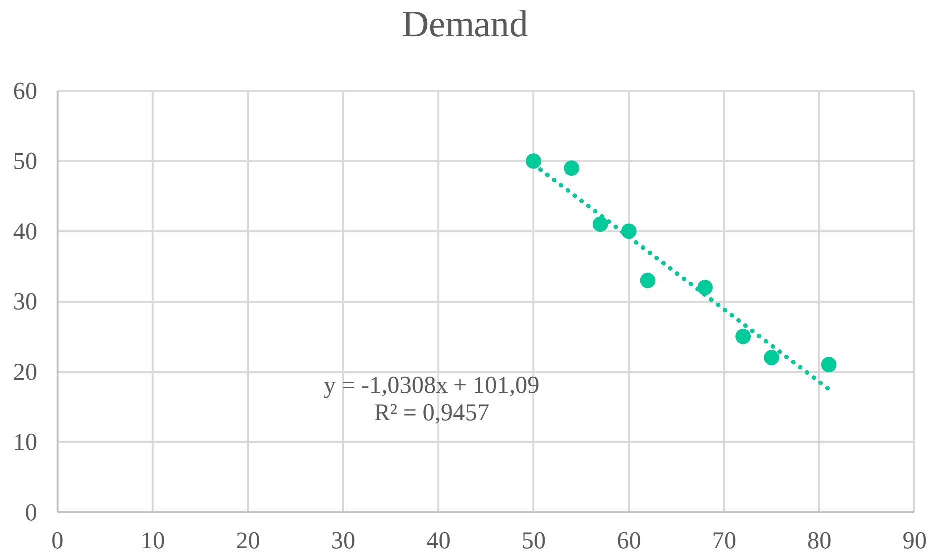

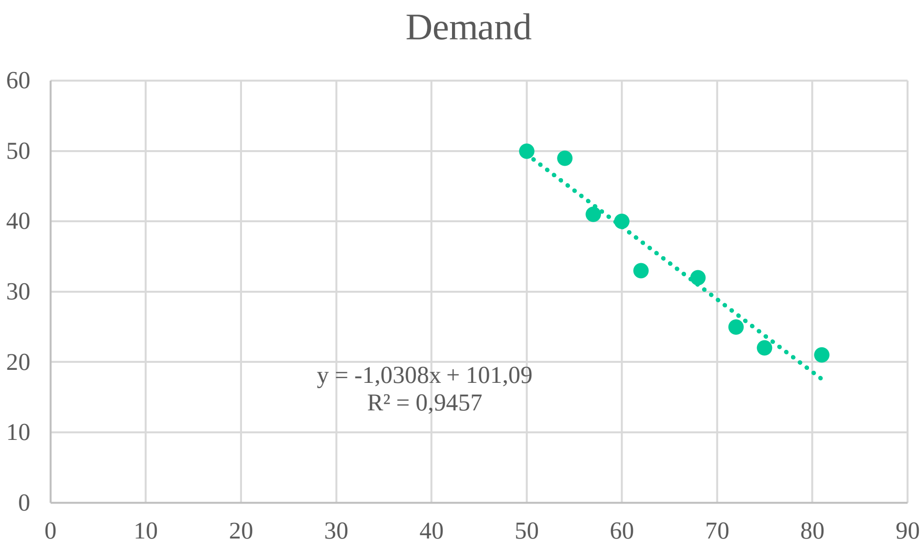

We have a predictive model for demand:

$$d = 101.09 - 1.0308 \cdot p$$

E.g., the predicted demand for price 70€ is

$$d = 101.09 - 1.0308 \cdot 70 = 28.93$$

Now we can use this equation in order to predict the utility of each price we can show to a customer on our site.

E.g., if offering a customer a price 70€, we expect to sell 28.93 unit of our product, i.e. to gain 28.93 customers buying.

Once we have built a predictive model demand/price for each segment, we can find the best pricing policy.

This is the topic of the next section.

Of course, the linear regression method is not the only possible one.

We could use logistic regression as well, or any other kind of predictive modeling technique.

Single segment problem

In its simplest form, the pricing optimization problem can be solved with the classic method of derivatives.

Given a function y = f(x), the derivative of y on x is

$$\lim_{\delta \to 0} \frac{f(x+\delta) - f(x)}{\delta}$$

The derivative can be approximated with the *incremental ratio* when the variation δ is very small:

$$\frac{f(x+\delta) - f(x)}{\delta}$$

If x is a stimulus and y a response, we use the variation in response when the stimulus is incremented of one small quantity.

E.g. the variation in sales of a product when we increment the price of 1 $.

This is the *marginal variation.*

E.g. saying that the derivative of the demand *y* at the price *x* = 50 is -3 means that if we increase the price from 50 to 51 then we sell 3 less product units.

$$\frac{f(x+\delta)-f(x)}{\delta}$$

In order to find the optimal price, we write an equation for the revenue as function of the price.

Then we impose the derivative must be equal to zero. The solution gives us the optimal price.

[To be more rigorous, we have to impose that the derivative of the derivative must be negative, but let us omit this.]

This is a special case of the two-segments problem we are going to see.

Pricing optimization problem

We have two segments, say *Females* and *Males* or any other couple.

The demand predictive models for these segments are:

$$d_1 = 100 - p_1$$

$$d_2 = 120 - 2p_2$$

Say we obtained those models with a linear regression.

Note in both segments the demand is a decreasing function of the price, as intuitively it should be.

We want to find two prices p_1 and p_2 which maximize overall revenue:

$$r = p_1 d_1 + p_2 d_2$$

This problem is interesting if we have a limited *capacity*, i.e. a limited number of product units we can sell. Let us assume that capacity is 40:

$$d_1 + d_2 <= 40$$

Our optimization problem can be formulated in this way

$$\text{maximize } r = p_1 d_1 + p_2 d_2$$

Such that

$$d_1 = 100 - p_1$$

$$d_2 = 120 - 2p_2$$

$$d_1 + d_2 <= 40$$

We have an *objective* (an expression to be maximized) and three *constraints* (equations or inequations to be satisfied).

We are going to see how it can be solved.

Analytical solution

We rewrite the revenue as a function of the segment demands, which are in turn functions of the segment prices.

[Remember that prices are our decision levers: demands are consequences of prices].

Using the demand/price equations

$$d_1 = 100 - p_1$$
$$d_2 = 120 - 2p_2$$

we write

$$r = (100 - d_1)d_1 + \tfrac{(120 - d_2)}{2} d_2 = 100 d_1 - d_1^2 + 60 d_2 - \tfrac{1}{2} d_2^2$$

We have expressed prices as functions of demands.

It can sounds counterintuitive, but it is a normal convention.

$$r = (100 - d_1)d_1 + \tfrac{(120 - d_2)}{2} d_2 = 100 d_1 - d_1^2 + 60 d_2 - \tfrac{1}{2} d_2^2$$

The derivatives of revenue as function of demands are:

$$J_1(d_1) = 100 - 2d_1$$
$$J_2(d_2) = 60 - d_2$$

(we used well-known formulas form the differential calculus).

These are the *marginal utilities* of the demand in the two segments: they measure the impact of selling one more product unit in a segment on revenue of that segment.
To add a demand unit in a segment, we have to decrement the price of a certain amount.
Then, J_1 measures the marginal revenue of the decision of changing the price p_1 of the amount which causes a demand d_1+1 in the first segment.
Looking at the demand-price equation, we see that this price change is -1.

$$d_1 = 100 - p_1$$

$$d_2 = 120 - 2p_2$$

$$J_1(d_1) = 100 - 2d_1$$

$$J_2(d_2) = 60 - d_2$$

$$r = (100 - d_1)d_1 + \frac{(120 - d_2)}{2} d_2 = 100d_1 - d_1^2 + 60d_2 - \tfrac{1}{2} d_2^2$$

Let $p_1 = 70$. Then $d_1 = 100 - 70 = 30$.

If we set p_1 to 69 we get d_1 raising to 31.

What is the impact of this decision on first segment?

It is $J_1(30) = 100 - 2 * 30 = 40$.

Indeed revenue grows from 70 * 30 = 2100 to

69 * 31 = 2339, with an increment of 39.

The result is not exactly 40 because we are using the incremental ratio, which is only an approximation of the derivative.

Yet the concept is the same.

Two demands d_1 and d_2 are optimal if they give us the greatest possible revenue.

From the differential calculus we know that if the two demands are optimal, then their derivatives must be equal:

$$100 - 2d_1 = 60 - d_2$$

or, in simplified form,

$$2d_1 - d2 = 40$$

But it is also

$$d_1 + d_2 = 40$$

Let us solve the two equations linear system

$$d_1 = 26.67 \text{ and } d_2 = 13.33$$

This is the optimal couple of demands.

$$d_1 = 26.67 \text{ and } d_2 = 13.33$$

Let us round the result

$$d_1 = 27 \text{ and } d_2 = 13.$$

These values are really optimal. No other couple of demands can give a greater revenue.

Therefore, the optimal prices are those associated with the optimal demands:

$$p_1 = 100 – 27 = 73$$

$$p_2 = (120 – 13) / 2 = 53.5$$

For $d_1 = 27$ and $d_2 = 13$ we have $r = 2666.5$

This is the greatest possible revenue given our demand/price models and the capacity constraint.

Intuitive explanation

We do not know demands generating maximum revenue, i.e. optimal d_1 and d_2.

Yet we know that at those points the two derivatives have to be equal.

Let us think of derivatives as *marginal utilities* of the two segments.

The optimal couple of demands must have two equal marginal utilities.

Indeed, if they were not, then we could maneuver the two prices and manage to move a unit demand from the segment with smaller derivative to the segment with larger derivative.

But in this case the couple d_1, d_2 would not be optimal!

This is an <u>extremely</u> important principle, applicable in a large variety of economic problems.

Optimality conditions

$$\max \sum_{t=1}^{T} r(d(t))$$

$$\sum_{t=1}^{T} d(t) \leq C$$

$$d(t) \geq 0$$

The optimality conditions can be formulated as follows (using the method of Lagrange)

$$J(d^*(t)) = \pi^*$$

$$\pi^*(C - \sum_{t=1}^{T} d^*(t)) = 0$$

$$\pi^* \geq 0$$

$$\max \sum_{t=1}^{T} r(d(t))$$

$$\sum_{t=1}^{T} d(t) \leq C$$

$$d(t) \geq 0$$

Let us rephrase in plain language.

Revenue *r* of a certain segment *t* is a function of *t* and of the demand *d(t)* of that segment.
Indeed, revenue is the product of price and demand, but the price itself is a function of the demand, because of the demand/price equation.

The sum of demand over all segments must not exceed the capacity *C*.

The demand of each segment must be zero or positive, not negative.

$$\max \sum_{t=1}^{T} r(d(t))$$

$$\sum_{t=1}^{T} d(t) \leq C$$

$$d(t) \geq 0$$

$$J(d^*(t)) = \pi^*$$

$$\pi^*(C - \sum_{t=1}^{T} d^*(t)) = 0$$

$$\pi^* \geq 0$$

The star symbol * denotes optimality, e.g. d^* denotes the optimal demand.
Remember that we do not know which is the optimal demand.
Optimality conditions describe the unknown optimal demands with some features they must have.

First condition.
All the derivatives of demands must be equal to a certain number π^* which therefore is the *"equilibrium" marginal utility*.
It is also interpretable as an *opportunity cost of capacity* when the demands are optimal. If we add an unit more to a segment, we decrement our capacity for the other segment, so we incur in a loss which is the other side of the benefit for the first segment.

$$\max \sum_{t=1}^{T} r(d(t))$$

$$\sum_{t=1}^{T} d(t) \le C$$

$$d(t) \ge 0$$

$$J(d^*(t)) = \pi^*$$

$$\pi^*(C - \sum_{t=1}^{T} d^*(t)) = 0$$

$$\pi^* \ge 0$$

Second condition.
If the demands are optimal, then

- either the common marginal utility is zero (we have no interest in selling more to any segment),
- or the capacity is saturated by the overall demand (we are not able to sell one more unit),
- or both.

These are very intuitive features of the (unknown) optimal demands.

$$\max \sum_{t=1}^{T} r(d(t))$$

$$\sum_{t=1}^{T} d(t) \le C$$

$$d(t) \ge 0$$

$$J(d^*(t)) = \pi^*$$

$$\pi^*(C - \sum_{t=1}^{T} d^*(t)) = 0$$

$$\pi^* \ge 0$$

Third condition.
The common marginal utility cannot be negative (otherwise, we have already sold too much units).

Note that it is not necessary to saturate the capacity.
Indeed, in order to sell more to a segment, we have to decrement the price for that segment. The smaller price applies to every unit of product, so increasing sales can decrement revenue as well as increment it.

The Hill-Climbing algorithm

Let

$$p_1 = 10 - d_1$$
$$p_2 = 10 - 2d_2$$

Then

$$J_1(d_1) = 10 - 2d_1$$
$$J_2(d_2) = 10 - 4d_2$$

We have C = 6 capacity unit.

What is the best way to allocate them?

The hill-climbing algorithm assigns a demand unit to the segment with the greater utility at each step.

This makes the marginal utility to decrease for the chosen segment, making it less attractive at the next step.

If the marginal utility is decreasing, then the algorithm converges to an optimal solution.

In practice, very often this hypothesis holds: the marginal utility is really decreasing.

k	$\mathbf{d_1}$	$\mathbf{d_2}$	$\mathbf{J_1(d_1)}$	$\mathbf{J_2(d_2)}$
0	0	0	10	10
1	1	0	8	10
2	1	1	8	6
3	2	1	6	6
4	3	1	4	6
5	3	2	4	2
6	4	2	2	2

At step 0 we start with the temporary solution

$$d_1 = 0 \text{ and } d_2 = 0.$$

At step1 we assign an additional unit to the segment with the greater J. Being J_1 and J_2 are equal, we can choose arbitrarily.

Say we choose segment 1.

This choice increments d_1, but decrements J_1, changing it from 10 to 8.

The reason is that decreasing the price to increase the offer, we make further increments less profitable.

J_1 decreases of 2 units because equation $J_1 = 100 – 2\ d_1$.

Under the hypothesis that marginal utilities are decreasing, the choice makes similar future choice less attractive.

At step 2 it is $J_2 > J_1$, then we assign the next demand unit to segment 2.

We continue until we are allowed (because capacity is not saturated yet) and until we gain something (because the marginal utilities are positive).

When one of these conditions is no more satisfied, we stop.

We cannot do better anymore.