ORIE 5355

Lecture 8: Algorithmic pricing: capacity, price differentiation, and competition

Nikhil Garg

Announcements

- HW3 will be released soon

Summary up to now

We want to sell an item

- Only one product type
- No capacity constraints
- No competition from other sellers
- No over-time dynamics
- Allowed to explicitly give different prices to different users

Then: revenue-maximizing price(s) and demand estimation

Maximizing revenue

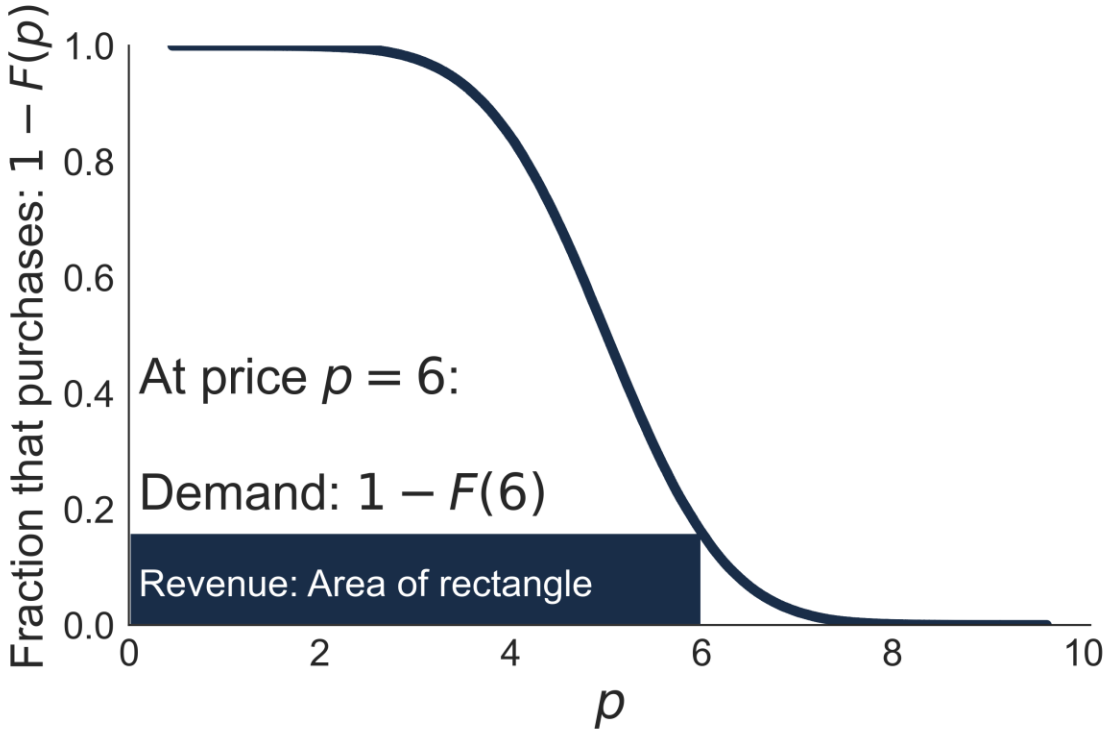
- Expected revenue at price p:

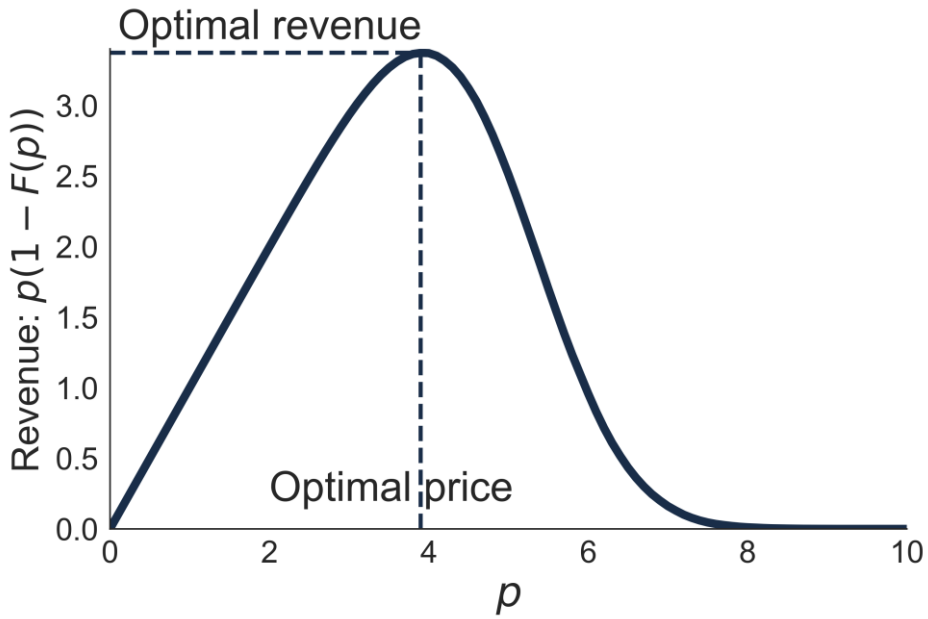
  [Revenue from each sale] x [Demand at price p]

  $p(d(p))$

- Revenue maximizing price:

  $\text{argmax}_p\ p(d(p))$





Demand estimation comments

- Demand estimation and forecasting is probably the *most important and difficult* challenge in revenue management
- Unlike most machine learning challenges, we need to estimate a *function* $F(p)$ [or treat price as a covariate]
- We made a *substantial* assumption that almost never holds in practice: that you have historical data at many different prices p

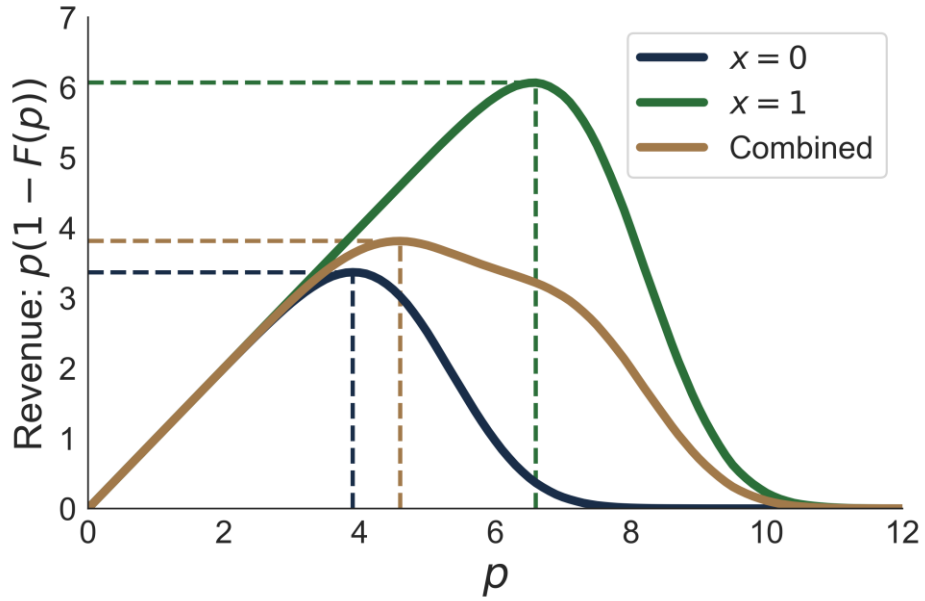
  Requires experimentation!

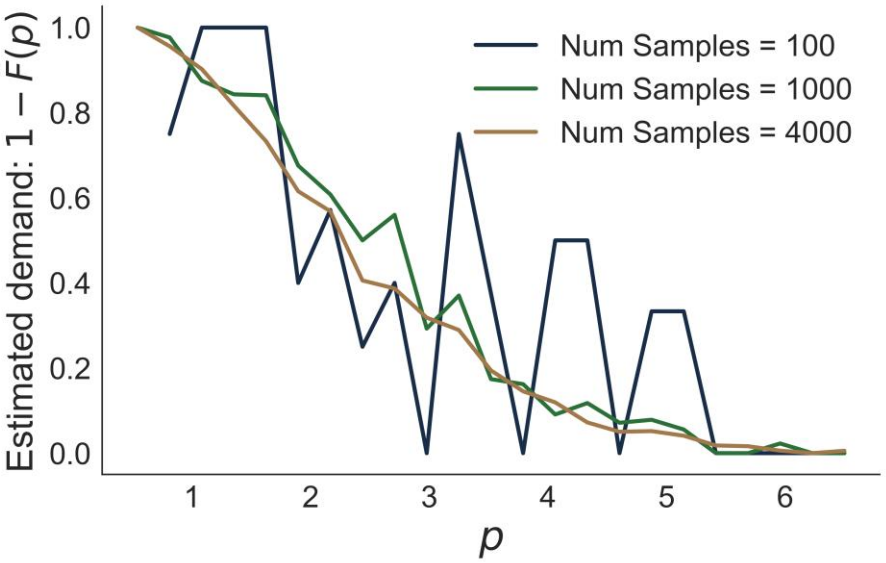
Up to now

- Given a demand distribution $\mathrm{d(p)} = 1 - F(p)$, how to calculate optimal prices

$$\arg\max_{\mathrm{p}} [\mathrm{p} \times \mathrm{d(p)}]$$

- How to estimate demand distributions, potentially as a function of covariates





Plan for rest of today

Many assumptions last time:

- No capacity constraints
- No competition from other sellers
- Only one product type
- Allowed to explicitly give different prices to different users
- No over-time dynamics

We'll peel back some of these assumptions today

Capacity constraints and pricing over time

Setting and examples

You often are trying to sell limited quantities of a good, to many potential customers over time

- Airline tickets – the airline “wastes” a seat that’s unsold
  - Same for concerts, sports, any event with a fixed date
  - Clothes that are going out of season/fashion
  - Electronics that become obsolete over time
- Any retail setting with inventory constraints
- Often 2 competing effects:
  - The items become less valuable over time, or you have a deadline to sell them
  - You have less stock over time

Simplified example

- You have 1 copy of the item to sell
- There are 2 time periods, today and tomorrow
  - One customer will come in today, a different one tomorrow
- No covariates
- No “discounting” (a dollar tomorrow is as valuable as a dollar today)
- You already have a good estimate of $\mathrm{d}(\mathrm{p})$

What price p_0 do you set today? What price p_1 do you set tomorrow?

A couple of observations

What I do today depends on what I can/will do tomorrow.

- I can't set p_0 unless I know how I will set p_1 in *each scenario*. (whether I sold the item today, or whether I didn't).
- I have to "simulate" the future

If I don't sell the item today, then *tomorrow* I am solving the same problem that we solved in class last time:

- Maximizing revenue for a single buyer/without capacity considerations
- => The price for tomorrow will be same as simple revenue maximizing price

$$p_1 = \arg\max_p [p \times d(p)]$$

Not true for the price today:

- If I sell the item today, then I lose out on a potential sale tomorrow
- If I don't sell the item today, I get another chance tomorrow

=> I should "take a risk" today to try to sell at a higher price

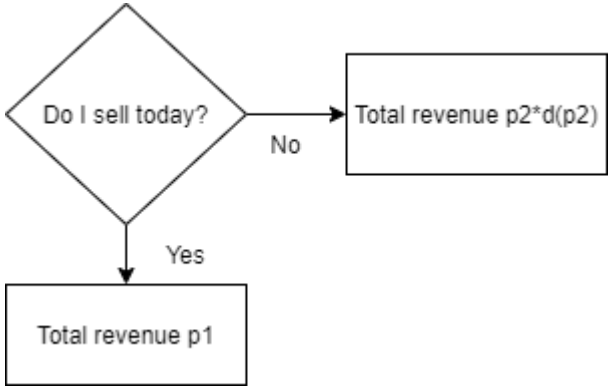
Backward Induction – solve the last day first, go backwards

Solving the example: "Bellman equation"

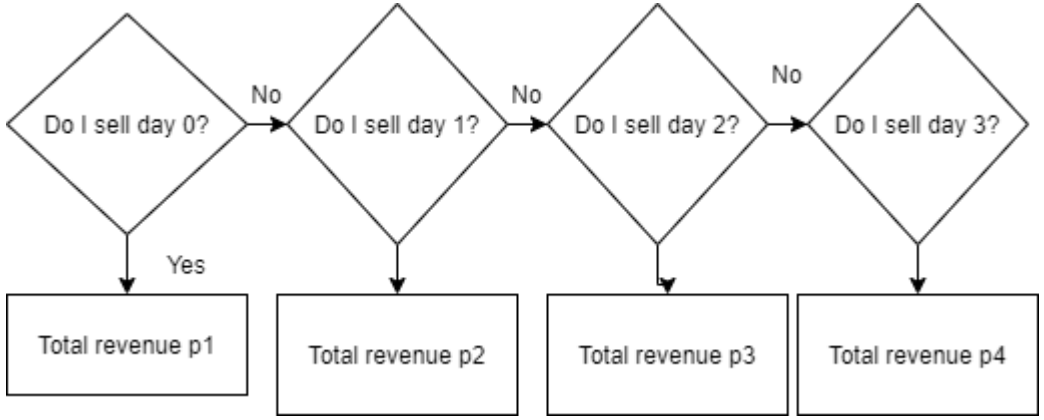
- If I don't sell today: (happens with probability $1 - d(p_0)$)
  - Then my revenue today is 0
  - Then the expected revenue tomorrow is: $p_1 d(p_1)$
- If I *do* sell today: (happens with probability $d(p_0)$)
  - My revenue today is p_0
  - Then the expected revenue tomorrow is 0
- So, my overall expected revenue is:

$$d(p_0)(p_0 + 0) + (1 - d(p_0))(0 + p_1 d(p_1))$$

- p_1 easy to solve – does not depend on p_0
- *Given* p_1, the above revenue function is only a function of p_0 *=> Can optimize* p_0



Bellman equation generally



- You can generalize this idea to selling any number of items sequentially for T days
- Start from Day $\mathrm{T}-1$: If you still have an item, do single-shot maximization
- Day $\mathrm{T}-2$: Given Day $\mathrm{T}-1$ price, you know expected reward if you still have an item to be sold after Day $\mathrm{T}-2$. And so, you can calculate optimal price for Day $\mathrm{T}-2$.
- Now, you have the expected reward if you still have an item to be sold after Day $\mathrm{T}-3$…

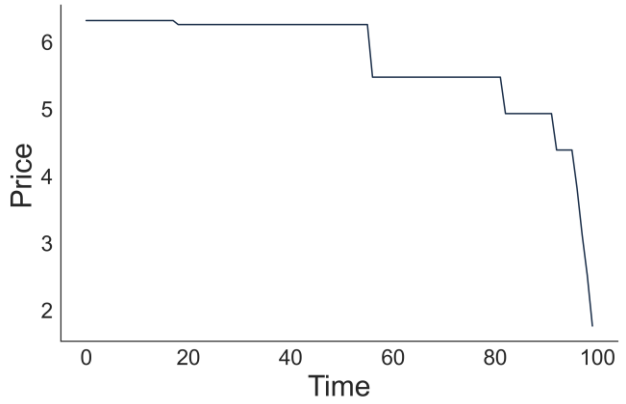
More Bellman equation

- Let V_t denote: "Expected revenue if I still have an item to sell on day t"

$$V_{T-1} = p_{T-1} \times d(p_{T-1})$$

$$V_{T-2} = [p_{T-2} \times d(p_{T-2})] + (1 - d(p_{T-2}))\, V_T$$

- Above means: "Value today is revenue today if I sell the item today, or tomorrow's expected revenue if I don't sell the item today"
- For each t, given V_{t+1} we can calculate optimal price p_t
- Keep iterating until you have prices $p_0 \ldots p_{T-1}$
- Resulting V_0 is my expected revenue given these prices



Questions pause

Bellman equations: a general idea

- Constructing a tree to reason about what happens tomorrow, and then iterating backwards, is a powerful + flexible algorithmic technique: "dynamic programming"
- Example: What if you have 5 copies of the same item?

  Let k denote how many copies of the item I have. Then:

$$V_{t,0} = 0 \text{ for all } t$$

$$V_{t,k} = \max_{p_{t,k}} d(p_{t,k})\left[p_{t,k} + V_{t+1,k-1}\right] + \left(1 - d(p_{t,k})\right) V_{t+1,k}$$

  If I sell an item today: Revenue today, plus future revenue from 1 less item

  If I don't sell: Future revenue from same number of items

  Competing effects: Now, less capacity over time → prices should go up (but less time to sell, so prices should go down).

Capacity constraints + over-time pricing in practice

- Dynamic programs/bellman equations are powerful, but often the real world is too complicated
  - Uncertainty in future capacity
  - Future actions of competitors
  - Future demand distributions
  - “Long time horizons” (T is big)
- In theory, dynamic programming can handle the above. In practice, hard to know how to calculate future value.

Approximating dynamic programming

- In the recommendations module, we created “score”(or “index”) functions:
  - Consider future users, through capacity and avg ratings terms in the score function
- With 1 item: V_{t+1} represents my “opportunity cost” if I sell an item today that I could have sold tomorrow.

  Also interpret as “safety net”: if fail to sell the item today, still earn V_{t+1} in expectation
- Instead of doing a full Bellman equation, estimate V_{t+1} through some other means, then plug into the decision problem for today (finding price p_t)
  - Can construct it like we did score functions for recommendations
  - AlphaGo to play Go: V_{t+1} is partially estimated via a neural network

Pricing with capacity summary

- Just like in recommendations, have to think about potential future demand
- Here, potential future demand lets us be “more aggressive” by pricing higher today
- If I can summarize future revenue (V_{t+1}) effectively, then I can optimize today’s prices
- Dynamic programming: start from the end!
- We assumed that customers can’t strategize on when to come – not true!

Questions?