

Revenue management under demand uncertainty: the case for a robust optimization approach

Aurélie Thiele
Lehigh University
aurelie.thiele@lehigh.edu

Air Products
Allentown, PA
August 10, 2006

What is revenue management?

- **Revenue management** refers to the collection of techniques available to the decision-maker to maximize revenue.
- It addresses 3 categories of **demand-management** decisions:
 - ① Structural decisions: Which selling format to use (posted prices, auctions), which market segmentation mechanisms to use, if any (e.g., impose Saturday-night stay to distinguish between business and leisure travelers in airline RM)...
 - ② Price decisions: how to price over time, how/when to markdown (discount)...
 - ③ Quantity decisions: whether to accept or reject an offer to buy, how to allocate capacity to different segments (e.g., how many student tickets to sell on a flight)...

Decision-Making under Uncertainty

- The fundamental hurdle in maximizing revenue is that demand is **uncertain**.
- Should we maximize the *expected* revenue?
 - 1 That is the approach taken by traditional decision-making under uncertainty.
 - 2 That would mean that we value equally a revenue of \$1,000,000 for sure and a revenue of \$10,000,000 with probability 0.1 and 0 with probability 0.9. (Same expected value.)
 - 3 In reality, decision-makers are **averse to risk**. They prefer to achieve the same revenue, or a revenue just slightly lower, with less risk.

Decision-Making under Uncertainty

Three drawbacks to traditional decision-making under uncertainty:

- 1 To incorporate risk (the few times it is done), decision-makers introduce a **utility function** (concave and increasing).
 - *In theory*, such a function would be elicited by a series of questions about an individual's preferences.
 - *But in practice*, nobody knows what their own utility function is, and except in finance, no one can help them figure it out.
- 2 **Probability distributions** are difficult to estimate (not enough historical data, no guarantee that future demand will obey past distribution...)
Example: “underdog” basketball team which makes the March Madness tournament for the first time in the last 25 years.
- 3 Problem quickly becomes intractable when multiple time periods (curse of dimensionality in dynamic programming).

Revenue Management under Demand Uncertainty

Can we develop an approach to revenue management under demand uncertainty that:

- 1 Takes risk into account, i.e., tries to avoid very adverse outcomes?
- 2 Does not require the precise knowledge of the probability distributions?
- 3 Is computationally tractable?

Revenue Management under Demand Uncertainty

Can we develop an approach to revenue management under demand uncertainty that:

- 1 Takes risk into account, i.e., tries to avoid very adverse outcomes?
- 2 Does not require the precise knowledge of the probability distributions?
- 3 Is computationally tractable?

YES using robust optimization

What is robust optimization?

A number of academics (Nahmias, Simchi-Levi, Sheffi in their respective books) have identified two principles as fundamental to the practice of decision-making under uncertainty:

- 1 Point forecasts are meaningless (because they are always wrong) and should be replaced by *range forecasts*.
- 2 *Aggregate forecasts* are more accurate than disaggregate ones.

Example: It is easier (yields more accurate forecasts) to estimate the demand for ties during any week than the demand for red ties with yellow stripes that same week.

Similarly, it is easier to estimate quarterly or yearly demand than weekly demand for any product.

Example

- HP has a plant in Singapore and a plant in Vancouver to produce printers.
- It produces the sure part of the demand in Northern America (lower bound of the range forecast) in Singapore, where costs are cheaper but transportation to the market takes longer.
- Then when the uncertain part of the demand is realized, it produces the missing printers in the Vancouver plant, which is more expensive but is closer to the market.
- Moreover, to take advantage of aggregate forecasts, it produces generic European printers and customizes them for specific countries as late as possible.

Fundamental insight

Canceling-out effect of uncertainty

In sums of independent random variables (r.v.), some r.v. will be higher than their mean and some will be lower, but if the number of r.v. is big enough, it is very unlikely that they will all be higher than expected or all lower!

The realistic range taken by a sum of r.v. is much smaller than the sum of the ranges of the individual r.v.

This is illustrated by the law of large numbers in the case of i.i.d. random variables.

Robust optimization

- **Robust optimization** can be defined as “reasonable-worst-case optimization.”
- Random variables are modeled as **uncertain parameters** belonging to a known uncertainty set.
- The decision-maker protects revenue against the worst-case values taken by the parameters within that set.

$$\max_{\mathbf{x} \in \mathcal{X}} \min_{\mathbf{d} \in \mathcal{D}} R(\mathbf{x}, \mathbf{d})$$

- The choice of the uncertainty set is critical to the performance of the approach (in particular to avoid over-conservatism): you only want the set to include reasonable outcomes.

The Origins of Robust Optimization

- The concept of robust optimization (protecting a system against unknown but bounded disturbances) finds its origins in engineering, specifically, *control theory*.
- Motivation: unmanned robots should not topple over just because there was a little bump on the road. Rockets should not go off trajectory just because the mass of fuel decreased due to consumption.
- The first work of that kind in optimization is due to Soyster (1973), but it was overly conservative and therefore was not adopted in practice (*range forecasts without aggregation*.)
- In the mid-1990s, research teams led independently by Ben-Tal and Nemirovski and El Ghaoui addressed the issue of *overconservatism*, making robust optimization much more attractive.

Literature Review

- These researchers focused on **ellipsoidal uncertainty sets**. However, this increases the complexity of the problem considered (the robust counterparts of linear problems are second order cone problems, for instance).
- Starting in 2001, Bertsimas and his co-authors developed a robust optimization approach based on **polyhedral uncertainty sets**, which does not change the class of the problem considered.
- Robust optimization with polyhedral uncertainty sets has been successfully applied to several classes of management problems (portfolio, inventory and revenue management).

Technical details

- The polyhedral uncertainty set proposed by Bertsimas et. al. consists of:
 - 1 *range forecasts* for all uncertain parameters (e.g., demand d_i of product i is in $[\bar{d}_i - \hat{d}_i, \bar{d}_i + \hat{d}_i]$ for each i),
 - 2 *a budget of uncertainty constraint* that limits the number of parameters that can deviate from their mean:

$$\sum_{i=1}^n \frac{|d_i - \bar{d}_i|}{\hat{d}_i} \leq \Gamma$$

- $\Gamma = 0$ yields $d_i = \bar{d}_i$ for all i : back to the nominal problem.
- $\Gamma = n$ makes the constraint redundant with $\bar{d}_i - \hat{d}_i \leq d_i \leq \bar{d}_i + \hat{d}_i$ for all i . Very conservative.
- $\Gamma \in (0, n)$ allows to adjust the degree of conservatism.

Technical details

- Robust optimization problems are max-min problems.
- When the objective function is bilinear in decision variables and uncertain parameters, use strong duality in linear programming.
- That allows us to rewrite the max-min problem as a max-max problem.
- Big linear maximization problem that can be solved in a **single step**.
- Need to pick uncertainty sets (and sometimes relax the problem) to keep that bilinearity, or use other tricks.

Production-distribution systems, Example 1

- A warehouse manager must decide how many products to order to minimize cost.
- The warehouse supplies n stores and it is only possible to order once for the whole planning period.
- The warehouse has an initial inventory of zero.
- It incurs a unit shortage cost s per unfilled item and a unit holding cost h per item remaining in the warehouse.
- Store demands are assumed to be i.i.d. with a symmetric distribution around the mean.
- They have all the same range forecast $[\bar{w} - \hat{w}, \bar{w} + \hat{w}]$ with \bar{w} the nominal forecast, common to each store.

Example 1

- Let x be the number of items ordered by the decision-maker, whose goal is to minimize the total cost:

$$\max \left\{ h\left(x - \sum_{i=1}^n w_i\right), s\left(\sum_{i=1}^n w_i - x\right) \right\}$$

- The robust problem for a given budget of uncertainty Γ is:

$$\begin{aligned} \min \quad & Z \\ \text{s.t.} \quad & Z \geq h(x - n\bar{w} + \Gamma\hat{w}), \\ & Z \geq s(-x + n\bar{w} + \Gamma\hat{w}), \\ & x \geq 0. \end{aligned}$$

- The solution to this problem is available in closed form:

$$x_{\Gamma} = n\bar{w} + \frac{s - h}{s + h} \Gamma\hat{w}.$$

Example 1

- The optimal objective is then: $C_{\Gamma} = \frac{2hs}{s+h} \Gamma \hat{w}$.
- If shortage is more penalized than holding, the decision-maker will order **more** than the nominal aggregate forecast.
- The excess amount will then be proportional to the maximum deviation $\Gamma \hat{w}$, as well as the ratio $(s-h)/(s+h)$.
- Assume that the variance of each store demand is equal to σ^2 .
- C_{Γ} is an **upper bound** to the true cost with probability $1 - \epsilon$ when Γ is at least equal to $(\sigma/\hat{w})\sqrt{n} \Phi^{-1}(1 - \epsilon/2)$.
(Independent of the cost parameters h and s .)
- With $n = 100$ and $\hat{w} = 2\sigma$, the actual cost falls below C_{10} with probability 0.95.

Example 1

- Two dangers: (i) “not worrying enough” about uncertainty (i.e., only considering the nominal values) and (ii) “worrying too much” (i.e., only considering their worst-case values).
- We compute the expected cost for the worst-case probability distribution of the aggregate demand W .
- Distribution of W is symmetric with mean $n\bar{w}$ and support $[n(\bar{w} - \hat{w}), n(\bar{w} + \hat{w})]$, and W falls within $[n\bar{w} - \Gamma\hat{w}, n\bar{w} + \Gamma\hat{w}]$ with probability $2\phi - 1$ where $\phi = \Phi((\Gamma - 1)/\sqrt{n})$.
- Let \mathcal{W} be the set of such probability distributions.
- We have the following bound:

$$\max_{W \in \mathcal{W}} E[\max\{h(x - W), s(W - x)\}] =$$

$$\hat{w}(s + h) \left[n(1 - \phi) + \Gamma \left\{ \phi - \frac{s^2 + h^2}{(s + h)^2} \right\} \right].$$

Example 1

- The budget of uncertainty minimizing this bound does not appear to be sensitive to the value of the cost parameters.
- Accounting for a *limited* amount of uncertainty via the robust optimization framework leads to significant cost benefits.
- A decision-maker implementing the nominal strategy will be penalized for not planning at all for randomness, i.e., the aggregate demand deviating from its point forecast.
- But protecting the system against the most negative outcome will also result in lost profit opportunities.
- The robust optimization approach achieves a trade-off between these two extremes.

Comments

- In practice, there is an important **time dimension** to revenue management that is not taken into account in the previous model.
- Framework can be modified to incorporate time.
- However, mostly you do it in an open-loop manner (don't use that you'll get more data later).
- Example about the underdog basketball team: You only want to order more T-Shirts when the team starts winning.
- **Fundamental idea:** Multi-period robust optimization should be adaptive (incorporate past observations).

Problem Setup

- Customer orders arrive sequentially.
- They can request multiple products (e.g., chemicals).
- If the product is not in stock it needs to be produced.
- Setup times between runs makes it more efficient to produce all orders of a same product at the same time.
- If plant is currently producing chemical A and customer 1 arrives and wants A, 2 arrives later and wants B, 3 arrives even later and wants A (while plant is still producing A), it makes more sense to fulfill customer 3's request before customer 2.
- Customers should not necessarily be served on a First-Come-First-Serve basis.

Problem Setup, Cont'd

- Lead time needs to be quoted (approx. order fulfillment time).
- “Ships within 3 to 5 business days” on Amazon.com is a *range forecast*!
- You can fulfill the order early even if the customer paid less for that longer lead time, if another customer paid for the same product and a short lead time.
- Example: I pay for Standard Shipping on Amazon.com, but if I order on Sunday, my order usually ships the next day.
- **Key**: Range forecasts are no longer just inputs to the robust optimization problem, but also outputs.
- The fact that customers can “jump” in line in front of others makes this a very complex problem in stochastic optimization (queues with *changing* priorities and multiple classes).

Robust Optimization Methodology

- Whatever you do will depend on the current state of the system now:
 - 1 which plants are working on which products,
 - 2 what is the estimated completion time of the run (end of production of that chemical),
 - 3 what are the lead times already quoted to customers.
- Robust optimization is used to:
 - 1 Model the random inputs: random order arrivals and quantities,
 - 2 Quote outputs: the range forecast for the current request,
 - 3 Update the production schedule.
- State is updated at each request arrival.

Discuss

- Inventory constraints: raw materials, preemption...
- Production process (dedicated production lines?)
- Price-dependent lead time quotation?
- Dedicated plants? Distribution system?