Selected Readings in Micro: Newsvendor Problem

Zhentong Lu

SOE and IAR, SUFE

April 16, 2017

The Newsvendor Problem (Gallego's Lecture Notes)

- problem: controlling the inventory of a single item with stochastic demands over a single period
  - date back to Edgeworth (1888), who used the Central Limit Theorem to determine the amount of cash to keep at a bank to satisfy random cash withdraws from depositors

The Newsvendor Problem (Gallego's Lecture Notes)

- problem: controlling the inventory of a single item with stochastic demands over a single period
  - date back to Edgeworth (1888), who used the Central Limit Theorem to determine the amount of cash to keep at a bank to satisfy random cash withdraws from depositors
- prototype problem faced by a newsvendor trying to decide how many newspapers to stock on a newsstand before observing demand
  - the newsvendor faced both overage and underage costs if he orders too much or if he orders too little

The Newsvendor Problem (Gallego's Lecture Notes)

- problem: controlling the inventory of a single item with stochastic demands over a single period
  - date back to Edgeworth (1888), who used the Central Limit Theorem to determine the amount of cash to keep at a bank to satisfy random cash withdraws from depositors
- prototype problem faced by a newsvendor trying to decide how many newspapers to stock on a newsstand before observing demand
  - the newsvendor faced both overage and underage costs if he orders too much or if he orders too little
- the Newsvendor Problems is therefore the problem of deciding the size of a single order that must be placed before observing demand when there are overage and underage costs
  - the problem is particularly important for items with significant demand uncertainty and large overage and underage costs

Basic Setup

- D: one period random demand
  - mean and variance: $\mu = E[D]$ and $\sigma^2 = Var[D]$

Basic Setup

- D: one period random demand
  - mean and variance: $\mu = E[D]$ and $\sigma^2 = Var[D]$
- c: unit ordering cost

Basic Setup

- D: one period random demand
  - mean and variance: $\mu = E[D]$ and $\sigma^2 = Var[D]$
- c: unit ordering cost
- $p(> c)$: selling price

Basic Setup

- D: one period random demand
  - mean and variance: $\mu = E[D]$ and $\sigma^2 = Var[D]$
- c: unit ordering cost
- $p(> c)$: selling price
- $s(< c)$: salvage value

Basic Setup

- D: one period random demand
  - mean and variance: $\mu = E[D]$ and $\sigma^2 = Var[D]$
- c: unit ordering cost
- $p(> c)$: selling price
- $s(< c)$: salvage value
- if Q units are ordered, then $\min(Q, D)$ units are sold and $(Q - D)_+ \equiv \max(Q - D, 0)$ units are salvaged

Profit and Cost Function

- expected profit function

$$\pi(Q) = pE[\min(Q, D)] + sE\left[(Q - D)_+\right] - cQ$$

Profit and Cost Function

- expected profit function

$$\pi(Q) = pE[\min(Q, D)] + sE\left[(Q - D)_+\right] - cQ$$

- equivalently, we can rewrite the expected profit as

$$\pi(Q) = (p - c)\,\mu - G(Q)$$

where the expected cost function G is

$$G(Q) = (c - s)\,E(Q - D)_+ + (p - c)\,E(D - Q)_+ \geq 0$$

  - let $h = c - s$, it could be interpreted as the per unit overage cost (holding cost)
  - let $b = p - c$, it could be interpreted as the per unit underage cost (penalty cost)
    - sometimes inflated to take into account the ill-will cost associated with unsatisfied demand

Remark

- maximizing $\pi(Q)$ is identical to minimizing $G(Q)$, which is the standard formulation in Newsvendor literature

Remark

- maximizing $\pi(Q)$ is identical to minimizing $G(Q)$, which is the standard formulation in Newsvendor literature
- suppose there is no demand uncertainty, i.e., $D = \mu$, then the optimal Q is
$$Q^*_{det} = \mu$$
  - in this case, $\pi = (p - c)\,\mu$ because $G(Q^*_{det}) = 0$

Remark

- maximizing $\pi(Q)$ is identical to minimizing $G(Q)$, which is the standard formulation in Newsvendor literature
- suppose there is no demand uncertainty, i.e., $D = \mu$, then the optimal Q is
$$Q^*_{det} = \mu$$
  - in this case, $\pi = (p - c)\mu$ because $G(Q^*_{det}) = 0$
- underlying assumption: price taking, no market power
  - need extensions when thinking about IO applications

Remark

- let $g(x) = hx_+ + bx_-$ and then we have

$$G(Q) = E[g(Q - D)]$$

  - g is convex, so G is convex become expectation is just a linear transformation

Remark

- let $g(x) = hx_+ + bx_-$ and then we have

$$G(Q) = E[g(Q-D)]$$

  - g is convex, so G is convex become expectation is just a linear transformation

- by Jensen's inequality, $G(Q) \geq g(Q-\pi)$, thus the expected profit is lower than it would be in the case of deterministic demand
  - intuition: the loss from incomplete information

Solution

- suppose D is continuous, the FOC of minimizing expected cost is

$$\begin{aligned} G^{'}(Q) &= hE\left[\delta\left(Q-D\right)\right] - bE\left[\delta\left(D-Q\right)\right] \\ &= h\Pr\left(Q>D\right) - b\Pr\left(D>Q\right) \\ &= 0 \end{aligned}$$

  - $\delta(x) = 1$ if $x > 0$ and zero otherwise

Solution

- suppose D is continuous, the FOC of minimizing expected cost is

$$\begin{aligned} G^{'}(Q) &= hE\left[\delta\left(Q-D\right)\right] - bE\left[\delta\left(D-Q\right)\right] \\ &= h\Pr\left(Q>D\right) - b\Pr\left(D>Q\right) \\ &= 0 \end{aligned}$$

  - $\delta(x) = 1$ if $x > 0$ and zero otherwise

- suppose the CDF of D, denoted as F, is continuous and strictly increasing, then we have

$$F(Q) = \Pr(D \leq Q) = \beta \equiv \frac{b}{b+h} = \frac{p-c}{p-s}$$

which implies

$$Q^{*} = F^{-1}(\beta)$$

  - this fractile solution appeared in the classical 1951 paper by Arrow, Harris and Marshak ("Optimal Inventory Policy", *Econometrica*)

Service Level as a Rule-of-Thumb

- the newsvendor solution can be interpreted as providing the smallest supply quantity that guarantees that all demand will be satisfied with probability at least $100\beta\%$
  - the profit maximizing solution results in a service level $100\beta\%$

Service Level as a Rule-of-Thumb

- the newsvendor solution can be interpreted as providing the smallest supply quantity that guarantees that all demand will be satisfied with probability at least $100\beta\%$
  - the profit maximizing solution results in a service level $100\beta\%$
- in practice, managers often specify β and then find Q accordingly (rule-of-thumb)
  - another notion: fraction of demand served/fill-rate, which is defined as $\alpha = E\left[\min\left(D, Q\right)\right]/E\left(D\right)$

Special Cases of Demand Distribution

- normal distribution: $D \sim N(\mu, \sigma)$

$$Q^* = \mu + z_\beta \sigma$$

where $z_\beta = \Phi^{-1}(\beta)$ (Φ is standard normal CDF) is called "safety factor" and the quantity $Q^* - \mu = z_\beta \sigma$ is known as "safety stock"

Special Cases of Demand Distribution

- normal distribution: $D \sim N(\mu, \sigma)$

$$Q^* = \mu + z_\beta \sigma$$

where $z_\beta = \Phi^{-1}(\beta)$ (Φ is standard normal CDF) is called "safety factor" and the quantity $Q^* - \mu = z_\beta \sigma$ is known as "safety stock"

- lognormal distribution: $\log(D) \sim N(\nu, \tau^2)$, with $\mu = E(D) = \exp(\nu + \tau^2/2)$ and $\sigma^2 = Var(D) = \mu^2 \left[\exp(\tau^2) - 1\right]$

$$Q^* = \exp(\nu + \tau z_\beta)$$

and

$$\pi(Q^*) = (p - c)\mu - (h + b)\mu\Phi(\tau - z_\beta) + h\mu$$

Worst Case Distribution

- sometimes there is no theoretical justification for demand to follow a particular distribution, such as Normal or Poisson, one has to work with guess-estimates of the mean and the forecast error/standard deviation

Worst Case Distribution

- sometimes there is no theoretical justification for demand to follow a particular distribution, such as Normal or Poisson, one has to work with guess-estimates of the mean and the forecast error/standard deviation
- Herbert Scarf (1958): there is a close-form solution that minimizes the objective function against the worst possible distribution with a given mean and standard deviation

$$Q^S = \mu + \frac{\sigma}{2}\left(\sqrt{\frac{b}{h}} - \sqrt{\frac{h}{b}}\right)$$

  - implication: order more (less) than the mean demand when $b > h$ ($b < h$)
  - $\left|Q^S - \mu\right|$ increases linearly in σ for $h \neq b$

Worst Case Distribution

- sometimes there is no theoretical justification for demand to follow a particular distribution, such as Normal or Poisson, one has to work with guess-estimates of the mean and the forecast error/standard deviation
- Herbert Scarf (1958): there is a close-form solution that minimizes the objective function against the worst possible distribution with a given mean and standard deviation

$$Q^S = \mu + \frac{\sigma}{2}\left(\sqrt{\frac{b}{h}} - \sqrt{\frac{h}{b}}\right)$$

  - implication: order more (less) than the mean demand when $b > h$ $(b < h)$
  - $\left|Q^S - \mu\right|$ increases linearly in σ for $h \neq b$
- implication: an interesting bound on optimal profit

$$1 - \sqrt{\frac{h}{b}}\frac{\sigma}{\mu} \leq \frac{\pi\left(Q^*\right)}{(p-c)\,\mu} \leq 1$$

  - the slackness of the left bound measures how much gain from knowing the parametric form of F

Multi-Period Models: Setup

- specification: finite horizon, no setup costs, zero lead times

Multi-Period Models: Setup

- specification: finite horizon, no setup costs, zero lead times
- let $D_1, \ldots, D_T$ be the demand for T periods, D_t's are independent random variables

Multi-Period Models: Setup

- specification: finite horizon, no setup costs, zero lead times
- let $D_1, ..., D_T$ be the demand for T periods, D_t's are independent random variables
- assumption: all stockouts are backordered
  - problematic when there is "lost sales", i.e., unsatisfied demand will be lost

Multi-Period Models: Setup

- specification: finite horizon, no setup costs, zero lead times
- let $D_1, ..., D_T$ be the demand for T periods, D_t's are independent random variables
- assumption: all stockouts are backordered
  - problematic when there is "lost sales", i.e., unsatisfied demand will be lost
- c_t: the unit cost in period t

Multi-Period Models: Setup

- specification: finite horizon, no setup costs, zero lead times
- let $D_1, ..., D_T$ be the demand for T periods, D_t's are independent random variables
- assumption: all stockouts are backordered
  - problematic when there is “lost sales”, i.e., unsatisfied demand will be lost
- c_t: the unit cost in period t
- x_t: inventory level at the beginning of t
  - a positive (negative) x_t indicates that x_t ($-x_t$) units of inventory are carried from the previous period

Law of Motion and (Per-Period) Cost Function

- let $y_t - x_t \geq 0$ denote the size of the order in period t, i.e., y_t is the "order-up-to" level (or "inventory position")
  - procurement cost: $c_t \cdot (y_t - x_t)$

Law of Motion and (Per-Period) Cost Function

- let $y_t - x_t \geq 0$ denote the size of the order in period t, i.e., y_t is the "order-up-to" level (or "inventory position")
  - procurement cost: $c_t \cdot (y_t - x_t)$

- inventory's law of motion

$$x_{t+1} = y_t - D_t$$

Law of Motion and (Per-Period) Cost Function

- let $y_t - x_t \geq 0$ denote the size of the order in period t, i.e., y_t is the "order-up-to" level (or "inventory position")
  - procurement cost: $c_t \cdot (y_t - x_t)$

- inventory's law of motion

$$x_{t+1} = y_t - D_t$$

- "end-of-period" cost function

$$G_t(y) = h_t E(y - D_t)_+ + p_t E(D_t - y)_+$$

  - h_t is the overage/holding cost and p_t is the underage/backorder penalty cost

Salvage ($T + 1$) Period

- the interpretation of h_t and p_t for period $t \leq T$ is different from that of period $T + 1$: in period $T + 1$, we would typically salvage remaining items and either produce or reimburse customers if there are backlogs

Salvage $(T+1)$ Period

- the interpretation of h_t and p_t for period $t \leq T$ is different from that of period $T+1$: in period $T+1$, we would typically salvage remaining items and either produce or reimburse customers if there are backlogs
- salvage value: $-C_{T+1}(x_{T+1})$, where C_{T+1} is some cost function

Salvage $(T+1)$ Period

- the interpretation of h_t and p_t for period $t \leq T$ is different from that of period $T+1$: in period $T+1$, we would typically salvage remaining items and either produce or reimburse customers if there are backlogs
- salvage value: $-C_{T+1}(x_{T+1})$, where C_{T+1} is some cost function
- why it is necessary? because it is the terminal condition for the dynamic programming problem

To be continued...

Next Class Reading

- Victor Aguirregabiria: "The Dynamics of Markups and Inventories in Retailing Firms," *The Review of Economic Studies*, 1999.