Course Notes On Dynamic Optimization (Fall 2023)
Lecture 2A: Inventory Control

Instructor: Daniel Russo

Email: djr2174@gsb.columbia.edu

Graduate Instructor: David Cheikhi
Email: d.cheikhi@columbia.edu

September, 2023

These notes are based of scribed notes from a previous edition of the class. I have done some follow up light editing, but there may be typos or errors.

Consider the following problem. A firm starts with zero inventory of a single product ($x_0 = 0$) and decides upon ordering replenishment $u_k \geqslant 0$ after observing stochastic demand $w_k \geqslant 0$ in the current time period $k = 0, 1, 2, \ldots, N-1$. The evolution of the inventory follows the recursive relation

$$x_{k+1} = x_k + u_k - w_k, \qquad k = 0, 1, \ldots, N-1,$$

and unfulfilled orders are allowed to be backlogged ($x_k < 0$) until replenishment products fulfill them. The firm's objective is to minimize the expected overall cost

$$\mathbb{E}\left[\sum_{k=0}^{N-1} g_k(x_k, u_k, w_k) + g_N(x_N)\right],$$

where expectation is computed with regards to the independent and identically distributed (iid) random demands w_k's. Per-stage cost functions take the form

$$\begin{aligned} g_k(x_k, u_k, w_k) &= cu_k + r(x_k + u_k - w_k), \qquad k = 0, 1, \ldots, N-1 \\ g_N(x_N) &= 0, \end{aligned}$$

where $r(x) = px^- + hx^+ = p\max\{0, -x\} + h\max\{0, x\}$ consists of possible backlogging costs and holding costs.

Assumption 1. *We assume that $p > c$ so as to exclude the trivial decision to constantly backlog unfulfilled orders and not place any replenishment orders.*

The main result in this section establishes the optimality of a very special class of policies called "base stock policies." The firm maintain a base-stock level – a kind of of ideal inventory position – and orders however much inventory is needed to replenish depleted inventory upto the base-stock level.

Proposition 1. *(Base-stock policies are optimal) An optimal policy* $\pi^* = (\mu_0^*, \ldots, \mu_{N-1}^*)$ *exists where*

$$\mu_k^*(x_k) = (S_k - x_k)^+ = \begin{cases} S_k - x_k, & x_k \leqslant S_k \\ 0, & \text{otherwise}, \end{cases}$$

for some scalars $S_0, S_1, \ldots, S_{N-1}$.

Remark. In the finite horizon model, the base stock levels vary depending on the number of periods remaining in the selling horizon. It is possible to prove that when the selling horizon N is very large, $S_0 \approx S_1 \approx S_2$, etc. If inventory is replenished daily, then inventory planner's optimal decision does not depend (meaningfully) on whether there is a year remaining in the selling horizon or 100 years. This leads to a great conceptual simplification as the optimal policy is described by a single base-stock level.

1 Proof of the proposition

In order to prove the claim we work on target inventory positions instead: $y_k := x_k + u_k$ for each k. Restating Proposition 1 in terms of y's rather than u's, the goal is to show that inventory position $y_k = \max\{S_k, x_k\}$ is optimal.

Define the function

$$Q_k^*(x,y) = \mathbb{E}\left[c(y-x) + r(y-w_k) + J_{k+1}^*(y-w_k)\right] = \underbrace{\mathbb{E}\left[cy + r(y-w_k) + J_{k+1}^*(y-w_k)\right]}_{G_k(y)} - cx,$$

for all possible inventory positions x and all feasible $y \geqslant x$. Then, the DP algorithm yields for each k that

$$J_k^*(x) = \min_{y \geqslant x} Q_k(x,y) = \min_{y \geqslant x} G_k(y) - cx \tag{1}$$

The constraint that $y \geqslant x_k$ is due to the fact that we cannot order negative inventory. It is immediate that there is an optimal policy is of the form

$$\mu_k^*(x_k) = \left[\arg\min_{y \geqslant x_k} G_k(y)\right] - x_k,$$

The next lemma establishes properties about G_k that imply that a minimizer exists *(Convex coercive functions attain their infimum)*. The convexity of G_k has a more striking implication. Choose the base-stock level S_k to be *global* minimizer

$$S_k \in \arg\min_{y \in \mathbb{R}} G_k(y).$$

Then, by convexity,

$$\max\{S_k, x_k\} \in \arg\min_{y \geqslant x_k} G_k(y).$$

If the current inventory position is below the base-stock level S_k, it is optimal to move to inventory position S_k. If inventory is already above S_k, they should stay order nothing.

Definition 1. A function $G : \Re \to \Re$ is coercive if $G(x) \xrightarrow{x \to \pm\infty} +\infty$.

Lemma 1. *For each $k \in \{0, \dots, N-1\}$, J_k^* and G_k are convex functions. Moreover, G_k is coercive.*

Proof.

- G_k is convex.

  Notice that $J_N^*(\cdot) = 0$ is convex and $r(\cdot)$ is convex, so $G_{N-1}(\cdot)$ is convex because it is a weighted sum (expectation as integral) of the above along with a linear term cy. Note the fact that given any F convex, the function $g(x) = \min_{y \geqslant x} F(y)$ is also convex. (Draw picture[1]) Hence, J_{N-1}^* is convex. The proof is then concluded by backward induction.

- G_k is coercive.

  Since the cost J_{k+1}^* is always nonnegative, we know that

$$\begin{aligned} G_k(y) &\geqslant \mathbb{E}[cy + r(y - w_k)] \\ &= \mathbb{E}[c(y^+ - y^-) + p(y - w_k)^- + h(y - w_k)^+] \\ &\geqslant \mathbb{E}[-cy^- + py^- + h(y - w_k)^+] \qquad (y^+ \geqslant 0 \text{ and } w_k \geqslant 0) \\ &= (p - c) \underbrace{y^-}_{\text{coercive}} + h\mathbb{E}[\underbrace{(y - w_k)^+}_{\text{coercive}}] \end{aligned}$$

  is coercive because $p > c$ and $h > 0$.

□

1.1 Background: Operations Conserving Convexity

See also Boyd and Vandenberghe, *Convex Optimization*.

- Nonnegative weighted sums
  - If $f_1, \dots, f_m : \mathcal{D} \to \Re$ are convex and $w_!, \dots, w_m \geqslant 0$, then $w_1 f_1 + \cdots + w_m f_m$ is convex.
  - Given $f : \mathcal{X} \times \mathcal{Y} \to \Re$,

$$g(x) = \int f(x, y) w(y) \, \mathrm{d}y$$

    is convex if $w(y) \geqslant 0$ and the mapping $x \mapsto f(x, y)$ is convex for all $y \in \mathcal{Y}$.
- Composition with an affine map
  - $g(x) = f(Ax + b)$ is convex if f is convex.
- Pointwise supremum
  - $g(x) = \sup_{y \in \mathcal{Y}} f(x, y)$ is convex if the mapping $x \mapsto f(x, y)$ is convex for all $y \in \mathcal{Y}$.

2 Extension to fixed ordering costs

There are many possible extensions to the inventory control problem formulated above. One important one introduces a fixed ordering cost $K > 0$. In this case, the per-stage cost would be

$$g_k(x, u, w) = \begin{cases} r(x + u - w), & u = 0 \\ K + cu + r(x + u - w), & u > 0 \end{cases}.$$

[1] One formal way to prove this is a follows. Define $\tilde{F}(x, y) = F(y)$ if $y \geqslant x$ and $\tilde{F}(x, y) = \infty$ otherwise. Then $g(x) = \min_y \tilde{F}(x, y)$ is convex using the properties listed in Subsec 1.1.

- If G_k were convex in this case, a similar argument can show the optimal policy is the so-called **multiperiod** (s, S) **policy**:

$$\mu_k^*(x_k) = \begin{cases} 0, & x_k \geqslant s_k \\ S_k - x_k, & x_k < s_k \end{cases},$$

where $S_k = \arg\min G_k(y)$ and $s_k = \arg\min\{y \mid G_k(y) = K + G_k(S_k)\} < S_k$ for each k.

- Unfortunately, G_k may not be convex. However, it is "close enough" to convex and the multiperiod (s, S) policy remains optimal. Scarf first developed a notion of K**-convexity** and completed the proof by showing that G_k is K-convex, i.e., G_k satisfies the following property:

$$K + G_k(z + y) \geqslant G_k(y) + z\left(\frac{G_k(y) - G_k(y-b)}{b}\right), \quad \text{for all } z \geqslant 0, b > 0, y.$$

To get some intuition for this, imagine taking $b \to 0$. Then, this a relaxation of the defining fact of convex functions: a K-convex function falls below its tangent by no more than K.