Dynamic pricing and inventory control with large replenishment lead times

Xin Chen

University of Illinois at Urbana-Champaign

Joint work with Sasha Stolyar and Linwei Xin

IMA Workshop on Data-Driven Supply Chain Management
October 4, 2018

Funding support: NSF, JD.com, UIUC-ZJU Institute

Outline

1. Model description
2. Prior work
3. Main result
   - Asymptotic optimality of constant-order list-price policies
4. Proof sketch
   - Three steps in the proof
5. Conclusion

Model and notation

- Single item, periodic-review, backorder model, long-run average profit
- Unit ordering, holding and backorder costs: c, h and b
- Demand $D_t \triangleq \gamma_t D(p_t) + \beta_t$, $D(p_t)$ strictly decreasing
- $\{\gamma_t\}$ i.i.d. with mean one
- $\{\beta_t\}$ i.i.d. with zero mean
- $p_t \in [p_{min}, p_{max}]$, where $p_{min} < p_{max}$
- Deterministic lead time $L > 0$

Model and notation

- Single item, periodic-review, backorder model, long-run average profit
- Unit ordering, holding and backorder costs: c, h and b
- Demand $D_t \triangleq \gamma_t D(p_t) + \beta_t$, $D(p_t)$ strictly decreasing
- $\{\gamma_t\}$ i.i.d. with mean one
- $\{\beta_t\}$ i.i.d. with zero mean
- $p_t \in [p_{min}, p_{max}]$, where $p_{min} < p_{max}$
- Deterministic lead time $L > 0$

Model and notation

- Single item, periodic-review, backorder model, long-run average profit
- Unit ordering, holding and backorder costs: c, h and b
- Demand $D_t \triangleq \gamma_t D(p_t) + \beta_t$, $D(p_t)$ strictly decreasing
- $\{\gamma_t\}$ i.i.d. with mean one
- $\{\beta_t\}$ i.i.d. with zero mean
- $p_t \in [p_{min}, p_{max}]$, where $p_{min} < p_{max}$
- Deterministic lead time $L > 0$

Model and notation

- Single item, periodic-review, backorder model, long-run average profit
- Unit ordering, holding and backorder costs: c, h and b
- Demand $D_t \triangleq \gamma_t D(p_t) + \beta_t$, $D(p_t)$ strictly decreasing
- $\{\gamma_t\}$ i.i.d. with mean one
- $\{\beta_t\}$ i.i.d. with zero mean
- $p_t \in [p_{min}, p_{max}]$, where $p_{min} < p_{max}$
- Deterministic lead time $L > 0$

Model and notation

- Single item, periodic-review, backorder model, long-run average profit
- Unit ordering, holding and backorder costs: c, h and b
- Demand $D_t \triangleq \gamma_t D(p_t) + \beta_t$, $D(p_t)$ strictly decreasing
- $\{\gamma_t\}$ i.i.d. with mean one
- $\{\beta_t\}$ i.i.d. with zero mean
- $p_t \in [p_{min}, p_{max}]$, where $p_{min} < p_{max}$
- Deterministic lead time $L > 0$

Model and notation

- Single item, periodic-review, backorder model, long-run average profit
- Unit ordering, holding and backorder costs: c, h and b
- Demand $D_t \triangleq \gamma_t D(p_t) + \beta_t$, $D(p_t)$ strictly decreasing
- $\{\gamma_t\}$ i.i.d. with mean one
- $\{\beta_t\}$ i.i.d. with zero mean
- $p_t \in [p_{min}, \ p_{max}]$, where $p_{min} < p_{max}$
- Deterministic lead time $L > 0$

Model and notation

- Single item, periodic-review, backorder model, long-run average profit
- Unit ordering, holding and backorder costs: c, h and b
- Demand $D_t \triangleq \gamma_t D(p_t) + \beta_t$, $D(p_t)$ strictly decreasing
- $\{\gamma_t\}$ i.i.d. with mean one
- $\{\beta_t\}$ i.i.d. with zero mean
- $p_t \in [p_{min}, \ p_{max}]$, where $p_{min} < p_{max}$
- Deterministic lead time $L > 0$

Inventory dynamics (in period t)

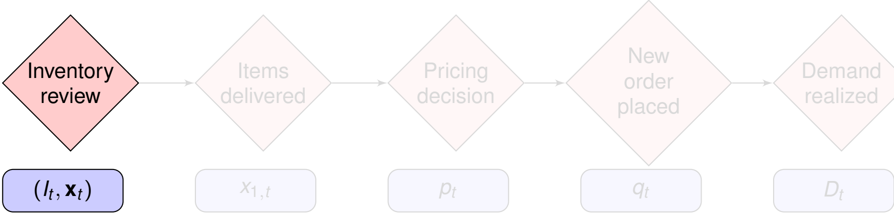

- On-hand inventory I_t
- Pipeline vector $\mathbf{x}_t = (x_{1,t}, x_{2,t}, \ldots, x_{L,t})$
  - Orders already placed but not yet received
- Decision variables p_t, q_t
  - p_t: pricing deicsion
  - q_t: new order placed
- Inventory update:
  - $I_{t+1} = I_t + x_{1,t} - D_t$, $\mathbf{x}_{t+1} = (x_{2,t}, \ldots, x_{L,t}, q_t)$

Inventory dynamics (in period t)

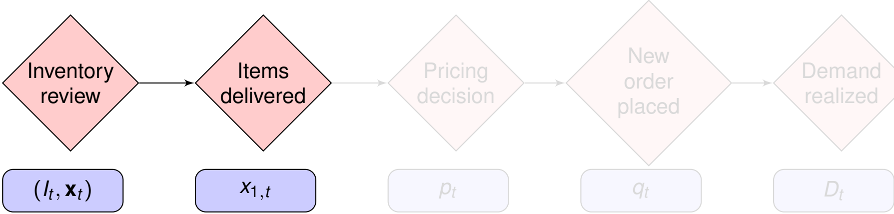

- On-hand inventory I_t
- Pipeline vector $\mathbf{x}_t = (x_{1,t}, x_{2,t}, \ldots, x_{L,t})$
  - Orders already placed but not yet received
- Decision variables p_t, q_t
- Inventory update:

Inventory dynamics (in period t)

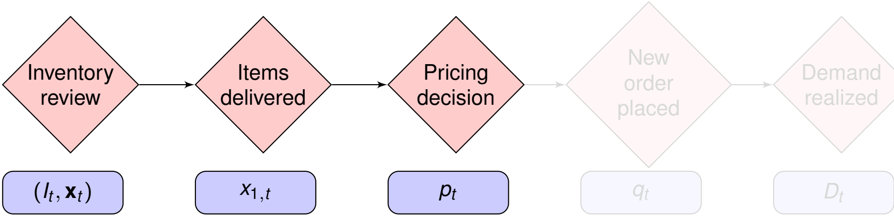

- On-hand inventory I_t
- Pipeline vector $\mathbf{x}_t = (x_{1,t}, x_{2,t}, \ldots, x_{L,t})$
  - Orders already placed but not yet received
- Decision variables p_t, q_t
  - p_t: pricing deicsion
  - q_t: new order placed
- Inventory update:

Inventory dynamics (in period t)

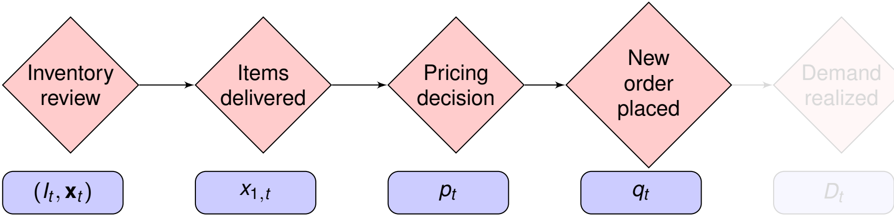

- On-hand inventory I_t
- Pipeline vector $\mathbf{x}_t = (x_{1,t}, x_{2,t}, \ldots, x_{L,t})$
  - Orders already placed but not yet received
- Decision variables p_t, q_t
  - p_t: pricing deicsion
  - q_t: new order placed
- Inventory update:

Inventory dynamics (in period t)

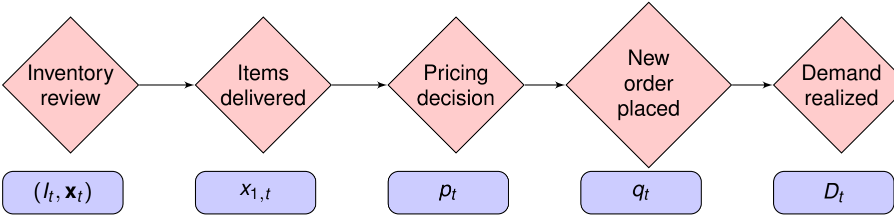

- On-hand inventory I_t
- Pipeline vector $\mathbf{x}_t = (x_{1,t}, x_{2,t}, \ldots, x_{L,t})$
  - Orders already placed but not yet received
- Decision variables p_t, q_t
  - p_t: pricing deicsion
  - q_t: new order placed
- Inventory update:

Inventory dynamics (in period t)

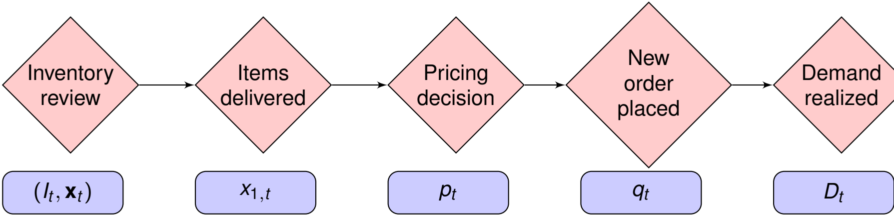

- On-hand inventory I_t
- Pipeline vector $\mathbf{x}_t = (x_{1,t}, x_{2,t}, \dots, x_{L,t})$
  - Orders already placed but not yet received
- Decision variables p_t, q_t
  - p_t: pricing deicsion
  - q_t: new order placed
- Inventory update:
  - $I_{t+1} = I_t + x_{1,t} - D_t, \quad \mathbf{x}_{t+1} = (x_{2,t}, \dots, x_{L,t}, q_t)$

Inventory dynamics (in period t)

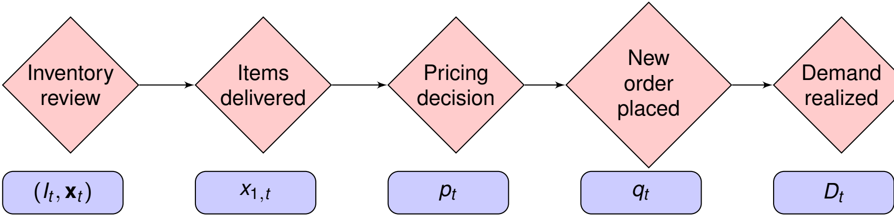

- On-hand inventory I_t
- Pipeline vector $\mathbf{x}_t = (x_{1,t}, x_{2,t}, \ldots, x_{L,t})$
  - Orders already placed but not yet received
- Decision variables p_t, q_t
  - p_t: pricing deicsion
  - q_t: new order placed
- Inventory update:
  - $I_{t+1} = I_t + x_{1,t} - D_t, \quad \mathbf{x}_{t+1} = (x_{2,t}, \ldots, x_{L,t}, q_t)$

Inventory dynamics (in period t)

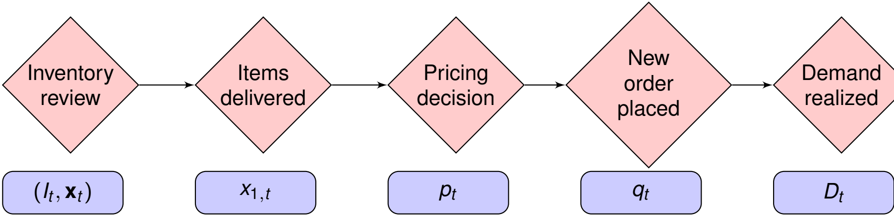

- On-hand inventory I_t
- Pipeline vector $\mathbf{x}_t = (x_{1,t}, x_{2,t}, \ldots, x_{L,t})$
  - Orders already placed but not yet received
- Decision variables p_t, q_t
  - p_t: pricing deicsion
  - q_t: new order placed
- Inventory update:
  - $I_{t+1} = I_t + x_{1,t} - D_t, \quad \mathbf{x}_{t+1} = (x_{2,t}, \ldots, x_{L,t}, q_t)$

Performance measure and optimal policy

- $G(x) \triangleq hx^+ + bx^-$
- Profit in period t:

$$C_t \triangleq p_t D_t - [c x_{1,t} + G(I_t + x_{1,t} - D_t)]$$

- Long run average profit of policy π:

$$C(\pi) \triangleq \liminf_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} E[C_t^{\pi}]$$

- Optimal long run average profit:

$$\mathrm{OPT}(L) \triangleq \sup_{\pi} C(\pi)$$

Performance measure and optimal policy

- $G(x) \triangleq hx^+ + bx^-$
- Profit in period t:

$$C_t \triangleq p_t D_t - [cx_{1,t} + G(I_t + x_{1,t} - D_t)]$$

- Long run average profit of policy π:

$$C(\pi) \triangleq \liminf_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} E[C_t^\pi]$$

- Optimal long run average profit:

$$\mathrm{OPT}(L) \triangleq \sup_{\pi} C(\pi)$$

Performance measure and optimal policy

- $G(x) \triangleq \; hx^+ + bx^-$

- Profit in period t:

$$C_t \triangleq \; p_t D_t - \left[c x_{1,t} + G\left(I_t + x_{1,t} - D_t \right) \right]$$

- Long run average profit of policy π:

$$C(\pi) \triangleq \; \liminf_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} E\left[C_t^{\pi} \right]$$

- Optimal long run average profit:

$$\mathrm{OPT}(L) \triangleq \; \sup_{\pi} \; C(\pi)$$

Performance measure and optimal policy

- $G(x) \triangleq hx^+ + bx^-$
- Profit in period t:

$$C_t \triangleq p_t D_t - \left[c x_{1,t} + G\left(I_t + x_{1,t} - D_t \right) \right]$$

- Long run average profit of policy π:

$$C(\pi) \triangleq \liminf_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} E\left[C_t^{\pi} \right]$$

- Optimal long run average profit:

$$\mathrm{OPT}(L) \triangleq \sup_{\pi} C(\pi)$$

Literature review

- First studied in Whitin (1955)
- Optimality of a base-stock list price policy in the zero lead-time setting (Federgruen and Heching (1999))
- Extension to the setting with setup costs (e.g. Chen and Simchi-Levi (2004a, 2004b), Yao et al. (2007), Huh and Janakarimian (2008)
- Setting with lead-times
  - base-stock list price policy is no longer optimal in general
  - Curse of dimensionality
- *"...it remains a significant challenge to incorporate lead time into stochastic models. Indeed, the zero lead time assumption is required for all the multi-period models reviewed here..."* (Chen and Simchi-Levi 2012)

Literature review

- First studied in Whitin (1955)
- Optimality of a base-stock list price policy in the zero lead-time setting (Federgruen and Heching (1999))
- Extension to the setting with setup costs (e.g. Chen and Simchi-Levi (2004a, 2004b), Yao et al. (2007), Huh and Janakarimian (2008)
- Setting with lead-times
  - base-stock list price policy is no longer optimal in general
  - Curse of dimensionality
- *"... it remains a significant challenge to incorporate lead time into stochastic models. Indeed, the zero lead time assumption is required for all the multi-period models reviewed here..."* (Chen and Simchi-Levi 2012)

Literature review cont.

- Selected literature: Thomas (1974), Petruzzi and Dada (1999), Agrawal and Seshadri (2000), Elmaghraby and Keskinocak (2003), Chen, Xu and Zhang (2009), Li, Lim and Rodrigues (2009), Chen, Pang and Pan (2014), Chen, Chao and Ahn (2015), Chen, Chao and Shi (2016)
- Bernstein, Li and Shang (2015): positive lead time, focusing on designing effective heuristics

Literature review cont.

- Selected literature: Thomas (1974), Petruzzi and Dada (1999), Agrawal and Seshadri (2000), Elmaghraby and Keskinocak (2003), Chen, Xu and Zhang (2009), Li, Lim and Rodrigues (2009), Chen, Pang and Pan (2014), Chen, Chao and Ahn (2015), Chen, Chao and Shi (2016)
- Bernstein, Li and Shang (2015): positive lead time, focusing on designing effective heuristics

Constant-order policies

- First studied in a lost-sales inventory model [Reiman (2004)]
- Always order the same amount of inventory regardless of what on hands and in-transit
- Example:
  - Constant-order quantity: **100**
  - If oh-hand=**0**, order **100**
  - If oh-hand=**1000**, order **100**

Performance in a lost-sales model

- Can beat base-stock as the lead time grows [Reiman (2004)]
- Surprising computational results of [Zipkin (2008)]
  - Compare to several heuristics
  - Constant-order policy did surprisingly well even when $L = 4$

Asymptotic optimality

- Lost-sales model
  - constant-order is **asymptotically optimal** as the lead time grows [Goldberg, Katz-Rogozhnikov, Lu, Sharma, Squillante (2016)]
  - **exponential** convergence [Xin and Goldberg (2016)]
- Dual-sourcing model
  - Tailored Base-Surge policy (constant-order + base-stock) is **asymptotically optimal** as the lead time difference grows [Xin and Goldberg (2017)]

Asymptotic optimality

- Lost-sales model
  - constant-order is **asymptotically optimal** as the lead time grows [Goldberg, Katz-Rogozhnikov, Lu, Sharma, Squillante (2016)]
  - **exponential** convergence [Xin and Goldberg (2016)]
- Dual-sourcing model
  - Tailored Base-Surge policy (constant-order + base-stock) is **asymptotically optimal** as the lead time difference grows [Xin and Goldberg (2017)]

Assumptions

Assumption

- The inverse function D^{-1} of D is continuous and strictly decreasing.
- The revenue $dD^{-1}(d)$ is a concave function of the expected demand d.
- $dD^{-1}(d)$ is Lipschitz continuous with a constant $\kappa > 0$.

Constant-order list-price policy

- $d_{min} \triangleq D(p_{max})$, $d_{max} \triangleq D(p_{min})$
- Compute the best constant-order policy:

$$\max_{x\in[d_{min},d_{max}]} \underbrace{\max_{\pi_p} C(\pi_{x,\pi_p})}_{\text{concave in } x}$$

- The best constant $x^* \in (d_{min}, d_{max})$

Theorem

$$\lim_{L\to\infty} \mathrm{OPT}(L) = \max_{x\in[d_{min},d_{max}]} \underbrace{\max_{\pi_p} C(\pi_{x,\pi_p})}_{\text{concave in } x}$$

Constant-order list-price policy

- $d_{min} \triangleq D(p_{max})$, $d_{max} \triangleq D(p_{min})$
- Compute the best constant-order policy:

$$\max_{x\in[d_{min},d_{max}]} \underbrace{\max_{\pi_p} C(\pi_{x,\pi_p})}_{\text{concave in } x}$$

- The best constant $x^* \in (d_{min}, d_{max})$

Theorem

$$\lim_{L\to\infty} \mathrm{OPT}(L) = \max_{x\in[d_{min},d_{max}]} \underbrace{\max_{\pi_p} C(\pi_{x,\pi_p})}_{\text{concave in } x}$$

Constant-order list-price policy

- $d_{min} \triangleq D(p_{max})$, $d_{max} \triangleq D(p_{min})$

- Compute the best constant-order policy:

$$\max_{x \in [d_{min}, d_{max}]} \underbrace{\max_{\pi_p} C(\pi_{x,\pi_p})}_{\text{concave in } x}$$

- The best constant $x^* \in (d_{min}, d_{max})$

Theorem

$$\lim_{L \to \infty} \mathrm{OPT}(L) = \max_{x \in [d_{min}, d_{max}]} \underbrace{\max_{\pi_p} C(\pi_{x,\pi_p})}_{\text{concave in } x}$$

Constant-order list-price policy

- $d_{min} \triangleq D(p_{max})$, $d_{max} \triangleq D(p_{min})$
- Compute the best constant-order policy:

$$\max_{x \in [d_{min}, d_{max}]} \underbrace{\max_{\pi_p} C(\pi_{x,\pi_p})}_{\text{concave in } x}$$

- The best constant $x^* \in (d_{min}, d_{max})$

Theorem

$$\lim_{L \to \infty} \mathrm{OPT}(L) = \max_{x \in [d_{min}, d_{max}]} \underbrace{\max_{\pi_p} C(\pi_{x,\pi_p})}_{\text{concave in } x}$$

Proof overview

- Step I: existence of a steady-state
  - **perturbative approaches**
- Step II: an upper bound of the optimal value
  - concavity argument
- Step III: match constant-order to the upper bound
  - vanishing discount approach

Proof overview

- Step I: existence of a steady-state
  - **perturbative approaches**
- Step II: an upper bound of the optimal value
  - **concavity argument**
- Step III: match constant-order to the upper bound
  - **vanishing discount approach**

Proof overview

- Step I: existence of a steady-state
  - **perturbative approaches**
- Step II: an upper bound of the optimal value
  - **concavity argument**
- Step III: match constant-order to the upper bound
  - **vanishing discount approach**

Proof overview

- Step I: existence of a steady-state
  - **perturbative approaches**
- Step II: an upper bound of the optimal value
  - **concavity argument**
- Step III: match constant-order to the upper bound
  - **vanishing discount approach**

Proof overview

- Step I: existence of a steady-state
  - **perturbative approaches**
- Step II: an upper bound of the optimal value
  - **concavity argument**
- Step III: match constant-order to the upper bound
  - **vanishing discount approach**

Proof overview

- Step I: existence of a steady-state
  - **perturbative approaches**
- Step II: an upper bound of the optimal value
  - **concavity argument**
- Step III: match constant-order to the upper bound
  - **vanishing discount approach**

Step I: existence of a steady-state

Lemma

Without loss of generality, there exists a stationary measure $\left(\mathcal{I}^{L,*}, \chi_1^{L,*}, \ldots, \chi_L^{L,*}\right)$ of the Markov chain under an optimal stationary policy, and it satisfies

$$\mathrm{OPT}(L) = \mathbb{E}\left[d_1^{L,*} D^{-1}\left(d_1^{L,*}\right)\right] - c\mathbb{E}[\chi_1^{L,*}] - \mathbb{E}\left[G\left(\mathcal{I}^{L,*}\right)\right].$$

Step II: upper bound of $OPT(L)$

Use concavity to obtain an upper bound:

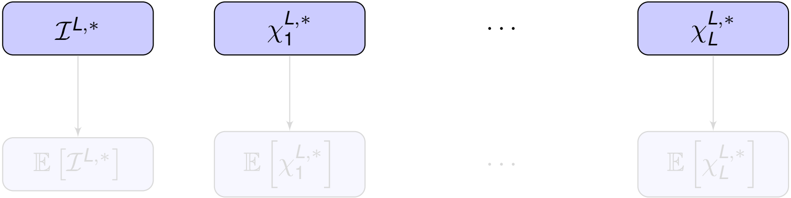

constant-order!

Step II: upper bound of $OPT(L)$

Use concavity to obtain an upper bound:

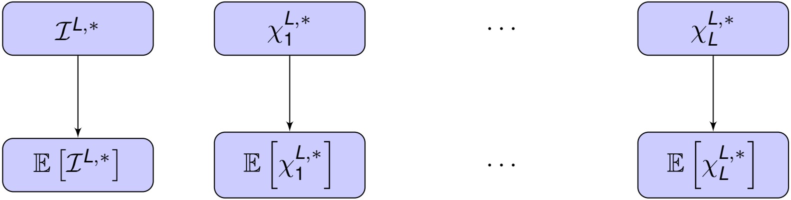

constant-order!

Step II: upper bound of $OPT(L)$

Use concavity to obtain an upper bound:

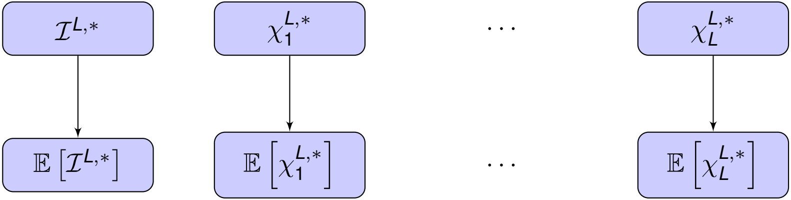

constant-order!

OPT(L) is at most

$$
\begin{aligned}
&\mathbb{E}\left[d_1^{L,*} D^{-1}\left(d_1^{L,*}\right)\right] - c\mathbb{E}[\chi_1^{L,*}] - \mathbb{E}\left[G\left(\mathcal{I}^{L,*}\right)\right] \\
&= \frac{1-\alpha}{1-\alpha^L} \sum_{k=1}^{L} \alpha^{k-1} \mathbb{E}\Bigg[d_k^{L,*} D^{-1}(d_k^{L,*}) - c\chi_1^{L,*} \\
&\qquad - G\left(\mathcal{I}^{L,*} + \sum_{t=1}^{k}\left(\chi_t^{L,*} - \gamma_t d_t^{L,*} - \beta_t\right)\right)\Bigg].
\end{aligned}
$$

Applying Jensen's inequalty, OPT(L) is at most

$$\frac{1-\alpha}{1-\alpha^L}\sum_{k=1}^{L}\alpha^{k-1}\mathbb{E}\Bigg[\mathbb{E}[d_k^{L,*}|\epsilon_{[k-1]}]D^{-1}\left(\mathbb{E}[d_k^{L,*}|\epsilon_{[k-1]}]\right)-c\mathbb{E}\left[\chi_1^{L,*}\right]$$
$$-G\left(\mathbb{E}[\mathcal{I}^{L,*}]+\sum_{t=1}^{k}\left(\mathbb{E}[\chi_t^{L,*}]-\gamma_t\mathbb{E}[d_t^{L,*}|\epsilon_{[t-1]}]-\beta_t\right)\right)\Bigg],$$

Thus,

$$\mathrm{OPT}(L)\leq \frac{1-\alpha}{1-\alpha^L}\max_{S\in[-\overline{S},\overline{S}]}V_L^{\alpha}(x_L,S)\quad\text{for each }\alpha\in(0,1),$$

and

$$\liminf_{L\to\infty}\mathrm{OPT}(L)\leq (1-\alpha)\limsup_{L\to\infty}\max_{S\in[-\overline{S},\overline{S}]}V_L^{\alpha}(x_L,S)$$
$$\leq (1-\alpha)\max_{S\in[-\overline{S},\overline{S}]}V_{\infty}^{\alpha}(x_{\infty},S).$$

Step III: match constant-order to the upper bound

Upper bound

Total discounted profit over an infinite horizon with initial on-hand inventory S and constant-order x_∞.

Constant-order policy

Long-run average profit under the best constant-order policy

convergence of a discounted problem to its long-run average counterpart

Step III: match constant-order to the upper bound

Upper bound

Total discounted profit over an infinite horizon with initial on-hand inventory S and constant-order x_∞.

Constant-order policy

Long-run average profit under the best constant-order policy

convergence of a discounted problem to its long-run average counterpart

Step III: match constant-order to the upper bound

Upper bound

Total discounted profit over an infinite horizon with initial on-hand inventory S and constant-order x_∞.

Constant-order policy

Long-run average profit under the best constant-order policy

convergence of a discounted problem to its long-run average counterpart

Schäl's conditions

Consider the following MDP with

- state space $\mathcal{S}$,
- action spaces $\mathcal{A}(s)$ for each $s \in \mathcal{S}$,
- probability transition function $q(.|s, a)$,
- deterministic and nonnegative single-period cost function $c(s, a)$.

Given a feasible policy π, a discount factor $\alpha \in (0, 1)$, and an initial state s, the expected long-run average cost and total discounted cost are denoted as $J^{\pi}(s)$ and $J^{\pi}_{\alpha}(s)$ respectively.

Schäl's conditions

Consider the following MDP with

- state space $\mathcal{S}$,
- action spaces $\mathcal{A}(s)$ for each $s \in \mathcal{S}$,
- probability transition function $q(.|s, a)$,
- deterministic and nonnegative single-period cost function $c(s, a)$.

Given a feasible policy π, a discount factor $\alpha \in (0, 1)$, and an initial state s, the expected long-run average cost and total discounted cost are denoted as $J^{\pi}(s)$ and $J_{\alpha}^{\pi}(s)$ respectively.

Schäl's conditions

1. $\mathcal{S}$ is a locally compact space with a countable base, i.e., there exists a countable collection $\mathcal{B}$ of open sets in a locally compact space $\mathcal{S}$ such that any open set containing $x \in \mathcal{S}$ contains at least one of the open sets in $\mathcal{B}$.
2. For each $s \in \mathcal{S}$, $\mathcal{A}(s)$ is nonempty and compact. Furthermore, $\mathcal{A}(.)$ is upper semicontinuous, i.e., for every open set $B \subseteq \mathbb{R}$, the set $\{s : \mathcal{A}(s) \subseteq B\}$ is open in $\mathcal{S}$.
3. The probability transition function $q : \{(s, a) : a \in \mathcal{A}(s)\} \to \mathbb{P}(\mathcal{S})$ is continuous with respect to weak convergence on $\mathbb{P}(\mathcal{S})$, where $\mathbb{P}(\mathcal{S})$ denotes the set of all probability measures on $\mathcal{S}$.
4. The single-period cost function c is lower semicontinuous, i.e., $\{(s, a) : c(s, a) > \gamma\}$ is an open set for all $\gamma \in \mathbb{R}$.
5. There exists a policy π and an initial state $s \in \mathcal{S}$ such that $J^\pi(s) < \infty$.
6. $\sup_{\alpha<1} \left(\inf_\pi J^\pi_\alpha(s) - \inf_{s' \in \mathcal{S}} \inf_\pi J^\pi_\alpha(s') \right) < \infty$ for all $s \in \mathcal{S}$.

Schäl's conditions

1. $\mathcal{S}$ is a locally compact space with a countable base, i.e., there exists a countable collection $\mathcal{B}$ of open sets in a locally compact space $\mathcal{S}$ such that any open set containing $x \in \mathcal{S}$ contains at least one of the open sets in $\mathcal{B}$.
2. For each $s \in \mathcal{S}$, $\mathcal{A}(s)$ is nonempty and compact. Furthermore, $\mathcal{A}(.)$ is upper semicontinuous, i.e., for every open set $B \subseteq \mathbb{R}$, the set $\{s : \ \mathcal{A}(s) \subseteq B\}$ is open in $\mathcal{S}$.
3. The probability transition function $q : \{(s, a) : \ a \in \mathcal{A}(s)\} \to \mathbb{P}(\mathcal{S})$ is continuous with respect to weak convergence on $\mathbb{P}(\mathcal{S})$, where $\mathbb{P}(\mathcal{S})$ denotes the set of all probability measures on $\mathcal{S}$.
4. The single-period cost function c is lower semicontinuous, i.e., $\{(s, a) : \ c(s, a) > \gamma\}$ is an open set for all $\gamma \in \mathbb{R}$.
5. There exists a policy π and an initial state $s \in \mathcal{S}$ such that $J^\pi(s) < \infty$.
6. $\sup_{\alpha<1} \left(\inf_\pi J^\pi_\alpha(s) - \inf_{s' \in \mathcal{S}} \inf_\pi J^\pi_\alpha(s') \right) < \infty$ for all $s \in \mathcal{S}$.

Vanishing discount approach

Theorem (Schäl 1993)

Under the above conditions, there exists an optimal stationary policy π^* such that for all $s \in \mathcal{S}$,

$$J^{\pi^*}(s) = \inf_{s' \in \mathcal{S}} \inf_{\pi} J^{\pi}(s') = \lim_{\alpha \uparrow 1} \left[(1-\alpha) \inf_{s' \in \mathcal{S}} \inf_{\pi} J^{\pi}_{\alpha}(s') \right].$$

In our setting, we need to prove

$$\liminf_{\alpha \uparrow 1} \left[(1-\alpha) \max_{S \in [-\overline{S}, \overline{S}]} V^{\alpha}_{\infty}(x_\infty, S) \right] = C(\pi_{x_\infty}).$$

Verifying Conditions

It suffices to verify condition

$$\sup_{\alpha \in (0,1)} \left[\max_{S' \in \mathbb{R}} V_\infty^\alpha(x_\infty, S') - V_\infty^\alpha(x_\infty, S) \right] < \infty.$$

Assume S_α^* solves $\max_{S' \in \mathbb{R}} V_\infty^\alpha(x_\infty, S')$ with an optimal policy π^*. In the inventory system starting with S_α^* following policy π^*,

$$I_n^* = I_{n-1}^* + x_\infty - \gamma d_n^* - \beta.$$

For the system starting with S, we want to construct a policy π to pursue I^* so that the profit difference is bounded (independent of α),

$$I_n = I_{n-1} + x_\infty - \gamma d_n - \beta.$$

Mapping to Random Yield Model with Capacity Constraint

- By treating supply as "demand" and demand as "supply", the pricing and inventory control problem becomes a random yield model with capacity constraint on orders

$$\tilde{I}_n = \tilde{I}_{n-1} + \gamma d_n - (x_\infty - \beta).$$

- Federgruen and Yang (2014) address infinite horizon random yield model without capacity constraint using the vanishing discount approach
  - $|I_n - I_n^*|$ decreases geometrically
  - the idea cannot be extended to the case with capacity constraint
- We show for the infinite horizon random yield model with capacity constraint
  - order to capacity when inventory is too low (uniformly on α)
  - carefully bound the cost difference

Mapping to Random Yield Model with Capacity Constraint

- By treating supply as “demand” and demand as “supply”, the pricing and inventory control problem becomes a random yield model with capacity constraint on orders

$$\tilde{I}_n = \tilde{I}_{n-1} + \gamma d_n - (x_\infty - \beta).$$

- Federgruen and Yang (2014) address infinite horizon random yield model without capacity constraint using the vanishing discount approach
  - $|I_n - I_n^*|$ decreases geometrically
  - the idea cannot be extended to the case with capacity constraint
- We show for the infinite horizon random yield model with capacity constraint
  - order to capacity when inventory is too low (uniformly on α)
  - carefully bound the cost difference

Mapping to Random Yield Model with Capacity Constraint

- By treating supply as “demand” and demand as “supply”, the pricing and inventory control problem becomes a random yield model with capacity constraint on orders

$$\tilde{I}_n = \tilde{I}_{n-1} + \gamma d_n - (x_\infty - \beta).$$

- Federgruen and Yang (2014) address infinite horizon random yield model without capacity constraint using the vanishing discount approach
  - $|I_n - I_n^*|$ decreases geometrically
  - the idea cannot be extended to the case with capacity constraint
- We show for the infinite horizon random yield model with capacity constraint
  - order to capacity when inventory is too low (uniformly on α)
  - carefully bound the cost difference

Mapping to Random Yield Model with Capacity Constraint

- By treating supply as “demand” and demand as “supply”, the pricing and inventory control problem becomes a random yield model with capacity constraint on orders

$$\tilde{I}_n = \tilde{I}_{n-1} + \gamma d_n - (x_\infty - \beta).$$

- Federgruen and Yang (2014) address infinite horizon random yield model without capacity constraint using the vanishing discount approach
  - $|I_n - I_n^*|$ decreases geometrically
  - the idea cannot be extended to the case with capacity constraint
- We show for the infinite horizon random yield model with capacity constraint
  - order to capacity when inventory is too low (uniformly on α)
  - carefully bound the cost difference

Mapping to Random Yield Model with Capacity Constraint

- By treating supply as “demand” and demand as “supply”, the pricing and inventory control problem becomes a random yield model with capacity constraint on orders

$$\tilde{I}_n = \tilde{I}_{n-1} + \gamma d_n - (x_\infty - \beta).$$

- Federgruen and Yang (2014) address infinite horizon random yield model without capacity constraint using the vanishing discount approach
  - $|I_n - I_n^*|$ decreases geometrically
  - the idea cannot be extended to the case with capacity constraint
- We show for the infinite horizon random yield model with capacity constraint
  - order to capacity when inventory is too low (uniformly on α)
  - carefully bound the cost difference

Mapping to Random Yield Model with Capacity Constraint

- By treating supply as "demand" and demand as "supply", the pricing and inventory control problem becomes a random yield model with capacity constraint on orders

$$\tilde{I}_n = \tilde{I}_{n-1} + \gamma d_n - (x_\infty - \beta).$$

- Federgruen and Yang (2014) address infinite horizon random yield model without capacity constraint using the vanishing discount approach
  - $|I_n - I_n^*|$ decreases geometrically
  - the idea cannot be extended to the case with capacity constraint
- We show for the infinite horizon random yield model with capacity constraint
  - order to capacity when inventory is too low (uniformly on α)
  - carefully bound the cost difference

Mapping to Random Yield Model with Capacity Constraint

- By treating supply as "demand" and demand as "supply", the pricing and inventory control problem becomes a random yield model with capacity constraint on orders

$$\tilde{I}_n = \tilde{I}_{n-1} + \gamma d_n - (x_\infty - \beta).$$

- Federgruen and Yang (2014) address infinite horizon random yield model without capacity constraint using the vanishing discount approach
    - $|I_n - I_n^*|$ decreases geometrically
    - the idea cannot be extended to the case with capacity constraint
- We show for the infinite horizon random yield model with capacity constraint
    - order to capacity when inventory is too low (uniformly on α)
    - carefully bound the cost difference

Conclusion

Main contribution

- Establish asymptotic optimality of constant-order policies for joint pricing and inventory control with large replenishment lead times

Conclusion

Main contribution

- Establish asymptotic optimality of constant-order policies for joint pricing and inventory control with large replenishment lead times

Future research directions

- Other inventory models
  - not universally held
  - counter-example: single-sourcing backlogged model
- Open problems
  - fixed ordering costs
  - general MDPs
  - ...

Future research directions

- Other inventory models
  - not universally held
  - counter-example: single-sourcing backlogged model
- Open problems
  - fixed ordering costs
  - general MDPs
  - ...

Future research directions

- Other inventory models
  - not universally held
  - counter-example: single-sourcing backlogged model
- Open problems
  - fixed ordering costs
  - general MDPs
  - ...

Future research directions

- Other inventory models
  - not universally held
  - counter-example: single-sourcing backlogged model
- Open problems
  - fixed ordering costs
  - general MDPs
  - ...

Extension and Implications

- Finite state, finite action MDPs
  - First type decisions: takes effect right away
  - Second type decisions: takes effect after a long lead time
- Claim: Open-loop control is asymptotically optimal for the second type decisions
- Implications for data-driven models
  - Ignore real time information when making the second type decisions
  - other examples?

Extension and Implications

- Finite state, finite action MDPs
  - First type decisions: takes effect right away
  - Second type decisions: takes effect after a long lead time
- Claim: Open-loop control is asymptotically optimal for the second type decisions
- Implications for data-driven models
  - Ignore real time information when making the second type decisions
  - other examples?

Extension and Implications

- Finite state, finite action MDPs
  - First type decisions: takes effect right away
  - Second type decisions: takes effect after a long lead time
- Claim: Open-loop control is asymptotically optimal for the second type decisions
- Implications for data-driven models
  - Ignore real time information when making the second type decisions
  - other examples?

Extension and Implications

- Finite state, finite action MDPs
  - First type decisions: takes effect right away
  - Second type decisions: takes effect after a long lead time
- Claim: Open-loop control is asymptotically optimal for the second type decisions
- Implications for data-driven models
  - Ignore real time information when making the second type decisions
  - other examples?

Extension and Implications

- Finite state, finite action MDPs
  - First type decisions: takes effect right away
  - Second type decisions: takes effect after a long lead time
- Claim: Open-loop control is asymptotically optimal for the second type decisions
- Implications for data-driven models
  - Ignore real time information when making the second type decisions
  - other examples?