A Weak Convergence Approach to Inventory Control Using a Long-term Average Criterion

Chao Zhu

(Joint with K.L Helmes and R.H. Stockbridge)

University of Wisconsin - Milwaukee

July, 2018

The 14th Workshop on Markov Processes and Related Topics

Outline

Formulation: The Inventory Process

- **The inventory process** (in the absence of orders):

  $$\mathrm{d}X_0(t) = \mu(X_0(t))\mathrm{d}t + \sigma(X_0(t))\mathrm{d}W(t), \; X(0) = x_0 \in \mathcal{I} = (a, b),$$

  where $-\infty \leq a < b \leq +\infty$.
- In the absence of restocking
  - demands tend to reduce the inventory

Formulation: The Inventory Process

- **The inventory process** (in the absence of orders):

  $$\mathrm{d}X_0(t) = \mu(X_0(t))\mathrm{d}t + \sigma(X_0(t))\mathrm{d}W(t),\ X(0) = x_0 \in \mathcal{I} = (a, b),$$

  where $-\infty \le a < b \le +\infty$.
- In the absence of restocking
  - demands tend to reduce the inventory $\rightsquigarrow$ a is an *attracting* point:

    $$\mathbb{P}_x\{\tau_{a+} \le \tau_r\} > 0, \forall a < x < r < b.$$

    a may be a regular, exit or natural boundary point.

Formulation: The Inventory Process

- **The inventory process** (in the absence of orders):

  $$\mathrm{d}X_0(t) = \mu(X_0(t))\mathrm{d}t + \sigma(X_0(t))\mathrm{d}W(t),\ X(0) = x_0 \in \mathcal{I} = (a, b),$$

  where $-\infty \leq a < b \leq +\infty$.
- In the absence of restocking
  - demands tend to reduce the inventory $\rightsquigarrow$ a is an *attracting* point:
    $$\mathbb{P}_x\{\tau_{a+} \leq \tau_r\} > 0, \forall a < x < r < b.$$
    a may be a regular, exit or natural boundary point.
  - reasonable restrictions on "returns": the inventory level can never reach b in finite time

Formulation: The Inventory Process

- **The inventory process** (in the absence of orders):

  $$\mathrm{d}X_0(t) = \mu(X_0(t))\mathrm{d}t + \sigma(X_0(t))\mathrm{d}W(t),\ X(0) = x_0 \in \mathcal{I} = (a, b),$$

  where $-\infty \leq a < b \leq +\infty$.
- In the absence of restocking
  - demands tend to reduce the inventory $\rightsquigarrow$ a is an *attracting* point:

    $$\mathbb{P}_x\{\tau_{a+} \leq \tau_r\} > 0, \forall a < x < r < b.$$

    a may be a regular, exit or natural boundary point.
  - reasonable restrictions on "returns": the inventory level can never reach b in finite time $\rightsquigarrow$ b is a *non-attracting* point:

    $$\mathbb{P}_x\{\tau_{b-} \leq \tau_r\} = 0, \forall a < r < x < b.$$

    b may be an entrance or natural point.

Inventory Management

The manager can decide *WHEN* and *HOW MUCH* to restock the inventory:

$$
\begin{array}{lccccccccc}
\text{WHEN:} & 0 & \leq & \tau_1 & \leq & \tau_2 & \leq & \tau_3 & \leq & \dots \\
 & & & \downarrow & & \downarrow & & \downarrow & & \\
\text{HOW MUCH:} & & & Y_1 & & Y_2 & & Y_3 & & \dots
\end{array}
$$

Inventory Management

The manager can decide *WHEN* and *HOW MUCH* to restock the inventory:

$$
\begin{array}{lccccccccc}
\text{WHEN:} & 0 & \le & \tau_1 & \le & \tau_2 & \le & \tau_3 & \le & \dots \\
 & & & \downarrow & & \downarrow & & \downarrow & & \\
\text{HOW MUCH:} & & & Y_1 & & Y_2 & & Y_3 & & \dots
\end{array}
$$

For each $k = 1, 2, \dots,$

- τ_k, the kth order time, is an $\{\mathcal{F}_t\}$-stopping time, and
- Y_k, the kth order size, is an $\{\mathcal{F}_{\tau_k}\}$-measurable nonnegative random variable.
- $X(\tau_k-)$: the inventory level just before the kth order,
- $X(\tau_k)$: the inventory level at the kth order,

$$X(\tau_k) = X(\tau_k-) + Y_k \ge X(\tau_k-)$$

Admissible Policies

- $\mathcal{A} = \{(\tau, Y) = (\tau_k, Y_k), k = 1, 2, \dots\}$, in which
  - $\{\tau_k\}$ is an increasing sequence of $\{\mathcal{F}_t\}$-stopping times,
  - Y_k is an $\{\mathcal{F}_{\tau_k}\}$-measurable nonnegative random variable,
  - $X(\tau_k) = X(\tau_k-) + Y_k \in \mathcal{E}$ (the state space).
- For models in which a is a reflecting boundary point, the class $\mathcal{A}_0 \subset \mathcal{A}$ consists of those policies (τ, Y) for which

$$\lim_{t\to\infty} t^{-1}\mathbb{E}[L_a(t)] = 0$$

Inventory Management

- The controlled inventory

$$X(t) = X(0-) + \int_0^t \mu(X(s))\mathrm{d}s + \int_0^t \sigma(X(s))\mathrm{d}W(s) + \sum_{k=1}^{\infty} I_{\{\tau_k \le t\}} Y_k,$$

$$X(t) \in \mathcal{E},$$

- The state space $\mathcal{E}$:

$$\mathcal{E} = \begin{cases} (a, b), & \text{if } a \text{ and } b \text{ are natural boundaries,} \\ [a, b), & \text{if } a \text{ is attainable and } b \text{ is natural,} \\ (a, b], & \text{if } a \text{ is natural and } b \text{ is entrance,} \\ [a, b], & \text{if } a \text{ is attainable and } b \text{ is entrance.} \end{cases}$$

The (s, S) Policy

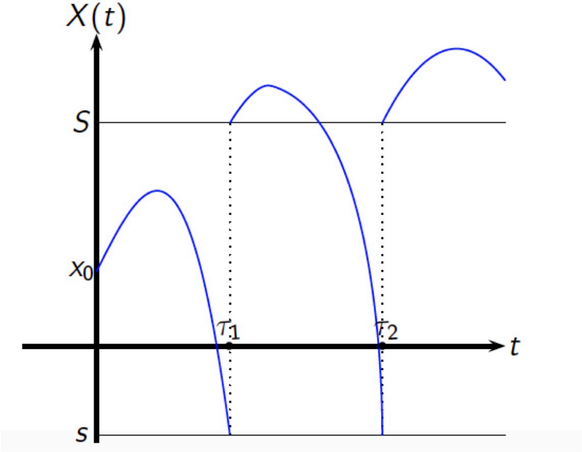

Question: Is the (s, S)-policy optimal? In what sense?

Long-term Average Cost

- $c_0 : \mathcal{I} \to \mathbb{R}^+$: holding/back-order cost rate.
- $c_1 : \overline{\mathcal{R}} \to \mathbb{R}^+$: ordering cost function, where

$$\mathcal{R} = \{(y, z) \in \mathcal{E}^2 : y < z\}, \ \ \overline{\mathcal{R}} = \{(y, z) \in \mathcal{E}^2 : y \le z\}.$$

- $\exists k_1 > 0$ s.t. $c_1 \ge k_1$; thus k_1 is the fixed cost for each order.

Long-term Average Cost

- $c_0 : \mathcal{I} \to \mathbb{R}^+$: holding/back-order cost rate.
- $c_1 : \overline{\mathcal{R}} \to \mathbb{R}^+$: ordering cost function, where

$$\mathcal{R} = \{(y,z) \in \mathcal{E}^2 : y < z\}, \quad \overline{\mathcal{R}} = \{(y,z) \in \mathcal{E}^2 : y \le z\}.$$

- $\exists k_1 > 0$ s.t. $c_1 \ge k_1$; thus k_1 is the fixed cost for each order.

Long-term Average Cost:

$$J(\tau, Y) := \limsup_{t\to\infty} \frac{1}{t} \mathbb{E}_{x_0}\Bigg[\int_0^t c_0(X(s))\mathrm{d}s + \sum_{k=1}^{\infty} I_{\{\tau_k \le t\}} c_1(X(\tau_k-), X(\tau_k))\Bigg].$$

Selected Literature

Discounted Criterion

- **A. Bensoussan and J.L. Lions** (1984). *Impulse control and quasi-variational inequalities*. Gauthier-Villars, Montrouge; Heyden & Son, Inc., Philadelphia, PA.
- **L.H.R. Alvarez and J. Lempa** (2008). On the Optimal Stochastic Impulse Control of Linear Diffusions, *SIAM J. Control Optim.*, **47**, 703–742.
- **M. Egami** (2008). A Direct Solution Method for Stochastic Impulse Control Problems of One-Dimensional Diffusions, *SIAM J. Control Optim.*, **47**, 1191–1218.
- **K.J. Arrow, S. Karlin and H. Scarf** (1958). Studies in the Mathematical Theory of Inventory and Production, SUP, Stanford.
- **A. Bensoussan, R.H. Liu, and S.P. Sethi** (2005). Optimality of an (s, S) policy with compound Poisson and diffusion demands: a quasi-variational inequalities approach, *SIAM J. Control Optim.*, **44**, 1650–1676.
- **L. Benkherouf and A. Bensoussan** (2009). Optimality of an (s, S) policy with compound Poisson and diffusion demands: a quasi-variational inequalities approach, *SIAM J. Control Optim.*, **48**, 756–762.

Selected Literature

Long-term Average Criterion

- **A. Bensoussan** (2011). *Dynamic programming and inventory control*, Studies in Probability, Optimization and Statistics, IOS Press, Amsterdam.
- **J. G. Dai and D. Yao** (2013). Brownian inventory models with convex holding cost, part 1: Average-optimal controls. *Stoch. Syst.*, 3(2):442–499.
- **S. He, D. Yao and H. Zhang** (2017). Optimal Ordering Policy for Inventory Systems with Quantity-Dependent Setup Costs. *Math. Oper. Res.*, 42(4): 979–1006.
- **K. L. Helmes, R.H. Stockbridge, and C.Z.** (2017). Continuous inventory models of diffusion type: long-term average cost criterion. *Annals Appl. Probab.*, 27(3): 1831–1885.
- **M. Ormeci, J.G. Dai, and J.V. Vate** (2008). Impulse control of Brownian motion: the constrained average cost case, *Operations Research*, **56(3)**, 618–629.
- **D. Yao, X. Chao and J. Wu** (2015). Optimal control policy for a Brownian inventory system with concave ordering cost. *J. Appl. Probab.*, 52:909–925.

Long-term Average Control Problems

The Usual Approaches

- Vanishing discount method ([**?**, **?**]):
  - Discount problem

$$\alpha v^\alpha(x) + F(x, Dv^\alpha(x), D^2 v^\alpha(x)) = 0,$$

  - the limiting behavior of $-\alpha v^\alpha$ and $v^\alpha(x) - v^\alpha(x_0)$ as $\alpha \downarrow 0$.
- Long-term average HJB equation [**?**, **?**, **?**]:

$$F(x, Du(x), D^2 u(x)) = \lambda.$$

Long-term Average Control Problems

The Usual Approaches

- Vanishing discount method ([**?**, **?**]):
  - Discount problem

$$\alpha v^\alpha(x) + F(x, Dv^\alpha(x), D^2 v^\alpha(x)) = 0,$$

  - the limiting behavior of $-\alpha v^\alpha$ and $v^\alpha(x) - v^\alpha(x_0)$ as $\alpha \downarrow 0$.
- Long-term average HJB equation [**?**, **?**, **?**]:

$$F(x, Du(x), D^2 u(x)) = \lambda.$$

- Often the state space needs to be bounded and/or controls need to be in a restricted class.

What We Intend to Do?

1. Formulation: general 1-dimensional diffusion model with general boundary behavior.
2. Holding/backorder and ordering costs: general and relaxed conditions.
3. New approach: weak convergence.
4. Result: the (s, S) policy is optimal in the admissible class of general impulse controls under very mild conditions.

Assumptions: Boundary Points

(a) Both the *speed measure M* and the *scale function S* of the process X_0 are absolutely continuous with respect to Lebesgue measure.
The operator of X_0 is

$$Af(x) = \frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}M}\left[\frac{\mathrm{d}f(x)}{\mathrm{d}S}\right], \qquad \forall x \in \mathcal{I}.$$

(b) The left boundary a is attracting and the right boundary b is non-attracting.
- Moreover, when b is a natural boundary, $M[y, b) < \infty$ for each $y \in \mathcal{I}$.
- The boundaries $a = -\infty$ and $b = \infty$ are required to be natural.

Assumptions: Boundary Points

(a) Both the *speed measure M* and the *scale function S* of the process X_0 are absolutely continuous with respect to Lebesgue measure.
The operator of X_0 is

$$Af(x) = \frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}M}\left[\frac{\mathrm{d}f(x)}{\mathrm{d}S}\right], \qquad \forall x \in \mathcal{I}.$$

(b) The left boundary a is attracting and the right boundary b is non-attracting.
- Moreover, when b is a natural boundary, $M[y, b) < \infty$ for each $y \in \mathcal{I}$.
- The boundaries $a = -\infty$ and $b = \infty$ are required to be natural.

Remark

These conditions are more general than the negative drift condition commonly used in the literature.

Assumptions: Costs

(a) The holding/back-order cost function $c_0 : \mathcal{I} \to \mathbb{R}^+$ is continuous. Moreover, at the boundaries

$$\lim_{x \to a} c_0(x) =: c_0(a) \in \overline{\mathbb{R}^+}, \quad \lim_{x \to b} c_0(x) =: c_0(b) \in \overline{\mathbb{R}^+};$$

we require $c_0(\pm\infty) = \infty$. Finally, for each $y \in \mathcal{I}$,

$$\int_y^b c_0(v) \mathrm{d}M(v) < \infty. \tag{1}$$

(b) The function $c_1 : \overline{\mathcal{R}} \to \overline{\mathbb{R}^+}$ is continuous with $c_1 \geq k_1 > 0$ for some constant k_1.

Assumptions: Costs

(a) The holding/back-order cost function $c_0 : \mathcal{I} \to \mathbb{R}^+$ is continuous. Moreover, at the boundaries

$$\lim_{x \to a} c_0(x) =: c_0(a) \in \overline{\mathbb{R}^+}, \quad \lim_{x \to b} c_0(x) =: c_0(b) \in \overline{\mathbb{R}^+};$$

we require $c_0(\pm\infty) = \infty$. Finally, for each $y \in \mathcal{I}$,

$$\int_y^b c_0(v) \mathrm{d}M(v) < \infty. \tag{1}$$

(b) The function $c_1 : \overline{\mathcal{R}} \to \overline{\mathbb{R}^+}$ is continuous with $c_1 \geq k_1 > 0$ for some constant k_1.

Some examples:

$$c_1(y, z) = k_1 + k_2(z - y), \quad c_1(y, z) = k_1 + k_2\sqrt{z - y}.$$

Some Remarks Concerning the Costs

In the literature ([**?**, **?**, **?**]),

- Usually the holding/back-order cost function is assumed to be convex, monotone, and/or with certain growth conditions.
- The ordering cost function usually takes the form of "fixed plus proportional cost" or is assumed to be concave.

Some Remarks Concerning the Costs

In the literature ([**?**, **?**, **?**]),

- Usually the holding/back-order cost function is assumed to be convex, monotone, and/or with certain growth conditions.
- The ordering cost function usually takes the form of "fixed plus proportional cost" or is assumed to be concave.
- Compared with our previous work ([**?**]), this work also removes many structural assumptions on c_0 and c_1.

  For example, the requirement that c_0 approaches ∞ at each boundary and the rather restrictive modularity condition of that paper on c_1 are unnecessary.

The Basic Strategy

1. First examine the inventory process under the (s, S)-policy with $s = y$ and $S = z$:
   The cost of such a policy is given by a nonlinear function $F(y, z), y < z \in \mathcal{E}$ and an optimal (s_*, S_*)-policy exists under certain conditions.

The Basic Strategy

1. First examine the inventory process under the (s, S)-policy with $s = y$ and $S = z$:
   The cost of such a policy is given by a nonlinear function $F(y, z), y < z \in \mathcal{E}$ and an optimal (s_*, S_*)-policy exists under certain conditions.
2. Next we construct a particular auxiliary function $G_0(x), x \in \mathcal{I}$. The function G_0 solves a system similar to but different from the long-term QVI.

The Basic Strategy

1. First examine the inventory process under the (s, S)-policy with $s = y$ and $S = z$:
   The cost of such a policy is given by a nonlinear function $F(y, z), y < z \in \mathcal{E}$ and an optimal (s_*, S_*)-policy exists under certain conditions.
2. Next we construct a particular auxiliary function $G_0(x), x \in \mathcal{I}$. The function G_0 solves a system similar to but different from the long-term QVI.
3. Finally we establish optimality of the (s_*, S_*) ordering policy in the general class of admissible policies via weak convergence and appropriate approximation of G_0.

The Inventory Process under the (s, S)-Policy

Let $(y, z) \in \mathcal{R}$. Define the ordering policy (τ, Y) by

$$\begin{cases} \tau_1 = \inf\{t \geq 0 : X(t-) = y\}, \\ \tau_k = \inf\{t > \tau_{k-1} : X(t-) = y\}, \; k \geq 2, \end{cases} \quad \text{and} \quad Y_k = z - y, \; k \geq 1.$$

- The process X has a unique stationary distribution with density π

$$\pi(x) = \begin{cases} 0, & a < x \leq y, \\ 2\kappa m(x) S[y, x], & y \leq x \leq z, \\ 2\kappa m(x) S[y, z], & z \leq x < b, \end{cases}$$

  in which $\kappa = \left(\int_y^z 2S[y, x] \mathrm{d}M(x) + 2S[y, z] M[z, b)\right)^{-1}$.

- The constant κ gives the expected frequency of orders.

The Inventory Process under the (s, S)-Policy (cont'd)

- The cost of such a policy is given by

$$J_0(\tau, Y) = \frac{c_1(y,z) + g_0(z) - g_0(y)}{\zeta(z) - \zeta(y)}, \qquad (2)$$

where

$$g_0(x) := 2\int_{x_0}^{x}\int_{u}^{b} c_0(v)\mathrm{d}M(v)\mathrm{d}S(u),$$
$$\zeta(x) := 2\int_{x_0}^{x} M[u,b)\mathrm{d}S(u),$$

and x_0 is the initial inventory.

- Note that for any $y < z \in \mathcal{E}$,

$$\mathbb{E}_z\left[\int_0^{\tau_y} c_0(X_0(s))\mathrm{d}s\right] = g_0(z) - g_0(y), \quad \mathbb{E}_z[\tau_y] = \zeta(z) - \zeta(y).$$

- Also $c_1(y, z)$ is the ordering cost.

- Note that for any $y < z \in \mathcal{E}$,

$$\mathbb{E}_z \left[\int_0^{\tau_y} c_0(X_0(s)) \mathrm{d}s \right] = g_0(z) - g_0(y), \quad \mathbb{E}_z[\tau_y] = \zeta(z) - \zeta(y).$$

- Also $c_1(y, z)$ is the ordering cost.
- Thus the right-hand side of Eq. (2) has a natural interpretation:

$$\frac{c_1(y,z) + g_0(z) - g_0(y)}{\zeta(z) - \zeta(y)} = \frac{\text{expected total cost per cycle}}{\text{expected cycle length}}.$$

The Function F

Define

$$F(y,z) := \frac{c_1(y,z) + g_0(z) - g_0(y)}{\zeta(z) - \zeta(y)}, \quad (y,z) \in \mathcal{R},$$

Suppose

(a) The boundary a is regular; or exit; or a is a natural boundary for which either (i) or (ii) hold:

  (i) $c_0(a) = \infty$;

  (ii) $c_0(a) < \infty$, the function $F_0(\cdot, z)$ is strictly decreasing in a neighborhood of a for each $z \in \mathcal{E}$ and there exists some $(\widehat{y}, \widehat{z}) \in \mathcal{R}$ such that $F_0(\widehat{y}, \widehat{z}) < c_0(a)$.

(b) The boundary b is entrance; or b is natural for which either (i) or (ii) hold:

  (i) $c_0(b) = \infty$;

  (ii) $c_0(b) < \infty$, $F_0(y, \cdot)$ is strictly increasing in a neighborhood of b for every $y \in \mathcal{E}$ and there exists some $(\widetilde{y}, \widetilde{z}) \in \mathcal{R}$ such that $F_0(\widetilde{y}, \widetilde{z}) < c_0(b)$.

Nonlinear Optimization and Optimal (s, S) Policy

Under these conditions:

- There exists a pair $(y_*, z_*) \in \mathcal{R}$ such that

$$F(y_*, z_*) = F_* := \inf \{F(y, z) : (y, z) \in \mathcal{R}\}. \quad (3)$$

- The (y_*, z_*)-policy is optimal in the class of all (s, S) ordering policies

$$F_* = F(y_*, z_*) = J_0(\tau^*, Y^*).$$

Nonlinear Optimization and Optimal (s, S) Policy

Under these conditions:

- There exists a pair $(y_*, z_*) \in \mathcal{R}$ such that

$$F(y_*, z_*) = F_* := \inf \{F(y, z) : (y, z) \in \mathcal{R}\}. \tag{3}$$

- The (y_*, z_*)-policy is optimal in the class of all (s, S) ordering policies

$$F_* = F(y_*, z_*) = J_0(\tau^*, Y^*).$$

The operators:

$$Af(x) = \frac{1}{2}\frac{\mathrm{d}}{\mathrm{d}M}\left[\frac{\mathrm{d}f(x)}{\mathrm{d}S}\right], \qquad \forall x \in \mathcal{I},$$

$$Bf(y, z) = f(z) - f(y), \qquad \forall (y, z) \in \overline{\mathcal{R}}.$$

The Auxiliary Function G_0

Define

$$G_0(x) = g_0(x) - F_*\zeta(x), \quad x \in \mathcal{E}.$$

- $G_0 \in C(\mathcal{E}) \cap C^2(\mathcal{I})$ and G_0 extends uniquely to $\overline{\mathcal{E}}$.
- G_0 is a solution of the system

$$\begin{cases} Af(x) + c_0(x) - F_* &= 0, \qquad x \in \mathcal{I}, \\ Bf(y,z) + c_1(y,z) &\geq 0, \qquad (y,z) \in \overline{\mathcal{R}} \\ f(x_0) &= 0, \\ Bf(y_0^*, z_0^*) + c_1(y_0^*, z_0^*) &= 0. \end{cases} \tag{4}$$

- The system (4) is similar to but different from the usual QVI for long-term average impulse control problem:

$$\min\left\{Au(x) + c_0(x) - F_*, \min_{z\in\mathcal{E}}[Bu(x,z) + c_1(x,z)]\right\} = 0, \ \forall x \in \mathcal{E}.$$

The Verification?

Let β_n be a localizing sequence. Then

$$\begin{aligned}
&\mathbb{E}_{x_0}[G_0(X(t \wedge \beta_n))] - G_0(x_0) \\
&= \mathbb{E}_{x_0}\left[\int_0^{t\wedge\beta_n} AG_0(X(s))\mathrm{d}s + \sum_{k=1}^{\infty} I_{\{\tau_k \leq t\wedge\beta_n\}} BG_0(X(\tau_k-), X(\tau_k))\right] \\
&\geq \mathbb{E}_{x_0}\left[\int_0^{t\wedge\beta_n} [F_* - c_0(X(s))]\mathrm{d}s - \sum_{k=1}^{\infty} I_{\{\tau_k \leq t\wedge\beta_n\}} c_1(X(\tau_k-), X(\tau_k))\right]
\end{aligned}$$

The Verification?

Let β_n be a localizing sequence. Then

$$\begin{aligned}
&\mathbb{E}_{x_0}[G_0(X(t\wedge\beta_n))] - G_0(x_0)\\
&= \mathbb{E}_{x_0}\left[\int_0^{t\wedge\beta_n} AG_0(X(s))\mathrm{d}s + \sum_{k=1}^{\infty} I_{\{\tau_k\le t\wedge\beta_n\}} BG_0(X(\tau_k-),X(\tau_k))\right]\\
&\ge \mathbb{E}_{x_0}\left[\int_0^{t\wedge\beta_n} [F_* - c_0(X(s))]\mathrm{d}s - \sum_{k=1}^{\infty} I_{\{\tau_k\le t\wedge\beta_n\}} c_1(X(\tau_k-),X(\tau_k))\right]
\end{aligned}$$

and hence

$$\begin{aligned}
&F_* - \liminf_{t\to\infty}\liminf_{n\to\infty}\frac{1}{t}\mathbb{E}_{x_0}[G_0(X(t\wedge\beta_n))]\\
&\le \limsup_{t\to\infty}\frac{1}{t}\mathbb{E}\left[\int_0^t c_0(X(s))\mathrm{d}s + \sum_{k=1}^{\infty} I_{\{\tau_k\le t\}} c_1(X(\tau_k-),X(\tau_k))\right] = J_0(\tau,Y)
\end{aligned}$$

The Verification?

Let β_n be a localizing sequence. Then

$$\mathbb{E}_{x_0}[G_0(X(t \wedge \beta_n))] - G_0(x_0)$$
$$= \mathbb{E}_{x_0}\left[\int_0^{t\wedge\beta_n} AG_0(X(s))\mathrm{d}s + \sum_{k=1}^{\infty} I_{\{\tau_k \le t\wedge\beta_n\}} BG_0(X(\tau_k-), X(\tau_k))\right]$$
$$\ge \mathbb{E}_{x_0}\left[\int_0^{t\wedge\beta_n} [F_* - c_0(X(s))]\mathrm{d}s - \sum_{k=1}^{\infty} I_{\{\tau_k \le t\wedge\beta_n\}} c_1(X(\tau_k-), X(\tau_k))\right]$$

and hence

$$F_* - \liminf_{t\to\infty} \liminf_{n\to\infty} \frac{1}{t}\mathbb{E}_{x_0}[G_0(X(t \wedge \beta_n))]$$
$$\le \limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t c_0(X(s))\mathrm{d}s + \sum_{k=1}^{\infty} I_{\{\tau_k \le t\}} c_1(X(\tau_k-), X(\tau_k))\right] = J_0(\tau, Y)$$

NOT SO FAST if the state space is unbounded.

Previous Attempts

- Use smooth pasting to construct a solution to the QVI.
  Many structural assumptions on the costs functions c_0, c_1 are required; not natural.

- Optimality in a smaller class:
  Optimality in the restricted class of impulse controls so that the transversality condition is satisfied.

- Ad hoc comparison result:
  For specific models, it is sufficient to consider policies whose order-to locations are uniformly bounded above.

Expected Occupation and Ordering Measures

Define

$$\mu_{0,t}(\Gamma_0) = \frac{1}{t}\mathbb{E}\left[\int_0^t I_{\Gamma_0}(X(s))\mathrm{d}s\right], \qquad \Gamma_0 \in \mathcal{B}(\mathcal{E}),$$

$$\mu_{1,t}(\Gamma_1) = \frac{1}{t}\mathbb{E}\left[\sum_{k=1}^{\infty} I_{\{\tau_k \le t\}} I_{\Gamma_1}(X(\tau_k-), X(\tau_k))\right], \qquad \Gamma_1 \in \mathcal{B}(\overline{\mathcal{R}}).$$

If a is a reflecting boundary, define the average expected local time measure $\mu_{2,t}$

$$\mu_{2,t}(\{a\}) = \tfrac{1}{t}\mathbb{E}[L_a(t)], \quad t > 0,$$

in which L_a denotes the local time of X at a.

Some Observations

We have the following observations:

- Let $\{t_i : i \in \mathbb{N}\}$ be such that $\lim_{i\to\infty} t_i = \infty$. If $\{\mu_{0,t_i} : i \in \mathbb{N}\}$ is not tight, then $J_0(\tau, Y) = \infty$.
- If $J_0(\tau, Y) < \infty$, then $\{\mu_{0,t}\}$ is tight.
- But $J_0(\tau, Y) < \infty$ does not necessarily imply that $\{\mu_{1,t}\}$ is tight.
- For each μ_0 attained as a weak limit of some sequence $\{\mu_{0,t_j}\}$ as $t_j \to \infty$, we have

$$\int_{\overline{\mathcal{E}}} c_0(x)\, \mu_0(dx) \leq J_0(\tau, Y) < \infty.$$

Approximation of the Function G_0

Define

$$G_n(x) = \frac{G_0(x)}{1 + \frac{1}{n} h(G_0(x))}, \quad x \in \mathcal{E}, \tag{5}$$

where

$$h(x) = \begin{cases} -\frac{1}{8}x^4 + \frac{3}{4}x^2 + \frac{3}{8}, & \text{for } |x| \leq 1, \\ |x|, & \text{for } |x| \geq 1. \end{cases}$$

Under some technical assumptions on G_0, we can show that

$$\lim_{n \to \infty} AG_n(x) = AG_0(x) \qquad \forall x \in \mathcal{I}$$

$$\lim_{n \to \infty} BG_n(y, z) = BG_0(y, z) \qquad \forall (y, z) \in \overline{\mathcal{R}}.$$

Key Observations

Let $(\tau, Y) \in \mathcal{A}_0$ with $J_0(\tau, Y) < \infty$. Let $\{t_j : j \in \mathbb{N}\}$ be a sequence such that $\lim_{j\to\infty} t_j = \infty$ and

$$J_0(\tau, Y) = \lim_{j\to\infty} \frac{1}{t_j} \mathbb{E}\left[\int_0^{t_j} c_0(X(s))\mathrm{d}s + \sum_{k=1}^{\infty} I_{\{\tau_k \le t_j\}} c_1(X(\tau_k-), X(\tau_k))\right]$$

$$= \lim_{j\to\infty} \left(\int_{\overline{\mathcal{E}}} c_0(x)\, \mu_{0,t_j}(\mathrm{d}x) + \int_{\overline{\mathcal{R}}} c_1(y, z)\, \mu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z)\right).$$

Key Observations

Let $(\tau, Y) \in \mathcal{A}_0$ with $J_0(\tau, Y) < \infty$. Let $\{t_j : j \in \mathbb{N}\}$ be a sequence such that $\lim_{j\to\infty} t_j = \infty$ and

$$J_0(\tau, Y) = \lim_{j\to\infty} \frac{1}{t_j}\mathbb{E}\left[\int_0^{t_j} c_0(X(s))\mathrm{d}s + \sum_{k=1}^{\infty} I_{\{\tau_k \le t_j\}} c_1(X(\tau_k-), X(\tau_k))\right]$$
$$= \lim_{j\to\infty}\left(\int_{\overline{\mathcal{E}}} c_0(x)\,\mu_{0,t_j}(\mathrm{d}x) + \int_{\overline{\mathcal{R}}} c_1(y, z)\,\mu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z)\right).$$

We have

$$0 = \lim_{j\to\infty}\left(\int_{\overline{\mathcal{E}}} AG_n(x)\,\mu_0(\mathrm{d}x) + \int_{\overline{\mathcal{R}}} BG_n(y, z)\,\mu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z)\right),\ \forall n \in \mathbb{N}$$

and

$$\liminf_{n\to\infty}\liminf_{j\to\infty}\int_{\overline{\mathcal{R}}}(BG_n(y, z) + c_1(y, z))\,\mu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z) \ge 0,$$
$$\liminf_{n\to\infty}\int_{\overline{\mathcal{E}}}(AG_n(x) + c_0(x))\,\mu_0(\mathrm{d}x) \ge \int_{\overline{\mathcal{E}}}(AG_0(x) + c_0(x))\,\mu_0(\mathrm{d}x) \ge F_*.$$

Optimality

$$
\begin{aligned}
&J_0(\tau, Y) \\
&\quad = \liminf_{n\to\infty} \lim_{j\to\infty} \bigg(\int_{\overline{\mathcal{E}}} (AG_n(x) + c_0(x))\, \mu_{0,t_j}(\mathrm{d}x) \\
&\qquad\qquad + \int_{\overline{\mathcal{R}}} (BG_n(y,z) + c_1(y,z))\, \mu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z) \bigg) \\
&\quad \geq \liminf_{n\to\infty} \liminf_{j\to\infty} \int_{\overline{\mathcal{E}}} (AG_n(x) + c_0(x))\, \mu_{0,t_j}(\mathrm{d}x) \\
&\qquad\qquad + \liminf_{n\to\infty} \liminf_{j\to\infty} \int_{\overline{\mathcal{R}}} (BG_n(y,z) + c_1(y,z))\, \mu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z) \\
&\quad \geq \liminf_{n\to\infty} \int_{\overline{\mathcal{E}}} (AG_n(x) + c_0(x))\, \mu_0(\mathrm{d}x) \\
&\qquad\qquad + \liminf_{n\to\infty} \liminf_{j\to\infty} \int_{\overline{\mathcal{R}}} (BG_n(y,z) + c_1(y,z))\, \mu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z) \\
&\quad \geq F_0^*.
\end{aligned}
$$

The Main Result

Theorem

(a) *Let* $(\tau, Y) \in \mathcal{A}_0$. *Then*

$$J_0(\tau, Y) \geq F_*.$$

(b) *Moreover, the* (s, S)*-policy with* $s = y_*$ *and* $S = z_*$ *is an optimal impulse policy.*

Some Remarks

- The ergodicity of the inventory process under the (s, S)-policy gives $J(\tau^*, Y^*) = F_*$.
- To establish the optimality of the (s_*, S_*)-policy in the general admissible class $\mathcal{A}_0$, the usual approach needs to solve the associated QVI and verify the transversality condition. These are not easy; in particular if the state space is unbounded.
- This work proposes a weak convergence approach together with an appropriate approximation of the function G_0 to establish the optimality.

Drifted Brownian motion inventory models

- The classical model drifted Brownian motion on $(-\infty, \infty)$
  - State:

$$X_0(t) = x_0 - \mu t + \sigma W(t), \quad x_0 \in (-\infty, \infty)$$

  - Costs:

$$c_0(x) = \begin{cases} -c_b\, x, & x < 0, \\ c_h\, x, & x \geq 0, \end{cases} \quad \text{and} \quad c_1(y,z) = k_1 + k_2(z-y)$$

  in which $c_b, c_h, k_1, k_2 > 0$.

- Drifted Brownian motion with reflection at $\{0\}$
  - State:

$$X_0(t) = x_0 - \mu t + \sigma W(t) + L_0(t), \; x_0 \in [0, \infty)$$

  - Costs:

$$c_0(x) = k_3 x + k_4 e^{-x}, \quad \text{and} \quad c_1(y,z) = k_1 + k_2\sqrt{z-y},$$

- In both models, the (s, S)-policy is optimal.

A Counter-intuitive Example

- Inventory level (in absence if ordering) is $X_0(t) = W(t) - t$, for $t \geq 0$,
- The cost functions are specified

$$c_0(x) = 2|x|, \ \forall x \in \mathbb{R}, \ \text{and} \ c_1(y,z) = k_1 + (z - y), \ \forall (y,z) \in \overline{\mathcal{R}}.$$

- A special ordering policy (τ, Y):
  - It runs in cycles, each of which is composed of two phases.
  - For cycle $i = 1, 2, 3, \ldots$,
    - Phase 1 consists of using the $(0, 1)$-ordering policy a total of 2^{i-1} times; the length of each sub-cycle is a random variable having mean 1.
    - Phase 2 involves a single $(0, 2^{(i-1)/2})$-ordering policy followed immediately by using the $(2^{(i-1)/2}, 2^{(i-1)/2})$-ordering policy 2^{i-1} times.
- Then (a) $J_0(\tau, Y) < \infty$; (b) $\{\mu_{0,t} : t > 0\}$ is tight as $t \to \infty$; and (c) $\{\mu_{1,t} : t > 0\}$ is not tight.

Geometric Brownian Motion Inventory Model

- In the absence of ordering, the inventory process X_0 satisfies

$$dX_0(t) = -\mu X_0(t)\mathrm{d}t + \sigma X_0(t)\mathrm{d}W(t), \quad X(0) = x_0 \in \mathcal{I} = (0, \infty),$$

in which $\mu, \sigma > 0$ and W is a standard Brownian motion.

- The scale and speed measures are: for $[l, x] \subset \mathcal{I}$

$$S[l, x] = \tfrac{\sigma^2}{2\mu+\sigma^2}\left[x^{1+2\mu/\sigma^2} - l^{1+2\mu/\sigma^2}\right] \quad \text{and}$$
$$M[l, x] = \tfrac{1}{2\mu+\sigma^2}\left[l^{-1-2\mu/\sigma^2} - x^{-1-2\mu/\sigma^2}\right].$$

- Note that $S(0, x] < \infty$ and $S[x, \infty) = \infty$ and thus 0 is attracting and ∞ is non-attracting. Both boundaries are natural.

Geometric Brownian Motion Inventory Model

- State:

$$dX_0(t) = -\mu X_0(t)\,dt + \sigma X_0(t)\,dW(t), \quad X(0) = x_0 \in \mathcal{I} = (0, \infty),$$

  in which $\mu, \sigma > 0$.
- Costs:
  - Case 1

$$c_0(x) = k_3 x + k_4 x^{\beta}, \quad \text{and} \quad c_1(y, z) = k_1 + k_2\sqrt{z - y},$$

  - Case 2

$$c_0(x) = \begin{cases} k_4(1 - x), & 0 \le x \le 1, \\ k_3(x - 1), & x \ge 1 \end{cases}$$

$$c_1(y, z) = k_1 + \frac{k_2}{2}(y^{-1/2} - z^{-1/2}) + \frac{k_2}{2}(z - y)$$

- In both models, the (s, S)-policy is optimal.

Examples when the (s, S)-policy is NOT optimal

- State:

$$dX_0(t) = -\mu X_0(t)\, dt + \sigma X_0(t)\, dW(t), \quad X(0) = x_0 \in \mathcal{I} = (0, \infty),$$

  in which $\mu, \sigma > 0$.

- Costs:

$$c_1(y, z) = k_1 + k_2(z^\eta - y^\eta), \quad c_0(x) = k_3 x + k_4 x^\beta.$$

  where $0 < \eta \leq 1$ and $\beta < 0$.

- $k_4 = 0$, $k_2, k_3 > 0$: "no-order" policy is optimal
- $k_2 = k_3 = 0$, $k_4 > 0$: no optimal inventory control policy.

Summary

In this work, we

- formulated an impulse inventory control problem for a general one-dimensional diffusion with general boundary conditions under the long-term average cost criterion,
- used a weak convergence approach together with an appropriate approximation of the function G_0 to establish the optimality of the (s_*, S_*) policy in the general admissible class of impulse controls.
- provided a nonlinear optimization approach to determine the optimal levels s_* and S_*,
- studied geometric and drifted Brownian motion inventory examples for illustration.

Thank you!