On the Modelling of Impulse Control with Random Effects for Continuous Markov Processes with Application to Ergodic Inventory Control Models

K.L. Helmes, R.H. Stockbridge and C. Zhu

Humboldt Universität zu Berlin and University of Wisconsin - Milwaukee

August 2, 2023

The 18th Workshop on Markov Processes and Related Topics

Outline

Impulse-Controlled Process: Intuitive Description

Example:

$$X_t = x_0 + \int_0^t \mu(X_s)\,\mathrm{d}s + \int_0^t \sigma(X_s)\,\mathrm{d}W_s + \sum_{k:\tau_k \le t} Z_k.$$

Here $(\tau, Z) = \{(\tau_k, Z_k) : k \in \mathbb{N}\}$ is an impulse policy satisfying

- each τ_k is a stopping time;
- each Z_k is a measurable r.v. with respect to the information available at time τ_k; and
- the sequence $\{\tau_k\}$ is non-decreasing.

Question: What is a precise model for such a process?

The Usual Approach

Robin (1978), Stettner (1983), Lepeltier and Marchal (1984), etc.

Define the countable product measurable space

$$\widetilde{\Omega} = \prod_{k=0}^{\infty} \Omega_k = \prod_{k=0}^{\infty} \Omega \qquad \widetilde{\mathcal{G}} = \bigotimes_{k=0}^{\infty} \mathscr{F}_k = \bigotimes_{k=0}^{\infty} \mathcal{F}.$$

Intuition. Use component Ω_k to determine the impulse-controlled process X over the interval $[\tau_k, \tau_{k+1})$.

Question. What restrictions are imposed on the policy

$$(\tau, Z) = \{(\tau_k, Z_k) : k \in \mathbb{N}\}?$$

The Usual Approach

Robin (1978), Stettner (1983), Lepeltier and Marchal (1984), etc.

Define the countable product measurable space

$$\widetilde{\Omega} = \prod_{k=0}^{\infty} \Omega_k = \prod_{k=0}^{\infty} \Omega \qquad \widetilde{\mathcal{G}} = \bigotimes_{k=0}^{\infty} \mathscr{F}_k = \bigotimes_{k=0}^{\infty} \mathcal{F}.$$

Intuition. Use component Ω_k to determine the impulse-controlled process X over the interval $[\tau_k, \tau_{k+1})$.

Question. What restrictions are imposed on the policy

$$(\tau, Z) = \{(\tau_k, Z_k) : k \in \mathbb{N}\}?$$

Each τ_k must be a stopping time ... but *with respect to which filtration*?

Different Filtrations for Different Interventions

- τ_1 must be a stopping time w.r.t. the filtration $\{\mathcal{F}_t^{(0)}\}$ with

$$\mathcal{F}_t^{(0)} = \sigma(X_0(u) : 0 \leq u \leq t)$$

  in which X_0 is the coordinate process on Ω_0. The impulse Z_1 is $\mathcal{F}_{\tau_1}^{(0)}$-measurable.

- τ_2 must be a stopping time w.r.t. the filtration $\{\mathcal{F}_t^{(1)}\}$ where

$$\mathcal{F}_t^{(1)} = \sigma(X_1(\tau_1 + u) : 0 \leq u \leq t)$$

  with X_1 being the coordinate process on Ω_1. The impulse Z_2 is $\mathcal{F}_{\tau_2}^{(1)}$-measurable.

- In general, τ_k must be a stopping time w.r.t. the filtration $\{\mathcal{F}_t^{(k-1)}\}$ having

$$\mathcal{F}_t^{(k-1)} = \sigma(X_k(\tau_{k-1} + u) : 0 \leq u \leq t);$$

  again, X_k denotes the coordinate process on component Ω_k. The impulse Z_k is required to be $\mathcal{F}_{\tau_k}^{(k-1)}$-measurable.

Impulse Control Model for Continuous Processes

Our Contribution

- $\Omega = D_{\mathcal{E}}[0, \infty)$.
- All decisions are made relative to *the natural filtration* generated by the coordinate process X.
- The interventions have random effects; namely, each intervention selects a distribution of the new location following the impulse.
- Identify a class of policies for which the controlled process is Markov.
- Determine a class of policies for which the controlled process has independent (and identically distributed) cycles.

Model Fundamentals

- $\mathcal{E}$, the state space (a complete separable metric space).
- $\Omega := D_{\mathcal{E}}[0,\infty)$, the space of càdlàg functions.
- $X : \Omega \to D_{\mathcal{E}}[0,\infty)$, the coordinate process so $X(t,\omega) = \omega(t)$ for all $t \geq 0$ and $\omega \in \Omega$.
- $\mathcal{F} = \sigma(X(t) : t \geq 0)$.
- $\{\mathcal{F}_t\}$ is the natural filtration: $\mathcal{F}_t := \sigma(X(s), 0 \leq s \leq t)$.
- $\{\mathbb{P}_x, x \in \mathcal{E}\}$ is a family of probability measures on $(\Omega, \mathcal{F})$ so that

$$(\Omega, \mathcal{F}, X, \{\mathcal{F}_t\}, \{\mathbb{P}_x, x \in \mathcal{E}\})$$

  is a Markov family.

Standing Assumption

For each $x \in \mathcal{E}$, $\mathbb{P}_x$ has its support in $C_{\mathcal{E}}[0,\infty) \subset \Omega$.

Model Fundamentals: Uncertain Impulse Mechanism

- Let $(\mathcal{Z}, \mathfrak{Z})$ be a measurable space representing the impulse control decisions.
- Let $\mathbb{Q} = \{Q_{(y,z)} : (y, z) \in \mathcal{E} \times \mathcal{Z}\}$ be a given family of probability measures on $\mathcal{E}$ such that

$$\text{for each } \Gamma \in \mathcal{B}(\mathcal{E}), \text{ the mapping}$$
$$(y, z) \mapsto Q_{(y,z)}(\Gamma) \text{ is } \mathcal{B}(\mathcal{E}) \otimes \mathfrak{Z}\text{-measurable.}$$

To Accommodate a Possible First Jump at Time 0 ...

- Every $\omega \in D_{\mathcal{E}}[0,\infty)$ is right continuous at 0
  $\rightsquigarrow$ This precludes the possibility of an intervention occurring at time 0.

To Accommodate a Possible First Jump at Time 0 ...

- Every $\omega \in D_{\mathcal{E}}[0,\infty)$ is right continuous at 0
  $\rightsquigarrow$ This precludes the possibility of an intervention occurring at time 0.
- We need to *augment* the space $D_{\mathcal{E}}[0,\infty)$ so that it contains the location from which the intervention occurs.
  1. Set $\check{\Omega} = \mathcal{E} \times D_{\mathcal{E}}[0,\infty)$ and $\check{\mathcal{F}} = \mathcal{B}(\mathcal{E}) \otimes \mathcal{F}$.
     Denote elements $\check{\omega} \in \check{\Omega}$ by $\check{\omega} = (\check{\omega}(0-), \check{\omega}(\cdot))$.
  2. Extend the coordinate process X on $D_{\mathcal{E}}[0,\infty)$ to $\check{\Omega}$ by defining $X(0-,\check{\omega}) = \check{\omega}(0-)$ while keeping $X(s) = \check{\omega}(s)$ for $s \geq 0$.
  3. Set $\check{\mathcal{F}}_t = \mathcal{B}(\mathcal{E}) \otimes \mathcal{F}_t$, $\check{\mathcal{F}}_{t-} = \mathcal{B}(\mathcal{E}) \otimes \mathcal{F}_{t-}$, for $t \geq 0$.
  4. For each $x \in \mathcal{E}$, extend the measure $\mathbb{P}_x$ on $(\Omega, \mathcal{F})$ to a measure $\check{\mathbb{P}}_x$ on $(\check{\Omega}, \check{\mathcal{F}})$ by putting full mass on the subset $\{\check{\omega} \in \check{\Omega} : \check{\omega}(0-) = x\}$.

To Accommodate a Possible First Jump at Time 0 ...

- Every $\omega \in D_{\mathcal{E}}[0, \infty)$ is right continuous at 0
  $\rightsquigarrow$ This precludes the possibility of an intervention occurring at time 0.
- We need to *augment* the space $D_{\mathcal{E}}[0, \infty)$ so that it contains the location from which the intervention occurs.
  1. Set $\check{\Omega} = \mathcal{E} \times D_{\mathcal{E}}[0, \infty)$ and $\check{\mathcal{F}} = \mathcal{B}(\mathcal{E}) \otimes \mathcal{F}$.
     Denote elements $\check{\omega} \in \check{\Omega}$ by $\check{\omega} = (\check{\omega}(0-), \check{\omega}(\cdot))$.
  2. Extend the coordinate process X on $D_{\mathcal{E}}[0, \infty)$ to $\check{\Omega}$ by defining $X(0-, \check{\omega}) = \check{\omega}(0-)$ while keeping $X(s) = \check{\omega}(s)$ for $s \geq 0$.
  3. Set $\check{\mathcal{F}}_t = \mathcal{B}(\mathcal{E}) \otimes \mathcal{F}_t$, $\check{\mathcal{F}}_{t-} = \mathcal{B}(\mathcal{E}) \otimes \mathcal{F}_{t-}$, for $t \geq 0$.
  4. For each $x \in \mathcal{E}$, extend the measure $\mathbb{P}_x$ on $(\Omega, \mathcal{F})$ to a measure $\check{\mathbb{P}}_x$ on $(\check{\Omega}, \check{\mathcal{F}})$ by putting full mass on the subset $\{\check{\omega} \in \check{\Omega} : \check{\omega}(0-) = x\}$.
- $(\check{\Omega}, \check{\mathcal{F}}, X, \{\check{\mathcal{F}}_t\}, \{\check{\mathbb{P}}_x : x \in \mathbb{E}\})$ is still a Markov family.

Nominal Impulse Policy

A nominal impulse policy $(\tau, Z) = \{(\tau_k, Z_k) : k \in \mathbb{N}\}$ is a sequence of pairs defined on $(\check{\Omega}, \check{\mathcal{F}})$ in which:

- τ_1 is an $\{\check{\mathcal{F}}_{t-}\}$-stopping time and for $k \geq 2$, τ_k is an $\{\mathcal{F}_{t-}\}$-stopping time;
- for each $k \in \mathbb{N}$, on the set $\{\tau_k < \infty\}$, $\tau_{k+1} > \tau_k$;
- $\lim_{k\to\infty} \tau_k = \infty$;
- for each $k \in \mathbb{N}$, Z_k is a $\mathcal{Z}$-valued, $\mathcal{F}_{\tau_k-}/\mathfrak{Z}$-measurable random variable (Z_1 being $\check{\mathcal{F}}_{\tau_1-}/\mathfrak{Z}$-measurable).

The Existence Result

Theorem 1

Let (τ, Z) be a nominal impulse policy. For each $k \in \mathbb{N}$, define the pre-impulse location $Y_k = X(\tau_k -)$ with the nominal impulse being Z_k on the set $\{\tau_k < \infty\}$. Set $\tau_0 = 0$. Then there exists a family of probability measures $\{\mathbb{P}_x^{(\tau,Z)} : x \in \mathcal{E}\}$ on $(\check{\Omega}, \check{\mathcal{F}})$ under which the coordinate process X satisfies the following properties:

(a) *under $\mathbb{P}_x^{(\tau,Z)}$ for each $x \in \mathcal{E}$, $X(0-) = x$ a.s. and moreover, for each $k \in \mathbb{N}$,*
   (i) *X is the fundamental Markov process on the interval $[\tau_{k-1}, \tau_k)$;*
   (ii) *on the set $\{\tau_k < \infty\}$, $Q_{(Y_k, Z_k)}$ is a regular conditional distribution of $X(\tau_k)$ given $\mathcal{F}_{\tau_k -}$; and*

(b) *for each $F \in \check{\mathcal{F}}$, the mapping $x \mapsto \mathbb{P}_x^{(\tau,Z)}(F)$ is universally measurable.*

Stationary Markov nominal impulse policy

A *stationary Markov nominal impulse policy* is a nominal impulse policy $(\tau, Z) = \{(\tau_k, Z_k) : k \in \mathbb{N}\}$ for which there exist measurable functions $\sigma : \Omega \to (0, \infty]$ and $\mathfrak{z} : \mathcal{E} \to \mathcal{Z}$ such that:

(a) for each $k \geq 1$, $\tau_k = \tau_{k-1} + \sigma \circ \theta_{\tau_{k-1}}$; and on the event $\{\tau_{k-1} < \infty\}$, for each $u \geq 0$,

$$\{\sigma \circ \theta_{\tau_{k-1}} > u\} \subset \{\sigma \circ \theta_{\tau_{k-1}} = u + \sigma \circ \theta_{\tau_{k-1}+u}\};$$

(b) $Z_k = \mathfrak{z}(X(\tau_k -))$.

Stationary Markov nominal impulse policy

A *stationary Markov nominal impulse policy* is a nominal impulse policy $(\tau, Z) = \{(\tau_k, Z_k) : k \in \mathbb{N}\}$ for which there exist measurable functions $\sigma : \Omega \to (0, \infty]$ and $\mathfrak{z} : \mathcal{E} \to \mathcal{Z}$ such that:

(a) for each $k \geq 1$, $\tau_k = \tau_{k-1} + \sigma \circ \theta_{\tau_{k-1}}$; and on the event $\{\tau_{k-1} < \infty\}$, for each $u \geq 0$,

$$\{\sigma \circ \theta_{\tau_{k-1}} > u\} \subset \{\sigma \circ \theta_{\tau_{k-1}} = u + \sigma \circ \theta_{\tau_{k-1}+u}\};$$

(b) $Z_k = \mathfrak{z}(X(\tau_k-))$.

Theorem 2

For a stationary Markov nominal impulse policy (τ, Z), $(\check{\Omega}, \check{\mathcal{F}}, X, \{\check{\mathcal{F}}_t\}, \{\mathbb{P}_x^{(\tau,Z)} : x \in \mathbb{E}\})$ *is a Markov family.*

Policies with Independent Cycles

An *independent-cycles nominal impulse policy* is a nominal impulse policy (τ, Z) for which for each $k \in \mathbb{N}$:

(a) there exists a random time σ_k such that $\tau_k = \tau_{k-1} + \sigma_k \circ \theta_{\tau_{k-1}}$ on the set $\{\tau_{k-1} < \infty\}$; and

(b) on the event $\{\tau_k < \infty\}$, the intervention (τ_k, Z_k) is such that $Q_{(Y_k(\omega), Z_k(\omega))}$ does not depend on ω.

Policies with Independent Cycles

An *independent-cycles nominal impulse policy* is a nominal impulse policy (τ, Z) for which for each $k \in \mathbb{N}$:

(a) there exists a random time σ_k such that $\tau_k = \tau_{k-1} + \sigma_k \circ \theta_{\tau_{k-1}}$ on the set $\{\tau_{k-1} < \infty\}$; and

(b) on the event $\{\tau_k < \infty\}$, the intervention (τ_k, Z_k) is such that $Q_{(Y_k(\omega), Z_k(\omega))}$ does not depend on ω.

Proposition 3

Let $\nu \in \mathcal{P}(\mathbb{E})$, (τ, Z) be an independent cycles nominal impulse policy and let $(\check{\Omega}, \check{\mathcal{F}}, \mathbb{P}_\nu^{(\tau,Z)})$ and X be the probability space and coordinate process, respectively. Then the cycles $\{X(t) : \tau_k \leq t < \tau_{k+1}\}$ for $k \in \mathbb{N}_0$ are independent.

An Example

- Let $\mathcal{E} = \mathbb{R}$ and fix $y \in \mathbb{R}$.
- Select the target location $z > y$ to which the controlled process aims to jump.
- Suppose $\text{supp}(Q_{(y,z)}) = [y_1, z]$, with $y_1 > y$.

An Example

- Let $\mathcal{E} = \mathbb{R}$ and fix $y \in \mathbb{R}$.
- Select the target location $z > y$ to which the controlled process aims to jump.
- Suppose $\text{supp}(Q_{(y,z)}) = [y_1, z]$, with $y_1 > y$.

Challenge: To define a nominal policy (τ, Z) for every $\check{\omega} \in \check{\Omega}$, not just those with specific discontinuities.

An Example

- Let $\mathcal{E} = \mathbb{R}$ and fix $y \in \mathbb{R}$.
- Select the target location $z > y$ to which the controlled process aims to jump.
- Suppose $\text{supp}(Q_{(y,z)}) = [y_1, z]$, with $y_1 > y$.

Challenge: To define a nominal policy (τ, Z) for every $\check{\omega} \in \check{\Omega}$, not just those with specific discontinuities.

- $Z_k(\check{\omega}) = z$, $\forall k \in \mathbb{N}, \check{\omega} \in \check{\Omega}$.
- $\tau_1(\check{\omega}) := \inf\{t \geq 0 : \check{\omega}(t-) \leq y, \check{\omega}(s) > y \text{ for } 0 \leq s < t\} =: \sigma(\check{\omega})$.
- for $k \in \mathbb{N}$, define $\tau_{k+1}(\check{\omega}) = \infty$ if $\check{\omega}(\tau_k) \notin [y_1, z]$; otherwise, define

$$\begin{aligned}\tau_{k+1}(\check{\omega}) &= \inf\{t \geq \tau_k(\check{\omega}) : \check{\omega}(t-) = y, \check{\omega}(s) > y, \forall s \in [\tau_k(\check{\omega}), t)\} \\ &= \tau_k(\check{\omega}) + \sigma \circ \theta_{\tau_k}(\check{\omega}).\end{aligned}$$

Ergodic Inventory Control with Random Effects

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 $-\infty \leq a < b \leq +\infty$.

Ergodic Inventory Control with Random Effects

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 $-\infty \leq a < b \leq +\infty$.

- demands tend to reduce the inventory $\rightsquigarrow$ a is an *attracting* point:

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

  a may be a regular (reflective or sticky), 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:

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

  b may be an entrance or natural point.

- Let $\mathcal{R} = \{(y, z) \in \mathcal{E}^2 : y < z\}$, where y denotes the pre-order and z the *nominal* post-order inventory levels, resp.

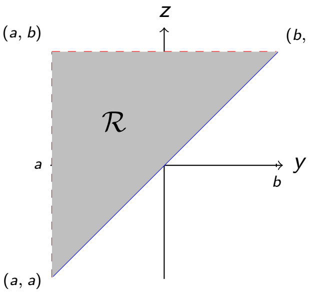

Formulation: Uncertain Impulse Mechanism

Let $\mathcal{Q} = \{Q(\cdot\,; y, z) : (y, z) \in \overline{\mathcal{R}}\}$ denote the collection of probability measures so that

(i) for each $(y, z) \in \overline{\mathcal{R}}$, $Q(\cdot\,; y, z) \in \mathcal{P}(\mathcal{E})$;

(ii) for each $\Gamma \in \mathcal{B}(\mathcal{E})$, the mapping

$$(y, z) \mapsto Q(\Gamma; y, z) \text{ is } \mathcal{B}(\mathcal{R})\text{-measurable.}$$

Formulation: Uncertain Impulse Mechanism

Let $\mathcal{Q} = \{Q(\cdot\,; y, z) : (y, z) \in \overline{\mathcal{R}}\}$ denote the collection of probability measures so that

(i) for each $(y, z) \in \overline{\mathcal{R}}$, $Q(\cdot\,; y, z) \in \mathcal{P}(\mathcal{E})$;

(ii) for each $\Gamma \in \mathcal{B}(\mathcal{E})$, the mapping

$$(y, z) \mapsto Q(\Gamma; y, z) \text{ is } \mathcal{B}(\mathcal{R})\text{-measurable.}$$

Examples:

1. $Q(\cdot\,; y, z) = \delta_{\{z\}}(\cdot)$
2. $Q(\cdot\,; y, z) = \theta\delta_{\{z\}}(\cdot) + (1-\theta)\text{Unif}(y, z)$, $\theta \in (0, 1)$.
3. $Q(\cdot\,; y, z) = \text{Unif}((1 - (z/k)^j)y + (z/k)^j z, z)$.
4. ...

Formulation: Admissible Policies

- $\mathcal{A} = \{(\tau, Z) = (\tau_k, Z_k), k = 1, 2, \ldots\}$, in which
  - $\{\tau_k\}$ is a strictly increasing sequence of $\{\mathcal{F}_{t-}\}$-stopping times with $\lim_{k\to\infty} \tau_k = \infty$,
  - $Z_k \in \mathcal{E}$ is $\{\mathcal{F}_{\tau_k -}\}$-measurable with $Z_k > X(\tau_k -)$.
- 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

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

  where L_a denotes the local time of X at a.

Remarks:

- Z_k is the *nominal "order-to"* location.
- The actual post-jump location $X(\tau_k)$ is determined by $Q(\cdot; X(\tau_k -), Z_k) \in \mathcal{P}(\mathcal{E})$ and may be different from Z_k.

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

Long-term Average Cost:

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

The (s, S) Policy

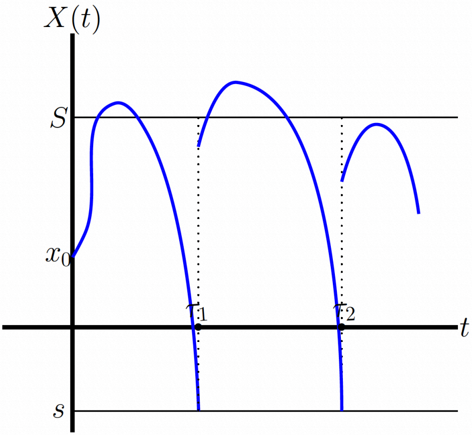

Questions: Does an optimal policy exist? Is the (s, S)-policy optimal? How to find an optimal (s, S) policy?

The 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 $H_0(y, z), y < z$.
   - An optimal (s_*, S_*)-policy exists under certain conditions.

The 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 $H_0(y, z), y < z$.
   - An optimal (s_*, S_*)-policy exists under certain conditions.
2. Then we establish optimality of the (s_*, S_*) ordering policy in the general class of admissible policies via weak convergence.

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

- The process X has a unique stationary distribution; and the long-run frequency of orders can be found.
- Then, the cost of such a policy $(s = y, S = z)$ is given by

$$J(\tau, Y) = \frac{\widehat{c_1}(y, z) + \widehat{Bg_0}(y, z)}{\widehat{B\zeta}(y, z)}, \tag{1}$$

where

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

and $x_0 \in \mathcal{I}$ is the initial inventory.

For any f and $(y, z) \in \overline{\mathcal{R}}$,

$$Bf(y, z) := f(z) - f(y), \qquad \widehat{f}(y, z) := \int_{y}^{z} f(y, v)Q(\mathrm{d}v; y, z).$$

Nonlinear Optimization of H_0

- Define

$$H_0(y,z) := \frac{\widehat{c}_1(y,z) + \widehat{Bg_0}(y,z)}{\widehat{B\zeta}(y,z)}, \forall (y,z) \in \mathcal{R}.$$

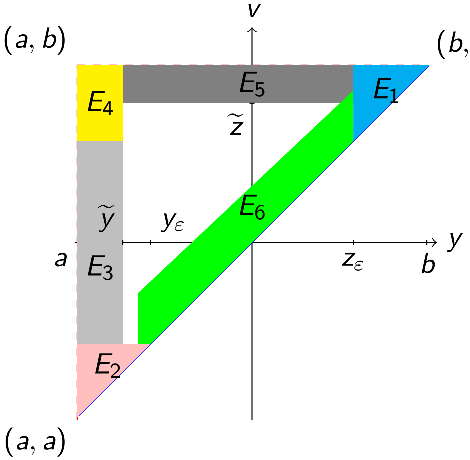

Nonlinear Optimization and Optimal (s, S) Policy

Under certain conditions:

- H_0 is lower semicontinuous on compact subsets of $\mathcal{R}$.
- There exists a pair $(y_0^*, z_0^*) \in \mathcal{R}$ such that

$$H_0(y_0^*, z_0^*) = H_0^* := \inf \left\{ H_0(y, z) : (y, z) \in \overline{\mathcal{R}} \right\}. \quad (2)$$

- The (y_0^*, z_0^*)-policy is optimal in the class of all (s, S) ordering policies

$$H_0^* = H_0(y_0^*, z_0^*) = J(\tau^*, Z^*).$$

Remark: If $H_0^* = 0$, then there is no optimal policy.

Expected Occupation and Ordering Measures

Define for $t > 0$

$$\begin{aligned} \mu_{0,t}(\Gamma_0) &:= \frac{1}{t}\mathbb{E}\left[\int_0^t I_{\Gamma_0}(X(s))\,ds\right], \quad \Gamma_0 \in \mathcal{B}(\mathcal{E}), \\ \nu_{1,t}(\Gamma_2) &:= \frac{1}{t}\mathbb{E}\left[\sum_{k=1}^{\infty} I_{\{\tau_k \le t\}} I_{\Gamma_2}(X(\tau_k-), Z_k)\right], \ \Gamma_2 \in \mathcal{B}(\overline{\mathcal{R}}). \end{aligned} \tag{3}$$

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.

$$J(\tau, Z) := \limsup_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\int_0^t c_0(X(s))\,ds + \sum_{k=1}^{\infty} I_{\{\tau_k\le t\}} c_1(X(\tau_k-), X(\tau_k))\right]$$

$$= \limsup_{t\to\infty} \int c_0(x)\mu_{0,t}(dx) + \int \widehat{c}_1(y,z)\nu_{1,t}(dy\times dz).$$

The Auxiliary Function U_0 and its Approximation U_n

Define

$$U_0(x) = g_0(x) - H_0^* \zeta(x), \quad x \in \mathcal{E}.$$

- $U_0 \in C(\mathcal{E}) \cap C^2(\mathcal{I})$ is a solution of the system

$$\begin{cases} Af(x) + c_0(x) - H_0^* = 0, & x \in \mathcal{I}, \\ \widehat{B}f(y,z) + \widehat{c}_1(y,z) \geq 0, & (y,z) \in \overline{\mathcal{R}} \\ f(x_0) = 0, & \\ \widehat{B}f(y_0^*, z_0^*) + \widehat{c}_1(y_0^*, z_0^*) = 0. & \end{cases}$$

The Auxiliary Function U_0 and its Approximation U_n

Define

$$U_0(x) = g_0(x) - H_0^* \zeta(x), \quad x \in \mathcal{E}.$$

- $U_0 \in C(\mathcal{E}) \cap C^2(\mathcal{I})$ is a solution of the system

$$\begin{cases} Af(x) + c_0(x) - H_0^* = 0, & x \in \mathcal{I}, \\ \widehat{B}f(y,z) + \widehat{c}_1(y,z) \geq 0, & (y,z) \in \overline{\mathcal{R}} \\ f(x_0) = 0, & \\ \widehat{B}f(y_0^*, z_0^*) + \widehat{c}_1(y_0^*, z_0^*) = 0. & \end{cases}$$

Define for $n \in \mathbb{N}$

$$U_n(x) = \frac{U_0(x)}{1 + \frac{1}{n} h(U_0(x))}, \quad x \in \mathcal{E},$$

where $h(x) = (-\frac{1}{8}x^4 + \frac{3}{4}x^2 + \frac{3}{8}) I_{\{|x| \leq 1\}} + |x| I_{\{|x|>1\}}$.

Key Observations

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

$$
\begin{aligned}
J(\tau, Y) &= \lim_{j\to\infty} \frac{1}{t_j} \mathbb{E}\left[\int_0^{t_j} c_0(X(s))\mathrm{d}s + \sum_{k=1}^{\infty} I_{\{\tau_k \le t_j\}} c_1(X(\tau_k-), X(\tau_k))\right] \\
&= \lim_{j\to\infty} \left(\int_{\overline{\mathcal{E}}} c_0(x)\, \mu_{0,t_j}(\mathrm{d}x) + \int_{\overline{\mathcal{R}}} \widehat{c}_1(y,z)\, \nu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z) \right) \\
&= \lim_{j\to\infty} \left(\int_{\overline{\mathcal{E}}} (AU_n(x) + c_0(x))\mu_0(\mathrm{d}x) \right. \\
&\qquad \left. + \int_{\overline{\mathcal{R}}} (\widehat{BU_n}(y,z) + \widehat{c}_1(y,z))\, \nu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z) \right), \ \forall n \in \mathbb{N},
\end{aligned}
$$

where $\mu_{0,t_j} \Rightarrow \mu_0$ as $j \to \infty$.

Key Observations (cont'd)

Because U_0 satisfies

$$AU_0 + c_0(x) - H_0^* = 0 \text{ and } \widehat{BU_0}(y,z) + \widehat{c}_1(y,z) \geq 0,$$

we can show by weak convergence that

$$\liminf_{n\to\infty} \liminf_{j\to\infty} \int_{\overline{\mathcal{R}}} (\widehat{BU_n}(y,z) + \widehat{c}_1(y,z))\, \nu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z) \geq 0,$$

and

$$\begin{aligned} \liminf_{n\to\infty} \int_{\overline{\mathcal{E}}} (AU_n(x) + c_0(x))\, \mu_0(\mathrm{d}x) &\geq \int_{\overline{\mathcal{E}}} (AU_0(x) + c_0(x))\, \mu_0(\mathrm{d}x) \\ &\geq H_0^*. \end{aligned}$$

Optimality

$$
\begin{aligned}
&J(\tau, Y)\\
&\quad = \liminf_{n\to\infty} \lim_{j\to\infty} \bigg(\int_{\overline{\mathcal{E}}} (AU_n(x) + c_0(x))\, \mu_{0,t_j}(\mathrm{d}x)\\
&\qquad\qquad + \int_{\overline{\mathcal{R}}} (\widehat{BU_n}(y,z) + \widehat{c}_1(y,z))\, \nu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z) \bigg)\\
&\quad \geq \liminf_{n\to\infty} \liminf_{j\to\infty} \int_{\overline{\mathcal{E}}} (AU_n(x) + c_0(x))\, \mu_{0,t_j}(\mathrm{d}x)\\
&\qquad\qquad + \liminf_{n\to\infty} \liminf_{j\to\infty} \int_{\overline{\mathcal{R}}} (\widehat{BU_n}(y,z) + \widehat{c}_1(y,z))\, \nu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z)\\
&\quad \geq \liminf_{n\to\infty} \int_{\overline{\mathcal{E}}} (AU_n(x) + c_0(x))\, \mu_0(\mathrm{d}x)\\
&\qquad\qquad + \liminf_{n\to\infty} \liminf_{j\to\infty} \int_{\overline{\mathcal{R}}} (\widehat{BU_n}(y,z) + \widehat{c}_1(y,z))\, \nu_{1,t_j}(\mathrm{d}y \times \mathrm{d}z)\\
&\quad \geq H_0^*.
\end{aligned}
$$

Optimality of the (s, S)-Policy

Theorem 4

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

$$J(\tau, Y) \geq H_0^*.$$

(b) *Moreover, the* (s, S)*-policy with* $s = y_0^*$ *and* $S = z_0^*$ *is an optimal impulse policy.*

Example: Logistic Storage Model

- Inventory level (in the absence of control)

$$dX_0(t) = -\mu X_0(t)(1 - X_0(t))\,dt + \sigma X_0(t)(1 - X_0(t))\,dW(t),$$

- For each $(y, z) \in \mathcal{R}$, $Q(\cdot; y, z)$ is the *'z-shifted uniform distributions'* on the interval $[(1 - (z/k)^j)y + (z/k)^j z, z]$ for some $j \in \mathbb{N}$.
- Assume

$$c_0(x) = k_0(x - \bar{x})^2 \text{ and } c_1(y, z) = k_1 + k_2(z - y),$$

for some $k_0, k_1, k_2 > 0$ and $\bar{x} \in (0, 1)$.

Numerical Results

Model 1: $\sigma = 0$ and $Q(\cdot; y, z) = \delta_{\{z\}}(\cdot)$;

Model 2: $\sigma = \frac{1}{10}$ and $Q(\cdot; y, z) = \delta_{\{z\}}(\cdot)$;

Model 3: $\sigma = \frac{1}{10}$ and $Q(\cdot; y, z) =$ the z-shifted uniform distribution;

Model	From	To	Mean Supply	Cost	Mean Cycle Length
Model 1	0.40567	0.59433	0.188661	0.938043	15.2759
Model 2	0.381724	0.56993	0.188206	1.00067	15.2779
Model 3	0.384973	0.6575	0.138321	1.33092	11.2843

Table 1: Comparison of Three Logisitic Inventory Models.

Optimal Harvesting Problem with Mean Field Interactions

- The growth of the forest:

$$dX_0(t) = \mu(X_0(t))\mathrm{d}t + \sigma(X_0(t))\mathrm{d}W(t), \quad X_0(0) \in (0, \infty).$$

- A particular forest owner's harvesting policy: $Q := \{(\tau_k, Y_k), k = 1, 2, \dots\}$. The resulting controlled forest:

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

where $Y_k := X(\tau_k-) - X(\tau_k) \geq 0$.

- The other agents in the market adopt the harvesting strategy $R := \{(\sigma_k, Z_k), k = 1, 2, \dots\}$ and the resulting average supply of log is

$$\kappa^R := \lim_{t\to\infty} \frac{1}{t}\mathbb{E}\left[\sum_{k=1}^{\infty} I_{\{\sigma_k \le t\}}(X(\sigma_k -) - X(\sigma_k))\right].$$

- Given $K > 0$ and a payoff function $\gamma : \mathcal{R} \times (0, \infty) \mapsto \mathbb{R}_+$, the expected reward for the forest owner is

$$J_x(Q, R) := \liminf_{t\to\infty} \frac{1}{t}\mathbb{E}_x\left[\sum_{k=1}^{\infty} I_{\{\tau_k \le t\}}(\gamma(X(\tau_k), X(\tau_k -), \kappa^R) - K)\right],$$

where $\mathcal{R} := \{(w, y) \in (0, \infty) \times (0, \infty) : w < y\}$.

Equilibrium

To find an admissible harvesting policy Q^* so that

$$J_x(Q, Q^*) \leq J_x(Q^*, Q^*), \quad \forall Q \text{ admissible.}$$

Assume

- Both the speed measure M and the scale function S of the process X_0 are absolutely continuous with respect to the Lebesgue measure.
- The left boundary 0 is a non-attracting point and the right boundary ∞ is a natural point. Moreover, $M[0, \infty) < \infty$.
- The drift function μ is continuously differentiable. Moreover, there exists a $y_1 > 0$ so that μ is strictly increasing on $(0, y_1]$ and strictly decreasing on $[y_1, \infty)$.
- The scale density s satisfies $\lim_{x \to \infty} s(x) = \infty$.
- $\gamma(w, y, z) = \varphi(z)(y - w)$ for a decreasing and strictly positive function φ.

Theorem 5

An equilibrium harvesting strategy Q^* exists. Moreover, Q^* is of (s,S)-type, which can be found by the fixed point of the mapping $\Phi : \mathcal{R} \to \mathcal{R}$ defined by $\Phi(w,y) := g \circ f(w,y)$, where

$$f(w,y) := \frac{y-w}{\xi(y)-\xi(w)}, \quad g(z) := \arg\max_{w<y} \frac{\varphi(z)(y-w)-K}{\xi(y)-\xi(w)}$$

and

$$\xi(x) = \int_{x_0}^{x} M[0,v]\mathrm{d}S(v).$$

Thank you!