Optimal Control of (McKean-Vlasov) Branching Diffusion

Julien Claisse with J. Kang, T. Lan and X. Tan

CEREMADE, Université Paris Dauphine–PSL

Séminaire FDD-FiME-MIRTE

16 octobre 2026

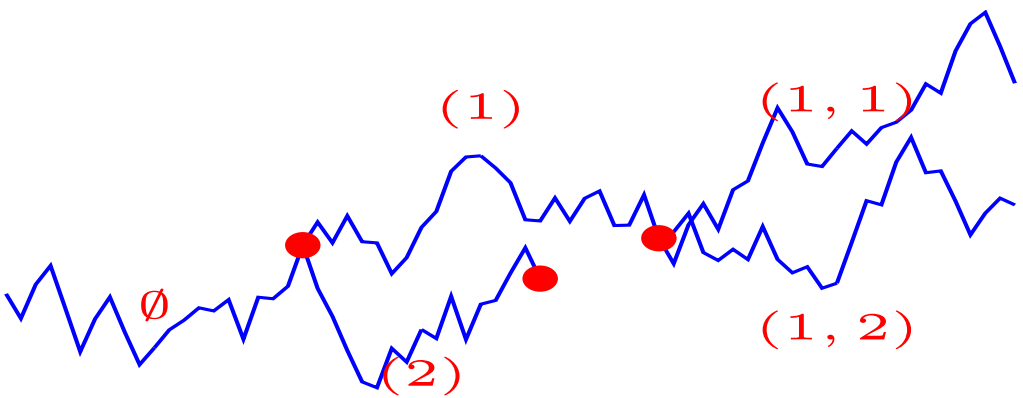
Outline

1. Branching Diffusion and PDE
   - Branching Diffusion
   - Nonlinear PDE
2. Control without Interaction
   - Finite Horizon Problem
   - Stopping Problems
3. Control with Interaction
   - Finite Population Setting
   - McKean-Vlasov Setting

Branching Diffusion

- Consider a population where each particle
  - moves according to a diffusion (b, σ)
  - dies at rate γ and gives birth to $\ell \in \mathbb{N}$ identical particles with probability p_ℓ
- Then each child particle follows the same but an independent dynamic as the mother particle $\longrightarrow$ *branching property*
- To describe the genealogy of the population we introduce the set of labels

$$\mathbb{K} := \{\emptyset\} \cup \bigcup_{n \geq 1} \mathbb{N}^n$$



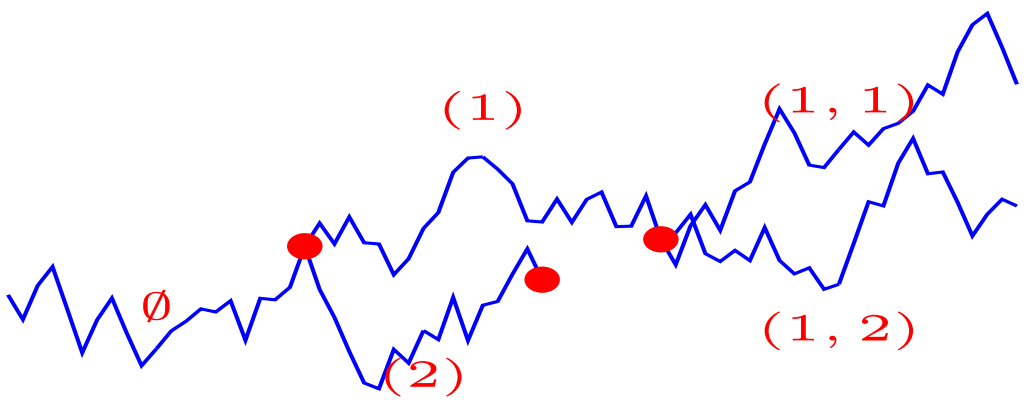
Branching Diffusion

- Let $(\Omega, (\mathcal{F}_s)_{s\geq 0}, \mathbb{P})$ be a filtered probability space
  - $(B^k)_{k\in\mathbb{K}}$ independent Brownian motions
  - $(Q^k)_{k\in\mathbb{K}}$ independent Poisson random measures
- Provided that the particle k is alive at time s
  - its position X^k is given by

$$dX_s^k = b(s, X_s^k)\, ds + \sigma(s, X_s^k)\, dB_s^k$$

  - its reproduction is driven by Q^k so that, given $\mathcal{F}_s$, it branches into $\ell \in \mathbb{N}$ particles in $[s, s+h]$ with probability

$$\gamma(s, X_s^k) p_\ell(s, X_s^k) h + o(h)$$



Nonlinear PDE

Nonlinear Feynman-Kac Formula (Skorohod '64)

The solution to the PDE

$$\begin{cases} \partial_t u + b \cdot \nabla_x u + \frac{1}{2} \operatorname{Tr}\left(\sigma\sigma^\top \nabla_x^2 u\right) + \gamma\Big(\sum_{\ell \geq 0} p_\ell u^\ell - u\Big) = 0 \\ u(T, \cdot) = g \end{cases}$$

admits the probabilistic representation

$$u(t,x) = \mathbb{E}\Big[\prod_{k \in K_T^{t,x}} g(X_T^k)\Big]$$

where $(X_T^k)_{k \in K_T^{t,x}}$ is a branching diffusion starting at time t from one particle at position x

Proof

- In the diffusion case, it holds by Markov property for τ stopping time in $[t, T]$

$$u(t,x) = \mathbb{E}\left[g\left(X_T^{t,x}\right)\right] = \mathbb{E}\left[u\left(\tau, X_\tau^{t,x}\right)\right]$$

Together with Itô formula, it gives

$$\mathbb{E}\left[\int_t^\tau \left(\partial_t u + b\cdot\nabla_x u + \frac{1}{2}\operatorname{Tr}\left(\sigma\sigma^\top \nabla_x^2 u\right)\right)(s, X_s^{t,x})\, dt\right] = 0$$

- The idea is the same but we combine Markov property with branching property to show that

$$u(t,x) = \mathbb{E}\Big[\prod_{k\in K_T^{t,x}} g(X_T^k)\Big] = \mathbb{E}\Big[\prod_{k\in K_\tau^{t,x}} u\left(\tau, X_\tau^k\right)\Big]$$

Indeed we have

$$\prod_{k\in K_T^{t,x}} g(X_T^k) = \prod_{k\in K_\tau^{t,x}} \psi^k \quad \text{with } \psi^k = \prod_{\substack{k' \text{ descendant of } k \\ \text{alive at time } T}} g(X_T^{k'})$$

and

$$\mathbb{E}\Big[\prod_{k\in K_\tau^{t,x}} \psi^k \,\Big|\, \mathcal{F}_\tau\Big] = \prod_{k\in K_\tau^{t,x}} \mathbb{E}\left[\psi^k \,|\, \mathcal{F}_\tau\right] = \prod_{k\in K_\tau^{t,x}} u\left(\tau, X_\tau^k\right)$$

Extensions

- Zero-Order Nonlinearity: McKean '75, Henry-Labordère '12, Henry-Labordère, Tan & Touzi '13, Bouchard, Tan, Warin & Zhou '16
- First-Order Nonlinearity: Henry-Labordère, Oudjane, Tan, Touzi & Warin '16, Bouchard, Tan & Warin '17
- Fully Nonlinear: Still open problem
- Beyond Parabolic PDE
  - Elliptic: Watanabe '64, Bossy, Champagnat, Leman, Maire, Violeau & Yvinec '13, Agarwal and Claisse '17
  - Hyperbolic: Kac '74, Dalang, Mueller & Tribe '07, Henry-Labordère & Touzi '18
  - Fractional: Privault and Penent '21, '23,

Outline

1. Branching Diffusion and PDE
   - Branching Diffusion
   - Nonlinear PDE
2. Control without Interaction
   - Finite Horizon Problem
   - Stopping Problems
3. Control with Interaction
   - Finite Population Setting
   - McKean-Vlasov Setting

Controlled Branching Diffusion

- Consider a controlled population where each particle
  - moves according to a diffusion $(b(t,x,a),\sigma(t,x,a))$
  - dies at rate $\gamma(t,x,a)$ and gives birth to $\ell \in \mathbb{N}$ identical particles with probability $p_\ell(t,x,a)$
- The class of admissible controls is

$$\mathcal{A} := \left\{(\alpha^k)_{k\in\mathbb{K}} : \ \alpha_k \text{ predictable valued in } A \text{ for all } k \in \mathbb{K}\right\}$$

- Provided that the particle k is alive at time s
  - its position X^k is given by

$$dX_s^k = b(s, X_s^k, \alpha_s^k)\, ds + \sigma(s, X_s^k, \alpha_s^k)\, dB_s^k$$

  - its reproduction is driven by Q^k so that, given $\mathcal{F}_s$, it branches into $\ell \in \mathbb{N}$ particles in $[s, s+h]$ with probability

$$\gamma(s, X_s^k, \alpha_s^k) p_\ell(s, X_s^k, \alpha_s^k) h + o(h)$$

Finite Horizon Problem

Consider the following optimal control problem

$$v(t,x) = \inf_{\alpha \in \mathcal{A}} \mathbb{E}\Big[\prod_{k \in K_T^{t,x,\alpha}} g(X_T^k) \Big]$$

- If $g = 0$, the goal is to minimize the probability of extinction before time T
- If $g(x) = e^{-\eta x}$ with $\eta > 0$, the goal is to maximize the expected CARA utility of the overall wealth of the population $\sum_{k \in K_T} X_T^k$

HJB Equation (Nisio '85, C. '15)

The value function $v : [0,T] \times \mathbb{R}^d \to \mathbb{R}$ is the unique viscosity solution to

$$\begin{cases} \partial_t u + \inf_{a \in A} \left\{ b^a \cdot \nabla_x u + \frac{1}{2} \operatorname{Tr}\left(\sigma^a \sigma^{a\top} \nabla_x^2 u \right) + \gamma^a \Big(\sum_{\ell \geq 0} p_\ell^a u^\ell - u \Big) \right\} = 0 \\ u(T, \cdot) = g \end{cases}$$

Idea of Proof

- This result follows by standard arguments from a dynamic programming principle of the form

$$v(t,x) = \inf_{\alpha \in \mathcal{A}} \mathbb{E}\Big[\prod_{k \in K_\tau^{t,x,\alpha}} v(\tau, X_\tau^k) \Big]$$

  BUT the branching property does not hold for controlled branching diffusion
- The key idea is that the HJB equation entails that one can restrict to feedback controls $\alpha_s^k = \alpha(s, X_s^k)$ which preserve the branching property
- We can adapt the ideas of proof from *Fleming & Soner '93*
  - Assume that there exists a smooth solution to the HJB equation and use a verification argument
  - Regularize the coefficients of the problem to ensure existence of a smooth solution to HJB
  - Use the stability of viscosity solution to pass to the limit

Optimal Stopping Problem

Optimal Stopping (Kharroubi & Ocello '24)

The value function $v : \mathbb{R}^d \to \mathbb{R}$ defined as

$$v(x) = \inf_{(\tau^k)_k \text{ stopping lines}} \mathbb{E}\Big[\prod_{k \text{ stopped before branching}} g(X^k_{\tau^k}) \Big]$$

is the unique viscosity solution to

$$\min\Big\{ b \cdot \nabla_x u + \frac{1}{2} \operatorname{Tr}\left(\sigma\sigma^\top \nabla_x^2 u\right) + \gamma\Big(\sum_{\ell \geq 0} p_\ell u^\ell - u \Big), g - u \Big\} = 0$$

Optimal Branching Problem

Optimal Branching (Guo, Lu & Zhang '24)

The value function $v : [0,T] \times \mathbb{R}^d \to \mathbb{R}$ defined as

$$v(t,x) = \inf_{(\tau^k)_k \text{ branching times}} \mathbb{E}\Big[\exp\Big(-\int_t^T \sum_{k \in K_s^{t,x}} c(X_s^k)\, ds\Big)\Big] \quad \text{with } c > 0$$

is the unique viscosity solution to

$$\begin{cases} \min\Big\{\partial_t u + b \cdot \nabla_x u + \frac{1}{2}\operatorname{Tr}\big(\sigma\sigma^\top \nabla_x^2 u\big) - cu, \sum_{\ell \geq 0} p_\ell u^\ell - u\Big\} = 0 \\ u(T,\cdot) = \text{ probability of extinction} \end{cases}$$

Outline

1. Branching Diffusion and PDE
   - Branching Diffusion
   - Nonlinear PDE
2. Control without Interaction
   - Finite Horizon Problem
   - Stopping Problems
3. Control with Interaction
   - Finite Population Setting
   - McKean-Vlasov Setting

Branching Diffusion with Interaction

- Consider a controlled population with mean-field interaction where each particle
  - moves according to a diffusion $(b(t,x,m,a),\sigma(t,x,m,a))$
  - dies at rate $\gamma(t,x,m,a)$ and gives birth to $\ell \in \mathbb{N}$ identical particles with probability $p_\ell(t,x,m,a)$
- The mean-field interaction term corresponds to the empirical measure of the population given by

$$Z_s = \sum_{k \in K_s} \delta_{X_s^k}$$

- Provided that the particle k is alive at time s
  - its position X^k is given by

$$dX_s^k = b(s, X_s^k, Z_s, \alpha_s^k)\, ds + \sigma(s, X_s^k, Z_s, \alpha_s^k)\, dB_s^k$$

  - its reproduction is driven by Q^k so that, given $\mathcal{F}_s$, it branches into $\ell \in \mathbb{N}$ particles in $[s, s+h]$ with probability

$$\gamma(s, X_s^k, Z_s, \alpha_s^k) p_\ell(s, X_s^k, Z_s, \alpha_s^k) h + o(h)$$

Finite Horizon Problem

Consider the following optimal control problem

$$v_n(t,\mathbf{x}) = \inf_{\alpha\in\mathcal{A}} \mathbb{E}\Big[G\Big(\sum_{k\in K_T^{t,\mathbf{x},\alpha}} \delta_{X_T^k}\Big)\Big], \quad \text{with } \mathbf{x} = (\mathbf{x}_1,..,\mathbf{x}_n) \in (\mathbb{R}^d)^n$$

Taking $G(m) = \exp(\langle m, \log(g)\rangle)$, we recover the multiplicative cost $\prod_{k\in K_T} g(X_T^k)$

HJB Equation (Ocello '23)

The family of value functions $v_n : [0,T]\times(\mathbb{R}^d)^n \to \mathbb{R}$ is the unique viscosity solution to the system of PDE

$$\begin{cases} \partial_t u_n(t,\mathbf{x}) + \inf_{\mathbf{a}\in A^n} \left\{\mathcal{L}_n^{\mathbf{a}} u(t,\mathbf{x})\right\} = 0, \quad n \geq 1 \\ u_n(T,\mathbf{x}) = G(\mu[\mathbf{x}]) \end{cases}$$

where $\mu[\mathbf{x}] = \sum_{i=1}^n \delta_{\mathbf{x}_i}$ and

$$\mathcal{L}_n^{\mathbf{a}} u(t,\mathbf{x}) = \boldsymbol{b}^{\boldsymbol{a}}(t,\mathbf{x})\cdot\nabla_{\mathbf{x}} u_n(t,\mathbf{x}) + \frac{1}{2}\operatorname{Tr}\left(\boldsymbol{\sigma}^{\boldsymbol{a}}\boldsymbol{\sigma}^{\boldsymbol{a}\,\top}(t,\mathbf{x})\nabla_{\mathbf{x}}^2 u_n(t,\mathbf{x})\right)$$

$$+\sum_{i=1}^n \gamma^{\mathbf{a}_i}(t,\mathbf{x}_i,\mu[\mathbf{x}])\Big(\sum_{\ell\geq 0} p_\ell^{\mathbf{a}_i}(t,\mathbf{x}_i,\mu[\mathbf{x}]) u_{n+\ell-1}(t,\mathbf{x}_1,..,\underbrace{\mathbf{x}_i,..,\mathbf{x}_i}_{\ell-1 \text{ times}},..,\mathbf{x}_n) - u_n(t,\mathbf{x})\Big)$$

where $\boldsymbol{b}^{\boldsymbol{a}}(t,\mathbf{x}) = (b^{\mathbf{a}_i}(t,\mathbf{x}_i,\mu[\mathbf{x}]))_{1\leq i\leq n}$ and $\boldsymbol{\sigma}^{\boldsymbol{a}}(t,\mathbf{x}) = \operatorname{diag}(\sigma^{\mathbf{a}_i}(t,\mathbf{x}_i,\mu[\mathbf{x}]))_{1\leq i\leq n}$.

Outline

1. Branching Diffusion and PDE
   - Branching Diffusion
   - Nonlinear PDE
2. Control without Interaction
   - Finite Horizon Problem
   - Stopping Problems
3. Control with Interaction
   - Finite Population Setting
   - McKean-Vlasov Setting

The Limit Case

- The aim is to study the limit in large population of the control problem with interaction
- Letting the initial size of the population $n \to \infty$, we expect that the interaction term converges as follows

$$\mu_t^n := \frac{1}{n} \sum_{i=1}^{n} \sum_{k \in K_t^i} \delta_{X_t^k} \longrightarrow \mathbb{E}\Big[\sum_{k \in K_t} \delta_{X_t^k} \Big] =: \mu_t$$

in the sense that for any $\varphi : \mathbb{R}^d \mapsto \mathbb{R}$,

$$\langle \mu_t^n, \varphi \rangle = \frac{1}{n} \sum_{i=1}^{n} \sum_{k \in K_t^i} \varphi(X_t^k) \longrightarrow \mathbb{E}\Big[\sum_{k \in K_t} \varphi(X_t^k) \Big] = \langle \mu_t, \varphi \rangle$$

where $\sum_{k \in K_t^i} \delta_{X_t^k}$ is the population stemming from the i-th particle

- Another point of view is to study the impact of population size variations in mean-field control problems

McKean-Vlasov Branching Diffusion

- Consider a controlled population with MKV interaction, *i.e.*, the interaction term $\mu_s \in \mathcal{M}(\mathbb{R}^d)$ corresponds to the population measure given by

$$\mu_s = \mathbb{E}\Big[\sum_{k\in K_s} \delta_{X_s^k}\Big] \quad \text{i.e.} \quad \langle \mu_s, \varphi\rangle = \mathbb{E}\Big[\sum_{k\in K_s} \varphi(X_s^k)\Big], \ \text{for } \varphi \in C_b(\mathbb{R}^d)$$

- Provided that the particle k is alive at time s
  - its position X^k is given by

$$dX_s^k = b(s, X_s^k, \mu_s, \alpha_s^k)\, ds + \sigma(s, X_s^k, \mu_s, \alpha_s^k)\, dB_s^k$$

  - its reproduction is driven by Q^k so that, given $\mathcal{F}_s$, it branches into $\ell \in \mathbb{N}$ particles in $[s, s+h]$ with probability

$$\gamma(s, X_s^k, \mu_s, \alpha_s^k) p_\ell(s, X_s^k, \mu_s, \alpha_s^k) h + o(h)$$

- The class of admissible controls $\mathcal{A}$ is restricted to closed–loop strategies $(\alpha^k)_{k\in\mathbb{K}}$ of the form

$$\alpha_s^k = \alpha(s, X_s^k, \mu_s)$$

McKean-Vlasov Branching Diffusion

Existence and Uniqueness (C., Kang and Tan '24)

Assuming that all the coefficients are Lipschitz in (x, m) with γ and $\sum_{\ell \geq 0} \ell\, p_\ell$ bounded. Then the MKV branching diffusion starting at time $t \geq 0$ from (random) initial state ξ and controlled by $\alpha \in \mathcal{A}$ is uniquely determined as a measure-valued càdlàg process

$$Z_s^{t,\xi,\alpha} = \sum_{k \in K_s^{t,\xi,\alpha}} \delta_{X_s^k}$$

where $K_s^{t,\xi,\alpha}$ contains the labels of particles alive at time s and the induced population measure $\mu_s^{t,\xi,\alpha} \in \mathcal{M}(\mathbb{R}^d)$ is given by

$$\mu_s^{t,\xi,\alpha} = \mathbb{E}\big[Z_s^{t,\xi,\alpha}\big] \quad \text{i.e.} \quad \langle \mu_s, \varphi \rangle = \mathbb{E}\Big[\sum_{k \in K_s} \varphi(X_s^k)\Big], \ \text{for } \varphi \in C_b(\mathbb{R}^d)$$

Optimal Control Problem

- Inspired by *Pham and Wei ' 17*, we aim to minimize

$$\mathbb{E}\Big[\int_t^T \sum_{k\in K_s^{t,\xi,\alpha}} L\big(s, X_s^k, \mu_s^{t,\xi,\alpha}, \alpha_s^k\big)ds + \sum_{k\in K_T^{t,\xi,\alpha}} g(X_T^k, \mu_T^{t,\xi,\alpha})\Big]$$

  or equivalently

$$\int_t^T \langle \mu_s^{t,\xi,\alpha}, L^\alpha(s,\cdot,\mu_s^{t,\xi,\alpha})\rangle ds + \langle \mu_T^{t,\xi,\alpha}, g(\cdot,\mu_T^{t,\xi,\alpha})\rangle$$

  with $L^\alpha(s,x,m) = L(s,x,\mu,\alpha(s,x,m))$

- **Key observation:** This problem depends on the *deterministic* flow of measures $(\mu_s^{t,\xi,\alpha})_{s\in[t,T]}$ rather than the stochastic process $(Z_s^{t,\xi,\alpha})_{s\in[t,T]}$
- **Main difficulty:** We need to prove an *invariance principle* of the form

$$\mu_s^{t,\xi_1,\alpha} = \mu_s^{t,\xi_2,\alpha}, \quad \text{for } \xi_1,\xi_2 \text{ with identical measure } \mathbb{E}[\xi_1] = \mathbb{E}[\xi_2]$$

Invariance Principle

Invariance Principle

It holds for all $s \in [t, T]$,

$$\mu_s^{t,\xi_1,\alpha} = \mu_s^{t,\xi_2,\alpha}, \quad \text{for } \xi_1, \xi_2 \text{ with identical measure } \mathbb{E}[\xi_1] = \mathbb{E}[\xi_2]$$

The result comes from uniqueness to the nonlinear Fokker-Planck equation

$$\begin{cases} \partial_t \mu + \partial_x \left(b^\alpha(s, x, \mu_s) \mu \right) - \frac{1}{2} \partial_{xx}^2 \left(\sigma^\alpha(s, x, \mu_s)^2 \mu \right) + \kappa^\alpha(s, x, \mu_s) \mu = 0 \\ \mu_t = m \end{cases}$$

that is satisfied by both $\mu^{t,\xi_1,\alpha}$ and $\mu^{t,\xi_2,\alpha}$ with $\mathbb{E}[\xi_1] = \mathbb{E}[\xi_2] =: m \in \mathcal{M}(\mathbb{R}^d)$ and

$$\kappa^\alpha := \gamma^\alpha \sum_{\ell \geq 0} (\ell - 1) p_\ell^\alpha$$

Dynamic Programming Principle

Let us introduce the the value function $v : [0,T] \times \mathcal{M}(\mathbb{R}^d) \to \mathbb{R}$ as

$$v(t,m) = \inf_{\alpha \in \mathcal{A}} \left\{ \int_t^T \langle \mu_s^{t,m,\alpha}, L^\alpha(s,\cdot,\mu_s^{t,m,\alpha}) \rangle ds + \langle \mu_T^{t,m,\alpha}, g(\cdot,\mu_T^{t,m,\alpha}) \rangle \right\}$$

where

$$\mu_s^{t,m,\alpha} := \mu_s^{t,\xi,\alpha}, \qquad \text{for any } \xi \text{ with measure } \mathbb{E}[\xi] = m$$

Dynamic Programming Principle

It holds for any $t \le s \le T$,

$$v(t,m) = \inf_{\alpha \in \mathcal{A}} \left\{ \int_t^s \langle \mu_\theta^{t,m,\alpha}, L^\alpha(\theta,\cdot,\mu_\theta^{t,m,\alpha}) \rangle d\theta + v(s,\mu_s^{t,m,\alpha}) \right\}$$

Differential Calculus on $\mathcal{M}(\mathbb{R}^d)$

Let $F : \mathcal{M}(\mathbb{R}^d) \to \mathbb{R}$, the linear derivative $\delta_m F : \mathbb{R}^d \times \mathcal{M}(\mathbb{R}^d) \to \mathbb{R}$ is defined as

$$\frac{F(m + \varepsilon(m' - m)) - F(m)}{\varepsilon} \xrightarrow[\varepsilon \to 0^+]{} \int_{\mathbb{R}^d} \delta_m F(x, m)\,(m' - m)(dx)$$

Itô Formula (Cao, Ren and Tan ' 26)

If $\delta_m F(\cdot, m)$ is of class C^2, it holds for $s \geq t$,

$$F(\mu_s^{t,m,\alpha}) = F(m) + \int_t^s \langle \mu_\theta^{t,m,\alpha}, \mathcal{G}_\theta^\alpha F(\cdot, \mu_\theta^{t,m,\alpha}) \rangle \, d\theta$$

where $(\mu_s)_{s \geq 0}$ is the population distribution and

$$\mathcal{G}_s^\alpha F(x, m) := \gamma^\alpha(s, x, m) \sum_{\ell \geq 0} (\ell - 1) p_\ell^\alpha(s, x, m) \delta_m F(x, m)$$
$$+ b^\alpha(s, x, m) \cdot \nabla_x \delta_m F(x, m) + \frac{1}{2} \operatorname{Tr}\left(\sigma^\alpha \sigma^{\alpha\top}(s, x, m) \nabla_x^2 \delta_m F(x, m)\right)$$

Hamilton-Jacobi-Bellman Equation

HJB Equation: Smooth Case

If the value function v is smooth, then it satisfies

$$\partial_t v(t,m) + \inf_{\alpha \in \mathcal{A}} \left\{ \langle m, \mathcal{G}_t^\alpha v(t,\cdot,m) + L^\alpha(t,\cdot,m) \rangle \right\} = 0$$

Conversely, if there exist a smooth function u satisfying

$$\begin{cases} \partial_t u(t,m) + \inf_{\alpha \in \mathcal{A}} \left\{ \langle m, \mathcal{G}_t^\alpha u(t,\cdot,m) + L^\alpha(t,\cdot,m) \rangle \right\} = 0 \\ u(T,m) = \langle m, g(\cdot,m) \rangle \end{cases}$$

and a control $\hat{\alpha} \in \mathcal{A}$ such that

$$\hat{\alpha}(t,\cdot,m) = \arg\min_{\alpha \in \mathcal{A}} \left\{ \langle m, \mathcal{G}_t^\alpha u(t,\cdot,m) + L^\alpha(t,\cdot,m) \rangle \right\}$$

then it holds

$$u(t,m) = v(t,m) = J(t,m,\hat{\alpha})$$

Linear Quadratic Example

Consider the one-dimensional setting with

- linear drift

$$b(t,x,m,a) = b_1(t)\,x + b_2(t)\,m(\mathbb{R}) + b_3(t)\,a$$

- σ, γ, $(p_\ell)_{\ell\in\mathbb{N}}$ constants
- quadratic costs

$$L(t,x,m,a) = L_1(t)\,x^2 + L_2(t)\,m(\mathbb{R}) + L_3(t)\,a^2$$
$$g(x,m) = g_1 x^2 + g_2\,m(\mathbb{R})$$

Then the value function can be computed in the form

$$v(t,m) = v_1(t)\,m_2 + v_2(t)\,m_1 m(\mathbb{R}) + r(t,m(\mathbb{R}))$$

and the optimal control as

$$\hat{\alpha}(t,x,m) = -\frac{b_3(t)}{2L_3(t)}\nabla_x \delta_m v(t,x,m) = a_1(t)\,x + a_2(t)\,m(\mathbb{R})$$

where

$$m_1 := \int x\,m(dx), \quad m_2 := \int x^2\,m(dx)$$

Hamilton-Jacobi-Bellman Equation

HJB Equation: Viscosity Case (Ekren, He, Lan & Tan '26)

The value function $v : [0,T] \times \mathcal{M}(\mathbb{R}^d) \to \mathbb{R}$ is the unique viscosity solution in an appropriate sense to

$$\begin{cases} \partial_t u(t,m) + \inf_{\alpha \in \mathcal{A}} \left\{ \langle m, \mathcal{G}_t^\alpha u(t,\cdot,m) + L^\alpha(t,\cdot,m) \rangle \right\} = 0 \\ u(T,m) = \langle m, g(\cdot,m) \rangle \end{cases}$$

Thank you for your attention!