Deep Signature Approach for McKean-Vlasov FBSDEs in a Random Environment

Mathieu Laurière

NYU Shanghai

joint work with:
Ruimeng Hu *(UCSB), Botao* Jin *(UCSB), Jiacheng* Zhang *(CUHK)*

Advances in Financial Mathematics 2026
Paris, January 28, 2026

Outline

Outline

Context: Mean-Field Games with Common Noise

- **Mean-Field Game (MFG) Theory** [Lasry and Lions, 2007, Huang et al., 2006]
  - Framework for strategic interactions among large numbers of rational agents.
  - Individual influence is negligible.
  - Interactions through collective distribution
  - Applications: Finance, economics, engineering, social sciences, . . .

Context: Mean-Field Games with Common Noise

- **Mean-Field Game (MFG) Theory** [Lasry and Lions, 2007, Huang et al., 2006]
  - Framework for strategic interactions among large numbers of rational agents.
  - Individual influence is negligible.
  - Interactions through collective distribution
  - Applications: Finance, economics, engineering, social sciences, . . .

- **The Role of Common Noise**
  - Represents systemic shocks or shared uncertainties.
  - *Examples:*
    - Growth theory [Lasry et al., 2008]
    - Macroeconomic fluctuations [Vu and Ichiba, 2025].
    - Systemic risk [Carmona et al., 2015]
    - Bank runs [Carmona et al., 2017].
    - Energy transition [Dumitrescu et al., 2024]
    - Climate variability [Lavigne and Tankov, 2023].
    - Crowd motion in uncertain environments [Achdou and Lasry, 2018].

State of the Art and Challenges

- **MFG *without* Common Noise**: Solutions characterized by
  - **Hamilton-Jacobi-Bellman – Fokker-Planck (HJB-FP)** system [Lasry and Lions, 2007, Huang et al., 2006].
  - or **McKean-Vlasov (MKV) forward-backward stochastic differential equations (FBSDE)** [Carmona and Delarue, 2013].

- **Challenges with Common Noise**
  - Distribution of states is stochastic.
  - Stochastic HJB-FP / Master equation [Cardaliaguet et al., 2019].
  - MV-FBSDE in random environment [Carmona et al., 2016].
  - Existing literature is primarily theoretical.

For more details on MFGs, see e.g. [Bensoussan et al., 2013, Carmona and Delarue, 2018a, Carmona and Delarue, 2018b, Achdou et al., 2021]

Existing Numerical Approaches

Relatively broad literature without common noise [. . .]. With common noise:

1. PDE & Finite Difference Schemes

- [Achdou and Lasry, 2018]: Studied crowd motion with common noise.
- Common noise taking a **finite number of values** at discrete times.

2. Deep Learning (DL) & Reinforcement Learning (RL)

- [Carmona and Laurière, 2022]: similar discrete models using DL.
- [Perrin et al., 2020, Wu et al., 2024]: RL with pop.-dependent policies.
- [Gu et al., 2024]: Master Equations for specific macro-economic models.

3. Specific Interaction Structures

- [Min and Hu, 2021]: Interaction via **conditional moments** only.
- [Gomes et al., 2023]: Interaction through a **1D stochastic price**.

Goal of this work:

- A tractable algo. for coef. depending on the **full conditional distribution** in nonlinear ways, conditioned on **common noise**

Probabilistic Characterization

- We consider a **generic McKean–Vlasov FBSDE** in a **random environment**:

$$\begin{cases} \mathrm{d}X_t = B(t, \Theta_t, Z_t^0, \mathcal{L}(\Theta_t|\mathcal{F}_t^0))\,\mathrm{d}t + \Sigma(\dots)\,\mathrm{d}W_t + \Sigma^0(\dots)\,\mathrm{d}W_t^0, \\ \mathrm{d}Y_t = -H(t, \Theta_t, Z_t^0, \mathcal{L}(\Theta_t|\mathcal{F}_t^0))\,\mathrm{d}t + Z_t\,\mathrm{d}W_t + Z_t^0\,\mathrm{d}W_t^0, \\ X_0 \sim \mu_0, \qquad Y_T = G(X_T, \mathcal{L}(X_T|\mathcal{F}_T^0)), \\ \Theta_t = (X_t, Y_t, Z_t). \end{cases} \tag{1}$$

- **Key Components:**
  - $\mathcal{L}(\cdot|\mathcal{F}^0)$: Marginal law conditional on common noise.

Probabilistic Characterization

- We consider a **generic McKean–Vlasov FBSDE** in a **random environment**:

$$\begin{cases} \mathrm{d}X_t = B(t, \Theta_t, Z_t^0, \mathcal{L}(\Theta_t|\mathcal{F}_t^0))\,\mathrm{d}t + \Sigma(\dots)\,\mathrm{d}W_t + \Sigma^0(\dots)\,\mathrm{d}W_t^0, \\ \mathrm{d}Y_t = -H(t, \Theta_t, Z_t^0, \mathcal{L}(\Theta_t|\mathcal{F}_t^0))\,\mathrm{d}t + Z_t\,\mathrm{d}W_t + Z_t^0\,\mathrm{d}W_t^0, \\ X_0 \sim \mu_0, \qquad Y_T = G(X_T, \mathcal{L}(X_T|\mathcal{F}_T^0)), \\ \Theta_t = (X_t, Y_t, Z_t). \end{cases} \tag{1}$$

- **Key Components:**
  - $\mathcal{L}(\cdot|\mathcal{F}^0)$: Marginal law conditional on common noise.
  - Coefficients depend on distribution via an embedding m.
  - Allows handling general nonlinear distributional dependence beyond moments.

Main Contribution

- **Methodology: Fixed point iterations**
  - Players react to the last iterate.
  - Encode conditional distribution via signature [Lyons and Qian, 2002] .
  - Update embedding representation via supervised learning.
  - Update controls by DeepBSDE solver [Han et al., 2018].

- **Methodology: Fixed point iterations**
  - Players react to the last iterate.
  - Encode conditional distribution via signature [Lyons and Qian, 2002] .
  - Update embedding representation via supervised learning.
  - Update controls by DeepBSDE solver [Han et al., 2018].
- **Advantages**
  - Handles **high-dimensional** problems efficiently.
  - Goes **beyond conditional-moment** structures [Min and Hu, 2021].
  - Does not require solving Master equation on proba. space [Gu et al., 2024].

Main Results

- **Theoretical Convergence**
  - Proof of convergence for the iterative scheme.
  - Up to a supervised learning error.
- **Numerical Experiments**
  - Supervised learning step validation.
  - MV-FBSDE with an explicit analytic benchmark.
  - Flocking Model with Common Noise:
    - Nonlinear distribution-dependent benchmark.
    - Demonstrates accuracy and performance.

Outline

Outline

General Notation

- **Spaces and Measures**
  - $\mathcal{P}^2(\mathbb{R}^n)$: Probability measures on $\mathbb{R}^n$ with finite second moment.
  - $\mathcal{W}_2$: 2-Wasserstein distance defined by:

$$\mathcal{W}_2^2(\mu, \nu) := \inf_{\pi \in \Pi(\mu,\nu)} \int_{\mathbb{R}^n \times \mathbb{R}^n} |x - y|^2 \pi(dx, dy)$$

- **Laws and Filtrations**
  - $\mathcal{L}(X)$: Law of random variable X.
  - $\mathcal{L}(X|\mathcal{G})$: Conditional law given σ-field $\mathcal{G}$.
  - W, W^0: Independent q-dimensional Brownian motions.
  - $\mathcal{F}, \mathcal{F}^0$: Their respective natural filtrations.

- **Dimensions**
  - d: Forward/Backward dimension; q: Noise dimension.
  - ℓ: Dimension of the measure embedding.

Definition of the Problem

We consider the FBSDE system with state $\Theta_t = (X_t, Y_t, Z_t) \in \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{R}^{d\times q}$.

$$\begin{cases} \mathrm{d}X_t = B(t, \Theta_t, Z_t^0, \mathcal{L}(\Theta_t|\mathcal{F}_t^0))\,\mathrm{d}t + \Sigma(t, \Theta_t, \mathcal{L}(\Theta_t|\mathcal{F}_t^0))\,\mathrm{d}W_t \\ \qquad\qquad + \Sigma^0(t, \Theta_t, \mathcal{L}(\Theta_t|\mathcal{F}_t^0))\,\mathrm{d}W_t^0, \\ \mathrm{d}Y_t = -H(t, \Theta_t, Z_t^0, \mathcal{L}(\Theta_t|\mathcal{F}_t^0))\,\mathrm{d}t + Z_t\,\mathrm{d}W_t + Z_t^0\,\mathrm{d}W_t^0, \\ X_0 \sim \mu_0, \qquad Y_T = G(X_T, \mathcal{L}(X_T|\mathcal{F}_T^0)), \qquad \Theta_t = (X_t, Y_t, Z_t). \end{cases} \tag{2}$$

Intuitively:

- Consider an MFG (or MFC with common noise)
- Apply stochastic Pontryagin's maximum principle
- Y represents the derivative of the value function

Note: We could also consider Y = value function (Bellman principle).

Assumption: Measure Dependence

Assumption A

Embedding Dependence: Coefficients depend on an ℓ-dimensional embedding.

$$
\begin{aligned}
B(t,x,y,z,z^0,\nu) &= b(t,x,y,z,z^0,m_1(t,x,y,z,\nu)),\\
\Sigma(t,x,y,z,\nu) &= \sigma(t,x,y,z,m_2(t,x,y,z,\nu)),\\
\Sigma^0(t,x,y,z,\nu) &= \sigma^0(t,x,y,z,m_3(t,x,y,z,\nu)),\\
H(t,x,y,z,z^0,\nu) &= h(t,x,y,z,z^0,m_4(t,x,y,z,\nu)),\\
G(x,\mu) &= g(x,m_5(x,\mu)).
\end{aligned} \tag{3}
$$

Based on Assumption A, we rewrite the FBSDE using the embeddings m_i:

$$\begin{cases} \mathrm{d}X_t = b(t, \Theta_t, Z_t^0, m_1(t, \Theta_t, \mathcal{L}(\Theta_t|\mathcal{F}_t^0)))\,\mathrm{d}t \\ \qquad\qquad +\sigma(t, \Theta_t, m_2(t, \Theta_t, \mathcal{L}(\Theta_t|\mathcal{F}_t^0)))\,\mathrm{d}W_t \\ \qquad\qquad +\sigma^0(t, \Theta_t, m_3(t, \Theta_t, \mathcal{L}(\Theta_t|\mathcal{F}_t^0)))\,\mathrm{d}W_t^0, \\ \mathrm{d}Y_t = -h(t, \Theta_t, Z_t^0, m_4(t, \Theta_t, \mathcal{L}(\Theta_t|\mathcal{F}_t^0)))\,\mathrm{d}t + Z_t\,\mathrm{d}W_t + Z_t^0\,\mathrm{d}W_t^0, \\ X_0 \sim \mu_0, \qquad Y_T = g(X_T, m_5(X_T, \mathcal{L}(X_T|\mathcal{F}_T^0))), \\ \Theta_t = (X_t, Y_t, Z_t), \quad t \in [0, T]. \end{cases} \tag{4}$$

Outline

- **Challenge:** The embeddings m_i depend on ∞-dimensional proba. distributions.
- **Solution:** Replace laws with finite-dim. approximations using **Path Signatures**
- See [Lyons and Qian, 2002, Lyons et al., 2007].
- **Path Space** $\mathcal{V}^p([0,T],\mathbb{R}^d)$
  - Continuous mappings with finite p-variation; Norm: $\|\cdot\|_{\mathcal{V}^p} := \|\cdot\|_\infty + \|\cdot\|_p$.
  - p-variation defined by partitions $D = \{0 \le t_0 < \cdots < t_r \le T\}$:

$$\|X\|_p := \left(\sup_{D \subset [0,T]} \sum_{i=0}^{r-1} \left\| X_{t_{i+1}} - X_{t_i} \right\|^p \right)^{1/p}.$$

Definition 1 (Signature)

Let $X \in \mathcal{V}^p([0,T],\mathbb{R}^d)$. The signature $\mathcal{S}(X)$ is the sequence of iterated integrals:

$$\mathcal{S}(X) = (1, X^1, \cdots, X^k, \cdots) \in \oplus_{k=0}^{\infty}(\mathbb{R}^d)^{\otimes k},$$

where the k-th level is defined as: $X^k := \int_{0<t_1<\cdots<t_k<T} \mathrm{d}X_{t_1} \otimes \cdots \otimes \mathrm{d}X_{t_k}$.

Motivation and Definition

- **Challenge:** The embeddings m_i depend on ∞-dimensional proba. distributions.
- **Solution:** Replace laws with finite-dim. approximations using **Path Signatures**
- See [Lyons and Qian, 2002, Lyons et al., 2007].
- **Path Space** $\mathcal{V}^p([0,T],\mathbb{R}^d)$
  - Continuous mappings with finite p-variation; Norm: $\|\cdot\|_{\mathcal{V}^p} := \|\cdot\|_\infty + \|\cdot\|_p$.
  - p-variation defined by partitions $D = \{0 \le t_0 < \cdots < t_r \le T\}$:

$$\|X\|_p := \left(\sup_{D \subset [0,T]} \sum_{i=0}^{r-1} \left\| X_{t_{i+1}} - X_{t_i} \right\|^p \right)^{1/p}.$$

Definition 1 (Signature)

Let $X \in \mathcal{V}^p([0,T],\mathbb{R}^d)$. The signature $\mathcal{S}(X)$ is the sequence of iterated integrals:

$$\mathcal{S}(X) = (1, X^1, \cdots, X^k, \cdots) \in \oplus_{k=0}^{\infty} (\mathbb{R}^d)^{\otimes k},$$

where the k-th level is defined as: $X^k := \int_{0<t_1<\cdots<t_k<T} \mathrm{d}X_{t_1} \otimes \cdots \otimes \mathrm{d}X_{t_k}$.

- **Truncated Signature:**
  - We denote $\mathcal{S}^M(X) := (1, X^1, \cdots, X^M)$.
  - Dimension: $\frac{d^{M+1}-1}{d-1}$.

Definition 2 (Log-Signature [Liao et al., 2019])

The log-signature, denoted by $\log \mathcal{S}(X)$, provides a more compact representation:

$$\log \mathcal{S}(X) := \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n} (\mathcal{S}(X) - I)^{\otimes n}, \tag{5}$$

where $I = (1, 0, 0, \cdots)$ is the multiplicative identity in the tensor algebra.

Why Signatures?

Several properties make signatures suitable for our framework:

- **Uniqueness via Time Augmentation**
  - Standard signatures are unique only up to "tree-like equivalence."
  - Fix: Augment path with time $\hat{X}_t = (t, X_t)$.
  - $\mathcal{S}(\hat{X})$ uniquely characterizes X [Boedihardjo et al., 2016].
- **Efficient Representation**
  - Factorial decay allows accurate representation w. low truncation order M.
- **Universality**
  - Sig. serve as a universal feature map for sequential data [Bonnier et al., 2019].
- See also [Min and Hu, 2021]
- Many applications in finance, see e.g. [Cuchiero et al., 2025]

Approximating the Embedding

- We use the signature of the common noise $\hat{W}^0$ (augmented with time) to capture the information in the conditional law.

- **Approximation Strategy:**
  - We seek functions m_i to approximate the original embeddings m_i.
  - Replace
    - a function of the conditional law $\mathcal{L}(\Theta|\mathcal{F}^0)$
    - with a function of the signature $\mathcal{S}^M(\hat{W}^0)$

$$m_i\big(t, \Theta_t, \mathcal{L}(\Theta_t|\mathcal{F}^0_t)\big) \approx \mathrm{m}_i\big(t, \Theta_t, \mathcal{S}^M(\hat{W}^0_{[0,t]})\big), \quad i = 1, .., 4$$
$$m_5\big(X_T, \mathcal{L}(X_T|\mathcal{F}^0_T)\big) \approx \mathrm{m}_5\big(X_T, \mathcal{S}^M(\hat{W}^0_{[0,T]})\big).$$

Approximating the Embedding

- We use the signature of the common noise $\hat{W}^0$ (augmented with time) to capture the information in the conditional law.

- **Approximation Strategy:**
  - We seek functions m_i to approximate the original embeddings m_i.
  - Replace
    - a function of the conditional law $\mathcal{L}(\Theta|\mathcal{F}^0)$
    - with a function of the signature $\mathcal{S}^M(\hat{W}^0)$

$$m_i\big(t, \Theta_t, \mathcal{L}(\Theta_t|\mathcal{F}_t^0)\big) \approx \mathsf{m}_i\big(t, \Theta_t, \mathcal{S}^M(\hat{W}^0_{[0,t]})\big), \quad i = 1, .., 4$$
$$m_5\big(X_T, \mathcal{L}(X_T|\mathcal{F}_T^0)\big) \approx \mathsf{m}_5\big(X_T, \mathcal{S}^M(\hat{W}^0_{[0,T]})\big).$$

- Note: We can also replace $\mathcal{S}^M$ with $\log \mathcal{S}^M$.

Outline

Main Ideas

- **MC Simulation:** system of (non-interacting) FBSDEs to get the distribution
  - N_1: Sample paths of individual noise W.
  - N_2: Sample paths of common noise W^0.
  - Partition π: $0 = t_0 < \cdots < t_{N_T} = T$.
- **Supervised learning:** capture dependence on mean field through embeddings
  - Parameterized embeddings: m_i^0.
  - M: Level of signature truncation (dim d_{sig}).
- **DeepBSDE:** reformulate FBSDEs as a control problem and solve it using DL
  - Parameterized initial cond. u^0 and decoupling fields: $v^0, v^{0,0}$.
- Repeat

Algorithm Overview

- **Iterative Process:** At iteration k, based on outputs from $k-1$:

1. **MC Simulation:**
   - Simulate (X, Y, Z, Z^0) using previous controls and embeddings.
   - Compute empirical distributions.

Algorithm Overview

- **Iterative Process:** At iteration k, based on outputs from $k-1$:

1. **MC Simulation:**
   - Simulate (X, Y, Z, Z^0) using previous controls and embeddings.
   - Compute empirical distributions.
2. **Supervised Learning:** Update the embedding functions m_i^k.
   - Targets: True m_i evaluated on empirical distributions.
   - Inputs: State variables and signatures of common noise.

Algorithm Overview

- **Iterative Process:** At iteration k, based on outputs from $k-1$:

1. **MC Simulation:**
   - Simulate (X, Y, Z, Z^0) using previous controls and embeddings.
   - Compute empirical distributions.
2. **Supervised Learning:** Update the embedding functions m_i^k.
   - Targets: True m_i evaluated on empirical distributions.
   - Inputs: State variables and signatures of common noise.
3. **DeepBSDE:** using the updated embeddings m_i^k, train new controls $u^k, v^k, v^{0,k}$.

Step 1: Generate Samples

- $N_1 \times N_2$ samples: simulate discrete-time dyn. using embeddings & controls from iter. $k-1$:

$$\tilde{Y}_0^{k,n_1,n_2} = u^{k-1}(\tilde{X}_0^{k,n_1,n_2}),$$

$$\tilde{Z}_{t_i}^{k,n_1,n_2} = v^{k-1}\Big(t_i, \tilde{X}_{t_i}^{k,n_1,n_2}, \mathcal{S}^M(\hat{W}_{[0,t_i]}^{0,k,n_2})\Big),$$

$$\tilde{Z}_{t_i}^{0,k,n_1,n_2} = v^{0,k-1}\Big(t_i, \tilde{X}_{t_i}^{k,n_1,n_2}, \mathcal{S}^M(\hat{W}_{[0,t_i]}^{0,k,n_2})\Big),$$

$$\begin{aligned}\tilde{X}_{t_{i+1}}^{k,n_1,n_2} &= \tilde{X}_{t_i}^{k,n_1,n_2} + b\Big(t_i, \tilde{\Theta}_{t_i}^{k,n_1,n_2}, \tilde{Z}_{t_i}^{0,k,n_1,n_2}, \mathsf{m}_1^{k-1}\big(t_i, \tilde{\Theta}_{t_i}^{k,n_1,n_2}, \mathcal{S}^M(\hat{W}_{[0,t_i]}^{0,k,n_2})\big)\Big)\Delta t_i \\ &+ \sigma\Big(t_i, \tilde{\Theta}_{t_i}^{k,n_1,n_2}, \mathsf{m}_2^{k-1}(\dots)\Big)\Delta W_{t_i}^{k,n_1,n_2} \\ &+ \sigma^0\Big(t_i, \tilde{\Theta}_{t_i}^{k,n_1,n_2}, \mathsf{m}_3^{k-1}(\dots)\Big)\Delta W_{t_i}^{0,k,n_2},\end{aligned}$$

$$\begin{aligned}\tilde{Y}_{t_{i+1}}^{k,n_1,n_2} &= \tilde{Y}_{t_i}^{k,n_1,n_2} - h\Big(t_i, \tilde{\Theta}_{t_i}^{k,n_1,n_2}, \mathsf{m}_4^{k-1}\big(t_i, \tilde{\Theta}_{t_i}^{k,n_1,n_2}, \mathcal{S}^M(\hat{W}_{[0,t_i]}^{0,k,n_2})\big)\Big)\Delta t_i \\ &+ \tilde{Z}_{t_i}^{k,n_1,n_2}\Delta W_{t_i}^{k,n_1,n_2} + Z_{t_i}^{0,k,n_1,n_2}\Delta W_{t_i}^{0,k,n_2}.\end{aligned}$$

- Compute empirical measures ("conditioned" on common noise):

$$\nu_{t_i}^{k,n_2} := \frac{1}{N_1}\sum_{n_1=1}^{N_1} \delta_{\tilde{\Theta}_{t_i}^{k,n_1,n_2}}, \qquad \mu_T^{k,n_2} := \frac{1}{N_1}\sum_{n_1=1}^{N_1} \delta_{\tilde{X}_T^{k,n_1,n_2}}.$$

Step 2: Learn Distribution Dependence

- Approximate the mapping signatures $\mapsto$ embeddings via Supervised Learning.

- **Optimization Objective:**

$$\mathsf{m}_i^k := \operatorname{arginf}_{\mathsf{m}} \sum_{n_1,n_2,i} \left\| m_i(t_i, \tilde{\Theta}_{t_i}^{k,n_1,n_2}, \nu_{t_i}^{n_2,k}) - \mathsf{m}\Big(t_i, \tilde{\Theta}_{t_i}^{k,n_1,n_2}, \mathcal{S}^M(\hat{W}_{[0,t_i]}^{0,k,n_2})\Big) \right\|_2^2,$$

$$\mathsf{m}_5^k := \operatorname{arginf}_{\mathsf{m}} \sum_{n_1,n_2} \left\| m_5(\tilde{X}_T^{k,n_1,n_2}, \mu_T^{n_2,k}) - \mathsf{m}\Big(\tilde{X}_{t_{N_T}}^{k,n_1,n_2}, \mathcal{S}^M(\hat{W}_{[0,t_{N_T}]}^{0,k,n_2})\Big) \right\|_2^2.$$

- **Function Classes:**
  1. Linear functional of signature: $\mathsf{m}(\cdot, \mathcal{S}) = \langle \varphi(\cdot), \mathcal{S} \rangle$, where φ is a NN.
  2. Neural network: m as NN directly.

- Update $u^k, v^k, v^{0,k}$ by minimizing the terminal loss (**DeepBSDE** method):

$$u^k, v^k, v^{0,k} = \text{arginf}_{u,v,v^0} \sum_{n=1}^{N} \left\| g\left(\check{X}_T^{k,n}, \mathsf{m}_5^k(\dots)\right) - \check{Y}_T^{k,n} \right\|_2^2$$

- **Constraints (Forward/Backward Dynamics):**

$$\check{Y}_0^{k,n} = u^k(\tilde{X}_0^{k,n}), \quad \tilde{Z}_{t_i}^{k,n} = v^k(\dots), \quad \tilde{Z}_{t_i}^{0,k,n} = v^{0,k}(\dots),$$

$$\check{X}_{t_{i+1}}^{k,n} = \check{X}_{t_i}^{k,n} + b\left(t_i, \check{\Theta}_{t_i}^{k,n}, Z_{t_i}^{0,k,n}, \mathsf{m}_1^k(t_i, \tilde{\Theta}_{t_i}^{k,n_1,n_2}, \mathcal{S}^M)\right)\Delta t_i$$
$$+ \sigma\left(\dots, \mathsf{m}_2^k(\dots)\right)\Delta\check{W}_{t_i}^{k,n} + \sigma^0\left(\dots, \mathsf{m}_3^k(\dots)\right)\Delta\check{W}_{t_i}^{0,k,n},$$

$$\check{Y}_{t_{i+1}}^{k,n} = \check{Y}_{t_i}^{k,n} - h\left(t_i, \check{\Theta}_{t_i}^{k,n}, \mathsf{m}_4^k(t_i, \tilde{\Theta}_{t_i}^{k,n_1,n_2}, \mathcal{S}^M)\right)\Delta t_i$$
$$+ Z_{t_i}^{k,n}\Delta\check{W}_{t_i}^{k,n} + Z_{t_i}^{0,k,n}\Delta\check{W}_{t_i}^{0,k,n}.$$

Outline

Theoretical Setup: Decoupled MV-FBSDE

We analyze the convergence in the following situation.

Assumption B

- Decoupled system where (Y, Z) does not appear in the forward dynamics,
- Volatilities depend only on (t, x).

Similar assumption as [Han et al., 2024].

Target MKV FBSDE System becomes:

$$\begin{cases} \mathrm{d}X_t = b(t, X_t, m_1(t, X_t, \mathcal{L}(X_t|\mathcal{F}_t^0)))\,\mathrm{d}t + \sigma(t, X_t)\,\mathrm{d}W_t + \sigma^0(t, X_t)\,\mathrm{d}W_t^0, \\ \mathrm{d}Y_t = -h(t, \Theta_t, Z_t^0, m_4(t, X_t, \mathcal{L}(X_t|\mathcal{F}_t^0)))\,\mathrm{d}t + Z_t\,\mathrm{d}W_t + Z_t^0\,\mathrm{d}W_t^0, \\ X_0 \sim \mu_0, \quad Y_T = g(X_T, m_5(X_T, \mathcal{L}(X_T|\mathcal{F}_T^0))). \end{cases} \tag{6}$$

The Iterative System (Iteration k)

At **iteration** k, given learned functions $(\mathsf{m}_1^k, \mathsf{m}_4^k, \mathsf{m}_5^k) \in \mathcal{M}_1 \times \mathcal{M}_1 \times \mathcal{M}_2$,

$$\begin{cases} \mathrm{d}X_t^k = b(t, X_t^k, \mathsf{m}_1^k(t, X_t^k, \mathcal{S}^M(\hat{W}^0_{[0,t]}))) \,\mathrm{d}t + \sigma(t, X_t^k) \,\mathrm{d}W_t + \sigma^0(t, X_t^k) \,\mathrm{d}W_t^0, \\ \mathrm{d}Y_t^k = -h(t, \Theta_t^k, Z_t^{0,k}, \mathsf{m}_4^k(t, X_t^k, \mathcal{S}^M(\hat{W}^0_{[0,t]}))) \,\mathrm{d}t + Z_t^k \,\mathrm{d}W_t + Z_t^{0,k} \,\mathrm{d}W_t^0, \\ X_0^k \sim \mu_0, \quad Y_T^k = g(X_T^k, \mathsf{m}_5^k(X_T^k, \mathcal{S}^M(\hat{W}^0_{[0,t]}))), \end{cases} \tag{7}$$

where we use the function sets $\mathcal{M}_1, \mathcal{M}_2$: Lip. functions with linear growth in x:

At **iteration** k, given learned functions $(\mathsf{m}_1^k, \mathsf{m}_4^k, \mathsf{m}_5^k) \in \mathcal{M}_1 \times \mathcal{M}_1 \times \mathcal{M}_2$,

$$\begin{cases} \mathrm{d}X_t^k = b(t, X_t^k, \mathsf{m}_1^k(t, X_t^k, \mathcal{S}^M(\hat{W}^0_{[0,t]}))) \,\mathrm{d}t + \sigma(t, X_t^k)\,\mathrm{d}W_t + \sigma^0(t, X_t^k)\,\mathrm{d}W_t^0, \\ \mathrm{d}Y_t^k = -h(t, \Theta_t^k, Z_t^{0,k}, \mathsf{m}_4^k(t, X_t^k, \mathcal{S}^M(\hat{W}^0_{[0,t]}))) \,\mathrm{d}t + Z_t^k\,\mathrm{d}W_t + Z_t^{0,k}\,\mathrm{d}W_t^0, \\ X_0^k \sim \mu_0, \quad Y_T^k = g(X_T^k, \mathsf{m}_5^k(X_T^k, \mathcal{S}^M(\hat{W}^0_{[0,t]}))), \end{cases} \tag{7}$$

where we use the function sets $\mathcal{M}_1, \mathcal{M}_2$: Lip. functions with linear growth in x:

$$\mathcal{M}_1 = \Big\{ m : [0,T] \times \mathbb{R}^d \times \mathbb{R}^{d_{\mathsf{sig}}} \to \mathbb{R}^\ell,$$
$$\|m(t,x,s) - m(t',x',s)\|^2 \le M\|x - x'\|^2, \quad \|m(t,x,s)\| \le M[1 + \|x\|] \Big\},$$

$\mathcal{M}_2 =$ same without time,

where $d_{\mathsf{sig}} =$ dim. of truncated signature of $\hat{W}_t = (t, W_t^0)$, i.e, $d_{\mathsf{sig}} = \frac{(q+1)^{M+1}-1}{q}$.

Assumption C

To prove convergence, we assume:

Assumption C

(a) **Lipschitz Coefficients:** $b, \sigma, \sigma^0, h, g$ are Lipschitz w.r.t all variables (except time) with constant L.

$$\|b(t,x,m) - b(t,x',m')\|^2 + \cdots \leq L[\|x - x'\|^2 + \cdots + \|m - m'\|^2].$$

(b) **Lipschitz Embeddings:** m_1, m_4, m_5 are Lipschitz w.r.t spatial variable and Wasserstein distance:

$$\|m_1(t,x,\mu) - m_1(t,x',\mu')\|^2 + \cdots \leq L[\|x - x'\|^2 + \mathcal{W}_2^2(\mu,\mu')].$$

(c) **Boundedness at 0:** Coefficients and embeddings at origin (and initial moment) are bounded by constant K.

Lemma 3: Well-posedness

Lemma 3

Under the standing assumptions,

1. *The target system* (6) *has a unique solution.*
2. *The iterative system* (7) *has a unique solution for all* k*.*
3. *Uniform bounds hold (indep. of* k*):*

$$\mathbb{E}\left[\sup_{0\le t\le T}\|X_t^k\|^2+\|Y_t^k\|^2+\int_0^T(\|Z_t^k\|_F^2+\|Z_t^{0,k}\|_F^2)\,\mathrm{d}t\right]\le C. \tag{8}$$

Proof Sketch: Well-posedness

For system (7) **(Iterative):**

- The term $\mathcal{S}^M(\hat{W}^0_{[0,t]})$ is $\mathcal{F}_t$-measurable.
- The coefficients $\tilde{b}^k(t,x,\omega) = b(t,x,\mathsf{m}_1^k(\dots))$ satisfy Lipschitz conditions.
- Apply standard BSDE results [Zhang, 2017].

For system (6) **(Target):**

- Use Picard iterations on the law term. Define sequence $\tilde{X}^n$:
$$\mathrm{d}\tilde{X}_t^{n+1} = b\Big(t, \tilde{X}_t^n, m_1\big(t, \tilde{X}_t^n, \mathcal{L}(\tilde{X}_t^n|\mathcal{F}_t^0)\big)\Big)\,\mathrm{d}t + \dots$$
- Analyze difference ΔX^{n+1}. By Itô's formula and Lipschitz assumptions:
$$\lambda\mathbb{E}\Big[\int_0^T e^{-\lambda t}\big\|\Delta X_t^{n+1}\big\|^2\,\mathrm{d}t\Big] \le \mathbb{E}\Big[\int_0^T \dots\Big(\dots + \mathcal{W}_2^2\Big)\dots\Big].$$
- Choosing large λ yields contraction. $\tilde{X}^n \to \mathcal{X}$ in $\mathbb{S}^2(\mathbb{F})$.

$\square$

Theorem 4

Under the standing assumptions, there exist constants $C > 0$ and $0 < q < 1$ such that:

$$\sup_{0\le t\le T}\left[\mathbb{E}\|X_t - X_t^k\|^2 + \mathbb{E}\|Y_t - Y_t^k\|^2\right] + \int_0^T \left[\mathbb{E}\|Z_t - Z_t^k\|^2 + \mathbb{E}\|Z_t^0 - Z_t^{0,k}\|^2\right] \mathrm{d}t$$

$$\le C\Bigg\{ q^k + \sum_{j=0}^{k-1} q^{k-j} \int_0^T \mathbb{E}\|\mathsf{m}_1^{j+1}(t, X_t^{j+1}, \mathcal{S}^M) - m_1(t, X_t^{j+1}, \mathcal{L}(X_t^j|\mathcal{F}_t^0))\|^2 \,\mathrm{d}t$$

$$+ \int_0^T \mathbb{E}\|\mathsf{m}_4^k(t, X_t^k, \mathcal{S}^M) - m_4(t, X_t^k, \mathcal{L}(X_t^{k-1}|\mathcal{F}_t^0))\|^2 \,\mathrm{d}t$$

$$+ \mathbb{E}\|\mathsf{m}_5^k(X_T^k, \mathcal{S}^M) - m_5(X_T^k, \mathcal{L}(X_T^{k-1}|\mathcal{F}_T^0))\|^2 \Bigg\}.$$

Proof: Forward Component Estimate

Let $\delta X^k = X - X^k$. Define the approximation error I_1^k:

$$I_1^k := \int_0^T \mathbb{E}\|m_1(t, X_t^k, \mathcal{L}(X_t^{k-1}|\mathcal{F}_t^0)) - m_1^k(t, X_t^k, \mathcal{S}^M)\|^2 \, \mathrm{d}t.$$

Using Lipschitz properties and $\mathcal{W}_2$ distance:

$$\begin{aligned}
\sup_{0\le s\le t} \mathbb{E}\|\delta X_s^k\|^2 &\le C \int_0^t \mathbb{E}\|m_1(s, X_s^k, \mathcal{L}(X_s|\mathcal{F}_s^0)) - m_1^k(s, X_s^k, \mathcal{S}^M)\|^2 \, \mathrm{d}s \\
&\le C \int_0^t \mathbb{E}\left[\mathcal{W}_2^2(\mathcal{L}(X_s|\mathcal{F}_s^0), \mathcal{L}(X_s^{k-1}|\mathcal{F}_s^0))\right] \mathrm{d}s + C I_1^k \\
&\le C \int_0^t \mathbb{E}\|\delta X_s^{k-1}\|^2 \, \mathrm{d}s + C I_1^k.
\end{aligned}$$

By induction, this leads to a term decaying with q^k plus accumulated errors.

Proof: Backward Component Estimate

Similarly for (Y, Z, Z^0), using errors I_4^k, I_5^k:

$$
\begin{aligned}
&\sup \mathbb{E}\|\delta Y^k\|^2 + \dots \\
&\leq C \int_t^T \mathbb{E}\|h(\dots) - h(\dots)\|^2 \,\mathrm{d}s + C\mathbb{E}\|g(\dots) - g(\dots)\|^2 \\
&\leq C \int_t^T \mathbb{E}\|\delta X_s^k\|^2 + \mathbb{E}\|m_4(\dots \mathcal{L}(X) \dots) - m_4(\dots \mathcal{L}(X^{k-1}) \dots)\|^2 \,\mathrm{d}s + CI_4^k \dots \\
&\leq C \int_0^T \mathbb{E}\|\delta X_s^{k-1}\|^2 \,\mathrm{d}s + CI_4^k + \dots + CI_1^{k-1} + CI_1^k.
\end{aligned}
$$

Combining forward and backward estimates yields the result.

□

Lemma : Change of Measure

Issue: The error terms in Theorem 5 depend on current state, e.g.

$$\mathsf{m}_4^k(t, X_t^k, \mathcal{S}^M)$$

but we only observe the state generated with *previous* embeddings: X_t^{k-1}.

Solution: Use Girsanov theorem to relate them.

Lemma : Change of Measure

Issue: The error terms in Theorem 5 depend on current state, e.g.

$$\mathsf{m}_4^k(t, X_t^k, \mathcal{S}^M)$$

but we only observe the state generated with *previous* embeddings: X_t^{k-1}.

Solution: Use Girsanov theorem to relate them.

Lemma 5

Assume the function b *is of the form* $b(t,x,m) = \sigma(t,x)\phi(t,x,m) + b^0(t,x)$*, s.t.*

$$\|\sigma\|_S \le K,\ \|\phi\|_\infty \le K,\ \|b^0(t,x) - b^0(t,x')\|^2 \le L\|x - x'\|^2,\ \|b^0(t,0)\|^2 \le K,$$
$$\|\phi(t,x,m) - \phi(t,x',m')\|^2 \le L\left[\|x - x'\|^2 + \|m - m'\|^2\right],$$

for all $t \in [0,T]$.
Let $\bar{m}^1, \bar{m}^2 \in \mathcal{M}_1$. *Let* $\bar{X}^i$ *solve the associated SDE:*

$$\mathrm{d}\bar{X}_t^i = b(t, \bar{X}_t^i, \bar{m}^i(t, \bar{X}_t^i, \mathcal{S}^M(\hat{W}_{[0,t]}^0)))\,\mathrm{d}t + \sigma(t, \bar{X}_t^i)\,\mathrm{d}W_t + \sigma^0(t, \bar{X}_t^i)\,\mathrm{d}W_t^0.$$

Then for any functions $(x, W^0) \mapsto m(x, W^0)$ *and* $(x, W^0) \mapsto m'(x, W^0)$ *that are at most linearly growing in* x*, and any* $\epsilon \in (0,1)$*, there exists a constant* $C(\epsilon)$ *s.t.*

$$\mathbb{E}\|m(\bar{X}_t^1, W^0) - m'(\bar{X}_t^1, W^0)\|^2 \le C(\epsilon)\left[\mathbb{E}\|m(\bar{X}_t^2, W^0) - m'(\bar{X}_t^2, W^0)\|^2\right]^\epsilon,\ \forall t \in [0,T].$$

Final Convergence Theorem

Combining the first convergence theorem with the previous lemma:

Theorem 6

Under the standing assumptions, there exist constants $\epsilon \in (0,1)$, $C(\epsilon) > 0$ *and* $0 < q < 1$ *such that the total error is bounded by:*

$$C(\epsilon)\Big\{ q^k + \sum_{j=0}^{k-1} q^{k-j} \int_0^T \mathbb{E}\|\mathsf{m}_1^{j+1}(t, X_t^j, \mathcal{S}^M) - m_1(t, X_t^j, \mathcal{L}(X_t^j|\mathcal{F}_t^0))\|^2 \, \mathrm{d}t$$
$$+ \int_0^T \mathbb{E}\|\mathsf{m}_4^k(t, X_t^{k-1}, \mathcal{S}^M) - m_4(t, X_t^{k-1}, \mathcal{L}(X_t^{k-1}|\mathcal{F}_t^0))\|^2 \, \mathrm{d}t$$
$$+ \mathbb{E}\|\mathsf{m}_5^k(X_T^{k-1}, \mathcal{S}^M) - m_5(X_T^{k-1}, \mathcal{L}(X_T^{k-1}|\mathcal{F}_T^0))\|^2 \Big\}^{\epsilon}.$$

This bound depends on the **supervised learning error** evaluated on the **previous iteration's paths** (which is available data).

Outline

Implementation Details:

- Framework: PyTorch + `signatory`.
- Discretization: $T = 1$, $N_T = 120$ steps ($4N_T$ for common noise signature).
- Optimizer: Adam.

Performance Metrics:

1. **Mean Absolute Error (MAE)** at time t_n:

$$\text{MAE}_{t_n} = \underset{j}{\text{mean}} \Big| m\big(t_n, X^j_{t_n}, \mathcal{L}(X^j_{t_n} | \mathcal{F}^0_{t_n})\big) - \mathsf{m}\big(t_n, X^j_{t_n}, \mathcal{S}^M(\hat{W}^{0,j}_{[0,t_n]})\big) \Big|.$$

2. **Mean Euclidean Error (MEE)**:

$$\text{MEE}(\theta, \hat{\theta}) = \underset{j,\, t_n}{\text{mean}} \|\theta^j_{t_n} - \hat{\theta}^j_{t_n}\|.$$

Exp. 1: Supervised Learning of m

Objective: Let $X_t = W_t + W_t^0$. Approximate the conditional expectation function:

$$m(t, x, \mathcal{L}(X_t|\mathcal{F}_t^0)) = \left(\frac{q}{q+2t}\right)^{q/2} e^{-\frac{\|x - W_t^0\|^2}{q+2t}}.$$

Comparison:

- **Archi. 1:** Linear functional of truncated signature $\mathcal{S}^M(\hat{W}^0_{[0,t]})$.
- **Archi. 2:** Feedforward NN with input $\mathcal{S}^M(\hat{W}^0_{[0,t]})$.

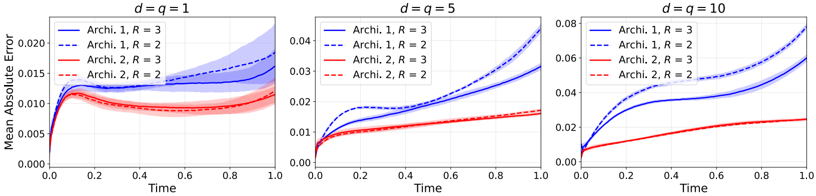

Figure: MAE: Archi 2 (NN) consistently outperforms Archi 1 (Linear) across dims.

Exp. 1: Effect of Signature Truncation (M)

We evaluate robustness by varying the truncation order M and NN depth R.

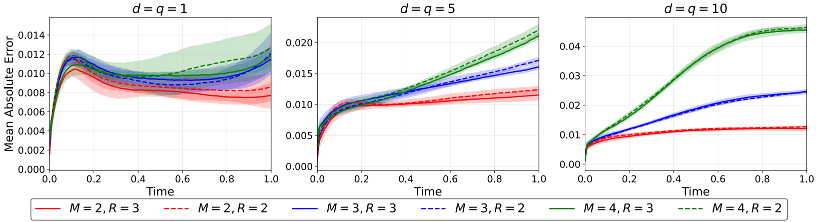

Figure: MAE increases with M and dimension.

- **Observation:** Performance degrades as M increases.
- **Reasoning:** Higher-order terms add variance/complexity without capturing necessary drift structure for this specific Brownian motion problem.
- **Choice:** $M = 2$ selected for balance.

Exp. 1: Log-Signatures Scalability

For high dimensions ($d = 15$), we test **Log-Signatures** to reduce feature dimension.

Configuration	$M = 2$		$M = 3$	
	log-sig	**sig**	**log-sig**	**sig**
$d = 5, H = 2$	9.58e-3 (6.84e-4)	1.05e-2 (3.76e-4)	1.03e-2 (3.74e-4)	1.25e-2 (4.45e-4)
$d = 5, H = 3$	9.04e-3 (3.10e-4)	1.00e-2 (3.39e-4)	1.02e-2 (2.34e-4)	1.25e-2 (3.70e-4)
$d = 10, H = 2$	9.61e-3 (4.07e-4)	1.11e-2 (2.25e-4)	1.15e-2 (1.70e-4)	1.74e-2 (2.22e-4)
$d = 10, H = 3$	9.56e-3 (3.13e-4)	1.07e-2 (3.57e-4)	1.14e-2 (3.36e-4)	1.77e-2 (4.76e-4)
$d = 12, H = 2$	9.80e-3 (2.67e-4)	1.16e-2 (1.67e-4)	1.22e-2 (1.25e-4)	2.33e-2 (5.70e-4)
$d = 12, H = 3$	9.62e-3 (1.57e-4)	1.12e-2 (3.53e-4)	1.23e-2 (3.58e-4)	2.47e-2 (9.69e-4)
$d = 15, H = 2$	9.83e-3 (2.29e-4)	1.23e-2 (4.81e-4)	1.62e-2 (4.79e-4)	3.73e-2 (1.92e-3)
$d = 15, H = 3$	9.76e-3 (2.38e-4)	1.19e-2 (3.05e-4)	1.59e-2 (6.14e-4)	3.49e-2 (8.11e-4)

Table: Time-averaged performance measured by MAE of M for dimensions $d = 5, 10, 12, 15$. Results are aggregated over 10^3 independent trajectories. Each entry reports the mean MAE with standard deviation (in parentheses), computed from 5 repeated experiments with different random seeds.

$\Rightarrow$ Log-signatures are more robust to truncation order in high dimensions.

Exp. 2: MV-FBSDE in Random Environment

Next, we consider the following MV-FBSDE in random environment for (X_t, Y_t, Z_t, Z_t^0):

$$
\begin{aligned}
\mathrm{d}X_t^i = &\Bigg[\sin\left(\mathbb{E}_{x'_t\sim\mathcal{L}(X_t|\mathcal{F}_t^0)} e^{-\frac{\|X_t - x'_t\|^2}{d}} - e^{-\frac{\|X_t - \mathbb{E}[X_t|\mathcal{F}_t^0]\|^2}{d+2t}}\left(\frac{d}{d+2t}\right)^{\frac{d}{2}}\right) \\
&+ \frac{1}{2}\left(\mathbb{E}[Y_t|\mathcal{F}_t^0] - \sin\left(t + \frac{1}{\sqrt{d}}\sum_{i=1}^{d}\mathbb{E}[X_t^i|\mathcal{F}_t^0]\right)e^{-\frac{t}{2}}\right)\Bigg]\mathrm{d}t + \mathrm{d}W_t^i + \mathrm{d}W_t^{0,i}, \quad 1 \le i \le d \\
\mathrm{d}Y_t = &\left[\frac{\sum_{i=1}^{d}(Z_t^i + Z_t^{0,i})}{2\sqrt{d}} - Y_t + \sqrt{2Y_t^2 + \|Z_t\|^2 + \|Z_t^0\|^2 + 1} - \sqrt{3}\right]\mathrm{d}t \\
&+ Z_t \cdot \mathrm{d}W_t + Z_t^0 \cdot \mathrm{d}W_t^0,
\end{aligned}
$$

with initial and terminal conditions $X_0^i = 0$ and $Y_T = \sin\left(T + \frac{\sum_{i=1}^{d} X_T^i}{\sqrt{d}}\right)$, where X_t^i is the i^{th} entry of the d-dimensional process X_t.

Exp. 2: MV-FBSDE in Random Environment

Next, we consider the following MV-FBSDE in random environment for (X_t, Y_t, Z_t, Z_t^0):

$$\mathrm{d}X_t^i = \Big[\sin\Big(\mathbb{E}_{x_t'\sim\mathcal{L}(X_t|\mathcal{F}_t^0)} e^{-\frac{\|X_t - x_t'\|^2}{d}} - e^{-\frac{\|X_t - \mathbb{E}[X_t|\mathcal{F}_t^0]\|^2}{d+2t}}\Big(\frac{d}{d+2t}\Big)^{\frac{d}{2}}\Big)$$

$$+\frac{1}{2}\Big(\mathbb{E}[Y_t|\mathcal{F}_t^0] - \sin\Big(t + \frac{1}{\sqrt{d}}\sum_{i=1}^{d}\mathbb{E}[X_t^i|\mathcal{F}_t^0]\Big)e^{-\frac{t}{2}}\Big)\Big]\,\mathrm{d}t + \mathrm{d}W_t^i + \mathrm{d}W_t^{0,i}, \quad 1 \le i \le d$$

$$\mathrm{d}Y_t = \Big[\frac{\sum_{i=1}^{d}(Z_t^i + Z_t^{0,i})}{2\sqrt{d}} - Y_t + \sqrt{2Y_t^2 + \|Z_t\|^2 + \|Z_t^0\|^2 + 1} - \sqrt{3}\Big]\,\mathrm{d}t$$

$$+ Z_t \cdot \mathrm{d}W_t + Z_t^0 \cdot \mathrm{d}W_t^0,$$

with initial and terminal conditions $X_0^i = 0$ and $Y_T = \sin\Big(T + \frac{\sum_{i=1}^{d} X_T^i}{\sqrt{d}}\Big)$, where X_t^i is the i^{th} entry of the d-dimensional process X_t.

One can check that the **solution** to the above MV-FBSDE is

$$X_t = W_t + W_t^0, \quad Y_t = \sin\Big(t + \frac{\sum_{i=1}^{d} X_t^i}{\sqrt{d}}\Big), \quad Z_t^i = Z_t^{0,i} = \frac{1}{\sqrt{d}}\cos\Big(t + \frac{\sum_{i=1}^{d} X_t^i}{\sqrt{d}}\Big).$$

The corresponding m functions are $m_2 \equiv m_3 \equiv m_4 \equiv m_5 \equiv 0$, and

$$m_1 = \Big(\tilde{\mathbb{E}}\Big[e^{-\frac{\|x - \tilde{x}_t\|^2}{d}}\Big], \tilde{\mathbb{E}}[\tilde{x}_t], \tilde{\mathbb{E}}[\tilde{y}_t]\Big) \in \mathbb{R}^{1+d+1},$$

where the expected value $\tilde{\mathbb{E}}$ is with respect to $(\tilde{x}_t, \tilde{y}_t) \sim \mathcal{L}(X_t, Y_t|\mathcal{F}_t^0)$.

Exp. 2: Results

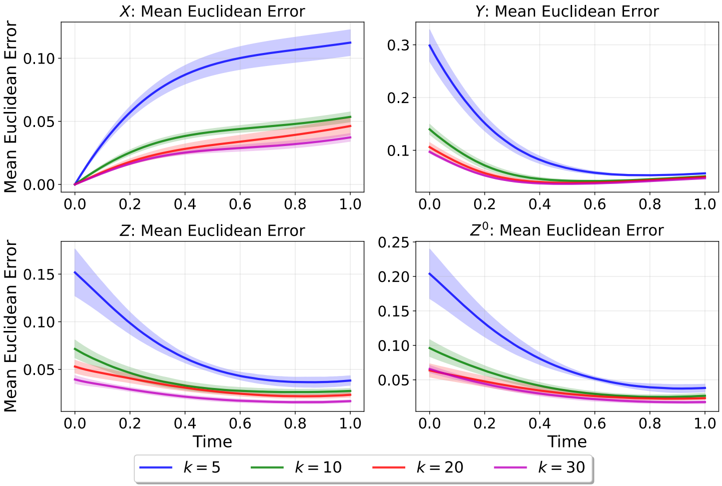

Figure: MEE convergence over iterations ($d = 5$).

$\Rightarrow$ Error stabilizes after $k = 20$

Exp. 2: Analytical vs Numerical Solution

Comparison after 30 iterations ($d = 5$).

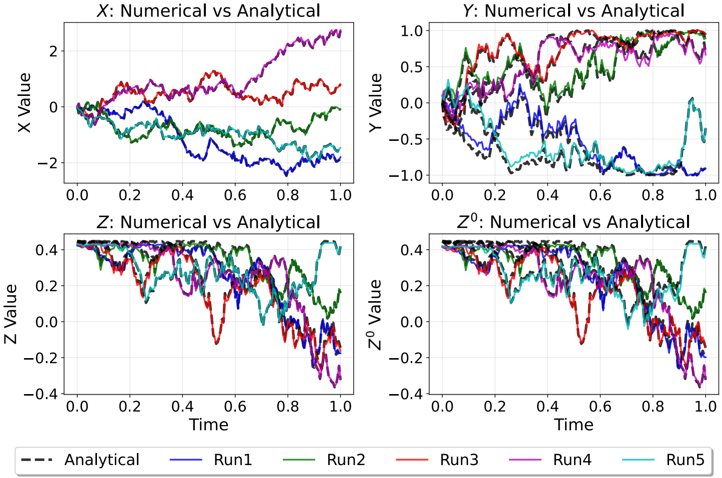

Figure: Solid: Numerical, Dashed: Analytical.

$\Rightarrow$ The algorithm accurately captures the pathwise dynamics.

Exp. 3: Flocking Model (Cucker-Smale type)

Representative Agent Dynamics: State $(x_t, v_t) \in \mathbb{R}^{2d}$ evolves as:

$$\begin{cases} \mathrm{d}x_t = v_t \ \mathrm{d}t, \\ \mathrm{d}v_t = u_t \ \mathrm{d}t + C \ \mathrm{d}W_t + D \ \mathrm{d}W_t^0, \end{cases} \tag{9}$$

where u_t is acceleration, W_t is idiosyncratic, and W_t^0 is common.

Objective: Minimize expected cost

$$J(u) = \mathbb{E} \int_0^T \Big(\underbrace{\|u_t\|_R^2}_{\text{energy}} + \underbrace{\mathcal{C}(x_t, v_t; f_t)}_{\text{misalignment}} \Big) \ \mathrm{d}t. \tag{10}$$

The misalignment cost $\mathcal{C}$ depends on the population distribution f:

$$\mathcal{C}(x, v; f) = \left\| \tilde{\mathbb{E}}_{(x', v') \sim f} \left[w(\|x - x'\|)(v' - v) \right] \right\|_Q^2 , \tag{11}$$

- Weight function: $w(r) := (1 + r^2)^{-\beta}$ ($\beta \geq 0$).
- Q, R: Symmetric positive definite weight matrices.

Equilibrium: Find $\hat{u}_t$ such that $f_t = \mathcal{L}(\hat{x}_t, \hat{v}_t | \mathcal{F}_t^0)$.

Exp. 3: Training Results

Training:

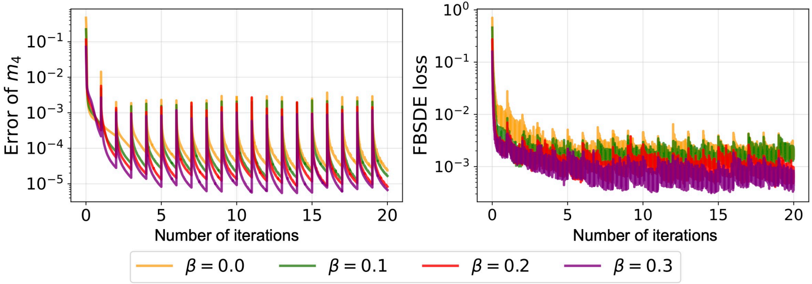

Figure: Left: Supervised loss (m_4). Right: DeepBSDE loss.

$\Rightarrow$ We observe fast convergence.

Exp. 3: Conditional Densities (Common Noise Effect)

Terminal density under different Common Noise realizations (Top vs Bottom rows).

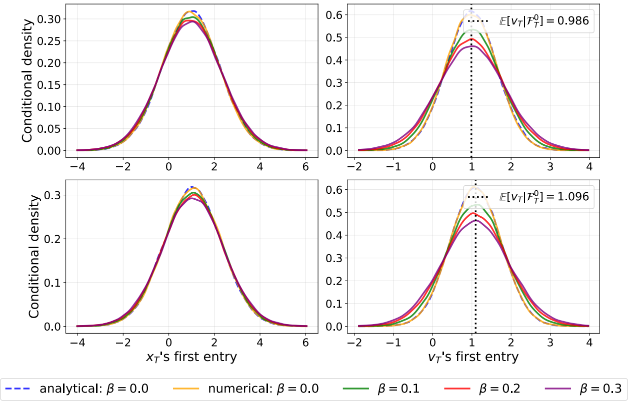

We observe:

- Larger β (weaker interaction) implies more dispersed density.
- Common noise shifts the mean.

Exp. 3: Benchmark Case ($\beta = 0$)

For $\beta = 0$ (Linear-Quadratic case), we have an analytical solution.

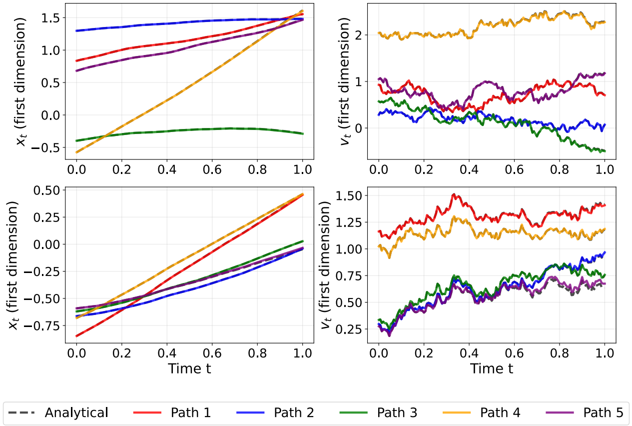

Figure: Top: Unconditional paths. Bottom: Conditional paths (fixed W^0).

$\rightarrow$ Numerical solutions match analytical benchmarks (dashed)

Outline

Conclusion & Future Directions

Summary of Contributions:

- **Novel Framework:** Proposed a scalable algorithm for solving MV-FBSDEs and MFGs with common noise.
- **Signature-based Embedding:** Utilized path signatures to efficiently capture the infinite-dimensional dependence on the conditional law $\mathcal{L}(X_t|\mathcal{F}_t^0)$.
- **Decoupled Training:** Combined interactions with DeepBSDEs to separate the learning of distribution embeddings from the control problem.
- **Validation:** Demonstrated robustness and accuracy on high-dimensional benchmarks (up to $d = 15$) and flocking models.

Future Directions:

- **Further Convergence Analysis:** Establishing theoretical convergence rates for the signature approximation under more general assumptions.
- **Broader Applications:** Extending to other MFGs (e.g., finite states, optimal stopping, infinite horizon, etc.).

Thank you for your attention!

References I

[Achdou et al., 2021] Achdou, Y., Cardaliaguet, P., Delarue, F., Porretta, A., and Santambrogio, F. (2021).
Mean Field Games: Cetraro, Italy 2019, volume 2281.
Springer Nature.

[Achdou and Lasry, 2018] Achdou, Y. and Lasry, J.-M. (2018).
Mean field games for modeling crowd motion.
In *Contributions to partial differential equations and applications*, pages 17–42. Springer.

[Bensoussan et al., 2013] Bensoussan, A., Frehse, J., Yam, P., et al. (2013).
Mean field games and mean field type control theory, volume 101.
Springer.

[Boedihardjo et al., 2016] Boedihardjo, H., Geng, X., Lyons, T., and Yang, D. (2016).
The signature of a rough path: uniqueness.
Advances in Mathematics, 293:720–737.

[Bonnier et al., 2019] Bonnier, P., Kidger, P., Arribas, I. P., Salvi, C., and Lyons, T. (2019).
Deep signature transforms.
In *Advances in Neural Information Processing Systems 32 (NeurIPS)*.

[Cardaliaguet et al., 2019] Cardaliaguet, P., Delarue, F., Lasry, J.-M., and Lions, P.-L. (2019).
The master equation and the convergence problem in mean field games.
Princeton University Press.

References II

[Carmona and Delarue, 2013] Carmona, R. and Delarue, F. (2013).
Probabilistic analysis of mean-field games.
SIAM Journal on Control and Optimization, 51(4):2705–2734.

[Carmona and Delarue, 2018a] Carmona, R. and Delarue, F. (2018a).
Probabilistic theory of mean field games with applications. I, volume 83 of *Probability Theory and Stochastic Modelling*.
Springer, Cham.
Mean field FBSDEs, control, and games.

[Carmona and Delarue, 2018b] Carmona, R. and Delarue, F. (2018b).
Probabilistic theory of mean field games with applications. II, volume 84 of *Probability Theory and Stochastic Modelling*.
Springer, Cham.
Mean field games with common noise and master equations.

[Carmona et al., 2016] Carmona, R., Delarue, F., and Lacker, D. (2016).
Mean field games with common noise.
Annals of probability: An official journal of the Institute of Mathematical Statistics, 44(6):3740–3803.

[Carmona et al., 2017] Carmona, R., Delarue, F., and Lacker, D. (2017).
Mean field games of timing and models for bank runs.
Applied Mathematics & Optimization, 76(1):217–260.

References III

[Carmona and Laurière, 2022] Carmona, R. and Laurière, M. (2022).
Convergence analysis of machine learning algorithms for the numerical solution of mean field control and games: II—the finite horizon case.
The Annals of Applied Probability, 32(6):4065–4105.

[Carmona et al., 2015] Carmona, R. A., Fouque, J. P., and Sun, L. H. (2015).
Mean field games and systemic risk.
Communications in Mathematical Sciences, 13(4):911–933.

[Cuchiero et al., 2025] Cuchiero, C., Gazzani, G., Möller, J., and Svaluto-Ferro, S. (2025).
Signature-Based Models in Finance.
In *Signature Methods in Finance: An Introduction with Computational Applications*, pages 227–264. Springer Nature Switzerland Cham.

[Dumitrescu et al., 2024] Dumitrescu, R., Leutscher, M., and Tankov, P. (2024).
Energy transition under scenario uncertainty: a mean-field game of stopping with common noise.
Mathematics and Financial Economics, 18(2):233–274.

[Gomes et al., 2023] Gomes, D., Gutierrez, J., and Laurière, M. (2023).
Machine learning architectures for price formation models with common noise.
In *2023 62nd IEEE conference on decision and control (CDC)*, pages 4345–4350. IEEE.

References IV

[Gu et al., 2024] Gu, Z., Lauriere, M., Merkel, S., and Payne, J. (2024).
Global solutions to master equations for continuous time heterogeneous agent macroeconomic models.
arXiv:2406.13726.

[Han et al., 2024] Han, J., Hu, R., and Long, J. (2024).
Learning high-dimensional McKean-Vlasov forward-backward stochastic differential equations with general distribution dependence.
SIAM Journal on Numerical Analysis, 62(1):1–24.

[Han et al., 2018] Han, J., Jentzen, A., and E, W. (2018).
Solving high-dimensional partial differential equations using deep learning.
Proceedings of the National Academy of Sciences, 115(34):8505–8510.

[Huang et al., 2006] Huang, M., Caines, P. E., and Malhamé, R. P. (2006).
Large population stochastic dynamic games: closed-loop mckean–vlasov systems and the nash certainty equivalence principle.
Communications in Information & Systems, 6(3):221–252.

[Lasry and Lions, 2007] Lasry, J.-M. and Lions, P.-L. (2007).
Mean field games.
Japanese journal of mathematics, 2(1):229–260.

[Lasry et al., 2008] Lasry, J.-M., Lions, P. L., and Guéant, O. (2008).
Application of mean field games to growth theory.

References V

[Lavigne and Tankov, 2023] Lavigne, P. and Tankov, P. (2023).
Decarbonization of financial markets: a mean-field game approach.
arXiv:2301.09163.

[Liao et al., 2019] Liao, S., Lyons, T., Yang, W., and Ni, H. (2019).
Learning stochastic differential equations using rnn with log signature features.
arXiv:1908.08286.

[Lyons and Qian, 2002] Lyons, T. and Qian, Z. (2002).
System control and rough paths.
Oxford University Press.

[Lyons et al., 2007] Lyons, T. J., Caruana, M., and Lévy, T. (2007).
Differential equations driven by rough paths.
Springer.

[Min and Hu, 2021] Min, M. and Hu, R. (2021).
Signatured deep fictitious play for mean field games with common noise.
In *International Conference on Machine Learning*, pages 7736–7747. PMLR.

[Perrin et al., 2020] Perrin, S., Pérolat, J., Laurière, M., Geist, M., Elie, R., and Pietquin, O. (2020).
Fictitious play for mean field games: Continuous time analysis and applications.
Advances in neural information processing systems, 33:13199–13213.

References VI

[Vu and Ichiba, 2025] Vu, H. and Ichiba, T. (2025).
Heterogenous macro-finance model: A mean-field game approach.
arXiv:2502.10666.

[Wu et al., 2024] Wu, Z., Laurière, M., Chua, S. J. C., Geist, M., Pietquin, O., and Mehta, A. (2024).
Population-aware online mirror descent for mean-field games by deep reinforcement learning.
In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, pages 2561–2563.

[Zhang, 2017] Zhang, J. (2017).
Backward stochastic differential equations.
Springer.