Solving dynamic portfolio selection problems via score-based diffusion models

Fengyi Yuan
(Joint work with Ahmad Aghapour and Erhan Bayraktar)

University of Michigan
fengyiy@umich.edu

SIFIN 2025
Miami, July 18th, 2025

Table of Contents

Mean-variance portfolio selection

- Static problem:

$$\sup_{w \in \mathcal{A}} \left\{ \mathbb{E}[w^\dagger R] - \frac{\gamma}{2} \mathrm{Var}(w^\dagger R) \right\}.$$

  - A simple quadratic optimization problem (with constraints), given $\mathbb{E}[R]$ and $\mathrm{Var}(R)$.
  - Numerical solvers are super efficient.

Mean-variance portfolio selection

- Dynamic problem:

$$\sup_{\vartheta \in \Theta} \mathbb{E}[(\vartheta \cdot S)_T] - \frac{\gamma}{2}\mathrm{Var}[(\vartheta \cdot S)_T]$$

  - $(\vartheta \cdot S)_T = \sum_{l=1}^{T-1} \vartheta_l^{\dagger}(S_{l+1} - S_l)$;
  - Given full information of $\{S_l\}_{l=1}^{T}$ (including marginals, transitions), it can be solved with DPP (backward induction);
  - With constraints, explicit solutions are usually unavailable, and numerical methods have high computation burdens...

Model-based v.s. Model-free

- Explicit solutions are model-based: it is optimal only with this specific model/structure; to implement the optimal strategy, one needs to know the information (e.g. parameters) of the model;
- Numerical solutions are semi-model-based: sometimes no need to assume specific model, or the model assumptions can be fairly general, but numerical solver needs samplers/oracles;
- Financial problems should be model-free: we don't know the data distribution, neither can we sample from it (adequately).

Generating data via diffusion model

Diffusion model is a type of generative model: after appropriate training, it can output samples with similar distributional properties as data.

Building blocks:

- A forward diffusion process:

$$\begin{cases} \mathrm{d}X_\tau = -\frac{1}{2}\beta(\tau)X_\tau\,\mathrm{d}\tau + \sqrt{\beta(\tau)}\,\mathrm{d}B_\tau, \\ X_0 \sim p_{\mathrm{data}}. \end{cases}$$

  - $p(\tau, \cdot)$: the density of X_τ;
  - As $\tau \to \infty$, the invariant measure is $\mathcal{N}(0, I)$ (pure noise).
  - We can not sample from p_{data}, but we can sample from $\mathcal{N}(0, I)$!

Generating data via diffusion model

How about starting from the noise and run the SDE backwardly, intuitively it gives us p_{data}!

- A reversed diffusion process:

$$\begin{cases} dY_\tau = & \Big(\frac{1}{2}\beta(\mathcal{T}-\tau)Y_\tau \\ & +\beta(\mathcal{T}-\tau)\nabla \log p(\mathcal{T}-\tau, Y_\tau)\Big)\mathrm{d}\tau + \sqrt{\beta(\mathcal{T}-\tau)}\mathrm{d}\bar{B}_\tau, \\ Y_0 \sim & \mathcal{N}(0, I), \end{cases}$$

  - Theory (FP equation): $Y_{\mathcal{T}-\tau}$ and X_τ have the same marginal distribution for any τ.
  - We do not know $\nabla_x p(\cdot,\cdot)$ (the score function), but score-matching techniques gives an approximator s_θ.
  - Theoretical error bounds: if s_θ and the score function are close, the distribution of Y_τ is close to p_{data} (in $\mathcal{W}_2$, TV...).

Diffusion model for time-series data

Challenges:

- $X = (X^1, X^2, \cdots, X^T) \sim \mathbb{P}$: data in the shape of time-series;
- We can still employ usual diffusion model, but this ignores the temporal structure (error only in $\mathcal{W}_2$);
- Dynamic problems are not stable in $\mathcal{W}_2$, but stable in $\mathcal{AW}_2$;
- We can not sample from $\mathbb{P}_{X^{1:t}}$: the conditional distribution.

Our results:
A conditional version of diffusion model, from which we sample adaptively, with $\mathcal{AW}_2$-bounds.

Diffusion model for time-series data

For $t \in \{1, 2, \cdots, T-1\}$, $x^{1:t} \in \mathbb{R}^{dt}$ consider the following forward processes:

$$\begin{cases} \mathrm{d}X_\tau^{t+1} = -X_\tau^{t+1}\mathrm{d}\tau + \sqrt{2}\mathrm{d}B_\tau^{t+1}, \\ X_0^{t+1} \sim \mathbb{P}_{x^{1:t}}. \end{cases}$$

Now the score function has three variables:

$$s_{t+1}(\tau, x^{1:t}, x) := \nabla_x p_{t+1}(\tau, x | x^{1:t}).$$

Assumption: score-matching errors are small

For any $\tau \in (0, \mathcal{T}]$ and $t = 1, 2, \cdots, T-1$, we have

$$\mathbb{E}_{X \sim p_1(\tau,\cdot)} |s_\theta^1(\tau, X) - \nabla_x \log p_1(\tau, X)|^2 \leq \varepsilon_{\text{score}}^2,$$

$$\mathbb{E}_{X^{1:t} \sim \mathbb{P}_{1:t}} \mathbb{E}_{X_\tau^{t+1} \sim p_{t+1}(\tau, \cdot | X^{1:t})} \left| s_\theta^{t+1}(\tau, X^{1:t}, X_\tau^{t+1}) - \nabla_x \log p_{t+1}(\tau, X_\tau^{t+1} | X^{1:t}) \right|^2 \leq \varepsilon_{\text{score}}^2.$$

Diffusion model for time-series data

The score-matching error gives the training objective, but it is not directly feasible:

$$\mathbb{E}_{X^{1:t}\sim\mathbb{P}_{1:t}}\mathbb{E}_{X_\tau^{t+1}\sim p_{t+1}(\tau,\cdot|X^{1:t})}\left|s_\theta^{t+1}(\tau,X^{1:t},X_\tau^{t+1})-\nabla_x \log p_{t+1}(\tau,X_\tau^{t+1}|X^{1:t})\right|^2 \leq \varepsilon_{\mathrm{score}}^2.$$

- We can not sample from $p_{t+1}(\tau,\cdot|X^{1:t})$;
- We do not know how to evaluate score function.

Diffusion model for time-series data

Denoising score-matching is feasible and equivalent:

Lemma

For any $t = 1, 2, \cdots, T-1$, ordinary score-matching is equivalent to the following:

$$\min_{\theta} \mathbb{E}_{X^{1:t+1} \sim \mathbb{P}_{1:t+1}} \mathbb{E}_{X_\tau^{t+1} \sim \phi(\tau, \cdot | X^{t+1})} \left| s_\theta^{t+1}(\tau, X^{1:t}, X_\tau^{t+1}) - \nabla_x \log \phi(\tau, X_\tau^{t+1} | X^{t+1}) \right|^2 .$$

Here, $\phi(\tau, \cdot | x_0)$ is the probability density function of the forward process, with initial condition $X_0 = x_0$. In particular,

$$\phi(\tau, x | x_0) = (2\pi(1 - e^{-2\tau}))^{-d/2} e^{-\frac{|x - x_0 e^{-\tau}|^2}{2(1 - e^{-2\tau})}} .$$

Diffusion model for time-series data

Assumption: the data distribution is reasonable

1. There exists a constant $L > 0$ such that for any $t \in \{1, 2, \cdots, T\}$, $\tau \in [0, \mathcal{T}]$ and $x^{1:t}, y^{1:t} \in \mathbb{R}^{dt}$,

$$|\nabla_x \log p_1(\tau, x) - \nabla_x \log p_1(\tau, y)| \leq L|x - y|,$$

$$|\nabla_x \log p_{t+1}(\tau, x|x^{1:t}) - \nabla_x \log p_{t+1}(\tau, y|y^{1:t})| \leq L(|x^{1:t} - y^{1:t}| + |x - y|).$$

2. For some $c > 0$, $\mathbb{E}_{\mathbb{P}_1} e^{c|X_1|} < \infty$, and for $t = \{1, 2, \cdots, T - 1\}$,

$$\sup_{x^{1:t} \in \mathbb{R}^{dt}} \mathbb{E}_{\mathbb{P}_{x^{1:t}}} e^{c|X^{t+1}|} < \infty.$$

Diffusion model for time-series data

Assumption: approximating network is good

1. There exists a constant $L > 0$ such that for any $t \in \{1, 2, \cdots, T\}$, $\tau \in [0, \mathcal{T}]$ and $x^{1:t}, y^{1:t} \in \mathbb{R}^{dt}$,

$$s_\theta^1(\tau, x) - s_\theta^1(\tau, y) \leq L|x - y|,$$
$$s_\theta^{t+1}(\tau, x^{1:t}, x) - s_\theta^{t+1}(\tau, y^{1:t}, y) \leq L(|x^{1:t} - y^{1:t}| + |x - y|).$$

2. There exists constants $C, R_0, \delta > 0$ such that for any for any $\theta \in \Theta$, $t \in \{1, 2, \cdots, T\}$, $\tau \in [0, \mathcal{T}]$, $x^{1:t} \in \mathbb{R}^{dt}$ and $x \in \mathbb{R}^d$ with $|x| > R_0$,

$$2x \cdot s_\theta^1(\tau, x) \leq -(1 + \delta)|x|^2 + C,$$
$$2x \cdot s_\theta^{t+1}(\tau, x^{1:t}, x) \leq -(1 + \delta)|x|^2 + C.$$

Diffusion model for time-series data: how to sample?

We propose an adaptive sampling scheme:

Algorithm

1. Starting from pure noise $z \sim \mathcal{N}(0, I)$, run the reversed diffusion process with approximated score function $s_\theta^1(\tau, \cdot)$ to get samples of y^1;
2. For $t \in 1, 2, \cdots, T-1$, for each generated sample $y^{1:t}$, use approximated score function $s^{t+1}(\tau, y^{1:t}, \cdot)$ and run the reversed diffusion process once to get a sample of y^{t+1};
3. Get the samples of whole path $y^{1:T}$.

The output joint measure is denoted by $\mathbb{Q}$.

Diffusion model for time-series data: main result

Theorem

If score-matching errors are small, the data distribution is reasonable, and the approximating network is good, then:

$$\mathcal{AW}_2^2(\mathbb{P},\mathbb{Q}) \leq C\left(\mathcal{T}^{\frac{5T}{2}}\varepsilon_{\text{score}}^{1/2^{T-1}} + \mathcal{T}^{\frac{5(T-1)}{2}}\alpha(\mathcal{T})^{1/2^{T-1}}\right),$$

where

$$\alpha(\mathcal{T}) = \mathcal{T}^2 e^{-\mathcal{T}} + e^{-c\mathcal{T}}.$$

Diffusion model for time-series data: remarks

- Note that $\mathcal{W}_2 \leq \mathcal{AW}_2$, this also gives a bound for classical Wasserstein metric (ignoring the temporal structure), keeping the same accumulation rate of score-matching error polynomial in $\mathcal{T}$;
- We have a new set of assumptions: drop the log-concavity of data distribution, add a dissipative condition to networks;
- Good thing: we can construct a s_θ satisfying the dissipative condition, while score-matching errors are small.

Diffusion model for time-series data: proof sketch

Step 1: A conditional distribution bound:

$$\mathcal{W}_2^2\left(\mathbb{P}_{x^{1:t}}, \mathbb{Q}^{\mathcal{T}}_{y^{1:t}}\right) \leq C\Big(\alpha(\mathcal{T}) + \mathcal{T}^2\mathcal{E}(x^{1:t})^{1/2} + \mathcal{T}^{5/2}|x^{1:t} - y^{1:t}|\Big).$$

1. Bound Wassertein by total variation:

$$\mathcal{W}_2^2(\mathbb{P}_{x^{1:t}}, \mathbb{Q}_{y^{1:t}}) \leq C\Bigg(R^2\mathrm{TV}(\mathbb{P}_{x^{1:t}}, \mathbb{Q}_{y^{1:t}}) + \mathbb{E}_{\mathbb{P}_{x^{1:t}}}\left[|X^{t+1}|^2 I_{\{|X^{t+1}|\geq R\}}\right] + \mathbb{E}_{\mathbb{Q}_{y^{1:t}}}\left[|Y^{t+1}|^2 I_{\{|Y^{t+1}|\geq R\}}\right]\Bigg).$$

2. TV term is bounded by classical approach (Girsanov's theorem);
3. $\mathbb{E}_{\mathbb{P}_{x^{1:t}}}\left[|X^{t+1}|^2 I_{\{|X^{t+1}|\geq R\}}\right]$ is bounded because data distribution is reasonable;
4. $\mathbb{E}_{\mathbb{Q}_{y^{1:t}}}\left[|Y^{t+1}|^2 I_{\{|Y^{t+1}|\geq R\}}\right]$ is bounded by analyzing reversed SDE with dissipative conditions.

Diffusion model for time-series data: proof sketch

Step 2: Constructing approximated coupling for each conditional distribution: $\pi_\varepsilon \in \Pi(\mathbb{P}_{x^{1:t}}, \mathbb{Q}_{y^{1:t}})$ is such that

$$\mathbb{E}_{\pi^{\varepsilon}_{1:t}}|X^{1:t} - Y^{1:t}|^2 \leq C\Big(\mathcal{T}^{\frac{5t}{2}}\varepsilon_{\text{score}}^{1/2^{t-1}} + \mathcal{T}^{\frac{5(t-1)}{2}}(\alpha(\mathcal{T}) + \varepsilon)^{1/2^{t-1}}\Big).$$

Step 3: Conclude by taking $\varepsilon \to 0$.

Model stability

We can now sample from a surrogate model $\mathbb{Q}$ and it is close to $\mathbb{P}$, will the optimal value of problem under $\mathbb{Q}$ near the optimal value under $\mathbb{P}$?

Theorem

Denote by $v(\mathbb{P})$ and $v(\mathbb{Q})$ the optimal values of the mean-variance problem under $\mathbb{P}$ and $\mathbb{Q}$, respectively. Then, under appropriate technical conditions, $|v(\mathbb{P}) - v(\mathbb{Q})| \leq C\mathcal{AW}_2(\mathbb{P}, \mathbb{Q})$.

Model stability

Establishing the $\mathcal{AW}_2$-stability requires dynamic programming principle!

$$v(\mathbb{P}) = \sup_{\vartheta} \left\{ \mathbb{E}_{\mathbb{P}}[(\vartheta \cdot S)_T] - \frac{\gamma}{2} \mathrm{Var}_{\mathbb{P}}[(\vartheta \cdot S)_T] \right\}$$

$v(\mathbb{P})$ is known to be time-inconsistent. We rely on its duality to quadratic hedging:

$$V(\mathbb{P}, c) = \min_{\vartheta} \mathbb{E}_{\mathbb{P}}\left[|c - (\vartheta \cdot S)_T|^2 \right].$$

$$v(\mathbb{P}) = \sup_{a \in \mathbb{R}_+} \left\{ -\frac{\gamma}{2} V\left(\mathbb{P}, \frac{1}{\gamma} + a \right) + \frac{1}{2\gamma} + a \right\}.$$

Idea: Prove the DPP and $\mathcal{AW}_2$-stability first for V, then prove the optimal multiplier a is uniform in $\mathbb{Q}$ as long as $\mathcal{AW}_2(\mathbb{P}, \mathbb{Q})$ is small.

Conclusions

Our results:

- An adaptive (training and) sampling scheme that facilitates conditional sampling, with $\mathcal{AW}_2$ error bounds;
- A model stability result that allows us to work under surrogate;
- (Not presented) A policy gradient algorithm that solves the mean-variance problem relying on surrogate model $\mathbb{Q}$;
- (Not presented) Numerical experiments confirm that the proposed algorithms have satisfactory performance on both synthetic and empirical data.

Workflow: theories

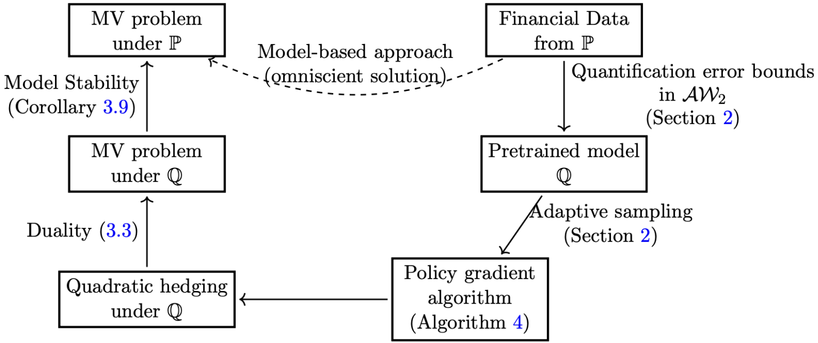

Future directions

- Relax assumptions on data distribution: with better approximation results, we expect merely Lipschitz continuity of score function will give similar bounds.
- Sample complexity, discretize error, low-dimension structure...
- Improve the algorithm, large-scale experiments.

Thank you!

Personal Webpage: https://fy-yuan.github.io

Paper on arXiv: 2507.09916.