Randomize instead of Regularize: Stable Time Integration for Poorly Conditioned Dynamical Systems

Yijun Dong
Joint work with: Benjamin Peherstorfer, Paul Schwerdtner

Courant Institute of Mathematical Sciences, New York University

SIAM CSE25, March 5, 2025, Fort Worth, Texas

NYU | COURANT INSTITUTE OF MATHEMATICAL SCIENCES

Time integration of discrete dynamical systems

Initialize with $\theta_0^* \in \mathbb{R}^p$. For each step $k = 0, 1, \cdots, K-1$, update $\theta_{k+1}^* = \theta_k^* + \frac{1}{K} F(\theta_k^*)$.

The increment function $F : \mathbb{R}^p \to \mathbb{R}^p$ involves a least-square problem:

$$F(\theta_k^*) = \arg\min_{\eta \in \mathbb{R}^p} \|J(\theta_k^*)\eta - f(\theta_k^*)\|_2^2,$$

where $J(\theta) \in \mathbb{R}^{n \times p}$ satisfies $\mathrm{rank}(J(\theta)) = p$ for all $\theta \in \mathbb{R}^p$, and $f : \mathbb{R}^p \to \mathbb{R}^n$.

- Goal: Given θ_0^* and F, we aim to **approximate the trajectory** $\theta_1^*, \cdots, \theta_K^*$.

Time integration of discrete dynamical systems

Initialize with $\theta_0^* \in \mathbb{R}^p$. For each step $k = 0, 1, \cdots, K-1$, update $\theta_{k+1}^* = \theta_k^* + \frac{1}{K} F(\theta_k^*)$.

The increment function $F : \mathbb{R}^p \to \mathbb{R}^p$ involves a least-square problem:

$$F(\theta_k^*) = \arg\min_{\eta \in \mathbb{R}^p} \| J(\theta_k^*)\eta - f(\theta_k^*) \|_2^2,$$

where $J(\theta) \in \mathbb{R}^{n \times p}$ satisfies $\mathrm{rank}(J(\theta)) = p$ for all $\theta \in \mathbb{R}^p$, and $f : \mathbb{R}^p \to \mathbb{R}^n$.

- Goal: Given θ_0^* and F, we aim to **approximate the trajectory** $\theta_1^*, \cdots, \theta_K^*$.
- Challenge: **Poorly conditioned** least-squares problems: $\kappa(J(\theta)) = \sigma_1(J(\theta))/\sigma_p(J(\theta)) \gg 1.$

Question: Under finite precision, how to approximate $\theta_1^*, \cdots, \theta_K^*$ with a low total error:

$$\mathcal{E}(\theta_1, \cdots, \theta_K) = \frac{1}{K} \sum_{i=1}^{K} \| \theta_i - \theta_i^* \|_2^2$$

Example: Neural Galerkin

- With initial condition $u_0 : \mathcal{X} \to \mathbb{R}$, aim to solve:

$$\partial_t u(t,x) = f(t,x),\ u(0,x) = u_0(x),\ \forall\ (t,x) \in [0,1] \times \mathcal{X}.$$

- Forward Euler:

$$u(t_{k+1},x) = u(t_k,x) + f(t_k,x)/K,\ u(t_0,x) = u_0(x),\ \forall\ k = 0,\cdots,K-1,\ x \in \mathcal{X}.$$

- **Nonlinear parametrization** at each time step: $u(t_k,\cdot\,) = U(\theta_k,\cdot\,)$ for every $k = 0,\cdots,K-1$ where $U(\theta_k,\cdot\,) : \mathcal{X} \to \mathbb{R}$ is a neural network parametrized by $\theta_k \in \mathbb{R}^p$.

- The **Neural Galerkin scheme** based on the **Dirac-Frenkel variational principle** seeks

$$\eta_k = \arg\min_{\eta \in \mathbb{R}^p} \left\| \nabla_\theta U(\theta_k, \mathcal{X})\eta - f(\theta_k, \mathcal{X}) \right\|_2^2$$

for all $k = 0,\cdots,K-1$ and updates $\theta_{k+1} = \theta_k + \eta_k/K$.

- The Jacobian matrix $\nabla_\theta U(\theta_k, \mathcal{X}) \in \mathbb{R}^{|\mathcal{X}|\times p}$ is **often poorly condition/numerically low-rank**.

[Lubich, 2005], [Sapsis, Lermusieux, 2009], [Du, Zaki, PhRvE2021], [Anderson, Farazmand, SISC2022], [Berman, Peherstorfer, NeurIPS2023], [Bruna, Peherstorfer, Vanden-Eijnden, JCP2024], …

Classical approach: deterministic regularization (truncated SVD)

Start with $\widehat{\theta}_0 = \theta_0^*$. For each $k = 0, \cdots, K-1$, update $\widehat{\theta}_{k+1} = \widehat{\theta}_k + \frac{1}{K} \widehat{F}(\widehat{\theta}_k)$ **well-conditioned** **regularized** increment function

Regularized integration with truncated SVD (TSVD) increments:

$$\widehat{F}(\widehat{\theta}_k) = [\![J(\widehat{\theta}_k)]\!]_r^\dagger f(\widehat{\theta}_k)$$

- $r = \text{rank}_\tau(J(\widehat{\theta}_k)) < p$ is the numerical rank w.r.t. a given precision $0 < \tau \ll 1$.

*$\text{rank}_\tau(J) = \min\{m \in [p] \,|\, \|J - [\![J]\!]_m\|_2 < \tau \|J\|_2\}$ where $[\![J]\!]_m$ denotes the optimal rank-m approximation of J from TSVD.

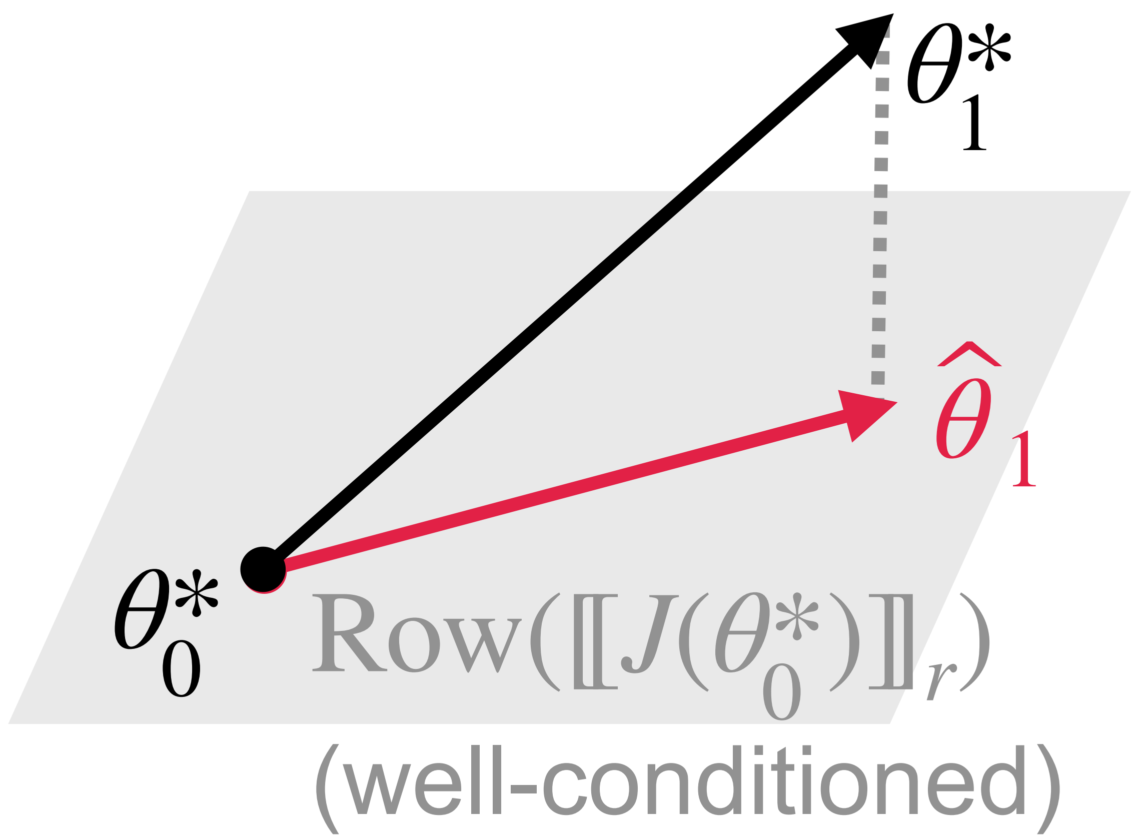

Limitation of deterministic regularization

Regularized integration with TSVD increments: $\widehat{\theta}_{k+1} = \widehat{\theta}_k + \frac{1}{K}[\![J(\widehat{\theta}_k)]\!]_r^\dagger f(\widehat{\theta}_k)$ for $k = 0, \cdots, K-1$

Deterministic regularization leads to accumulation of bias

- Consider toy dynamics with a constant $J(\cdot) \equiv J$ that admits a low numerical rank $\mathrm{rank}_\tau(J) = r < p$.
- Let $P_r \in \mathbb{R}^{p \times p}$ be the orthogonal projector onto $\mathrm{Row}([\![J]\!]_r)$.
- If $\left\| (I_p - P_r)\left(\frac{1}{K}\sum_{j=0}^{k-1} f(\theta_j^*)\right) \right\|_2 \geq b_{-r}$ for every $k = 0, \cdots, K-1$, then

$$\mathscr{E}(\widehat{\theta}_1, \cdots, \widehat{\theta}_K) \geq \frac{b_{-r}^2}{3\sigma_{r+1}^2(J)}$$

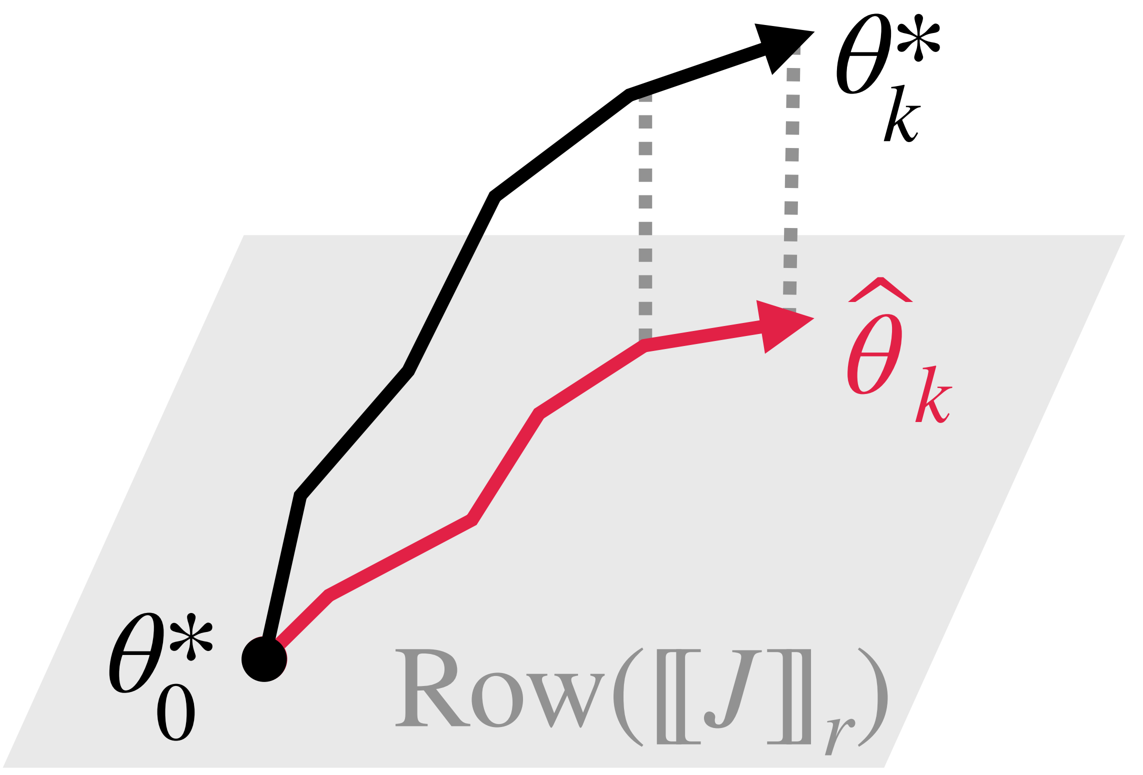

Limitation of deterministic regularization

Regularized integration with TSVD increments: $\hat{\theta}_{k+1} = \hat{\theta}_k + \frac{1}{K}[\![J(\hat{\theta}_k)]\!]_r^\dagger f(\hat{\theta}_k)$ for $k = 0, \cdots, K-1$

Deterministic regularization leads to accumulation of bias

- Consider toy dynamics with a constant $J(\cdot) \equiv J$ that admits a low numerical rank $\mathrm{rank}_\tau(J) = r < p$.
- Let $P_r \in \mathbb{R}^{p \times p}$ be the orthogonal projector onto $\mathrm{Row}([\![J]\!]_r)$.
- If $\left\| (I_p - P_r)\left(\frac{1}{K}\sum_{j=0}^{k-1} f(\theta_j^*)\right) \right\|_2 \geq b_{-r}$ for every $k = 0, \cdots, K-1$, then

$$\mathscr{E}(\hat{\theta}_1, \cdots, \hat{\theta}_K) \geq \frac{b_{-r}^2}{3\sigma_{r+1}^2(J)}$$

Lower bound of error due to accumulation of bias over time

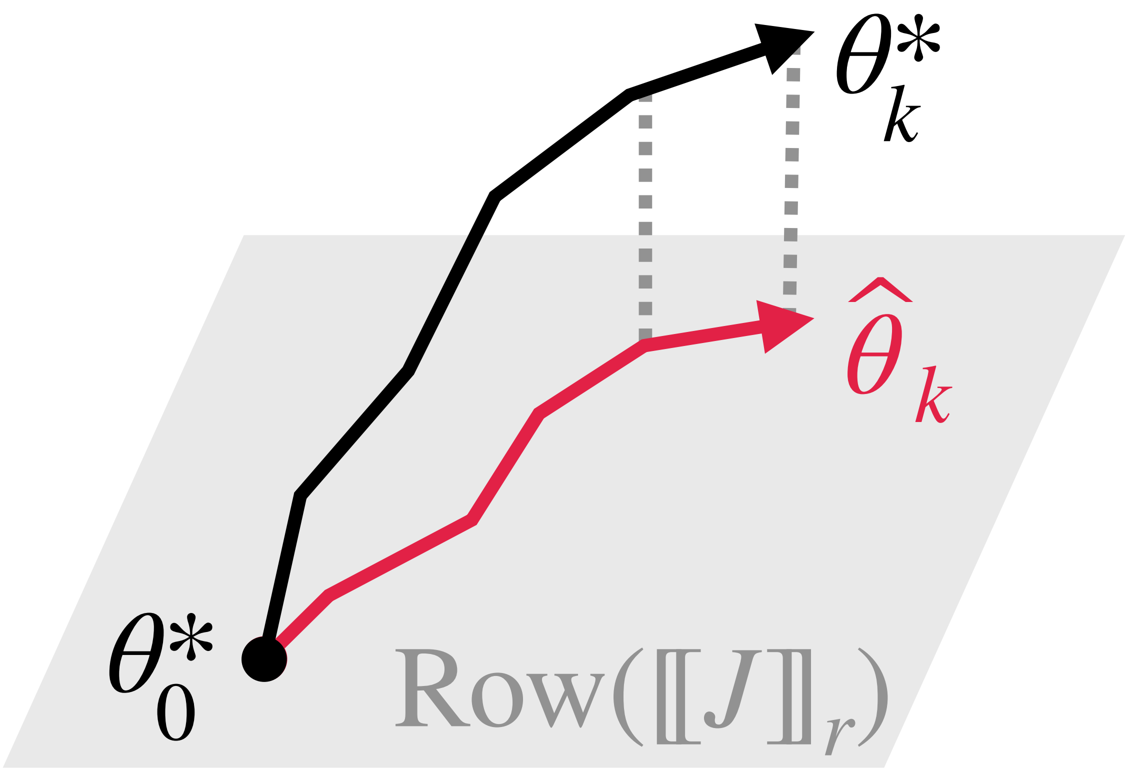

Limitation of deterministic regularization

Regularized integration with TSVD increments: $\widehat{\theta}_{k+1} = \widehat{\theta}_k + \frac{1}{K}[\![J(\widehat{\theta}_k)]\!]_r^\dagger f(\widehat{\theta}_k)$ for $k = 0, \cdots, K-1$

Deterministic regularization leads to accumulation of bias

- Consider toy dynamics with a constant $J(\cdot) \equiv J$ that admits a low numerical rank $\text{rank}_\tau(J) = r < p$.
- Let $P_r \in \mathbb{R}^{p \times p}$ be the orthogonal projector onto $\text{Row}([\![J]\!]_r)$.
- If $\left\| (I_p - P_r)\left(\frac{1}{K}\sum_{j=0}^{k-1} f(\theta_j^*)\right) \right\|_2 \geq b_{-r}$ for every $k = 0, \cdots, K-1$, then

$$\mathcal{E}(\widehat{\theta}_1, \cdots, \widehat{\theta}_K) \geq \frac{b_{-r}^2}{3\sigma_{r+1}^2(J)}$$

Lower bound of error due to accumulation of bias over time

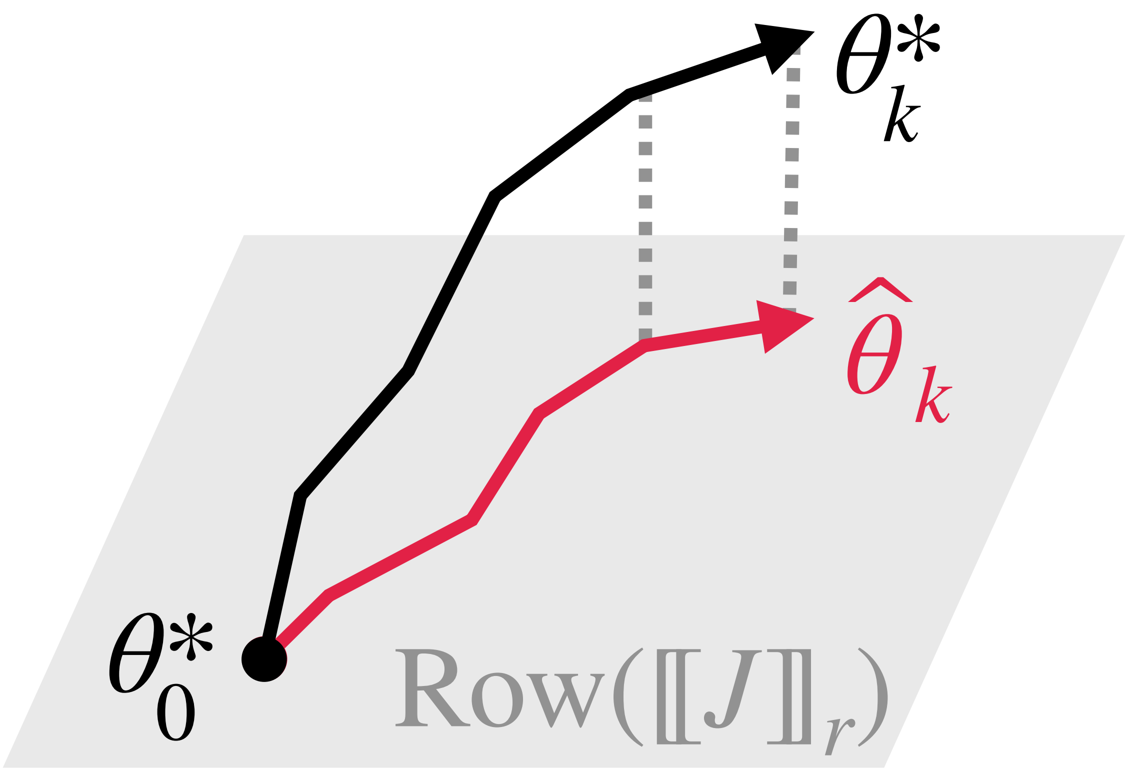

Accumulation of bias also occurs under other deterministic regularizations like **Tikhonov regularization**:

$$\widehat{F}(\widehat{\theta}_k) = \arg\min_{\eta \in \mathbb{R}^p} \|J(\widehat{\theta}_k)\eta - f(\widehat{\theta}_k)\|_2^2 + \alpha\|\eta\|_2^2$$

Randomized time integration

Start with $\widetilde{\theta}_0 = \theta_0^*$. For each $k = 0, \cdots, K-1$,

update $\widetilde{\theta}_{k+1} = \widetilde{\theta}_k + \frac{1}{K}\widetilde{\eta}_k$ **well-conditioned randomized** increments

Randomized increments: $\widetilde{\eta}_k = \frac{1}{q}\sum_{i=1}^{q} \widetilde{F}(\widetilde{\theta}_k; \Gamma_{k,i})$

- Randomized increment function: $\widetilde{F} : \mathbb{R}^p \times \mathcal{S} \to \mathbb{R}^p$
- $\{\Gamma_{k,i} \sim P_S \,|\, k = 0, \cdots, K-1,\ i \in [q]\}$ are drawn i.i.d. from a given distribution P_S supported on some set $\mathcal{S}$ (e.g., Gaussian random matrices).
- Local sample size: $q \in \mathbb{N}$

[Dong, Schwerdtner, Peherstorfer, "Randomize instead of Regularize: Stable Time Integration for Poorly Conditioned Dynamical Systems", in preparation.]

Randomized time integration

Start with $\widetilde{\theta}_0 = \theta_0^*$. For each $k = 0, \cdots, K-1$,

update $\widetilde{\theta}_{k+1} = \widetilde{\theta}_k + \frac{1}{K}\widetilde{\eta}_k$ — **well-conditioned randomized** increments

Randomized increments: $\widetilde{\eta}_k = \frac{1}{q}\sum_{i=1}^{q} \widetilde{F}(\widetilde{\theta}_k; \Gamma_{k,i})$

- Randomized increment function: $\widetilde{F} : \mathbb{R}^p \times \mathcal{S} \to \mathbb{R}^p$
- $\{\Gamma_{k,i} \sim P_S \,|\, k = 0, \cdots, K-1,\ i \in [q]\}$ are drawn i.i.d. from a given distribution P_S supported on some set $\mathcal{S}$ (e.g., Gaussian random matrices).
- Local sample size: $q \in \mathbb{N}$

[Dong, Schwerdtner, Peherstorfer, "Randomize instead of Regularize: Stable Time Integration for Poorly Conditioned Dynamical Systems", in preparation.]

<u>Assumption</u>: **unbiased** $\widetilde{F}$ with **low variance**

For all $\theta \in \mathbb{R}^p$, there exists some $C_v \geq 0$ such that

(1) $\mathbb{E}_\Gamma[\widetilde{F}(\theta, \Gamma)] = F(\theta)$

(2) $\mathbb{E}_\Gamma[\|\widetilde{F}(\theta, \Gamma) - F(\theta)\|_2^2] \leq C_v \|F(\theta)\|_2^2$

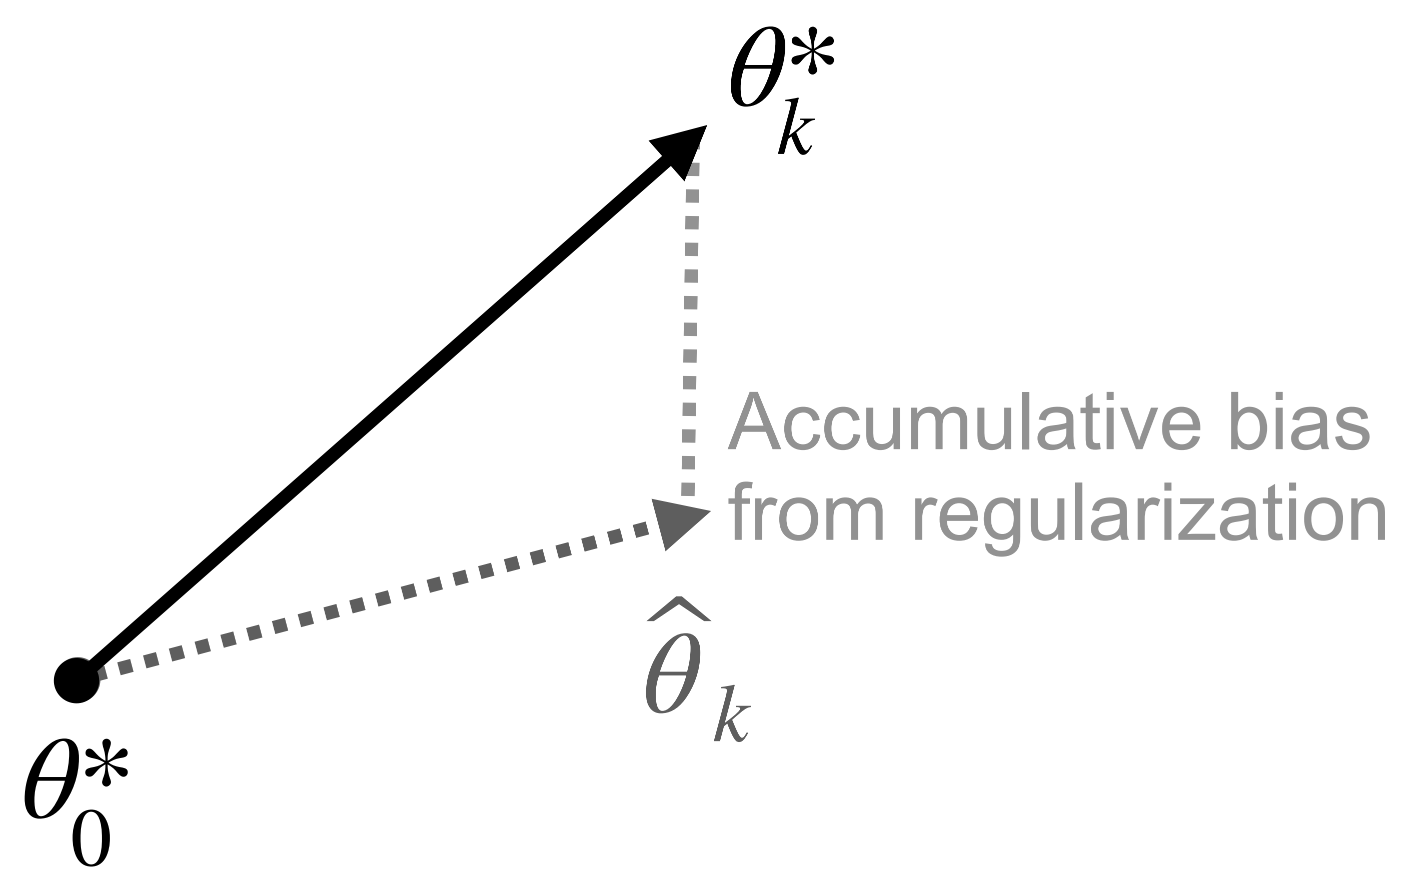

Randomized time integration

Start with $\widetilde{\theta}_0 = \theta_0^*$. For each $k = 0, \cdots, K-1$,

update $\widetilde{\theta}_{k+1} = \widetilde{\theta}_k + \boxed{\frac{1}{K}\widetilde{\eta}_k}$ **well-conditioned randomized** increments

Randomized increments: $\widetilde{\eta}_k = \frac{1}{q}\sum_{i=1}^{q} \widetilde{F}(\widetilde{\theta}_k; \Gamma_{k,i})$

- Randomized increment function: $\widetilde{F} : \mathbb{R}^p \times \mathcal{S} \to \mathbb{R}^p$
- $\{\Gamma_{k,i} \sim P_S \,|\, k = 0, \cdots, K-1,\ i \in [q]\}$ are drawn i.i.d. from a given distribution P_S supported on some set $\mathcal{S}$ (e.g., Gaussian random matrices).
- Local sample size: $q \in \mathbb{N}$

[Dong, Schwerdtner, Peherstorfer, "Randomize instead of Regularize: Stable Time Integration for Poorly Conditioned Dynamical Systems", in preparation.]

<u>Assumption</u>: **unbiased** $\widetilde{F}$ with **low variance**

For all $\theta \in \mathbb{R}^p$, there exists some $C_v \geq 0$ such that

(1) $\mathbb{E}_\Gamma[\widetilde{F}(\theta, \Gamma)] = F(\theta)$

(2) $\mathbb{E}_\Gamma[\|\widetilde{F}(\theta, \Gamma) - F(\theta)\|_2^2] \leq C_v \|F(\theta)\|_2^2$

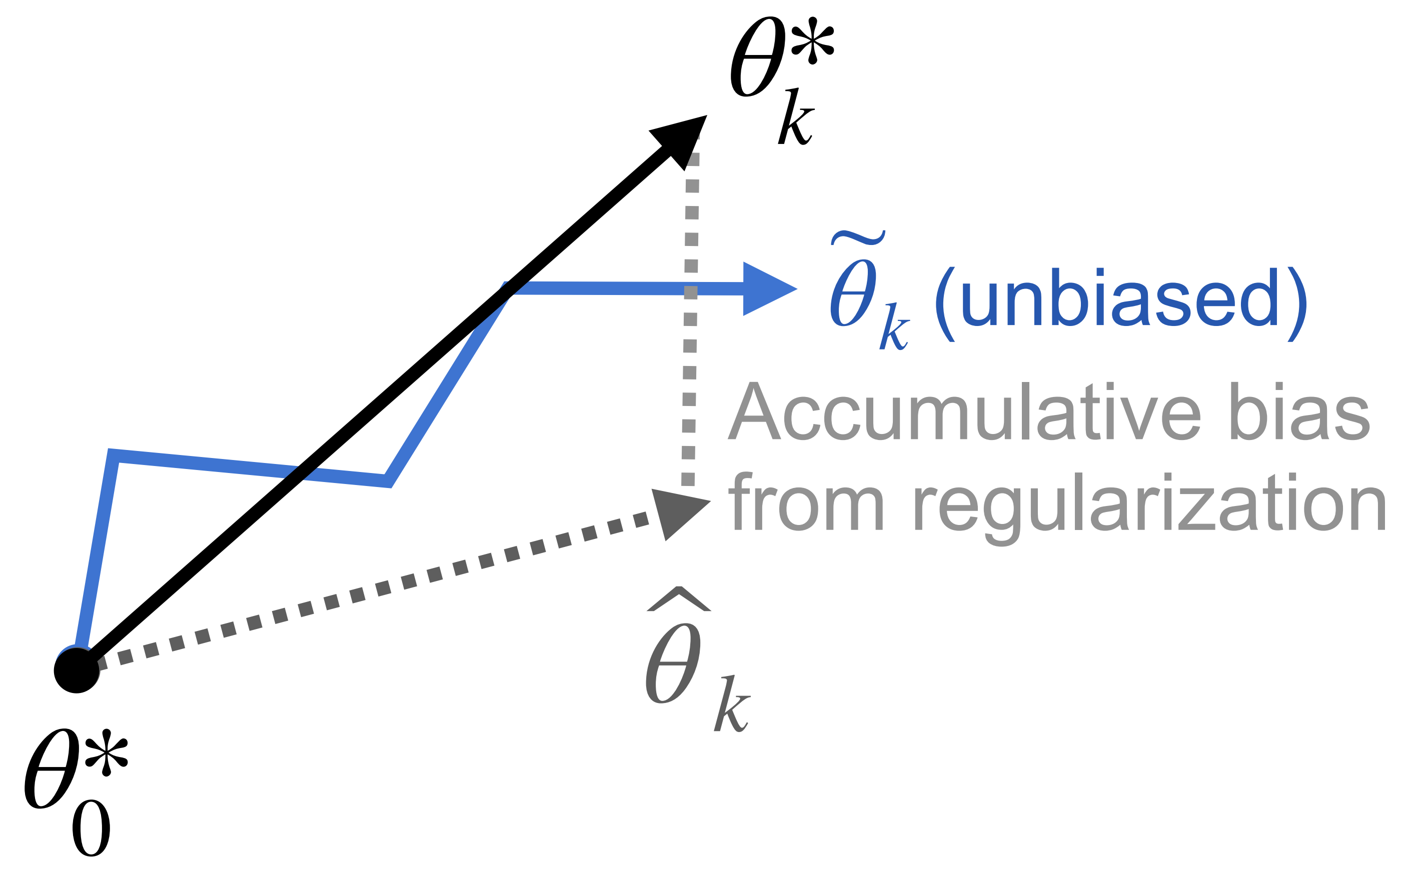

Randomized time integration

Start with $\widetilde{\theta}_0 = \theta_0^*$. For each $k = 0, \cdots, K-1$,

update $\widetilde{\theta}_{k+1} = \widetilde{\theta}_k + \boxed{\frac{1}{K}\widetilde{\eta}_k}$ **well-conditioned randomized** increments

Randomized increments: $\widetilde{\eta}_k = \frac{1}{q}\sum_{i=1}^{q} \widetilde{F}(\widetilde{\theta}_k; \Gamma_{k,i})$

- Randomized increment function: $\widetilde{F} : \mathbb{R}^p \times \mathcal{S} \to \mathbb{R}^p$
- $\{\Gamma_{k,i} \sim P_S \,|\, k = 0, \cdots, K-1,\ i \in [q]\}$ are drawn i.i.d. from a given distribution P_S supported on some set $\mathcal{S}$ (e.g., Gaussian random matrices).
- Local sample size: $q \in \mathbb{N}$

[Dong, Schwerdtner, Peherstorfer, "Randomize instead of Regularize: Stable Time Integration for Poorly Conditioned Dynamical Systems", in preparation.]

Assumption: **unbiased** $\widetilde{F}$ with **low variance**

For all $\theta \in \mathbb{R}^p$, there exists some $C_v \geq 0$ such that

(1) $\mathbb{E}_\Gamma[\widetilde{F}(\theta, \Gamma)] = F(\theta)$

(2) $\mathbb{E}_\Gamma[\|\widetilde{F}(\theta, \Gamma) - F(\theta)\|_2^2] \leq C_v \|F(\theta)\|_2^2$

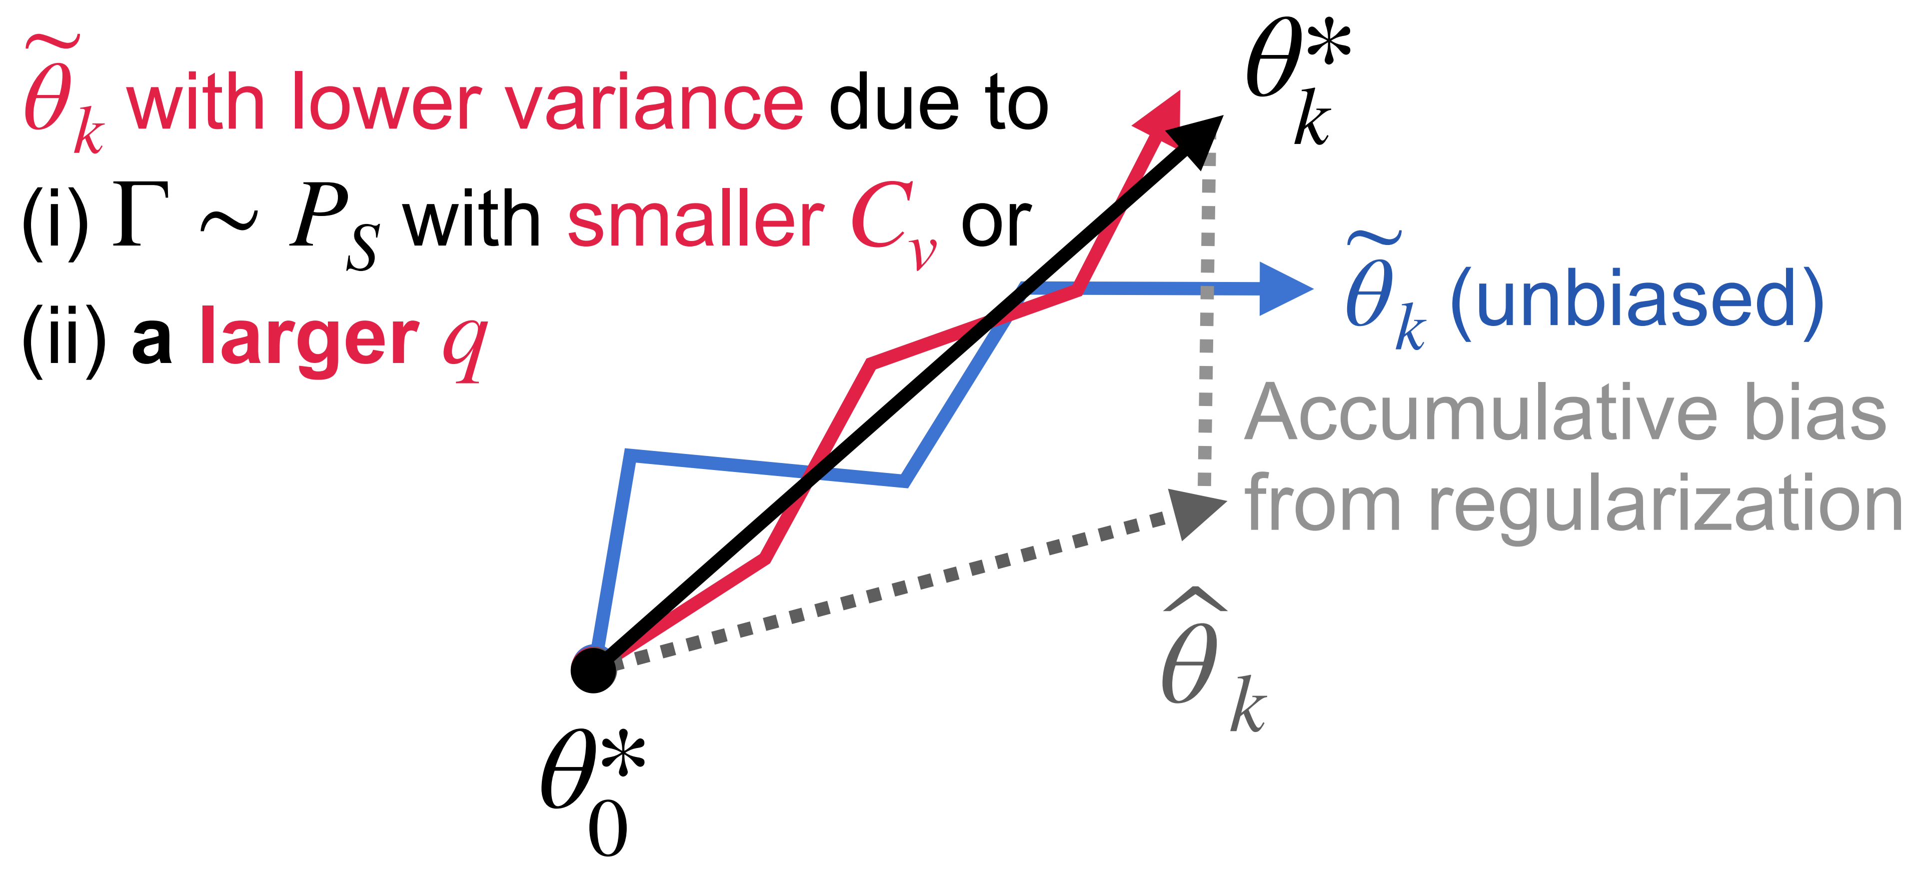

Randomized time integration: convergence

Assumptions:

- **Unbiased** $\widetilde{F}$ with **low variance**: $\mathbb{E}_{\Gamma}[\widetilde{F}(\theta,\Gamma)] = F(\theta)$, and (2) $\mathbb{E}_{\Gamma}[\|\widetilde{F}(\theta,\Gamma) - F(\theta)\|_2^2] \leq C_v \|F(\theta)\|_2^2$.
- Lipschitz and bounded F: (1) $\|F(\theta) - F(\theta')\|_2 \leq L_F \|\theta - \theta'\|_2 \ \forall \ \theta, \theta'$, and (2) $\|F(\theta)\|_2 \leq B_F \|\theta\|_2 \ \forall \ \theta$.

Theorem:

Start with $\widetilde{\theta}_0 = \theta_0^*$. Let $\widetilde{\theta}_{k+1} = \widetilde{\theta}_k + \frac{1}{K}\widetilde{\eta}_k$ where $\widetilde{\eta}_k = \frac{1}{q}\sum_{i=1}^{q} \widetilde{F}(\widetilde{\theta}_k; \Gamma_{k,i})$ for each $k = 0, \cdots, K-1$.

If $\max_{0 \leq k \leq K} \|\widetilde{\theta}_k\|_2 = B_\theta$, then

$$\mathbb{E}\left[\mathscr{E}(\widetilde{\theta}_1, \cdots, \widetilde{\theta}_K)\right] \leq \frac{B_F^2 B_\theta^2 e^{2L_F}}{K} \frac{C_v}{q}.$$

$\widetilde{\theta}_k$ with lower variance due to
(i) $\Gamma \sim P_S$ with smaller C_v or
(ii) a **larger** q

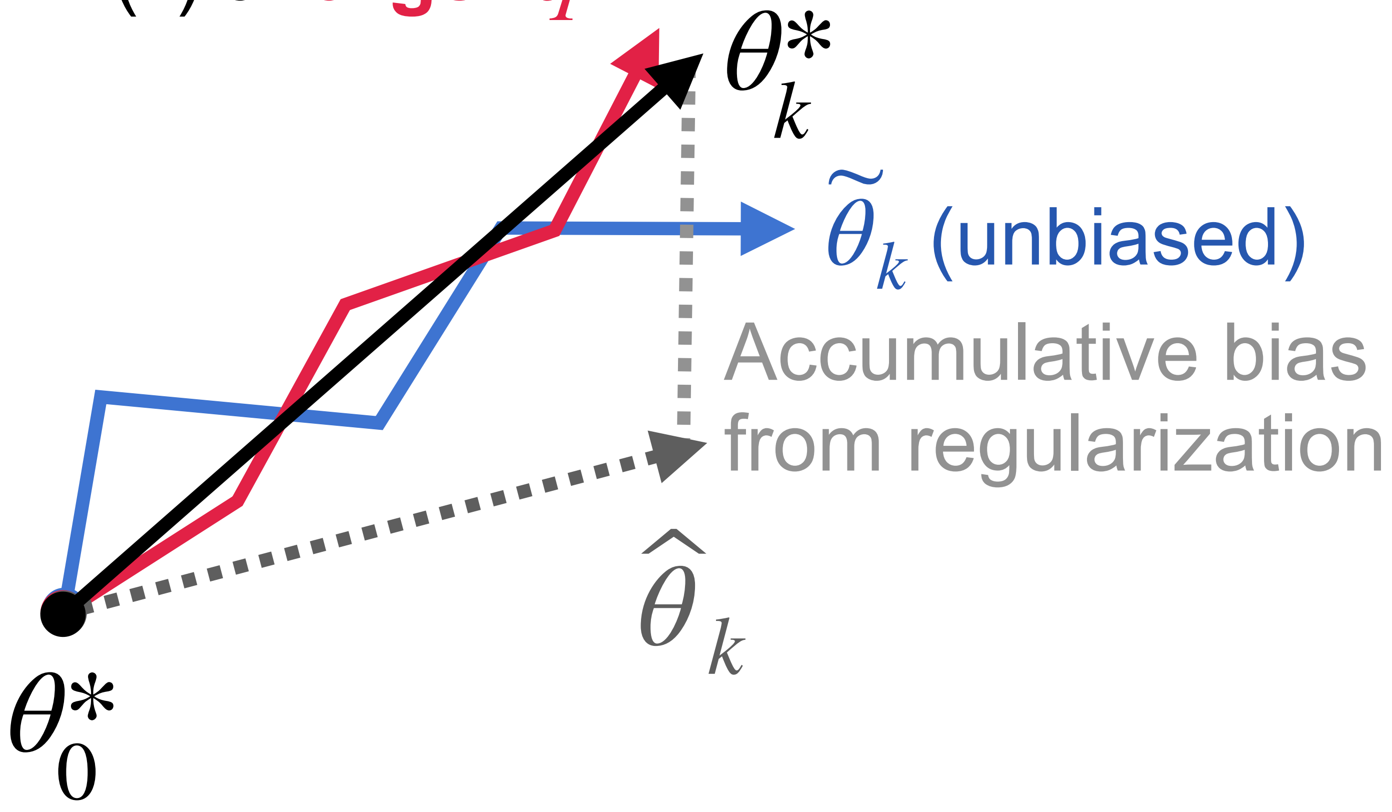

Randomized time integration: convergence

Assumptions:

- **Unbiased** $\widetilde{F}$ with **low variance**: $\mathbb{E}_\Gamma[\widetilde{F}(\theta,\Gamma)] = F(\theta)$, and (2) $\mathbb{E}_\Gamma[\|\widetilde{F}(\theta,\Gamma) - F(\theta)\|_2^2] \le C_v \|F(\theta)\|_2^2$.
- Lipschitz and bounded F: (1) $\|F(\theta) - F(\theta')\|_2 \le L_F \|\theta - \theta'\|_2 \ \forall\ \theta, \theta'$, and (2) $\|F(\theta)\|_2 \le B_F \|\theta\|_2 \ \forall\ \theta$.

Theorem:

Start with $\widetilde{\theta}_0 = \theta_0^*$. Let $\widetilde{\theta}_{k+1} = \widetilde{\theta}_k + \frac{1}{K}\widetilde{\eta}_k$ where

$$\widetilde{\eta}_k = \frac{1}{q}\sum_{i=1}^{q} \widetilde{F}(\widetilde{\theta}_k; \Gamma_{k,i}) \text{ for each } k = 0, \cdots, K-1.$$

If $\max_{0 \le k \le K} \|\widetilde{\theta}_k\|_2 = B_\theta$, then

$$\mathbb{E}\left[\mathscr{E}(\widetilde{\theta}_1, \cdots, \widetilde{\theta}_K)\right] \le \frac{B_F^2 B_\theta^2 e^{2L_F}}{K} \frac{C_v}{q}.$$

$\widetilde{\theta}_k$ with lower variance due to
(i) $\Gamma \sim P_S$ with smaller C_v or
(ii) a **larger** q

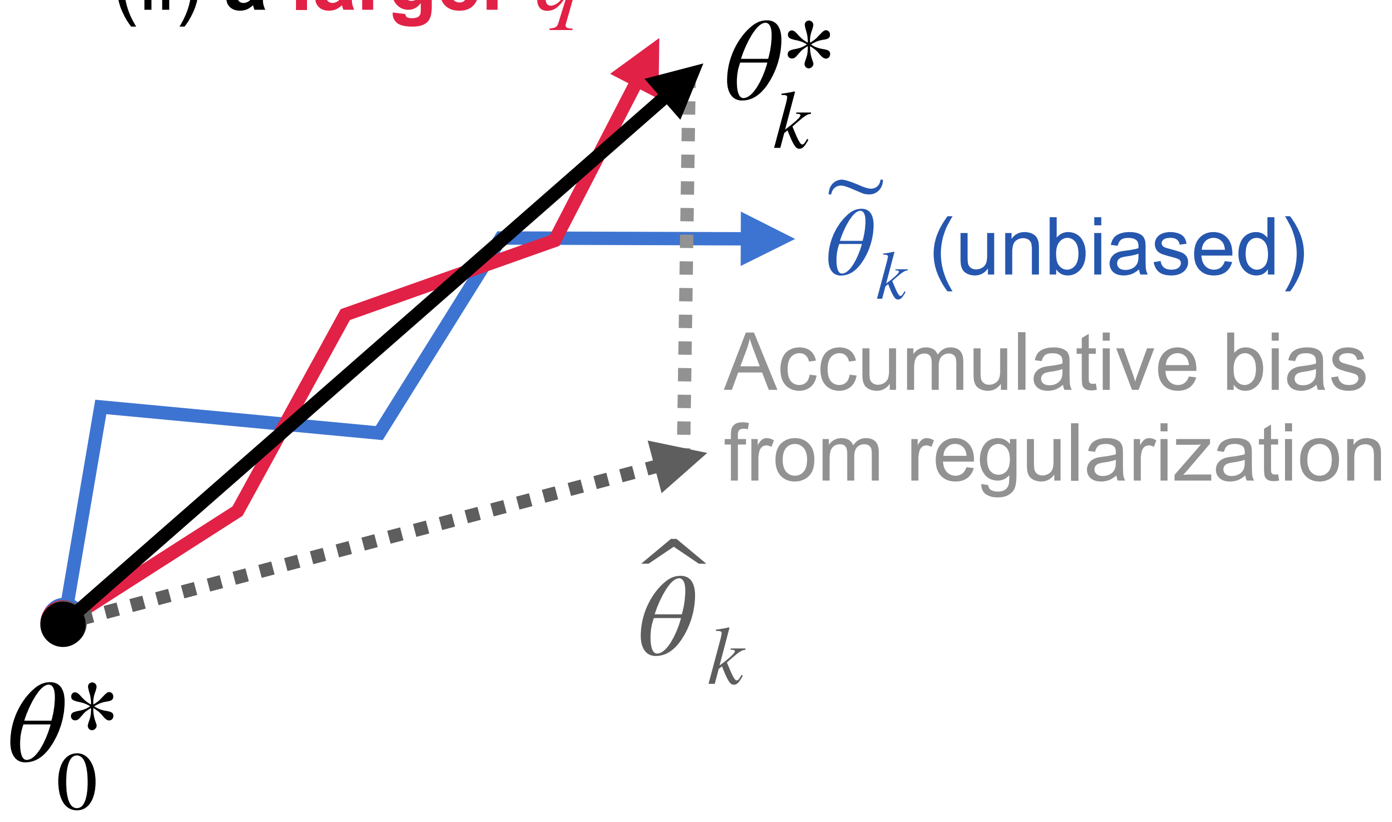

The expected total error converges as q increases.

Randomized time integration for least-squares increments

Question: How to construct an **unbiased** $\widetilde{F}$ with **low variance** for $F(\theta_k^*) = \arg\min_{\eta \in \mathbb{R}^p} \|J(\theta_k^*)\eta - f(\theta_k^*)\|_2^2$

s.t. $\mathbb{E}_\Gamma[\widetilde{F}(\theta, \Gamma)] = F(\theta)$; $\mathbb{E}_\Gamma[\|\widetilde{F}(\theta, \Gamma) - F(\theta)\|_2^2] \leq C_v \|F(\theta)\|_2^2$; and $\widetilde{F}$ **is well-conditioned**?

Randomized time integration for least-squares increments

Question: How to construct an **unbiased** $\widetilde{F}$ with **low variance** for $F(\theta_k^*) = \arg\min_{\eta \in \mathbb{R}^p} \|J(\theta_k^*)\eta - f(\theta_k^*)\|_2^2$

s.t. $\mathbb{E}_\Gamma[\widetilde{F}(\theta, \Gamma)] = F(\theta)$; $\mathbb{E}_\Gamma[\|\widetilde{F}(\theta, \Gamma) - F(\theta)\|_2^2] \leq C_v \|F(\theta)\|_2^2$; and $\widetilde{F}$ **is well-conditioned**?

Randomized least-squares increments:

Inspired by [Berman, Peherstorfer, NeurIPS2023]

$$\widetilde{F}(\theta; \Gamma) = \frac{p}{m} \Gamma \arg\min_{v \in \mathbb{R}^m} \| \ (J(\theta)\Gamma) \ v - f(\theta)\|_2^2 \quad (0 < m \leq p)$$

- $\Gamma \in \mathbb{R}^{p \times m}$ is a random matrix drawn from a **rotation-invariant** distribution with $\mathrm{rank}(\Gamma) = m$ **almost surely**.
  - Gaussian random matrix: $\Gamma = G$ with $G_{ij} \sim \mathcal{N}(0, 1/m)$ i.i.d.
  - Random unitary embedding: $\Gamma = \mathrm{ortho}(G)$

Randomized time integration for least-squares increments

Question: How to construct an **unbiased** $\widetilde{F}$ with **low variance** for $F(\theta_k^*) = \arg\min_{\eta \in \mathbb{R}^p} \|J(\theta_k^*)\eta - f(\theta_k^*)\|_2^2$

s.t. $\mathbb{E}_\Gamma[\widetilde{F}(\theta, \Gamma)] = F(\theta)$; $\mathbb{E}_\Gamma[\|\widetilde{F}(\theta, \Gamma) - F(\theta)\|_2^2] \leq C_v \|F(\theta)\|_2^2$; and $\widetilde{F}$ **is well-conditioned**?

Randomized least-squares increments: Inspired by [Berman, Peherstorfer, NeurIPS2023]

$$\widetilde{F}(\theta; \Gamma) = \frac{p}{m} \Gamma \arg\min_{v \in \mathbb{R}^m} \| (J(\theta)\Gamma)\, v - f(\theta)\|_2^2 \quad (0 < m \leq p)$$

- $\Gamma \in \mathbb{R}^{p \times m}$ is a random matrix drawn from a **rotation-invariant** distribution with $\mathrm{rank}(\Gamma) = m$ **almost surely**.
  - Gaussian random matrix: $\Gamma = G$ with $G_{ij} \sim \mathcal{N}(0, 1/m)$ i.i.d.
  - Random unitary embedding: $\Gamma = \mathrm{ortho}(G)$

$\widetilde{F}$ **is unbiased** with **low variance** associated with $C_v = \max\left\{1, \left(\frac{p-m}{m}\right)^2\right\}$.

Larger $m \Rightarrow$ smaller C_v & lower variance

Randomized time integration for least-squares increments

Question: How to construct an **unbiased** $\widetilde{F}$ with **low variance** for $F(\theta_k^*) = \arg\min_{\eta \in \mathbb{R}^p} \|J(\theta_k^*)\eta - f(\theta_k^*)\|_2^2$

s.t. $\mathbb{E}_\Gamma[\widetilde{F}(\theta, \Gamma)] = F(\theta)$; $\mathbb{E}_\Gamma[\|\widetilde{F}(\theta, \Gamma) - F(\theta)\|_2^2] \leq C_v \|F(\theta)\|_2^2$; and $\widetilde{F}$ **is well-conditioned**?

Randomized least-squares increments: Inspired by [Berman, Peherstorfer, NeurIPS2023]

$$\widetilde{F}(\theta; \Gamma) = \frac{p}{m} \Gamma \arg\min_{v \in \mathbb{R}^m} \| (J(\theta)\Gamma)\, v - f(\theta)\|_2^2 \quad (0 < m \leq p)$$

- $\Gamma \in \mathbb{R}^{p \times m}$ is a random matrix drawn from a **rotation-invariant** distribution with $\mathrm{rank}(\Gamma) = m$ **almost surely**.
  - Gaussian random matrix: $\Gamma = G$ with $G_{ij} \sim \mathcal{N}(0, 1/m)$ i.i.d.
  - Random unitary embedding: $\Gamma = \mathrm{ortho}(G)$

$\widetilde{F}$ **is unbiased** with **low variance** associated with $C_v = \max\left\{1, \left(\frac{p-m}{m}\right)^2\right\}$.

Larger $m \Rightarrow$ smaller C_v & lower variance

$\widetilde{F}$ **has better conditioning than** F: $\kappa(J(\theta)\Gamma) \leq \kappa(J(\theta))$ for Γ with orthonormal columns.

Is $\widetilde{F}$ well-conditioned?

Sketching improves conditioning

<u>Theorem</u>: For $J \in \mathbb{R}^{n\times p}$ with singular values $\sigma_1(J) \geq \cdots \geq \sigma_p(J) > 0$ and a random unitary embedding $\Gamma \in \mathbb{R}^{p\times m}$, if $\log(p) \ll m \ll p$ and $l = \Omega(m)$, with probability at least 0.98,

$$\kappa(J\Gamma) \leq \frac{\sigma_1(J)}{\sigma_l(J)}\, O\left(\sqrt{\frac{p}{m}}\right)\sqrt{\frac{1+O(\sqrt{m/p})}{1-O(\log(l)/m)}} \lesssim \frac{\sigma_1(J)}{\sigma_l(J)}\sqrt{\frac{p}{m}}$$

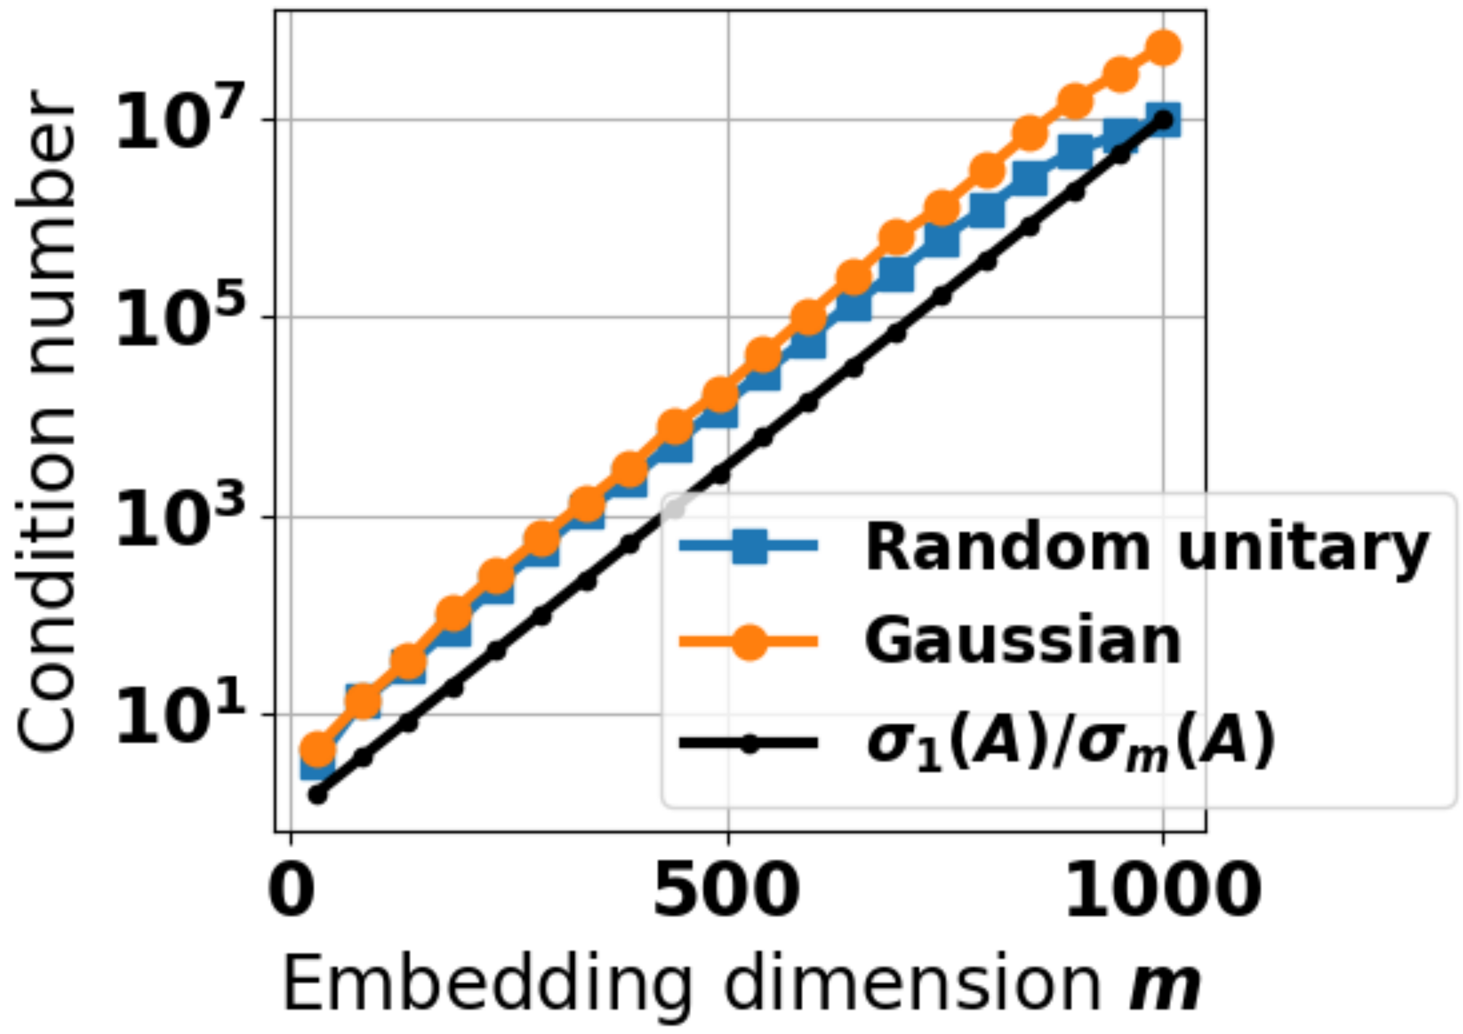

Sketching improves conditioning

Theorem: For $J \in \mathbb{R}^{n\times p}$ with singular values $\sigma_1(J) \geq \cdots \geq \sigma_p(J) > 0$ and a random unitary embedding $\Gamma \in \mathbb{R}^{p\times m}$, if $\log(p) \ll m \ll p$ and $l = \Omega(m)$, with probability at least 0.98,

$$\kappa(J\Gamma) \leq \frac{\sigma_1(J)}{\sigma_l(J)} \, O\left(\sqrt{\frac{p}{m}}\right) \sqrt{\frac{1 + O(\sqrt{m/p})}{1 - O(\log(l)/m)}} \lesssim \frac{\sigma_1(J)}{\sigma_l(J)} \sqrt{\frac{p}{m}}$$

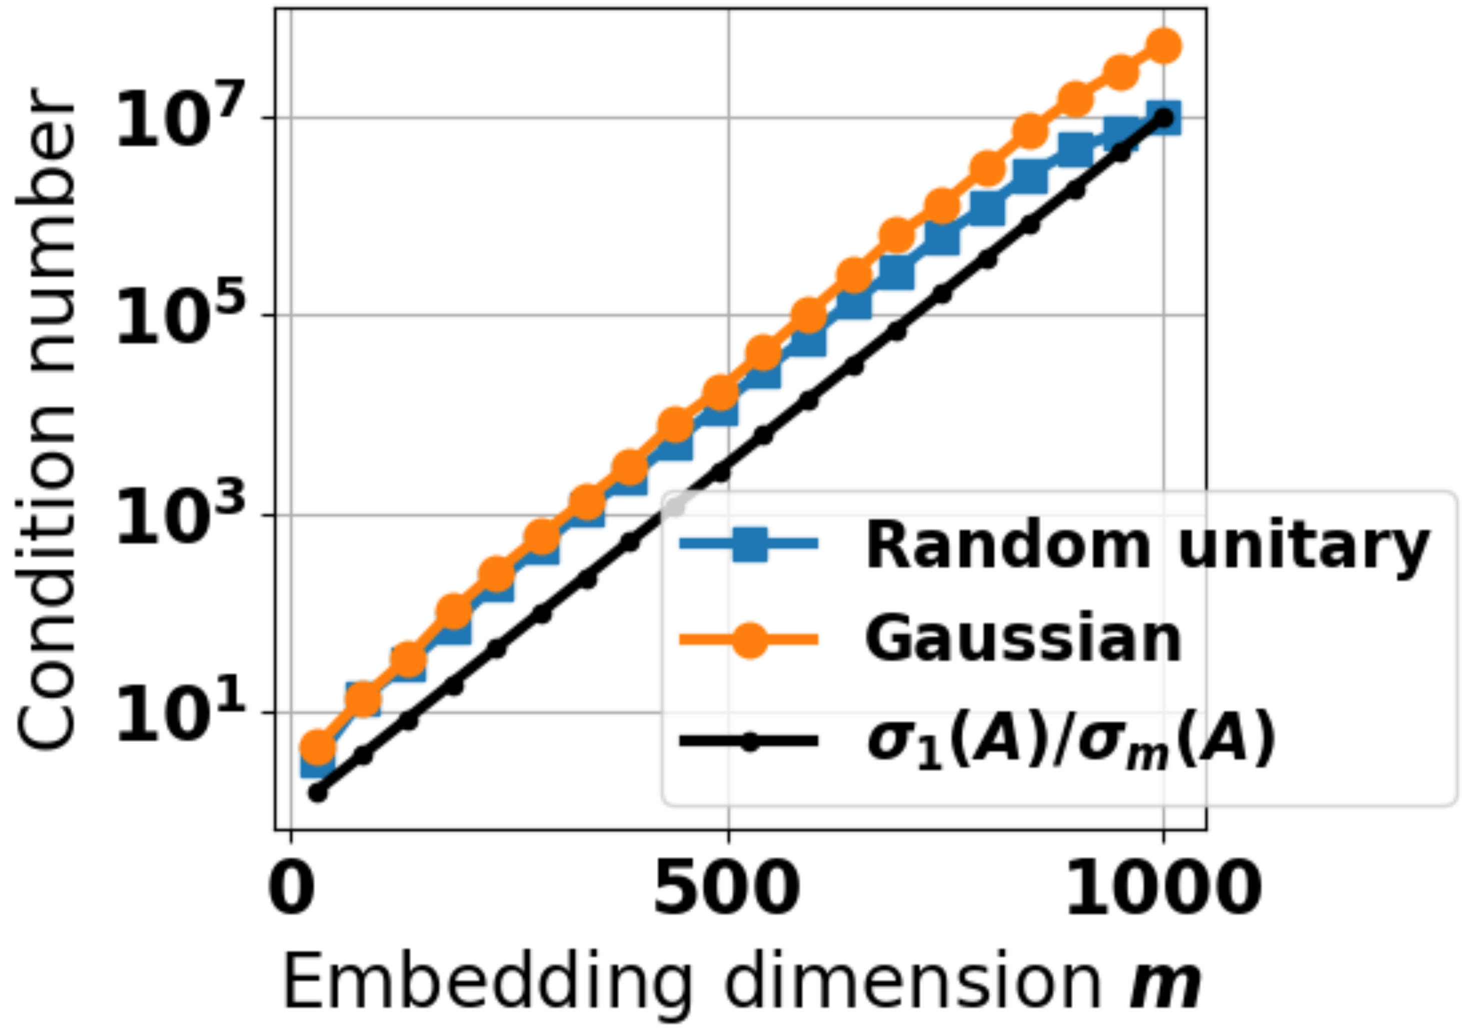

For any poorly conditioned $J(\theta)$, taking an embedding dimension m slightly smaller than numerical rank $r = \mathrm{rank}_\tau(J(\theta))$ brings a **well-conditioned** matrix with

$$\kappa(J(\theta)\Gamma) \leq 1/\tau$$

Sketching improves conditioning

Theorem: For $J \in \mathbb{R}^{n \times p}$ with singular values $\sigma_1(J) \geq \cdots \geq \sigma_p(J) > 0$ and a random unitary embedding $\Gamma \in \mathbb{R}^{p \times m}$, if $\log(p) \ll m \ll p$ and $l = \Omega(m)$, with probability at least 0.98,

$$\kappa(J\Gamma) \leq \frac{\sigma_1(J)}{\sigma_l(J)} O\left(\sqrt{\frac{p}{m}}\right) \sqrt{\frac{1 + O(\sqrt{m/p})}{1 - O(\log(l)/m)}} \lesssim \frac{\sigma_1(J)}{\sigma_l(J)} \sqrt{\frac{p}{m}}$$

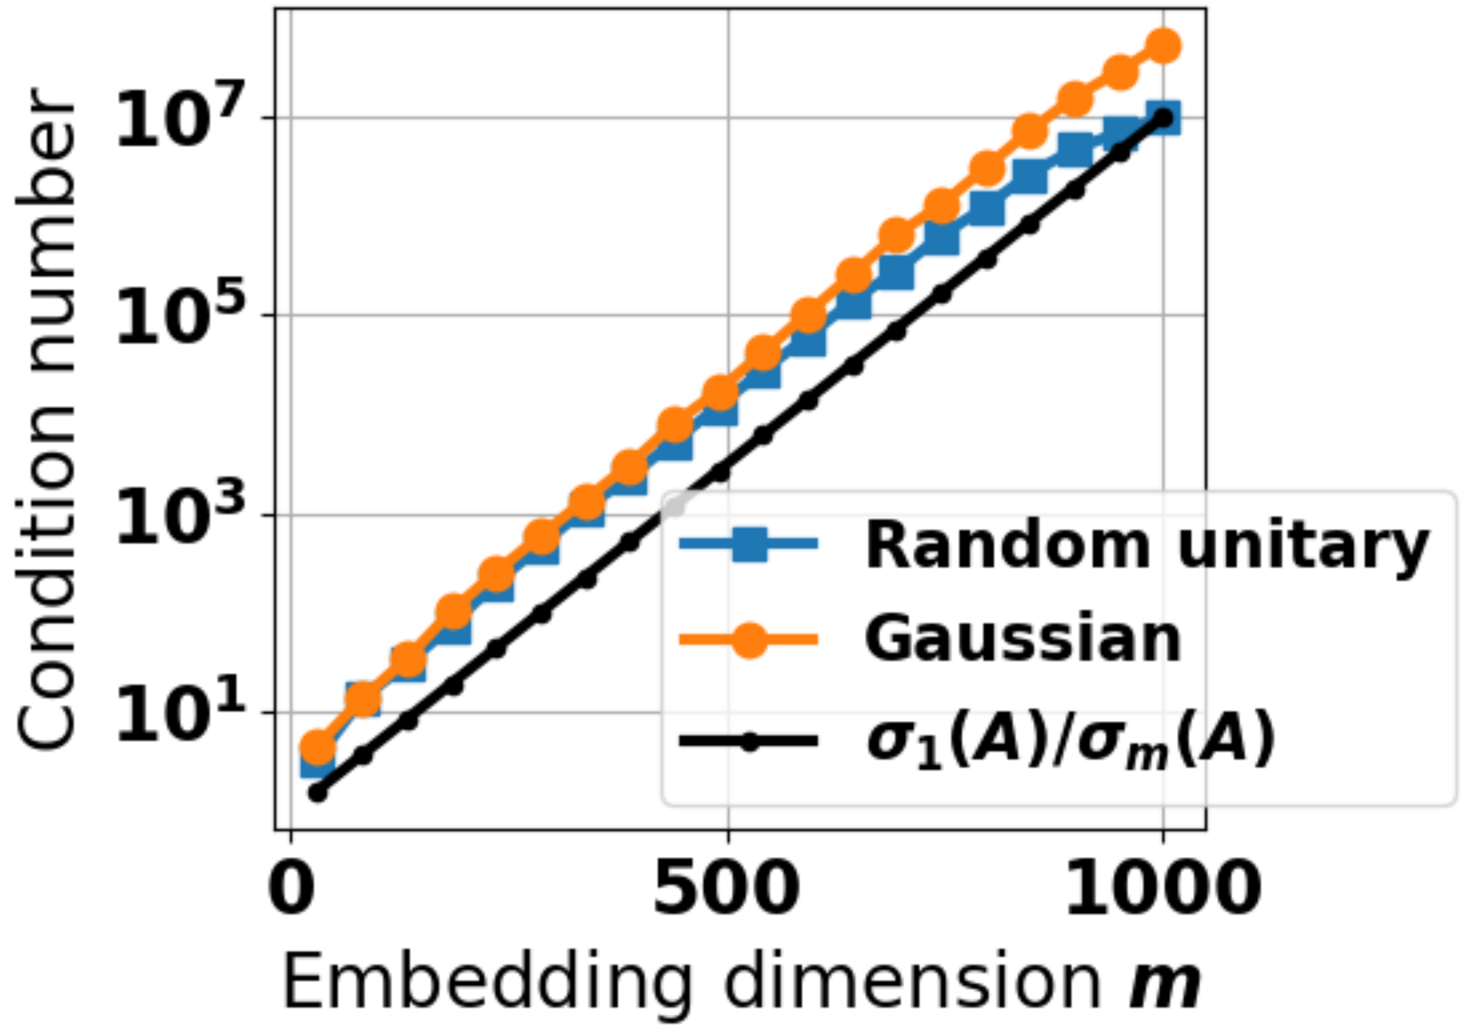

For any poorly conditioned $J(\theta)$, taking an embedding dimension m slightly smaller than numerical rank $r = \mathrm{rank}_\tau(J(\theta))$ brings a **well-conditioned** matrix with

$$\kappa(J(\theta)\Gamma) \leq 1/\tau$$

Variance-conditioning trade-off:

- Larger $m \Rightarrow$ smaller C_v
- Smaller $m \Rightarrow$ smaller $\kappa(J(\theta)\Gamma)$

Experiment: Synthetic dynamical system with $J(\cdot) \equiv J$

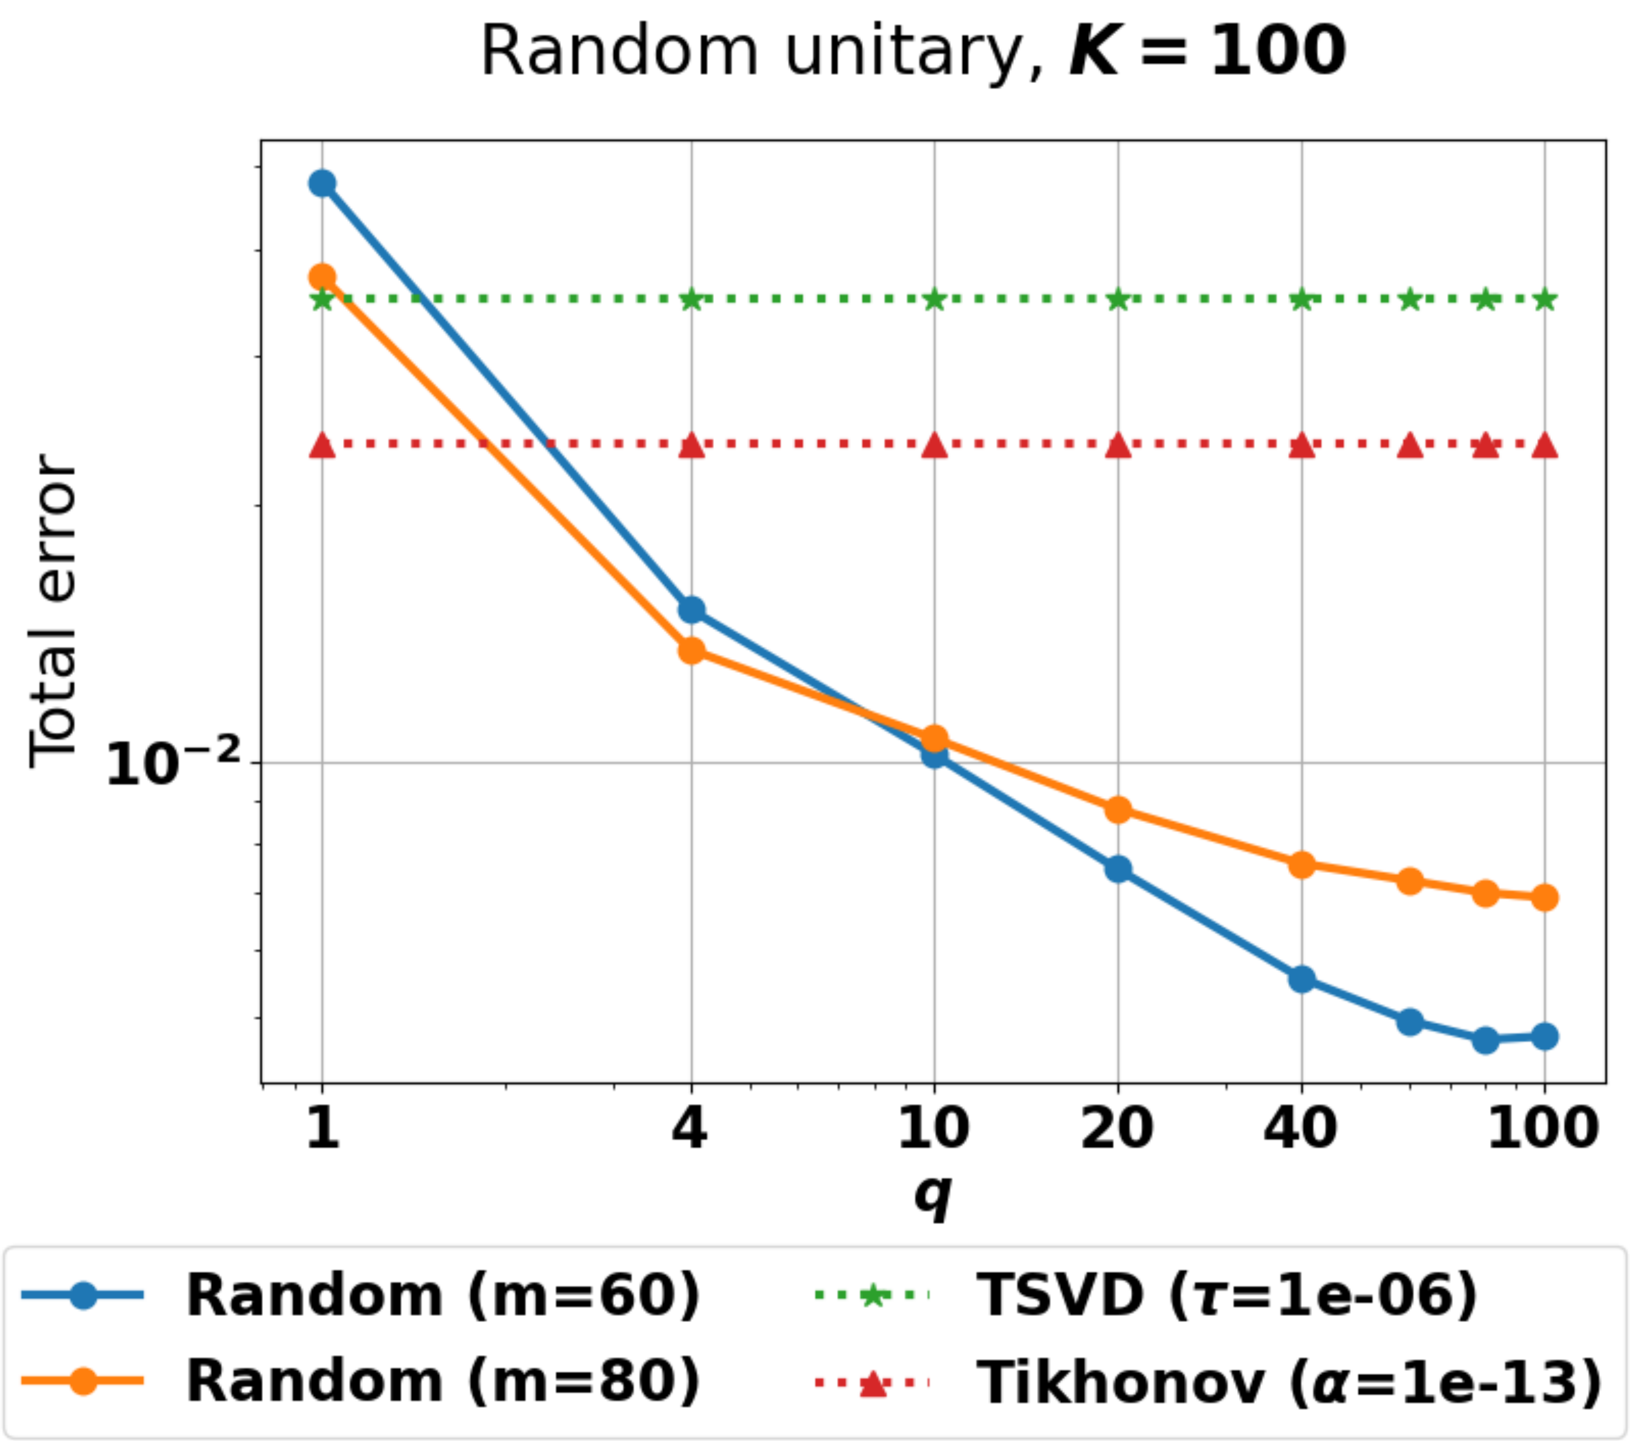

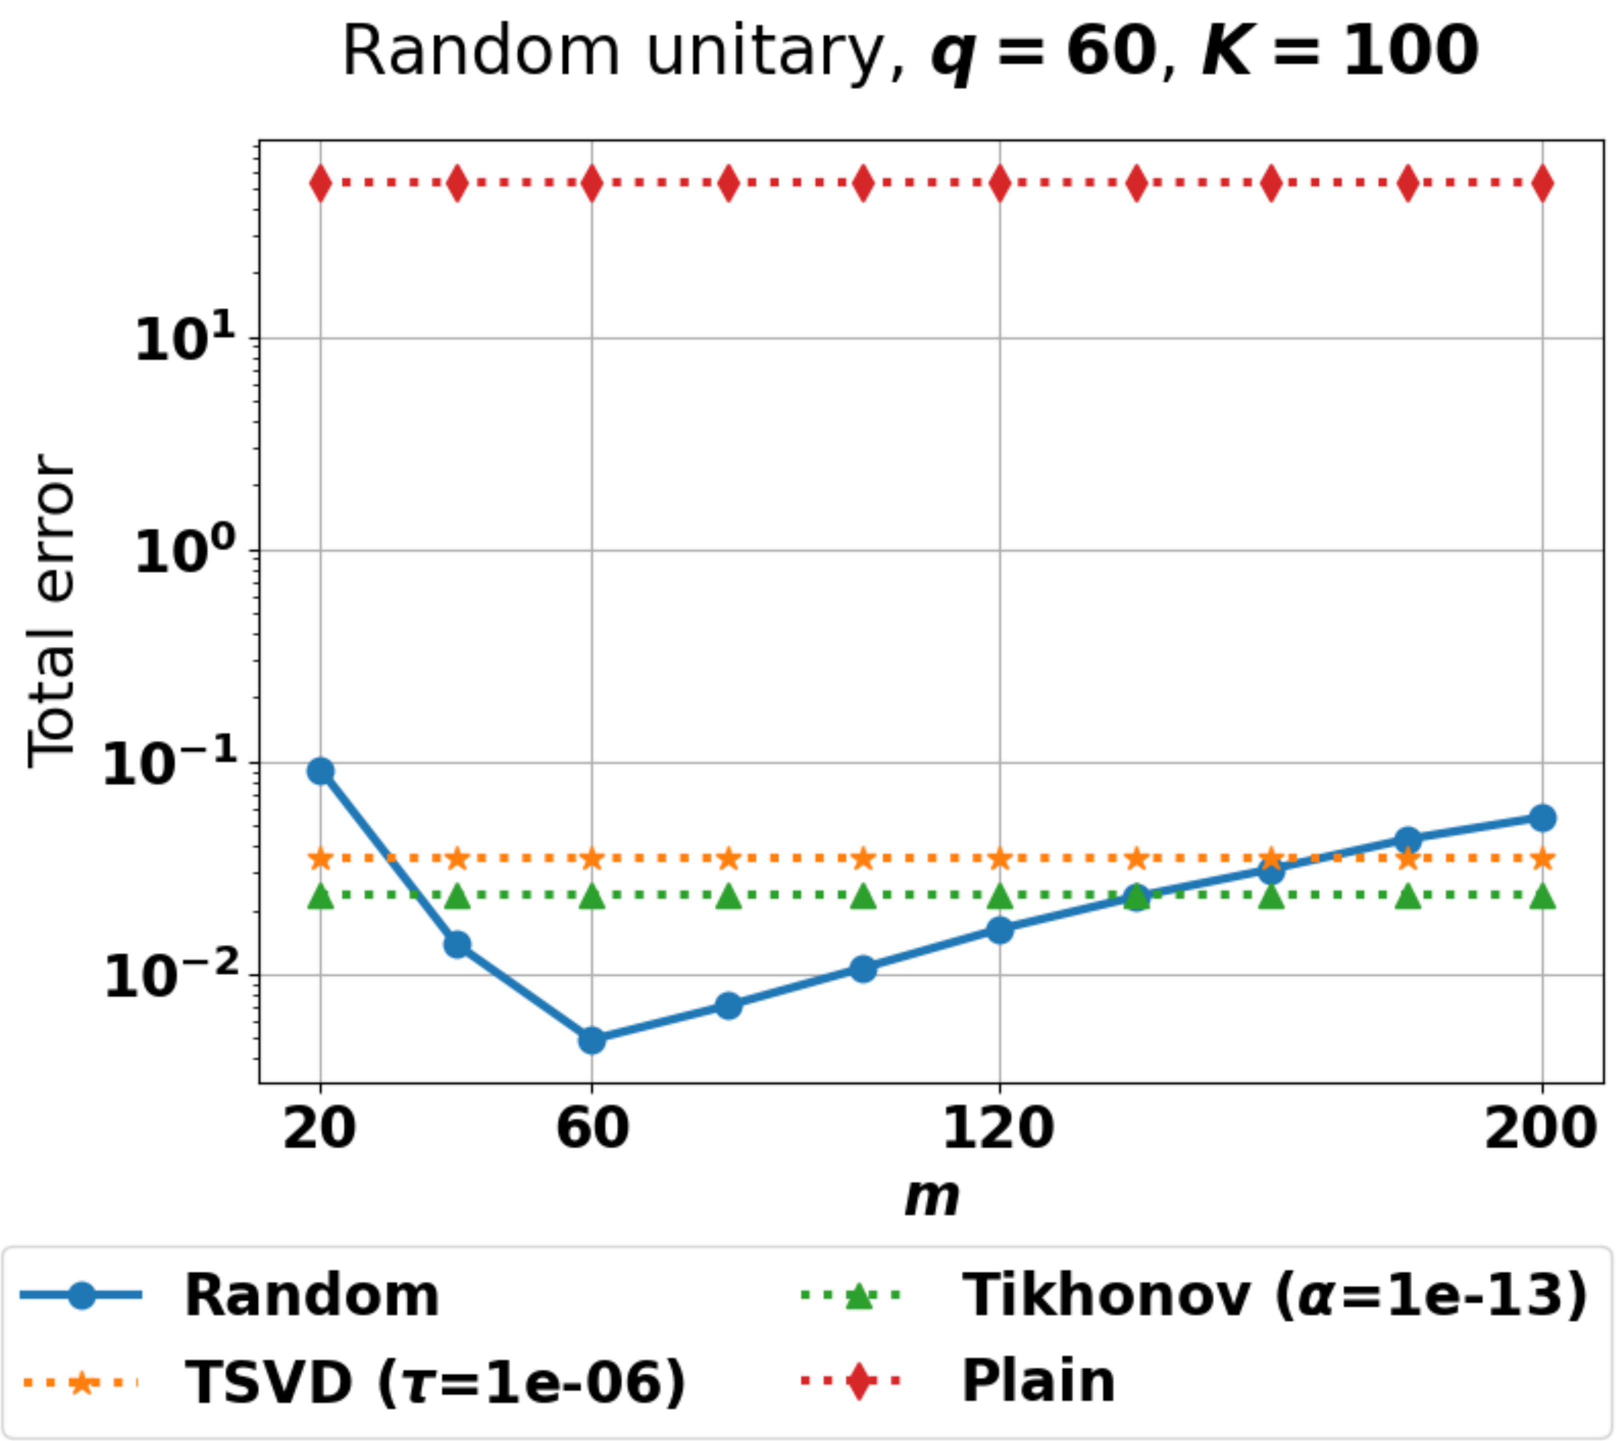

- Plain = direct time integration with poorly conditioned increments
- $\tau = 10^{-6}$ and $\alpha = 10^{-13}$ are the best regularization hyperparameters obtained from grid searches.

- **Regularized time integration** with TSVD or Tikhonov regularization leads to (large) accumulative bias.
- **Randomized time integration** becomes more accurate as q **increases (variance decreases)**.
- Optimal m under variance-conditioning trade-off

Experiment: Approximating double-well quantum dynamics

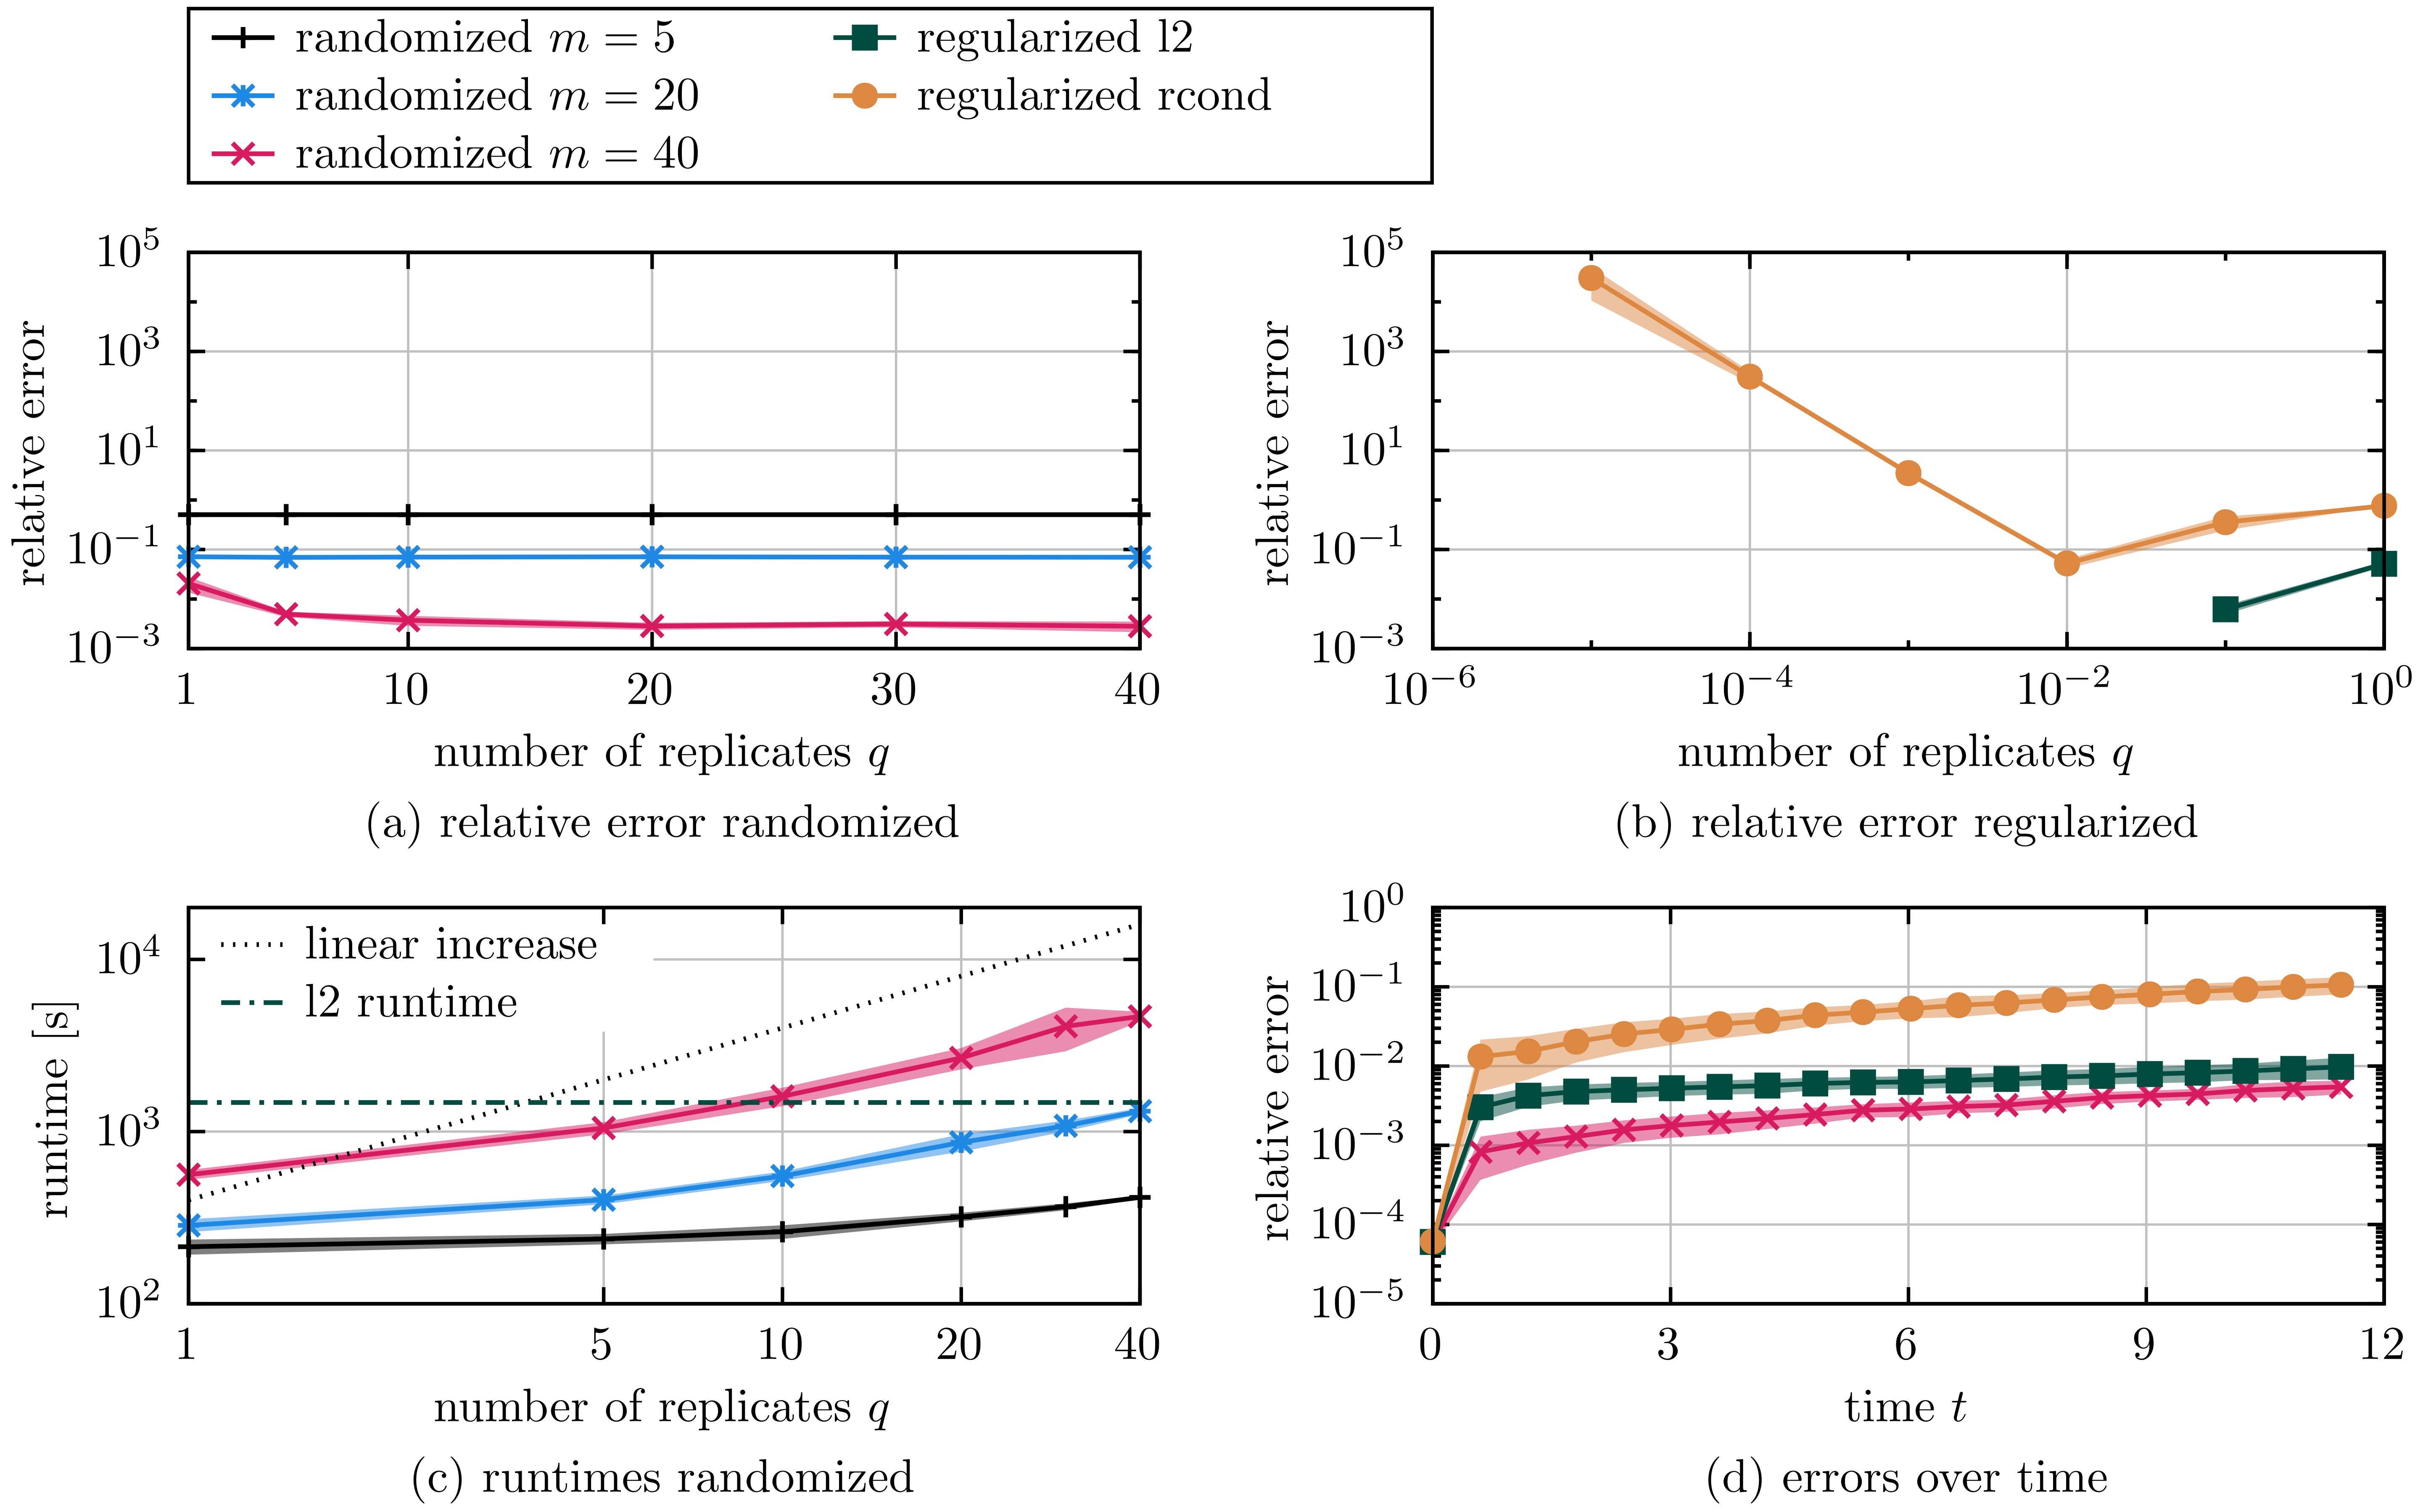

Following the setup in [Feischl, Lasser, Lubich, Nick, 2024]

- 1D-Schrödinger Equation with neural network parametrization, updated via Neural Galerkin scheme ([Bruna, Peherstorfer, Vanden-Eijnden, JCP2024], [Berman, Peherstorfer, NeurIPS2023])
- Randomized time integration ($m = 40$) outperforms the counterparts with TSVD and Tikhonov regularization.
- Randomized time integration ($m = 40$) becomes more accurate as q increases (variance decreases).
  - For $q > 10$, convergence of relative error in q plateaus due to numerical error in basic operations like addition.
- Regularized time integration is sensitive to hyperparameters and can fail under a bad choice.

Takeaways

Dong, Schwerdtner, Peherstorfer, “Randomize instead of Regularize: Stable Time Integration for Poorly Conditioned Dynamical Systems”, in preparation.

Problem: **Stable time integration** for discrete dynamical systems with **poorly conditioned least-squares increments**

Takeaways

Dong, Schwerdtner, Peherstorfer, "Randomize instead of Regularize: Stable Time Integration for Poorly Conditioned Dynamical Systems", in preparation.

Problem: **Stable time integration** for discrete dynamical systems with **poorly conditioned least-squares increments**

Challenge: Classical approaches involve **deterministic regularizations** like TSVD and Tikhonov regularization, which could **lead to bias accumulation over time**.

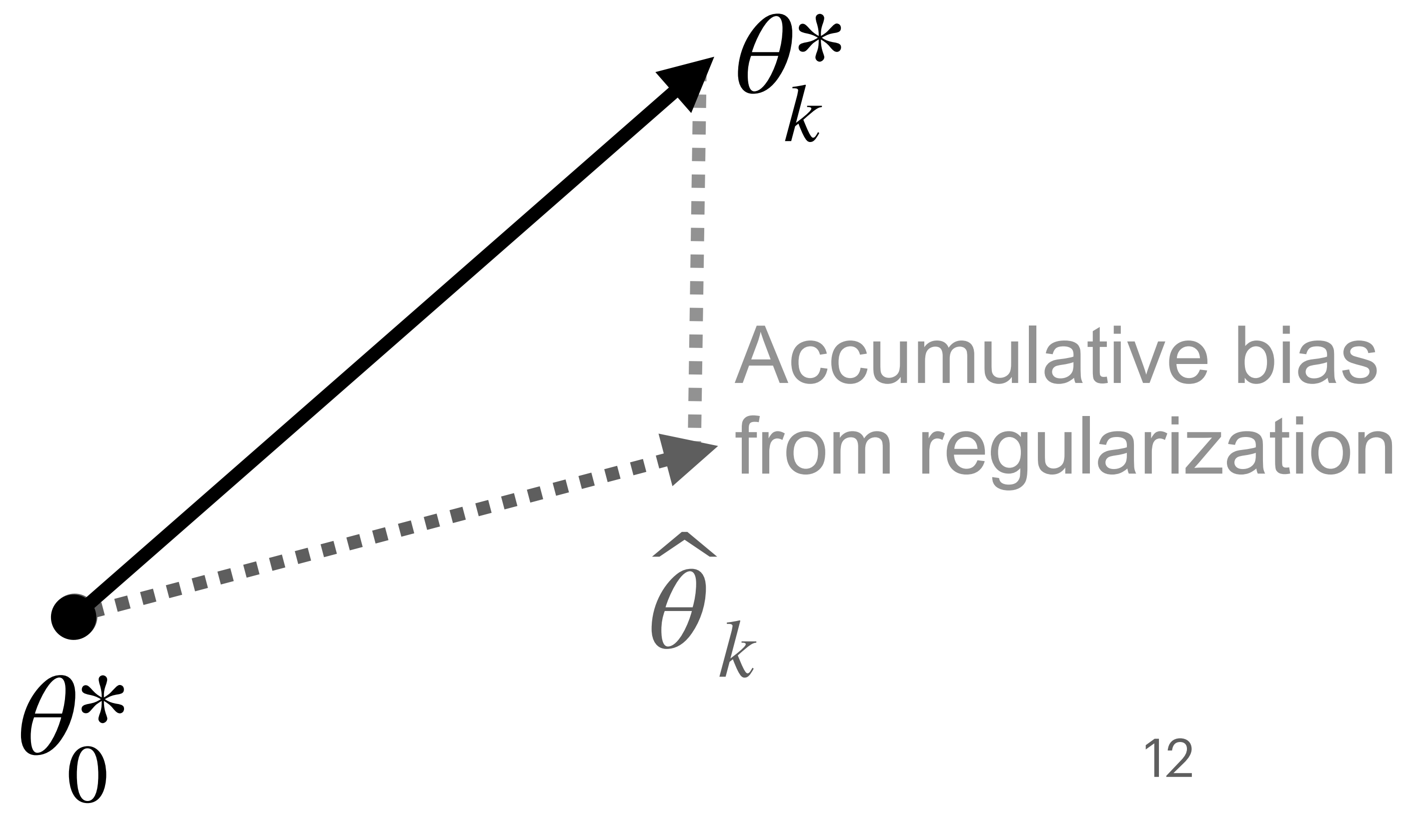

Takeaways

Dong, Schwerdtner, Peherstorfer, “Randomize instead of Regularize: Stable Time Integration for Poorly Conditioned Dynamical Systems”, in preparation.

Problem: **Stable time integration** for discrete dynamical systems with **poorly conditioned least-squares increments**

Challenge: Classical approaches involve **deterministic regularizations** like TSVD and Tikhonov regularization, which could **lead to bias accumulation over time**.

Our approach: **Randomized time integration** via **unbiased randomized increments**, with **variance** controlled by (i) the embedding dimension m and (ii) local sample size q.

- The randomized least-squares increment is **well-conditioned** when $m \lesssim \text{rank}_\tau(J(\theta))$.
- Trade-off between **variance (C_v)** and **conditioning** leads to an optimal choice of m.

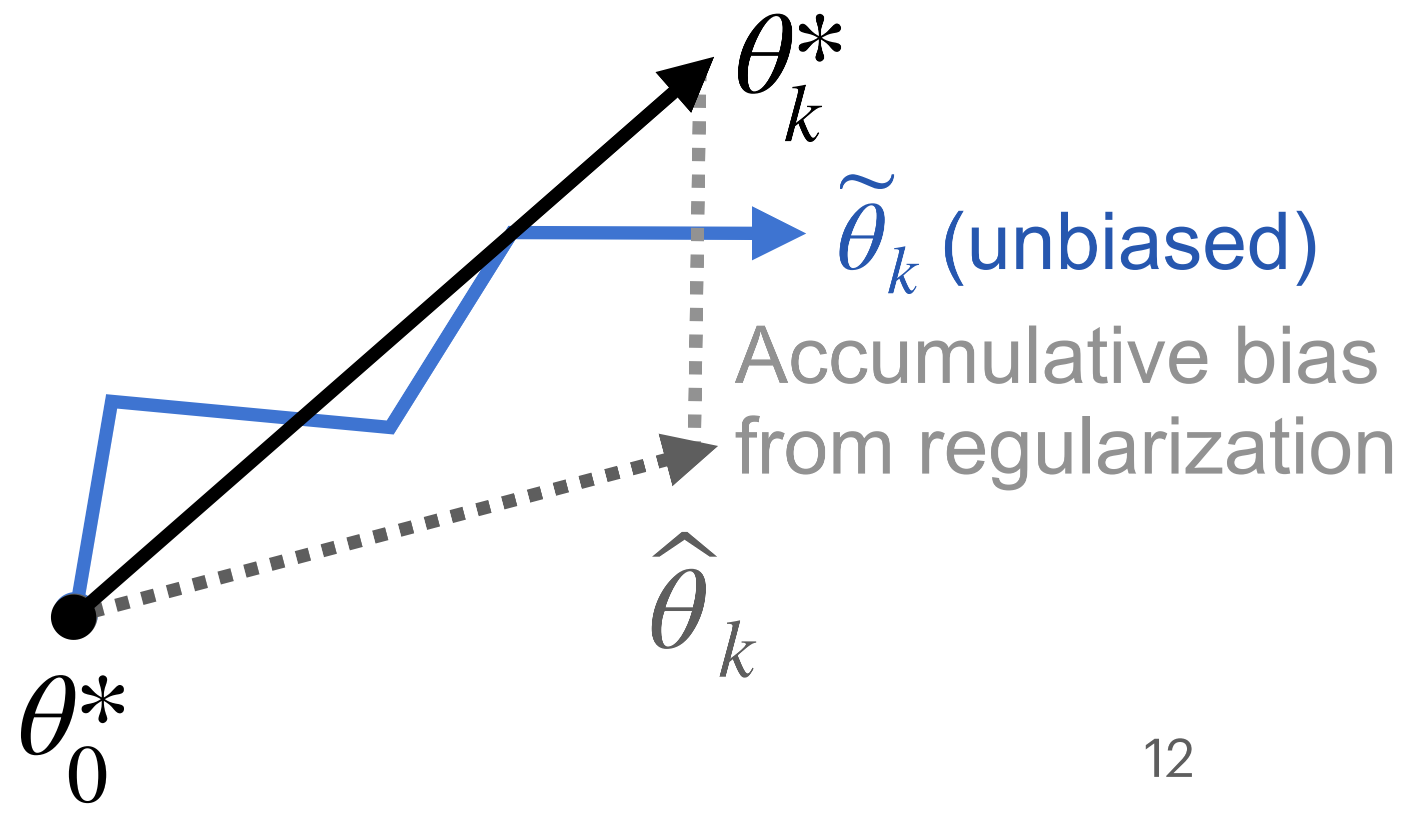

Takeaways

Dong, Schwerdtner, Peherstorfer, "Randomize instead of Regularize: Stable Time Integration for Poorly Conditioned Dynamical Systems", in preparation.

Problem: **Stable time integration** for discrete dynamical systems with **poorly conditioned least-squares increments**

Challenge: Classical approaches involve **deterministic regularizations** like TSVD and Tikhonov regularization, which could **lead to bias accumulation over time**.

Our approach: **Randomized time integration** via **unbiased randomized increments**, with **variance** controlled by (i) the embedding dimension m and (ii) local sample size q.

- The randomized least-squares increment is **well-conditioned** when $m \lesssim \mathrm{rank}_\tau(J(\theta))$.
- Trade-off between **variance (C_v)** and **conditioning** leads to an optimal choice of m.

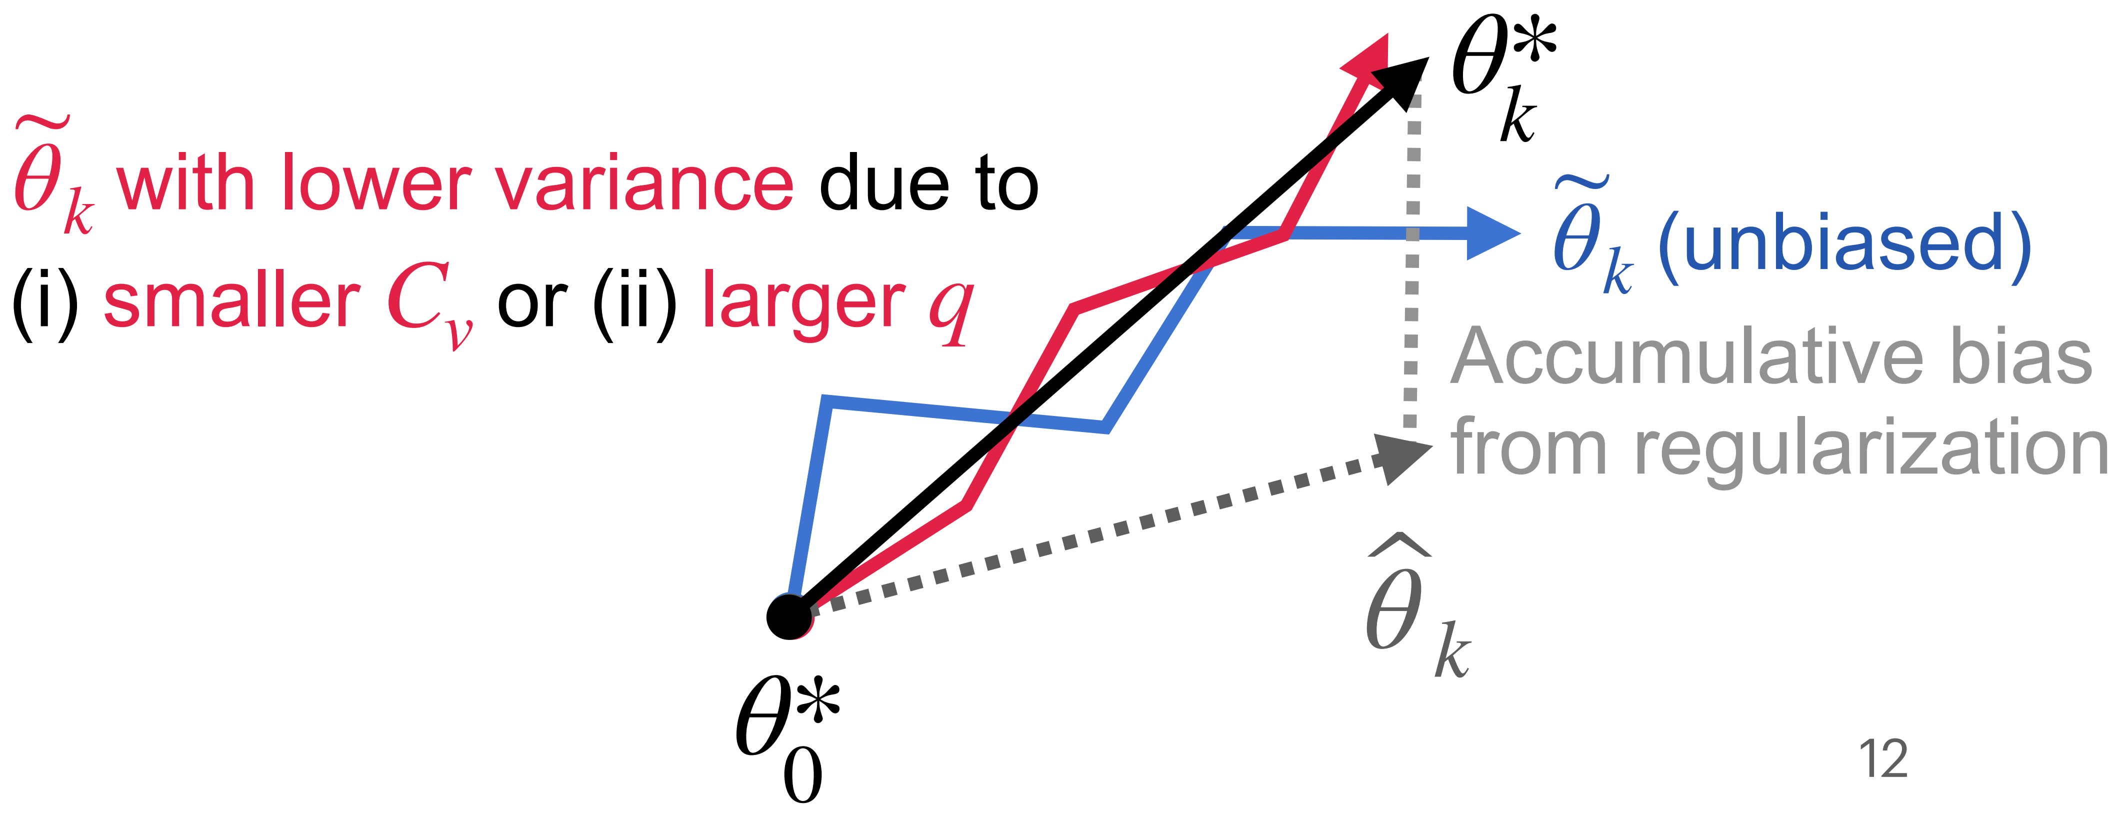

Takeaways

Dong, Schwerdtner, Peherstorfer, "Randomize instead of Regularize: Stable Time Integration for Poorly Conditioned Dynamical Systems", in preparation.

Problem: **Stable time integration** for discrete dynamical systems with **poorly conditioned least-squares increments**

Challenge: Classical approaches involve **deterministic regularizations** like TSVD and Tikhonov regularization, which could **lead to bias accumulation over time**.

Our approach: **Randomized time integration** via **unbiased randomized increments**, with **variance** controlled by (i) the embedding dimension m and (ii) local sample size q.

- The randomized least-squares increment is **well-conditioned** when $m \lesssim \text{rank}_\tau(J(\theta))$.
- Trade-off between **variance (C_v)** and **conditioning** leads to an optimal choice of m.

$\widetilde{\theta}_k$ with lower variance due to (i) smaller C_v or (ii) larger q

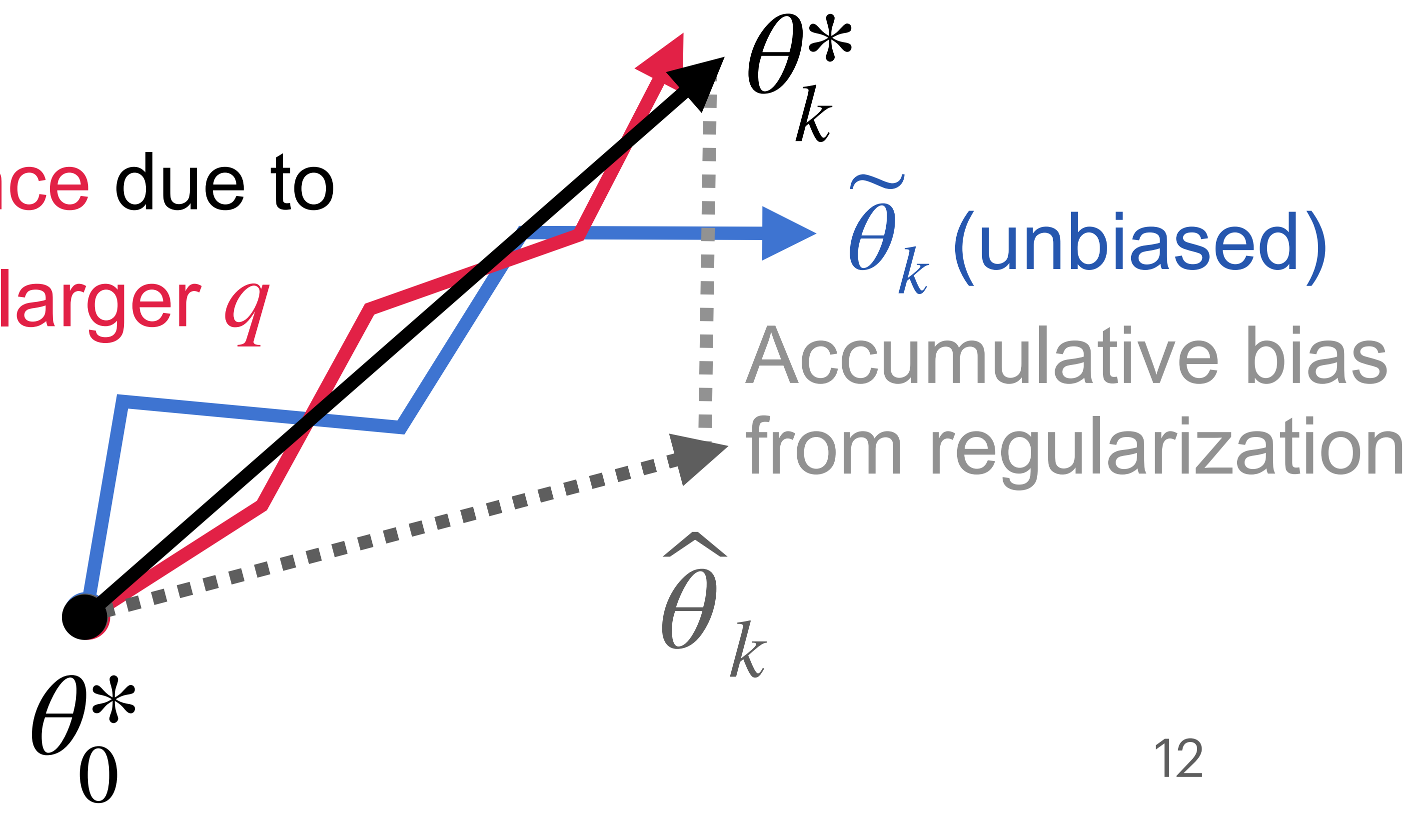

Thank you!