Collider Spin Tomography with Missing Neutrinos

Jia Liu (刘佳)
Peking University

Collaborators: Tengyu Ai, Qi Bi, Yuxin He, Zekun Li, Xiao-Ping Wang
Based on arXiv: 2607.28346

International Workshop on New Opportunities for Particle Physics 2026
2026-08-14

Outlines

- Quantum tomography and kinematic ambiguity
- Our goal and setup
- Identifiability analysis
- Self-consistent reconstruction
- Numerical results
- Summary

Why full spin tomography?

$$ee, pp \to f\bar{f}$$

- Tomography reconstructs the quantum state of the produced particle pair, not just one selected observable
  - Spin information is encoded in the production density matrix
  - Gives polarization $B_i^{\pm}$ and spin-correlation observables C_{ij}
  - Enables entanglement and CHSH tests at colliders
  - Probing BSM parameters in spin correlations

What does full tomography reconstruct?

- Differential spin density matrix $\rho(\hat{k})$ for two spin-half particles:

$$\rho(\hat{k}) = \frac{1}{4}\frac{\mathrm{d}\sigma}{\sigma \mathrm{d}\hat{k}}\left[\mathbb{I}_4 + B_i^+(\hat{k})\,\sigma_i \otimes \mathbb{I}_2 + B_i^-(\hat{k})\,\mathbb{I}_2 \otimes \sigma_i + C_{ij}(\hat{k})\,\sigma_i \otimes \sigma_j\right]$$

  - Differential production rate: $\mathscr{R}(\hat{k}) \equiv \frac{\mathrm{d}\sigma}{\sigma\,\mathrm{d}\hat{k}}$ **1**
  - Single-particle polarizations: $B_i^{\pm}(\hat{k})$ **3+3=6**
  - Spin-correlation matrix: $C_{ij}(\hat{k})$ in helicity basis **9**
  - In total: 1+3+3+9=**16** differential components

Normalization:

$$\int \mathrm{d}\hat{k}\,\mathrm{Tr}[\rho(\hat{k})] = 1$$

- Once $\rho(\hat{k})$ is reconstructed, any polarization, correlation, entanglement, or Bell observable can be derived from it.

Spin is measured through decay

- $e^+e^- \to \tau^+\tau^-$ production at $\hat{k}$ direction and its spin density matrix:

$$\rho^{\mathrm{prod}}_{\alpha\alpha';\beta\beta'} = \frac{1}{|\overline{M}|^2} \overline{\sum_{\text{ini}}} M(e^+e^- \to \tau^+_\alpha(\hat{k})\tau^-_\beta(-\hat{k}))M^*(e^+e^- \to \tau^+_{\alpha'}(\hat{k})\tau^-_{\beta'}(-\hat{k}))$$

- Colliders do not measure spin directly
- Spin information is transferred to decay-product momenta
  - Two-body τ decay: $\tau^\pm \to \pi^\pm \overset{(-)}{\nu}$ and spin measurement operator:

$$\hat{D}^\pm_{\alpha\alpha'}(\hat{q}_\pm) = \frac{f_\pi^2}{2m_\tau^3|q|} M(\tau^\pm_{\alpha'} \to \pi^\pm(\hat{q}_\pm)\nu)M^*(\tau^\pm_\alpha \to \pi^\pm(\hat{q}_\pm)\nu) = \frac{1}{2}(\mathbf{I}_2 \mp \hat{q}_\pm \cdot \vec{\sigma})_{\alpha\alpha'}$$

  - Differential $\pi^\pm$ distributions: $\frac{1}{\sigma}\frac{\mathrm{d}\sigma}{\mathrm{d}\hat{q}_+\mathrm{d}\hat{q}_-} = \mathrm{Tr}\left[\rho^{\mathrm{prod}}_{\tau^+\tau^-}\hat{E}\right] \equiv \mathrm{Tr}\left[\rho^{\mathrm{prod}}_{\tau^+\tau^-}\hat{D}^+(\hat{q}_+) \otimes \hat{D}^-(\hat{q}_-)\right]$
- Decay turns spin tomography into a momentum-distribution measurement

The real obstacle: invisible neutrinos

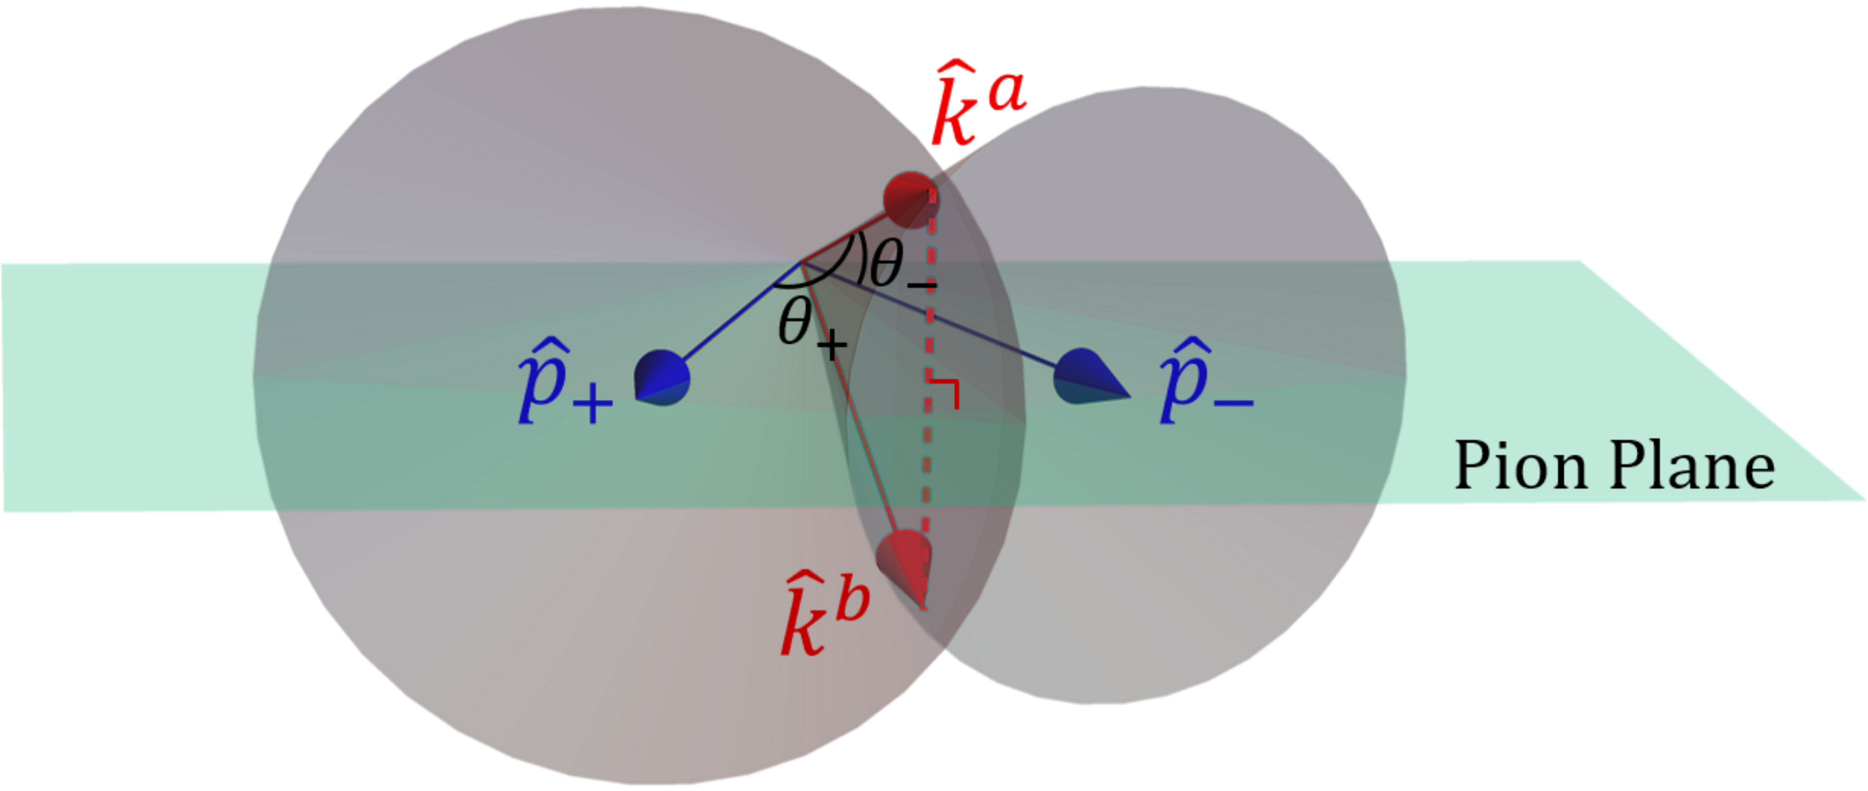

- Weak decay typically involves ν
  - E.g. τ, or clean leptonic decay of top
- We observe only the visible particle momenta, not p_ν
- The parent direction $\hat{k}$ is not uniquely reconstructed
  - 2-fold $\tau\bar{\tau}$: $e^+e^- \to \tau^+(\hat{k})\tau^-(-\hat{k}) \to \pi^+(\hat{p}_+)\bar{\nu}_\tau\pi^-(\hat{p}_-)\nu_\tau$

    dof = 8 ($p_{\nu,\bar{\nu}}$) - 4 ($E/\vec{p}$) - 2 (m_{τ^+,τ^-}) - 2 ($m_{\nu,\bar{\nu}}$) = 0
  - Multi-fold leptonic $t\bar{t}$: $pp \to t\bar{t} \to b\ell_1^+\nu_{\ell_1} + \bar{b}\ell_2^-\bar{\nu}_{\ell_2}$

    dof = 8 ($p_{\nu,\bar{\nu}}$) - 2 ($\vec{p}_T$) - 2 ($m_{t,\bar{t}}$) - 2 (m_{W^+,W^-}) - 2 ($m_{\nu,\bar{\nu}}$) = 0

One visible event $\Longrightarrow$ multiple possible kinematic configurations

Information incomplete at event-by-event level

Belle τ EDM: flat averaging over two solutions

BELLE
PLB 551 (2003) 16
JHEP 04 (2022) 110

- Belle searched for the τ EDM using classical optimal observables
- Each event admits two kinematically allowed τ directions $\hat{k}^{a,b}$
- The observables were evaluated for both solutions and averaged equally
  - **Flat average**: $O_{\text{flat}}(\Phi) = \dfrac{1}{2}\left[O\left(\phi^a\right) + O\left(\phi^b\right)\right]$
  - $\Phi = \{\vec{p}_+, \vec{p}_-\}$: visible $\pi^\pm$ momentum ; dof = 6;
  - $\phi = \{\hat{k}, \hat{q}_+, \hat{q}_-\}$ $\hat{k}$: τ direction in COM; $\hat{q}_\pm$: $\pi^\pm$ direction in tau rest frame
- Classical observable, not a reconstruction of the full spin density matrix

See also
Li et al. 2605.01233
Zhou et al. 2605.26724

Flat averaging avoids choosing one branch,
but does not use the spin-dependent probabilities of the two solutions

BTW: for BELLE II, it may use extra vertex information to deal with tau pair twofold issue.

Top measurements: reduce or choose

ATLAS
Nature 633 (2024) 542
CMS
ROPP 87(2024)117801

- Dileptonic $t\bar{t}$: two neutrinos create reconstruction ambiguities
- ATLAS/CMS measured the entanglement marker
  - $D = -\frac{1}{3}\mathrm{Tr}[C]$, using diagonals rather than the full ρ. No ρ_{SM} assumed.
- ATLAS: Ellipse method (one real solution in 85% of events) $\Longrightarrow$ Neutrino Weighting method (assign weights to each possible solutions, 5% of events) $\Longrightarrow$ Simple pairing method (ℓ+closest b-jet)
- CMS: Best fit solution method (lowest $m_{t\bar{t}}$, weighted by $m_{\ell b}$, 90% of events)
- These are powerful measurements, but not unresolved-fold tomography

Existing strategies either reduce the target observable or select best kinematic branch

Outlines

- Quantum tomography and kinematic ambiguity
- Our goal and setup
- Identifiability analysis
- Self-consistent reconstruction
- Numerical results
- Summary

Our Goal

- Keep all kinematic solutions $\{\phi^a, \phi^b, \ldots\}$
- Do not choose the “best” branch
- Do not assume a production template: $\rho^{\text{prod}}_{\text{SM}}$
- Reconstruct ρ via data

visible event $\Longrightarrow$ $\{\phi^a, \phi^b\}$ $\Longrightarrow$ visible-data tomography

Event-by-event level

Global property level

> Ambiguity is not the same as information loss!
> What part of $\rho(\hat{k})$ is fixed by visible data alone?

Our benchmark channel

- $e^+e^- \to \tau^+\tau^- \to \pi^+\pi^- + \nu\bar{\nu}$
  - Clean production kinematics in an ee collider
  - $\tau^\pm \to \pi^\pm\nu$ provides maximal spin-analyzing power
  - Two neutrinos produce an event-by-event twofold ambiguity
  - Simple enough for an analytic identifiability analysis
  - Directly relevant to Belle I/II, BESIII, and future STCF studies

A minimal process containing both complete spin information and unresolved kinematic ambiguity

Geometry of the twofold solution

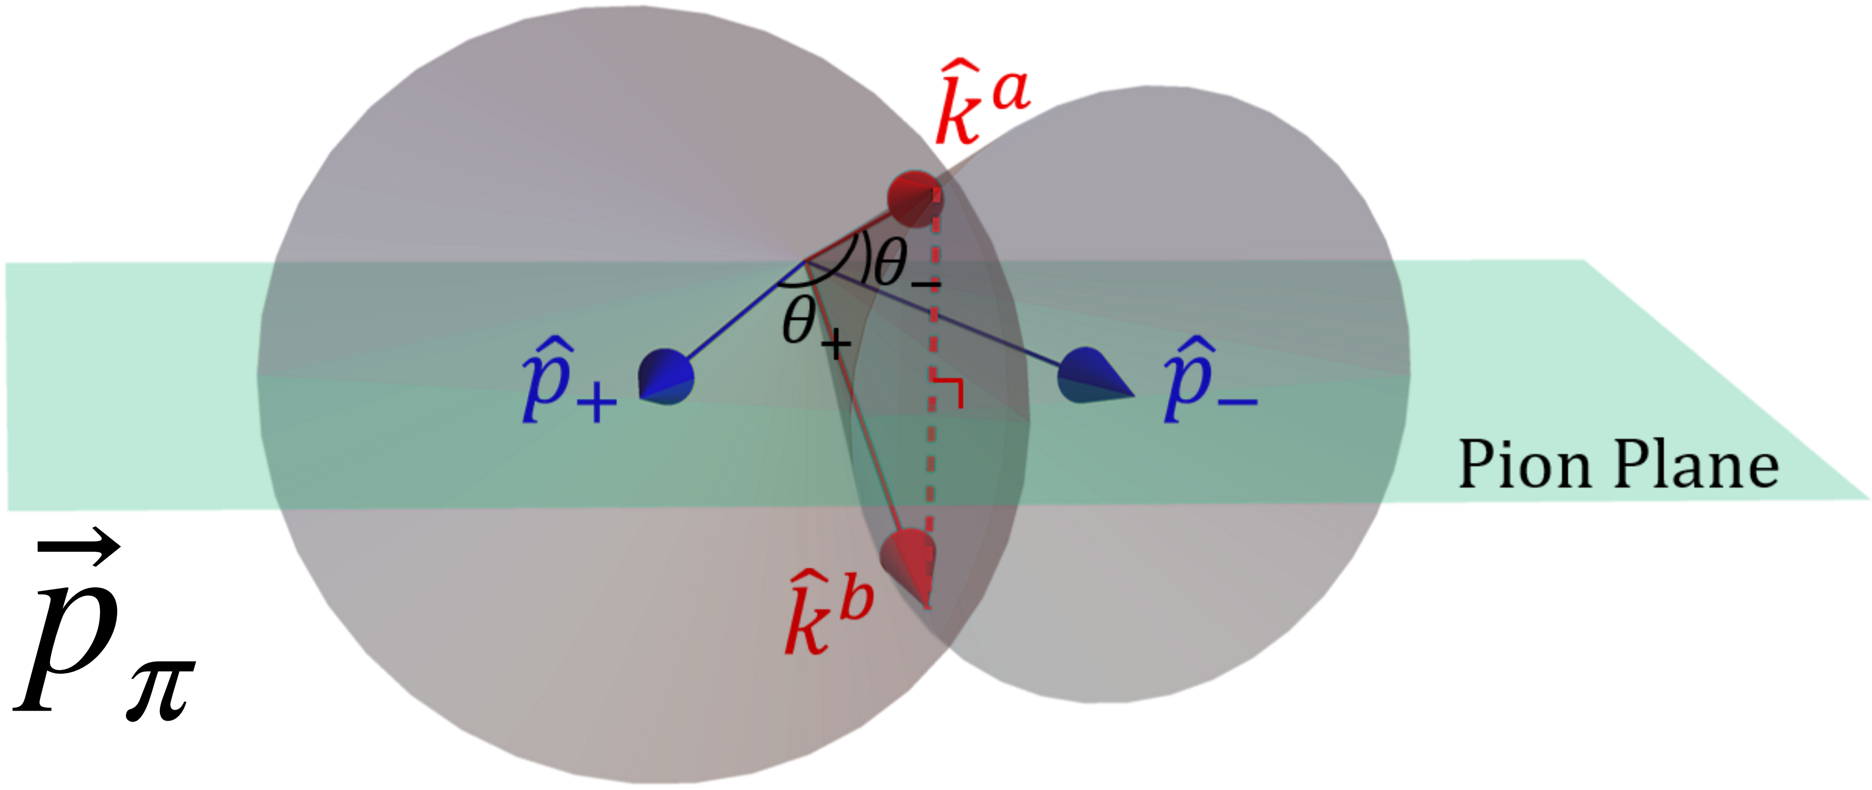

- The τ energy is fixed by $e^+e^- \to \tau^+\tau^-$
- For each pion, $\hat{k}$ must lie on a cone around $\vec{p}_\pi$
- The two pion cones intersect at two points
- These two points give the two allowed τ directions

$$\hat{k}^{a,b} = u\,\hat{p}_+ + v\,\hat{p}_- \pm w\,\frac{\vec{p}_+ \times \vec{p}_-}{|\vec{p}_+ \times \vec{p}_-|}$$

The two solutions differ only by the sign of the component normal to the pion plane.

The twofold ambiguity is a mirror ambiguity in the parent τ direction.

From a visible event to hidden folds

- The detector measures only visible variables:

$$\Phi_{\text{visible}} = \{\vec{p}_+, \vec{p}_-\}; \, \phi = \{\hat{k}, \hat{q}_+, \hat{q}_-\}$$

- The full configuration has two possible branches:

$$\phi \Longrightarrow \Phi \text{ (1-to-1)}; \quad \Phi \Longrightarrow \{\phi^a(\Phi), \phi^b(\Phi)\}$$

- Each branch has its own τ direction and spin analyzers, the branch label is not observed
- **Ideal case** (no kinematic ambiguity):
- $f(\phi) \equiv \dfrac{\mathrm{d}\sigma}{\sigma\,\mathrm{d}\phi} = \mathrm{Tr}\left[\rho_{\tau^+\tau^-}\,\hat{E}\right]$ internal full distribution;
    - The decay as general quantum measurement $\hat{E}$
    - Observed results from experiments: $\dfrac{\mathrm{d}\sigma}{\sigma\,\mathrm{d}\phi}$
    - Extracting the spin information: $\langle O \rangle \Longrightarrow \rho_{\tau^+\tau^-}$

From a visible event to hidden folds

- The detector measures only visible variables:

$$\Phi_{\text{visible}} = \{\vec{p}_+, \vec{p}_-\};\ \phi = \{\hat{k}, \hat{q}_+, \hat{q}_-\}$$

- The full configuration has two possible branches:

$$\Phi \Longrightarrow \{\phi^a(\Phi), \phi^b(\Phi)\}$$

- Variable relations:
  - $f(\phi) \equiv \dfrac{\mathrm{d}\sigma}{\sigma\,\mathrm{d}\phi} = \mathrm{Tr}\left[\rho_{\tau^+\tau^-}\,\hat{E}\right]$ internal distribution; $g(\Phi) \equiv \dfrac{\mathrm{d}\sigma}{\sigma\,\mathrm{d}\Phi}$ visible pion distribution
  - Phase-space Jacobian $J^{-1}(\Phi) \equiv \left|\det\left(\dfrac{\partial\{\vec{p}_+, \vec{p}_-\}}{\partial\{\hat{k}, \hat{q}_+, \hat{q}_-\}}\right)\right| = \left|\det\left(\dfrac{\partial\{\Phi\}}{\partial\{\phi\}}\right)\right|$
  - Distribution relation: $g(\Phi) = \sum_{t=a,b} J^t(\Phi) f(\phi^t) \equiv \int \mathrm{d}\hat{k}\,\mathrm{Tr}\left[\rho(\hat{k})\overline{E}\right]$, with $\overline{E}(\Phi) = \sum_t J^t(\Phi)\,\hat{E}(\phi^t)\delta(\hat{k} - \hat{k}^t)$

The visible data provides coarse-grained measurements

Outlines

- Quantum tomography and kinematic ambiguity
- Our goal and setup
- Identifiability analysis
- Self-consistent reconstruction
- Numerical results
- Summary

The theoretical question: what information is invisible?

- Define the visible-data map: $\mathscr{M} : \rho \longmapsto g(\Phi)$
- Consider a deformation of the production density matrix: $\rho \longrightarrow \rho + \delta\rho$
- It is unidentifiable only when: $\delta g(\Phi) = \mathscr{M}[\delta\rho](\Phi) = 0,$ for all visible Φ
  - Equivalently, $\delta\rho \in \ker \mathscr{M}$
- Multiple hidden solutions do not automatically imply information loss
- The visible map may remain informationally complete
- Only directions in $\ker \mathscr{M}$ are invisible to the data
- All directions outside the kernel remain reconstructible
  - **Two possible cases:**

$$\ker \mathscr{M} = 0 \quad \Longrightarrow \quad \text{full identifiability}$$

$$\ker \mathscr{M} \neq 0 \quad \Longrightarrow \quad \text{only the null directions are lost}$$

Information loss is determined by the null space—not by the number of kinematic solutions

Only one density-matrix direction is invisible

- Solving the null-space condition for $e^+e^- \to \tau^+\tau^- \to \pi^+\pi^- + \nu\bar{\nu}$

$$\rho(\hat{k}) = \frac{1}{4}\mathscr{R}(\hat{k})\left[\mathbb{I}_4 + B_i^+(\hat{k})\,\sigma_i \otimes \mathbb{I}_2 + B_i^-(\hat{k})\,\mathbb{I}_2 \otimes \sigma_i + C_{ij}(\hat{k})\,\sigma_i \otimes \sigma_j\right]$$

- Analytic result for $\delta g(\Phi) = 0$ when varying $\delta\rho$:

$$\delta\mathscr{R}(\hat{k}) = \delta\left(\frac{d\sigma}{\sigma d\hat{k}}\right) = 0, \qquad \delta B_i^\pm(\hat{k}) = 0,$$

$$\delta C_{ij}(\hat{k}) = 0 \quad \text{except for, } \delta C_{nr}(\hat{k}) = -\delta C_{rn}(\hat{k}) = d_{\rm null}/\mathscr{R}(\hat{k})$$

$$C = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & C_{rn} \\ 0 & -C_{nr} & 0 \end{pmatrix}$$

- where $d_{\rm null}$ is independent of $\hat{k}$.
- Therefore, the invisible combination is: $C_{nr}^- \equiv C_{nr} - C_{rn}$.
- The differential rate $\dfrac{d\sigma}{\sigma d\hat{k}}$, $B_i^\pm(\hat{k})$, 8 out of 9 C_{ij} remains identifiable
- Only one antisymmetric correlation direction C_{nr}^- is lost.

The twofold ambiguity produces a 1D null space—not a general failure of tomography

Outlines

- Quantum tomography and kinematic ambiguity
- Our goal and setup
- Identifiability analysis
- Self-consistent reconstruction
- Numerical results
- Summary

Without a production template, the fold probabilities are unknown

- For one visible event Φ, both folds are kinematically allowed: $\Phi \Longrightarrow \{\phi^a, \phi^b\}$
- Their conditional fold weights are

$$\omega_t(\Phi;\rho) = \frac{J^t(\Phi)\, f(\phi^t;\rho)}{\sum_{s=a,b} J^s(\Phi)\, f(\phi^s;\rho)}, \qquad t = a, b\,.$$

- The fold weights depend on the production density matrix ρ, but ρ is precisely what tomography aims to reconstruct
- Circular structure: $\rho \Longrightarrow f(\phi^t;\rho) \Longrightarrow \omega_t \Longrightarrow$ tomographic observables $\Longrightarrow \rho$

Theory template ρ_{template} $\Longrightarrow$ known weights,

Model-independent tomography ρ unknown $\Longrightarrow$ weights must be learned from data

Production-model independence turns fold assignment into a self-consistency problem

The fixed-point family and self-consistent unfolding

- Define the vector $\mathbf{x} = \{\frac{\mathrm{d}\sigma}{\sigma \mathrm{d}\hat{k}}|_\alpha, B^+_{i,\alpha}, B^-_{i,\alpha}, C_{ij,\alpha}\}$, 16 components binned for $\hat{k}$
- Data-driven iteration: $\mathbf{x}^{(n-1)} \Longrightarrow \omega_t^{(n)} \Longrightarrow \langle \widetilde{O} \rangle \Longrightarrow \mathbf{x}^{(n)}$

$$\mathbf{x}^{(n)} = T\left(\mathbf{x}^{(n-1)}\right)$$

  - Start from a trial density matrix
  - Use it to compute the fold probabilities
  - Reconstruct all density-matrix parameters
  - Repeat until the update becomes self-consistent

The fixed-point family and self-consistent unfolding

- Fixed-point condition: $\mathbf{x}_* = T(\mathbf{x}_*)$
- The null space produces a continuous fixed-point family:

$$\mathbf{x}_*(d_{\text{null}}) = \mathbf{x}_{\text{truth}} + d_{\text{null}}\,\mathbf{n}_{\text{null}}$$

- where $\mathbf{n}_{\text{null}}$ is the direction associated with $C_{nr} - C_{rn}$
- $\mathbf{x}_{\text{truth}}$ is the 15 components with their true value
  - Every member of the family gives the same visible distribution $\delta g(\Phi) = 0$
  - All members coincide with the truth in the identifiable subspace
  - They differ only along the unidentifiable null direction
  - Positivity restricts which members correspond to physical ρ

Local convergence toward the fixed-point family

- Linearized iteration
  - Near a fixed point $\mathbf{x}_*$,

$$\delta\mathbf{x}^{(n)} \simeq D_T(\mathbf{x}_*)\,\delta\mathbf{x}^{(n-1)}, \text{ with } (D_T)_{ml} \equiv \left.\frac{\partial T_m}{\partial x_l}\right|_{\mathbf{x}_*}$$

- For an eigenmode $\mathbf{v}$,

$$D_T\mathbf{v} = \lambda\mathbf{v}, \qquad \delta\mathbf{x}^{(n)} \propto \lambda^n$$

- Convergence criterion:

$$|\lambda| < 1 \quad \Longrightarrow \quad \text{perturbation is damped,}$$
$$|\lambda| = 1 \quad \Longrightarrow \quad \text{marginal direction .}$$

Local convergence toward the fixed-point family

- Null direction: analytic result
  - Because every point along the null family is a fixed point, $\mathbf{x}_*(d_{\text{null}}) = \mathbf{x}_{\text{truth}} + d_{\text{null}}\mathbf{n}_{\text{null}}$, differentiation $\mathbf{x}_* = T(\mathbf{x}_*)$ gives

$$D_T(\mathbf{x}_*)\,\mathbf{n}_{\text{null}} = \mathbf{n}_{\text{null}}$$

$$\lambda_{\text{null}} = 1$$

- Identifiable subspace: numerical result
  - After removing the null mode, $\rho_{\text{id}}(D_T) < 1$
  - For the least-damped identifiable mode, numeric results shows

$$\lambda_{\text{id, max}}^{\text{SM}} \simeq 0.9605, \quad \lambda_{\text{id, max}}^{\text{non-SM}} \simeq 0.9608$$

- Thus, all identifiable perturbations are locally damped, while the null direction has no restoring force.

The identifiable components converge uniquely; the null component does not

Outlines

- Quantum tomography and kinematic ambiguity
- Our goal and setup
- Identifiability analysis
- Self-consistent reconstruction
- Numerical results
- Summary

Closure test: Standard Model

- $e^+e^- \to \tau^+\tau^- \to \pi^+\pi^- + \nu\bar{\nu}$
- Generate MC events from the SM truth density matrix
- Reconstruct using only visible pion momenta
- Start from a spin-independent trial density matrix
- The first update reproduces the flat 50% average
- Iterations move the identifiable components toward the truth
- Differential cross-section is successfully reconstructed

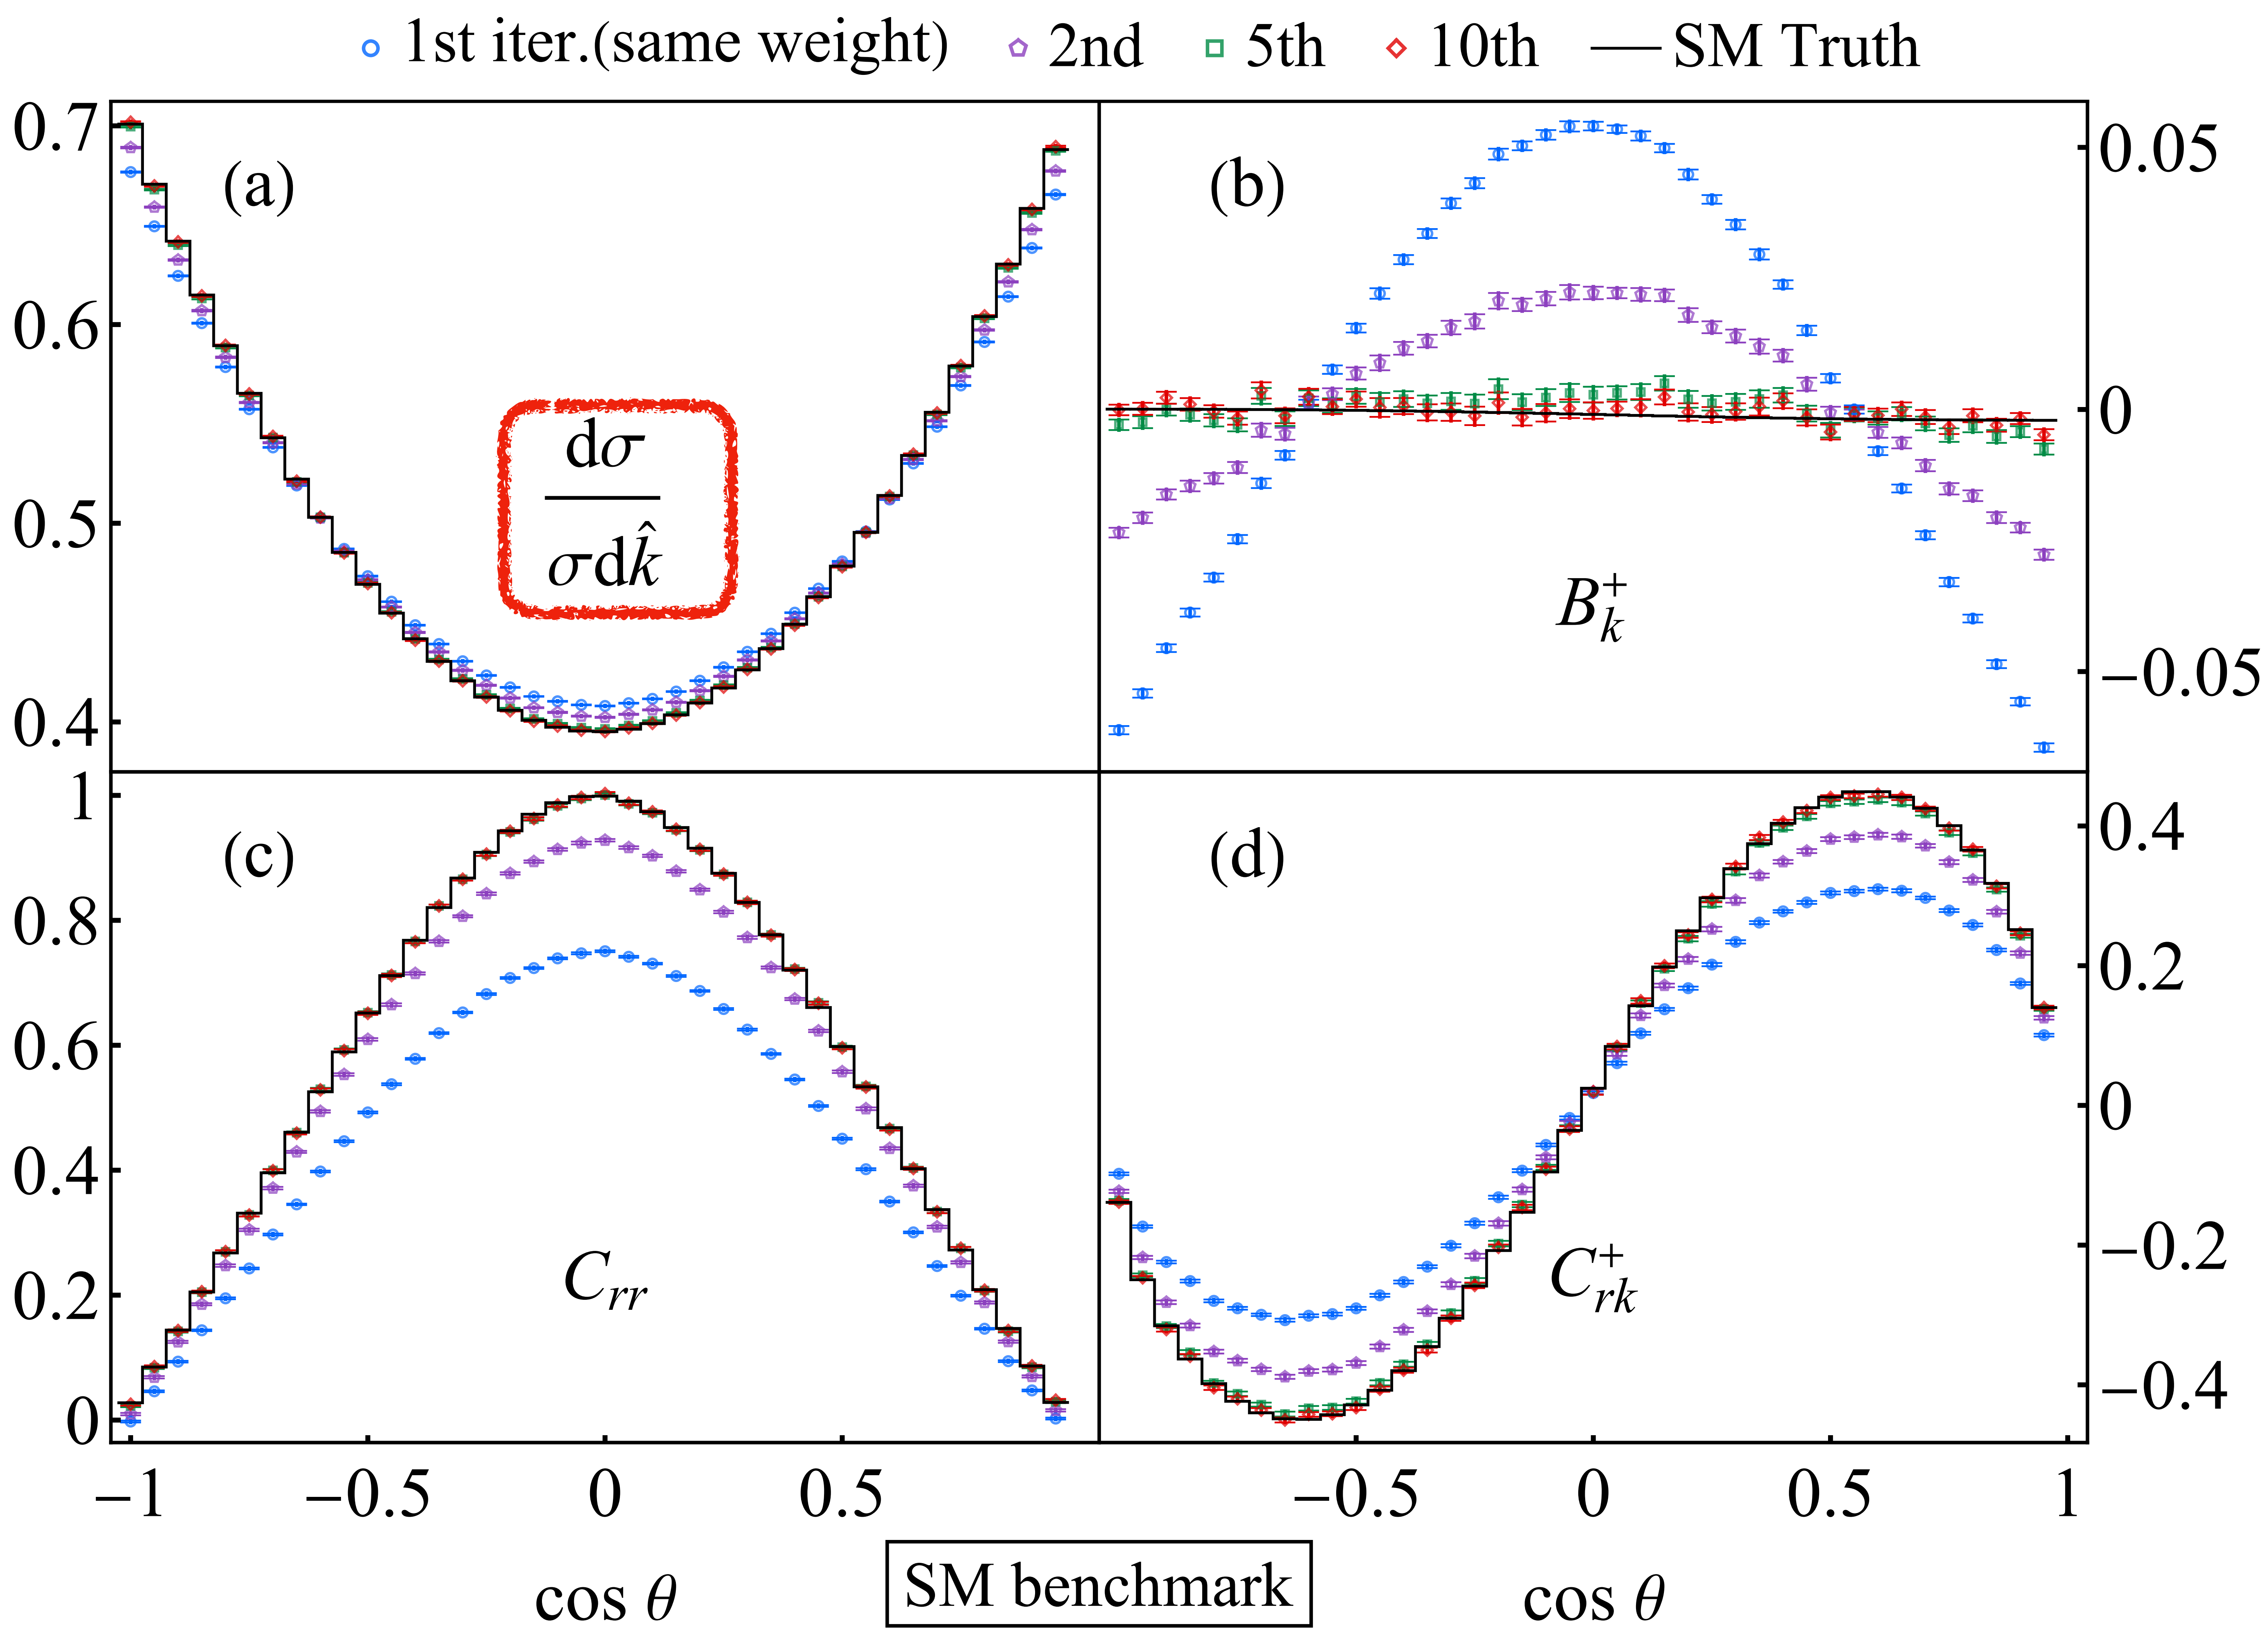

Flat averaging is not tomography; self-consistency removes the fold-induced bias

Closure test: Standard Model

- Other spin-correlations reconstructed successfully

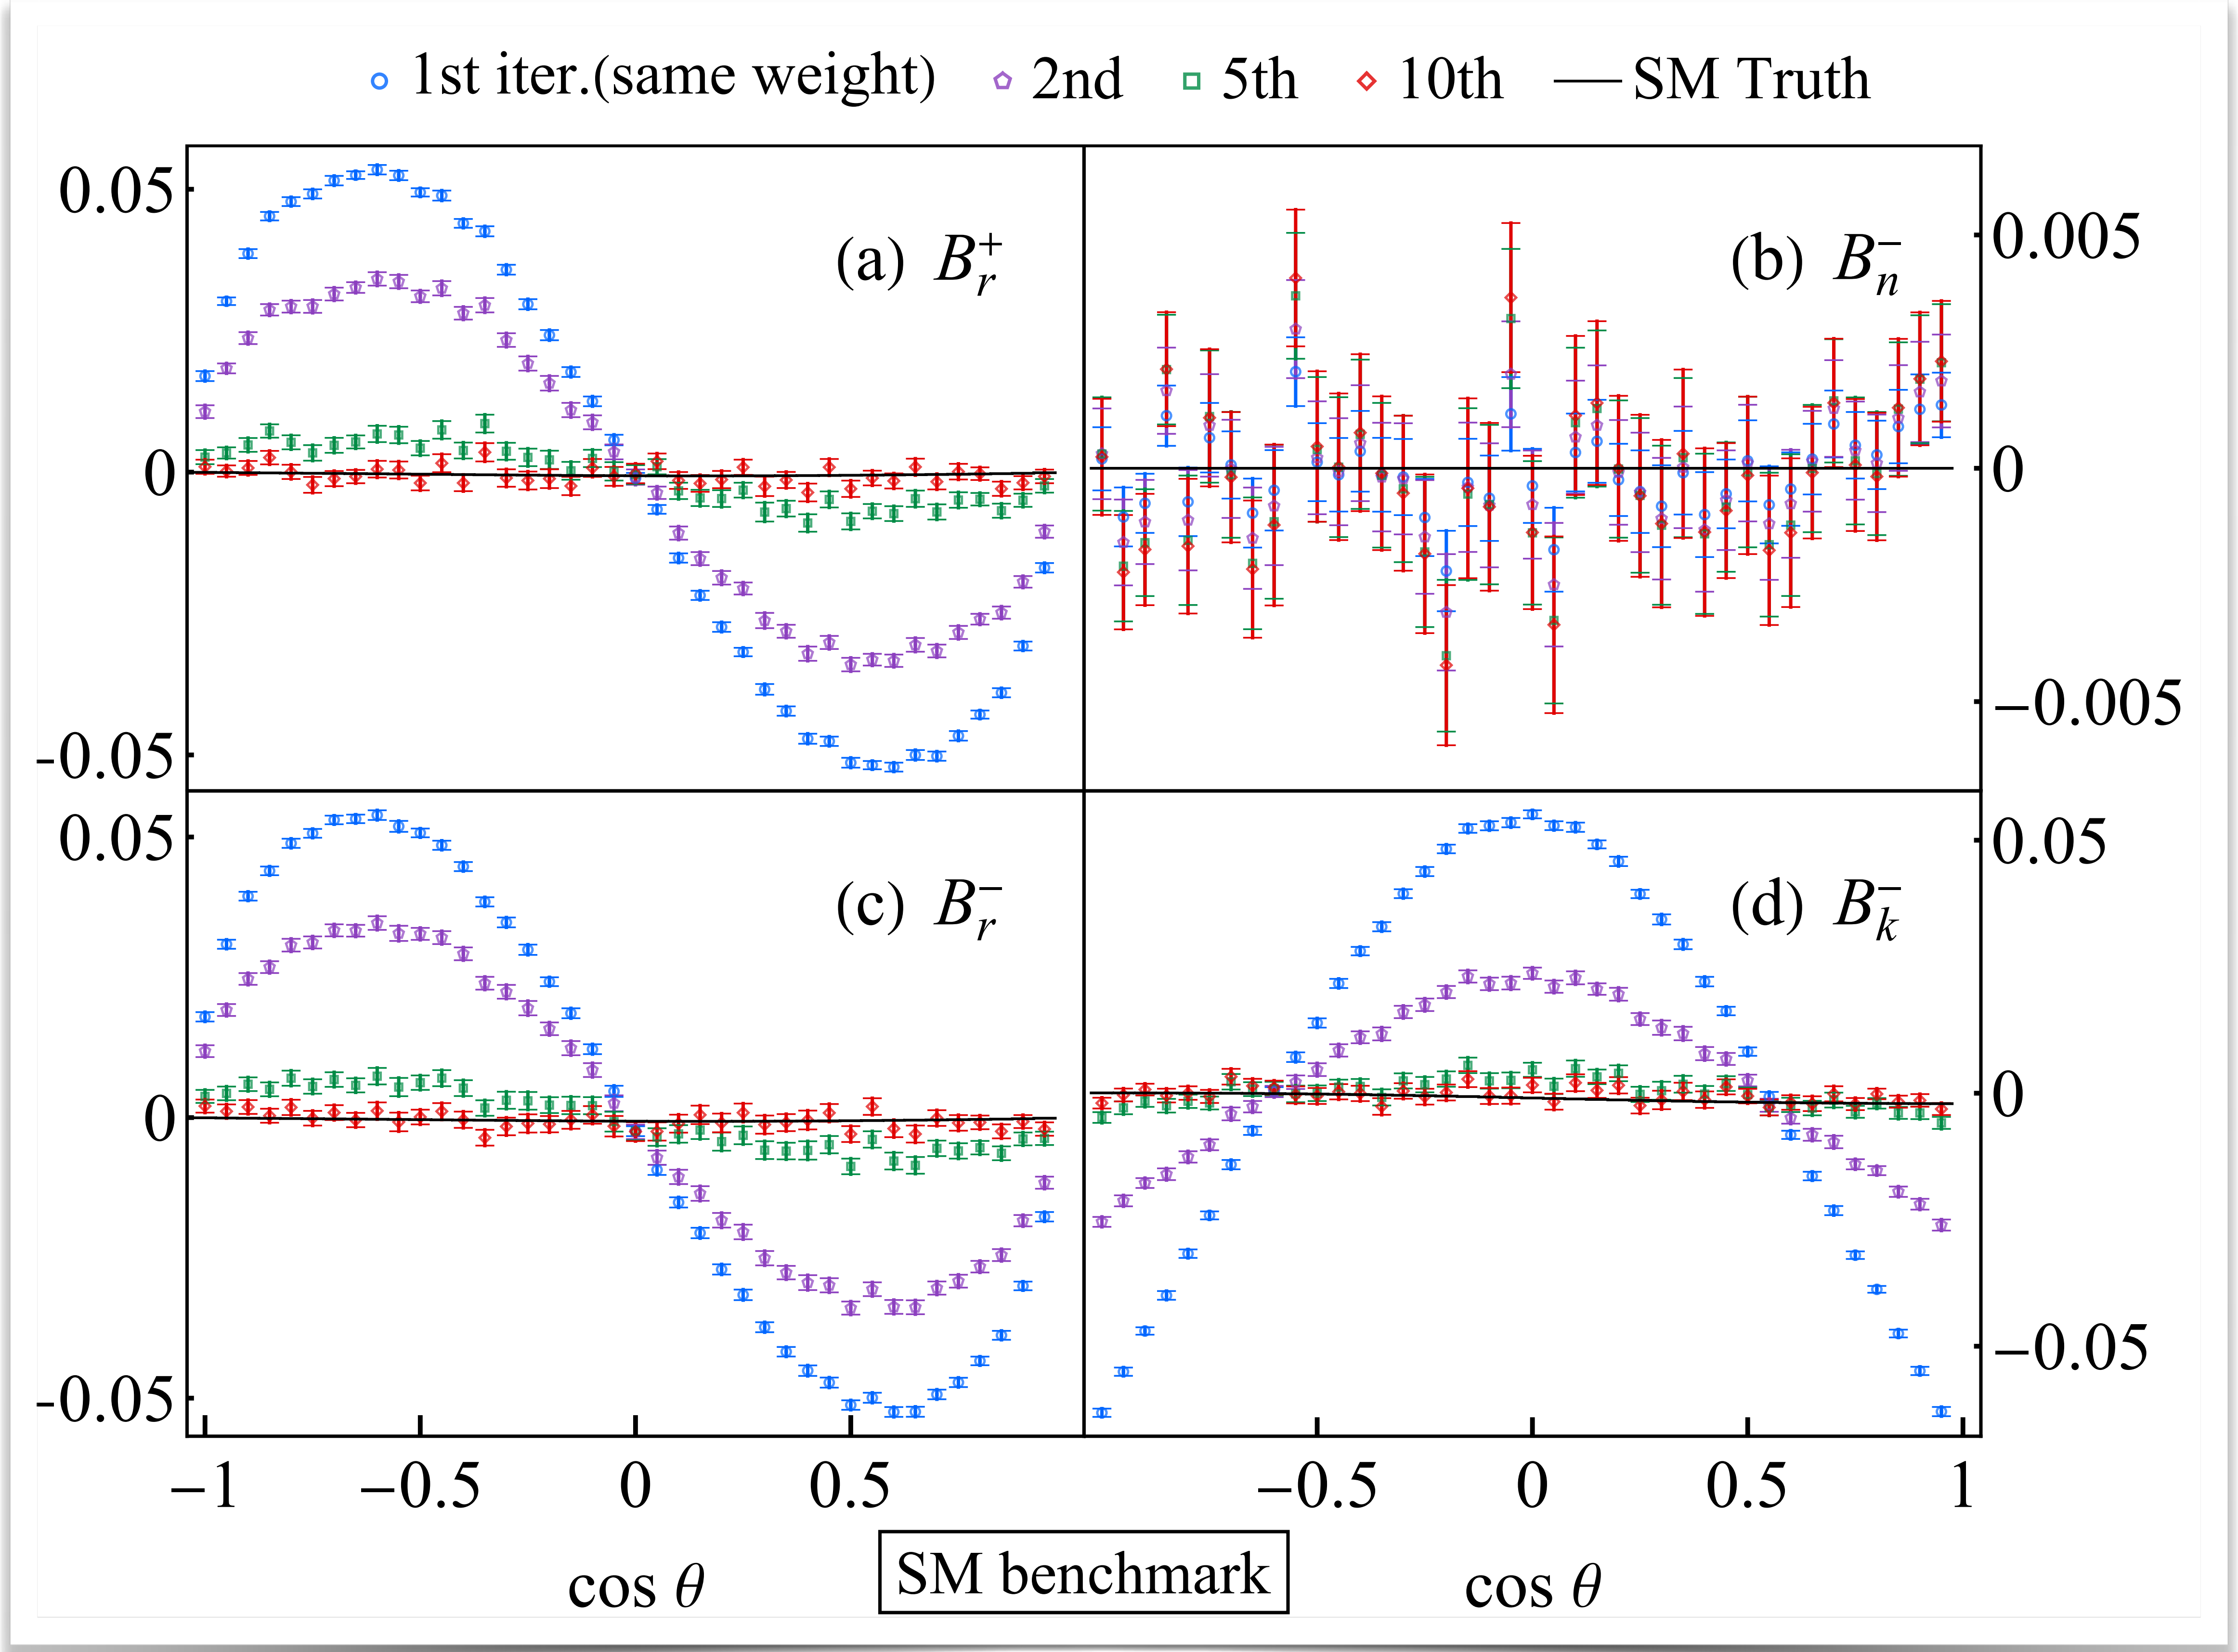

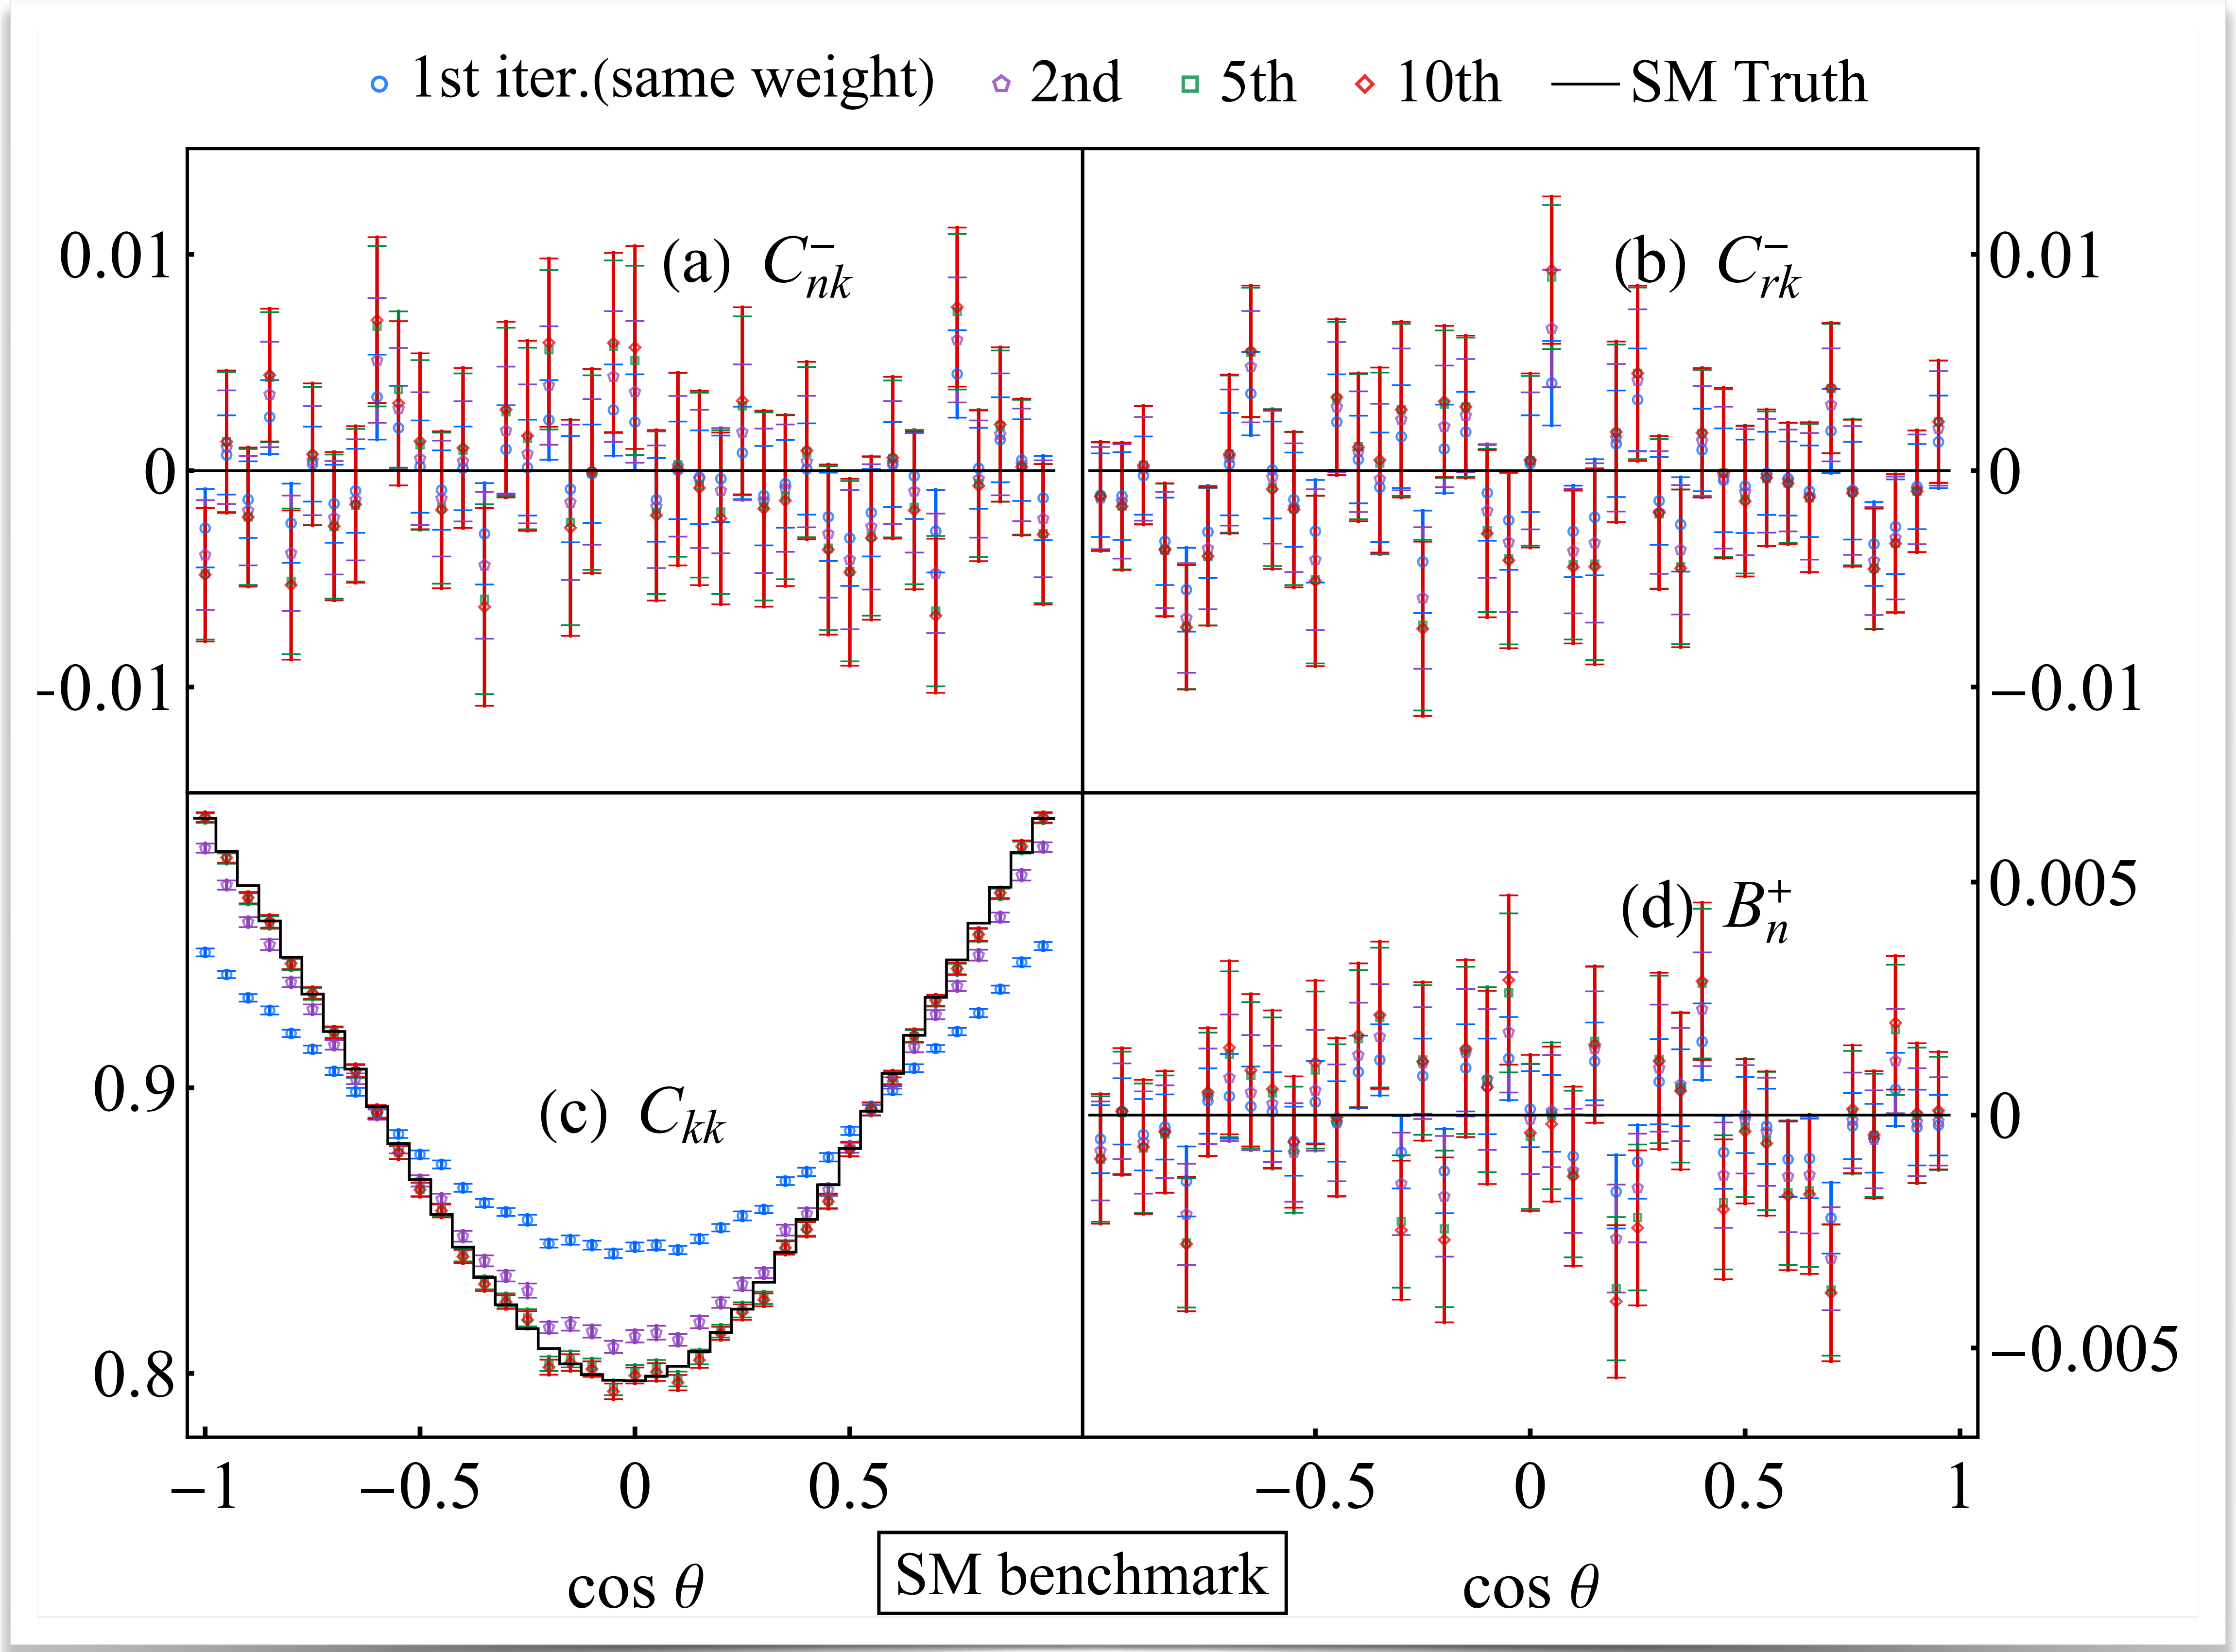

Flat averaging is not tomography; self-consistency removes the fold-induced bias

Closure test: Standard Model

- Others reconstructed successfully,
- but C^{-}_{nr} constructed?
  See non-SM model

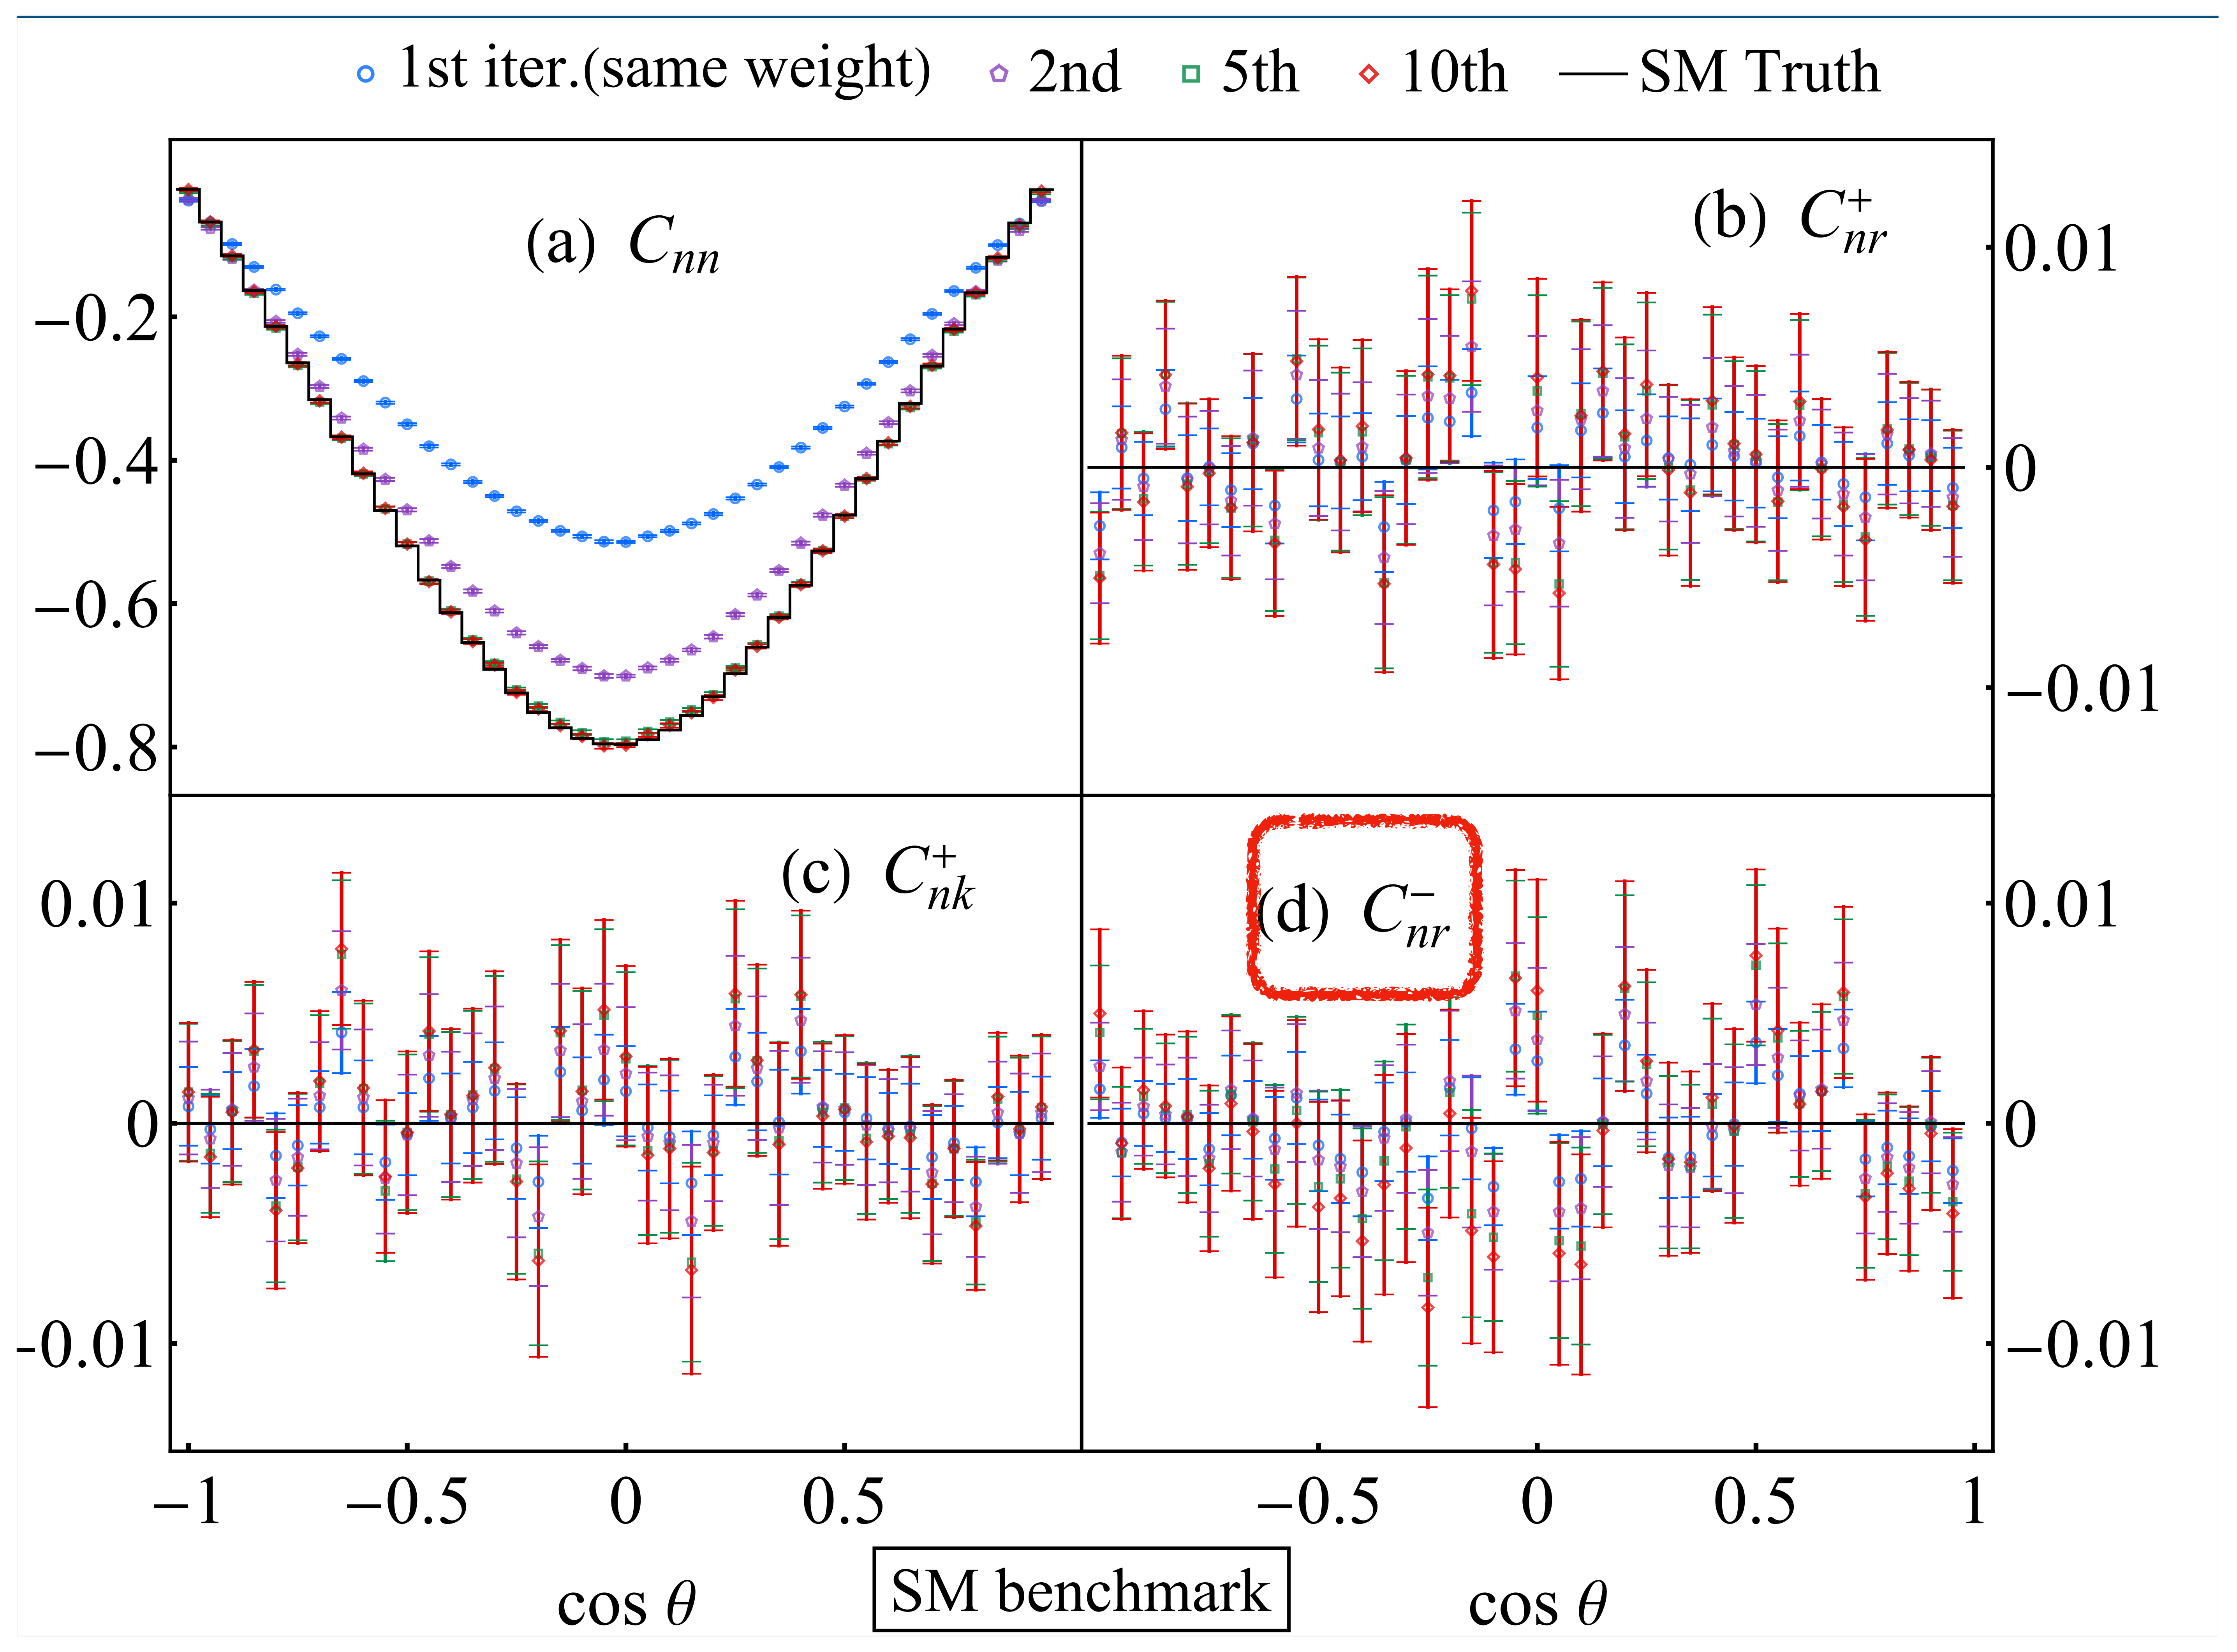

Flat averaging is not tomography; self-consistency removes the fold-induced bias

Stress test: anomalous τ dipoles

- Non-SM benchmark:

$$a_\tau = 0.01 - 0.02i, \quad \tilde{d}_\tau \equiv d_\tau \frac{2m_\tau}{e} = -0.05 + 0.03i\,.$$

- Generate events with a non-SM production density matrix
- Apply the same visible-only reconstruction and spin-independent initialization
- No SM or non-SM production template is supplied to the iteration
- Flat averaging remains biased
- All identifiable components evolve toward the non-SM truth

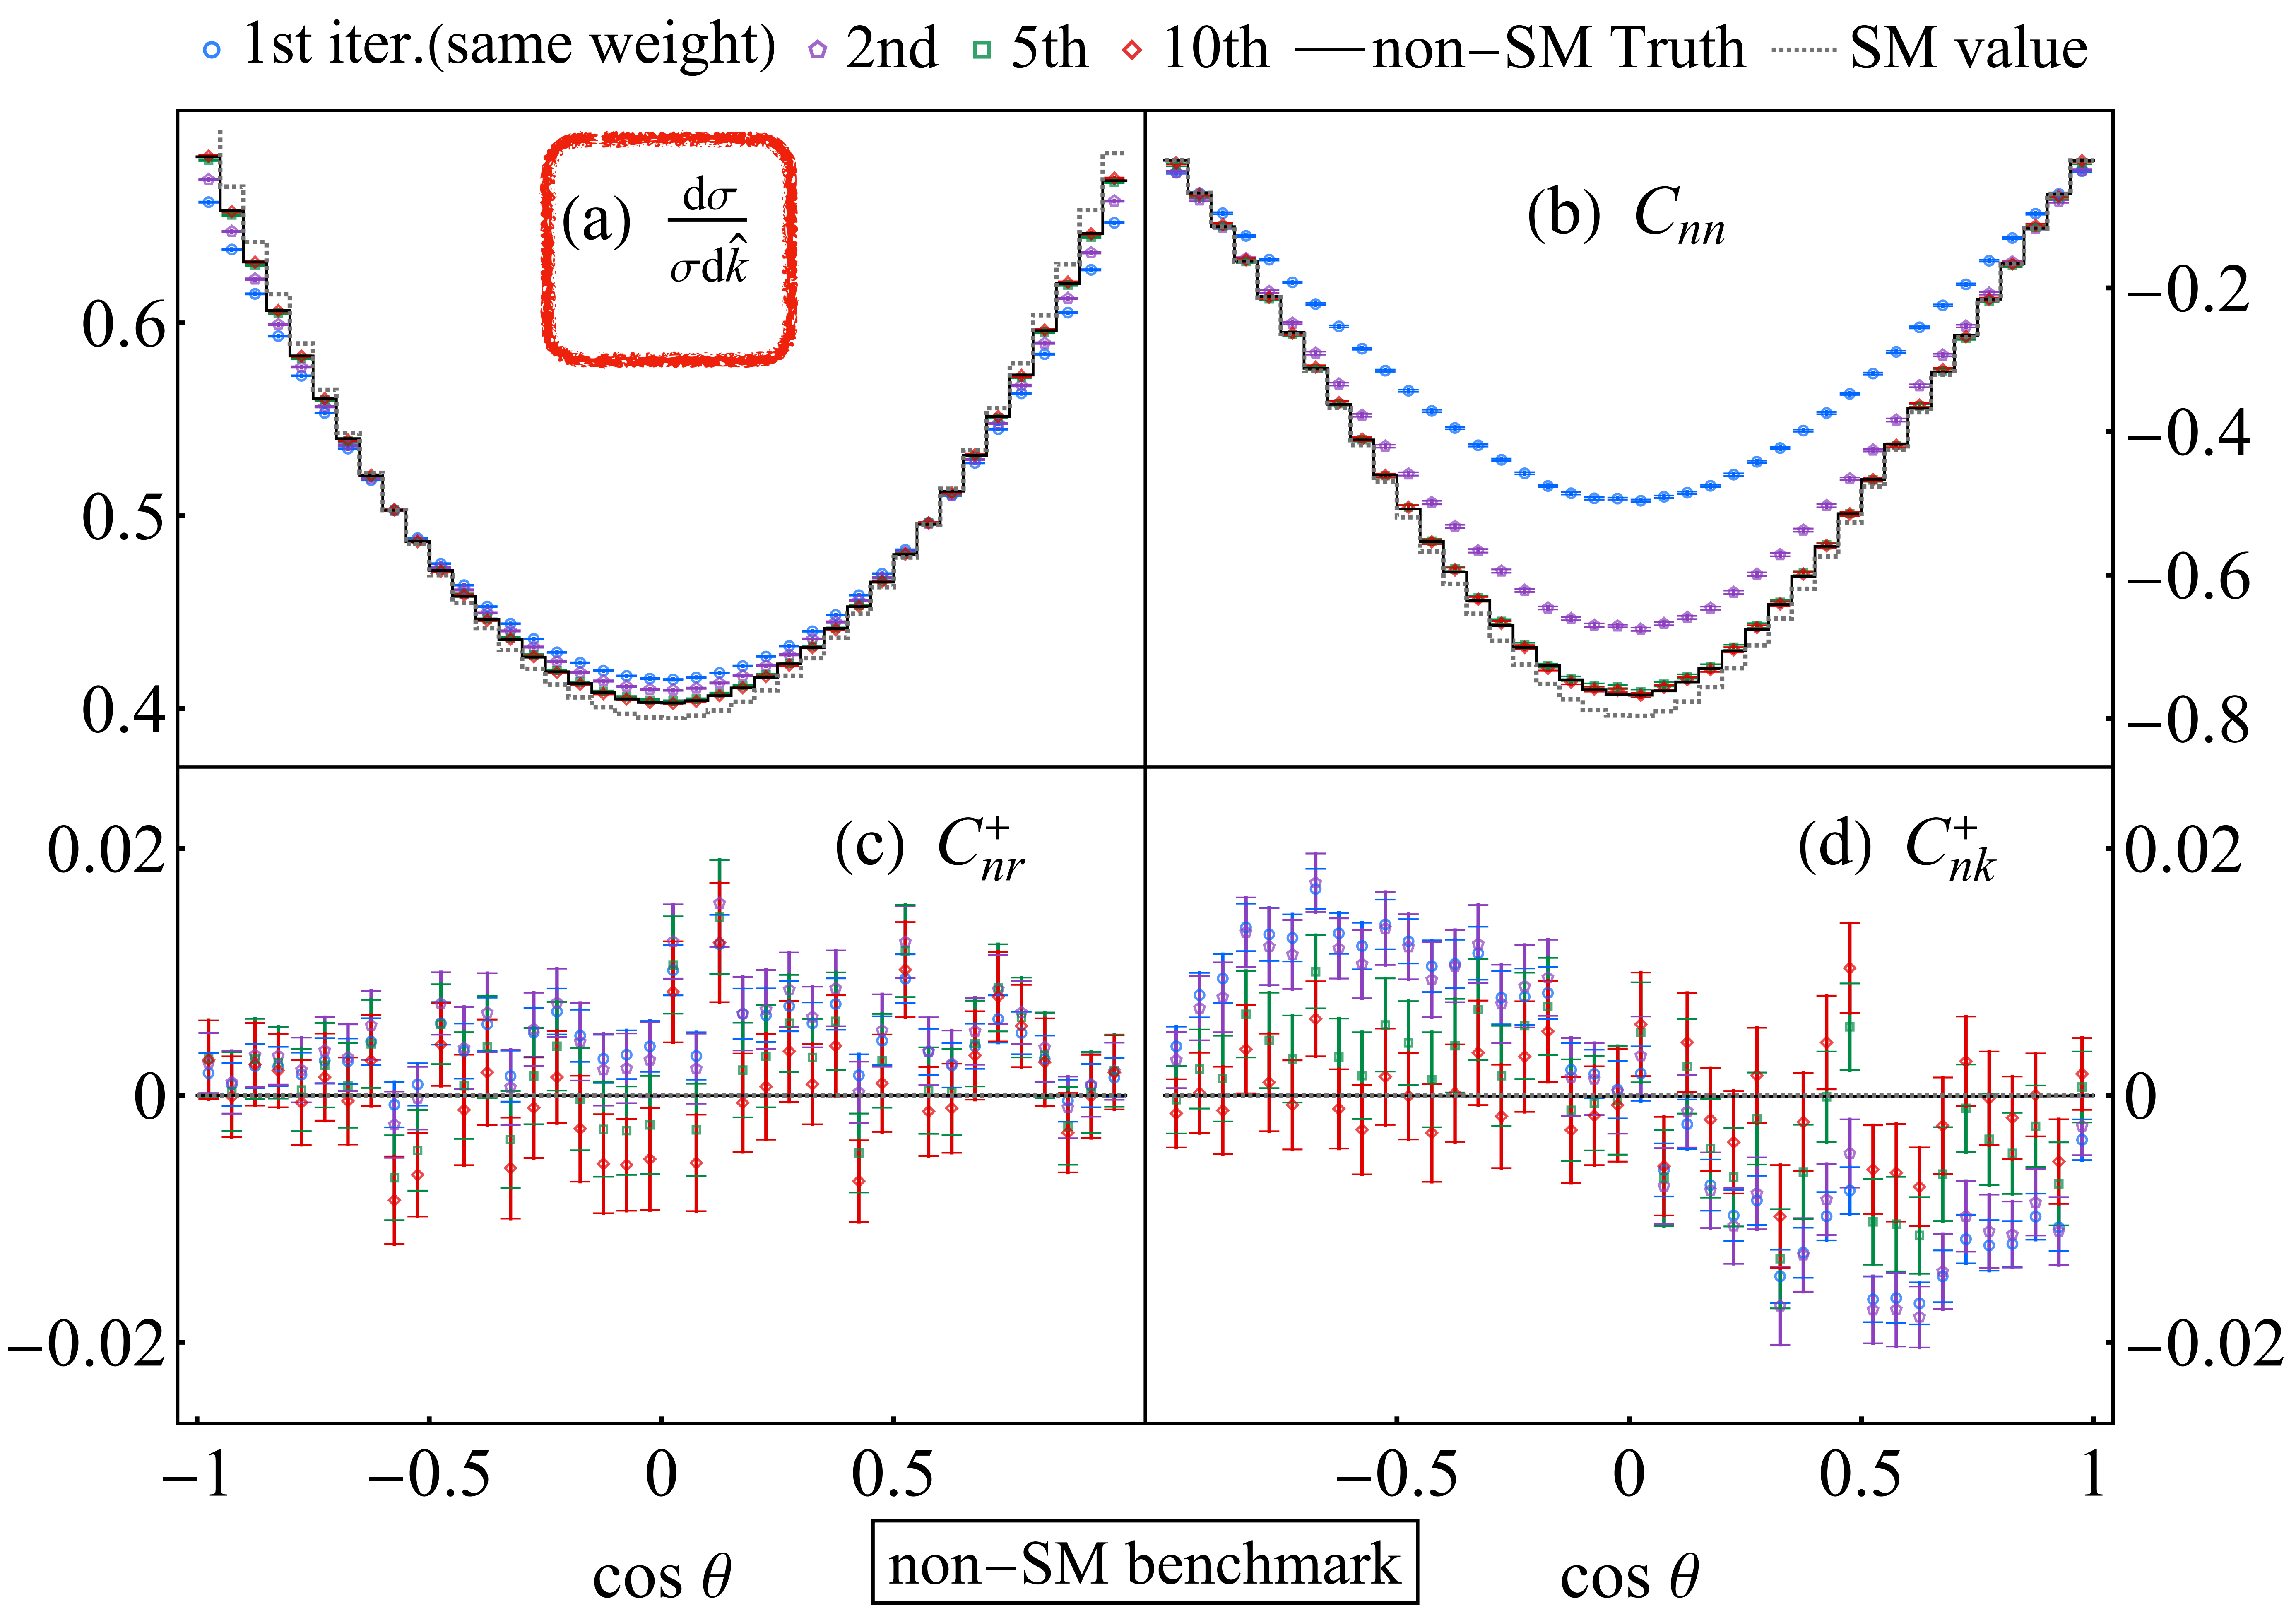

Model-independent algorithm + non-SM data $\Longrightarrow$ non-SM density matrix

The iteration follows the density matrix encoded in the data—not an assumed SM template

Stress test: anomalous τ dipoles

- Other spin-correlations reconstructed successfully

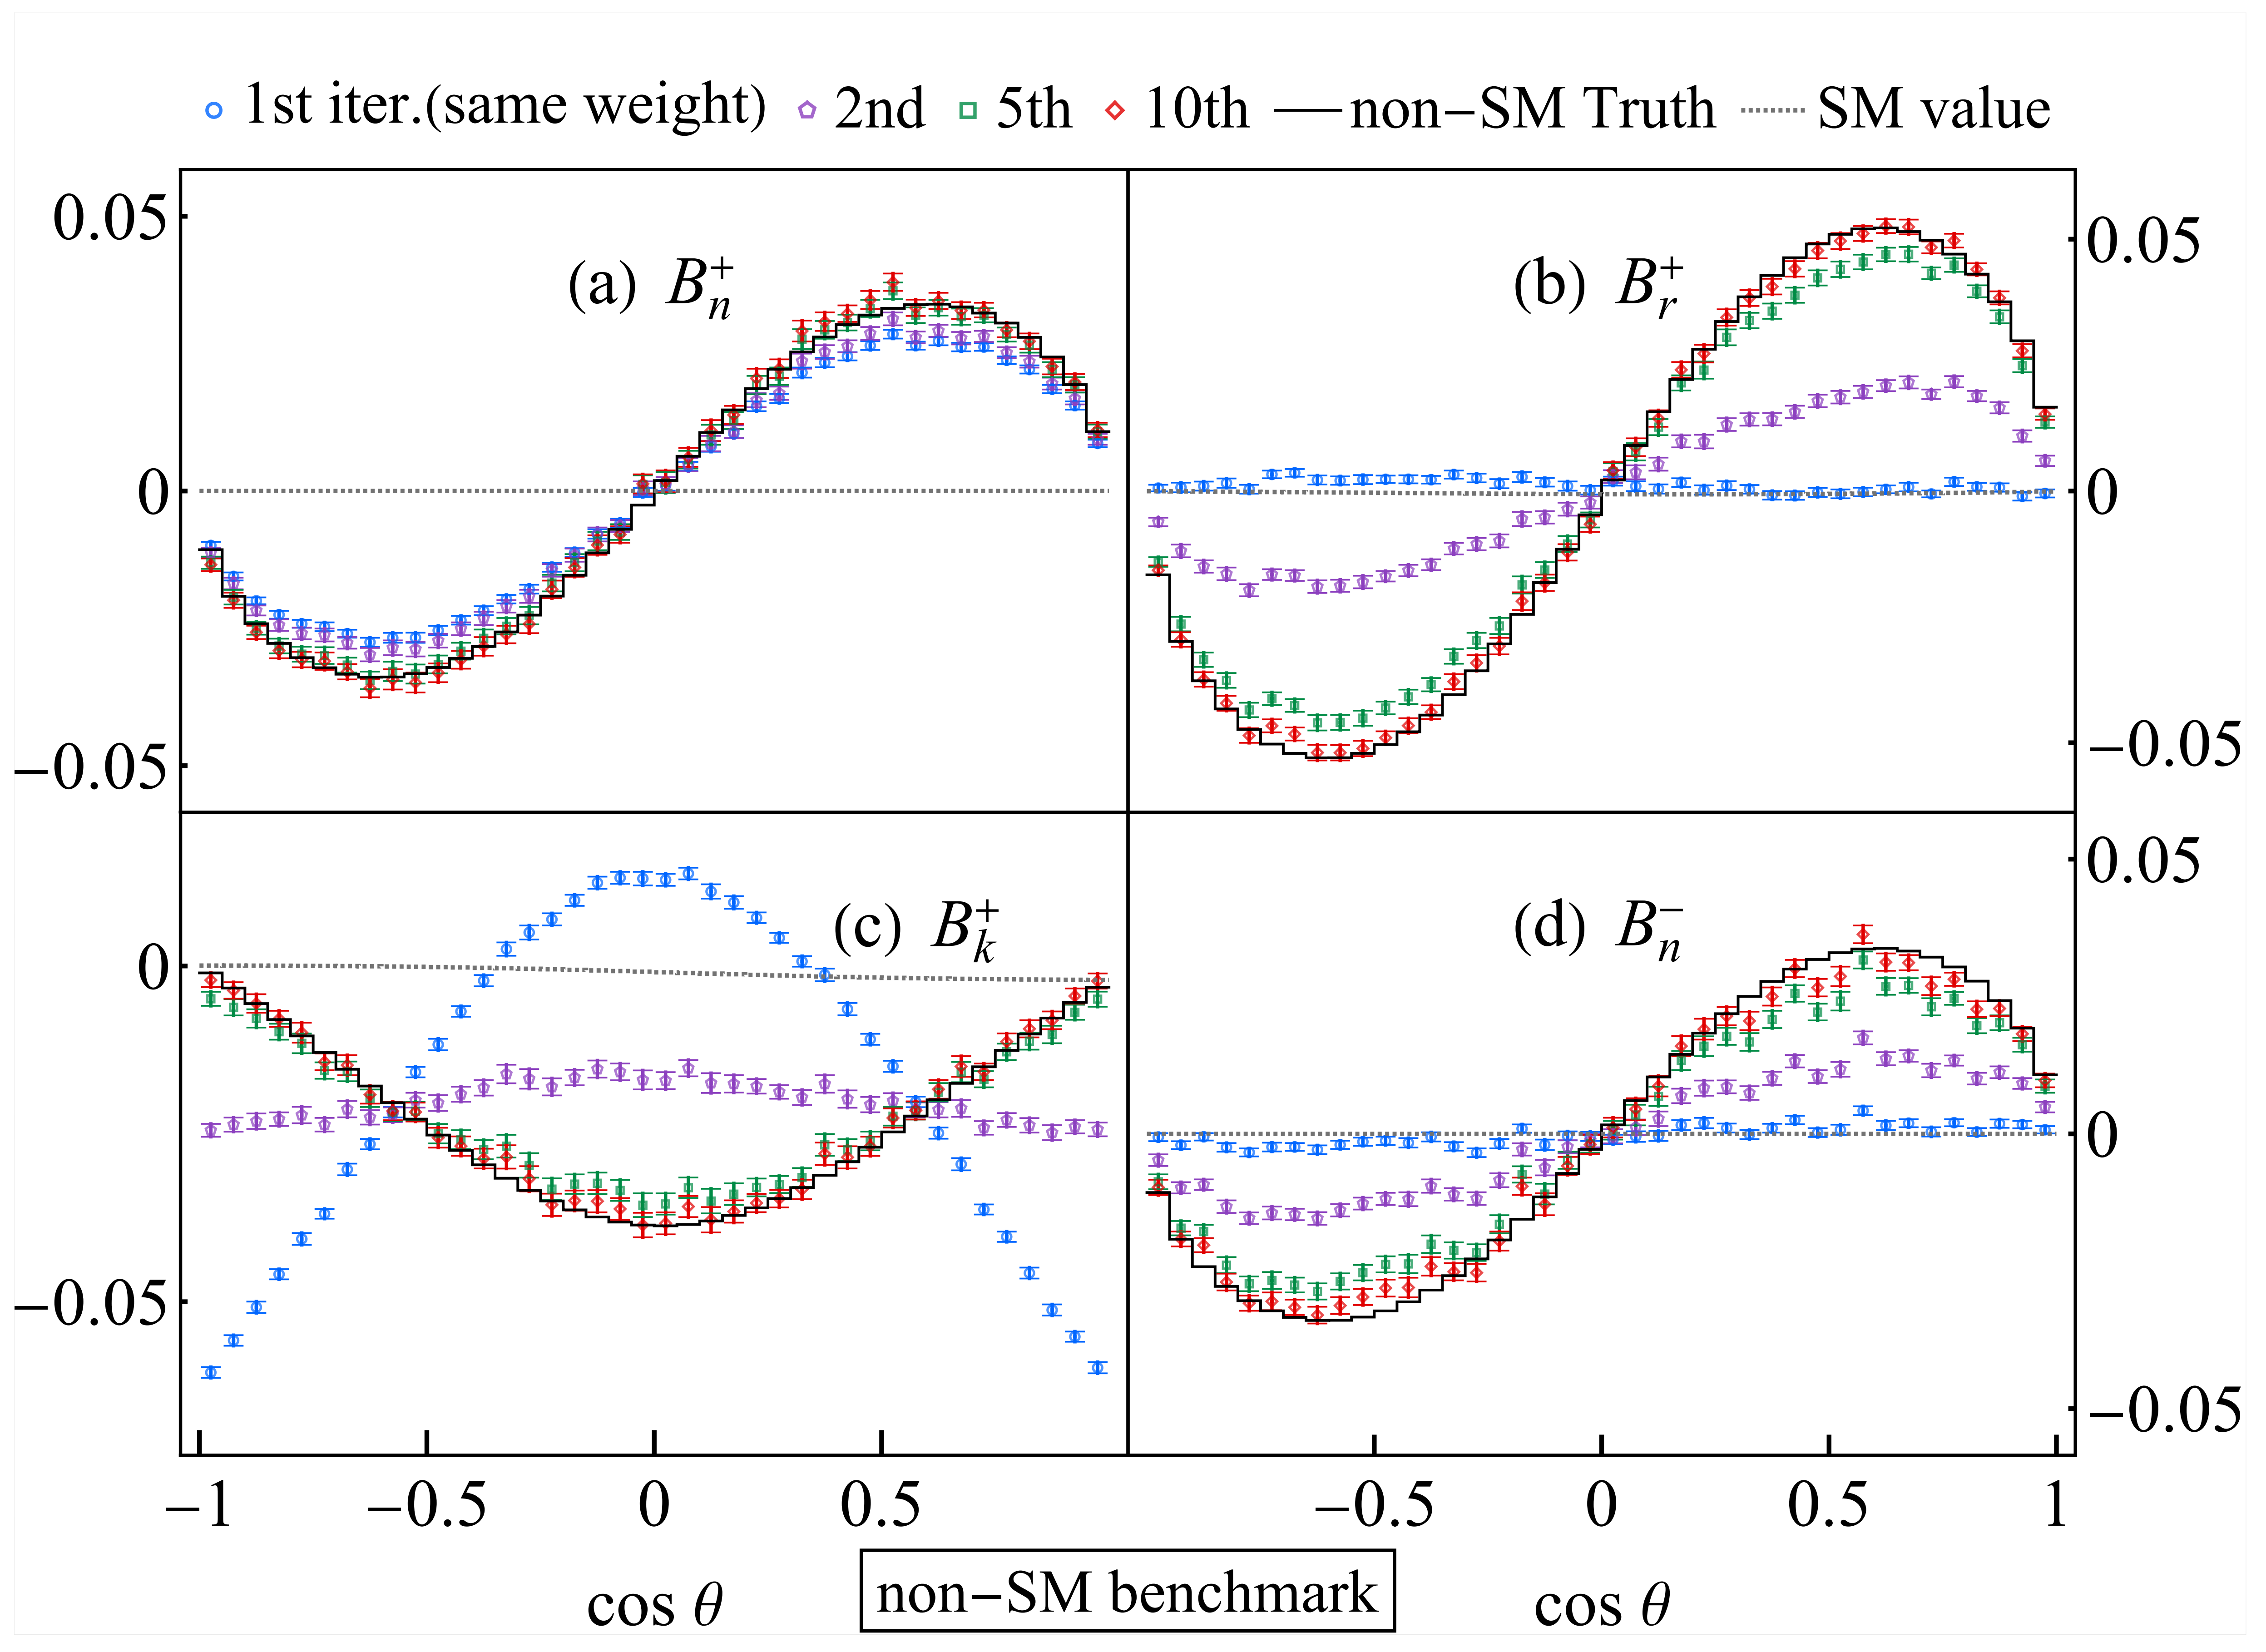

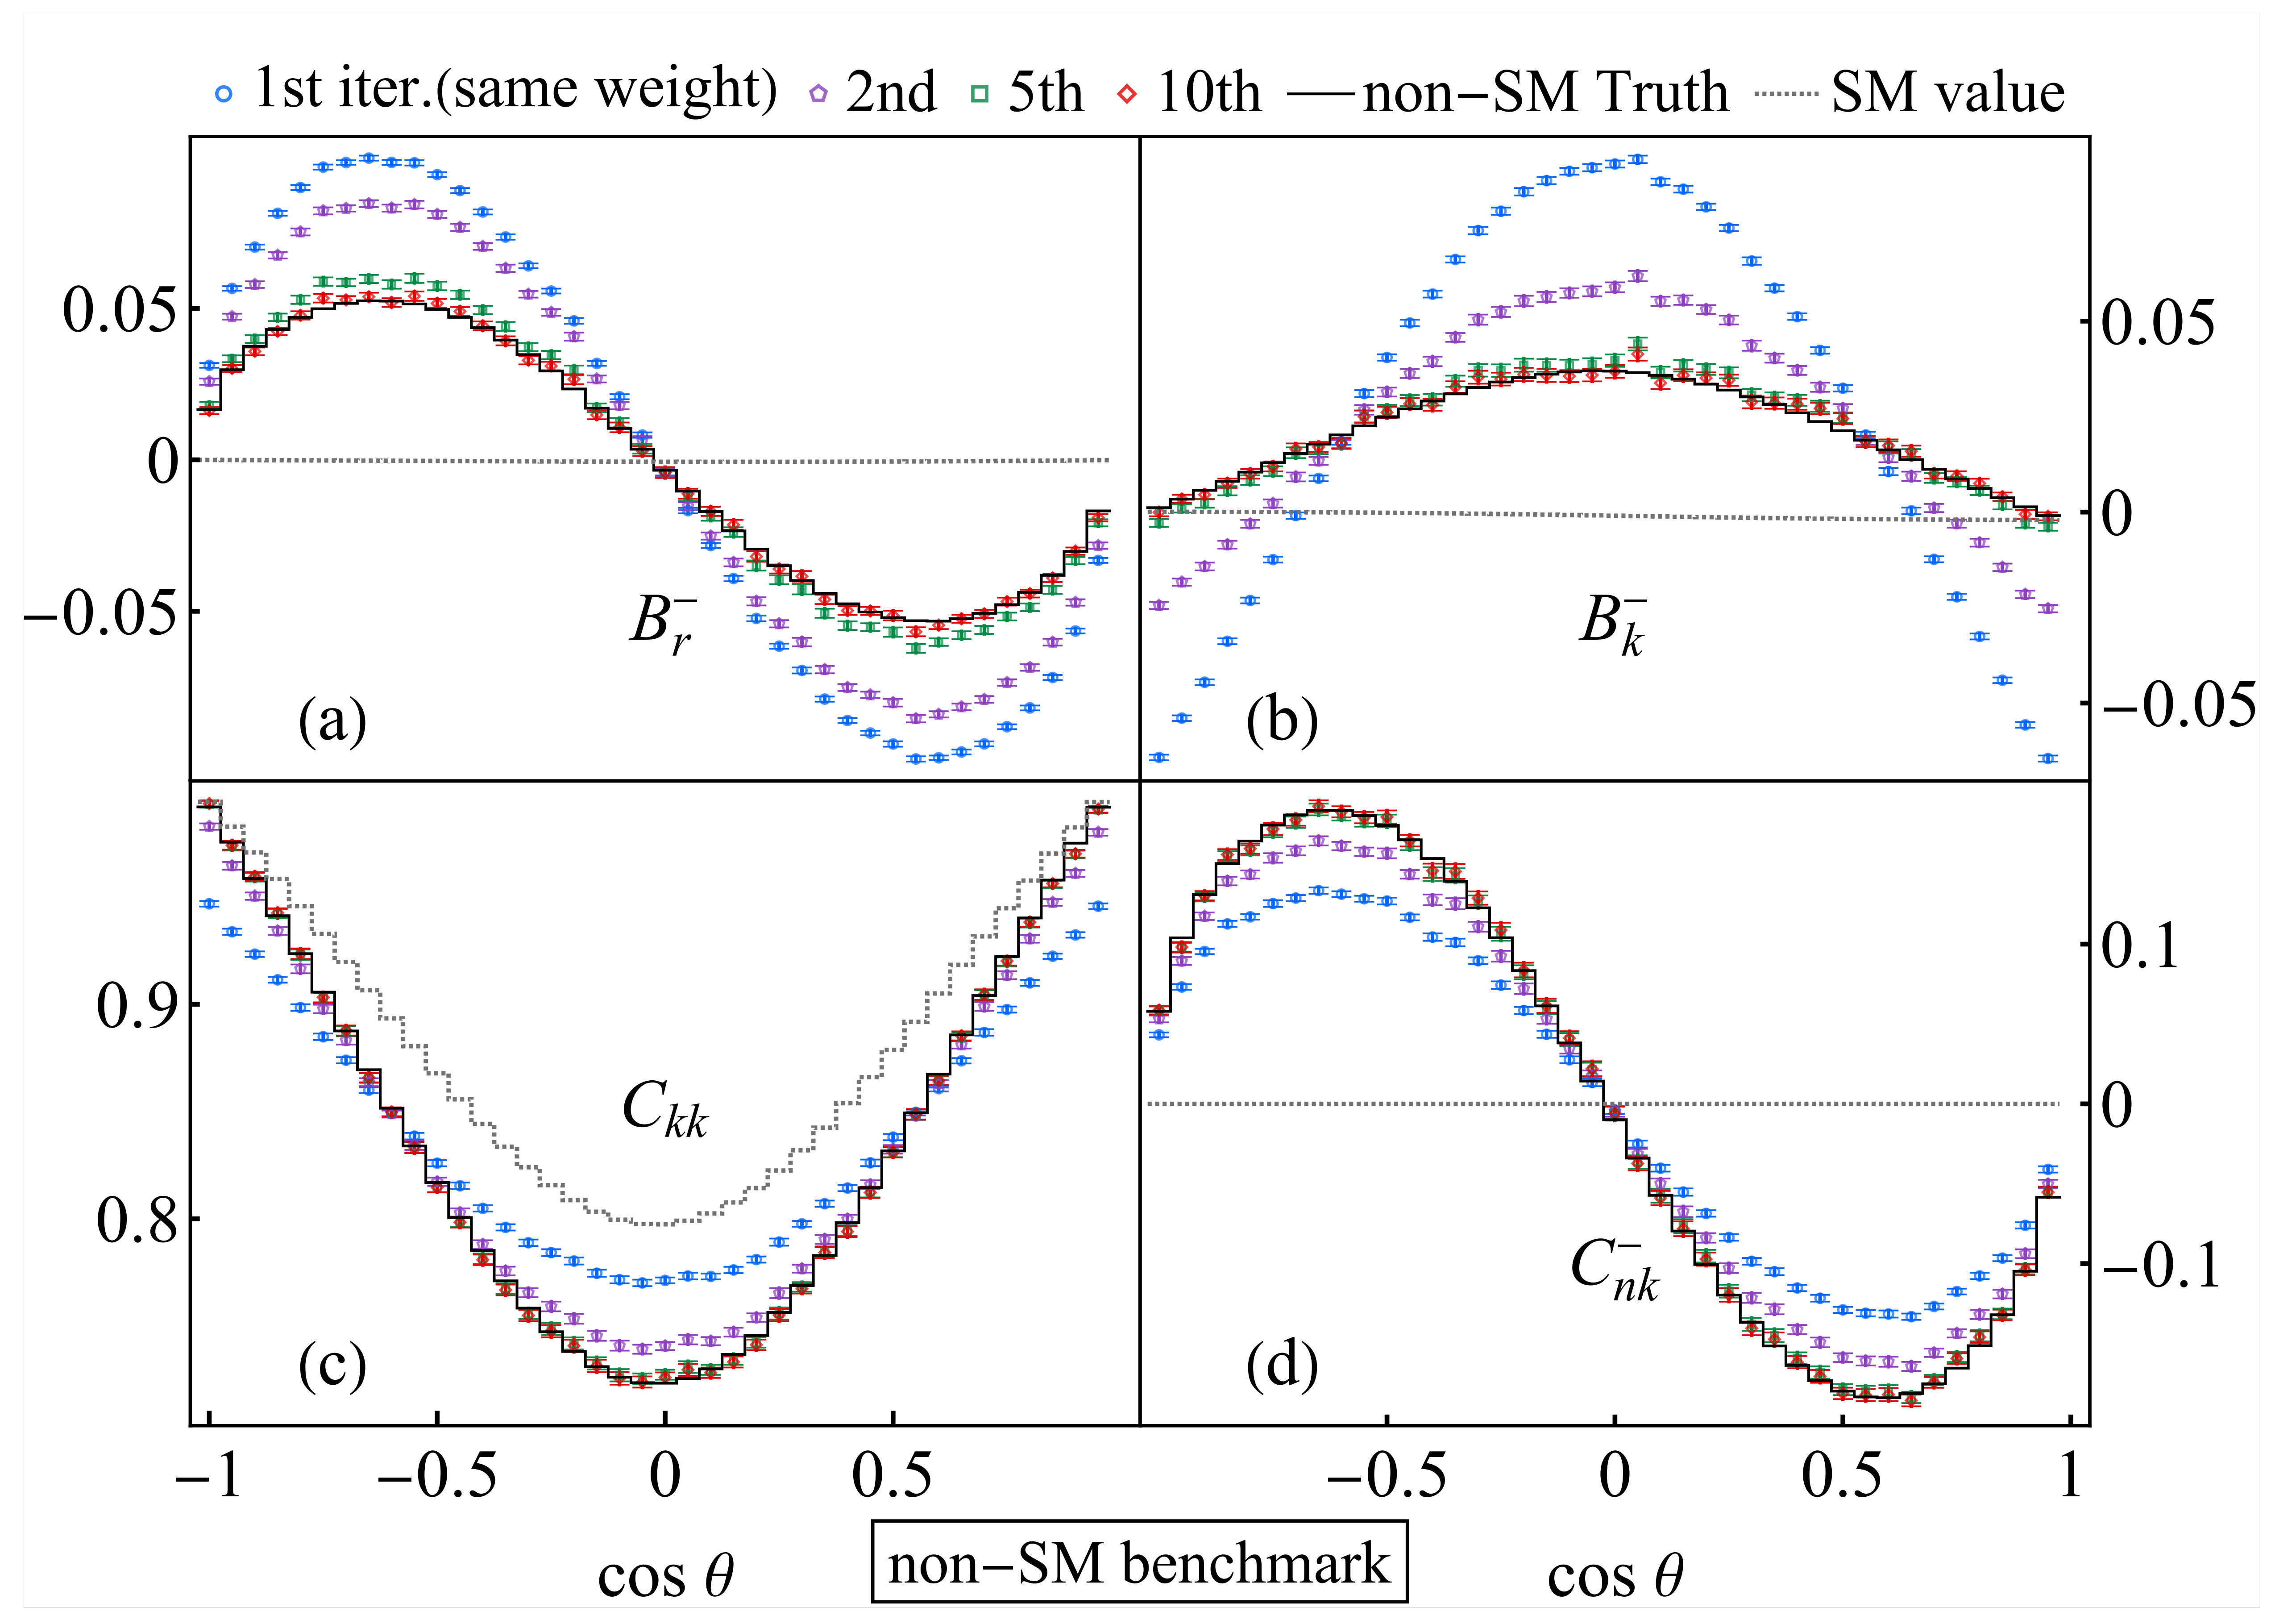

Model-independent algorithm + non-SM data $\Longrightarrow$ non-SM density matrix

The iteration follows the density matrix encoded in the data—not an assumed SM template

Stress test: anomalous τ dipoles

- Others reconstructed successfully
- But is C^{-}_{nr} constructed? In non-SM model, it does not converge to the truth value.
- What if more iterations?

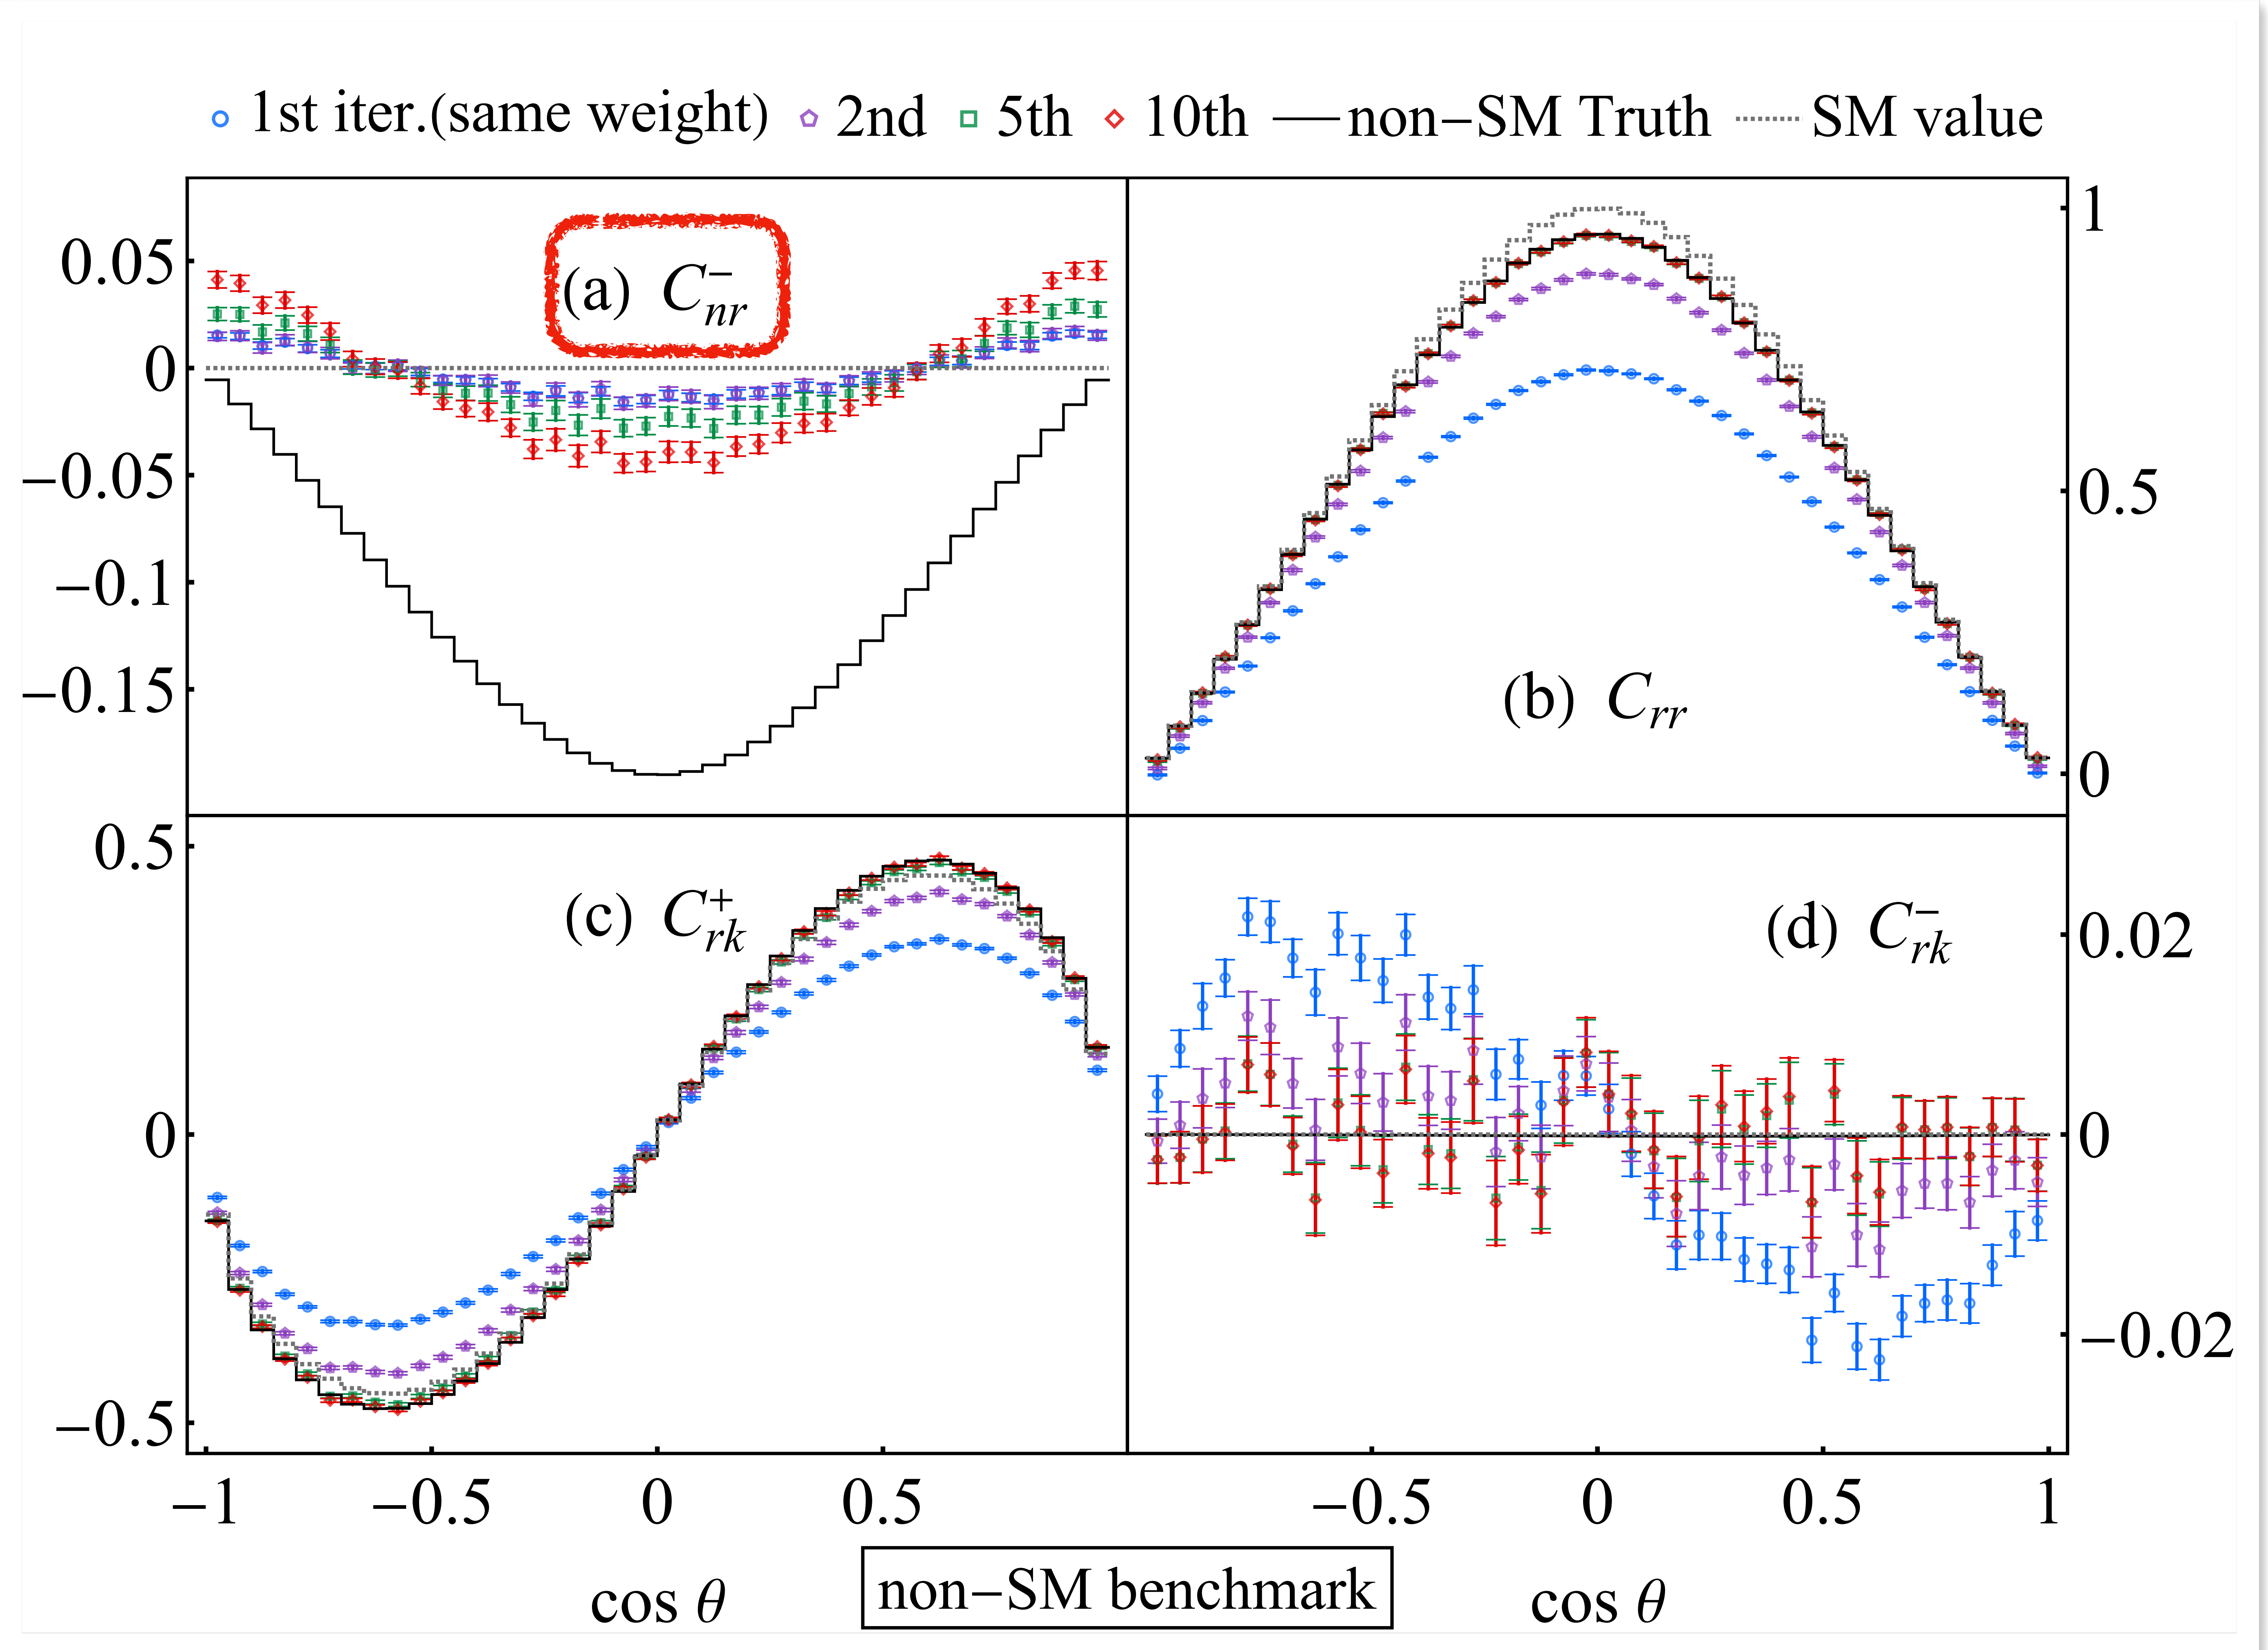

Model-independent algorithm + non-SM data $\Longrightarrow$ non-SM density matrix

The iteration follows the density matrix encoded in the data—not an assumed SM template

The null component C^{-}_{nr} behaves differently

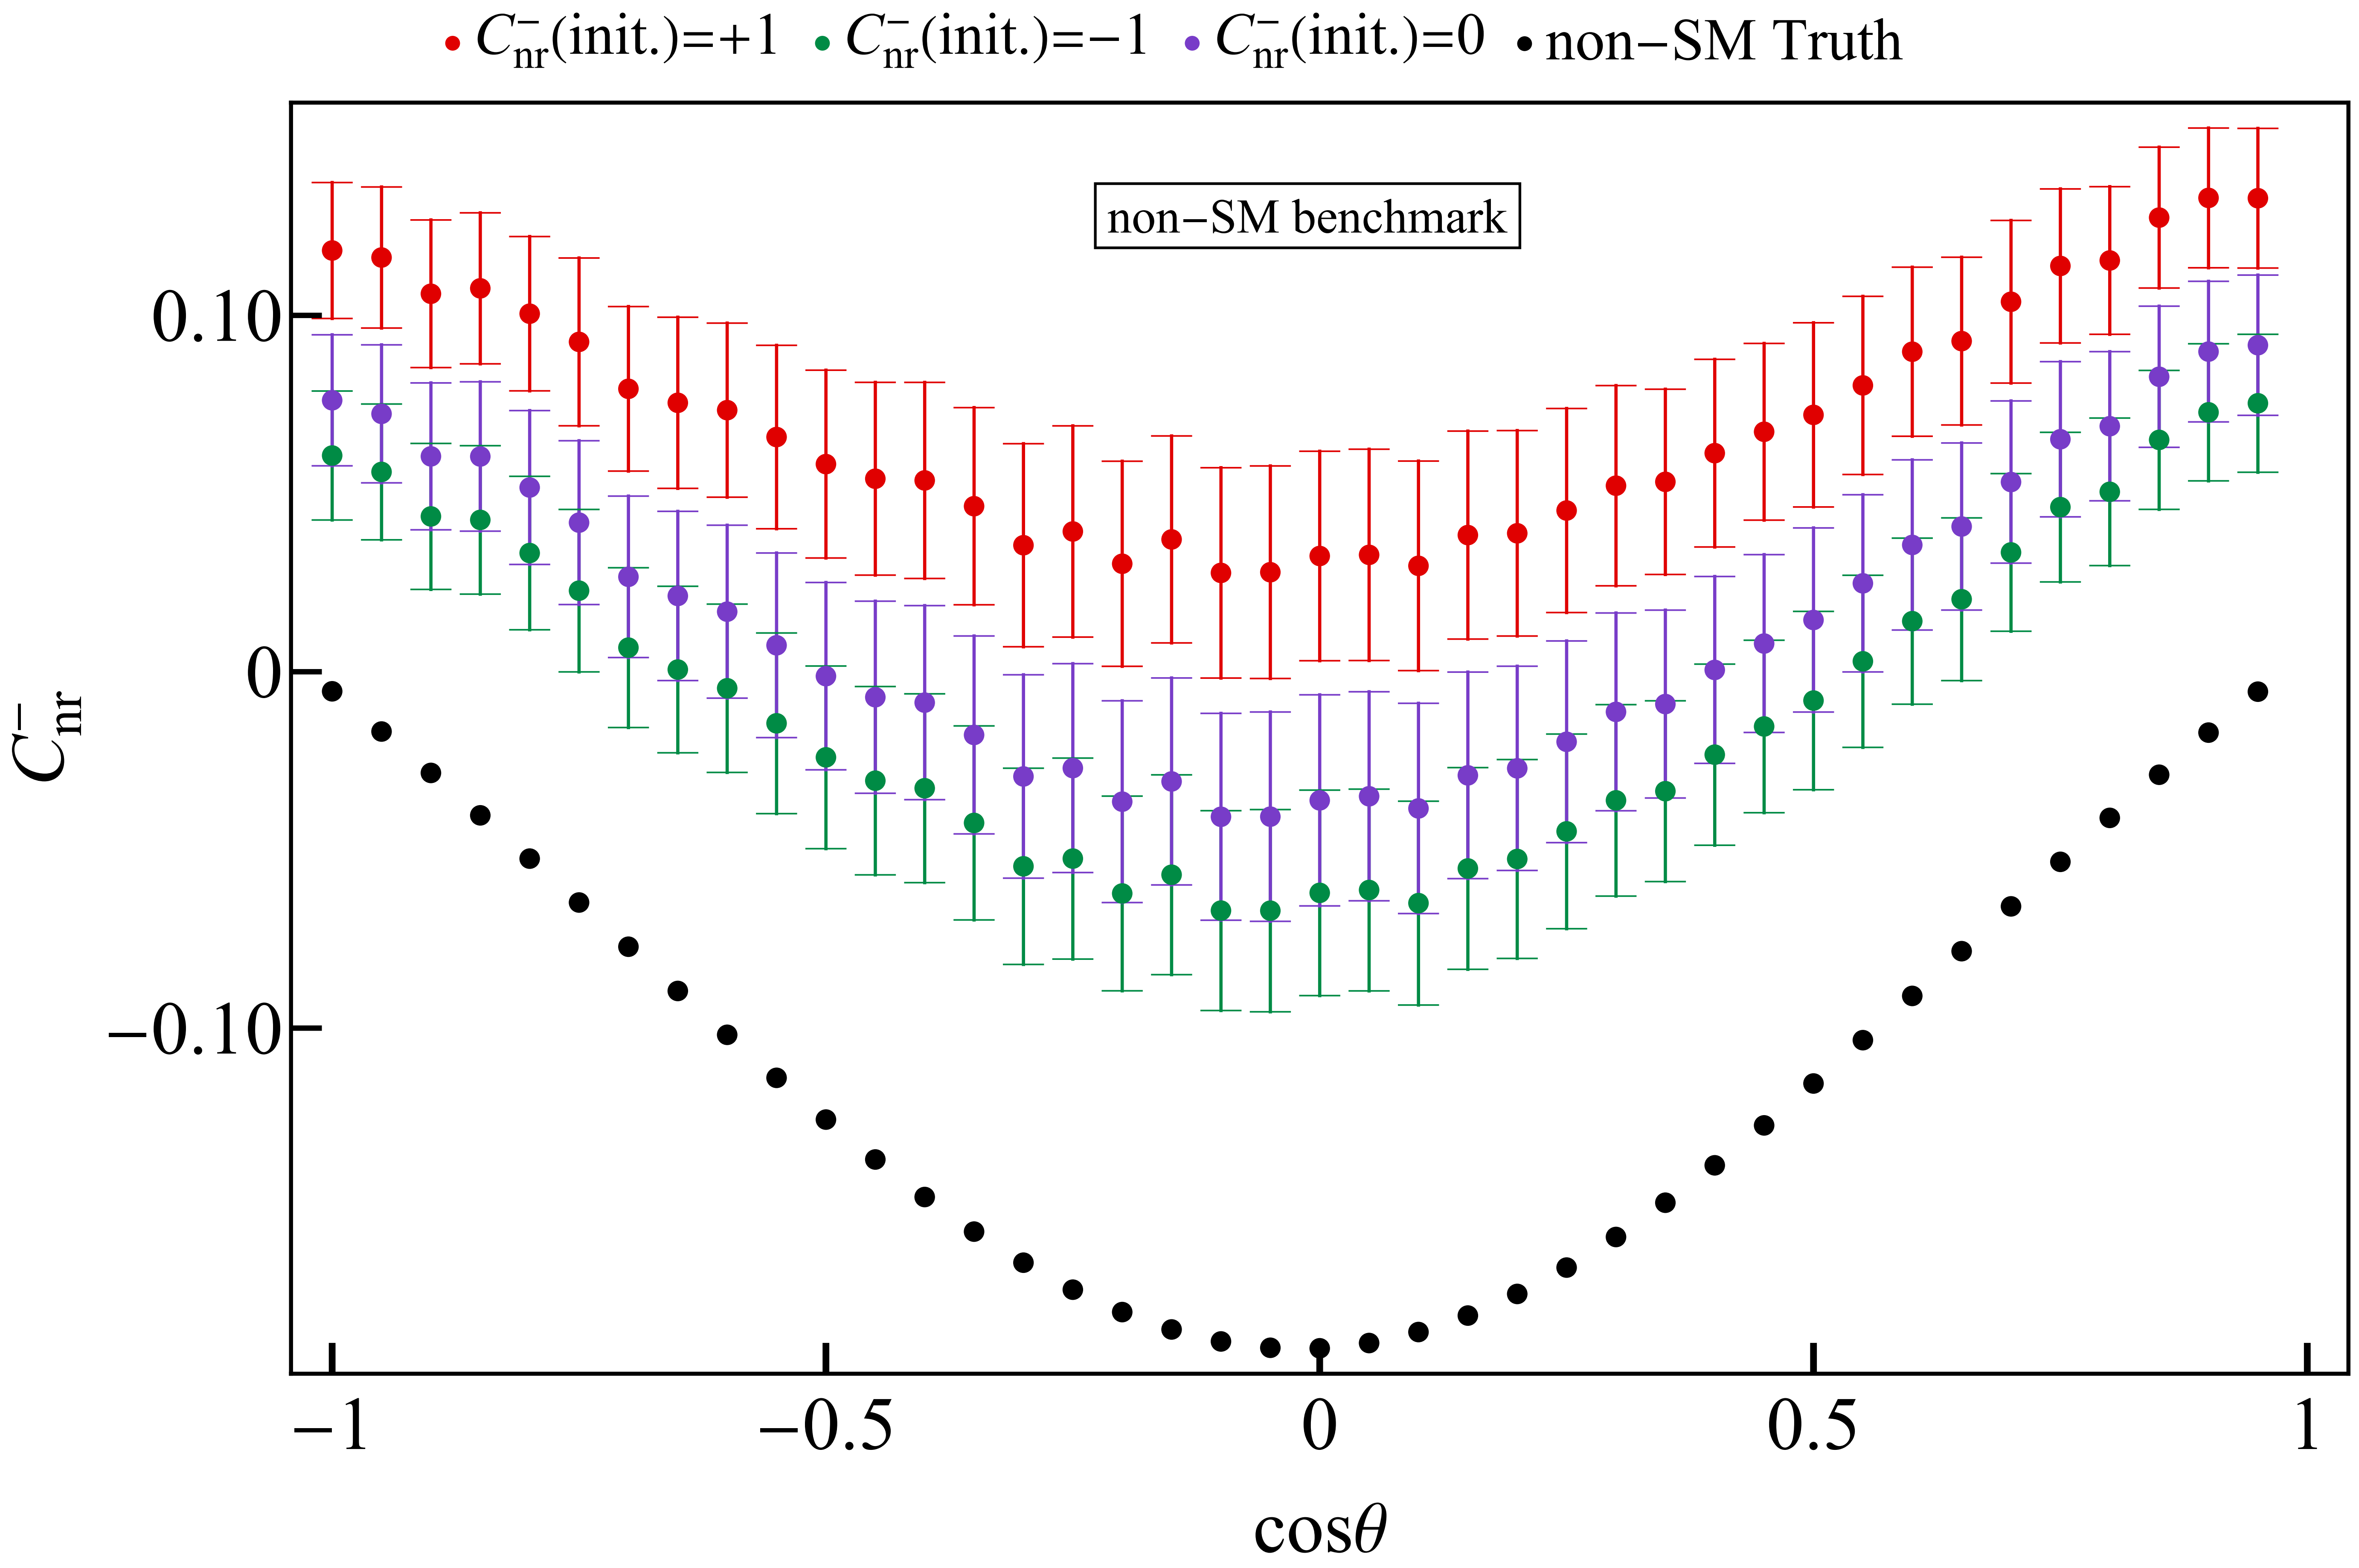

- Identifiable components converge toward the truth very quickly (< 10 iterations)
- The null component C^{-}_{nr} is not fixed by visible data
- Different initializations lead to different stabilizing point, but are consistent with the fixed-point family

$$\mathbf{x}_{*}(d_{\rm null}) = \mathbf{x}_{\rm truth} + d_{\rm null}\,\mathbf{n}_{\rm null}$$
$$\delta C_{nr}(\hat{k}) = -\delta C_{rn}(\hat{k}) = d_{\rm null}/\mathscr{R}(\hat{k}),$$

- 4000 iterations stabilizes
- Iter - truth = $\delta C^{-}_{nr}(\hat{k}) \simeq d_{\rm null} \times \left(\dfrac{{\rm d}\sigma}{\sigma {\rm d}\hat{k}}\right)^{-1}$; The numeric difference is consistent with analytic results.

Positivity still constrains quantum observables

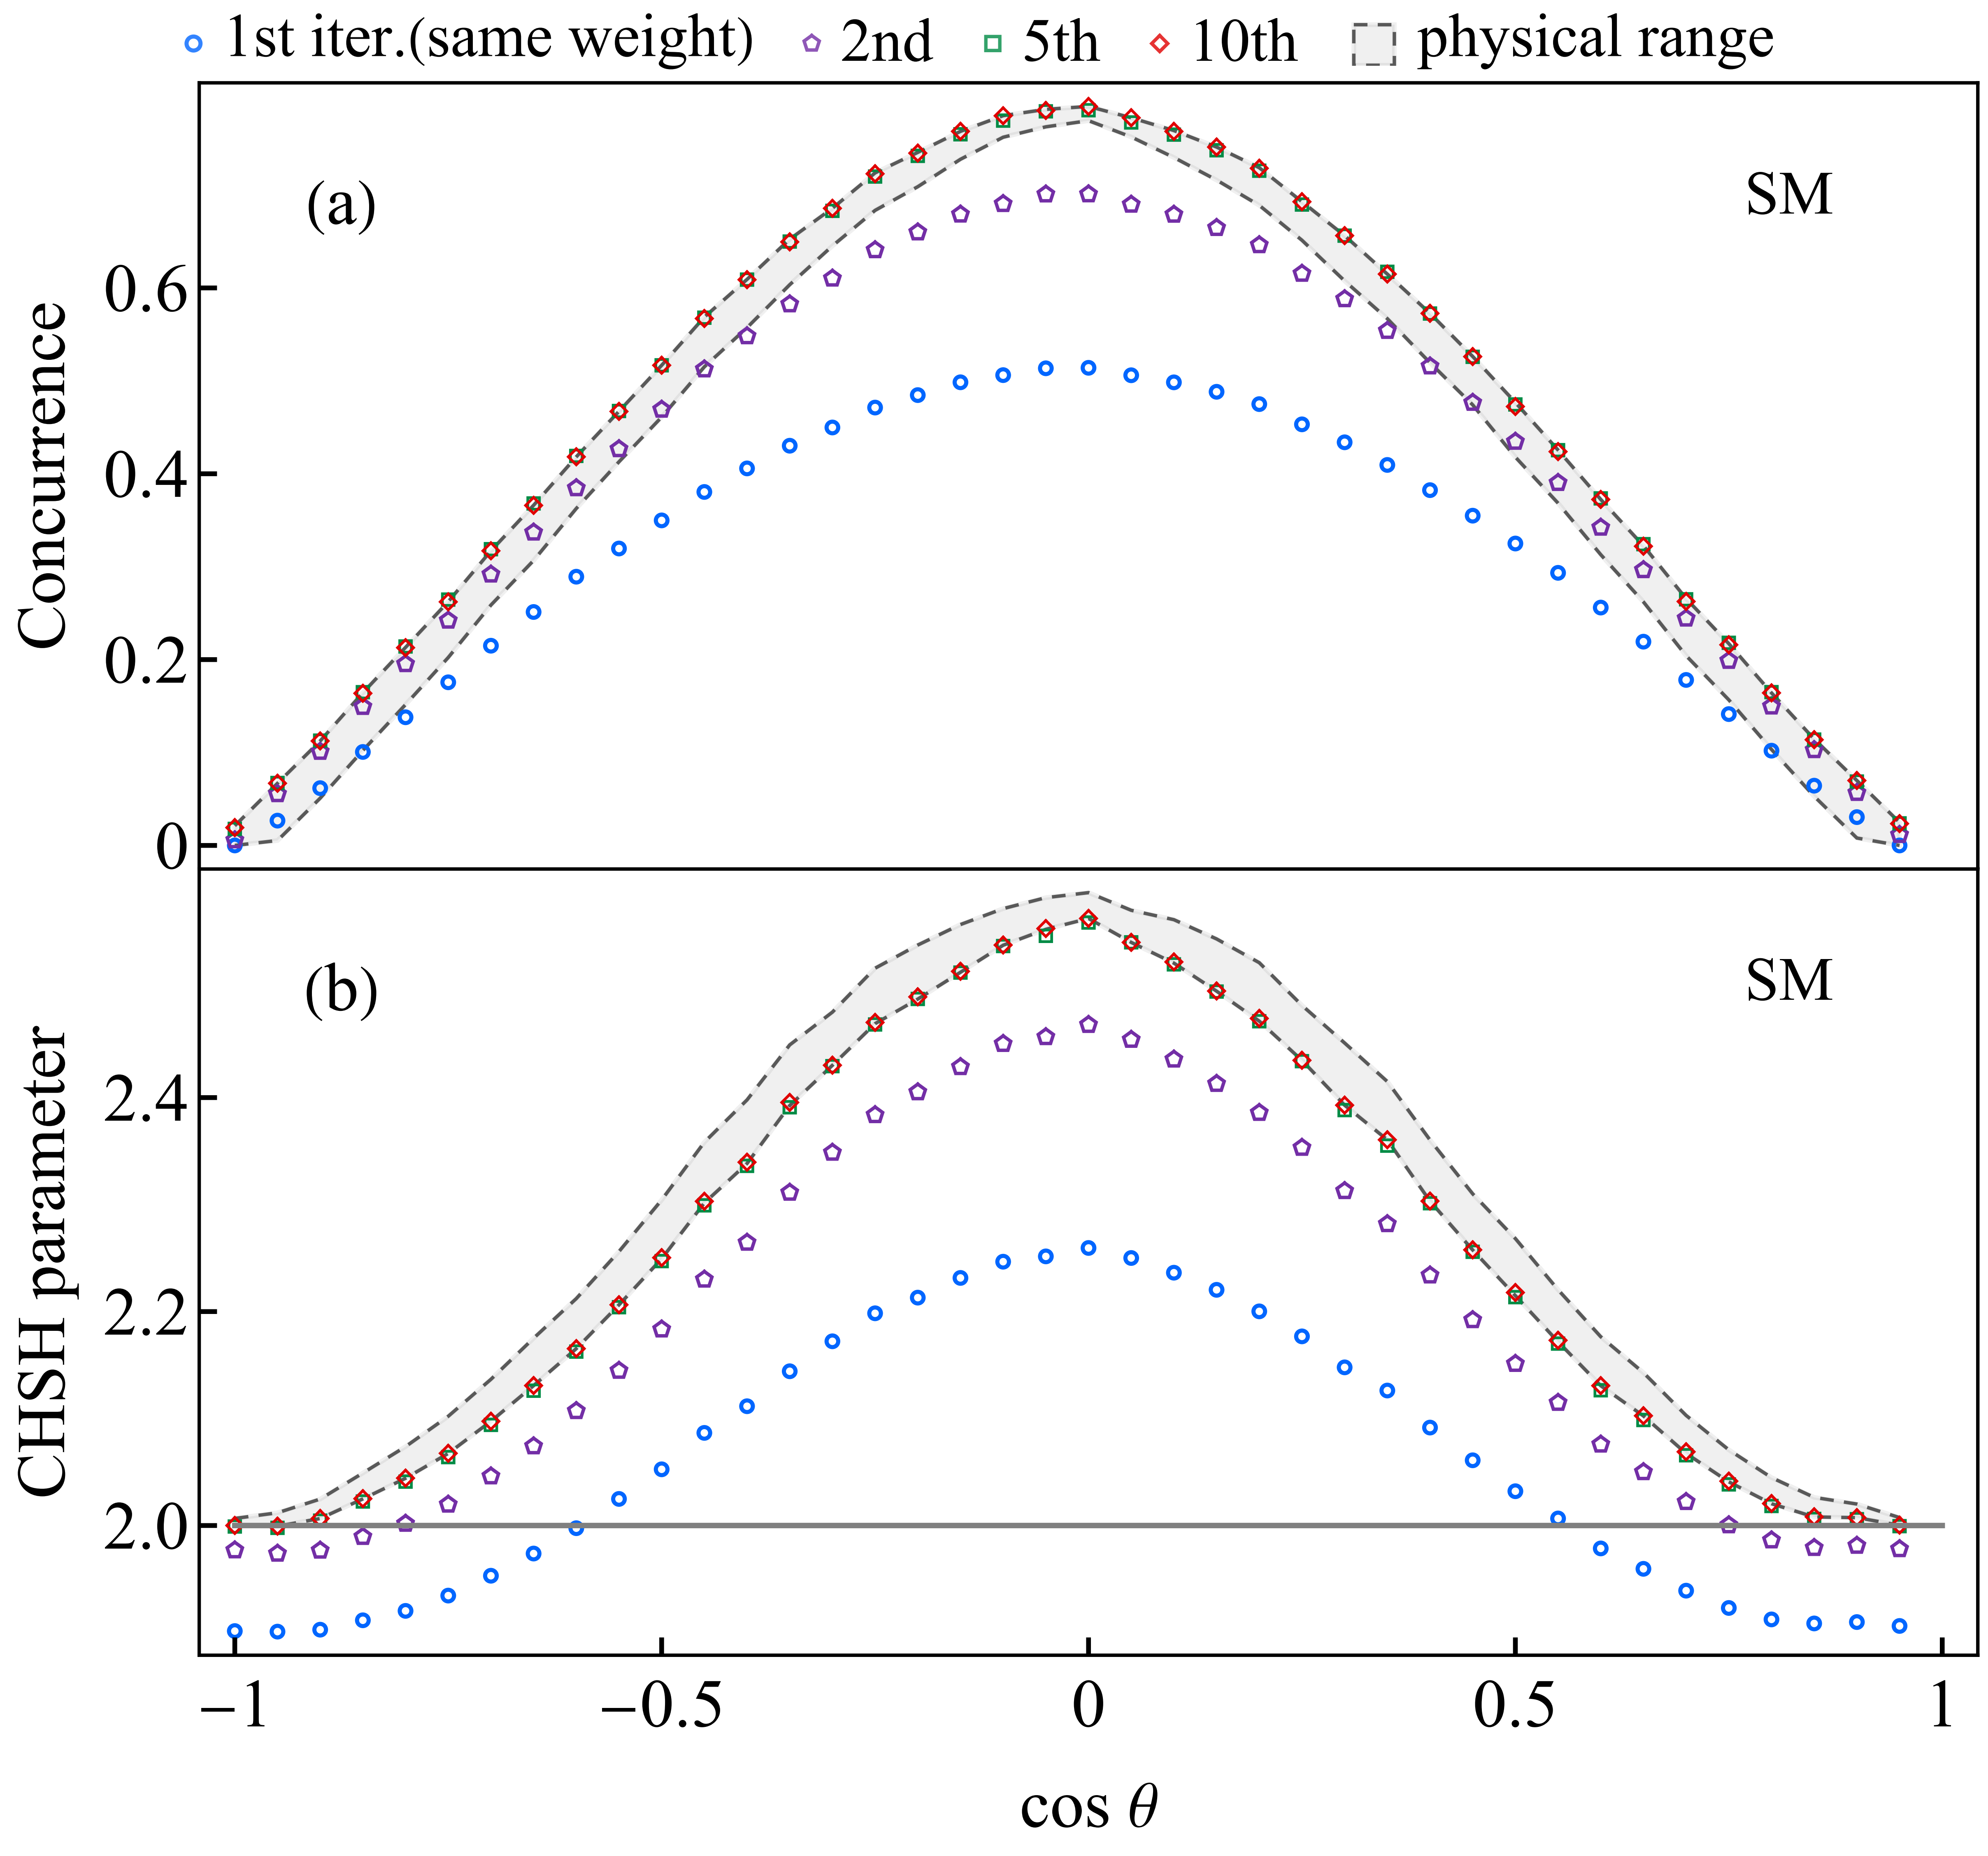

- Positivity restricts the fixed-point family
- Concurrence and CHSH need not have unique reconstructed values
- They can nevertheless be bounded using visible data and $\rho(\hat{k}) \succeq 0$
- The positivity-constrained intervals contain the truth in the closure tests
- Flat averaging can underestimate the quantum observables

Unidentifiable does not mean unconstrained.

Implications for BELLE I/II, BESIII and STCF

- The procedure of model-independent reconstruction under kinetic ambiguity

$$\textbf{visible event } \Phi \Longrightarrow \{\phi^a, \phi^b\} \Longrightarrow \textbf{visible-map tomography} \Longrightarrow \begin{cases} \textbf{identifiable } \rho \\ \textbf{positivity ranges} \end{cases}$$

- Practical next steps:
  - Include detector acceptance and efficiency
  - Include momentum resolution and backgrounds
  - Combine additional τ decay channels
  - Re-evaluate the null space for the full experimental measurement

Kinematic ambiguity can be modeled,
not simply averaged away or branch-selected.

Summary

- Missing neutrinos turn ideal spin measurements into coarse-grained visible measurements
- Information loss is determined by the null space of the visible map
- For $e^+e^- \to \tau^+\tau^- \to \pi^+\pi^-\bar{\nu}\nu$, only $C_{nr} - C_{rn}$ is unidentifiable
- All other differential spin information can be reconstructed without a production template
- Positivity gives controlled ranges for null-sensitive quantum observables

Kinematic ambiguity does not destroy tomography; it defines which part of tomography is identifiable

Keep all folds $\Longrightarrow$ find the null space $\Longrightarrow$ reconstruct the identifiable density matrix

Thanks you!