Supporting Information for

Randomized Iterative Trajectory Reweighting for Steady-State Distributions Without Discretization Error

Sagar Kania, Robert J. Webber, Gideon Simpson, David Aristoff, Daniel M. Zuckerman

Daniel M. Zuckerman
E-mail: zuckermd@ohsu.edu

This PDF file includes:

Supporting Information Text

Analysis of RiteWeight fixed point

This section identifies the fixed point of the RiteWeight algorithm, which is defined as follows.

Definition 1 (Fixed point) *A collection of trajectory weights $(w_i)_{i=1}^N$ with $\sum_i w_i = 1$ is a fixed point of RiteWeight if the weights stay the same during each RiteWeight iteration for each positive probability choice of partition.*

In practice, we typically apply RiteWeight in a continuous state space with clusters defined by a random Voronoi tesselation. Here for simplicity our analysis considers a finite state space with clusters defined by a random hyperplane tesselation. See (1) for background on random hyperplane tesselations, which are called "stable under iteration tesselations" in stochastic geometry.

Assumption 1 (Finite state space) *The state space consists of a finite number of distinct microstates $\alpha \in \mathbb{R}^d$.*

Assumption 2 (Random hyperplane tesselation) *The RiteWeight clusters are generated through one or more iterations of the following procedure. Initially, there is a single cluster containing all the microstates. At each iteration, any cluster C that contains at least two microstates is randomly split by a hyperplane into two new clusters as follows. First, we choose a uniformly random direction*

$$u \sim \mathrm{Unif}\{v \in \mathbb{R}^d \,:\, \|v\| = 1\}.$$

Then, conditional on u, we choose a uniformly random offset

$$\gamma \sim \mathrm{Unif}\left\{\eta \in \mathbb{R} \,:\, \min_{\alpha \in C} \alpha^\top u < \eta < \max_{\alpha \in C} \alpha^\top u\right\}.$$

The normal direction and the offset define two new clusters of microstates,

$$\{\alpha \in C : \alpha^\top u < \gamma\} \quad \textit{and} \quad \{\alpha \in C : \alpha^\top u > \gamma\}$$

that are split by the hyperplane.

Under Assumptions 1 and 2, the following main result shows that the RiteWeight fixed point is completely determined by the microstate transition matrix. The microstate transition matrix does not change with the RiteWeight iterations. Therefore, assuming RiteWeight converges, it can only converge to a single fixed point.

Theorem 1 (RiteWeight fixed point) *Consider a finite state space (Assumption 1) and a collection of trajectories with weights $(w_i)_{i=1}^N$ satisfying $\sum_i w_i = 1$. Define the associated microstate transition matrix $\boldsymbol{P}$ with entries*

$$P_{\alpha\beta} = \frac{\sum_{i_1 \in \alpha,\, i_2 \in \beta} w_i}{\sum_{i_1 \in \alpha} w_i}, \qquad [1]$$

and assume $\boldsymbol{P}$ has a unique stationary measure. Then, under the random hyperplane model (Assumption 2), $(w_i)_{i=1}^N$ is a fixed point of RiteWeight if and only if the vector $\boldsymbol{\mu}$ with entries

$$\mu_\alpha = \sum_{i_1 \in \alpha} w_i \qquad [2]$$

is the unique fixed point of $\boldsymbol{P}$, that is,

$$\boldsymbol{\mu}^\top \boldsymbol{P} = \boldsymbol{\mu}^\top. \qquad [3]$$

First we check that equation Eq. (3) implies $(w_i)_{i=1}^n$ is a fixed point of RiteWeight. This happens if the weights $w_I = \sum_{i_1 \in I} w_i$ appear as entries of the left leading eigenvector of the cluster transition matrix $\boldsymbol{T}$. By assumption, the microstate transition matrix $\boldsymbol{P}$ has a unique stationary measure, so it has one closed communicating class. It follows that the cluster transition matrix $\boldsymbol{T}$ has one closed communicating class, so $\boldsymbol{T}$ has a unique stationary measure also. Next, equation Eq. (3) implies

$$\sum_I w_I T_{IJ} = \sum_\alpha \sum_{\beta \in J} \mu_\alpha P_{\alpha\beta} = \sum_{\beta \in J} \mu_\beta = w_J, \quad \text{for each cluster } J.$$

This confirms that the left leading eigenvector has the correct entries, so $(w_i)_{i=1}^n$ is a fixed point of RiteWeight.

Next, we assume that $(w_i)_{i=1}^n$ is a fixed point of RiteWeight and check that equation Eq. (3) holds. Since $(w_i)_{i=1}^n$ is a fixed point, each positive-probability cluster J must satisfy

$$\sum_{\alpha \in J} \mu_\alpha = w_J = \sum_I w_I T_{IJ} = \sum_\alpha \sum_{\beta \in J} \mu_\alpha P_{\alpha\beta}.$$

By considering all positive probability clusters, we arrive at the identity

$$\boldsymbol{\mu}^\top \boldsymbol{A} = \boldsymbol{\mu}^\top \boldsymbol{P} \boldsymbol{A}$$

where $\boldsymbol{A}$ is the matrix whose columns are characteristic functions for each positive probability random cluster: 1 values indicate membership in the cluster and 0 values indicate non-membership. To complete the proof, we will show that $\boldsymbol{A}$ has full column rank and therefore Eq. (3) holds.

We observe that the first iteration of the random hyperplane tesselation generates a uniformly random direction $u \in \mathbb{R}^d$ that leads to distinct values $\alpha^\top u$ for every distinct microstate α with probability one. Hence, there is an ordering of microstates, say $\alpha_1, \ldots, \alpha_n$, so that the event

$$\alpha_1^\top u < \cdots < \alpha_n^\top u,$$

occurs with positive probability. It follows that a division into clusters

$$\{\alpha_1, \ldots, \alpha_i\} \quad \text{and} \quad \{\alpha_{i+1}, \ldots, \alpha_n\}$$

occurs with positive probability for each $i = 1, \ldots, n-1$. With the microstates ordered in this way, the linear span of the columns of $\boldsymbol{A}$ includes all the vectors $\sum_{j=1}^{i} \boldsymbol{e}_j$, where $\boldsymbol{e}_i$ is a standard basis vector. This shows $\boldsymbol{A}$ has full column rank and equation Eq. (3) holds, completing the proof.

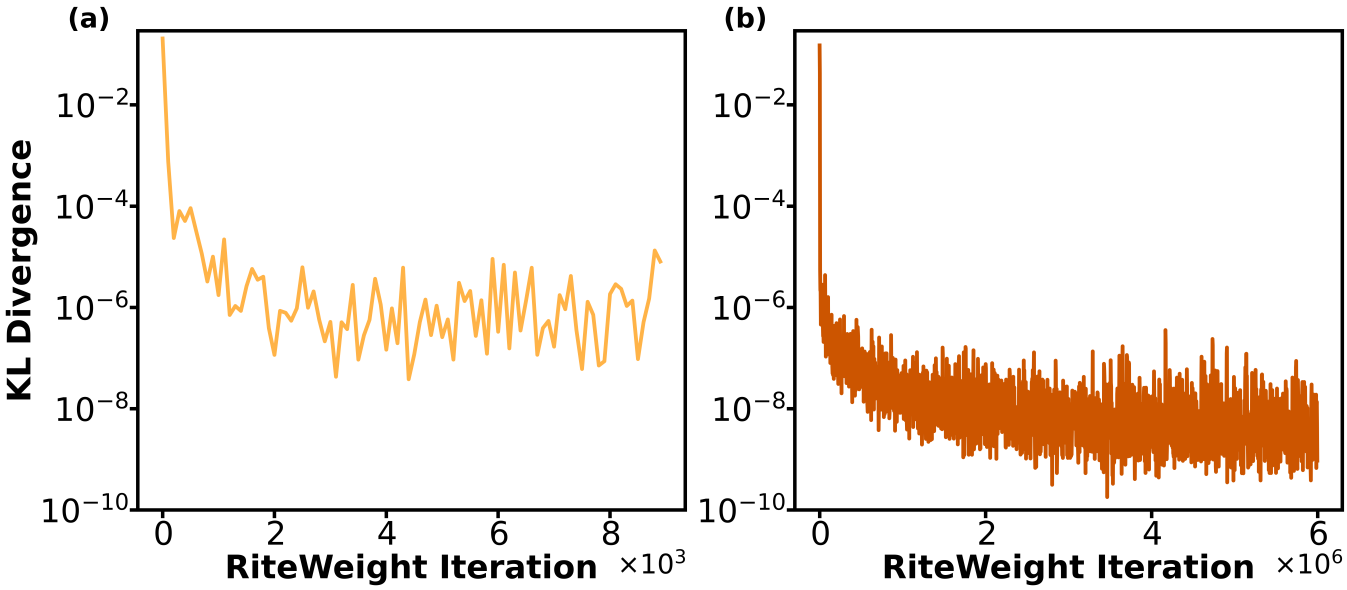

Fig. S1. Convergence of the SynMD Trp-cage equilibrium distribution estimated using RiteWeight. The vertical axes show the symmetric Kullback-Leibler (KL) divergence between the estimated stationary distribution in the current iteration vs. the estimated distribution 100 iterations prior. The horizontal axes show the number of iterations. (a) The KL divergence using 1000 clusters. (b) The KL divergence using 10 clusters.

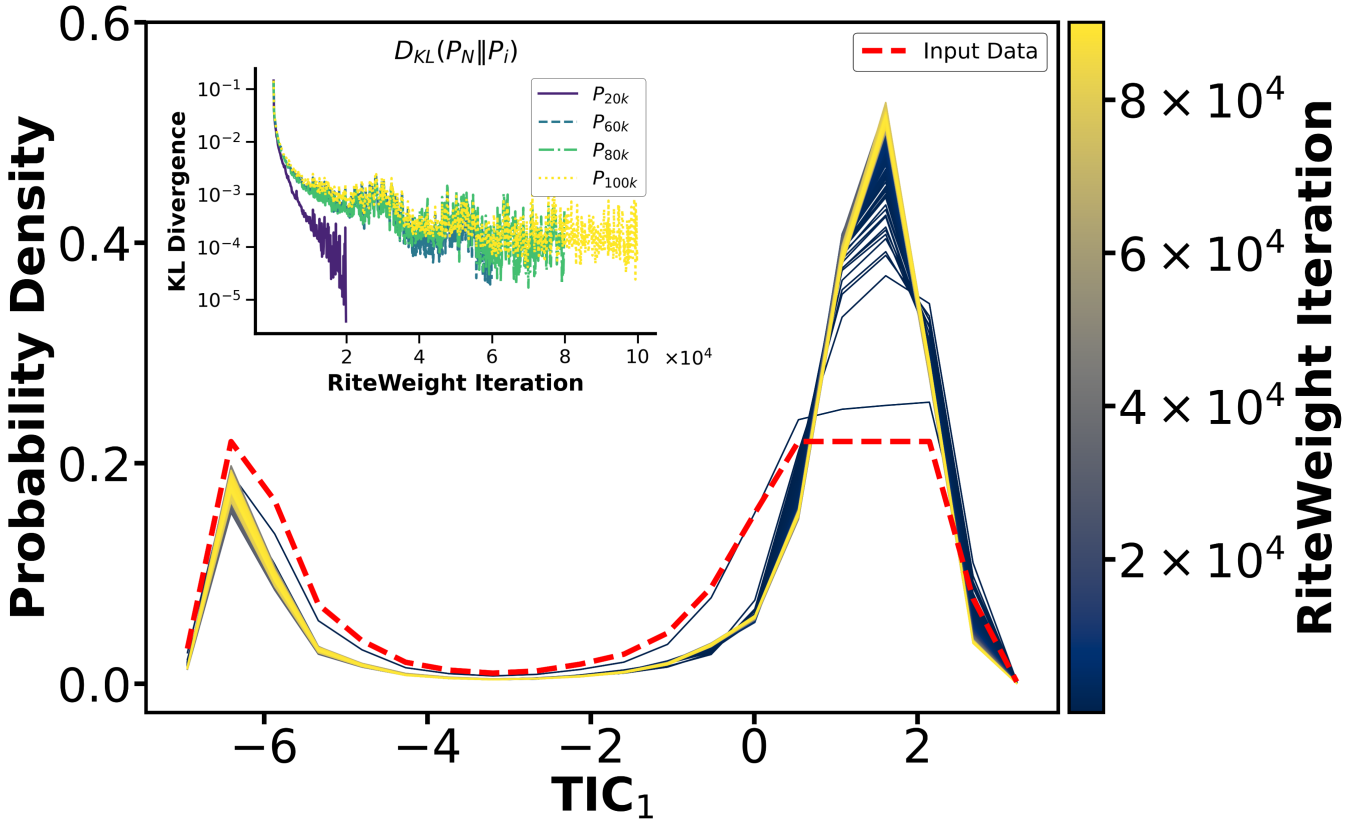

Fig. S2. Convergence of the atomistic Trp-cage equilibrium distribution using RiteWeight. The main plot shows the convergence of the initial (dashed red) distribution to the final (yellow) distribution using increments of 100 iterations. The inset shows the KL divergence of the RiteWeight distribution with reference to iteration 20,000, 60,000, 80,000 or 100,000, as shown in the legend. The similarity of the curves for 60,000 iterations and beyond suggests RiteWeight is converged.

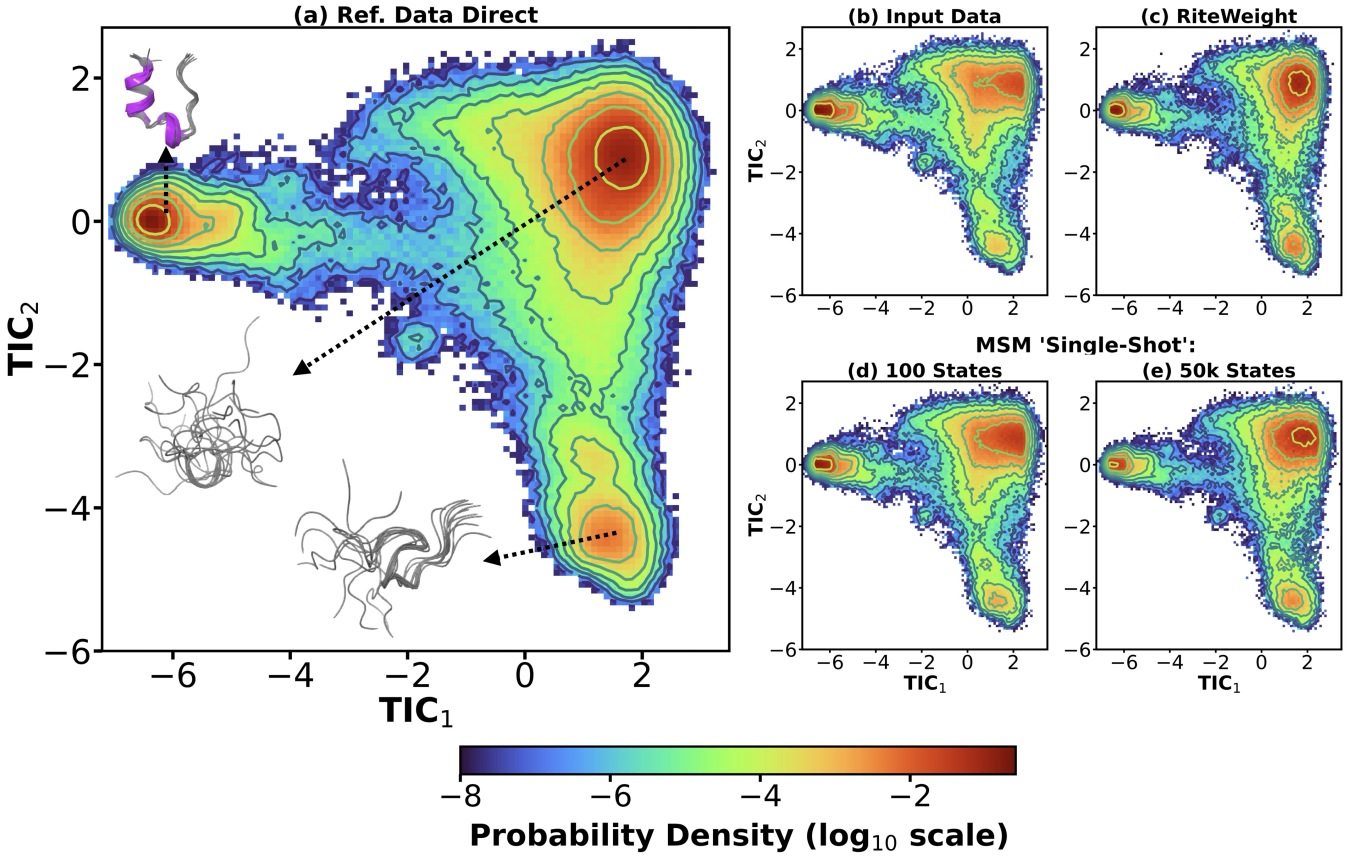

Fig. S3. Equilibrium probability densities projected onto the two slowest tICA (TIC) components. Contour lines indicate 0.82-unit increments in the base-10 logarithm of the probability density. (a) The reference density is calculated from a 208 μs MD trajectory. Ten randomly selected configurations are overlaid for each of the most probable states. (b) The input data is mis-distributed relative to the reference. (c) RiteWeight successfully recovers the true equilibrium distribution. (d–e) MSM 'single-shot' estimates using coarse (100 states) and fine (50,000 states) resolutions fail to recover features of the reference distribution near the energy basins. The same data employed for panels (b–e) is shown projected onto TIC_1 in Fig. 3 of the main text.

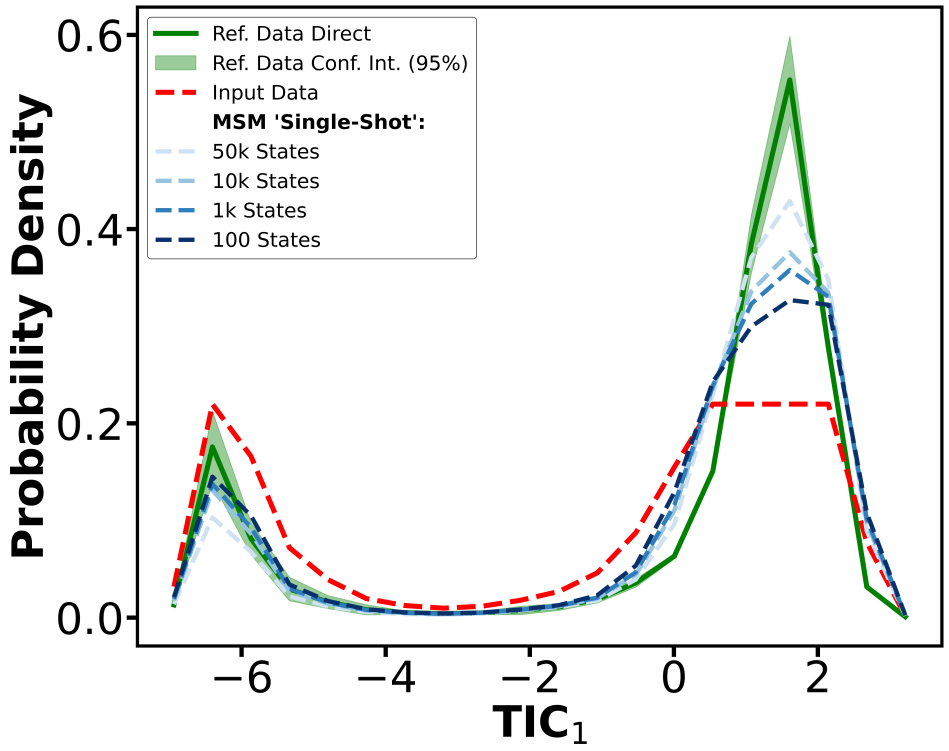

Fig. S4. Markov state model "single shot" estimates for the equilibrium stationary distribution based on different numbers of MSM states (blue dashed lines) using the lag time τ = 100 ns.

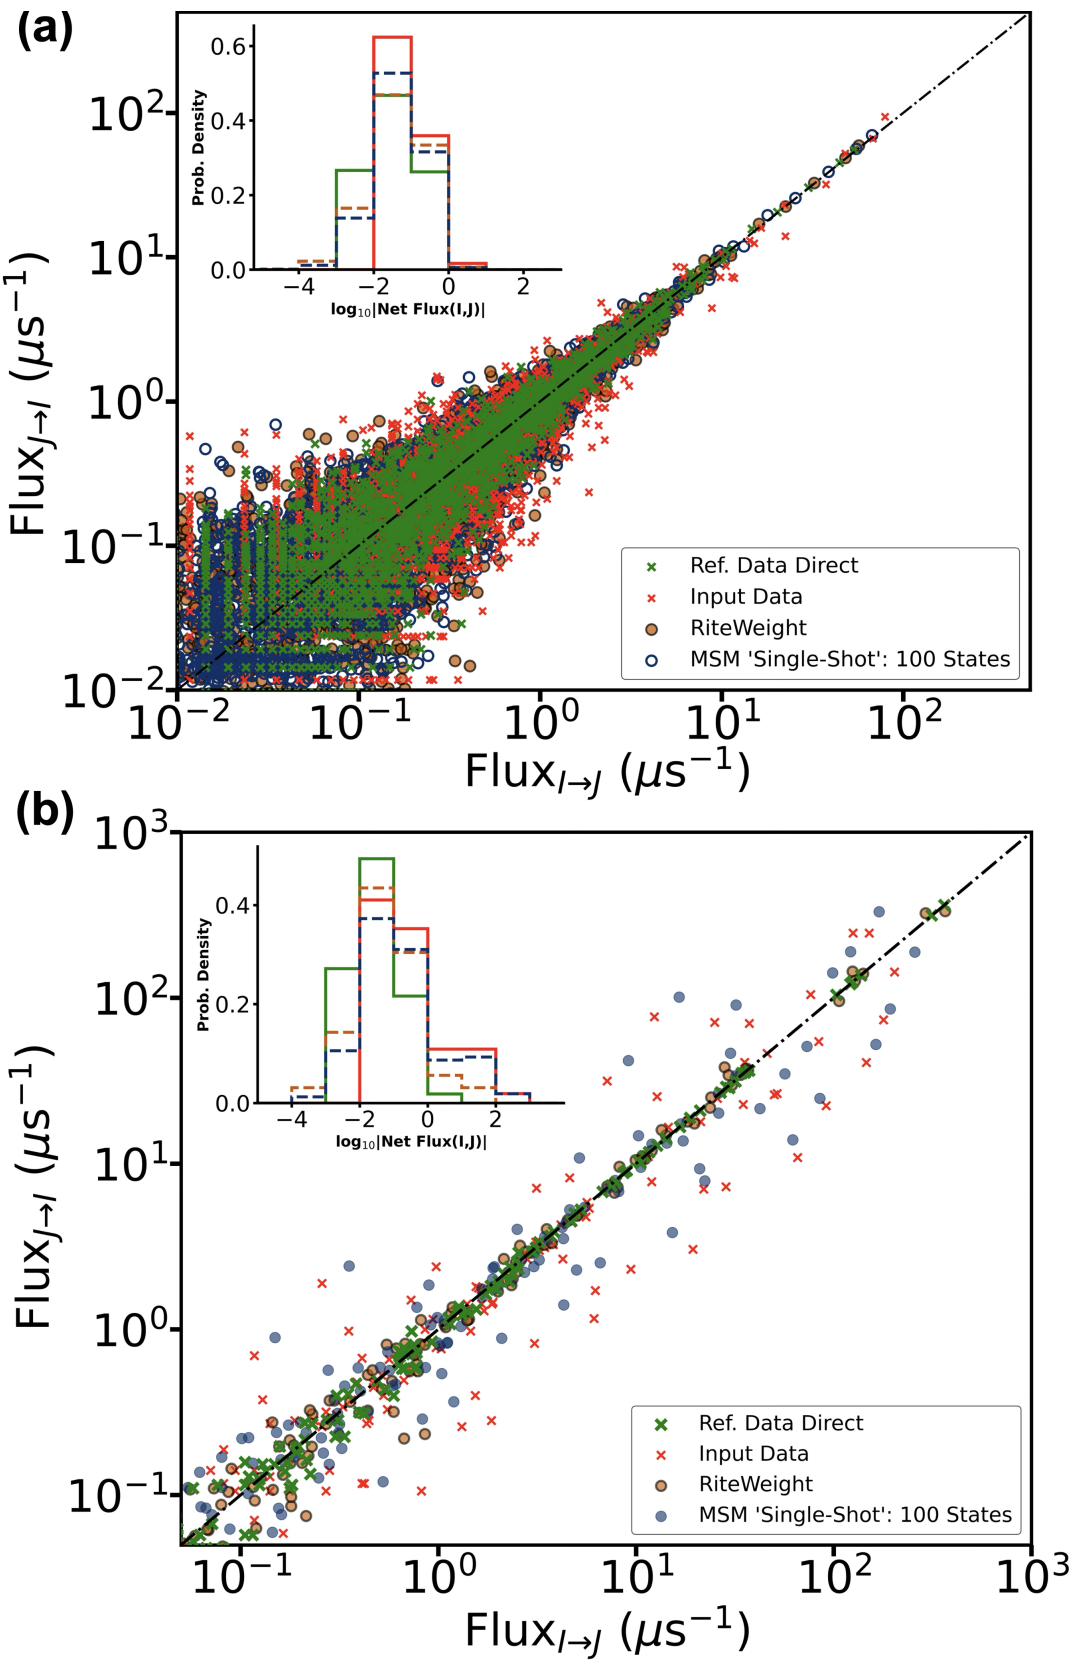
(a)
Prob. Density
log10|Net Flux(I,J)|
Ref. Data Direct
Input Data
RiteWeight
MSM 'Single-Shot': 100 States
Flux$_{J\to I}$ (μs^{-1})
Flux$_{I\to J}$ (μs^{-1})
(b)
Prob. Density
log10|Net Flux(I,J)|
Ref. Data Direct
Input Data
RiteWeight
MSM 'Single-Shot': 100 States
Flux$_{J\to I}$ (μs^{-1})
Flux$_{I\to J}$ (μs^{-1})

Fig. S5. Detailed-balance assessment at fine and coarse scales. Here, $\mathrm{Flux}_{I\to J}$ is the total weight of trajectory segments transitioning from state I to J over a lag time of τ = 10 ns. The violation of detailed balance is measured by the distance of data points from the diagonal. The insets show histograms of the net flux magnitudes $|\text{Net Flux}(I,J)| = |\text{Flux}_{I\to J} - \text{Flux}_{J\to I}|$ for the reference data (green solid line), mis-distributed input data (red solid line), RiteWeight estimates (dark orange dashed line), and MSM "single-shot" estimates based on 100 states (blue dashed line). (a) Fine-scale regions of configuration space are defined by the 100 MSM states used to discretize the equilibrium distribution in Fig. 3 of the main text. At the fine scale, MSMs and RiteWeight both shift to smaller violations compared to the input data, even though MSMs explicitly enforce detailed balance during stationary-distribution estimation (2) whereas RiteWeight does not. (b) Coarse-scale regions are defined by 20 uniform bins along TIC_1. The fluxes estimated by RiteWeight are similar to those from the reference data and lie close to the diagonal; hence RiteWeight largely restores detailed balance in this coarser analysis, whereas MSM estimates fail to recover detailed balance.

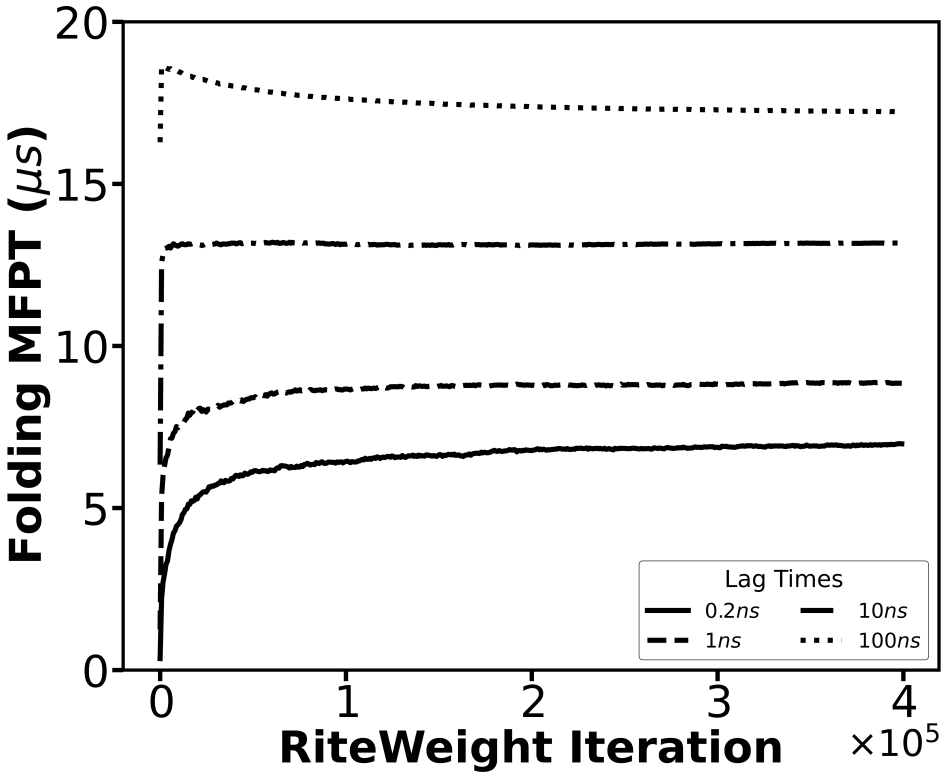

Fig. S6. Convergence of the atomistic Trp-cage folding MFPT using RiteWeight. Four different lag times are examined, and the MFPT is steady after 200,000 iterations for any lag time. The variation of MFPT with lag time is not an artifact but a physical consequence of the fact that first folding events sometimes are missed with longer lag times.

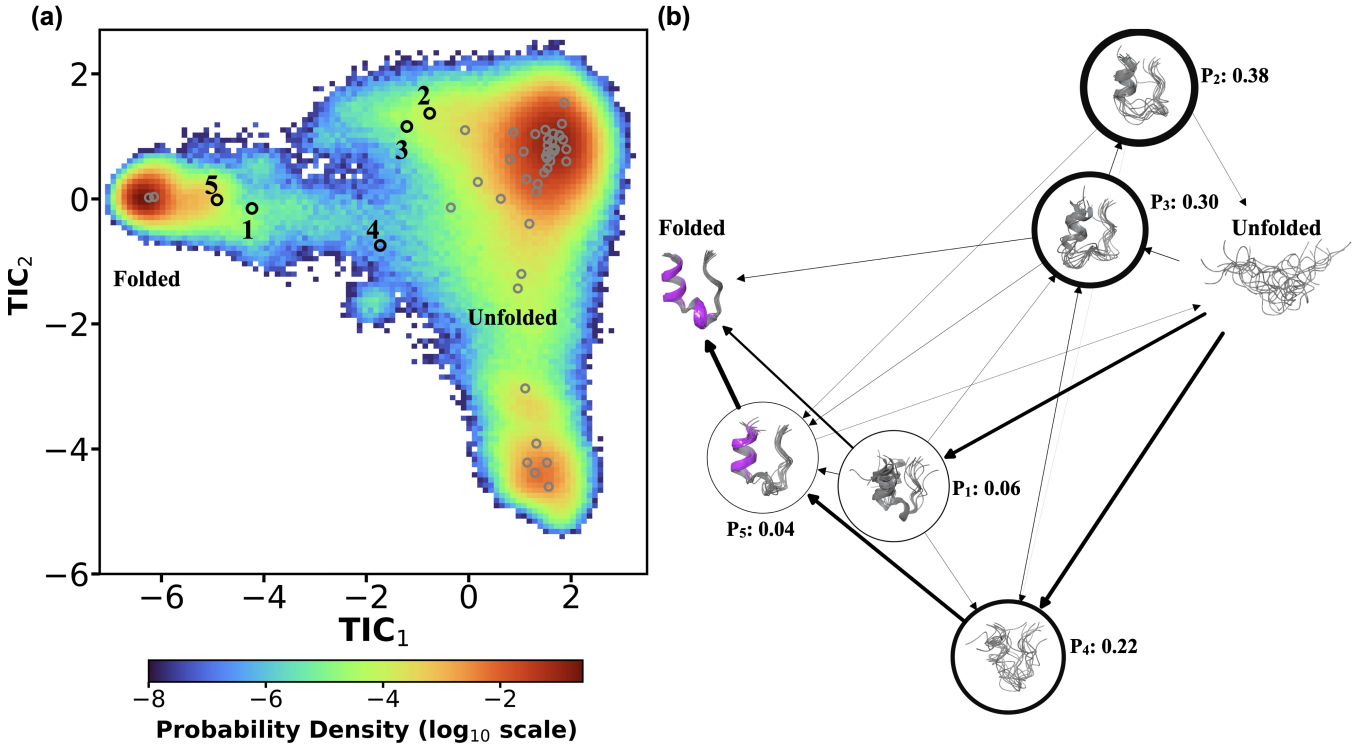

Fig. S7. Folding intermediates and sequences of transitions constructed from reference data at a 0.2 ns lag time. (a) The equilibrium probability density is projected onto the two slowest tICA components. Black circles indicate MSM states selected as intermediate macrostates, while gray circles denote MSM states assigned to the unfolded and folded macrostates. (b) In the nonequilibrium flux network, arrow thicknesses indicate net fluxes. Node thicknesses indicate non-equilibrium steady state probabilities for the intermediate macrostates, which are normalized to sum to one for clarity. These non-equilibrium probabilities are distinct from the equilibrium probability density in panel (a). Ten randomly selected configurations are overlaid for each macrostate, with purple regions highlighting folded helical segments. Cluster 2 is partially unfolded; clusters 1, 3, and 5 feature a partially folded N-terminus; and cluster 4 identifies the formation of the hydrophobic core.

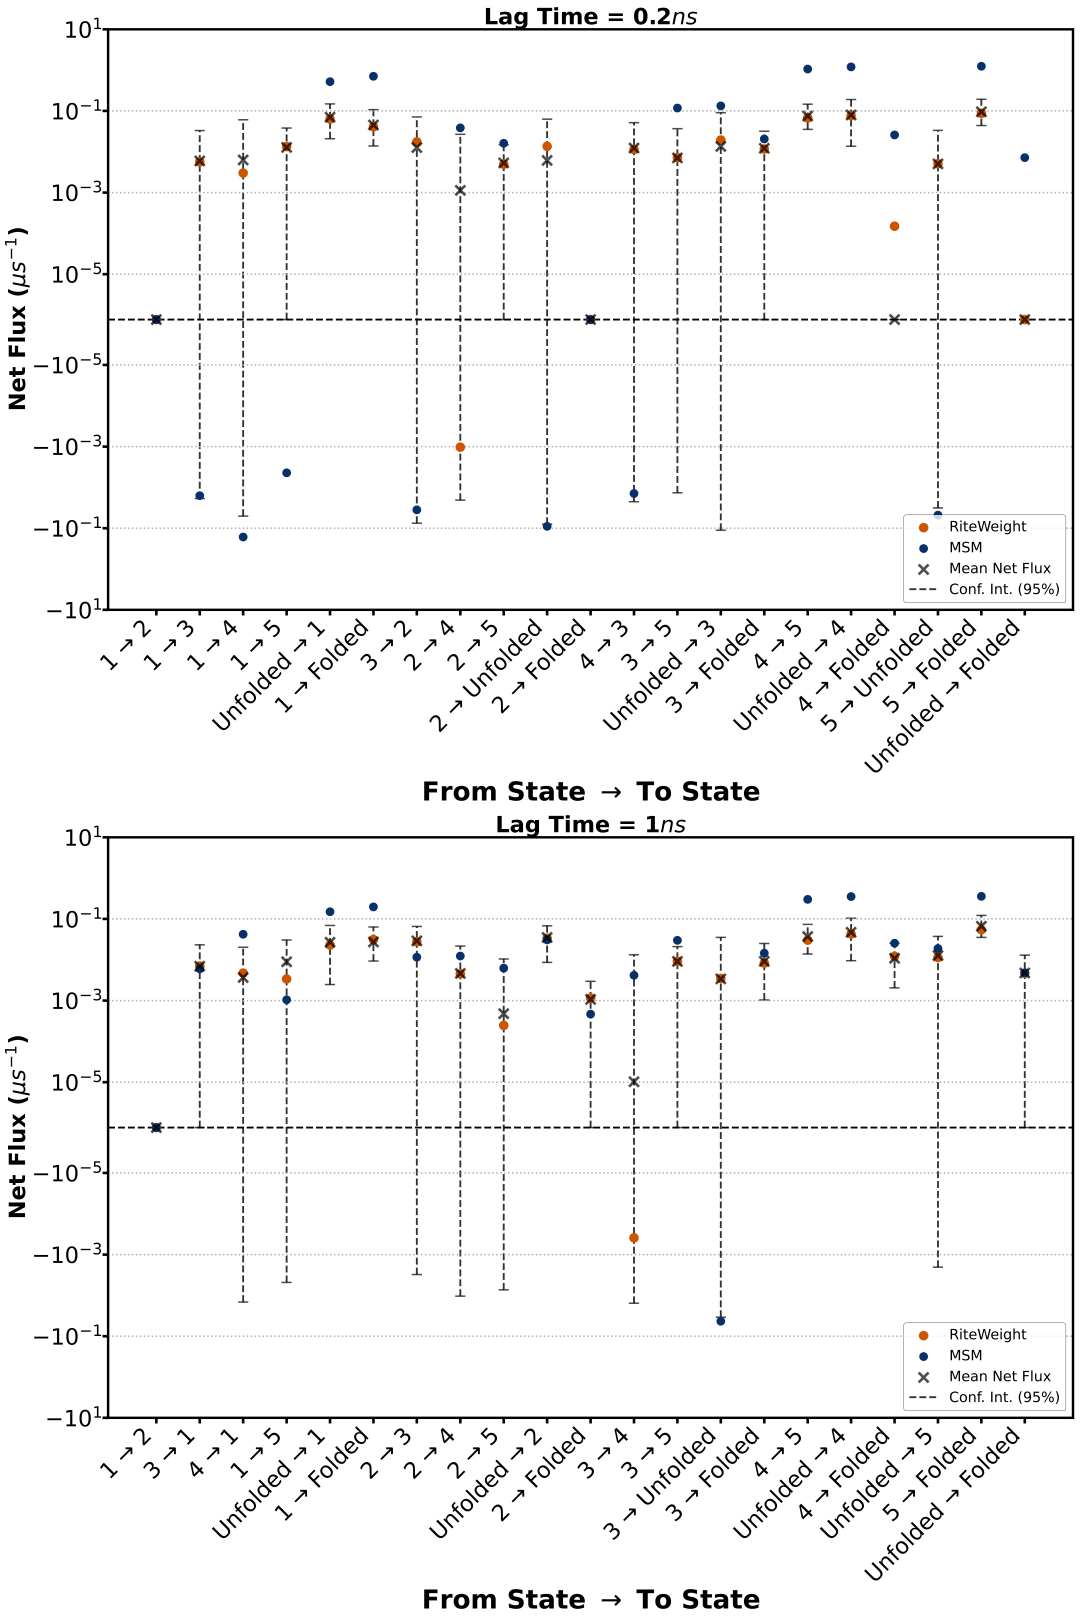

Fig. S8. Net flux analysis showing confidence intervals for MD-based estimates. For each pair of macrostates, the net flux is shown for RiteWeight (dark orange), Markov state models (MSM – blue), and molecular dynamics (MD – X and confidence interval). The MD estimates are calculated from the ensemble of reactive trajectories that proceed from the unfolded to the folded macrostate. The MD 95% confidence interval is derived by bootstrapping the 24 round-trip paths starting from the folded state, entering the unfolded state, and then ending back at the folded state. Two lag times are used for analysis: 0.2 ns (top) and 1.0 ns (bottom).

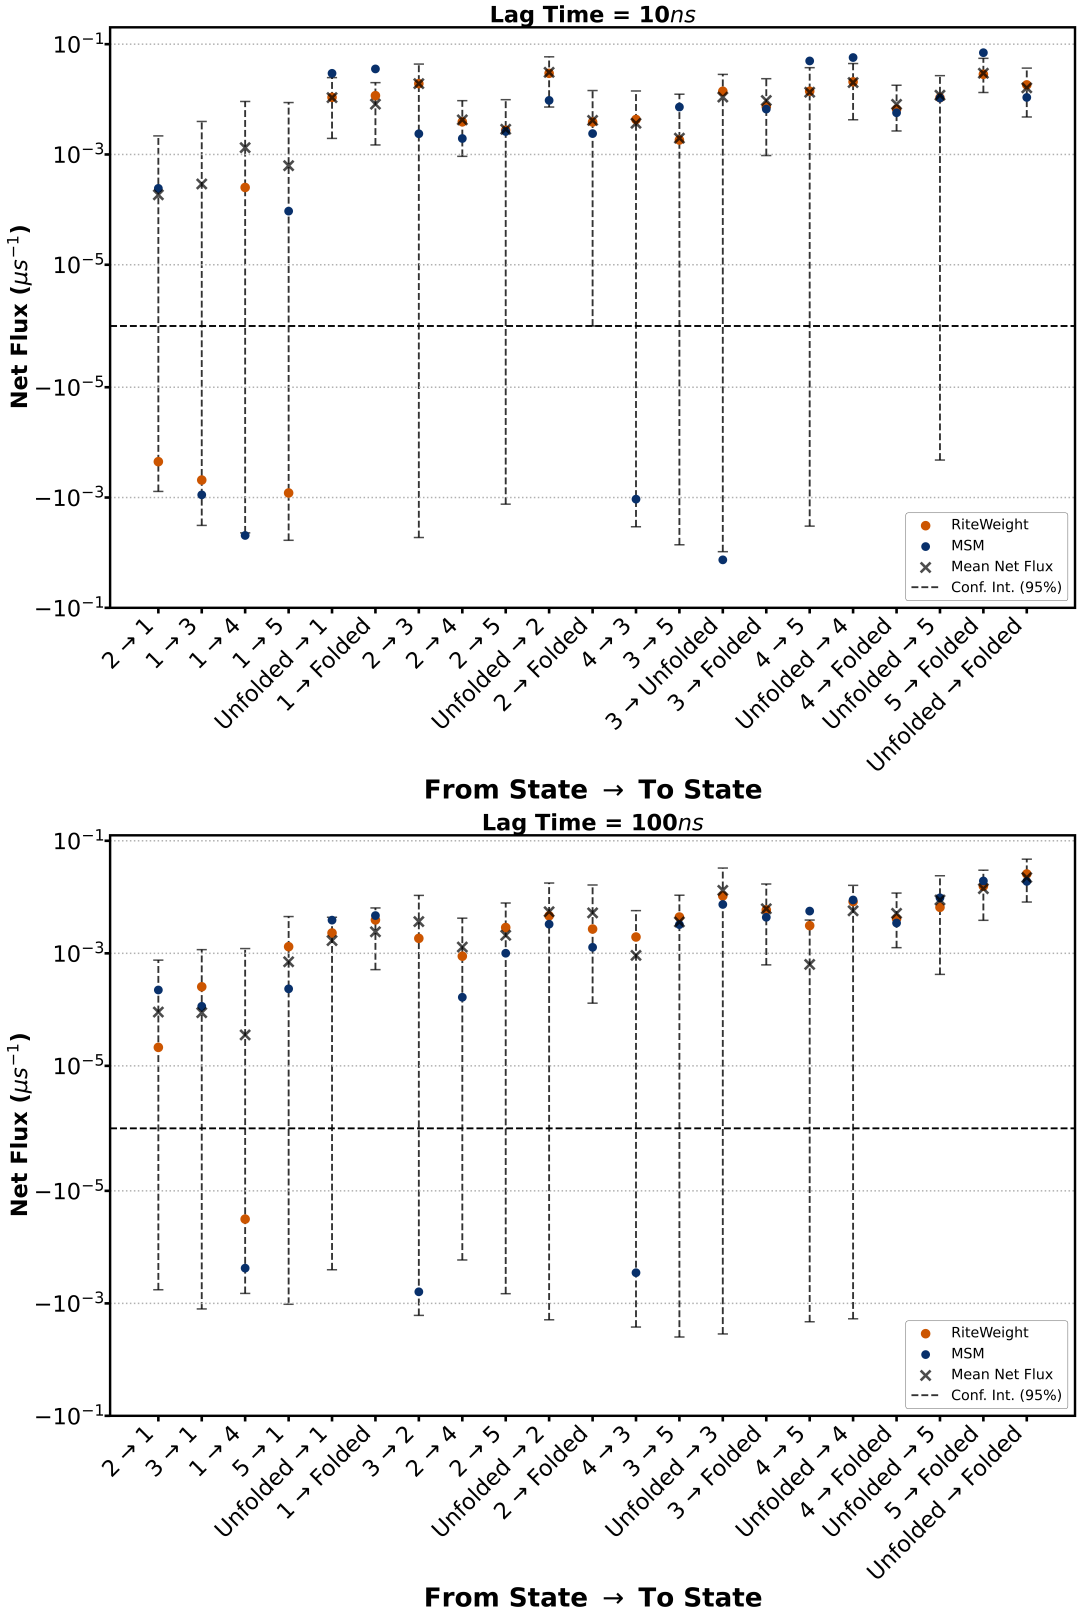

Fig. S9. Net flux analysis showing confidence intervals for MD-based estimates. Two lag times are used for analysis: 10 ns (top) and 100 ns (bottom).

References

1. E O'Reilly, NM Tran, Stochastic geometry to generalize the mondrian process. *SIAM J. on Math. Data Sci.* **4**, 531–552 (2022).
2. MK Scherer, et al., Pyemma 2: A software package for estimation, validation, and analysis of markov models. *J. chemical theory computation* **11**, 5525–5542 (2015).