Agent-guided variance reduction for importance-sampled protein kinetics: recovering the alanine dipeptide $\beta/C_5 \leftrightarrow C_7^{\mathrm{eq}}$ isomerization rate

Michael M. Kim*[1]

[1]Department of Mechanical Engineering, Stanford University, Stanford, CA 94305

July 13, 2026

Abstract

Estimating protein conformational kinetics from importance-sampled trajectories is limited by weight degeneracy: when biased transition paths are reweighted to the true dynamics, the variance of the path weights grows exponentially and a few paths dominate the estimate. We ask which weight-control scheme tames this variance for a real peptide, and we answer it as an agentic method search: Claude Code proposes an estimator variant, implements it, launches a rate calculation, reads back the weight statistics and effective sample size against a brute-force reference, and iterates. On the $\beta/C_5 \leftrightarrow C_7^{\mathrm{eq}}$ isomerization of alanine dipeptide, we compare sequential importance sampling (SMC), sequential importance resampling (SIR), and a branching resampled random walk (RRS). RRS is the only scheme that converges in both directions, recovering the brute-force rate to within 3–6% where the naive committor estimate is up to 6× wrong. Code and data are available from the author on request.

1 Introduction

The rates of protein conformational changes control biological function, yet they are among the hardest quantities to compute. The transitions that matter—a loop reorganizing, a dihedral flipping, a domain rotating—are rare on the timescale of the fastest molecular motions, so a direct molecular dynamics (MD) trajectory spends essentially all of its steps inside a metastable region and almost none crossing between regions [1–3]. Waiting for enough spontaneous crossings to estimate a rate can require billions of steps, and the cost grows exponentially with the barrier height [4].

Importance sampling attacks this problem directly. Instead of waiting, one biases the dynamics toward the transition, samples many reactive paths cheaply, and then reweights the biased trajectories back to the true, unbiased dynamics so that the recovered rate remains statistically exact [5–8]. A particularly clean route parameterizes the bias with the committor [9, 10]—the probability that a trajectory reaches the product region before the reactant—and reconstructs the rate as an escape flux times a committor-weighted reactive probability [11, 12]. Modern work approximates the committor with a neural network [13–16], which lets the approach scale to many degrees of freedom.

*Email: mckim2013@gmail.com

The catch is variance. When a biased path is reweighted to the true dynamics, it carries a statistical weight equal to a product of per-step normalization factors, so its variance grows exponentially in path length—the classic weight-degeneracy pathology of sequential importance sampling [17, 18]. Proteins are close to the worst case: the fast, stiff vibrational modes inject noise into the weight at every step without advancing the slow dihedral reorganization we actually care about. A handful of paths then acquire enormous weights, the effective sample size collapses [19], and the rate estimate becomes both noisy and unreliable at fixed compute.

The standard remedy is resampling: periodically prune low-weight paths and multiply high-weight ones, so the population stays diverse and the weights stay bounded [18, 20, 21]. But "resample" names a family, not a method—one must choose *what* triggers a resampling, *how* offspring counts are drawn, and *how* weights are reset [22, 23]—and the best choice is system dependent. Exploring that design space by hand is slow: each candidate is a small code change followed by a long run and a careful reading of variance diagnostics.

This is where we bring Claude into the loop. We treat the choice of estimator as a search problem and drive it agentically with Claude Code [24, 25]: the agent proposes a resampling variant, implements it in the sampler, launches the rate calculation, parses the resulting weight distribution, effective sample size, and rate against a brute-force ground truth, and then proposes the next variant. On the $\beta/C_5 \leftrightarrow C_7^{\mathrm{eq}}$ isomerization of alanine dipeptide [26]—the standard peptide benchmark, small enough that an exact brute-force rate is attainable—this loop compared three schemes and identified a branching resampled random walk (RRS) as the one that works for this protein. Our contributions are: (i) a committor-based, reweighted rate estimator for alanine dipeptide with an exact brute-force reference in both directions; (ii) a controlled comparison of SMC, SIR, and RRS weight control on that reference; and (iii) a demonstration that an agentic method search can select a variance-reduction scheme that a human would otherwise tune by trial and error.

2 Theory

2.1 Transition paths and the reactive probability

We evolve the system with overdamped Langevin dynamics and decompose a long trajectory into shorter paths. A path is a sequence of microscopic states

$$\Gamma = \{\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_N\}, \tag{1}$$

initiated at $\mathbf{x}_1$ whenever the trajectory leaves region A and terminated whenever it enters region A or B, i.e. $\mathbf{x}_N \in \mathrm{A} \cup \mathrm{B}$. Paths terminating in A are "failure" paths Γ_{AA}; those terminating in B are "success" paths Γ_{AB}. Successive states are separated by a fixed lag $\tau = k\Delta t$, and we write $\mathcal{K}_\tau(\mathbf{y} \,|\, \mathbf{x})$ for the stopped transition kernel—the probability density of visiting $\mathbf{y}$ as the next state given the current state $\mathbf{x}$, with the time separation $\tau \wedge \tau_{\mathrm{A}\cup\mathrm{B}}$,

$$\mathcal{K}_\tau(\mathbf{y} \,|\, \mathbf{x}) = \mathbb{E}\Big[\, \delta\big(\mathbf{y} - \mathbf{x}(\tau \wedge \tau_{\mathrm{A}\cup\mathrm{B}})\big) \,\Big|\, \mathbf{x}(0) = \mathbf{x} \Big]\,. \tag{2}$$

The probability density of a success path factorizes as

$$p(\Gamma_{\mathrm{AB}}) = \rho_{\mathrm{A}}(\mathbf{x}_1) \prod_{i=1}^{N-1} \mathcal{K}_\tau(\mathbf{x}_{i+1} \,|\, \mathbf{x}_i), \tag{3}$$

where $\rho_{\mathrm{A}}(\mathbf{x}_1)$ is the density of the initial state sampled at the moment of exiting region A. The success probability is the integral over all success paths, $P_{\mathrm{AB}} = \int_{\Gamma \in \{\Gamma_{\mathrm{AB}}\}} p(\Gamma)\, \mathrm{d}\Gamma$, and the A $\to$ B rate follows from the Hill relation [11, 12],

$$r_{\mathrm{AB}} = \Phi_{\mathrm{A}}\, P_{\mathrm{AB}}, \qquad \mathrm{MFPT}_{\mathrm{AB}} = 1/r_{\mathrm{AB}}, \tag{4}$$

where Φ_A is the flux of trajectories escaping region A (the inverse mean time between successive first exits). The flux is cheap—escapes from a region are not rare—so the difficulty is concentrated in P_{AB}, the probability that an escaped trajectory commits to the far region rather than returning.

2.2 Importance sampling and the path weight

For rare events the success probability is tiny compared with that of a failure path. We therefore modify the transition kernel with an importance function $I(\mathbf{x})$ [5–8],

$$\mathcal{K}'_\tau(\mathbf{y}\,|\,\mathbf{x}) = Z_\tau^{-1}(\mathbf{x})\,\frac{I(\mathbf{y})}{I(\mathbf{x})}\,\mathcal{K}_\tau(\mathbf{y}\,|\,\mathbf{x}), \qquad Z_\tau(\mathbf{x}) = \int_\Omega \mathrm{d}\mathbf{y}\,\frac{I(\mathbf{y})}{I(\mathbf{x})}\,\mathcal{K}_\tau(\mathbf{y}\,|\,\mathbf{x}), \tag{5}$$

where the normalization factor $Z_\tau(\mathbf{x})$, computed by Monte Carlo integration, keeps $\mathcal{K}'_\tau$ a probability density. Propagating a walker under $\mathcal{K}'_\tau$ and reweighting back to the true dynamics attaches to each path the correction factor

$$p'(\Gamma_{AB}) = \frac{I(\mathbf{x}_N)}{I(\mathbf{x}_1)}\left[\prod_{i=1}^{N-1} Z_\tau^{-1}(\mathbf{x}_i)\right] p(\Gamma_{AB}), \qquad W(\Gamma) \coloneqq \prod_{i=1}^{N-1} Z_\tau(\mathbf{x}_i), \tag{6}$$

with path weight $W(\Gamma)$. Setting $I(\mathbf{x}) = 0$ for $\mathbf{x} \in$ A and $I(\mathbf{x}) = 1$ for $\mathbf{x} \in$ B, the optimality (normalization) condition $\int_\Omega I(\mathbf{y})\mathcal{K}_\tau(\mathbf{y}\,|\,\mathbf{x})\,\mathrm{d}\mathbf{y} = I(\mathbf{x})$ is exactly the backward equation whose solution is the committor [9]. When I equals this optimal importance function, every $Z_\tau = 1$ and the bias enhances success paths by the known factor $I(\mathbf{x}_N)/I(\mathbf{x}_1)$; when I is imperfect, the weights $W(\Gamma)$ restore the unbiased statistics exactly. The reweighted estimator of the reactive probability is therefore

$$\widehat{P}_{AB} = \frac{1}{N_p}\sum_{i=1}^{N_p} I(\mathbf{x}_1^{(i)})\,W(\Gamma_i)\,\mathbb{1}[\Gamma_i \in \{\Gamma_{AB}\}]\,, \tag{7}$$

which is unbiased for *any* importance function, including an imperfect learned committor—errors in the guide degrade efficiency, not correctness [7].

2.3 The weight-variance problem

Unbiasedness is not enough. Because $W(\Gamma)$ in Eq. (6) is a product of $N-1$ fluctuating factors, its variance grows exponentially in path length,

$$\mathrm{Var}[W] \;\sim\; e^{\,cN}, \qquad c > 0. \tag{8}$$

The efficiency of Eq. (7) is governed by the effective sample size [19],

$$\mathrm{ESS} = \frac{\left(\sum_i W_i\right)^2}{\sum_i W_i^2}, \tag{9}$$

which collapses from N_p toward 1 as a single weight comes to dominate. For alanine dipeptide the fast bond and angle modes make the per-step factors noisy without advancing the slow (ϕ, ψ) reorganization, so weight degeneracy sets in quickly—the reason a raw estimator is unusable.

2.4 Resampling schemes as variance control

Resampling interrupts the exponential growth by periodically replacing the weighted population with a nearly equally weighted one of the same expectation. We compare three members of this family, all sharing Eq. (7) and differing only in how weights are controlled. *SMC* (sequential importance sampling, no resampling) lets weights accumulate unchecked; it is the baseline and the most exposed to Eq. (8) [17]. *SIR* (sequential importance resampling) resamples the whole population by low-variance systematic resampling whenever ESS $< \theta N_p$ (here $\theta = 0.5$), resetting every weight to the mean [18, 20]; global resampling is aggressive and can over-prune diversity. *RRS* (resampled random walk / branching) treats each walker individually: whenever its weight leaves a target band $[w_{\min}, w_{\max}]$, it is branched into an integer number of unit-weight offspring,

$$n_{\text{off}} = \lfloor W \rfloor + \mathbb{1}[\, u < W - \lfloor W \rfloor \,], \qquad u \sim \mathcal{U}(0,1), \tag{10}$$

so overweight walkers split and underweight walkers undergo Russian-roulette pruning, with survivors reset to unit weight [21–23]. Branching keeps the weight distribution concentrated near unity while preserving the estimator's expectation, and is the scheme the agentic search ultimately selected.

3 Methods

3.1 System and reference dynamics

As a molecular proof of principle, we consider alanine dipeptide (Ac-Ala-NHMe, 22 atoms) in vacuum, a standard benchmark for conformational sampling and rare-event dynamics [26]. The molecular-mechanics parameters are generated with the AMBER14 force field [27] and the dynamics are propagated with OpenMM tooling [28]. The system is evolved by constrained overdamped Langevin dynamics, integrated with the Euler–Maruyama scheme,

$$\mathbf{x}(t+\Delta t) = \mathbf{x}(t) - \gamma^{-1}\mathbf{M}^{-1}\nabla U(\mathbf{x}(t))\,\Delta t + \sqrt{2\,\gamma^{-1}\beta^{-1}\Delta t}\;\mathbf{M}^{-1/2}\boldsymbol{\xi}, \qquad \boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbb{I}), \tag{11}$$

at $T = 300\,\text{K}$, damping $\gamma = 1.0\,\text{ps}^{-1}$, and timestep $\Delta t = 0.01\,\text{fs}$, with hydrogen bond-length constraints imposed by SHAKE-like projection [29]. The conformation is characterized by the backbone dihedrals $\boldsymbol{\theta} = (\phi, \psi)^\top$, and metastable regions are defined as disks of radius 10° in Ramachandran space: A = β/C_5 centered at $(-150°, 160°)$ and B = C_7^{eq} centered at $(-75°, 60°)$, with the α_R region at $(-70°, -40°)$ acting as an absorbing guard so that indirect crossings are not miscounted (Fig. 3). All runs use random seed 2026.

3.2 Learned importance function

The importance function $I(\mathbf{x})$ is a SchNet-style graph neural network [30] acting on ten chemically distinct heavy atoms selected from the topology, with 64 features, three interaction layers, a 20-Gaussian radial basis (cutoff 0.9 nm), and translational/rotational invariance (Fig. 1). A single head yields a bidirectional committor via $I_+ = \sigma(f)$ and $I_- = \sigma(-f)$. It is trained by minimizing the log-space residual of the optimality condition of Sec. 2.2, $\ln[\int I(\mathbf{y})\mathcal{K}_\tau(\mathbf{y}|\mathbf{x})\mathrm{d}\mathbf{y}] - \ln I(\mathbf{x})$, over swarm-propagated lags, using a bootstrapped target network updated by exponential moving average ($\tau_{\text{ema}} = 5 \times 10^{-3}$), a replay buffer, and gradient clipping [31–33]. The production importance function is an ensemble average of checkpoints from epochs 1000–1100 [34], which reduces the estimator noise of any single network.

3.3 Rate estimator

The escape flux Φ in Eq. (4) is measured by counting successive first exits from a region during an equilibration run (10^4 escapes). The reactive probability P_{AB} is estimated by launching $N_p = 1000$ importance-guided walkers, each carrying a swarm of $n_{\mathrm{swarm}} = 100$ micro-steps used to estimate the normalization factor Z_τ, accumulating weights (6), applying the chosen resampling scheme (Sec. 2.4), and evaluating Eq. (7). The calculation is repeated for 1000 independent trials to obtain the running mean, standard deviation, and standard error. RRS uses the weight band $[1, \infty)$ (Russian-roulette pruning of sub-unit weights); SIR uses the threshold $\theta = 0.5$.

3.4 Brute-force ground truth

For an exact reference we launch 10^4 unbiased walkers from each region and record first-passage times to the other, giving $r_{\mathrm{AB}}^{\mathrm{BF}} = 1.08 \times 10^{-3}$ fs^{-1} (MFPT = 926 fs) and $r_{\mathrm{BA}}^{\mathrm{BF}} = 4.48 \times 10^{-4}$ fs^{-1} (MFPT = 2234 fs). These are the yardstick for every estimator.

3.5 Agentic method search with Claude Code

The comparison in this paper was conducted as a closed loop driven by Claude Code (Fig. 2). Each iteration: (1) Claude proposes a weight-control variant; (2) implements it as a branch of the sampler's `process_walkers` routine (the SMC/SIR/RRS logic lives in a single function, so a variant is a localized diff); (3) launches the rate calculation and the matched brute-force reference; (4) reads back the rate-versus-trial log, the path-weight histogram, and the effective sample size, and compares to ground truth; and (5) proposes the next variant informed by *why* the previous one failed—degeneracy, over-pruning, or non-convergence. Claude also generated the diagnostics (weight distributions, convergence traces, MFPT bar charts) that appear as our figures. The full propose–implement–run–analyze loop was logged end to end, so the search is auditable and reproducible.

4 Results

4.1 A benchmark with an exact answer

Alanine dipeptide places the two slow states in a two-dimensional dihedral plane whose free-energy surface we can draw explicitly (Fig. 3). The β/C_5 and C_7^{eq} regions sit in separate free-energy wells; because the peptide is small, unbiased MD still yields a converged first-passage rate in both directions, so every biased estimate can be checked against an exact number rather than against itself.

4.2 Weight degeneracy is severe

The path-weight distributions (Fig. 4) make the central difficulty concrete: although the vast majority of weights sit near the mean $\langle W(\Gamma)\rangle$, the distribution is extraordinarily heavy tailed, with rare paths reaching weights orders of magnitude larger. A single such path can outweigh the other 10^5 combined, precisely the effective-sample-size collapse anticipated by Eqs. (8)–(9). Any scheme that does not actively control this tail produces a rate dominated by a few paths.

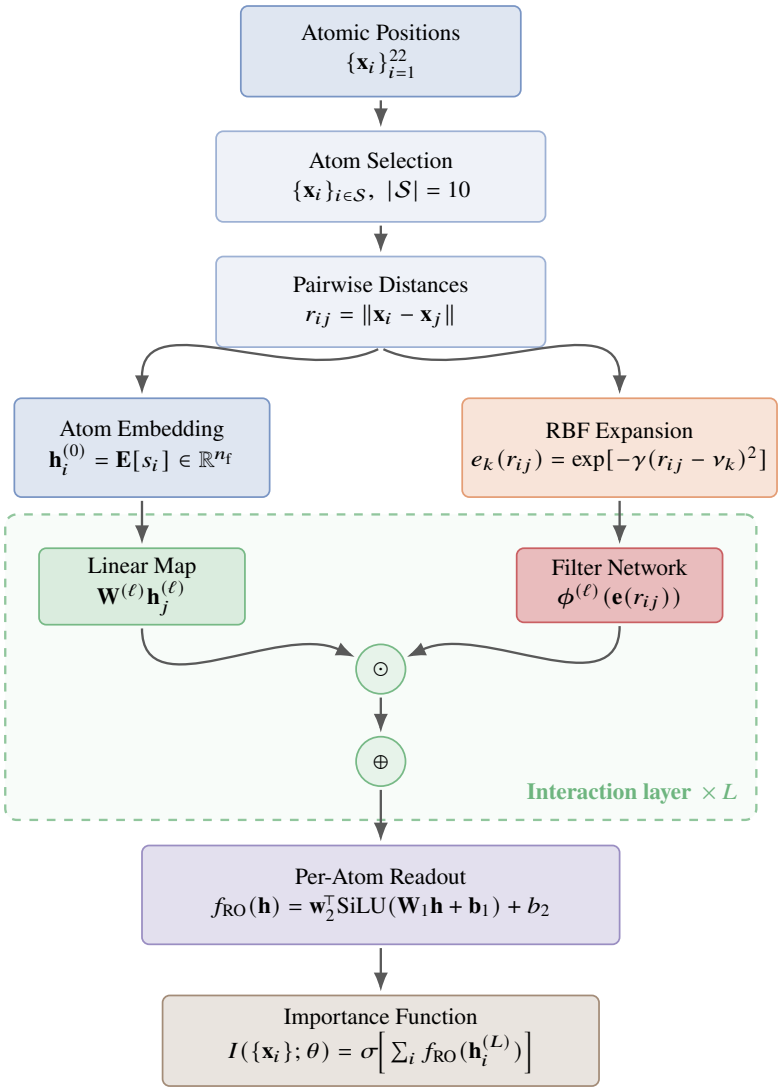

Figure 1: The learned importance function. SchNet-style graph neural network parameterizing the committor $I(\{\mathbf{x}_i\};\theta)$ for alanine dipeptide. Ten chemically distinct atoms are selected from the 22-atom topology; their pairwise distances feed a per-atom embedding and a radial-basis expansion. Each L interaction layer updates atom features through a continuous-filter convolution, $\mathbf{h}_i^{(\ell+1)} = \mathbf{h}_i^{(\ell)} + \sum_{j\neq i} \mathbf{W}^{(\ell)}\mathbf{h}_j^{(\ell)} \odot \phi^{(\ell)}(\mathbf{e}(r_{ij}))$. A per-atom readout, summed and passed through a sigmoid, yields the importance function.

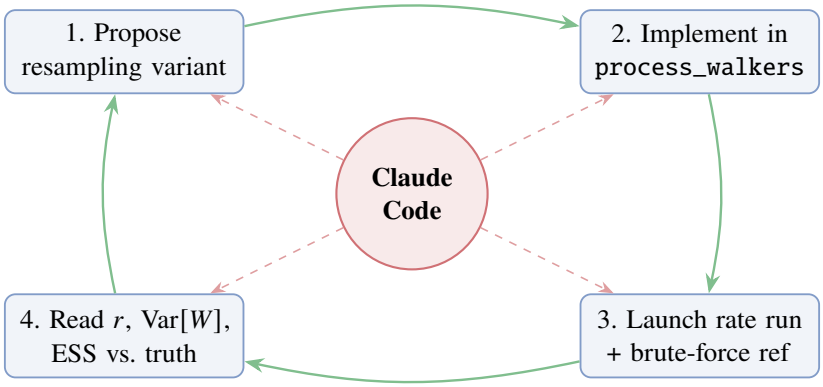

Figure 2: The agentic method search. Claude Code drives a closed loop over estimator designs: propose a weight-control scheme, implement it as a localized change to the sampler, run the biased rate calculation against a matched brute-force reference, and read back weight-variance and effective-sample-size diagnostics to decide the next variant. The loop selected RRS for this protein.

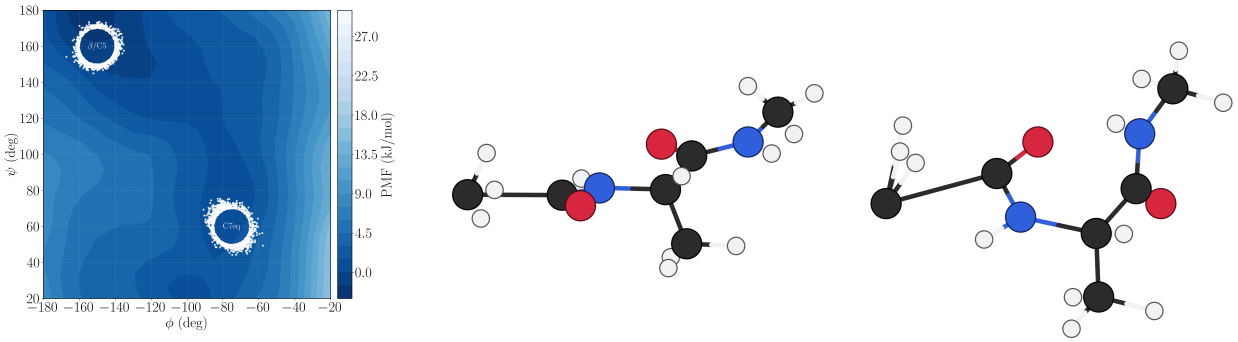

(a) Metastable regions **(b)** A = β/C_5 **(c)** B = C_7^{eq}

Figure 3: The $\beta/C_5 \leftrightarrow C_7^{\mathrm{eq}}$ benchmark. (a) Potential of mean force in the backbone dihedrals (ϕ, ψ); the two metastable regions β/C_5 and C_7^{eq} (dashed rings) are separated by a free-energy barrier. (b,c) The two metastable conformations of alanine dipeptide (carbon black, nitrogen blue, oxygen red, hydrogen white).

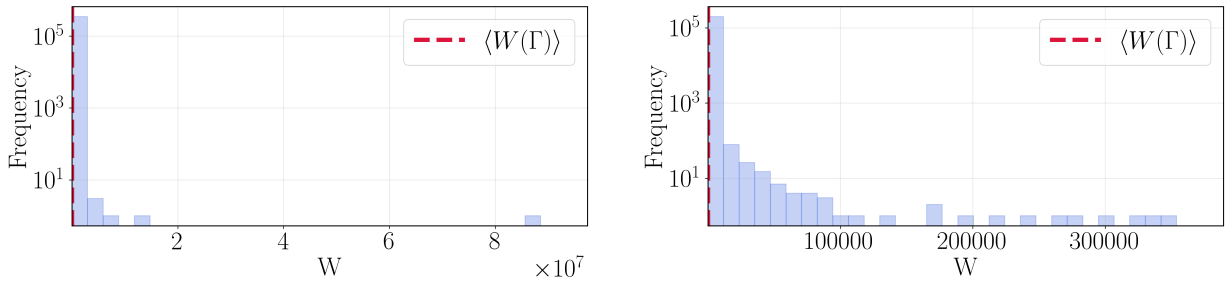

(a) $W(\Gamma)$, $\beta/C_5 \to C_7^{\mathrm{eq}}$ **(b)** $W(\Gamma)$, $C_7^{\mathrm{eq}} \to \beta/C_5$

Figure 4: The weight-variance problem. Histograms (log frequency) of path weights $W(\Gamma)$ under RRS. Even after branching control, the distribution is heavy tailed—rare paths carry weights far above the mean $\langle W(\Gamma) \rangle$ (dashed)—illustrating why naïve reweighting is unusable and why weight control is essential.

Table 1: Weight-control schemes versus brute force. Final estimates over up to 1000 trials ($N_p = 1000$ walkers, ensemble importance function, seed 2026). "Trials" is the number completed before convergence or termination; σ/μ is the per-trial relative standard deviation of the rate (a proxy for weight variance). Only RRS completes both directions and matches the reference. Rates in fs^{-1}.

Scheme	Direction	Trials	Rate	σ/μ	MFPT (fs)
SMC (no resampling)	$\beta/C_5 \to C_7^{\mathrm{eq}}$	1000	1.02×10^{-3}	0.70	978
	$C_7^{\mathrm{eq}} \to \beta/C_5$	31†	3.91×10^{-4}	0.38	2555
SIR (global resample)	$\beta/C_5 \to C_7^{\mathrm{eq}}$	118†	1.22×10^{-3}	2.18	822
	$C_7^{\mathrm{eq}} \to \beta/C_5$	533†	5.25×10^{-4}	3.15	1906
RRS (branching)	$\beta/C_5 \to C_7^{\mathrm{eq}}$	1000	$\mathbf{1.05 \times 10^{-3}}$	0.92	**950**
	$C_7^{\mathrm{eq}} \to \beta/C_5$	1000	$\mathbf{4.22 \times 10^{-4}}$	0.66	**2372**
Brute force (exact)	$\beta/C_5 \to C_7^{\mathrm{eq}}$	—	1.08×10^{-3}	—	926
	$C_7^{\mathrm{eq}} \to \beta/C_5$	—	4.48×10^{-4}	—	2234

†Run terminated before 1000 trials (degeneracy / non-convergence).

4.3 The method search: SMC degenerates, SIR is unstable, RRS converges

Table 1 summarizes the three schemes against the brute-force reference. The pattern is stark. SMC (no resampling) completes the A → B direction but stalls after only 31 of 1000 trials in the reverse direction, where weight degeneracy inflates the population and the compute budget is exhausted before convergence. SIR (global resampling) never settles: its per-trial relative standard deviation is 2–3× the mean, and both directions were terminated early (118 and 533 trials) as the estimate wandered rather than converged—aggressive global resampling destroys walker diversity here. RRS is the *only* scheme that runs all 1000 trials in both directions and converges cleanly, with per-trial spread comparable to or better than the alternatives and no catastrophic tail behavior.

4.4 RRS recovers the brute-force rate; the naive committor does not

The running-rate traces for the selected RRS estimator (Fig. 5a,b) show the reweighted estimate r^{IS} settling onto the brute-force line r^{BF} within a few hundred trials, far above the naive committor estimate r^{NS}. The headline validation is Fig. 5c,d. Using the same learned committor two ways gives very different answers. Applied naively—flux times the mean committor over escape configurations ("N.S.")—the estimate is badly biased: $\mathrm{MFPT}_{\mathrm{AB}} = 926$ fs against a naive 5412 fs, and $\mathrm{MFPT}_{\mathrm{BA}} = 2234$ fs against a naive 4028 fs, i.e. up to $\sim 6\times$ too slow, because an imperfect committor does not satisfy the exact flux relation. Using the same committor as an importance-sampling guide and reweighting with RRS ("I.S.") corrects this to 950 fs (+2.6%) and 2372 fs (+6.2%), both within error of the brute-force bars ("B.F."). The reweighting turns an imperfect learned object into a quantitatively correct rate—the property promised by Eq. (7)—provided the weight variance is controlled, which is exactly what RRS delivers.

5 Conclusions

We set out to answer a concrete question—which weight-control scheme makes importance-sampled protein kinetics converge—and to answer it agentically. On the $\beta/C_5 \leftrightarrow C_7^{\mathrm{eq}}$ isomerization of alanine dipeptide, with an exact brute-force rate in both directions, a Claude-driven method search compared sequential importance

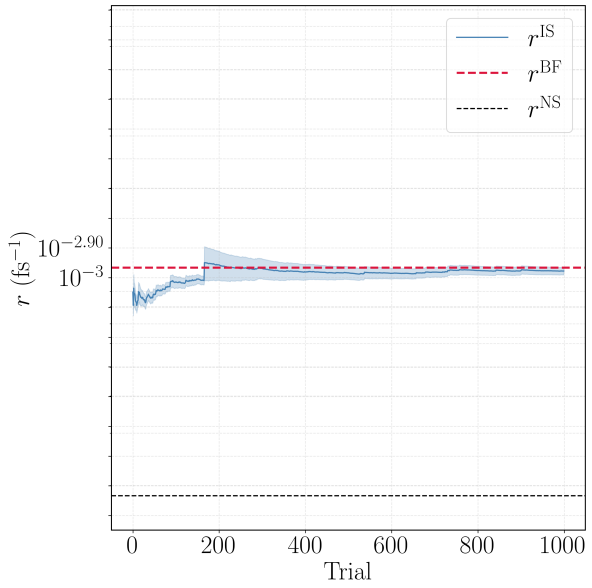

(a) Rate convergence, $\beta/C_5 \to C_7^{\mathrm{eq}}$

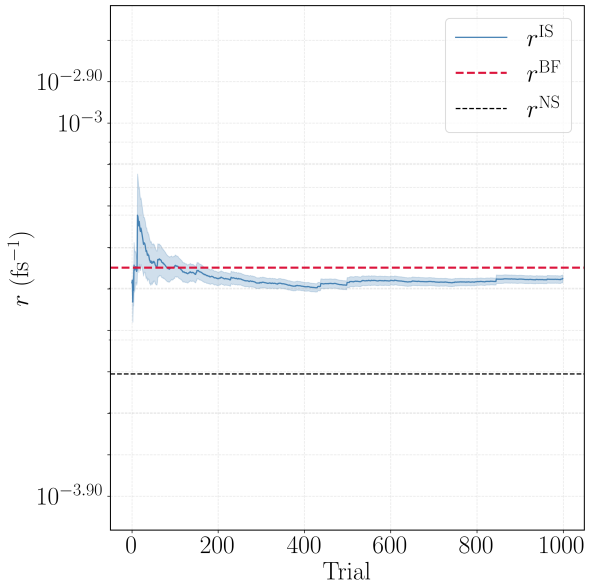

(b) Rate convergence, $C_7^{\mathrm{eq}} \to \beta/C_5$

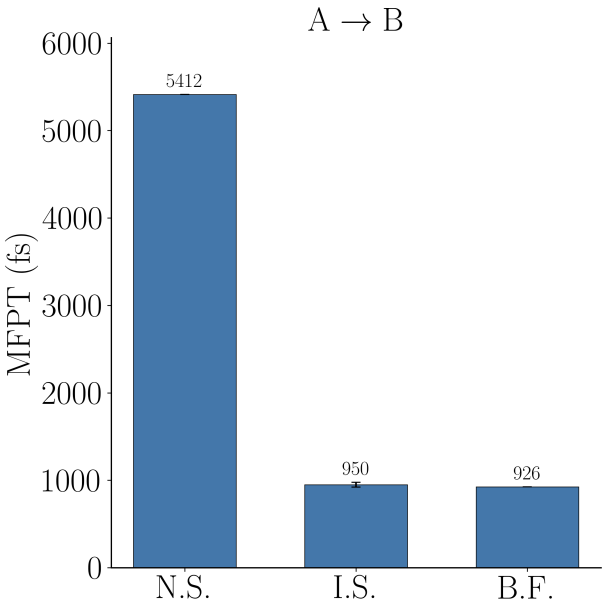

(c) MFPT, $\beta/C_5 \to C_7^{\mathrm{eq}}$

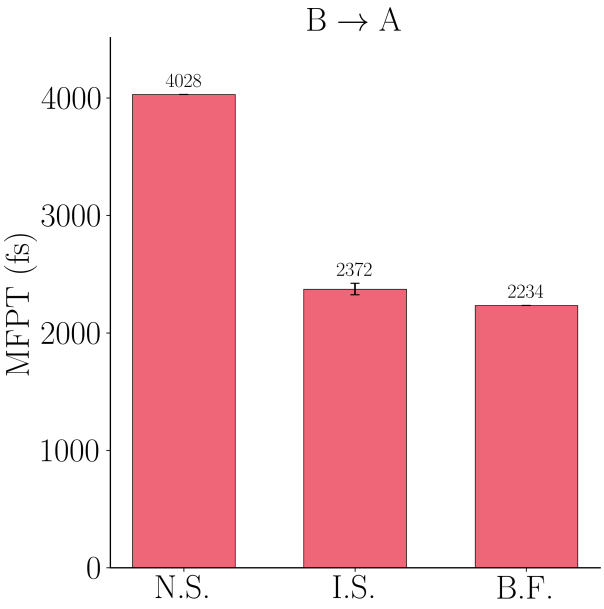

(d) MFPT, $C_7^{\mathrm{eq}} \to \beta/C_5$

Figure 5: RRS recovers the exact kinetics. (a,b) Running estimate of the reweighted rate r^{IS} (with confidence band) versus trial, compared to the brute-force reference r^{BF} (red dashed) and the naive committor estimate r^{NS} (black dashed); the importance-sampling estimate converges onto the exact rate in both directions. (c,d) Mean first-passage time from the naive committor (N.S.), the RRS importance-sampling estimate (I.S.), and brute-force MD (B.F.): the naive estimate is up to 6× too slow, while reweighted RRS matches brute force to within 3–6%.

sampling, global resampling, and a branching resampled random walk, and selected RRS: it is the only scheme that completes both directions, controls the heavy-tailed weight distribution, and recovers the true rate to within 3–6%, where the naive committor estimate is up to 6× wrong. The scientific content is the estimator and its validation; the methodological content is that the estimator was *found*, not hand-tuned, by an agent reading its own variance diagnostics against ground truth.

Several limitations frame the outlook. The benchmark is a single peptide in vacuum; the exponential weight growth of Eq. (8) is milder here than it will be for a larger, solvated system, so the real test of RRS is whether it scales. The B $\to$ A direction retains a slightly larger residual bias and per-trial variance than A $\to$ B, pointing to the weight band $[w_{\min}, w_{\max}]$ and the ESS trigger as hyperparameters worth autotuning—themselves natural targets for the agent to optimize. More broadly, the loop of Fig. 2 is not specific to resampling: an agent that can propose an estimator, implement it, run it, and read back a variance diagnostic against a trusted reference is a general engine for numerical-scheme discovery in molecular simulation. Turning Claude loose on the functional *form* of the bias, not just the resampling, is the direction we find most promising.

Acknowledgements

This work was carried out during the *Built with Claude: Life Sciences* hackathon, in partnership with Gladstone Institutes. We thank the organizers and the Anthropic team, and gratefully acknowledge Anthropic API credits provided for the event. The committor and rate-estimation framework builds on the authors' prior work on importance sampling for rare transition paths [7, 8].

Code and Data Availability

Code and data supporting this study—the trained importance-function checkpoints, the brute-force references, and the Claude Code agent transcript—are available from the author on reasonable request. Every result uses random seed 2026 and is reproducible from the accompanying pipeline (equilibration, brute-force MD, committor training, branching-random-walk rate estimation, and plotting).

License

References

[1] Hendrik Anthony Kramers. "Brownian motion in a field of force and the diffusion model of chemical reactions". In: *physica* 7.4 (1940), pp. 284–304.

[2] Peter Hänggi, Peter Talkner, and Michal Borkovec. "Reaction-rate theory: fifty years after Kramers". In: *Reviews of modern physics* 62.2 (1990), p. 251.

[3] Arthur F Voter, Francesco Montalenti, and Timothy C Germann. "Extending the time scale in atomistic simulation of materials". In: *Annual review of materials research* 32.1 (2002), pp. 321–346.

[4] Lydia Freddolino et al. "Challenges in protein-folding simulations". In: *Nature physics* 6.10 (2010), pp. 751–758.

[5] Wei Cai et al. "Importance sampling of rare transition events in Markov processes". In: *Physical Review E* 66.4 (2002), p. 046703.

[6] Maurice de Koning et al. "Adaptive importance sampling Monte Carlo simulation of rare transition events". In: *The Journal of chemical physics* 122.7 (2005).

[7] Michael Kim and Wei Cai. "Accelerated Markov Chain Monte Carlo Simulation via Neural Network-Driven Importance Sampling". In: *arXiv preprint arXiv:2602.12294* (2026).

[8] Michael Kim and Wei Cai. "Accelerated Langevin Dynamics Simulation via Neural Network-Driven Importance Sampling". In: *chemRxiv preprint* (2026).

[9] Eric Vanden-Eijnden. "Transition path theory". In: *Computer Simulations in Condensed Matter Systems: From Materials to Chemical Biology Volume 1*. Springer, 2006, pp. 453–493.

[10] Eric Vanden-Eijnden et al. "Transition-path theory and path-finding algorithms for the study of rare events." In: *Annual review of physical chemistry* 61 (2010), pp. 391–420.

[11] Terrell Hill. *Free energy transduction in biology: the steady-state kinetic and thermodynamic formalism*. Elsevier, 2012.

[12] Manon Baudel, Arnaud Guyader, and Tony Lelièvre. "On the Hill relation and the mean reaction time for metastable processes". In: *Stochastic Processes and their Applications* 155 (2023), pp. 393–436.

[13] Yuehaw Khoo, Jianfeng Lu, and Lexing Ying. "Solving for high-dimensional committor functions using artificial neural networks". In: *Research in the Mathematical Sciences* 6.1 (2019), p. 1.

[14] Qianxiao Li, Bo Lin, and Weiqing Ren. "Computing committor functions for the study of rare events using deep learning". In: *The Journal of Chemical Physics* 151.5 (2019).

[15] Andrew R Mitchell and Grant M Rotskoff. "Committor guided estimates of molecular transition rates". In: *Journal of Chemical Theory and Computation* 20.21 (2024), pp. 9378–9393.

[16] John Strahan et al. "Predicting rare events using neural networks and short-trajectory data". In: *Journal of computational physics* 488 (2023), p. 112152.

[17] Arnaud Doucet, Nando De Freitas, and Neil Gordon. "An introduction to sequential Monte Carlo methods". In: *Sequential Monte Carlo methods in practice*. Springer, 2001, pp. 3–14.

[18] Jun S Liu, Rong Chen, and Tanya Logvinenko. "A theoretical framework for sequential importance sampling with resampling". In: *Sequential Monte Carlo methods in practice*. Springer, 2001, pp. 225–246.

[19] Luca Martino, Víctor Elvira, and Francisco Louzada. "Effective sample size for importance sampling based on discrepancy measures". In: *Signal Processing* 131 (2017), pp. 386–401.

[20] Pierre Del Moral, Arnaud Doucet, and Ajay Jasra. "On adaptive resampling strategies for sequential Monte Carlo methods". In: (2012).

[21] Zhan Shi et al. *Branching random walks*. Vol. 2151. Springer, 2015.

[22] Jiří Vorba and Jaroslav Křivánek. "Adjoint-driven Russian roulette and splitting in light transport simulation". In: *ACM Transactions on Graphics (TOG)* 35.4 (2016), pp. 1–11.

[23] Justin F Talbot. *Importance resampling for global illumination*. Brigham Young University, 2005.

[24] Anthropic. *Claude Code: agentic coding at the command line*. https://code.claude.com/docs. 2025.

[25] Anthropic. *Building effective agents*. https://www.anthropic.com/engineering/building-effective-agents. 2024.

[26] Rubicelia Vargas et al. "Conformational study of the alanine dipeptide at the MP2 and DFT levels". In: *The Journal of Physical Chemistry A* 106.13 (2002), pp. 3213–3218.

[27] James A Maier et al. "ff14SB: improving the accuracy of protein side chain and backbone parameters from ff99SB". In: *Journal of chemical theory and computation* 11.8 (2015), pp. 3696–3713.

[28] Peter Eastman et al. "OpenMM 7: Rapid development of high performance algorithms for molecular dynamics". In: *PLoS computational biology* 13.7 (2017), e1005659.

[29] Jean-Paul Ryckaert, Giovanni Ciccotti, and Herman JC Berendsen. "Numerical integration of the cartesian equations of motion of a system with constraints: molecular dynamics of n-alkanes". In: *Journal of computational physics* 23.3 (1977), pp. 327–341.

[30] Kristof Schütt et al. "Schnet: A continuous-filter convolutional neural network for modeling quantum interactions". In: *Advances in neural information processing systems* 30 (2017).

[31] Richard S Sutton, Andrew G Barto, et al. *Reinforcement learning: An introduction*. Vol. 1. 1. MIT press Cambridge, 1998.

[32] Avishek Das et al. "Reinforcement learning of rare diffusive dynamics". In: *The Journal of Chemical Physics* 155.13 (2021).

[33] Diederik P Kingma. "Adam: A method for stochastic optimization". In: *arXiv preprint arXiv:1412.6980* (2014).

[34] Pavel Izmailov et al. "Averaging weights leads to wider optima and better generalization". In: *arXiv preprint arXiv:1803.05407* (2018).