Queen Mary University of London

Stochastic Path Sampler for Lattice Field Theory

S. Chen, M. Qian, G. Aarts, B. Lucini, K. Zhou, arXiv:2606.13790

Biagio Lucini

School of Mathematical Sciences, Queen Mary University of London

Bridging machine learning and statistical field theory through generative and diffusion processes

ECT*, Trento, 7–11 September 2026

Code and data:
https://github.com/Chern2/Stochastic-Path-Sampler-For-Lattice-Field-Theory

Supported by THE ROYAL SOCIETY, UKRI Science and Technology Facilities Council

Sampling in lattice field theory and where SPS sits

- Target known only up to normalisation: $\pi^*(\phi) = e^{-S(\phi)}/Z$, with Z intractable
- HMC: critical slowing down near phase transitions and the continuum limit, topological freezing between sectors

Neural samplers for LFT (reviews: [1])

Variational, data-free: normalizing flows [2], CNFs [3], autoregressive nets, stochastic NFs [4]. Mode collapse on multimodal targets [5].

Data-driven: GANs, diffusion / score models [6]. Need HMC reference samples.

Path-space samplers in ML [7]: CRAFT, NETS, PIS, DDS, CMCD. Path-space KL along a *prescribed* annealing path.

This work: Stochastic Path Sampler [8]

- Learn **both** forward and backward stochastic dynamics between a Gaussian prior and e^{-S}
- Intermediate marginals emerge implicitly, no annealing schedule
- Trained from $S(\phi)$ alone: **no reference data**
- **Exact**: trajectory-level Metropolis–Hastings correction

[1] Cranmer et al. 2309.01156; Aarts et al. 2501.05580; Tomiya JPSJ 2025. [2] Albergo, Kanwar, Shanahan 1904.12072; Nicoli et al. 2007.07115; Kanwar et al. 2003.06413; Boyda et al. 2008.05456; Albergo et al. 2106.05934, 2202.11712. [3] Gerdes et al. 2207.00283; de Haan et al. 2110.02673. [4] Wu, Köhler, Noé 2002.06707; Caselle et al. 2201.08862; Bulgarelli et al. 2412.00200. [5] Hackett et al. 2107.00734; Chen et al. 2502.02127. [6] Zhou et al. 1810.12879; Wang, Aarts, Zhou 2309.17082; Zhu et al. 2502.05504; Vega et al. 2510.26081; Aarts et al. 2601.19552; Alharazin et al. 2602.09045; Komijani et al. 2605.06134. [7] Matthews et al. 2201.13117; Albergo, Vanden-Eijnden 2410.02711; Zhang, Chen 2111.15141; Vargas et al. 2302.13834, 2307.01050; Berner et al. 2211.01364. [8] Chen, Qian, Aarts, Lucini, Zhou, arXiv:2606.13790.

Physical idea: reversibility as the training objective

- Clausius inequality: irreversibility $\Leftrightarrow$ *nonnegative entropy production*[1]
- Stochastic thermodynamics: entropy production measures the asymmetry between forward and time-reversed trajectory ensembles

$$\Sigma = \left\langle \int_0^T dt\, \delta A_\tau(t) \right\rangle + \left\langle \log \frac{\pi_0}{\pi^*} \right\rangle \geq 0$$

- **Strategy**: drive prior π_0 (A) to target π^* (B); learn the dynamics so the transport is as *reversible* as possible
- Reversible transport (trajectory-level balance[2]) $\Rightarrow$ endpoint of the forward process follows π^*

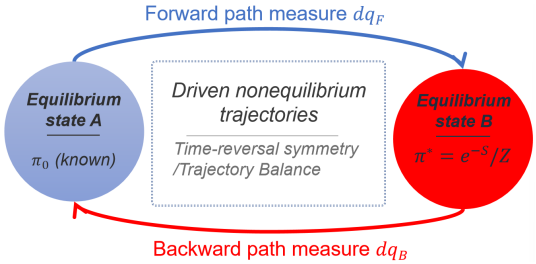

Forward process induces dq_{F}, backward process from π^* induces dq_{B} on the same trajectory space.

[1] Pachter, Yang, Dill, Nat. Rev. Phys. 2024, 2310.06070

[2] GFlowNets: Bengio et al. 2106.04399; Malkin et al. 2201.13259

Discrete Langevin paths with learnable kernels

Time interval $[0,1]$ discretised in T steps, $\Delta t = 1/T$; $\boldsymbol{s}_i \in \mathbb{R}^{|\Lambda|}$ is a lattice configuration.

Forward process (prior $\pi_0 = \mathcal{N}(0, I) \to$ target)

$$\boldsymbol{s}_{i+1} = \boldsymbol{s}_i + \sigma_\theta^2(t_i)\, \boldsymbol{K}_{\theta,\mathrm{F}}(\boldsymbol{s}_i, t_i)\, \Delta t + \sigma_\theta(t_i)\, \boldsymbol{\xi}_i \sqrt{\Delta t}, \qquad \boldsymbol{\xi}_i \sim \mathcal{N}(0, I)$$

$$dq_{\mathrm{F}}(\tau) = \pi_0(\boldsymbol{s}_0) \prod_{i=0}^{T-1} q_{\mathrm{F}}(\boldsymbol{s}_{i+1} \mid \boldsymbol{s}_i)\, d\tau$$

Backward process (target $\to$ prior), independent drift network

$$\boldsymbol{s}_i = \boldsymbol{s}_{i+1} + \sigma_\theta^2(t_i)\, \boldsymbol{K}_{\theta,\mathrm{B}}(\boldsymbol{s}_{i+1}, t_{i+1})\, \Delta t + \sigma_\theta(t_i)\, \boldsymbol{\eta}_{i+1} \sqrt{\Delta t}, \qquad dq_{\mathrm{B}}(\tau) = \pi^*(\boldsymbol{s}_T) \prod_{i=0}^{T-1} q_{\mathrm{B}}(\boldsymbol{s}_i \mid \boldsymbol{s}_{i+1})\, d\tau$$

- Stochastic quantisation[3] / MALA[4]: $\boldsymbol{K}_{\theta,\mathrm{F}} = -\nabla S$. Here **both drifts and $\sigma_\theta(t)$ are neural networks**; Gaussian kernels give path probabilities in closed form

[3] Parisi, Wu 1981; Damgaard, Hüffel, Phys. Rept. 1987; diffusion models as SQ: Wang, Aarts, Zhou 2309.17082

[4] Roberts, Rosenthal 1998

Path probability ratio as entropy production: the loss

- Accumulated affinity plus boundary term along a path τ

$$\Sigma(\tau) = \sum_{i=0}^{T-1} \log \frac{q_{\mathrm{F}}(\boldsymbol{s}_{i+1} \mid \boldsymbol{s}_i)}{q_{\mathrm{B}}(\boldsymbol{s}_i \mid \boldsymbol{s}_{i+1})} + \log \frac{\pi_0(\boldsymbol{s}_0)}{\tilde{\pi}^*(\boldsymbol{s}_T)}$$

- Its forward average is a **KL divergence on path space**, computable with the unnormalised target

$$\mathcal{L}(\theta) = \mathbb{E}_{dq_{\mathrm{F}}}[\Sigma(\tau)] = \int dq_{\mathrm{F}}(\tau)\, \log \frac{dq_{\mathrm{F}}(\tau)}{d\tilde{q}_{\mathrm{B}}(\tau)}, \qquad d\tilde{q}_{\mathrm{B}} = \tilde{\pi}^*(\boldsymbol{s}_T) \prod_i q_{\mathrm{B}}(\boldsymbol{s}_i \mid \boldsymbol{s}_{i+1})\, d\tau$$

- Jarzynski[5]: at the optimum $\mathcal{L} = -\log Z$ and the path ratio becomes a *global constant*

$$q_\theta(\boldsymbol{s}_T; \tau) \equiv \frac{\prod_i q_{\mathrm{F}}(\boldsymbol{s}_{i+1} \mid \boldsymbol{s}_i)}{\prod_i q_{\mathrm{B}}(\boldsymbol{s}_i \mid \boldsymbol{s}_{i+1})}\, \pi_0(\boldsymbol{s}_0) \;\longrightarrow\; q_\theta(\phi) \propto \tilde{\pi}^*(\phi)$$

  At finite training $q_\theta(\boldsymbol{s}_T; \tau)$ still depends on the path: residual irreversibility

- Minimised over $\boldsymbol{K}_{\theta,\mathrm{F}}$, $\boldsymbol{K}_{\theta,\mathrm{B}}$, σ_θ by backpropagating through the whole trajectory. Only $S(\phi)$ enters: **no training data**

[5] Jarzynski 2011; on the lattice: Caselle et al. 1604.05544; in SNFs: Caselle et al. 2201.08862

Exactness: trajectory-level Independence Metropolis–Hastings

Residual irreversibility from finite capacity, training and discretisation. Correct on the *extended trajectory space*:

$$P_{\text{accept}}(\boldsymbol{s}'_T \mid \boldsymbol{s}_T) = \min\left(1,\ \frac{\tilde{\pi}^*(\boldsymbol{s}'_T)}{\tilde{\pi}^*(\boldsymbol{s}_T)}\ \frac{\prod_i q_{\mathrm{B}}(\boldsymbol{s}'_i \mid \boldsymbol{s}'_{i+1})}{\prod_i q_{\mathrm{F}}(\boldsymbol{s}'_{i+1} \mid \boldsymbol{s}'_i)\,\pi_0(\boldsymbol{s}'_0)}\ \frac{\prod_i q_{\mathrm{F}}(\boldsymbol{s}_{i+1} \mid \boldsymbol{s}_i)\,\pi_0(\boldsymbol{s}_0)}{\prod_i q_{\mathrm{B}}(\boldsymbol{s}_i \mid \boldsymbol{s}_{i+1})}\right)$$

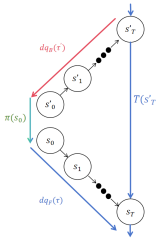

(a) $T(\boldsymbol{s}'_T \to \boldsymbol{s}_T)$

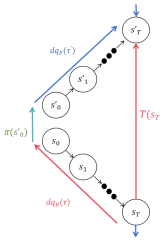

(b) $T(\boldsymbol{s}_T \to \boldsymbol{s}'_T)$

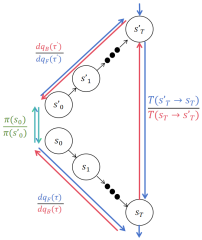

(c) ratio

- Independent proposals with weight $\tilde{\pi}^*/q_\theta$, as in flow-based IMH[6]
- Detailed balance over whole trajectories
- Needs full support coverage: checked via magnetisation histograms (later)

[6] Albergo, Kanwar, Shanahan 1904.12072

Networks, training and generation

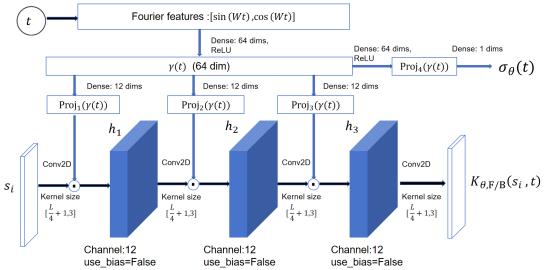

- Time-conditional CNN: three residual conv blocks, circular padding (exact PBC), kernels $[L/4 + 1, 3]$
- Fourier time embedding modulates each block
- Z_2-equivariant drift, $\boldsymbol{K}(-\phi, t) = -\boldsymbol{K}(\phi, t)$; kernels symmetrised over D_4 / D_2

Training

Sample $\boldsymbol{s}_0 \sim \mathcal{N}(0, I)$, run forward dynamics, accumulate $\log q_\theta(\boldsymbol{s}_T; \tau)$, loss $\mathcal{L} = \mathbb{E}_{q_F}[\log q_\theta(\boldsymbol{s}_T; \tau) - \log \tilde{\pi}^*(\boldsymbol{s}_T)]$. Adam, 15,000 steps, batch 12, $T_{\mathrm{train}} = 250$.

Generation

Same forward dynamics with $T_{\mathrm{gen}} = 2500$ (time embedding allows $T_{\mathrm{gen}} \neq T_{\mathrm{train}}$), then IMH.

Fully dataless: one network per (κ, L), trained from $S(\phi)$ alone

Benchmark: two-dimensional ϕ^4 theory

$$S(\phi) = \sum_{x\in\Lambda}\Big(-2\kappa \sum_{\mu=1}^{2} \phi_x\phi_{x+\hat{\mu}} + (1-2\lambda)\phi_x^2 + \lambda\phi_x^4 \Big)$$

- Geometry $L \times 8$ (finite-temperature setup of Nicoli et al.[7]), $L = 16, 32, 48, 64$
- $\lambda = 0.022$, scan $\kappa \in [0.20, 0.30]$, $\Delta\kappa = 0.01$
- Infinite volume: $\kappa_c \approx 0.239$. On these lattices the crossover and susceptibility peak sit at $\kappa \approx 0.27$: the pseudocritical region
- HMC reference: 2×10^6 trajectories, 100 leapfrog steps, measured every 500

Observables

$$m = \tfrac{1}{|\Lambda|}\textstyle\sum_x \phi_x, \quad |M| = \langle |m| \rangle$$
$$\chi = |\Lambda|\big(\langle m^2\rangle - \langle |m|\rangle^2\big)$$
$$F = -\tfrac{1}{|\Lambda|}\log Z = -\tfrac{1}{|\Lambda|}\log \mathbb{E}_{dq_F}\big[w(\phi;\tau)\big]$$

with importance weight $w(\phi;\tau) = \tilde{\pi}^*(\phi)/q_\theta(\phi;\tau)$

Three curves throughout: HMC, raw SPS, **SPS+IMH**

[7] Nicoli et al., PRL 2021, 2007.07115

Absolute magnetisation

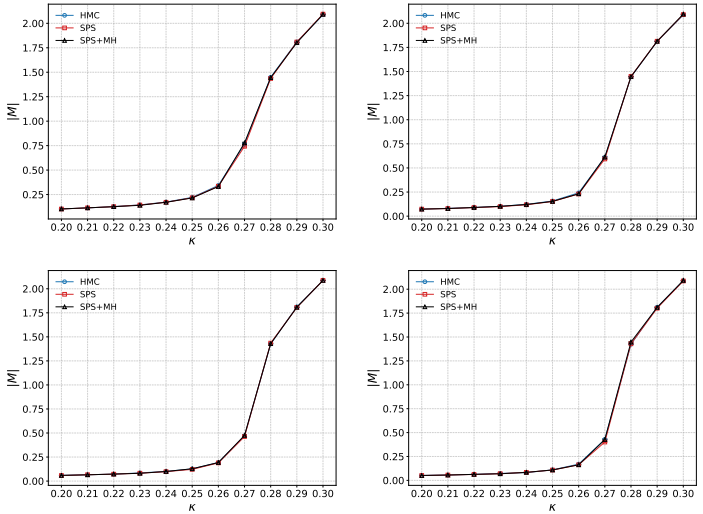

- HMC, SPS and SPS+IMH essentially indistinguishable over the full κ range
- Jump at $\kappa \simeq 0.27$ reproduced on every volume
- Raw SPS deviates up to $\sim 1.5\sigma$ for $\kappa = 0.26$–0.28, where the distribution becomes bimodal
- IMH removes the residual deviation
- Both Z_2 sectors populated in the broken phase, no mode collapse (histograms in backup)

Susceptibility

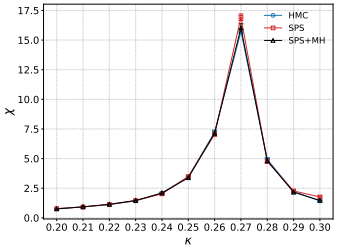

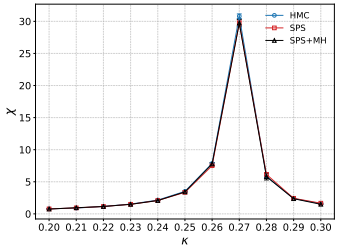

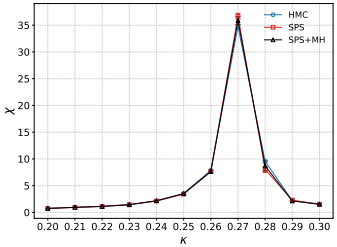

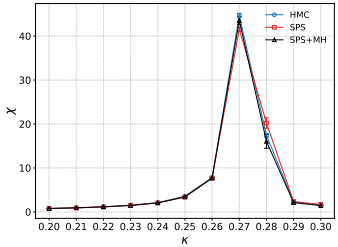

- Sharp peak at $\kappa \simeq 0.27$, growing with L, reproduced
- Symmetric phase: all three agree within 1σ
- Raw SPS off by $> 3\sigma$ for $\kappa \gtrsim 0.27$: nonzero diffusion at the last step keeps feeding rare, high-action tails
- **SPS+IMH agrees with HMC for all κ and L**

Free energy density

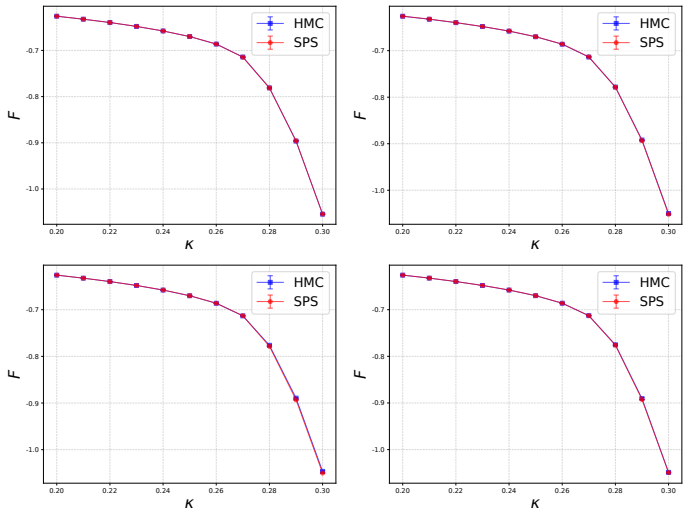

- SPS gives direct access to $\log Z$ by reweighting trajectories with $w = \tilde{\pi}^*/q_\theta$, as for flows[8]

$$F = -\frac{1}{|\Lambda|} \log \mathbb{E}_{dq_F}[w(\phi;\tau)]$$

- Unbiased even though w is path dependent
- SPS (10,240 samples) vs. HMC (160,000): agreement within errors for all κ, L

[8] Nicoli et al. 2007.07115

IMH acceptance rate

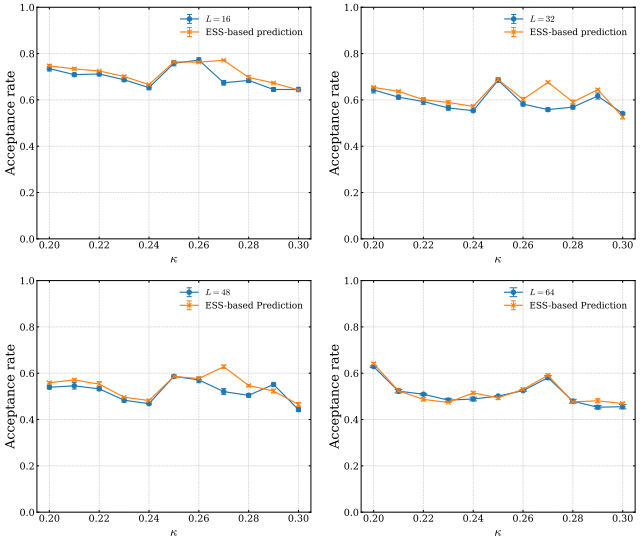

- Circles: measured IMH acceptance over 4096 proposals. Crosses: prediction from the effective sample size of the importance weights (formula in backup)
- Agreement is good; largest gap at $\kappa = 0.27$ from limited weight statistics
- $L = 16$: 0.60–0.76; $L = 64$: 0.45–0.62. Pseudocritical region no worse than the symmetric phase
- Log-weight variance is extensive, so the ESS deficit grows like $|\Lambda|$ but the acceptance deficit only like $\sqrt{|\Lambda|}$

Training dynamics and cost at $\kappa = 0.27$

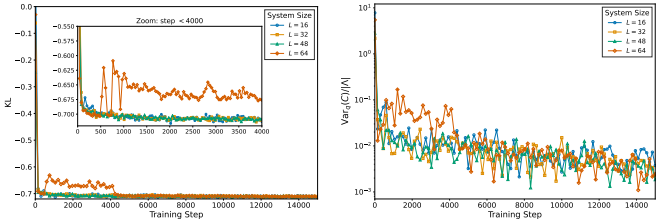

Left: path-space KL vs. training step. Right: $\mathrm{Var}_q(C)/|\Lambda|$, $C = \log \tilde{\pi}^* - \log q_\theta$ (log scale).

- KL drops steeply in the first $\sim$ 5,000 steps to ≈ -0.714 for every L
- Log-weight variance falls sharply in the first $\sim$ 1,000 steps, then decreases slowly: lower variance $\leftrightarrow$ higher acceptance
- Cost almost flat in the volume at fixed training length (single RTX 5090m GPU)

L	16	32	48	64
Training ($T = 250$, 15,000 steps, batch 12)	1.1 h	1.1 h	1.1 h	1.2 h
Generation ($T = 2500$, 4096 samples)	1.1 min	1.5 min	1.9 min	2.1 min

Autocorrelation: critical slowing down at $\kappa = 0.27$, $L = 64$

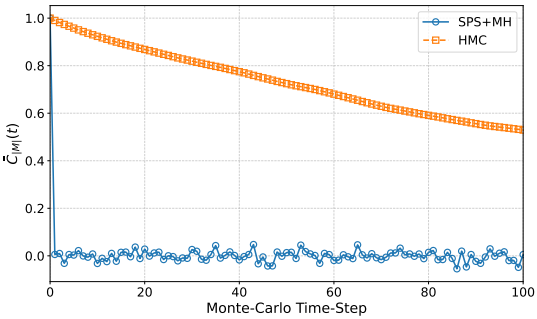

Normalised autocorrelation of $|M|$; SPS+IMH (circles) vs. HMC (squares).

$$\tau_{\rm int} = \tfrac{1}{2} + \frac{1}{C(0)} \sum_{t=1}^{t_{\rm max}} C(t)$$

automatic window: stop at $t > 6\,\tau_{\rm int}(t)$

Integrated autocorrelation time of $|M|$

$$\tau_{\rm int}^{\rm SPS+IMH} \simeq 0.5 \quad \text{vs.} \quad \tau_{\rm int}^{\rm HMC} \simeq 160$$

- Proposals are independent: IMH rejections are the only source of correlation
- Caveat: IMH steps vs. HMC trajectories. Not cost-matched; training and generation cost must be included

Conclusions and outlook

Stochastic Path Sampler

- Learn forward and backward Langevin dynamics between a Gaussian prior and e^{-S} by minimising **path-space entropy production** (a KL divergence). Data-free: only $S(\phi)$ is evaluated
- Exact sampling via a **trajectory-level Independence Metropolis–Hastings** correction

Results in 2D ϕ^4

- $|M|$, χ, F reproduce HMC in the symmetric, pseudocritical and broken phases up to $L = 64$
- Acceptance stays moderate ($\gtrsim 0.45$), cost grows only mildly with volume, $\tau_{\mathrm{int}}(|M|)$ at the pseudocritical point reduced by orders of magnitude

Next

- Compact $U(1)$ in 2D: does the stochastic path construction ease topological freezing?
- Smaller Δt to raise acceptance; local, scalable architectures for 3D and 4D (current kernels grow with L)

arXiv:2606.13790 · code and data: `github.com/Chern2/Stochastic-Path-Sampler-For-Lattice-Field-Theory`

Backup Slides

Mode coverage across the transition ($L = 64$, raw SPS vs. HMC)

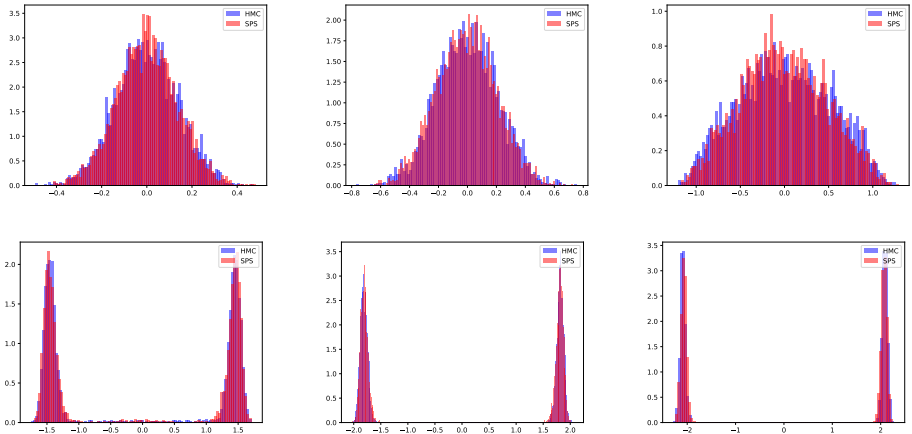

- HMC (blue) almost entirely hidden under raw SPS (red). Both Z_2 sectors populated: **no mode collapse**, so IMH is legitimate
- At $\kappa = 0.27$ nonnegligible weight sits between the modes: this is what stresses KL-based training

IMH acceptance rate and ESS (i), the exact relation

Proposal q (the forward path measure), weights $w = \tilde{\pi}^*/q > 0$ known up to $Z = \mathbb{E}_q[w]$; IMH proposes $y \sim q$, accepts with $\min(1, w(y)/w(x))$. ESS: $\mathrm{ESS}/N_s \simeq \mathbb{E}[w]^2/\mathbb{E}[w^2]$.

Derivation

At stationarity $x \sim \pi$, whose density relative to q is $w/\mathbb{E}[w]$; the proposal y is an independent draw from q:

$$\langle A \rangle = \frac{\mathbb{E}_{x,y\sim q}\big[\, w(x) \min\big(1, \tfrac{w(y)}{w(x)}\big)\big]}{\mathbb{E}[w]} = \frac{\mathbb{E}\big[\min(w, w')\big]}{\mathbb{E}[w]} \quad \text{since} \quad w \min(1, w'/w) = \min(w, w')$$

The reweighting factor is absorbed exactly and Z cancels. With $\min(a, b) = \frac{1}{2}(a + b - |a - b|)$,

$$\boxed{\langle A \rangle = 1 - \frac{\mathbb{E}\,|w - w'|}{2\,\mathbb{E}[w]} = 1 - G_w}$$

G_w = Gini coefficient of the weights $= \|\pi \otimes q - q \otimes \pi\|_{\mathrm{TV}}$

No assumption on the shape of the weight distribution. Estimated from the *same batch* used for the ESS, without running the chain: sort $w_{(1)} \le \cdots \le w_{(N_s)}$, then $\langle A \rangle \simeq \frac{2}{N_s(N_s-1)\bar{w}} \sum_k (N_s - k)\, w_{(k)}$, an all-pairs U-statistic ($O(N_s \log N_s)$, jackknife errors).

Follows from the symmetric representation of MH kernels, $\iint \min(\pi(dx)Q(x, dy),\ \pi(dy)Q(y, dx))$: Tierney, Ann. Appl. Probab. 8 (1998) 1; Billera, Diaconis, Statist. Sci. 16 (2001) 335. Metropolized independent sampling vs. weight distribution: Liu, Stat. Comput. 6 (1996) 113. Gini mean difference: David, *Order Statistics* (1981).

IMH acceptance rate and ESS (ii), log-normal weights

Model $X \equiv \log w \sim \mathcal{N}(\mu, \sigma^2)$. Then $\mathbb{E}[w^n] = e^{n\mu + n^2\sigma^2/2}$, so $\mathrm{ESS}/N_s = \mathbb{E}[w]^2/\mathbb{E}[w^2] = e^{-\sigma^2}$.

Acceptance rate

Current state $\sim \pi \propto w\,q$: exponential tilting of X by e^X gives $X_{\rm cur} \sim \mathcal{N}(\mu + \sigma^2, \sigma^2)$; proposal $X' \sim \mathcal{N}(\mu, \sigma^2)$. Hence $D \equiv X' - X_{\rm cur} \sim \mathcal{N}(m, s^2)$ with $m = -\sigma^2$, $s^2 = 2\sigma^2$, and

$$\begin{aligned}\langle A \rangle = \mathbb{E}\big[\min(1, e^D)\big] &= P(D>0) + \mathbb{E}\big[e^D\,\mathbf{1}_{D<0}\big] = \Phi\Big(\frac{m}{s}\Big) + e^{m+s^2/2}\,\Phi\Big(-\frac{m+s^2}{s}\Big) \\ &= 2\,\Phi\Big(-\frac{\sigma}{\sqrt{2}}\Big) = \mathrm{erfc}\Big(\frac{\sigma}{2}\Big) \qquad \text{since } m + s^2/2 = 0.\end{aligned}$$

Eliminating σ between the two closed forms: $\boxed{\langle A \rangle = \mathrm{erfc}\Big(\tfrac{1}{2}\sqrt{-\ln(\mathrm{ESS}/N_s)}\Big)}$

IMH analogue of the erfc formula for the HMC acceptance. Also from (i) with the classical result that the Gini coefficient of a log-normal is $\mathrm{erf}(\sigma/2)$.

HMC: Gupta, Irbäck, Karsch, Petersson, Phys. Lett. B 242 (1990) 437. Same Gaussian expectation in pseudo-marginal MCMC tuning: Pitt et al., J. Econometrics 171 (2012) 134; Doucet, Pitt, Deligiannidis, Kohn, Biometrika 102 (2015) 295. Log-normal Gini: Aitchison, Brown, *The Lognormal Distribution* (1957).

IMH acceptance rate and ESS (iii), volume scaling and a bound

For a sampler with extensive log-weight variance, $\sigma^2 = a|\Lambda|$, the two closed forms fix the volume dependence of both quantities.

Small variance, $\sigma^2 \ll 1$

$\mathrm{erfc}(x) = 1 - 2x/\sqrt{\pi} + O(x^3)$, $e^{-\sigma^2} = 1 - \sigma^2 + \dots$

$$1 - \langle A \rangle \simeq \sqrt{\sigma^2/\pi}, \qquad 1 - \frac{\mathrm{ESS}}{N_s} \simeq \sigma^2$$

ESS deficit linear in $|\Lambda|$, acceptance deficit $\propto \sqrt{|\Lambda|}$: a visible ESS drift with nearly flat acceptance is *consistent*, as observed.

$R \equiv \dfrac{(1-\langle A\rangle)^2}{1 - \mathrm{ESS}/N_s} = \dfrac{\mathrm{erf}(\sigma/2)^2}{1 - e^{-\sigma^2}} \to \dfrac{1}{\pi}$, in $[0.32, 0.39]$ for $\sigma^2 \lesssim 0.7$.

Large variance, $\sigma^2 \gg 1$

$\mathrm{erfc}(x) \simeq e^{-x^2}/(\sqrt{\pi}\, x)$

$$\langle A \rangle \simeq \sqrt{\tfrac{4}{\pi\sigma^2}}\, e^{-\sigma^2/4}, \qquad \frac{\mathrm{ESS}}{N_s} = e^{-\sigma^2}$$

$$\ln\langle A\rangle \simeq -\frac{a|\Lambda|}{4} - \frac{1}{2}\ln\frac{\pi a|\Lambda|}{4}, \qquad \ln\frac{\mathrm{ESS}}{N_s} = -a|\Lambda|$$

Both exponential in the volume, rates differing by a factor 4: $\langle A\rangle \propto (\mathrm{ESS}/N_s)^{1/4}$ up to a slow prefactor. They do *not* scale identically.

Bound and diagnostic: $\langle A\rangle/(\mathrm{ESS}/N_s) = e^{\sigma^2}\,\mathrm{erfc}(\sigma/2) \geq 0.909$, minimum at $\sigma^2 \approx 0.11$; the two cross at $\sigma^2 \approx 0.32$. A measured pair violating the bound signals non-log-normal weights that no σ^2 can absorb; the shape-independent relation (i) still applies and can be evaluated on the generated weights.