Infinite Dimensional Adjoint Sampler: Scalable Sampling on Function Spaces

Byoungwoo Park[1]**, Juho Lee**[1]**, Guan-Horng Liu**[2]
[1]KAIST, [2]FAIR at Meta
{bw.park, juholee}@kaist.ac.kr, ghliu@meta.com

Abstract

We present the adjoint sampler for infinite-dimensional function spaces, a stochastic optimal control (SOC)-based diffusion sampler that operates directly in function space and targets Gibbs-type distributions on infinite-dimensional Hilbert spaces. Our Function space Adjoint Sampler (FAS) generalizes Adjoint Sampling [16] to Hilbert spaces based on a SOC theory called stochastic maximum principle, yielding a simple and scalable matching-type objective for a functional representation. Through experiments, we show its effectiveness on transition-path sampling.

1 Infinite-Dimensional Sampling Problem

Sampling from a Gibbs distribution is a fundamental task across many fields. The goal is to draw samples from a target distribution π that is known only through an unnormalized energy $U : \mathcal{H} \to \mathbb{R}$:

$$\mathrm{d}\pi(\mathbf{x}) = \frac{1}{\mathcal{Z}} e^{-U(\mathbf{x})} \mathrm{d}\nu(\mathbf{x}), \quad \mathcal{Z} = \int_{\mathcal{H}} e^{-U(\mathbf{x})} \mathrm{d}\nu(\mathbf{x}), \tag{1}$$

where ν is the measure defined on some space $\mathcal{H}$. In finite dimensions, *e.g.*, $\mathcal{H} = \mathbb{R}^d$, the Gibbs density can be expressed with respect to Lebesgue measure *i.e.*, $\mathrm{d}\pi(\mathbf{x}) \propto e^{-U(\mathbf{x})}\mathrm{d}\mathbf{x}$. However, due to the absence of an equivalent formulation of Lebesgue measure for arbitrary Hilbert spaces, the expression in (1) is well-defined only after specifying an appropriate reference measure ν. The following lemma states a condition for the target law to be well defined.

Lemma 1 (Finite Partition)**.** *Let $U(\mathbf{x}) \geq \beta \|\mathbf{x}\|_{\mathcal{H}}^2 - C$ with $\beta, C > 0$. Then, for any $Tr(Q) < \infty$, choosing $\nu = \mathcal{N}(0, Q)$ in (1) results in a finite partition function $\mathcal{Z} < \infty$.*

Lemma 1 implies that a valid reference measure ν is the centered Gaussian $\mathcal{N}(0, Q)$ with any trace-class covariance Q, *i.e.*, $\mathrm{Tr}(Q) < \infty$. This choice guarantees $\mathcal{Z} < \infty$, so the target π in (1) is well-defined on $\mathcal{H}$. In this setting, a dynamical approach samples from (1) on $\mathcal{H}$ by evolving an $\mathcal{H}$-valued stochastic process that follows the gradient flow of U, in direct analogy overdamped Langevin dynamics [9, 3, 22]. Convergence guarantees rely on the choice of reference dynamics:

$$\mathrm{d}\mathbf{X}_t = \mathcal{A}\mathbf{X}_t \mathrm{d}t + \sigma_t d\mathbf{W}_t^Q, \quad \mathbf{X}_0 = \mathbf{x}_0 \in \mathcal{H}, \tag{2}$$

where $\mathcal{A} : \mathcal{H} \to \mathcal{H}$ is a linear operator, $\sigma_t : \mathbb{R}_{>0} \to \mathbb{R}_{>0}$ is a noise schedule, and $\mathbf{W}_t^Q$ is an infinite dimensional Wiener process. Under the mild assumptions, (2) admits an invariant Gaussian law

$$\nu_\infty = \mathcal{N}(0, Q_\infty), \quad Q_\infty = \sigma_\infty^2 \int_0^\infty e^{t\mathcal{A}} Q e^{t\mathcal{A}^\dagger} \mathrm{d}t, \tag{3}$$

where $\dagger$ denotes the Hermitian adjoint.[1] While convergence is guaranteed and direct simulation from (1) can be achieved via spatial and temporal discretization of the $\mathcal{H}$-valued processes in [9, 3, 22],

[1]An overview of relevant concepts on infinite-dimensional Hilbert spaces appears in Appendix A.

39th Conference on Neural Information Processing Systems (NeurIPS 2025) Workshop: Frontiers in Probabilistic Inference: Sampling Meets Learning.

the chains are of limited practical use as the mixing is effective only near equilibrium as $t \uparrow \infty$. They rely on asymptotic convergence to the invariant reference measure ν_∞ in (3) by Lemma 1.

Non Equilibrium Sampling To reach the target distribution (1) within a finite time horizon $T < \infty$, we recast the problem as a non-equilibrium sampling (NES) problem. In contrast to Langevin-type processes that use an invariant reference ν_∞ discussed above, we apply an importance-sampling (IS) via change of measure to sample the target in finite time:

$$\mathrm{d}\pi(\mathbf{x}) \propto e^{-U(\mathbf{x})} \frac{\mathrm{d}\nu_\infty(\mathbf{x})}{\mathrm{d}\nu_T^{\mathbf{x}_0}(\mathbf{x})} \mathrm{d}\nu_T^{\mathbf{x}_0}(\mathbf{x}) = e^{-U(\mathbf{x}) - \log \frac{\mathrm{d}\nu_T^{\mathbf{x}_0}(\mathbf{x})}{\mathrm{d}\nu_\infty(\mathbf{x})}} \mathrm{d}\nu_T^{\mathbf{x}_0}(\mathbf{x}), \tag{4}$$

where $\nu_T^{\mathbf{x}_0}$ is marginal law of the reference dynamics (2). Hence, for NES in infinite dimensions we need an extra correction term in the importance weight, given by the Radon–Nikodým derivative (RND). It corrects the mismatch between the two Gaussian measures so that after reweighting, samples drawn from $\nu_T^{\mathbf{x}_0}$ have the desired target π, while guarantees the π to be well-defined as stated in Lemma 1. However, the infinite dimensional RND exists only when the two measures are exactly compatible. Outside that narrow compatibility the measures are singular, and no density can be written. Luckily, we can derive the RND under Ornstein–Uhlenbeck semigroup *i.e.*, the dynamics has form in (2). Then, ν_t and ν_∞ becomes mutually absolutely continuous Gaussian measures [7]:

Theorem 2 (Explicit RND)**.** *Suppose Assumption B.1 holds. Then, For any $t > 0$ and initial condition $\mathbf{x}_0 \in \mathcal{H}$, $\nu_t^{\mathbf{x}_0} = \mathcal{N}(e^{t\mathcal{A}}\mathbf{x}_0, Q_t)$ and $\nu_\infty = \mathcal{N}(0, Q_\infty)$ are mutually absolutely continuous. Moreover, for any $\mathbf{x} \in \mathcal{H}$, the RND $q_t(\mathbf{x}_0, \mathbf{x}) := \frac{\mathrm{d}\nu_t^{\mathbf{x}_0}(\mathbf{x})}{\mathrm{d}\nu_\infty(\mathbf{x})}$ is given by:*

$$q_t(\mathbf{x}_0, \mathbf{x}) = C \exp\left[-\tfrac{1}{2}\langle \Theta_t^{-1}\mathbf{m}_t^{\mathbf{x}_0}, \mathbf{m}_t^{\mathbf{x}_0}\rangle_\infty + \langle \Theta_t^{-1}\mathbf{m}_t^{\mathbf{x}_0}, \mathbf{x}\rangle_\infty - \tfrac{1}{2}\langle e^{2t\mathcal{A}}\Theta_t^{-1}\mathbf{x}, \mathbf{x}\rangle_\infty \right], \tag{5}$$

where we denote $\mathbf{m}_t^{\mathbf{x}_0} := e^{t\mathcal{A}}\mathbf{x}_0$, $\Theta_t = 1 - e^{2t\mathcal{A}}$, $C = det(\Theta_t)^{-\frac{1}{2}}$ and $\langle u, v\rangle_\infty = \langle Q_\infty^{-\frac{1}{2}} u, Q_\infty^{-\frac{1}{2}} v\rangle_{\mathcal{H}}$.

Theorem 2 provides an explicit RND, so IS can in principle generate samples from π in (1). In practice, unfortunately, the importance weights are computed on a finite grid, and when that grid is refined the number of samples required grows exponentially. IS therefore becomes impractical.

2 Infinite Dimensional Adjoint Sampler

Inspired by recent advances in neural samplers [26, 2, 24, 25], we here to reformulate the NES problem as a variational inference problem on path space $\Omega = C([0, T], \mathcal{H})$ to build an efficient sampler by introducing a learnable *control* to steer the process (2) to the π in finite horizon $T < \infty$. Let us introduce a target path measure $\mathbb{P}^\star$, which is an extension of target π in (1) into path space Ω:

$$\mathrm{d}\mathbb{P}^\star(\mathbf{X}) \propto e^{-U(\mathbf{X}_T) - \log q_T(\mathbf{x}_0, \mathbf{X}_T)} \mathrm{d}\mathbb{P}(\mathbf{X}), \tag{6}$$

where $\mathbb{P}$ is path measure induced by the process in (2). We consider that a variational path measures $\mathbb{P}^\alpha$ induced by following infinite-dimensional additive *controlled* SDEs [14]:

$$\mathrm{d}\mathbf{X}_t^\alpha = \left[\mathcal{A}\mathbf{X}_t^\alpha + \sigma_t Q^{1/2}\alpha_t\right]\mathrm{d}t + \sigma_t \mathrm{d}\mathbf{W}_t^Q, \quad \mathbf{X}_0^\alpha = \mathbf{x}_0, \quad t \in [0, T], \tag{7}$$

where $\alpha_{(\cdot)} : [0, T] \times \mathcal{H} \to \mathcal{U}$ is infinite dimensional Markov control. Then, by invoking generalized Girsanov theorem [8, Theorem 10.14], we can represent our NES problem for (1) into the following infinite-dimensional SOC problem in $\mathcal{H}$ [15, 23][2]:

$$\min_\alpha D_{\mathrm{KL}}(\mathbb{P}^\alpha | \mathbb{P}^\star) = \mathbb{E}_{\mathbb{P}^\alpha}\left[\int_0^T \tfrac{1}{2}\left\|\alpha_t\right\|_{\mathcal{H}}^2 \mathrm{d}t + U(\mathbf{X}_T^\alpha) + \log q_T(\mathbf{x}_0, \mathbf{X}_T^\alpha)\right] \text{ s.t. } (7), \tag{8}$$

Applying the optimal control $\alpha^\star$ to the controlled SDE in (7) yields the desired distribution [20, Theorem 3.2]. In other word, it is possible to achieve exact sampling from π in (1), namely $\mathbf{X}_T^{\alpha^\star} \sim \pi$.

[2]Full proofs and necessary derivations are provided in Appendix B.

2.1 Infinite Dimensional Adjoint Matching

The SOC problem in (8) follows the *least-action* principle in path space, where the goal is to find the *global* optimal control minimizes the terminal cost $g(\mathbf{x}) := U(\mathbf{x}) + \log q_T(\mathbf{x}_0, \mathbf{x})$ while minimizing the running costs. In practice, solving this global problem by direct path-wise minimization is expensive because every update requires gradients through the full trajectory [16]. A practical workaround is to move to the dual formulation. *Adjoint Matching* [AM; 11] offers a scalable route by reformulating the global SOC problem into the local minimization problem.

Extending AM to infinite dimensions requires a rigorous definition of the *first-order adjoint process* that appears in the corresponding infinite-dimensional SOC formulation. Below, we will demonstrate that a specific mathematical concept in SOC called *Stochastic Maximum Principle* [SMP; 1] provides existence and uniqueness of the adjoint process via coupled *Forward-Backward SDEs* [FBSDEs; 12] system and an necessary condition for the optimal control.

Lemma 3 (Stochastic Maximum Principle [12]). *Consider the infinite-dimensional SOC problem:*

$$\min_{\alpha\in\mathbb{A}} \mathbb{E}_{\mathbb{P}^\alpha}\left[\int_0^T l(t, \mathbf{X}_t^\alpha, \alpha_t)\mathrm{d}s + h(\mathbf{X}_T^\alpha)\right], \tag{9}$$

$$\textit{s.t. } \mathrm{d}\mathbf{X}_t^\alpha = \left[\mathcal{A}\mathbf{X}_t^\alpha + f(t, \mathbf{X}_t^\alpha, \alpha_t)\right]\mathrm{d}t + g(t, \mathbf{X}_t^\alpha)\mathrm{d}\mathbf{W}_t^Q, \quad \mathbf{X}_0^\alpha = \mathbf{x}_0. \tag{10}$$

Now, let Assumption B.1-B.2 holds and define the Hamiltonian $H : [0,T]\times\mathcal{H}\times\mathcal{H}\times\mathcal{H}\times\mathcal{H}\to\mathbb{R}$*:*

$$H(t, \mathbf{X}, \alpha, \mathbf{Y}, \mathbf{Z}) := l(t, \mathbf{X}, \alpha) + \langle \mathbf{Y}, f(t, \mathbf{X}, \alpha)\rangle_{\mathcal{H}} + \langle \mathbf{Z}, g(t, \mathbf{X})\rangle_{\mathcal{H}}. \tag{11}$$

Suppose $\mathbf{X}^\star$ *is the optimally controlled process driven by optimal control* $\alpha^\star$*. Then, the following first-order adjoint infinite-dimensional backward SDEs has a unique (weak) solution:*

$$\mathrm{d}\mathbf{Y}_t^\star = -\left[\mathcal{A}^\dagger\mathbf{Y}_t^\star + D_{\mathbf{x}}H(t, \mathbf{X}_t^\star, \alpha_t^\star, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star)\right]\mathrm{d}t + \mathbf{Z}_t^\star\mathrm{d}\mathbf{W}_t^Q, \quad \mathbf{Y}_T^\star = D_{\mathbf{x}}h(\mathbf{X}_T^\star) \tag{12}$$

Moreover, we get necessary condition of optimal control $\alpha_t^\star = \arg\min_{\alpha\in\mathbb{A}} H(t, \mathbf{X}_t^\star, \alpha, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star)$.

Lemma 3 states that, for the controlled forward SDE in (10) associated with the general SOC problem (9), there exists a first-order adjoint pars $(\mathbf{Y}_t, \mathbf{Z}_t)$ solving a backward SDE in (12). This result provides a *necessary condition* for optimality: any control that violates the backward BSDE in (12) or the point-wise minimisation of Hamiltonian for all $t\in[0,T]$ cannot be optimal.

This applies because the Hamiltonian for our SOC problem (8) is a convex functional in the control, given the state control pair $(\mathbf{X}^{\bar\alpha}, \bar\alpha)$, and is defined by:

$$H(t, \mathbf{X}_t^{\bar\alpha}, \bar\alpha_t, \mathbf{Y}_t^{\bar\alpha}, \mathbf{Z}_t^{\bar\alpha}) := \frac{1}{2}\left\|\bar\alpha_t\right\|_{\mathcal{H}}^2 + \langle \mathbf{Y}_t^{\bar\alpha}, \sigma_t Q^{1/2}\bar\alpha_t\rangle_{\mathcal{H}} + \langle \mathbf{Z}_t^{\bar\alpha}, \sigma_t Q^{1/2}\rangle_{\mathcal{H}}. \tag{13}$$

Given convexity, we can find the optimal control as a critical point of H for any $t\in[0,T]$:

$$\nabla_\alpha H(t, \mathbf{X}_t^{\bar\alpha}, \alpha_t, \mathbf{Y}_t^{\bar\alpha}, \mathbf{Z}_t^{\bar\alpha}) = \alpha_t + \langle \mathbf{Y}_t^{\bar\alpha}, \sigma_t Q^{1/2}\rangle_{\mathcal{H}} = \nabla_\alpha \frac{1}{2}\left\|\alpha_t + \sigma_t Q^{1/2}\mathbf{Y}_t^{\bar\alpha}\right\|_{\mathcal{H}}^2. \tag{14}$$

Thus, the local Hamiltonian minimization can be viewed as a *matching* problem that is computationally convenient. Indeed, the Hamiltonian is the Fenchel conjugate of the original SOC problem. Hence, the primal SOC problem enforces path-wise optimality, while the Hamiltonian provides a point-wise optimality, allowing us to compute $\alpha_t^\star$ locally and avoid costly gradient propagation through entire trajectories. Now, the critical point in (14) is found while the state variables $(\mathbf{X}_t^{\bar\alpha}, \mathbf{Y}_t^{\bar\alpha}, \mathbf{Z}_t^{\bar\alpha})$ remain fixed for the chosen control $\bar\alpha$. Leveraging this principle, we therefore introduce a dual objective:

Proposition 4 (Adjoint Matching in $\mathcal{H}$). *Consider the infinite-dimensional matching objective:*

$$\mathcal{L}(\theta) = \int_0^T \mathbb{E}_{\mathbb{P}^{\bar\alpha}}\left[\frac{1}{2}\left\|\alpha^\theta(\mathbf{X}_t^{\bar\alpha}, t) + \sigma_t Q^{1/2}\mathbf{Y}_t^{\bar\alpha}\right\|_{\mathcal{H}}^2\right]\mathrm{d}t, \quad \bar\alpha := \textit{stopgrad}(\alpha^\theta) \tag{15}$$

$$\mathrm{d}\mathbf{Y}_t^{\bar\alpha} = -\mathcal{A}^\dagger\mathbf{Y}_t^{\bar\alpha}\mathrm{d}t + \mathbf{Z}_t^{\bar\alpha}\mathrm{d}\mathbf{W}_t^Q, \quad \mathbf{Y}_T^{\bar\alpha} = D_{\mathbf{x}}g(\mathbf{X}_T^{\bar\alpha}). \tag{16}$$

Then, we get the critical point of $\mathcal{L}$ *is optimal control* $\alpha^\star$ *for the SOC problem in (8).*

Note that when $\mathcal{H} := \mathbb{R}^d$ with $Q = \mathbf{I}_d$, (15) recovers the lean AM objective in [11]. In other words, our formulation is a generalization of AM into the arbitrary Hilbert spaces.

Yet, for $\mathcal{L}$ in (15), the stochastic term $\mathbf{Z}_t^{\bar\alpha}$ can incur a computational burden because this term requires simulating the BSDE in (16) to obtain the solution $\mathbf{Y}_t^{\bar\alpha}$, which markedly increases the cost of each optimization step. Since this BSDE is adapted to the filtration generated by the same Q-Wiener process $\mathbf{W}_t^Q$, we substitute the adjoint process $\mathbf{Y}_t^{\bar\alpha}$ with its conditional expectation $\mathbb{E}_{\mathbb{P}^{\bar\alpha}}[\mathbf{Y}_t^{\bar\alpha}|\mathbf{X}_t^{\bar\alpha}]$.

Proposition 5 (Unbiased Estimator)**.** *Let us define the sample-wise adjoint matching objective with conditional expectation* $\tilde{\mathbf{Y}}_t^{\bar{\alpha}} := \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}[\mathbf{Y}_t^{\bar{\alpha}} | \mathbf{X}_t^{\bar{\alpha}}]$ *for any sample trajectory* $\mathbf{X}_t^{\bar{\alpha}} \sim \mathbb{P}^{\bar{\alpha}}$*:*

$$\tilde{\mathcal{L}}(\theta) := \int_0^T \tfrac{1}{2} ||\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2} \tilde{\mathbf{Y}}_t^{\bar{\alpha}}||_{\mathcal{H}}^2 \mathrm{d}t \tag{17}$$

Then, we get unbiased gradient estimator $\frac{\mathrm{d}}{\mathrm{d}\theta} \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}[\tilde{\mathcal{L}}(\theta)] = \frac{\mathrm{d}}{\mathrm{d}\theta} \mathcal{L}(\theta)$ *with same critical point.*

In practice, we can estimate $\tilde{\mathbf{Y}}_t^{\bar{\alpha}}$ from the trajectory saved during simulation of $\mathbf{X}_T^{\bar{\alpha}}$ for efficient computation. This preserves the same critical point $\alpha^\star$ for the SOC problem in (8) while reducing the computational burden because the substitution yields a closed form solution of the BSDE:

$$\tilde{\mathbf{Y}}_t^{\bar{\alpha}} = \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}[\mathbf{Y}_t^{\bar{\alpha}} | \mathbf{X}_t^{\bar{\alpha}}] = e^{-(T-t)\mathcal{A}^\dagger} \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}[D_{\mathbf{x}} g(\mathbf{X}_T^{\bar{\alpha}}) | \mathbf{X}_t^{\bar{\alpha}}] \approx e^{-(T-t)\mathcal{A}^\dagger} D_{\mathbf{x}} g(\mathbf{X}_T^{\bar{\alpha}}). \tag{18}$$

Finally, since the optimal path measure $\mathbb{P}^\star$ in (1) admits the endpoint disintegration $\mathbb{P}^\star(\mathbf{X}_t) = \mathbb{P}^\star(\mathbf{X}_T)\mathbb{P}_{t|T}(\mathbf{X}_t \mid \mathbf{X}_T)$[3]. Substituting these relations into (15) yields our proposed training objective for the efficient and scalable sampling in function spaces:

$$\mathcal{L}_{\texttt{FAS}}(\theta) = \int_0^T \mathbb{E}_{\mathbb{P}_{t|T}\mathbb{P}_T^{\bar{\alpha}}} \left[\frac{1}{2} \left\| \alpha^\theta(\mathbf{X}_t, t) + \sigma_t Q^{1/2} e^{-(T-t)\mathcal{A}^\dagger} D_{\mathbf{x}} g(\mathbf{X}_T^{\bar{\alpha}}) \right\|_{\mathcal{H}}^2 \right] \mathrm{d}t. \tag{19}$$

Note that the current control $\bar{\alpha}$ may not satisfy this disintegration. We can project the path measure onto a conditional path measure that does satisfy the disintegration [16].

3 Experiments

We evaluate FAS on the Müller-Brown (MB) potential, a standard benchmark for transition path sampling between metastable states. Along three local minima, our goal is to sample paths that connect the top left state and the bottom right state. We report three metrics. Target hit percentage (THP) measures how often final positions reach the target metastable state. We also track the maximum energy along each sampled path, and the log-likelihood of each path. A strong method produces more probable paths (*i.e.*, high log-likelihood), while accurate transition states (*i.e.*, low maximum energy) and reaches the target reliably (*i.e.*, high THP)[4].

Table 1: Performance on MB potential.

Methods	THP (↑)	Max-Energy (↓)	Log-Likelihood (↑)
MCMC‡	-	-13.77±16.43	-
DL [13]	82%	-14.93±12.61	10.42±0.42
FAS **(Ours)**	**100%**	**-36.70**±3.09	**11.48**±0.05

‡ result from [13].

Figure 1: Sampled paths on energy landscape.

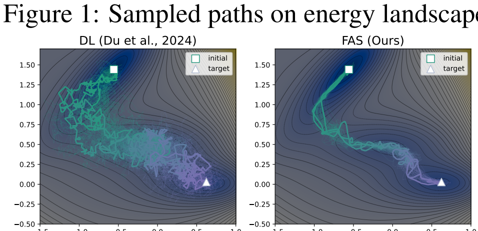

As shown in Table 1, FAS outperforms the baselines and produces the most reliable paths. In Figure 1, the sampled paths transition between the two metastable states along low-energy regions. Unlike finite-dimensional methods, DL, our infinite-dimensional formulation enables enforcement of fixed boundary conditions through an appropriate choice of Hilbert spaces. This yields paths that satisfy the required boundary constraints and keeping high values of THP.

4 Conclusion

We introduced the Function-space Adjoint Sampler (FAS), a scalable diffusion sampler for Gibbs-type distributions on separable Hilbert spaces. FAS generalizes adjoint matching algorithms [11, 16] to infinite-dimensional function spaces, enabling efficient sampling over functional representations such as paths evaluated on a one-dimensional time grid, and it demonstrates superior performance on transition path sampling tasks. A key limitation is the lack of validation on real-world molecular TPS and across a broader spectrum of functional representations.

[3] See Appendix B.8 for details

[4] Further experimental details are provided in Appendix D.

References

[1] Alain Bensoussan. Stochastic maximum principle for distributed parameter systems. *Journal of the Franklin Institute*, 315(5-6):387–406, 1983.

[2] Julius Berner, Lorenz Richter, and Karen Ullrich. An optimal control perspective on diffusion-based generative modeling. *Transactions on Machine Learning Research*, 2024.

[3] Charles-Edouard Bréhier. Lecture notes: Invariant distributions for parabolic spdes and their numerical approximations. 2017.

[4] René Carmona. *Lectures on BSDEs, stochastic control, and stochastic differential games with financial applications*. SIAM, 2016.

[5] Jaemoo Choi, Yongxin Chen, Molei Tao, and Guan-Horng Liu. Non-equilibrium annealed adjoint sampler. *arXiv preprint arXiv:2506.18165*, 2025.

[6] G. Da Prato and J. Zabczyk. Differentiability of the feynman-kac semigroup and a control application. *Atti della Accademia Nazionale dei Lincei. Classe di Scienze Fisiche, Matematiche e Naturali. Rendiconti Lincei. Matematica e Applicazioni*, 8(3):183–188, 10 1997.

[7] G. Da Prato and J. Zabczyk. *Second Order Partial Differential Equations in Hilbert Spaces*. London Mathematical Society Lecture Note Series. Cambridge University Press, 2002.

[8] G. Da Prato and J. Zabczyk. *Stochastic Equations in Infinite Dimensions*. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 2014.

[9] G. Da Prato, J. Zabczyk, P.J. Zabczyk, J.W.S. Cassels, and S.P.G.N.J. Hitchin. *Ergodicity for Infinite Dimensional Systems*. Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press, 1996.

[10] Christoph Dellago, Peter G Bolhuis, Félix S Csajka, and David Chandler. Transition path sampling and the calculation of rate constants. *The Journal of chemical physics*, 108(5):1964–1977, 1998.

[11] Carles Domingo-Enrich, Michal Drozdzal, Brian Karrer, and Ricky T. Q. Chen. Adjoint matching: Fine-tuning flow and diffusion generative models with memoryless stochastic optimal control. In *The Thirteenth International Conference on Learning Representations*, 2025.

[12] Kai Du and Qingxin Meng. A maximum principle for optimal control of stochastic evolution equations. *SIAM Journal on Control and Optimization*, 51(6):4343–4362, 2013.

[13] Yuanqi Du, Michael Plainer, Rob Brekelmans, Chenru Duan, Frank Noe, Carla P Gomes, Alan Aspuru-Guzik, and Kirill Neklyudov. Doob's lagrangian: A sample-efficient variational approach to transition path sampling. *Advances in Neural Information Processing Systems*, 37:65791–65822, 2024.

[14] Giorgio Fabbri, Fausto Gozzi, and Andrzej Swiech. Stochastic optimal control in infinite dimension. *Probability and Stochastic Modelling. Springer*, 2017.

[15] Marco Fuhrman. A class of stochastic optimal control problems in hilbert spaces: Bsdes and optimal control laws, state constraints, conditioned processes. *Stochastic processes and their applications*, 108(2):263–298, 2003.

[16] Aaron Havens, Benjamin Kurt Miller, Bing Yan, Carles Domingo-Enrich, Anuroop Sriram, Brandon Wood, Daniel Levine, Bin Hu, Brandon Amos, Brian Karrer, Xiang Fu, Guan-Horng Liu, and Ricky T. Q. Chen. Adjoint Sampling: Highly scalable diffusion samplers via Adjoint Matching. *International Conference on Machine Learning*, 2025.

[17] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35:26565–26577, 2022.

[18] Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Burigede liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.

[19] Guan-Horng Liu, Jaemoo Choi, Yongxin Chen, Benjamin Kurt Miller, and Ricky TQ Chen. Adjoint schr\" odinger bridge sampler. *arXiv preprint arXiv:2506.22565*, 2025.

[20] Byoungwoo Park, Jungwon Choi, Sungbin Lim, and Juho Lee. Stochastic optimal control for diffusion bridges in function spaces. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024.

[21] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.

[22] Taiji Suzuki. Generalization bound of globally optimal non-convex neural network training: Transportation map estimation by infinite dimensional langevin dynamics. *Advances in Neural Information Processing Systems*, 33:19224–19237, 2020.

[23] Evangelos A. Theodorou, George I. Boutselis, and Kaivalya Bakshi. Linearly solvable stochastic optimal control for infinite-dimensional systems. In *2018 IEEE Conference on Decision and Control (CDC)*, pages 4110–4116, 2018.

[24] Francisco Vargas, Andrius Ovsianas, David Fernandes, Mark Girolami, Neil D Lawrence, and Nikolas Nüsken. Bayesian learning via neural Schrödinger–Föllmer flows. *Statistics and Computing*, 33(1):3, 2023.

[25] Francisco Vargas, Shreyas Padhy, Denis Blessing, and Nikolas Nüsken. Transport meets variational inference: Controlled monte carlo diffusions. In *The Twelfth International Conference on Learning Representations*, 2024.

[26] Qinsheng Zhang and Yongxin Chen. Path integral sampler: A stochastic control approach for sampling. In *International Conference on Learning Representations*, 2022.

A Overview of Relevant Concepts

We first give a concise overview of core probabilistic notions in infinite dimensional spaces, then describe stochastic differential equations on Hilbert spaces [7, 8].

A.1 Gaussian Measure and Generalized Wiener Processes.

Gaussian measures are central because the infinite dimensional setting lacks a Lebesgue reference measure, and a Gaussian with a trace class covariance provides a natural baseline that keeps stochastic integrals square integrable and provides explicit RND, both essential for our framework.

Gaussian Measure. Let $(\Omega, \mathcal{F}, \mathbb{Q})$ be a probability space and $(\mathcal{H}, B(\mathcal{H}))$ be a measurable state spaces where $\mathcal{H}$ be a separable Hilbert space with inner product $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and norm $\|\cdot\|_{\mathcal{H}}$ and $\mathcal{B}(\mathcal{H})$ is Borel measure on $\mathcal{H}$. We consider the $\mathcal{H}$-valued random variable $\mathbf{X} : \Omega \to \mathcal{H}$. Now, define the push-forward measure $\nu := \mathbf{X}_{\#}\mathbb{Q}$ by $\nu(\mathbf{B}) = \mathbb{Q}\big(\mathbf{X}^{-1}(\mathbf{B})\big)$ for every Borel set $\mathbf{B} \subset B(\mathcal{H})$. We call ν Gaussian if every $\mathbb{R}$-valued projection $\langle u, \mathbf{X} \rangle_{\mathcal{H}}$ is a Gaussian random variable for all $u \in \mathcal{H}$. When $\mathcal{H} = \mathbb{R}^d$ with the standard inner product, this definition reproduces the usual Gaussian distribution.

For a Gaussian measure ν, there exists a unique mean function $\mathbf{m} \in \mathcal{H}$ that statisfies:

$$\mathbf{m} = \int_{\mathcal{H}} \mathbf{X} d\nu, \quad \langle \mathbf{m}, u \rangle_{\mathcal{H}} = \mathbb{E}_\nu \left[\langle \mathbf{X}, u \rangle_{\mathcal{H}} \right], \quad \forall u \in \mathcal{H}. \tag{20}$$

Moreover, there exists a unique bounded linear operator $Q : \mathcal{H} \to \mathcal{H}$ such that

$$Qu = \int_{\mathcal{H}} \langle \mathbf{X}, u \rangle_{\mathcal{H}} \mathbf{X} d\nu - \langle \mathbf{m}, u \rangle_{\mathcal{H}} \mathbf{m}, \quad \forall u \in \mathcal{H}. \tag{21}$$

A Gaussian measure is fully specified by its mean and covariance operator. If the mean is zero, so that $\nu = \mathcal{N}(0, Q)$, we say that ν is *centred*. The Fernique's integrability theorem [8, Theorem 2.7] states that for any centred Gaussian measure ν on a separable Banach space $\mathcal{B}$ (in particular on $\mathcal{H}$) there exists $r > 0$ such that $\int_{\mathcal{H}} \exp(r \|\mathbf{X}\|_{\mathcal{H}}^2) d\nu < \infty$. Consequently, it implies that every polynomial moment is finite *i.e.*, $\mathbb{E}_\nu[\|\mathbf{X}\|_{\mathcal{H}}^2] < \infty$. Using definition (21) and an arbitrary orthonormal basis of $\mathcal{H}$, $\{\phi^{(k)}\}_{k \geq 1}$ for each $\phi^{(k)} \in \mathcal{H}$ for every $k \in \mathbb{N}$:

$$\mathrm{Tr}(Q) = \sum_{k \geq 1} \langle Q\phi^{(k)}, \phi^{(k)} \rangle_{\mathcal{H}} = \sum_{k \geq 1} \mathbb{E}_\nu \left[\langle \mathbf{X}, \phi^{(k)} \rangle_{\mathcal{H}}^2 \right] = \mathbb{E}_\nu \left[\|\mathbf{X}\|_{\mathcal{H}}^2 \right] < \infty. \tag{22}$$

This implies that the symmetric and non-negative covariance operator Q is trace class, *i.e.*, $\mathrm{Tr}(Q) < \infty$, hence implies that Q is compact. Hence, there exists the eigen-system $\{(\lambda^{(k)}, \phi^{(k)}) \in \mathbb{R} \times \mathcal{H} : k \in \mathbb{N}\}$, where $\{\phi^{(k)}\}_{k \geq 1}$ is orthonormal eigen-basis and $\{\lambda^{(k)}\}_{k \geq 1}$ is positive eigenvalues such that

$$Q\phi^{(k)} = \lambda^{(k)} \phi^{(k)}, \quad \text{and} \quad \mathrm{Tr}(Q) = \sum_{k \geq 1} \langle Q\phi^{(k)}, \phi^{(k)} \rangle_{\mathcal{H}} = \sum_{k \geq 1} \lambda^{(k)} < \infty. \tag{23}$$

We now extend the notion of a Wiener process to an infinite dimensional SDE. In finite dimensions a standard $\mathbb{R}^d$ Wiener process has independent increments $W_{t+\Delta_t} - W_t \sim \mathcal{N}(0, \Delta_t \mathbf{I}_d)$. When d tends to infinity the covariance $\Delta_t \mathbf{I}_d$ blows up, so the operator is not trace class, no corresponding centred Gaussian measure exists in $\mathcal{H}$, and such increments cannot be defined directly.

Wiener Processes on Hilbert Spaces. A natural way to generalize the Wiener process to a Hilbert space $\mathcal{H}$ is through the *cylindrical Wiener process*. It is introduced as the formal series of independent $\mathbb{R}$-valued Wiener processes $\{W^{(k)}\}_{k \geq 1}$ along an orthonormal basis $\{\phi^{(k)}\}_{k \geq 1}$:

$$\mathbf{W}_t = \sum_{k \geq 1} \mathbf{W}_t^{(k)} \phi^{(k)}, \quad t \geq 0. \tag{24}$$

However, the series $\mathbf{W}_t$ dose not converge in $\mathcal{H}$ because

$$\mathbb{E} \left[\|\mathbf{W}_t\|_{\mathcal{H}}^2 \right] = \sum_{k \geq 1} \mathbb{E} \left[(\mathbf{W}_t^{(k)})^2 \right] = t \sum_{k \geq 1} 1 = \infty, \tag{25}$$

thereby $\mathbf{W}_t$ fails to be an $\mathcal{H}$-valued random variable. Hence, it can be used only under additional smoothing to obtain $\mathcal{H}$-valued noise process. This smoothing is done with trace class operator Q. Define $\mathbf{W}_t^Q = Q^{1/2}\mathbf{W}_t = \sum_{k\geq 1}\sqrt{\lambda^{(k)}}\mathbf{W}_t^{(k)}\phi^{(k)}$. Then, we have

$$\mathbb{E}\left[\left\|\mathbf{W}_t^Q\right\|_{\mathcal{H}}^2\right] = \sum_{k\geq 1}\mathbb{E}\left[\lambda^{(k)}(\mathbf{W}_t^{(k)})^2\right] = t\sum_{k\geq 1}\lambda^{(k)} < \infty, \tag{26}$$

where $\{(\lambda^{(k)}, \phi^{(k)}) \in \mathbb{R}\times\mathcal{H} : k \in \mathbb{N}\}$ is eigen-system of Q. Hence $\mathbf{W}_t^Q$ is $\mathcal{H}$-valued process. This smoothed process is called a *Q-Wiener process* or *coloured noise* because the noise covariance operator Q introduces a spatial correlation.

A.2 Stochastic Differential Equations in Hilbert Spaces

With the infinite dimensional Wiener process $\mathbf{W}_t^Q$ well defined $\mathcal{H}$, we then define the infinite-dimensional SDEs as follows:

$$d\mathbf{X}_t = \mathcal{A}\mathbf{X}_t \mathrm{d}t + \sigma_t \mathrm{d}\mathbf{W}_t^Q, \quad \mathbf{X}_0 = \mathbf{x}_0 \in \mathcal{H}, \tag{27}$$

where $\mathcal{A} : \mathcal{H} \to \mathcal{H}$ is a linear operator and $\sigma_t > 0$. We recall the classical mild formulation driven by a general Q-Wiener process on $\mathcal{H}$. We then restrict attention to a cylindrical Wiener process on $\mathcal{H}_0$, introduce an appropriate Hilbert–Schmidt embedding that can be naturally aligned with our framework, and outline the minimal assumptions that still guarantee existence of a solution for the infinite-dimensional SDE under the choice of cylindrical Wiener process.

Mild solution. Let a linear operator $\mathcal{A} : \mathcal{H} \to \mathcal{H}$ generate a strongly continuous C_0 semi-group $\{e^{t\mathcal{A}}\}_{t\geq 0}$ on $\mathcal{H}$. Assume $\mathbf{W}^Q$ is a centred Q-Wiener process with trace-class covariance operator Q and that $e^{t\mathcal{A}}Q^{1/2}$ is Hilbert-Schmidt for any $t > 0$. Under these conditions, the infinite-dimensional SDE (or stochastic PDE) in (2) admit the unique mild solution:

$$\mathbf{X}_t = e^{t\mathcal{A}}\mathbf{x}_0 + \sigma_t\int_0^t e^{(t-s)\mathcal{A}} d\mathbf{W}_s^Q, \quad t \geq 0. \tag{28}$$

The resulting stochastic process $\{\mathbf{X}_t\}_{t\geq 0}$ is then $\mathcal{H}$-valued random variable since the stochastic convolution $\int_0^t e^{(t-s)\mathcal{A}} d\mathbf{W}_s^Q = \int_0^t e^{(t-s)\mathcal{A}}Q^{1/2}d\mathbf{W}_s$ is lies on $\mathcal{H}$ because $e^{(t-s)\mathcal{A}}Q^{1/2}$ is Hilbert-Schmidt. We can compute its mean $\mathbf{m}_t$ and covariance operator Q_t for any $t > 0$ as follows:

$$\mathbf{m}_t = e^{t\mathcal{A}}\mathbf{x}_0, \quad Q_t = \sigma_t^2\int_0^t e^{(t-s)\mathcal{A}}Qe^{(t-s)\mathcal{A}^\dagger}ds. \tag{29}$$

Due to the Hille-Yosida Theorem [8, Theorem A.3], $\mathcal{A}$ satisfies $\left\|e^{t\mathcal{A}}\right\| \leq Me^{-wt}$ for some $M \geq 1$, $w \geq 0$. Hence, for every $0 \leq s \leq t$, $T_{st} := e^{(t-s)\mathcal{A}}Qe^{(t-s)\mathcal{A}^\dagger}$ satisfies:

$$\mathrm{Tr}(T_{st}) \leq \left\|e^{(t-s)\mathcal{A}}\right\|\mathrm{Tr}(Q)\left\|e^{(t-s)\mathcal{A}^\dagger}\right\| \leq M^2e^{-2w(t-s)}\mathrm{Tr}(Q) < \infty. \tag{30}$$

Implies that T_{st} is a trace class operator. Therefore, since the bound is integrable on $[0, t]$,

$$\mathrm{Tr}(Q_t) = \mathrm{Tr}(\sigma_t^2\int_0^t T_{st}ds) \leq \sigma_t^2\int_0^t \mathrm{Tr}(T_{st})ds < \infty. \tag{31}$$

Hence Q_t is a trace-class operator for every $t < \infty$. Hence, we can conform that the $\mathcal{H}$-valued stochastic process $\mathbf{X}_t$ has a Gaussian law $\nu_t = \mathcal{N}(\mathbf{m}_t, Q_t)$.

Invariant Measure. Let $\mathcal{A}$ is dissipative (*i.e.*, $\langle\mathcal{A}\mathbf{x}, \mathbf{x}\rangle_{\mathcal{H}} \leq -\lambda\left\|\mathbf{x}\right\|_{\mathcal{H}}^2$ for some $\lambda > 0$), the integral converges as $t \to \infty$ and we define the noise schedule such that $\sup_t \sigma_t < \infty$, $\sigma_t \to \sigma_\infty$. Then, the operator Q_t solves the following differential Lyapunov equation:

$$\dot{Q}_t = \mathcal{A}Q_t + Q_t\mathcal{A}^\dagger + \sigma_t^2 Q, \quad Q_t = e^{t\mathcal{A}}Q_0e^{t\mathcal{A}^\dagger} + \int_0^t \sigma_{t-r}^2 e^{r\mathcal{A}}Qe^{r\mathcal{A}^\dagger}\mathrm{d}r. \tag{32}$$

Because $Q_0 = 0$ since we assume Dirac delta initial condition $\mathbf{x}_0$ in (2), we obtain

$$Q_t \overset{t\to\infty}{\to} Q_\infty := \sigma_\infty^2\int_0^\infty e^{t\mathcal{A}}Qe^{t\mathcal{A}^\dagger}\mathrm{d}t \tag{33}$$

Now, let $F(t) = \sigma_\infty^2 e^{t\mathcal{A}} Q e^{t\mathcal{A}^\dagger}$. Then, we get:

$$\tfrac{\mathrm{d}}{\mathrm{d}t} F(t) = \mathcal{A} F(t) + F(t)\mathcal{A}^\dagger = \mathcal{A}\sigma_\infty^2 e^{t\mathcal{A}} Q e^{t\mathcal{A}^\dagger} + \sigma_\infty^2 e^{t\mathcal{A}} Q e^{t\mathcal{A}^\dagger} \mathcal{A}^\dagger. \tag{34}$$

Since $F(t) \overset{t\to\infty}{\to} 0$ and $F(0) = \sigma_\infty^2 Q$, by integrating (34) from 0 to ∞, we get:

$$\mathcal{A} Q_\infty + Q_\infty \mathcal{A}^\dagger = \int_0^\infty \tfrac{\mathrm{d}}{\mathrm{d}t} F(t)\mathrm{d}t = \lim_{t\infty} F(t) - F(0) = -\sigma_\infty^2 Q \tag{35}$$

Hence, we get the operator Lyapunov equation [7, Proposition 10.1.4]:

$$\mathcal{A} Q_\infty + Q_\infty \mathcal{A}^\dagger + \sigma_\infty^2 Q = 0, \tag{36}$$

where Q_∞ is trace class operator. so $\mathcal{N}(0, \infty)$ is the unique invariant Gaussian measure of the process in (2). Finally, if $\mathcal{A}$ is trace-class, self-adjoint operator ($i.e.$, $\mathcal{A} = \mathcal{A}^\dagger$) and $\mathcal{A}$ and Q are commute ($i.e.$, $Q\mathcal{A} = \mathcal{A}Q$), we get $Q_\infty = -\frac{1}{2}\sigma_\infty^2 \mathcal{A}^{-1} Q$ [7, Proposition 10.1.6].

B Proofs and Derivations

B.1 Assumptions

Assumption B.1. *We assume a linear operator $\mathcal{A} : \mathcal{H} \to \mathcal{H}$ that generates a C_0-semigroup and is exponentially stable, dissipative i.e., $\langle \mathcal{A}\mathbf{x}, \mathbf{x}\rangle_\mathcal{H} \le -\lambda \|\mathbf{x}\|_\mathcal{H}^2$ for some $\lambda > 0$, trace class i.e., $\mathrm{Tr}(\mathcal{A}) < \infty$, and self-adjoint i.e., $\mathcal{A} = \mathcal{A}^\dagger$. We also take a Q-Wiener process with a self-adjoint trace-class operator Q that commutes with $\mathcal{A}$ i.e., $Q\mathcal{A} = \mathcal{A}Q$.*

Assumption B.2. *For each $(\mathbf{x}, \alpha) \in \mathcal{H} \times \mathcal{U}$, the maps $t \mapsto f(t, \mathbf{x}, \alpha)$, $t \mapsto g(t, \mathbf{x}, \alpha)$, and $t \mapsto l(t, \mathbf{x}, \alpha)$ are predictable. The function $h(\mathbf{x})$ is $\mathcal{F}_T$-measurable. For each $(t, \alpha, \omega) \in [0, T] \times \mathcal{U} \times \Omega$, the functions f, g, l, and h are globally twice Fréchet differentiable with respect to $\mathbf{x}$. The derivatives $D_\mathbf{x} f$, $D_\mathbf{x} g$, $D_{\mathbf{xx}} f$, $D_{\mathbf{xx}} g$, $D_{\mathbf{xx}} l$, and $D_{\mathbf{xx}} h$ are continuous in $\mathbf{x}$ and uniformly bounded by a constant K. Moreover, we assume the growth bounds:*

$$|f(t, \mathbf{x}, \alpha)| + \|g(t, \mathbf{x}, \alpha)\| + \|D_\mathbf{x} l(t, \mathbf{x}, \alpha)\| + \|D_\mathbf{x} h(\mathbf{x})\| \le K(1 + \|\mathbf{x}\|_\mathcal{H} + |\alpha|_\mathcal{U}), \tag{37}$$

$$|l(t, \mathbf{x}, \alpha)| + |h(\mathbf{x})| \le K(1 + \|\mathbf{x}\|_\mathcal{H} + |\alpha|_\mathcal{U}^2). \tag{38}$$

Assumption B.3. *The function $\mathcal{V} : [0, T] \times \mathcal{H} \to \mathbb{R}$ and its derivatives $D_\mathbf{x}\mathcal{V}, D_{\mathbf{xx}}\mathcal{V}, \partial_t \mathcal{V}$ are uniformly continuous on bounded subsets of $[0, T] \times \mathcal{H}$ and $(0, T) \times \mathcal{H}$, respectively. Moreover, for all $(t, \mathbf{x}) \in (0, T) \times \mathcal{H}$, there exists $C_1, C_2 > 0$ such that*

$$|\mathcal{V}(t, \mathbf{x})| + |D_\mathbf{x}\mathcal{V}(t, \mathbf{x})| + |\partial_t \mathcal{V}(t, \mathbf{x})| + \|D_{\mathbf{xx}}\mathcal{V}(t, \mathbf{x})\| + |\mathcal{A}^\dagger D_\mathbf{x}\mathcal{V}(t, \mathbf{x})| \le C_1(1 + |\mathbf{x}|)^{C_2}, \tag{39}$$

where $\mathcal{A}^\dagger$ is adjoint operator of $\mathcal{A}$.

B.2 Preliminary

We recall two key ingredients. The first is the product rule for Hilbert space semimartingales, also called stochastic integration by parts. We use this rule to derive the adjoint matching objective. The second is a verification theorem that we use to establish the optimality.

Lemma B.4 (Stochastic integration by parts in $\mathcal{H}$)**.** *Assume A, C are $\mathcal{H}$–valued predictable processes with $\int_0^T \mathbb{E}\|A_s\|^2 \,\mathrm{d}s < \infty$ and $\int_0^T \mathbb{E}\|C_s\|^2 \,\mathrm{d}s < \infty$ for every $T > 0$ and B, D are predictable processes with values in the Hilbert–Schmidt space $\mathcal{L}_2(\mathcal{U}_Q, \mathcal{H})$ and $\int_0^T \mathbb{E}\|B_s\|_{\mathcal{L}_2}^2 \,\mathrm{d}s < \infty$, $\int_0^T \mathbb{E}\|D_s\|_{\mathcal{L}_2}^2 \,\mathrm{d}s < \infty$ for every $T > 0$. Then, let us define two stochastic convolutions:*

$$\mathbf{X}_t = \mathbf{X}_0 + \int_0^t A_s ds + \int_0^t B_s d\mathbf{W}_s^Q, \quad \mathbf{Y}_t = \mathbf{Y}_0 + \int_0^t C_s ds + \int_0^t D_s d\mathbf{W}_s^Q \tag{40}$$

Then, for all $t \ge 0$, almost surely, we get:

$$\langle \mathbf{X}_t, \mathbf{Y}_t\rangle_\mathcal{H} = \langle \mathbf{X}_0, \mathbf{Y}_0\rangle_\mathcal{H} + \int_0^t \langle \mathbf{Y}_s, A_s\rangle_\mathcal{H} \mathrm{d}s + \int_0^t \langle \mathbf{X}_s, C_s\rangle_\mathcal{H} \mathrm{d}s \tag{41}$$

$$+ \int_0^t \langle \mathbf{Y}_s, B_s d\mathbf{W}_s^Q\rangle_\mathcal{H} + \int_0^t \langle \mathbf{X}_s, D_s d\mathbf{W}_s^Q\rangle_\mathcal{H} + \int_0^t Tr\left[B_s Q D_s^\dagger\right] ds. \tag{42}$$

Equivalently, in differential form,

$$\mathrm{d}\langle \mathbf{X}_t, \mathbf{Y}_t\rangle_\mathcal{H} = \langle \mathbf{Y}_t, A_t\rangle_\mathcal{H} \mathrm{d}t + \langle \mathbf{X}_t, C_t\rangle_\mathcal{H} \mathrm{d}t + \langle \mathbf{Y}_t, B_t \mathrm{d}\mathbf{W}_t^Q\rangle_\mathcal{H} + \langle \mathbf{X}_t, D_t \mathrm{d}\mathbf{W}_t^Q\rangle_\mathcal{H} + Tr\left[B_t Q D_t^\dagger\right] \mathrm{d}t.$$

Proof. Let us define $\mathbb{R}$-valued process $\mathbf{Z}_t := \langle \mathbf{X}_t, \mathbf{Y}_t \rangle_{\mathcal{H}}$ and C^2 mapping $f : \mathcal{H} \times \mathcal{H} \to \mathbb{R}$ such that $f(\mathbf{x}, \mathbf{y}) = \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{H}}$ for any $\mathbf{x}, \mathbf{y} \in \mathcal{H}$. Since

$$D_{\mathbf{x}} f(\mathbf{x}, \mathbf{y}) = \mathbf{y}, D_{\mathbf{y}} f(\mathbf{x}, \mathbf{y}) = \mathbf{x}, D_{\mathbf{xx}} f(\mathbf{x}, \mathbf{y}) = 0, D_{\mathbf{yy}} f(\mathbf{x}, \mathbf{y}) = 0, \tag{43}$$

and $D_{\mathbf{xy}} f(\mathbf{x}, \mathbf{y})$ is the continuous bilinear form $(h, k) \to \langle h, k \rangle_{\mathcal{H}}$. Now, write the semi-martingale decomposition on $\mathcal{H} \times \mathcal{H}$:

$$(\mathbf{X}_t, \mathbf{Y}_t) = (\mathbf{X}_0, \mathbf{Y}_0) + \int_0^t (A_s, B_s) ds + \int_0^t \tilde{B}_s d\mathbf{W}_s^Q, \tag{44}$$

where $\tilde{B}_s : \mathcal{U} \to \mathcal{H} \times \mathcal{H}$ such that $\tilde{B}_s u = (B_s u, D_s u)$. By applying the infinite-dimensional Itô formula [8, Chapter 4.4] for f on $\mathcal{H} \times \mathcal{H}$ gives:

$$d\mathbf{Z}_t = \langle D_{\mathbf{x}} f(\mathbf{X}_t, \mathbf{Y}_t), d\mathbf{X}_t \rangle_{\mathcal{H}} + \langle D_{\mathbf{y}} f(\mathbf{X}_t, \mathbf{Y}_t), d\mathbf{Y}_t \rangle_{\mathcal{H}} + \tfrac{1}{2} Tr \left[D_{\mathbf{xy}} f(\mathbf{X}_t, \mathbf{Y}_t)(\tilde{B}_t Q \tilde{B}_t^\dagger) \right] \mathrm{d}t, \tag{45}$$

where we compute the trace by expanding the function in any orthonormal basis $\{\phi^{(k)}\}_{k \geq 1}$ of $\mathcal{U}$:

$$\tfrac{1}{2} Tr \left[D_{\mathbf{xy}} f(\mathbf{X}_t, \mathbf{Y}_t)(\tilde{B}_t Q \tilde{B}_t^\dagger) \right] = \sum_{k \geq 1} \lambda^{(k)} \left(\langle B_t, \phi^{(k)} \rangle_{\mathcal{H}} + \langle D_t, \phi^{(k)} \rangle_{\mathcal{H}} \right) = Tr \left[B_t Q D_t^\dagger \right], \tag{46}$$

where $\lambda^{(k)}$ is eigen-value corresponding eigen-basis $\phi^{(k)}$ such that $Q\phi^{(k)} = \lambda^{(k)} \phi^{(k)}$. Therefore,

$$d\langle \mathbf{X}_t, \mathbf{Y}_t \rangle_{\mathcal{H}} = \langle \mathbf{Y}_t, d\mathbf{X}_t \rangle_{\mathcal{H}} + \langle \mathbf{X}_t, d\mathbf{Y}_t \rangle_{\mathcal{H}} + Tr \left[B_t Q D_t^\dagger \right] \mathrm{d}t. \tag{47}$$

By plugging $d\mathbf{X}_t = A_t \mathrm{d}t + B_t d\mathbf{W}_t^Q$ and $d\mathbf{Y}_t = C_t \mathrm{d}t + D_t d\mathbf{W}_t^Q$, we get the desired results. □

Lemma B.5 (Verification Theorem). *Let $\mathcal{V}$ be a solution of Hamilton-Jacobi-Bellman (HJB) equation:*

$$\partial_t \mathcal{V}_t + \mathcal{A} \mathcal{V}_t + \inf_{\alpha \in \mathbb{A}} \left[\langle \alpha, \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}_t \rangle + \tfrac{1}{2} \|\alpha\|_{\mathcal{H}}^2 \right] = 0, \quad \mathcal{V}(T, \mathbf{x}) = g(\mathbf{x}) \tag{48}$$

which satisfying the Assumption B.3. Then, we have $\mathcal{V}(0, \mathbf{x}_0) \leq \mathcal{J}(\alpha, \mathbb{P}^\alpha)$ for every $\alpha \in \mathbb{A}$. Now, let $(\alpha^*, \mathbf{X}^\star)$ be an admissible pair such that

$$\alpha_t^*(t, \mathbf{X}_t^\star) = \arg\inf_{\alpha \in \mathbb{A}} \left[\langle \alpha_t, \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}_t \rangle + \frac{1}{2} \|\alpha_t\|_{\mathcal{H}}^2 \right] = -\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\star) \tag{49}$$

for almost every $t \in [0, T]$ and $\mathbb{P}$-almost surely. Then $(\alpha^*, \mathbf{X}^{\alpha^*})$ satisfying $\mathcal{V}(0, \mathbf{x}_0) = \mathcal{J}(\alpha^\star, \mathbb{P}^\star)$.

Proof. To start the proof, we formally minimize $F(D_{\mathbf{x}} \mathcal{V})$. Using results from [6], we obtain the candidate minimizer and proceed accordingly:

$$F(\mathbf{x}) = \inf_{\alpha \in \mathbb{A}} \left[\langle \mathbf{x}, \alpha \rangle + \frac{1}{2} \|\alpha\|_{\mathcal{H}}^2 \right] = -\frac{1}{2} \|\mathbf{x}\|_{\mathcal{H}}^2 . \tag{50}$$

Thus, $F(\sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V})$ is minimized at at $\alpha^* = -\sigma Q^{1/2} D_{\mathbf{x}} \mathcal{V}$. Next, applying Itô's formula [8, Chapter 4.4] to $\mathcal{V}$ and take the expectation on both sides to obtain

$$\mathbb{E}_{\mathbb{P}^\alpha} [\mathcal{V}(T, \mathbf{X}_T^\alpha)] = \mathbb{E}_{\mathbb{P}^\alpha} [g(\mathbf{X}_T^\alpha)] = \tag{51}$$

$$= \mathcal{V}(0, \mathbf{x}_0) + \mathbb{E}_{\mathbb{P}^\alpha} \left[\int_0^T \left(\partial_t \mathcal{V}(t, \mathbf{X}_t^\alpha) + \mathcal{A} \mathcal{V}(t, \mathbf{X}_t^\alpha) + \langle \sigma_t Q^{1/2} D_{\mathbf{x}} \mathcal{V}(t, \mathbf{X}_t^\alpha), \alpha_t \rangle_{\mathcal{H}} \right) \mathrm{d}t \right], \tag{52}$$

Using that $\mathcal{V}$ satisfies (49), add $\mathbb{E}_{\mathbb{P}^\alpha} \left[\int_0^T \frac{1}{2} \|\alpha_t\|_{\mathcal{H}}^2 \mathrm{d}t \right]$ to both sides. Then the LHS of (51) becomes:

$$\mathbb{E}_{\mathbb{P}^\alpha} \left[g(\mathbf{X}_T^\alpha) + \int_0^T \tfrac{1}{2} \|\alpha_t\|_{\mathcal{H}}^2 \mathrm{d}t \right] = \mathcal{J}(\alpha, \mathbb{P}^\alpha). \tag{53}$$

Now for the RHS of the equation (51):

$$\mathcal{V}(0,\mathbf{x}_0)+\mathbb{E}_{\mathbb{P}^\alpha}\left[\int_0^T\left(\tfrac{1}{2}\|\alpha_t\|_{\mathcal{H}}^2+\partial_t\mathcal{V}(t,\mathbf{X}_t^\alpha)+\mathcal{A}\mathcal{V}(t,\mathbf{X}_t^\alpha)+\langle\sigma_t Q^{1/2}D_{\mathbf{x}}\mathcal{V}(t,\mathbf{X}_t^\alpha),\alpha_t\rangle_{\mathcal{H}}\right)\mathrm{d}t\right] \tag{54}$$

$$\overset{(i)}{=}\mathcal{V}(0,\mathbf{x}_0)+\mathbb{E}_{\mathbb{P}^\alpha}\left[\int_0^T\left(\langle\sigma_t Q^{1/2}D_{\mathbf{x}}\mathcal{V}(t,\mathbf{X}_t^\alpha),\alpha_t\rangle_{\mathcal{H}}+\tfrac{1}{2}\|\alpha_t\|_{\mathcal{H}}^2-F(\sigma_t Q^{1/2}D_{\mathbf{x}}\mathcal{V}(t,\mathbf{X}_t^\alpha))\right)\mathrm{d}t\right]. \tag{55}$$

Here (i) holds because we add and subtract $\mathbb{E}_{\mathbb{P}^\alpha}\left[\int_0^T F(D_{\mathbf{x}}\mathcal{V}(t,\mathbf{X}_t^\alpha))\mathrm{d}t\right]$ and use again that $\mathcal{V}$ satisfies (49). Therefore, we obtain the following equation:

$$\mathcal{J}(\alpha,\mathbb{P}^\alpha)=\mathcal{V}(0,\mathbf{x}_0) \tag{56}$$

$$+\mathbb{E}_{\mathbb{P}^\alpha}\left[\int_0^T\left(\langle\sigma_t Q^{1/2}D_{\mathbf{x}}\mathcal{V}(t,\mathbf{X}_t^\alpha),\alpha_t\rangle_{\mathcal{H}}+\tfrac{1}{2}\|\alpha_t\|_{\mathcal{H}}^2-F(\sigma_t Q^{1/2}D_{\mathbf{x}}\mathcal{V}(t,\mathbf{X}_t^\alpha))\right)\mathrm{d}t\right]. \tag{57}$$

By definition, $\left[\langle\sigma_t Q^{1/2}D_{\mathbf{x}}\mathcal{V}(t,\mathbf{X}_t^\alpha),\alpha_t\rangle_{\mathcal{H}}+\tfrac{1}{2}\|\alpha_t\|_{\mathcal{H}}^2\right]-F(D_{\mathbf{x}}\mathcal{V}(t,\mathbf{X}_t^\alpha))\geq 0$. Thus, taking the infimum over $\alpha\in\mathbb{A}$ on the RHS of (56) yields $\mathcal{J}(\alpha,\mathbb{P}^\alpha)\geq\mathcal{V}(0,\mathbf{x}_0)$. We already verified that $F(\sigma_t Q^{1/2}D_{\mathbf{x}}\mathcal{V})$ attains the infimum at $\alpha_t^\star=-\sigma_t Q^{1/2}D_{\mathbf{x}}\mathcal{V}$, so choose $u=\alpha^\star$. Then,

$$\left[\langle\sigma Q^{1/2}D_{\mathbf{x}}\mathcal{V}(s,\mathbf{X}_s^u),u_s\rangle+\frac{1}{2}\|u_s\|^2\right]-F(\sigma Q^{1/2}D_{\mathbf{x}}\mathcal{V}(s,\mathbf{X}_s^u))=0. \tag{58}$$

Thus, we get $\mathcal{J}(u,\mathbb{P}^u)=\mathcal{V}(t,\mathbf{x})$. Combining with (56), we conclude that $(\alpha^\star,\mathbf{X}^{\alpha^\star})$ is optimal for $(t,\mathbf{x})\in[0,T]\times\mathcal{H}$. This completes the proof. □

B.3 Derivation of Divergence Between Path Measures

Here we present an infinite-dimensional generalization of Girsanov theorem [8, Theorem 10.14], which plays a crucial role in estimating the divergence between two path.

Theorem B.6 (Generalized Girsanov)**.** *Let $\{\zeta_t\}_{0\leq t\leq T}$ be a $\mathcal{H}_0$-valued predictable process such that*

$$\mathbb{P}\left(\int_0^T\|\zeta_t\|_{\mathcal{H}_0}^2\,dt<\infty\right)=1,\quad\mathbb{E}\left[\exp\left(\int_0^T\langle\zeta_t,d\mathbf{W}_t^Q\rangle_{\mathcal{H}_0}-\frac{1}{2}\int_0^T\|\zeta_t\|_{\mathcal{H}_0}^2\,dt\right)\right]=1. \tag{59}$$

Then the process $\tilde{\mathbf{W}}_t^Q=\mathbf{W}_t^Q-\int_0^t\zeta_s ds$ is a Q-Wiener process with respect to the filtration on the probability space $(\Omega,\mathcal{F},\mathbb{Q})$ where

$$d\mathbb{Q}=\exp\left(\int_0^T\langle\zeta_t,d\mathbf{W}_t^Q\rangle_{\mathcal{H}_0}-\frac{1}{2}\int_0^T\|\zeta_t\|_{\mathcal{H}_0}^2\,dt\right)d\mathbb{P}. \tag{60}$$

Or alternatively, by substituting $\mathbf{W}_t^Q=\tilde{\mathbf{W}}_t^Q+\int_0^t\zeta_s ds$ to (65), we obtain

$$d\mathbb{Q}=\exp\left(\int_0^T\langle\zeta_t,d\tilde{\mathbf{W}}_t^Q\rangle_{\mathcal{H}_0}+\frac{1}{2}\int_0^T\|\zeta_t\|_{\mathcal{H}_0}^2\,dt\right)d\mathbb{P}. \tag{61}$$

Proof. Proof can be found in [8, Theorem 10.14] □

Let us define $\zeta_t=Q^{1/2}\alpha_t$. For a $\mathbb{P}^\alpha$ Q-Wiener process $\tilde{\mathbf{W}}_t^Q$, $\mathbf{W}_t^Q=\tilde{\mathbf{W}}_t^Q+\int_0^t Q^{1/2}\alpha_s ds$ is a Q-Wiener process on $\mathbb{P}$ because

$$d\mathbf{X}_t=\mathcal{A}\mathbf{X}_t dt+\sigma_t Q^{1/2}\alpha_t dt+\sigma_t d\tilde{\mathbf{W}}_t^Q \tag{62}$$

$$=\mathcal{A}\mathbf{X}_t dt+\sigma_t Q^{1/2}\alpha_t dt+\sigma_t\left[d\mathbf{W}_t^Q-Q^{1/2}\alpha_t dt\right] \tag{63}$$

$$=\mathcal{A}\mathbf{X}_t dt+\sigma_t dW_t^Q \tag{64}$$

Hence, we get

$$d\mathbb{P}^{\alpha} = \exp\left(\int_0^T \langle \zeta_s, d\tilde{\mathbf{W}}_t^Q \rangle_{\mathcal{H}_0} + \frac{1}{2}\int_0^T \|\zeta_s\|_{\mathcal{H}_0}^2 \, dt\right) d\mathbb{P}. \tag{65}$$

Now, by change of measure in (6), we get the following relation:

$$d\mathbb{P}^{\star} = \frac{1}{\tilde{\mathcal{Z}}} e^{-U(\mathbf{X}_T) - \log q_T(\mathbf{X}_T)} d\mathbb{P} \tag{66}$$

$$= \frac{1}{\tilde{\mathcal{Z}}} e^{-U(\mathbf{X}_T^{\alpha}) - \log q_T(\mathbf{X}_T^{\alpha})} \frac{d\mathbb{P}}{d\mathbb{P}^{\alpha}} d\mathbb{P}^{\alpha}, \tag{67}$$

where $\tilde{\mathcal{Z}} = \mathbb{E}\left[\frac{1}{\tilde{\mathcal{Z}}} e^{-U(\mathbf{X}_T) - \log q_T(\mathbf{X}_T)}\right]$ is normalization constant. Combining (65) and (67) and take logarthm on both sides, we get

$$\log \frac{d\mathbb{P}^{\star}}{d\mathbb{P}^{\alpha}} = \log \frac{1}{\tilde{\mathcal{Z}}} e^{-U(\mathbf{X}_T^{\alpha}) - \log q_T(\mathbf{X}_T^{\alpha})} + \log \frac{d\mathbb{P}}{d\mathbb{P}^{\alpha}} \tag{68}$$

$$= -U(\mathbf{X}_T^{\alpha}) - \log q_T(\mathbf{X}_T^{\alpha}) - \int_0^T \langle \zeta_t, d\tilde{\mathbf{W}}_t^Q \rangle_{\mathcal{H}_0} - \frac{1}{2}\int_0^T \left\| Q^{1/2} \alpha_t \right\|_{\mathcal{H}_0}^2 dt - \log \tilde{\mathcal{Z}}. \tag{69}$$

Since $\tilde{\mathbf{W}}_t^Q$ is $\mathbb{P}^{\alpha}$ Q-Wiener process, taking expecatation with respect to $\mathbb{P}^{\alpha}$ on (69) we get

$$\mathbb{E}_{\mathbb{P}^{\alpha}}\left[\frac{d\mathbb{P}^{\alpha}}{d\mathbb{P}^{\star}}\right] = \mathbb{E}_{\mathbb{P}^{\alpha}}\left[\frac{1}{2}\int_0^T \|\alpha_t\|_{\mathcal{H}}^2 \, dt + (U(\mathbf{X}_T^{\alpha}) + \log q_T(\mathbf{X}_T^{\alpha}))\right] + \log \tilde{\mathcal{Z}}. \tag{70}$$

Then, since $\tilde{\mathcal{Z}}$ is constant, we get the desired KL-divergence minimization objective:

$$\min_{\alpha} D_{\mathrm{KL}}(\mathbb{P}^{\alpha} | \mathbb{P}^{\star}) = \min_{\alpha} \mathbb{E}_{\mathbb{P}^{\alpha}}\left[\frac{1}{2}\int_0^T \|\alpha_t\|_{\mathcal{H}}^2 \, dt + (U(\mathbf{X}_T^{\alpha}) + \log q_T(\mathbf{X}_T^{\alpha}))\right]. \tag{71}$$

B.4 Proof of Lemma 1

Lemma 1 (Finite Partition)**.** *Let $U(\mathbf{x}) \geq \beta \|\mathbf{x}\|_{\mathcal{H}}^2 - C$ with $\beta, C > 0$. Then, for any $Tr(Q) < \infty$, choosing $\nu = \mathcal{N}(0, Q)$ in (1) results in a finite partition function $\mathcal{Z} < \infty$.*

Proof. The Fernique's integrability theorem [8, Theorem 2.7] states that for any centred Gaussian measure $\nu = \mathcal{N}(0, Q)$ with any $\mathrm{Tr}(Q) < \infty$ on a separable Banach space $\mathcal{B}$ (in particular on $\mathcal{H}$) there exists $r > 0$ such that $\int_{\mathcal{H}} e^{r\|\mathbf{x}\|_{\mathcal{H}}^2} d\nu(\mathbf{x}) < \infty$. Because the energy satisfies quadratic lower bound $U(\mathbf{x}) \geq \beta \|\mathbf{x}\|_{\mathcal{H}}^2 - C$, we obtain $e^{-U(\mathbf{x})} \leq \tilde{C} e^{-\beta\|\mathbf{x}\|^2}$, where $\tilde{C} = e^C$. Hence, we get:

$$\mathcal{Z} = \int_{\mathcal{H}} e^{-U(\mathbf{x})} d\nu(\mathbf{x}) \leq \tilde{C} \int_{\mathcal{H}} e^{-\beta\|\mathbf{x}\|^2} d\nu(\mathbf{x}) \leq \tilde{C} \int_{\mathcal{H}} e^{r\|\mathbf{x}\|^2} d\nu(\mathbf{x}) < \infty, \tag{72}$$

which confirms that the partition function $\mathcal{Z}$ is finite. □

B.5 Proof of Theorem 2

Theorem 2 (Explicit RND)**.** *Suppose Assumption B.1 holds. Then, For any $t > 0$ and initial condition $\mathbf{x}_0 \in \mathcal{H}$, $\nu_t^{\mathbf{x}_0} = \mathcal{N}(e^{t\mathcal{A}}\mathbf{x}_0, Q_t)$ and $\nu_\infty = \mathcal{N}(0, Q_\infty)$ are mutually absolutely continuous. Moreover, for any $\mathbf{x} \in \mathcal{H}$, the RND $q_t(\mathbf{x}_0, \mathbf{x}) := \frac{\mathrm{d}\nu_t^{\mathbf{x}_0}(\mathbf{x})}{\mathrm{d}\nu_\infty(\mathbf{x})}$ is given by:*

$$q_t(\mathbf{x}_0, \mathbf{x}) = C \exp\left[-\tfrac{1}{2}\langle \Theta_t^{-1} \mathbf{m}_t^{\mathbf{x}_0}, \mathbf{m}_t^{\mathbf{x}_0} \rangle_\infty + \langle \Theta_t^{-1} \mathbf{m}_t^{\mathbf{x}_0}, \mathbf{x} \rangle_\infty - \tfrac{1}{2}\langle e^{2t\mathcal{A}} \Theta_t^{-1} \mathbf{x}, \mathbf{x} \rangle_\infty\right], \tag{5}$$

where we denote $\mathbf{m}_t^{\mathbf{x}_0} := e^{t\mathcal{A}}\mathbf{x}_0$, $\Theta_t = 1 - e^{2t\mathcal{A}}$, $C = det(\Theta_t)^{-\frac{1}{2}}$ and $\langle u, v \rangle_\infty = \langle Q_\infty^{-\frac{1}{2}} u, Q_\infty^{-\frac{1}{2}} v \rangle_{\mathcal{H}}$.

Proof. From (32), we get the covariance idendity:

$$Q_t = \int_0^t \sigma_{t-r}^2 e^{r\mathcal{A}} Q e^{r\mathcal{A}^\dagger} \mathrm{d}r \tag{73}$$

$$= \int_0^t \sigma_\infty^2 e^{r\mathcal{A}} Q e^{r\mathcal{A}^\dagger} \mathrm{d}r + \int_0^t (\sigma_{t-r}^2 - \sigma_\infty^2) e^{r\mathcal{A}} Q e^{r\mathcal{A}^\dagger} \mathrm{d}r \tag{74}$$

$$= Q_\infty - \int_t^\infty e^{r\mathcal{A}} Q e^{r\mathcal{A}^\dagger} \mathrm{d}r + \int_0^t (\sigma_{t-r}^2 - \sigma_\infty^2) e^{r\mathcal{A}} Q e^{r\mathcal{A}^\dagger} \mathrm{d}r \tag{75}$$

$$= Q_\infty - e^{t\mathcal{A}} Q_\infty e^{t\mathcal{A}^\dagger} + \int_0^t (\sigma_{t-r}^2 - \sigma_\infty^2) e^{r\mathcal{A}} Q e^{r\mathcal{A}^\dagger} \mathrm{d}r. \tag{76}$$

Since, $\int_0^t (\sigma_{t-r}^2 - \sigma_\infty^2) e^{r\mathcal{A}} Q e^{r\mathcal{A}^\dagger} \mathrm{d}r \to o(1)$ as we choose $\sigma_t \to \sigma_\infty$, we get the relation:

$$Q_t \approx Q_\infty - e^{t\mathcal{A}} Q_\infty Q^{t\mathcal{A}^\dagger} = Q_\infty^{1/2} \Theta_t Q_\infty^{1/2}. \tag{77}$$

Hence, by applying [7, Proposition 1.3.11], for two centered Gaussians $\mathcal{N}(0, Q_t)$ and $\mathcal{N}(0, Q_\infty)$,

$$q_t(0, \mathbf{x}) = \frac{d\mathcal{N}(0, Q_t)}{d\mathcal{N}(0, Q_\infty)}(\mathbf{x}) = \mathit{det}(\Theta)^{-\frac{1}{2}} \exp\left[-\tfrac{1}{2}\langle e^{2t\mathcal{A}}\Theta^{-1}\mathbf{x}, \mathbf{x}\rangle_\infty\right]. \tag{78}$$

Now, for the shift from centered case to the general mean function $e^{t\mathcal{A}}\mathbf{x}_0$, we utilize the Cameron-Martin theorem [7, Theorem. 1.3.6]. It implies that:

$$\frac{d\mathcal{N}(e^{t\mathcal{A}}\mathbf{x}_0, Q_t)}{d\mathcal{N}(0, Q_t)}(\mathbf{x}) = \exp\left[\langle Q_t^{-\frac{1}{2}} e^{t\mathcal{A}}\mathbf{x}_0, Q_t^{-\frac{1}{2}}\mathbf{x}\rangle_\mathcal{H} - \tfrac{1}{2}\langle Q_t^{-\frac{1}{2}} e^{t\mathcal{A}}\mathbf{x}_0, Q_t^{-\frac{1}{2}} e^{t\mathcal{A}}\mathbf{x}_0\rangle_\mathcal{H}\right] \tag{79}$$

$$= \exp\left[\langle Q_\infty^{\frac{1}{2}} Q_t^{-1} e^{t\mathcal{A}}\mathbf{x}_0, Q_\infty^{-\frac{1}{2}}\mathbf{x}\rangle_\mathcal{H} - \tfrac{1}{2}\langle Q_\infty^{\frac{1}{2}} Q_t^{-1} e^{t\mathcal{A}}\mathbf{x}_0, Q_\infty^{-\frac{1}{2}} e^{t\mathcal{A}}\mathbf{x}_0\rangle_\mathcal{H}\right] \tag{80}$$

$$\overset{(i)}{=} \exp\left[\langle \Theta^{-1} Q_\infty^{-\frac{1}{2}} e^{t\mathcal{A}}\mathbf{x}_0, Q_\infty^{-\frac{1}{2}}\mathbf{x}\rangle_\mathcal{H} - \tfrac{1}{2}\langle \Theta^{-1} Q_\infty^{-\frac{1}{2}} e^{t\mathcal{A}}\mathbf{x}_0, Q_\infty^{-\frac{1}{2}} e^{t\mathcal{A}}\mathbf{x}_0\rangle_\mathcal{H}\right] \tag{81}$$

$$= \exp\left[\langle \Theta^{-1}\mathbf{m}_t^{\mathbf{x}_0}, \mathbf{x}\rangle_\infty - \tfrac{1}{2}\langle \Theta^{-1}\mathbf{m}_t^{\mathbf{x}_0}, \mathbf{m}_t^{\mathbf{x}_0}\rangle_\infty\right], \tag{82}$$

where (i) follows from (77) such that $\Theta^{-1}Q_\infty^{-1/2} = (Q_t^{-1/2}Q_\infty^{1/2})^\dagger Q_t^{-1/2} = Q_\infty^{1/2}Q_t^{-1}$. Hence, by chain-rule, we get the desired result:

$$q_t(\mathbf{x}_0, \mathbf{x}) := \frac{d\mathcal{N}(e^{t\mathcal{A}}\mathbf{x}_0, Q_t)}{d\mathcal{N}(0, Q_t)} \frac{d\mathcal{N}(0, Q_t)}{d\mathcal{N}(0, Q_\infty)}(\mathbf{x}) \tag{83}$$

$$= \mathit{det}(\Theta)^{-\frac{1}{2}} \exp\left[-\tfrac{1}{2}\langle \Theta^{-1}\mathbf{m}_t^{\mathbf{x}_0}, \mathbf{m}_t^{\mathbf{x}_0}\rangle_\infty + \langle \Theta^{-1}\mathbf{m}_t^{\mathbf{x}_0}, \mathbf{x}\rangle_\infty - \tfrac{1}{2}\langle e^{2t\mathcal{A}}\Theta^{-1}\mathbf{x}, \mathbf{x}\rangle_\infty\right] \tag{84}$$

It concludes the proof. ☐

B.6 Proof of Proposition 4

Proposition 4 (Adjoint Matching in $\mathcal{H}$)**.** *Consider the infinite-dimensional matching objective:*

$$\mathcal{L}(\theta) = \int_0^T \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\tfrac{1}{2}\left\|\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}}\right\|_\mathcal{H}^2\right] \mathrm{d}t, \quad \bar{\alpha} := \mathit{stopgrad}(\alpha^\theta) \tag{15}$$

$$\mathrm{d}\mathbf{Y}_t^{\bar{\alpha}} = -\mathcal{A}^\dagger \mathbf{Y}_t^{\bar{\alpha}}\mathrm{d}t + \mathbf{Z}_t^{\bar{\alpha}}\mathrm{d}\mathbf{W}_t^Q, \quad \mathbf{Y}_T^{\bar{\alpha}} = D_\mathbf{x} g(\mathbf{X}_T^{\bar{\alpha}}). \tag{16}$$

Then, we get the critical point of $\mathcal{L}$ is optimal control $\alpha^\star$ for the SOC problem in (8).

Proof. Let $\mathbf{X}^\theta := \mathbf{X}^{\alpha^\theta}$ be a controlled process with parametrized control policy α_θ and $\xi_t := \partial_\theta \mathbf{X}_t^\theta$ be the variation process of $\mathbf{X}^\theta$ with respect to θ. Since diffusion coefficient σ_t is independent to θ, we get the differential equation for the variation process by differentiate (7) with respect to θ as follows:

$$d\xi_t = \mathcal{A}\xi_t \mathrm{d}t + \sigma_t Q^{1/2}\left(D_\mathbf{x}\alpha^\theta(\mathbf{X}_t^\theta, t)\xi_t + \partial_\theta\alpha^\theta(\mathbf{X}_t^\theta, t)\right)\mathrm{d}t, \quad \xi_t = 0. \tag{85}$$

Now, let us define the cost functional with $\mathbb{P}^\theta := \mathbb{P}^{\alpha^\theta}$:

$$\mathcal{J}(\alpha^\theta, \mathbb{P}^\theta) = \mathbb{E}_{\mathbb{P}^\theta}\left[\int_0^T \frac{1}{2}\left\|\alpha^\theta(t, \mathbf{X}_t^\theta)\right\|_{\mathcal{H}}^2 dt + g(\mathbf{X}_T^\theta)\right]. \tag{86}$$

Then, differentiating cost functional (86) with respect to θ is given by,

$$\tfrac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{J}(\alpha^\theta, \mathbb{P}^\theta) = \mathbb{E}_{\mathbb{P}^\theta}\left[\int_0^T \langle \alpha^\theta(\mathbf{X}_t^\theta, t), D_{\mathbf{x}}\alpha^\theta(\mathbf{X}_t^\theta, t)\xi_t + \partial_\theta\alpha^\theta(\mathbf{X}_t^\theta, t)\rangle_{\mathcal{H}}\mathrm{d}t + \langle D_{\mathbf{x}}g(\mathbf{X}_T^\theta), \xi_T\rangle_{\mathcal{H}}\right]. \tag{87}$$

Next, by applying the Lemma B.4 to $\langle \mathbf{Y}_t^\theta, \xi_t\rangle_{\mathcal{H}}$,we obtain:

$$\mathrm{d}\langle \mathbf{Y}_t^\theta, \xi_t\rangle_{\mathcal{H}} = \langle \mathrm{d}\mathbf{Y}_t^\theta, \xi_t\rangle_{\mathcal{H}} + \langle \mathbf{Y}_t^\theta, \mathrm{d}\xi_t\rangle_{\mathcal{H}} \tag{88}$$

$$\overset{(i)}{=} \langle -\mathcal{A}^\dagger\mathbf{Y}_t^\theta, \xi_t\rangle_{\mathcal{H}}\mathrm{d}t + \langle \mathbf{Z}_t^\theta d\mathbf{W}_t^Q, \xi_t\rangle_{\mathcal{H}} \tag{89}$$

$$+ \langle \mathbf{Y}_t^\theta, \mathcal{A}\xi_t + \sigma_t Q^{1/2}\left(D_{\mathbf{x}}\alpha^\theta(\mathbf{X}_t^\theta, t)\xi_t + \partial_\theta\alpha^\theta(\mathbf{X}_t^\theta, t)\right)\rangle_{\mathcal{H}} \tag{90}$$

$$\overset{(ii)}{=} \langle \mathbf{Y}_t^\theta, \sigma_t Q^{1/2}\left(D_{\mathbf{x}}\alpha^\theta(\mathbf{X}_t^\theta, t)\xi_t + \partial_\theta\alpha^\theta(\mathbf{X}_t^\theta, t)\right)\rangle_{\mathcal{H}}\mathrm{d}t + \langle \mathbf{Z}_t^\theta d\mathbf{W}_t^Q, \xi_t\rangle_{\mathcal{H}} \tag{91}$$

where (i) follows since the quadratic covariation $[\xi, \mathbf{Y}^\theta]_t$ canceled out because ξ has no martingale term, and (ii) follows from the adjointness of $\mathcal{A}$ and $Q^{1/2}$ such that *i.e.*, $\langle \mathcal{A}^\dagger u, v\rangle_{\mathcal{H}} = \langle u, \mathcal{A}v\rangle_{\mathcal{H}}$. Since $\xi_0 = 0$ and $\mathbf{Y}_T^\theta = D_{\mathbf{x}}g(\mathbf{X}_T^\theta)$, by integrating (88) and takes expectation, we obtain:

$$\mathbb{E}_{\mathbb{P}^\theta}\left[\langle D_{\mathbf{x}}g(\mathbf{X}_T^\theta), \xi_T\rangle\right] = \mathbb{E}_{\mathbb{P}^\theta}\left[\int_0^T \mathrm{d}\langle \mathbf{Y}_t^\theta, \xi_t\rangle_{\mathcal{H}}\right] \tag{92}$$

$$= \mathbb{E}_{\mathbb{P}^\theta}\left[\int_0^T \langle \mathbf{Y}_t^\theta, \sigma_t Q^{1/2}\left(D_{\mathbf{x}}\alpha^\theta(\mathbf{X}_t^\theta, t)\xi_t + \partial_\theta\alpha^\theta(\mathbf{X}_t^\theta, t)\right)\rangle_{\mathcal{H}}\mathrm{d}t\right]. \tag{93}$$

Finally, substituting (92) into (87), we get:

$$\tfrac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{J}(\alpha^\theta, \mathbb{P}^\theta) \overset{(i)}{=} \mathbb{E}_{\mathbb{P}^\theta}\left[\int_0^T \langle \alpha^\theta(\mathbf{X}_t^\theta, t) + \sigma_t Q^{1/2}\mathbf{Y}_t^\theta, D_{\mathbf{x}}\alpha^\theta(\mathbf{X}_t^\theta, t)\xi_t + \partial_\theta\alpha^\theta(\mathbf{X}_t^\theta, t)\rangle_{\mathcal{H}}dt\right], \tag{94}$$

where (i) follows from the adjointness of $Q^{1/2}$ and scalar σ_t. Since $\bar{\alpha} = \text{stopgrad}(\alpha^\theta)$ only blocks the gradient, $\bar{\alpha}(t, \mathbf{x}) = \alpha^\theta(t, \mathbf{x})$ in pointwise manner. Hence, the forward law satisfies $\mathbb{P}^\theta = \mathbb{P}^{\bar{\alpha}}$. Therefore, from (94), we can define

$$\tfrac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{J}(\alpha^\theta, \mathbb{P}^{\bar{\alpha}}) \overset{(i)}{=} \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\int_0^T \langle \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}}, \partial_\theta\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t)\rangle_{\mathcal{H}}dt\right] \tag{95}$$

$$= \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\int_0^T \langle \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t), \partial_\theta\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t)\rangle_{\mathcal{H}} + \langle \sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}}, \partial_\theta\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t)\rangle_{\mathcal{H}}dt\right] \tag{96}$$

$$= \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\int_0^T \tfrac{1}{2}\partial_\theta\left\|\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}}\right\|_{\mathcal{H}}^2\mathrm{d}t\right] \tag{97}$$

$$\overset{(ii)}{=} \partial_\theta\int_0^T \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\tfrac{1}{2}\left\|\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}}\right\|_{\mathcal{H}}^2\right]\mathrm{d}t \tag{98}$$

$$= \tfrac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{L}(\theta), \tag{99}$$

where (i) follows by substituting $\mathbb{P}^{\bar{\alpha}}$ into (94) while $D_{\mathbf{x}}\alpha^\theta\xi_t$ disappears due to stopgrad operation blocks the gradient through $\mathbf{X}_t^{\bar{\alpha}}$ in (85) (*i.e.*, $\xi_t = 0$) and (ii) follows from dominate convergence. Therefore, from the SMP stationarity:

$$\alpha^\theta(\mathbf{X}_t^\theta, t) + \sigma_t Q^{1/2}\mathbf{Y}_t^\theta = \alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}} = 0, \quad t \in [0, T], \tag{100}$$

we obtain $\frac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{J}(\alpha^\theta, \mathbb{P}^\theta) = \frac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{L}(\theta) = 0$. This implies that the SOC objective in (8) and matching objective in (15) share the same critical points.

Now, let us verify that the critical point $\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) = -\sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}}$ in (100) is optimal control $\alpha^\star$ for the SOC problem in (8). To do so, let us assume that there exists $\mathcal{V}$ be a solution of HJB equation (49) satisfying Assumption B.3. Then, by applying Itô formula, we have:

$$d\mathcal{V}(t, \mathbf{X}_t^\theta) = \big[\partial_t \mathcal{V}_t + \langle \mathcal{A}\mathbf{X}_t^\theta, D_\mathbf{x}\mathcal{V}_t\rangle_\mathcal{H} \tag{101}$$

$$+ \tfrac{1}{2}\mathrm{Tr}(\sigma_t^2 Q D_{\mathbf{xx}}\mathcal{V}_t) + \langle \sigma_t Q^{1/2}\alpha_t^\theta, D_\mathbf{x}\mathcal{V}_t\rangle_\mathcal{H}\big]\mathrm{d}t + \langle D_\mathbf{x}\mathcal{V}, \sigma_t \mathrm{d}\mathbf{W}_t^Q\rangle_\mathcal{H} \tag{102}$$

$$\overset{(i)}{=} \big[\langle \alpha_t^\theta, \sigma_t Q^{1/2} D_\mathbf{x}\mathcal{V}_t\rangle_\mathcal{H} - F(\sigma_t Q^{1/2} D_\mathbf{x}\mathcal{V}_t)\big]\mathrm{d}t + \langle D_\mathbf{x}\mathcal{V}, \sigma_t \mathrm{d}\mathbf{W}_t^Q\rangle_\mathcal{H}, \tag{103}$$

where (i) follows from the definition of HJB equation (49) and we define $F(\sigma_t Q^{1/2} D_\mathbf{x}\mathcal{V}_t) := \inf_{\alpha\in\mathbb{A}}\left[\langle \alpha_t, \sigma_t Q^{1/2} D_\mathbf{x}\mathcal{V}_t\rangle_\mathcal{H} + \frac{1}{2}\left\|\alpha_t\right\|_\mathcal{H}^2\right]$. Now, define $\mathbf{Y}_t^\theta := D_\mathbf{x}\mathcal{V}_t$ and $\mathbf{Z}_t^\theta := D_{\mathbf{xx}}\mathbf{Z}_t$ and apply the Itô formula, then we get

$$d\mathbf{Y}_t^\theta = \big[\partial_t D_\mathbf{x}\mathcal{V}_t + \langle \mathcal{A}\mathbf{X}_t^\theta, D_{\mathbf{xx}}\mathcal{V}_t\rangle_\mathcal{H} \tag{104}$$

$$+ \tfrac{1}{2}\mathrm{Tr}(\sigma_t Q D_{\mathbf{xx}}(D_\mathbf{x}\mathcal{V}_t)) + \langle \sigma_t Q^{1/2}\alpha_t^\theta, D_{\mathbf{xx}}\mathcal{V}_t\rangle_\mathcal{H}\big] + \sigma_t \mathbf{Z}_t^\theta d\mathbf{W}_t^Q. \tag{105}$$

Here, by differentiate the HJB in (49) with respect to $\mathbf{x}$, we get:

$$\partial_t D_\mathbf{x}\mathcal{V}_t + \mathcal{A}^\dagger D_\mathbf{x}\mathcal{V} + \langle \mathcal{A}\mathbf{X}_t^\theta, D_{\mathbf{xx}}\mathcal{V}_t\rangle_\mathcal{H} \tag{106}$$

$$+ \tfrac{1}{2}\mathrm{Tr}(\sigma_t^2 Q D_{\mathbf{xx}}(D_\mathbf{x}\mathcal{V}_t)) + D_{\mathbf{xx}}\mathcal{V}_t\sigma_t Q^{1/2} D_\mathbf{x}F(\sigma_t Q^{1/2} D_\mathbf{x}\mathcal{V}) = 0. \tag{107}$$

Now, by substituting (106) into (104), we get the following:

$$\mathrm{d}\mathbf{Y}_t^\theta = \left[-\mathcal{A}^\dagger\mathbf{Y}_t + \langle \sigma_t Q^{1/2}\alpha_t^\theta - \sigma_t Q^{1/2} D_\mathbf{x}F(\sigma_t Q^{1/2}\mathbf{Y}_t^\theta), D_{\mathbf{xx}}\mathcal{V}\rangle\right]\mathrm{d}t + \sigma_t\mathbf{Z}_t^\theta d\mathbf{W}_t^Q \tag{108}$$

$$\overset{(i)}{=} \left[-\mathcal{A}^\dagger\mathbf{Y}_t^\theta + \langle \sigma_t Q^{1/2}\alpha_t^\theta + \sigma_t^2 Q\mathbf{Y}_t^\theta, D_{\mathbf{xx}}\mathcal{V}\rangle\right]\mathrm{d}t + \sigma_t\mathbf{Z}_t^\theta d\mathbf{W}_t^Q \tag{109}$$

$$\overset{(ii)}{=} -\mathcal{A}^\dagger\mathbf{Y}_t^\theta \mathrm{d}t + \sigma_t\mathbf{Z}_t^\theta d\mathbf{W}_t^Q, \tag{110}$$

where (i) follows from $F(\sigma_t Q^{1/2}\mathbf{Y}_t^\theta) = -\frac{1}{2}\left\|\sigma_t Q^{1/2}\mathbf{Y}_t^\theta\right\|_\mathcal{H}^2$ by definition of F and (ii) follows from $\alpha_t^\theta = -\sigma_t Q^{1/2}\mathbf{Y}_t^\theta$. Since the stopgrad operation does not affect any step in the derivation, therefore, if $\mathbf{Y}_t^\theta = D_\mathbf{x}\mathcal{V}(t, \mathbf{X}_t^\theta)$, then $\mathbf{Y}_t^\theta$ is the solution of (16) and the critical point in (100) satisfies

$$\alpha_t^\theta = -\sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}} = -\sigma_t Q^{1/2} D_\mathbf{x}\mathcal{V}_t. \tag{111}$$

Hence, by Lemma B.5, this control satisfies the HJB and is optimal, hence $\alpha^\theta = \alpha^\star$. Conversely, if we assume that the critical point $\alpha_t^\theta = -\sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}}$ is optimal where $\mathbf{Y}_t^{\bar{\alpha}}$ is the solution of (16), then the same verification immediately yields $\mathbf{Y}_t^{\bar{\alpha}} = D_\mathbf{x}\mathcal{V}(t, \mathbf{X}_t^{\bar{\alpha}})$. This shows that the critical point in (100) is the optimal control $\alpha^\star$ for the SOC problem in (8). This completes the proof. ☐

Sufficient Condition Indeed, in the special case *i.e.*, when g is convex, the converse becomes a sufficient condition. One may ask whether solving the adjoint equation and minimizing the Hamiltonian ensures optimality. In general, SOC sampling needs non convex g, so this condition often fails. We therefore do not apply it in our study. Instead, we formulate and prove a sufficient optimality condition tailored to specific SOC problem in (8).

Proposition B.7 (Stochastic Maximum Principle)**.** *Assume the terminal cost g is convex. Let us fix an admissible control $\bar{\alpha} \in \mathbb{A}$ and consider the corresponding infinite dimensional coupled FBSDEs:*

$$\mathrm{d}\mathbf{X}_t^{\bar{\alpha}} = \left[\mathcal{A}\mathbf{X}_t^{\bar{\alpha}} + \sigma_t Q^{1/2}\bar{\alpha}_t\right]\mathrm{d}t + \sigma_t \mathrm{d}\mathbf{W}_t^Q, \quad \mathbf{X}_0^{\bar{\alpha}} = \mathbf{x}_0, \tag{112}$$

$$\mathrm{d}\mathbf{Y}_t^{\bar{\alpha}} = -\left[\mathcal{A}^\dagger\mathbf{Y}_t^{\bar{\alpha}} + D_\mathbf{x}H(t, \mathbf{X}_t^{\bar{\alpha}}, \bar{\alpha}_t, \mathbf{Y}_t^{\bar{\alpha}}, \mathbf{Z}_t^{\bar{\alpha}})\right]\mathrm{d}t + \mathbf{Z}_t^{\bar{\alpha}}\mathrm{d}\mathbf{W}_t^Q, \quad \mathbf{Y}_T^{\bar{\alpha}} = D_\mathbf{x}g(\mathbf{X}_T^{\bar{\alpha}}). \tag{113}$$

Then, under the state variables is fixed $(\mathbf{X}_t^{\bar{\alpha}}, \mathbf{Y}_t^{\bar{\alpha}}, \mathbf{Z}_t^{\bar{\alpha}})$ for a choosen control $\bar{\alpha}$, if following holds:

$$H(t, \mathbf{X}_t^{\bar{\alpha}}, \alpha_t^\star, \mathbf{Y}_t^{\bar{\alpha}}, \mathbf{Z}_t^{\bar{\alpha}}) = \min_{\alpha\in\mathbb{A}} H(t, \mathbf{X}_t^{\bar{\alpha}}, \alpha, \mathbf{Y}_t^{\bar{\alpha}}, \mathbf{Z}_t^{\bar{\alpha}}), \quad \forall t \in [0, T], \tag{114}$$

the resulting sequence of controls $\{\alpha_t^\star\}_{t\in[0,T]}$ is the optimal control $\alpha^\star$ for a given SOC problem (8).

Proof. We adapt the proof of the sufficient condition in [4, Theorem 4.14] into our infinite-dimensional control problem. Fix $\beta \in \mathcal{U}$ be a admissible control, let us write $\mathbf{X}_t^\star := \mathbf{X}_t^{\alpha^\star}$ and define $\Delta\mathbf{X}_t := \mathbf{X}_t^\beta - \mathbf{X}_t^\star$. By convexity of g, we get

$$g(\mathbf{X}_T^\beta) - g(\mathbf{X}_T^\star) \geq \langle D_\mathbf{x} g(\mathbf{X}_T^\beta), \mathbf{X}_T^\beta - \mathbf{X}_T^\star \rangle_\mathcal{H} = \langle \mathbf{Y}_T^\star, \Delta\mathbf{X}_T \rangle_\mathcal{H}. \tag{115}$$

Now, apply the Hilbert space product rule B.4 to $\langle \mathbf{Y}_T^\star, \Delta\mathbf{X}_T \rangle_\mathcal{H}$. Then, by using the dynamics of $\mathbf{X}^\beta$ and $\mathbf{X}^\star$ and the adjoint equation for $\mathbf{Y}^\star$, we get:

$$\mathbb{E}\left[\langle \mathbf{Y}_T^\star, \Delta\mathbf{X}_T \rangle_\mathcal{H}\right] = \mathbb{E}\int_0^T \langle -\mathcal{A}^\dagger \mathbf{Y}_t^\star - D_\mathbf{x} H(t, \mathbf{X}_t^\star, \alpha^\star, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star, \Delta\mathbf{X}_t \rangle_\mathcal{H} \mathrm{d}t \tag{116}$$

$$+ \mathbb{E}\int_0^T \langle \mathbf{Y}_t^\star, \mathcal{A}\Delta\mathbf{X}_t + \sigma_t Q^{1/2}\beta_t - \sigma_t Q^{1/2}\alpha_t^\star \rangle_\mathcal{H} \mathrm{d}t \tag{117}$$

$$\overset{(i)}{=} -\mathbb{E}\int_0^T \langle D_\mathbf{x} H(t, \mathbf{X}_t^\star, \alpha^\star, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star, \Delta\mathbf{X}_t \rangle_\mathcal{H} \mathrm{d}t \tag{118}$$

$$+ \mathbb{E}\int_0^T \langle \mathbf{Y}_t^\star, \sigma_t Q^{1/2}\beta_t - \sigma_t Q^{1/2}\alpha_t^\star \rangle_\mathcal{H} \mathrm{d}t, \tag{119}$$

where (i) follows from the adjointness of $\mathcal{A}$. On the other hand, using the definition of $\mathcal{J}$ in (86):

$$\mathcal{J}(\beta) - \mathcal{J}(\alpha^\star) = \mathbb{E}\left[\int_0^T \left(\tfrac{1}{2}\|\beta_t\|_\mathcal{H}^2 - \tfrac{1}{2}\|\alpha_t^\star\|_\mathcal{H}^2\right)\mathrm{d}t + g(\mathbf{X}_T^\beta) - g(\mathbf{X}_T^\star)\right] \tag{120}$$

$$\geq \mathbb{E}\left[\int_0^T \left(\tfrac{1}{2}\|\beta_t\|_\mathcal{H}^2 - \tfrac{1}{2}\|\alpha_t^\star\|_\mathcal{H}^2\right) + \langle \mathbf{Y}_t^\star, \sigma_t Q^{1/2}\beta_t - \sigma_t Q^{1/2}\alpha_t^\star \rangle_\mathcal{H} \mathrm{d}t\right] \tag{121}$$

$$- \mathbb{E}\int_0^T \langle D_\mathbf{x} H(t, \mathbf{X}_t^\star, \alpha^\star, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star, \Delta\mathbf{X}_t \rangle_\mathcal{H} \mathrm{d}t \tag{122}$$

$$\overset{(i)}{=} \mathbb{E}\int_0^T \left[H(t, \mathbf{X}_t^\beta, \beta_t, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star) - H(t, \mathbf{X}_t^\star, \alpha^\star, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star) - \langle D_\mathbf{x} H(t, \mathbf{X}_t^\star, \alpha^\star, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star, \Delta\mathbf{X}_t \rangle_\mathcal{H}\right] \mathrm{d}t, \tag{123}$$

where (i) follows by the definition of Hamiltonian (13). By convexity of $\mathcal{H}$ with respect to $(\mathbf{x}, \alpha)$,

$$H(t, \mathbf{X}_t^\beta, \beta_t, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star) \geq H(t, \mathbf{X}_t^\star, \alpha^\star, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star) \tag{124}$$

$$+ \langle D_\mathbf{x} H(t, \mathbf{X}_t^\star, \alpha^\star, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star, \Delta\mathbf{X}_t \rangle_\mathcal{H} + \langle D_\alpha H(t, \mathbf{X}_t^\star, \alpha^\star, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star, \beta_t - \alpha_t^\star \rangle_\mathcal{H}. \tag{125}$$

Then, the point-wise minimality of $\alpha^\star$ implies that $\langle D_\alpha H(t, \mathbf{X}_t^\star, \alpha^\star, \mathbf{Y}_t^\star, \mathbf{Z}_t^\star, \beta_t - \alpha_t^\star \rangle_\mathcal{H} \geq 0$ for any admissible control β. Therefore, it implies that $\mathcal{J}(\beta) \geq \mathcal{J}(\alpha^\star)$ because of (114). Finally, since we choose β arbitrary, resulting $\alpha^\star$ is optimal control for all $t \in [0, T]$. This concludes the proof. □

B.7 Proof of Proposition 5

Proposition 5 (Unbiased Estimator)**.** *Let us define the sample-wise adjoint matching objective with conditional expectation* $\tilde{\mathbf{Y}}_t^{\bar{\alpha}} := \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}[\mathbf{Y}_t^{\bar{\alpha}} | \mathbf{X}_t^{\bar{\alpha}}]$ *for any sample trajectory* $\mathbf{X}_t^{\bar{\alpha}} \sim \mathbb{P}^{\bar{\alpha}}$*:*

$$\tilde{\mathcal{L}}(\theta) := \int_0^T \tfrac{1}{2} ||\alpha^\theta(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2} \tilde{\mathbf{Y}}_t^{\bar{\alpha}}||_\mathcal{H}^2 \mathrm{d}t \tag{17}$$

Then, we get unbiased gradient estimator $\frac{\mathrm{d}}{\mathrm{d}\theta}\mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}[\tilde{\mathcal{L}}(\theta)] = \frac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{L}(\theta)$ *with same critical point.*

Proof. Under the stopgrad operation, the forward law $\mathbb{P}^{\bar{\alpha}}$ does not depend on θ. By the dominated convergence theorem we pass the derivative through the expectation and the integral:

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\tilde{\mathcal{L}}(\theta)\right] = \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\frac{\mathrm{d}}{\mathrm{d}\theta}\int_0^T \frac{1}{2}||\alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2}\tilde{\mathbf{Y}}_t^{\bar{\alpha}}||_{\mathcal{H}}^2 \mathrm{d}t\right] \tag{126}$$

$$= \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\int_0^T \langle \alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2}\tilde{\mathbf{Y}}_t^{\bar{\alpha}},\ \partial_{\theta}\alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t)\rangle_{\mathcal{H}} \mathrm{d}t\right] \tag{127}$$

$$= \int_0^T \left[\mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\langle \alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t), \partial_{\theta}\alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t)\rangle_{\mathcal{H}}\right] + \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\langle \sigma_t Q^{1/2}\tilde{\mathbf{Y}}_t^{\bar{\alpha}}, \partial_{\theta}\alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t)\rangle_{\mathcal{H}}\right]\right] \mathrm{d}t \tag{128}$$

$$= \int_0^T \left[\mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\langle \alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t), \partial_{\theta}\alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t)\rangle_{\mathcal{H}}\right] + \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\langle \sigma_t Q^{1/2}\mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\mathbf{Y}_t^{\bar{\alpha}}|\mathbf{X}_t^{\bar{\alpha}}\right], \partial_{\theta}\alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t)\rangle_{\mathcal{H}}\right]\right] \mathrm{d}t \tag{129}$$

$$\overset{(i)}{=} \int_0^T \left[\mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\langle \alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t), \partial_{\theta}\alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t)\rangle_{\mathcal{H}}\right] + \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\langle \sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}}, \partial_{\theta}\alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t)\rangle_{\mathcal{H}}\right]\right] \mathrm{d}t \tag{130}$$

$$= \mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\int_0^T \langle \alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t) + \sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}},\ \partial_{\theta}\alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t)\rangle_{\mathcal{H}} \mathrm{d}t\right] \tag{131}$$

$$\overset{(ii)}{=} \frac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{L}(\theta), \tag{132}$$

$$\tag{133}$$

where (i) follows from the tower property and (ii) follows from (95). Hence, we get $\frac{\mathrm{d}}{\mathrm{d}\theta}\mathbb{E}_{\mathbb{P}^{\bar{\alpha}}}\left[\tilde{\mathcal{L}}(\theta)\right] = \frac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{L}(\theta) = 0$ at $\alpha^{\theta}(\mathbf{X}_t^{\bar{\alpha}}, t) = -\sigma_t Q^{1/2}\mathbf{Y}_t^{\bar{\alpha}}$. It shows that the critical point remains unchanged to $\frac{\mathrm{d}}{\mathrm{d}\theta}\mathcal{L}(\theta)$. It concludes the proof. □

B.8 Endpoint Disintegration

From (6), let us denote $\mathbf{X} := \{\mathbf{X}_t\}_{t\in[0,T]} \in \Omega$ and define the IS as follows:

$$w(\mathbf{X}) = e^{-U(\mathbf{X}_T) - \log q_T(\mathbf{x}_0, \mathbf{X}_T)}. \tag{134}$$

Then, the RND for $\mathbb{P}^{\star}$ with respect to reference measure $\mathbb{P}$ is given by:

$$\frac{d\mathbb{P}^{\star}}{d\mathbb{P}}(\mathbf{X}) = \frac{w(\mathbf{X})}{\mathbb{E}_{\mathbb{P}}[w(\mathbf{X})]}. \tag{135}$$

Then, for any bounded functional $F : \Omega \to \mathbb{R}$, we have the expectation:

$$\mathbb{E}_{\mathbb{P}^{\star}}[F(\mathbf{X}^{\star})] = \frac{\mathbb{E}_{\mathbb{P}}[F(\mathbf{X})\, w(\mathbf{X})]}{\mathbb{E}_{\mathbb{P}}[w(\mathbf{X})]}. \tag{136}$$

Since terminal law of $\mathbb{P}$ admits the density q_T with respect to ν_{∞} based on Theorem 2, we can disintegrate $\mathbb{P}$ with respect to $\mathbf{X}_T$ to obtain:

$$\mathbb{E}_{\mathbb{P}}[F(\mathbf{X})w(\mathbf{x})] = \int_{\mathcal{H}}\left(\int_{\mathcal{H}} F(\mathbf{X})\mathbb{P}_{\cdot|T}(d\mathbf{X}|\mathbf{X}_T = \mathbf{y})\right) e^{-U(\mathbf{y}) - \log q_T(\mathbf{x}_0, \mathbf{x}_T)} q_T(\mathbf{x}_0, \mathbf{y}) d\nu_{\infty}(\mathbf{y}) \tag{137}$$

$$= \int_{\mathcal{H}}\left(\int_{\mathcal{H}} F(\mathbf{X})\mathbb{P}_{\cdot|T}(d\mathbf{X}|\mathbf{X}_T = \mathbf{y})\right) e^{-U(\mathbf{y})} d\nu_{\infty}(\mathbf{y}) \tag{138}$$

$$\overset{(i)}{=} \int_{\mathcal{H}}\left(\int_{\mathcal{H}} F(\mathbf{X})\mathbb{P}_{\cdot|T}(d\mathbf{X}|\mathbf{X}_T = \mathbf{y})\right) \mathbb{E}_{\mathbb{P}}\left[w(\mathbf{X})\right] d\pi(\mathbf{y}), \tag{139}$$

where (i) follows from the definition of target distribution π in (1). Hence, we obtain that

$$\mathbb{E}_{\mathbb{P}^{\star}}[F(\mathbf{X}^{\star})] = \int \left(\int F(\mathbf{X})\mathbb{P}_{\cdot|T}(d\mathbf{X}|\mathbf{X}_T = \mathbf{y})\right) \pi(d\mathbf{y}) \tag{140}$$

$$= \mathbb{E}_{\mathbb{P}_{\cdot|T}\mathbb{P}_T^{\star}}\left[F(\mathbf{X})\right]. \tag{141}$$

C Numerical Computation

C.1 Enforcing Boundary Condition

To properly define a transition path between two metastable states, say $\mathbf{A}$ and $\mathbf{B}$, we must fix these states as the boundary conditions for a desired path. Our method achieves this by decomposing the path $\mathbf{X}_t$ into two parts: a fixed reference path $\mathbf{x}_0$ and a fluctuating residual path $\mathbf{R}_t$. Then, the total path is given by $\mathbf{X}_t = \mathbf{x}_0 + \mathbf{R}_t$, where $\mathbf{x}_0$ lifting the residual path (boundary values are zero) to the path that have desired boundary values. The reference path $\mathbf{x}_0$ provides the direct connection from $\mathbf{A}$ to $\mathbf{B}$, while the residual path $\mathbf{R}_t$ is forced to be zero at the boundaries. This construction guarantees that the full path $\mathbf{X}_t$ always begins at $\mathbf{A}$ and ends at $\mathbf{B}$. We define the lifting $\mathbf{x}_0 \in \mathcal{H}_0$ as a continuous function from a coordinate $u \in [0, L]$ to the configuration space:

$$\mathbf{x}_0[u] = c_0(u)\mathbf{A} + c_1(u)\mathbf{B}, \tag{142}$$

where $c_0(0) = 1$, $c_1(0) = 0$, $c_0(L) = 0$ and $c_1(L) = 1$ ensuring $\mathbf{x}_0[0] = \mathbf{A}$ and $\mathbf{x}_0[L] = \mathbf{B}$.

Dirichlet Space Now, consider the Dirichlet eigen-basis $\phi^{(k)}[u] = (\frac{2}{L})^{1/2} \sin\left(\frac{\pi k u}{L}\right)$ for all $k \in \mathbb{N}$ and define $\mathcal{H} = L^2([0, L])$ and $\mathcal{H}_0 = \overline{\text{span}}\{\phi^{(k)}\}$, which has zero boundary and $\mathcal{H}_0 \subset \mathcal{H}$. Then, we define the $\mathcal{H}_0$-valued residual process $\mathbf{R}_t$ on the Dirichlet space $\mathcal{H}_0$ as follows:

$$d\mathbf{R}_t = \Big[\mathcal{A}\mathbf{R}_t + \sigma_t Q^{1/2}\alpha(t, \mathbf{x}_0 + \mathbf{R}_t)\Big]\,\mathrm{d}t + \sigma_t Q^{1/2} d\mathbf{W}_t, \quad \mathbf{R}_0 = 0, \tag{143}$$

where $\mathcal{A}$ be a generator with homogeneous Dirichlet domain on $[0, L]$ and $Q^{1/2} : \mathcal{U} \to \mathcal{H}_0$ is Hilbert-Schmidt with noise schedule σ_t. Hence, we get

$$\sigma_t Q^{1/2}\alpha(t, \mathbf{X}_t) \in \mathcal{H}_0, \quad \sigma_t Q^{1/2}\mathrm{d}\mathbf{W}_t \in \mathcal{H}_0. \tag{144}$$

Then, since the space $\mathcal{H}_0$ is spanned by the Dirichlet eigen-basis, it implies that

$$\Big(\sigma_t Q^{1/2}u(t, \mathbf{X}_t)\Big)[0] = \Big(\sigma_t Q^{1/2}u(t, \mathbf{X}_t)\Big)[L] = 0, \ \Big(\sigma_t Q^{1/2}\mathrm{d}\mathbf{W}_t\Big)[0] = \Big(\sigma_t Q^{1/2}\mathrm{d}\mathbf{W}_t\Big)[L] = 0. \tag{145}$$

In other words, by expanding $\mathbf{R}_t = \sum_{k\geq 1} \mathbf{R}_t^{(k)}\phi^{(k)}$ results that the residual process $\mathbf{R}_t$ is guaranteed to be $\mathbf{R}_t[0] = 0$ and $\mathbf{R}_t[L] = 0$ for any $t \in [0, T]$ since $\phi^{(k)}[0] = 0$ and $\phi^{(k)}[L] = 0$ for any $k \in \mathbb{N}$. Consequently, since $\mathbf{X}_t = \mathbf{x}_0 + \mathbf{R}_t$, we can deduce the following $\mathcal{H}$-valued SDE:

$$d\mathbf{X}_t = \Big[\mathcal{A}(\mathbf{X}_t - \mathbf{x}_0) + \sigma_t Q^{1/2}\alpha(t, \mathbf{X}_t)\Big]\,\mathrm{d}t + \sigma_t Q^{1/2} d\mathbf{W}_t, \quad \mathbf{X}_t = \mathbf{x}_0. \tag{146}$$

To enable this construction, we choose $\mathcal{A}$ and $Q^{1/2}$ so that the operators satisfy Assumption B.1 for well-posedness and closed-form formulas used in FAS. We set $\mathcal{A} = \Delta$ with the Laplacian operator Δ on $L^2(\mathit{Dom})$ with Dirichlet boundary, where *Dom* is bounded and C^∞. Then $-\mathcal{A}$ is positive self-adjoint. We set $Q = (-\mathcal{A})^{-s}$ with $s > \frac{d}{2}$. We verify Assumption B.1 as follows.

- **(A) Trace-class** Let $\{\phi^{(k)}\}_{k\geq 1}$ be an eigen-basis with $\mathcal{A}\phi^{(k)} = -\lambda^{(k)}\phi^{(k)}$. Then, we have

$$Q\mathbf{x} = \sum_{k\geq 1} (\lambda^{(k)})^{-s}\mathbf{x}^{(k)}\phi^{(k)}. \tag{147}$$

  Hence the eigenvalues of Q are $(\lambda^{(k)})^{-s}$ and it implies that $\mathit{Tr}(Q) = \sum_{k\geq 1}(\lambda^{(k)})^{-s}$. By the Weyl law $\lambda^{(k)} \asymp k^{\frac{2}{d}}$ on a bounded C^∞ domain in $\mathbb{R}^d$, therefore $\mathit{Tr}(Q) < \infty$ if $s > \frac{d}{2}$.
- **(B) Self-adjoint** Since Laplacian Δ is self-adjoint, we get adjointness of $\mathcal{A}$ directly, and for Q:

$$Q^\dagger = ((-\mathcal{A})^{-s})^\dagger = ((-\mathcal{A})^{-s}) = Q \tag{148}$$

- **(C) Dissipative** Since $0 < \lambda_1 < \lambda_2, \cdots < \lambda_k < \cdots$ for $k \in \mathbb{N}$, we obtain:

$$\langle \mathcal{A}\mathbf{x}, \mathbf{x}\rangle_{\mathcal{H}} = -\sum_{k\geq 1} \lambda^{(k)} \left\| \langle \mathbf{x}, \phi^{(k)}\rangle_{\mathcal{H}} \right\|^2 \leq -\lambda^{(1)} \sum_{k\geq 1} \left\| \langle \mathbf{x}, \phi^{(k)}\rangle_{\mathcal{H}} \right\|^2 = -\lambda^{(1)} \left\|\mathbf{x}\right\|^2_{\mathcal{H}}. \tag{149}$$

- **(D) Commute** For any $\mathbf{x} \in \mathcal{H}$, we get:

$$\mathcal{A}Q\mathbf{x} = \mathcal{A}\left(\sum_{k\geq 1}(\lambda^{(k)})^{-s}\mathbf{x}^{(k)}\phi^{(k)}\right) = \sum_{k\geq 1}(\lambda^{(k)})^{-s}\mathbf{x}^{(k)}(\mathcal{A}\phi^{(k)}) \tag{150}$$

$$= \sum_{k\geq 1}(\lambda^{(k)})^{-s}\mathbf{x}^{(k)}(-\lambda^{(k)}\phi^{(k)}) = -\sum_{k\geq 1}(\lambda^{(k)})^{1-s}\mathbf{x}^{(k)}\phi^{(k)} \tag{151}$$

and

$$Q\mathcal{A}\mathbf{x} = -Q\left(\sum_{k\geq 1}\lambda^{(k)}\mathbf{x}^{(k)}\phi^{(k)}\right) = -\sum_{k\geq 1}\lambda^{(k)}\mathbf{x}^{(k)}(Q\phi^{(k)}) \tag{152}$$

$$= \sum_{k\geq 1}\lambda^{(k)}\mathbf{x}^{(k)}((\lambda^{(k)})^{-s}\phi^{(k)}) = -\sum_{k\geq 1}(\lambda^{(k)})^{1-s}\mathbf{x}^{(k)}\phi^{(k)}. \tag{153}$$

Therefore, we get $\mathcal{A}Q^{1/2}\mathbf{x} = Q^{1/2}\mathcal{A}\mathbf{x}$. It implies that $\mathcal{A}$ and $Q^{1/2}$ are commute.

Moreover, with this choice, we can effectively calculate the RND in Theorem 2 as follows:

$$\log q_t(\mathbf{x}_0, \mathbf{x}) = -\frac{1}{2}\sum_{k\geq 1}\frac{(q^{(k)})_\infty^{-\frac{1}{2}}}{e^{2t\lambda_k}-1}\left[\left\|\mathbf{x}_0^{(k)}\right\|^2 - 2e^{t\lambda_k}\langle\mathbf{x}_0^{(k)}, \mathbf{x}^{(k)}\rangle + \left\|\mathbf{x}^{(k)}\right\|^2\right], \tag{154}$$

where $q_\infty^{(k)}$ is diagonal component of Q_∞ in (3) such that $Q_\infty\phi^{(k)} = q_\infty^{(k)}\phi^{(k)}$. From (36), we get $Q_\infty = -\frac{1}{2}\sigma_\infty^2\mathcal{A}^{-1}Q$, therefore we get $q_\infty^{(k)} = \frac{1}{2}\sigma_\infty^2(\lambda^{(k)})^{(1+s)}$.

C.2 Finite Approximation

Numerical Simulation For Dirichlet boundary condition, we set $\phi_k[u] = (\frac{2}{L})^{\frac{1}{2}}\sin\left(\frac{\pi k u}{L}\right)$ for all $k \in \mathbb{N}$ with $\lambda_k = (\frac{\pi K}{L})^2$. Hence, we can express $\mathbf{R}_t$ as an infinite-system of real-valued SDEs:

$$\mathrm{d}\mathbf{R}_t^{(k)} = \left[-\lambda^{(k)}\mathbf{R}_t^{(k)} + \sigma_t\langle\alpha(t, \mathbf{x}_0 + \mathbf{R}_t), (Q^{1/2})^\dagger\phi^{(k)}\rangle\right]\mathrm{d}t + \sigma_t(\lambda^{(k)})^{-\frac{s}{2}}\mathrm{d}\beta_t^{(k)}, \quad \mathbf{R}_t^{(k)} = 0, \tag{155}$$

$$= \left[-\lambda^{(k)}\mathbf{R}_t^{(k)} + \sigma_t(\lambda^{(k)})^{-\frac{s}{2}}\langle\alpha(t, \mathbf{X}_t), \phi^{(k)}\rangle\right]\mathrm{d}t + \sigma_t(\lambda^{(k)})^{-\frac{s}{2}}\mathrm{d}\beta_t^{(k)} \tag{156}$$

where $\{\beta_t^{(k)}\}_{k\geq 1}$ are standard Wiener processes. Note that, to compute $(Q^{1/2})^\dagger\phi^{(k)}$, we choose $\mathcal{U} = \mathcal{H}$, means $\mathbf{W}_t = \sum_{k\geq 1}\phi_k\beta_t^{(k)}$. Even though $\mathbf{W}_t$ is not an $\mathcal{H}$-valued random variable, it is a cylindrical process indexed by $\mathcal{H}$, applying Hilbert-Schmidt $Q^{1/2}$ before stochastic integral yields the resulting process remain square integrable.

If we treat the finite path evaluations as samples on a uniform grid $[0, L]$, we project onto the sine basis with the discrete sine transform (`DST`). Since a finite-resolution path has a natural cutoff frequency K as a number of evaluation points, projection to and from the sine basis is exact. Concretely, approximate the Laplacian by $\Delta \triangleq \mathbf{E}\mathbf{D}\mathbf{E}^\top$, where $\mathbf{E}^\top$ is the `DST` projection so that $\tilde{\mathbf{R}}_t = \mathbf{E}^\top\mathbf{R}_t$ and $\mathbf{D}$ is a diagonal matrix containing the eigen values $\lambda^{(k)}$. Hence, stochastic evolution in (143) is described by the finite-dimensional model (configuration to spectral space):

$$d\tilde{\mathbf{R}}_t = \left[-\mathbf{D}\tilde{\mathbf{R}}_t + \sigma_t\mathbf{D}^{-\frac{s}{2}}\mathbf{E}^\top\alpha(t, \mathbf{X}_t)\right]\mathrm{d}t + \sigma_t\mathbf{D}^{-\frac{s}{2}}d\tilde{\mathbf{W}}_t. \tag{157}$$

The simulation algorithm for FAS is provided in detail in Algorithm 1.

Bridge Sampling Because the solution of the reference process in (2) defines an Ornstein-Uhlenbeck semigroup with Gaussian marginals for all times, the conditional law at time t given time T is:

$$\mathbb{P}_{t|T}(\mathbf{X}_t|\mathbf{x}_T) = \mathcal{N}(\mathbf{m}_{t|T}, Q_{t|T}), \tag{158}$$

where the conditional mean function and covariance operator is given by:

$$\mathbf{m}_{t|T} = \mathbf{x}_0 + Q_{t|T} Q_T^{-1}(\mathbf{x}_T - \mathbf{x}_0), \quad Q_{t|T} = Q_T - Q_{t|T} Q_T^{-1} Q_{T|t} \tag{159}$$

$$Q_t = \int_0^t e^{(t-s)\mathcal{A}} \sigma_s^2 Q e^{(t-s)\mathcal{A}^\dagger} \mathrm{d}s, \tag{160}$$

$$Q_{t|T} = \int_0^t e^{(t-s)\mathcal{A}} \sigma_s^2 Q e^{(T-s)\mathcal{A}^\dagger} \mathrm{d}s = Q_t e^{(T-t)\mathcal{A}^\dagger}, \quad Q_{T|t} = (Q_{t|T})^\dagger. \tag{161}$$

Moreover, the coordinate process $\mathbf{X}_t^{(k)} = \langle \mathbf{X}_t, \phi^{(k)} \rangle$ of the Ornstein-Uhlenbeck semigroup in (2) is also Gaussian and the coordinates are independent. Therefore we can compute coordinate-wise:

$$\mathbb{P}_{t|T}(\mathbf{X}_t^{(k)} | \mathbf{x}_T^{(k)}) = \mathcal{N}(\mathbf{m}_{t|T}^{(k)}, q_{t|T}^{(k)}) \tag{162}$$

where we denote $q_t^{(k)} = \langle Q_t, \phi^{(k)} \rangle$ for any $t \in [0, T]$ and

$$\mathbf{m}_{t|T}^{(k)} = \mathbf{x}_0^{(k)} + \tfrac{q_t^{(k)}}{q_\infty^{(k)}} e^{-(T-t)\lambda^{(k)}} (\mathbf{x}_T^{(k)} - \mathbf{x}_0^{(k)}), \tag{163}$$

$$q_{t|T}^{(k)} = q_t^{(k)} - \tfrac{(q_t^{(k)})^2}{q_\infty^{(k)}} e^{-2(T-t)\lambda^{(k)}} \tag{164}$$

$$q_t^{(k)} = \int_0^t \sigma_s^2 e^{-2(t-s)\lambda^{(k)}} (\lambda^{(k)})^{-s} \mathrm{d}s, \quad \forall t \in [0, T]. \tag{165}$$

Noise Schedule σ_t The closed-form simulation may be infeasible because computing $q_t^{(k)}$ in (165) requires numerical integration for arbitrary noise schedule σ_t. In this case, the two noise schedules studied for adjoint sampling in [16, 19] can be adapted to our setting.

- The constant noise schedule sets $\sigma(t) := \sigma$. Then, we can can compute the integral in (165):

$$Q_t^{(k)} = (\lambda^{(k)})^{-s} \frac{(1 - e^{-2t\lambda_k})}{2\lambda_k}. \tag{166}$$

- The geometric noise schedule [21, 17] decreases monotonically by predefined $\beta_{\min}$ and $\beta_{\max}$:

$$\sigma_t = \begin{cases} \beta_{\text{min}} \left(\frac{\beta_{\text{max}}}{\beta_{\text{min}}}\right)^{T-t} \sqrt{2 \log \frac{\beta_{\text{max}}}{\beta_{\text{min}}}}, & t < T \\ \beta_{\text{min}} \sqrt{2 \log \frac{\beta_{\text{max}}}{\beta_{\text{min}}}}, & t \geq T. \end{cases} \tag{167}$$

  Because choosing the noise schedule $\sigma_t = \beta e^{\alpha t}$ gives a closed-form for the integral in (165), we rearrange it as follows:

$$\sigma_t = \beta_{\text{min}} \left(\frac{\beta_{\text{max}}}{\beta_{\text{min}}}\right)^{T-t} \sqrt{2 \log \frac{\beta_{\text{max}}}{\beta_{\text{min}}}} = \beta_{\text{min}} (\tfrac{\beta_{\text{max}}}{\beta_{\text{min}}})^T \sqrt{2 \log \frac{\beta_{\text{max}}}{\beta_{\text{min}}}} e^{-\log\left(\frac{\beta_{\text{max}}}{\beta_{\text{min}}}\right) t}. \tag{168}$$

  Therefore, we get $\beta = \beta_{\text{min}} (\frac{\beta_{\text{max}}}{\beta_{\text{min}}})^T \sqrt{2 \log \frac{\beta_{\text{max}}}{\beta_{\text{min}}}}$, $\alpha = -\log\left(\frac{\beta_{\text{max}}}{\beta_{\text{min}}}\right)$. This choice yields that:

$$q_t^{(k)} = \beta^2 (\lambda^{(k)})^{-s} \frac{e^{2\alpha t} - e^{-2\lambda_k t}}{2(\alpha + \lambda_k)} \tag{169}$$

Replay buffer $\mathcal{B}$ Following previous diffusion samplers [16, 19, 5], we maintain a replay buffer $\mathcal{B}$ for the computation of the objective in (19). In this case, the expectation under $\mathbb{P}^\alpha$ is approximated by an average over $\mathcal{B}$, which holds the most recent $|\mathcal{B}|$ samples. We refresh the buffer $\mathcal{B}$ by inserting N new samples after every L gradient steps.

Parametrization fo control network We parameterize the control α^θ with a neural operator to preserve its functional form. Following [19], we parametrize the control as $\alpha^\theta(t, \mathbf{x}) := \sigma_t Q^{1/2} u^\theta(t, \mathbf{x})$, which absorbs the noise schedule into the control and removes its explicit appearance from the matching objective in (19), yielding more stable training. For MB potential, we take u^θ to be a Fourier Neural Operator (FNO) [18] with 6 layers, 16 Fourier modes, and width 64. We also condition on the time t and the grid coordinate $\mathbf{U}$ by providing Fourier-feature embeddings of both to $u(t, \mathbf{x}; \mathbf{U})$.

Clipping α_{max} We clip the energy gradient so its maximum norm does not exceed α_{max}

Time scaling λ_t As standard in AM, we apply a time scaling $\lambda_t = \frac{1}{\sigma_t^2}$ to improve numerical stability.

Algorithm 1 FAS Sampling with path lifting

Require: Initial condition $\mathbf{x}_0$, linear operator $\mathcal{A}$, trace-class operator Q, control α, discrete and inverse sine transforms $\texttt{DST}$, $\texttt{iDST}$, set evaluation points $\mathbf{U} = \{u_i\}_{i=1}^{K}$.

Set initial condition $\mathbf{X}_0^\alpha = \mathbf{x}_0$ and $\mathbf{R}_0^\alpha = \mathbf{0}$
for time t **in range** $[0, T]$ **do**
  Estimate control $\alpha_t = \alpha(\mathbf{X}_t^\alpha, t)$
  Sample Gaussian noise $\varepsilon \sim \mathcal{N}(0, \mathbf{I}_K)$
  Project state, control $\{\tilde{\mathbf{R}}_t^{(k)}\}_{k \in \mathbf{U}} = \texttt{DST}(\mathbf{R}_t^\alpha), \{\tilde{\alpha}_t^{(k)}\}_{k \in \mathbf{U}} = \texttt{DST}(\alpha_t)$
  for $k = 1, \cdots, |\mathbf{U}|$ **do in parallel**
    $\tilde{\mathbf{R}}_{t+\delta_t}^{(k)} = \left[-\lambda^{(k)} \tilde{\mathbf{R}}_t^{(k)} + \sigma_t (\lambda^{(k)})^{-\frac{s}{2}} \tilde{\alpha}_t^{(k)}\right] \delta_t + \sigma_t (\lambda^{(k)})^{-\frac{s}{2}} \varepsilon \sqrt{\delta_t}$
  end for
  Get residual $\mathbf{R}_{t+\delta_t}^\alpha = \texttt{iDST}(\{\tilde{\mathbf{R}}_{t+\delta_t}^{(k)}\}_{k \in \mathbf{U}})$
  Update $\mathbf{R}_t^\alpha = \mathbf{R}_{t+\delta_t}^\alpha$
  Get path $\mathbf{X}_t^\alpha = \mathbf{x}_0 + \mathbf{R}_t^\alpha$
end for
return $\mathbf{X}_T^\alpha \sim \mathbb{P}_T^\alpha$

Algorithm 2 Function Space Adjoint Sampler (FAS)

Require: Initial condition $\mathbf{x}_0$, terminal cost $g(x)$, control network $\alpha^\theta(\mathbf{x}, t)$, replay buffer $\mathcal{B}$, total epoch M, resample sizes N, gradient steps L, time scaling λ_t, maximum energy norm α_{max}.

for epoch m **in** $1, 2, \ldots, M$ **do**
  Sample N paths $\mathbf{X}_T^{\bar{\alpha}^m}$ from Algorithm 1 with $\bar{\alpha}^m = \text{stopgrad}(\alpha_\theta^{\theta_m})$
  Compute adjoint $\mathbf{Y}_T^{\bar{\alpha}^m} = \texttt{clip}\left(\nabla g(\mathbf{X}_T^{\bar{\alpha}^m}), \alpha_{\text{max}}\right)$
  Update buffer $\mathcal{B} \leftarrow \mathcal{B} \cup (\mathbf{X}_T^{\bar{\alpha}^m}, \mathbf{Y}_T^{\bar{\alpha}^m})$
  for step **in** $1, 2, \ldots, L$ **do**
    Sample $t \sim \mathcal{U}[0, T]$, $(\mathbf{X}_T, \mathbf{Y}_T) \sim \mathcal{B}$ and $\mathbf{X}_t \sim \mathbb{P}_{t|T}(\cdot|\mathbf{X}_T)$
    Compute the training objective:
$$\mathcal{L}_{\texttt{FAS}}(\theta) = \mathbb{E}_{t, (\mathbf{X}_T, \mathbf{Y}_T), \mathbf{X}_t} \left[\lambda_t \left\| \alpha^{\theta_s}(\mathbf{X}_t, t) + \sigma_t Q^{1/2} e^{-(T-t)\mathcal{A}^\dagger} \mathbf{Y}_T \right\|_{\mathcal{H}}^2\right]$$
    Update θ_m^{l+1} with $\nabla_\theta \mathcal{L}_{\texttt{FAS}}(\theta_m^l)$.
  end for
end for
return Approximated optimal control $\alpha^\star \approx \alpha^{\theta_M}$

Algorithms The overall algorithm of the implementation of FAS is summarized in the Algorithm 2.

D Experimental Details

D.1 Müller Brown Potential

The Müller–Brown potential $V : \mathbb{R}^2 \to \mathbb{R}$ used in our experiments is given by

$$\begin{aligned} V(x, y) = &-200 \cdot \exp(-(x_1 - 1)^2 - 10y^2) - 100 \cdot \exp(-x^2 - 10 \cdot (y - 0.5)^2) && (170) \\ &- 170 \cdot \exp(-6.5 \cdot (x + 0.5)^2 + 11 \cdot (x + 0.5) \cdot (y - 1.5) - 6.5 \cdot (y - 1.5)^2) && (171) \\ &+ 15 \cdot \exp(0.7 \cdot (x + 1)^2 + 0.6 \cdot (x + 1) \cdot (y - 1) + 0.7 \cdot (y - 1)^2). && (172) \end{aligned}$$

D.2 Baselines

For DL [13], we run their official codes using the default settings from their implementation.

D.3 Evaluation Metrics

Target hit percentage (THP): THP is the indicator-based success rate of trajectories reaching the target metastable state $\mathbf{B}$. Given endpoints $\{\mathbf{X}_T^{(i)}\}_{i=1}^N$ from N paths, we define THP as:

$$\text{THP} = 100 \times \frac{|\{i : \mathbf{X}_T^{(i)} \in \mathcal{B}\}|}{N}, \quad \text{where} \quad \mathcal{B} = \{\mathbf{X} | \mathbf{X} - \mathbf{B} < \epsilon\}. \tag{173}$$

We set $\epsilon = 0.1$ in this experiments.

Max-Energy: We measures how well the method find probable transition states when crossing the energy barrier. Formally, given a transition path $\mathbf{X}$ of length L that reaches $\mathbf{B}$, we define

$$\text{Max-Energy}(\mathbf{X}) = \max_{u \in [0,L]} V(\mathbf{X}[u]) \tag{174}$$

Log-Likelihood: To sample trajectories under fixed boundary conditions, we target the transition-path distribution [10]. Given an initial configuration $\mathbf{X}[0] \sim \pi_{\mathbf{A}}$ and a discrete MD/overdamped–Langevin evolution for L steps, a path $\mathbf{X} = \mathbf{X}[0], \ldots, \mathbf{X}[L]$ has probability:

$$\pi(\mathbf{X}) = \pi_0(\mathbf{A}) \cdot \prod_{i=0}^{N-1} \mathbf{K}(\mathbf{X}[t_{i+1}] | \mathbf{X}[t_i]) \pi_{\mathbf{B}}(\mathbf{X}[L]), \tag{175}$$

where the one-step kernel (Euler–Maruyama discretization) is

$$\mathbf{K}(\mathbf{x}, \mathbf{y}) = \mathcal{N}(\mathbf{x} | \mathbf{y} - \tfrac{1}{\gamma M} \nabla_{\mathbf{x}} V(\mathbf{y}) \mathrm{d}t, \tfrac{2}{\gamma} k_B T \mathrm{d}t). \tag{176}$$

The most probable path maximizes $\log \pi(\mathbf{X})$ under V while connecting the metastable basins. We therefore use $U(\mathbf{X}) = -\log \pi(\mathbf{X})$ as the potential in the target cost functional (86). For evaluation across different lengths, we report the length-normalized log-likelihood

$$\log \pi(\mathbf{X}) = \log \pi_0(\mathbf{X}[0]) + \tfrac{1}{N} \sum_{i=0}^{N-1} \mathbf{K}(\mathbf{X}[t_{i+1}] | \mathbf{X}[t_i]) + \log \pi_{\mathbf{B}}(\mathbf{X}[L]). \tag{177}$$