Discrete Markov Probabilistic Models: An Improved Discrete Score-Based Framework with sharp convergence bounds under minimal assumptions

Le Tuyet Nhi Pham^{*1} Dario Shariatian^{*2} Antonio Ocello¹ Giovanni Conforti³ Alain Durmus¹

Abstract

This paper introduces the Discrete Markov Probabilistic Model (DMPM), a novel algorithm for discrete data generation. The algorithm operates in discrete space, where the noising process is a continuous-time Markov chain that can be sampled exactly via a Poissonian clock that flips labels uniformly at random. The time-reversal process, like the forward noise process, is a jump process, with its intensity governed by a discrete analogue of the classical score function. Crucially, this intensity is proven to be the conditional expectation of a function of the forward process, strengthening its theoretical alignment with score-based generative models while ensuring robustness and efficiency. We further establish convergence bounds for the algorithm under minimal assumptions and demonstrate its effectiveness through experiments on low-dimensional Bernoulli-distributed datasets and high-dimensional binary MNIST data. The results highlight its strong performance in generating discrete structures. This work bridges theoretical foundations and practical applications, advancing the development of effective and theoretically grounded discrete generative modeling.

Introduction

Score-based Generative Models (SGMs) have become a key reference for generating complex data, such as images (see, e.g., Rombach et al., 2022; Ramesh et al., 2022; Saharia et al., 2022), audio (Chen et al., 2020; Kong et al., 2020), and video (Ho et al., 2022; Villegas et al., 2022; Bar-Tal

et al., 2024). In continuous time, this approach benefits from a strong theoretical framework, and a scalable, stable learning objective.

By contrast, discrete generative modeling continues to pose significant challenges. Multiple diffusion-based methods have recently been proposed for discrete spaces (Austin et al., 2021; Hoogeboom et al., 2021; Shi et al., 2024; Campbell et al., 2022; Holderrieth et al., 2024; Ren et al., 2024), or spaces of mixed type (Bertazzi et al., 2024), but there is still no consensus on which approach is theoretically sound or most practically efficient. Various formulations rely on complex forward kernels or computationally unstable ratio-based estimators for backward transitions, leading to limited convergence guarantees and high computational costs in high dimensions. Furthermore, recent analyses of discrete diffusions have introduced valuable theoretical tools (Campbell et al., 2022; Holderrieth et al., 2024; Ren et al., 2024), yet most methods remain either overly generic or require strong assumptions, making them difficult to scale or to deploy with simple, stable training objectives.

Contributions. In this paper, we introduce *Discrete Markov Probabilistic Models* (DMPMs), a score-based generative models for discrete data inspired by Sohl-Dickstein et al. (2015) and Austin et al. (2021) that bridges these gaps. Our framework specializes the forward noising process to a continuous-time Markov chain on the hypercube $\{0,1\}^d$. Leveraging theoretical insights on time-reversal Markov dynamics of this process, this choice preserves the key strengths and structure of continuous SGMs, addressing the issues raised in prior work. Our main results are summarized as follows:

- Forward-Backward Construction. We provide a principled derivation of the noising (forward) and denoising (backward) processes.
- Score and denoiser function and stable estimation. Our analysis reveals that the time-reversed process inherits a score function with an explicit conditional expectation form. By formulating the learning objective as an L² projection onto this score, we obtain a simple and principled regression loss term, as opposed to the high-variance estimators based on probability ratios

^{*}Equal contribution ¹Centre de Mathématiques Appliquées (CMAP), École Polytechnique, Institut Polytechnique de Paris, 91120 Palaiseau, France ²Sierra lab, Inria, Paris, France ³Department of Mathematics, University of Padova, Italy. Correspondence to: Le Tuyet Nhi Pham <le-tuyet-nhi.pham@polytechnique.edu>.

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

used in some prior work. Furthermore, we introduce a *denoiser reparameterization* of the score function, closely paralleling the continuous setting, which provides an interpretable and practical target for model training.

- **Theoretical Guarantees.** We prove that DMPMs converge to the underlying data distribution under minimal assumptions, providing non-asymptotic error bounds that underscore the method's reliability. In contrast to many earlier works, we prove that the sampling error grows linearly rather than exponentially with respect to dimension.
- Empirical Performance. We demonstrate that our approach attains competitive or superior performance on discrete datasets, including binarized MNIST, frequently with fewer function evaluations compared to existing discrete diffusion frameworks (e.g., 2.89 vs 7.34 FID compared to Discrete Flow Matching, Gat et al., 2024, with 2.5x fewer network calls).

Notation. Given a measurable space $(\mathsf{E}, \mathcal{E})$, we denote by $\mathcal{P}(\mathsf{E})$ the set of probability measures on E . Given two probability measures $\mu, \nu \in \mathcal{P}(\mathsf{E})$, the Kullback–Leibler divergence (also called relative entropy) of μ with respect to ν is defined as $\mathrm{KL}(\mu|\nu) := \int \log \frac{d\mu}{d\nu} d\mu$ if μ is absolutely continuous with respect to ν , and $\mathrm{KL}(\mu|\nu) = +\infty$ otherwise. The total variation distance between μ and ν is defined as $\|\mu - \nu\|_{\mathrm{TV}} = \sup_{\mathsf{A} \in \mathcal{E}} |\mu(\mathsf{A}) - \nu(\mathsf{A})|$. Consider a random variable *X*, we denote by $\mathrm{Law}(X)$ the law of *X*. We denote by δ_x the Dirac mass at *x*.

1. Forward and backward process of DMPMs

We propose a generative modeling framework that adapts diffusion-based methods to discrete spaces. Let $(\vec{X}_t)_{t \in [0,T]}$ be a forward Markov process on $\{0,1\}^d$, initialized from the data distribution μ^* , and evolving over a fixed time horizon $T_f > 0$ toward a simple base distribution. In the continuous setting, this role is often played by the Ornstein–Uhlenbeck process converging to a Gaussian. We define the corresponding backward process $(\vec{X}_t)_{t \in [0,T_f]}$ as $\vec{X}_t := \vec{X}_{T-t}$, which reconstructs μ^* from the base distribution. While the backward process can be Markovian too, its transition rates or drfit terms are typically intractable and must be approximated from forward trajectories, commonly done via score matching in continuous domains.

To adapt this idea, we introduce a forward CTMC on $\{0, 1\}^d$ where each bit flips via an independent Poisson clock. We show that its time-reversal remains a tractable CTMC with backward rates given by conditional expectations over the forward process, enabling efficient regression-based training in the discrete setting.

1.1. CTMCs on discrete state-spaces

A CTMC $(X_t)_{t \in [0,T_f]}$ defined over the discrete space X is a Markovian stochastic process which is piecewise constant with jump at random times following a non-homogeneous Poisson process. As we shall see, under mild assumption, a (non-homogeneous) CTMC is uniquely characterized by a family of rate matrices $(q_t)_{t \in [0,T_f]}$, $q_t : X \times X \to \mathbb{R}$, which constitutes the infinitesimal generator of the process, and should satisfy $\sum_{y \in X} q_t(x, y) = 0$, for any $x \in X$. In particular, this object allows to define a Markov process whose transitions are informally characterized as $h \to 0$ by:

$$\mathbb{P}(X_{t+h} = y | X_t = x) = \delta_x(y) + hq_t(x, y) + o(h) , \quad (1)$$

where o is the standard little-o Landau notation.

Definition and sampling procedure. When a simple characterization of the transition probability matrix is not available or does not exist, it is possible to simulate the process using the rate matrices, as suggested by equation (1). Popular sampling strategies include Gillespie's algorithm or τ -leaping (Gillespie, 2007). In our case, we introduce the former which uses the jump rate and jump kernel associated to the process, defined as:

$$\lambda_t(x) = \sum_{y \in \mathsf{X}} q_t(x, y) , \quad k_t(x, y) = \mathbb{1}_{x \neq y} \frac{q_t(x, y)}{\lambda_t(x)} .$$
 (2)

Informally, the jump rate governs the frequency of the random jumps of the process, and the jump kernel the next state at these jumps. More precisely, starting from a drawn X_0 from μ_0 and a sequence of i.i.d. random variables distributed according to the exponential distribution with parameter 1, $\{E_i : i \in \mathbb{N}\}$, we can define the jump times $(T_i)_{i \in \mathbb{N}}$ of the process and its transition by induction setting $T_0 = 0$. Given (T_i, X_{T_i}) , we define the next jump time as $T_{i+1} = T_i + \Delta T_{i+1}$, where

$$\Delta T_{i+1} = \inf\{t \ge 0 : \int_0^t \lambda_{T_i+r}(X_{T_i}) \mathrm{d}r \ge E_i\}.$$
 (3)

Then, set $X_t = X_{T_i}$ for $t \in (T_i, T_{i+1} \wedge T_f)$, and finally if $T_{i+1} < T_f, X_{T_{i+1}} = Y$ for Y distributed according to $Cate(\{k_{T_{i+1}}(X_{T_i}, y)\}_{y \in X})$. Note that in the case where $\lambda_t = \lambda > 0$ for any $t \in [0, T_f], (T_i)_{i \in \mathbb{N}}$ simply corresponds to the jump times a simple homogeneous Poisson process with rate λ , defined as $\mathfrak{N}_t = \sum_{i>1} \mathbb{1}_{T_i \leq t}$. Another equivalent procedure for simulating the process is provided in the supplement (see Appendix C.1). In practice, since the integral in (3) cannot be computed exactly, a time-discretized sampling strategy must be employed. We present these methodological considerations in Section 1.5 below.

Kolmogorov equation. As a final remark, the timemarginal distributions of the process $(X_t)_{t \in [0,T_f]}$ starting from $X_0, \nu_t^{(X_0)}(x) = \mathbb{P}(X_t = x)$ and identified as a tuple, satisfy the backward Kolmogorov equation:

$$\partial_t \nu_t^{(X_0)}(x) = [\nu_t^{(X_0)}]^{\mathrm{T}} q_t(x) = \sum_{y \in \mathsf{X}} \nu_t^{(X_0)}(y) q_t(y, x) .$$
(4)

Here, we identify q_t as a matrix, similarly to $\nu_t^{(X_0)}$. When the rate matrix is simple (e.g., as we use in our forward process in Section 1.2), we can use the Kolmogorov equation to derive the transition probability matrix $p_t(x_0, x_t) =$ $\nu_t^{(x_0)}(x_t) = \mathbb{P}(X_t = x_t | X_0 = x_0)$, for $x_0, x_t \in X, 0 \leq$ $t \leq T_f$, yielding a *simulation-free* procedure.

1.2. Simple case $X = \{0, 1\}$

To introduce the key ideas, we first consider the simple case where the state space is $X = \{0, 1\}$, *i.e.*, when d = 1. We will then extend our method to d > 1 by factorizing the forward process over dimensions.

Forward process. We define the forward process $(\vec{X}_t)_{t \in [0,T_f]}$ as the homogeneous CTMC starting from $\vec{X}_0 \sim \mu^*$, and driven by a simple bit-flip process associated with the rate matrix defined for any $x, y \in X$, as

$$\vec{q}^{1}(x,y) := \begin{cases} \lambda , & \text{if } y \neq x ,\\ -\lambda , & \text{otherwise }. \end{cases}$$
(5)

In the case of a constant forward rate matrix \overrightarrow{q}^1 , the transition probability matrix \overrightarrow{p}_t^1 , for $0 \le t \le T_f$, is known to be $\overrightarrow{p}_t^1 = \exp(t\overrightarrow{q}^1)$ (Liggett, 2010), which we compute as:

$$\overrightarrow{p}_{t}^{1}(x,y) = \begin{cases} \frac{1}{2} + \frac{1}{2}e^{-2\lambda t} , & \text{if } x = y ,\\ \frac{1}{2} - \frac{1}{2}e^{-2\lambda t} , & \text{otherwise} . \end{cases}$$
(6)

The detailed derivations are given in the supplementary material B.1 and rely on equation (4). We note that using a time-dependent rate λ_t is a straightforward extension, as defining $\overrightarrow{q}_t^1 = \lambda_t \overrightarrow{q}^1$ yields the transition matrix $\exp(\overrightarrow{q}^1 \int_0^t \lambda_s ds) = \overrightarrow{p}_{\int_0^t \lambda_s ds}^1$, and leave it to future work. Backward process. To recover the data distribution, we analyze the time-reversed process, which is denoted by $(\overleftarrow{X}_t)_{t \in [0, T_f]}$, and defined as $\overleftarrow{X}_t = \overrightarrow{X}_{T_f - t}$ for any $t \in [0, T_f]$. Conforti & Léonard (2022, Theorem 2.8) show that

 $(\overline{X}_t)_{t \in [0, T_f]}$ is also a non-homogeneous CTMC, associated with a family of generator matrices $(\overline{q}_t)_{t \in [0, T_f]}$ satisfying the time-reversal formula:

$$\mu_{T_f-t}(x)\overleftarrow{q}_t^1(x,y) = \mu_{T_f-t}(y)\overrightarrow{q}^1(y,x) , \qquad (7)$$

for any $0 \leq t \leq T_f$ and $x \neq y \in X$, where for any $t \in$

 $[0, T_f]$, we denote by μ_t the forward marginal distribution:

$$\mu_t(x) = \mathbb{P}(\vec{X}_t = x) . \tag{8}$$

Since our chosen rate matrix \overrightarrow{q} is symmetric (see (5)) and $\mu_{T_f-t}(x) > 0$ for all $x \in X$, $t \in [0, T_f)$, we deduce that the backward generator \overleftarrow{q}_t^1 for $0 \le t < T_f$ is given for any $x \ne y \in X$ by

$$\overleftarrow{q}_{t}^{1}(x,y) = \overrightarrow{q}^{1}(y,x)\frac{\mu_{T_{f}-t}(y)}{\mu_{T_{f}-t}(x)}.$$
(9)

Discrete score function. Define $s_t : X \to \mathbb{R}$ for any $x \in X$ by

$$s_t(x) := \frac{\mu_{T_f - t}(x) - \mu_{T_f - t}(1 - x)}{\mu_{T_f - t}(x)} .$$
(10)

 s_t acts as a discrete derivative in X of $\log \mu_t$, and thus serves as a discrete analogue of the score function in continuous models. With this notation, $\overleftarrow{q}_t(x, y)$ (9) can be expressed, for any $x \neq y \in X$, for $0 \leq t < T_f$, as:

$$\overleftarrow{q}_{t}^{1}(x,y) := \begin{cases} \lambda(1-s_{t}(x)) , & \text{if } y \neq x ,\\ -\lambda(1-s_{t}(x)) , & \text{otherwise }. \end{cases}$$
(11)

Access to the sequence $(s_t)_{t \in [0,T_f]}$ is equivalent to having access to $(\overleftarrow{q}_t^1)_{t \in [0,T_f]}$, and therefore allows to sample from \overleftarrow{X}_t for any $t \in [0, T_f]$ using the procedure described in Section 1.1. However, the score function s is generally intractable, as it depends on the unknown marginal distributions $(\mu_t)_{t \in [0,T_f]}$. To address this, we derive an alternative expression for s_t in terms of a conditional expectation over the forward process. For $x \in X$ and $t \in [0, T_f)$,

$$s_t(x) = \left[\frac{2\alpha_{T_f-t}}{1+\alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}\mathbb{1}_{\overrightarrow{X}_0}(\overrightarrow{X}_{T_f-t})}{1-\alpha_{T_f-t}^2} \middle| \overrightarrow{X}_{T_f-t} = x \right],$$
(12)

with

$$\mathbf{e}_t := \mathbf{e}^{-2\lambda t} \ . \tag{13}$$

 $\alpha_t := \mathrm{e}^{-2 \wedge t} \; .$ Indeed, the score function can be computed as

$$\begin{split} s_t(x) &= 1 - \frac{\mu_{T_f - t}(y)}{\mu_{T_f - t}(x)} = 1 - \sum_{x_0 \in \mathsf{X}} \frac{p_{T_f - t|0}(y|x_0)}{\mu_{T_f - t}(x)} \mu^\star(x_0) \\ &= 1 - \sum_{x_0 \in \mathsf{X}} \frac{p_{T_f - t|0}(y|x_0)}{p_{T_f - t|0}(x|x_0)} p_{0|T_f - t}(x_0|x) \\ &= \mathbb{E} \left[1 - \frac{p_{T_f - t|0}(y|\overrightarrow{X_0})}{p_{T_f - t|0}(x|\overrightarrow{X_0})} \ \middle| \ \overrightarrow{X}_{T_f - t} = x \right] \,, \end{split}$$

and plugging in the expression for our forward transition

matrix given in (6) we obtain equation (12). Therefore, the function s is an L²-projection and its approximation boils down to a regression problem.

1.3. General state space $X = \{0, 1\}^d$

Forward process. We generalize the previous results for the hypercube in \mathbb{R}^d , *i.e.*, the state space is $X = \{0, 1\}^d$ with $d \in \mathbb{N}^*$. We consider the homogeneous CTMC $(\overrightarrow{X}_t)_{t \in [0, T_f]}$ starting from $\overrightarrow{X}_0 \sim \mu^*$, defined with the following rate matrix:

$$\overrightarrow{q}(x,y) := \begin{cases} \lambda , & \text{if } \|y - x\|^2 = 1 ,\\ -\lambda d , & \text{if } y = x ,\\ 0 , & \text{otherwise }. \end{cases}$$
(14)

The corresponding sampling procedure following the one provided in Section 1.1 is given in Appendix B.2 for completeness. Similarly to the one-dimensional case, we can establish an explicit expression for the transition probability matrix \vec{p}_t for $0 \le t \le T_f$ as

$$\overrightarrow{p}_t(x,y) = \prod_{i=1}^d \overrightarrow{p}_t^1(x^i, y^i) , \qquad (15)$$

where $\overrightarrow{p}_{t}^{1}$ is defined in (6) and $x = (x^{i})_{i=1}^{d}$, $y = (y^{i})_{i=1}^{d} \in X$. The detailed computation is given in the supplementary material B.2.1 The factorization of the transition probability in (15) is of great practical interest, as this tells us that the dynamic of the forward process simply consists in the single-bit forward dynamic applied independently to each component, as described in Section 1.2. As a consequence, the forward marginal distribution μ_t of \overrightarrow{X}_t admits the formula

$$\mu_t(x) = \sum_{z \in \mathsf{X}} \mu_0(z) \prod_{i=1}^d \overrightarrow{p}_t(z, x) .$$
 (16)

Backward process and score function. Denote by $(\overline{X}_t)_{t\in[0,T_f]}$, the time-reversal process associated with $(\overline{X}_t)_{t\in[0,T_f]}$, and defined as $\overline{X}_t = \overline{X}_{T_f-t}$ for any $t \in [0,T_f]$. As in the case d = 1, Conforti & Léonard (2022, Theorem 2.8) shows that $(\overline{X}_t)_{t\in[0,T_f]}$ is also a *non-homogeneous* CTMC, with backward generator matrix $(\overline{q}_t)_{t\in[0,T_f]}$ that satisfies (7) and therefore (9), proceeding as before. As in the case d = 1, we show that $(\overline{q}_t)_{t\in[0,T_f]}$ depends only on a discrete score function, which we now introduce.

First, note that (9) and (14) yield $\overleftarrow{q}_t(x, y) = 0$, for $x, y \in X$ satisfying $||x - y||^2 \neq 1$ and $x \neq y$. Then, for $0 \leq t < T_f$, define $s_t : X \to \mathbb{R}^d$ for any $x \in X$, $s_t(x) = \{s_t^\ell(x)\}_{\ell=1}^d$ as the vector in \mathbb{R}^d , with components $\ell \in \{1, \ldots, d\}$,

$$s_t^{\ell}(x) := \frac{\mu_{T_f - t}(x) - \mu_{T_f - t}(\varphi^{(\ell)}(x))}{\mu_{T_f - t}(x)} , \qquad (17)$$

where $\varphi^{(\ell)} : X \to X$ is defined as $\varphi^{(\ell)}(x) = y$, with y obtained by flipping the ℓ -th bit of x, *i.e.*, $y^{\ell} = 1 - x^{\ell}$, and $y^i = x^i$ for $i \neq \ell$. Then, for $0 \leq t < T_f$, $x \neq y \in X$, we can write the backward generator $q_t(x, y)$, as given in (7), as:

$$\overleftarrow{q}_t(x,y) = \sum_{\ell=1}^a \lambda(1 - s_t^\ell(x)) \mathbb{1}_{y = \varphi^{(\ell)}(x)}$$

Score function. Note that the function *s* thus defined is an extension to the case $d \ge 1$ of the function *s* defined for d = 1 in (10). As a result, s_t is a conditional expectation over the forward process, where each of its components admits an expression similar to the 1d case (12).

Proposition 1.1. *The score function can be expressed as a conditional expectation:*

$$s_t^{\ell}(x) = \mathbb{E}\left[f_t^{\ell}(\overrightarrow{X}_0, \overrightarrow{X}_{T_f-t}) | \overrightarrow{X}_{T_f-t} = x\right] , \qquad (18)$$

where $t \in [0, T_f)$, $x \in X$, $\ell = 1, ..., d$, s_t^{ℓ} is the ℓ -th component of the score function s_t , and

$$f_t^{\ell}(\vec{X}_0, \vec{X}_{T_f-t}) = \frac{2\alpha_{T_f-t}}{1 + \alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}(\vec{X}_{T_f-t}^{\ell} - \vec{X}_0^{\ell})^2}{1 - \alpha_{T_f-t}^2}$$
(19)

The proof of this result is given in Appendix B.2.2. Similarly to the 1d case, access to the score allows to simulate the backward process following the procedure described in Section 1.1 since the the non-homogeneous jump rate $\overleftarrow{\lambda}_t$ and jump kernel \overleftarrow{k}_t of the backward process are given by

$$\begin{aligned} \overleftarrow{\lambda}_t(x) &= \lambda \sum_{\ell=1}^d (1 - s_t^\ell(x)) ,\\ \overleftarrow{k}_t(x, y) &= \sum_{\ell=1}^d \mathbb{1}_{y = \varphi^{(\ell)}(x)} \cdot \lambda (1 - s_t^\ell(x)) / \overleftarrow{\lambda}_t(x) , \end{aligned}$$
(20)

for $x \neq y \in X$ and $t \in [0, T_f)$.

1.4. Approximating the score function

In this section, we derive the training objective to estimate the score function $(s_t)_{t \in [0,T_f)}$ defined in (17), which governs the backward rate matrix. As in standard diffusion models, our goal is to sample from the time-reversed process, requiring approximation of the (intractable) score. Leveraging its conditional expectation structure (Proposition 1.1), we approximate s_t using a parameterized family $\{(t, x) \mapsto s_t^{\theta}(x)\}_{\theta \in \Theta}$, where θ is optimized via an adapted score-matching objective, defined as the function

$$\mathfrak{L}_{\mathrm{L}^{2}}:\theta\mapsto\int_{0}^{T_{f}}\mathbb{E}\left[\|s^{\theta}_{T_{f}-t}(\overrightarrow{X}_{t})-f_{T_{f}-t}(\overrightarrow{X}_{0},\overrightarrow{X}_{t})\|^{2}\right]\mathrm{d}t.$$
(21)

Another objective function to fit θ can be derived by using the fact that for any $x \in X$, $t \in [0, T_f)$, $\ell \in \{1, \ldots, d\}$, $1 - s_t^{\ell}(x)$ is non-negative. Thus, we introduce the following entropy-based term:

$$\theta \mapsto \int_0^{T_f} \mathbb{E}\left[\sum_{\ell=1}^d (1-s_{T_f-t}^{\theta,\ell})h\left(\frac{1-s_{T_f-t}^\ell}{1-s_{T_f-t}^{\theta,\ell}}\right)(\overrightarrow{X}_t)\right] \mathrm{d}t \;,$$

where $h(a) = a \log(a) - (a - 1)$. Minimizing this function is equivalent to minimizing:

$$\begin{aligned} \mathfrak{L}_{\mathbf{e}} &: \theta \mapsto \int_{0}^{T_{f}} \mathbb{E} \left[\sum_{\ell=1}^{d} \left(-s_{T_{f}-t}^{\theta,\ell}(\overrightarrow{X}_{t}) + \left(s_{T_{f}-t}^{\ell}(\overrightarrow{X}_{t}) - 1\right) \log(1 - s_{T_{f}-t}^{\theta,\ell}(\overrightarrow{X}_{t})) \right) \right] \mathrm{d}t \;. \end{aligned}$$

$$(22)$$

We further derive a discrete denoiser structure in Appendix C.3, rewriting the score function as

$$s_t^{\ell}(x) = \frac{2\alpha_{T_f-t}}{1+\alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}d_t^{\ell}(x)}{1-\alpha_{T_f-t}^2}, \qquad (23)$$

where $d_t^{\ell}(x) = \mathbb{P}(\vec{X}_0^{\ell} \neq x^{\ell} | \vec{X}_{T_f-t} = x)$ serves as a classifier referred to as a *discrete denoiser*. We leverage this structure by reparameterizing our score model as:

$$s_t^{\theta,\ell}(x) = \frac{2\alpha_{T_f-t}}{1+\alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}d_t^{\theta,\ell}(x)}{1-\alpha_{T_f-t}^2} .$$
(24)

As a result, we modify our objective \mathfrak{L}_{L^2} to $\mathfrak{L}_{L^2}^{den}$ to fit the conditional expectations $(d_t)_{t \in [0,T_f)}$ instead of the score functions $(s_t)_{t \in [0,T_f)}$, as follows:

$$\mathfrak{L}_{\mathrm{L}^{2}}^{\mathrm{den}}:\theta\mapsto\int_{0}^{T_{f}}\mathbb{E}\left[\|\sum_{\ell=1}^{d}d_{T_{f}-t}^{\theta,\ell}(\overrightarrow{X}_{t})-\mathbb{1}_{\overrightarrow{X}_{0}^{\ell}}(\overrightarrow{X}_{t}^{\ell})\|^{2}\right]\mathrm{d}t\,,$$
(25)

see Appendix C.4 for more details. This reparameterization moves the approximation from a space of ratios into probability space, which is smoother and more amenable to learning, mitigating the instability of direct score or rate estimation as reported in (Lou et al., 2024). To fit $d_t^{\theta}(x)$ to $d_t(x)$, we introduce an additional cross-entropy loss \mathfrak{L}_{CE} :

$$\begin{split} \mathfrak{L}_{\mathrm{CE}} &: \theta \mapsto \int_{0}^{T_{f}} \mathbb{E} \bigg[\sum_{\ell=1}^{d} -\mathbb{1}_{\overrightarrow{X}_{0}^{\ell} \neq \overrightarrow{X}_{T_{f}-t}^{\ell}} \log d_{t}^{\theta,\ell}(\overrightarrow{X}_{T_{f}-t}^{\ell}) \\ &+ (\mathbb{1}_{\overrightarrow{X}_{0}^{\ell} \neq \overrightarrow{X}_{T_{f}-t}^{\ell}} - 1) \log(1 - d_{t}^{\theta,\ell}(\overrightarrow{X}_{T_{f}-t}^{\ell})) \bigg] \mathrm{d}t \; . \end{split}$$

Based on the previous discussions, we consider a linear combination of the losses $\mathfrak{L}_{L^2}^{den}$, \mathfrak{L}_e , \mathfrak{L}_{CE} , respectively weighted by factors $\varpi_1, \varpi_2, \varpi_3$, which results in the loss \mathfrak{L}_{ϖ} :

$$\mathfrak{L}_{\varpi} = \varpi_1 \mathfrak{L}_{\mathrm{L}^2}^{\mathrm{den}} + \varpi_2 \mathfrak{L}_{\mathrm{e}} + \varpi_3 \mathfrak{L}_{\mathrm{CE}} .$$
 (26)

The expected value of d_t^{ℓ} is given by

$$w_t = \mathbb{E}\left[d_t^\ell(\overrightarrow{X}_{T_f-t})\right] = (1 - \alpha_{T_f-t})/2 , \qquad (27)$$

as detailed in Appendix C.4. Thus, we scale losses \mathcal{L}_{L^2} , \mathcal{L}_{CE} by $1/w_t$, ensuring a more balanced average magnitude across timesteps; see (50) and Figure 4 in Appendix C.4. This leads to the updated loss \mathcal{L}_{ϖ}^w (see (52)). Detailed derivations are provided in Appendix C.4. The final training procedure is outlined in Algorithm 2.

1.5. Generative process

Alike classical continuous diffusion models, exact simulations of the reverse process are not possible and face the same challenges: i) we do not have access to i.i.d. samples from μ_{T_f} , ii) the backward process characteristics depend on the score function of the forward process defined in (17), which is intractable, and iii) we have to discretize the continuous process.

Initialize the backward from the uniform distribution. We show that $(\vec{X}_t)_{t \in [0, T_f]}$ converges geometrically to γ^d , the uniform distribution over X (see Appendix B.2.3 in the supplementary document). This should be put in parallel with diffusion-based models, where the stochastic process at hand, e.g., Ornstein–Uhlenbeck, converges geometrically fast to some Gaussian distribution. The generative model can then be initialized to γ^d rather than μ_{T_f} .

Score approximation. We have access to a score approximation $(s_t^{\theta^*})_{t \in [0,T_f)}$, so the generative model can then be sampled analogously to the backward process, replacing $(s_t)_{t \in [0,T_f)}$ with $(s_t^{\theta^*})_{t \in [0,T_f)}$, leading to the non-homogeneous jump rate and kernel approximating (20):

$$\begin{split} \lambda_t^{\theta^\star}(x) &= \lambda \sum_{\ell=1}^d (1 - s_t^{\theta^\star,\ell}(x)) \;, \\ k_t^{\theta^\star}(x,y) &= \sum_{\ell=1}^d \mathbbm{1}_{y = \varphi^{(\ell)}(x)} \cdot \lambda (1 - s_t^{\theta^\star,\ell}(x)) / \lambda_t^{\theta^\star}(x) \;, \end{split}$$

for $x, y \in X$ and $t \in [0, T_f)$, where we denote by $s_t^{\theta^*, \ell}$ the ℓ -th component of $s_t^{\theta^*}$. For completeness, Algorithm 1 in Appendix C.2 provides the pseudo-code for an ideal, continuous-time approximation of the backward process.

Time discretization. Exact integration of jump rates is infeasible in practice, so we discretize time and approximate the backward rate and kernel using piecewise constant functions. Let $\{t_k\}_{k=0}^K$ be a time grid with step sizes $\tau_k = t_k - t_{k-1}$. Given $X_{t_k}^{\star}$, for $x \neq y \in X$ and $t \in [t_k, t_{k+1})$, we set:

$$\begin{aligned} \overleftarrow{q}_{t}^{\theta^{\star}}(x,y) &= \lambda \sum_{\ell=1}^{d} (1 - s_{t_{k}}^{\theta^{\star},\ell}(\overleftarrow{X}_{t_{k}}^{\star})) \mathbb{1}_{y=\varphi^{(\ell)}(x)} \\ &= (1 - \hat{s}_{t}^{\theta^{\star}}(x,y)) \overrightarrow{q}(x,y) \;, \end{aligned}$$

where $\hat{s}_t^{\theta^{\star}}(x,y) = \sum_{\ell=1}^d s_{t_k}^{\theta^{\star},\ell} (\overleftarrow{X}_{t_k}^{\star}) \mathbb{1}_{y=\varphi^{(\ell)}(x)}.$

This yields a tractable CTMC $(\overleftarrow{X}_t^*)_{t \in [0,T_f]}$, which can be simulated starting from γ^d . Under mild conditions, its final law converges to the target distribution. The associated DMPM sampler is given in Algorithm 3, Appendix C.5, with time-schedule choices listed in Table 2.

Flips sampler. The standard sampler updates one bit at each timestep. To improve parallelism and sample diversity, we propose a flip-schedule where M_{t_k} components are flipped simultaneously at step t_k , based on the probability distribution defined by \hat{k}^{θ} . We consider linear and cosine flip-schedules (Table 3), implemented in Algorithm 4, Appendix C.5.

Denoise-renoise sampler. We also introduce a denoiserenoise sampler based on the discrete denoiser from Equation (23). Inspired by multistep consistency models (Song et al., 2023), this method alternates denoising from $t_0 = 0$ to T_f and re-noising back to t_1 , and so on. The full procedure is detailed in Algorithm 5 and Appendix C.5.

2. Convergence of DMPMs algorithm

This section provides quantitative error estimates between the generated final distribution $\text{Law}(X_{T_f}^{\star})$ and our data distribution μ^{\star} via the Kullback–Leibler divergence KL. To this end, we consider the following assumptions on the parameterized score and the original data distribution:

Assumption 2.1. Let $h(a) := a \log(a) - (a-1)$ for a > 0. There exists $\epsilon \in (0, 1)$ such that

$$\max_{0 \leqslant k \leqslant K} \mathbb{E} \left[\sum_{\ell=1}^{d} (1 - s_{t_k}^{\theta^{\star},\ell}) h\left(\frac{1 - s_{t_k}^{\ell}}{1 - s_{t_k}^{\theta^{\star},\ell}}\right) (\overrightarrow{X}_{T_f - t_k}) \right] \leqslant \epsilon .$$
(28)

Note that Assumption 2.1 is induced by the entropic term

 \mathcal{L}_{e} defined in (22) of the loss function we consider in practice. This condition naturally appears as we bound the KL divergence of the path probability measures corresponding to the approximate score s^{θ^*} and the ideal one *s* respectively. Indeed, we prove a Girsanov type theorem which provide an explicit expression of the density between these two measures in Theorem F.12 in the supplement F.2.1. While standard Girsanov theorem for diffusion implies an L²-type approximation error condition for generative models (see, e.g., Conforti et al., 2025; Lee et al., 2023; Chen et al., 2022a), our result naturally involve the entropic-type condition (28) due to the discrete structure of our noising process.

Assumption 2.2. The data distribution does not admit any zero-value, *i.e.*, $\mu^*(x) \in (0, 1)$ for any $x \in X$.

Assumption 2.2 implies that the data distribution has the finite Fisher-like information

$$\beta_{\gamma^d}(\mu^\star) := \mathbb{E}\left[\sum_{\ell=1}^d h\left(\mathrm{e}^{g(\overrightarrow{X}_0) - g(\varphi^{(\ell)}(\overrightarrow{X}_0))}\right)\right] < +\infty ,$$
(29)

with $g := -\log(d\mu^*/d\gamma^d)$. Note that Assumption 2.2 is put in parallel with the finite relative Fisher information condition provided by Conforti et al. (2025). However, Assumption 2.2 is much simpler as the state space is finite, and the function h is only infinite if μ^* has not full support.

We are now ready to state the error's bound of the generated data using DMPMs given in Algorithm 3, noting that every measure on the hypercube possesses finite entropy.

Theorem 2.3. Under Assumption 2.1 and Assumption 2.2, the following bound holds

$$\operatorname{KL}(\mu^{\star}|\operatorname{Law}(\overline{X}_{T_{f}}^{\star})) \leqslant e^{-4\lambda T_{f}} \operatorname{KL}(\mu^{\star}|\gamma^{d}) + \lambda \tau \beta_{\gamma^{d}}(\mu^{\star}) + \lambda \epsilon T_{f} , \qquad (30)$$

with $\tau := \max\{\tau_k, k = 1, ..., K\}.$

Theorem 2.3 is one of our distinguishing results, which guarantees the convergence of DMPMs algorithm, and makes it stronger than other algorithms built before for discrete target distribution.

The term ϵT_f in (30) appears because the score function s_t is replaced in the discretization by its approximation $s_t^{\theta^*}$ satisfying Assumption 2.1. The term $e^{-4\lambda T_f} \text{KL}(\mu^*|\gamma^d)$ represents the initialization error, as our backward dynamic starts at γ^d instead of μ_{T_f} . Finally, the term $\tau \beta_{\gamma^d}(\mu^*)$ means that the data distribution μ^* cannot be peculiar, in the sense that μ^* does not admit any zero-value. The detailed proof of Theorem 2.3 is given in the supplementary material F.5.1.

Following Conforti et al. (2025, Theorem 3), a tighter

bound on the sampling error, one that scales logarithmically rather than linearly with the discrete Fisher information, is achieved with an appropriate sequence of step sizes.

Theorem 2.4. Let $c \in (0, 1/2]$ and $T_f \ge 1 + 2c$. Suppose Assumption 2.1 and Assumption 2.2 hold, and $L = d^{-1}\beta_{\gamma^d}(\mu^*) \ge 2$. Choose an exponentially decreasing stepsizes, i.e., $\tau_{k+1} = c \min \{\max\{T_f - t_k, a\}, 1\}$ for k < K, with a = 1/L, then we have that

$$\operatorname{KL}(\mu^{\star}|\operatorname{Law}(\overleftarrow{X}_{T_{f}}^{\star})) \lesssim e^{-4\lambda T_{f}} \operatorname{KL}(\mu^{\star}|\gamma^{d}) + \lambda \epsilon T_{f} + \lambda cd[1 + \log(L)].$$
(31)

The proof of Theorem 2.4 benefits from the choice of the step-size's scheme and is postponed to Appendix F.5.2. We deduce the following complexity result for DMPMs to achieve an $\epsilon > 0$ discretization error.

Corollary 2.5. *Consider the sequence of step-size as in Theorem 2.4 and suppose Assumption 2.1 and Assumption 2.2 hold. Choosing*

$$c = \frac{\epsilon}{\lambda d[1 + \log(L)]} \quad and \quad T_f = \frac{1}{4\lambda} \log \frac{\mathrm{KL}(\mu^* | \gamma^d)}{\epsilon},$$
(32)

we get

$$K \lesssim d[1 + \log(L)] \left[\log(\mathrm{KL}(\mu^* | \gamma^d) / \epsilon) + \lambda \log(L) \right] / \epsilon$$

and makes the approximation error $\tilde{O}(\epsilon \log(\mathrm{KL}(\mu^*|\gamma^d)))$, where the notation \tilde{O} means that logarithmic factors of d, ϵ have been dropped.

The proof of Corollary 2.5 is provided in Appendix F.5.3.

In our next result, we get rid of Assumption 2.2 using an early stopping strategy.

Theorem 2.6. Under Assumption 2.1, for any $\eta \in (0, T_f)$, let $c \in (0, 1/2]$ and $T_f - \eta \ge 1 + 2c$. Set $L = d^{-1}\beta_{\gamma^d}(\mu_\eta)$ and assume that $L \ge 2$. Choose the constant and exponentially decreasing sequence of step-size, i.e., satisfying $\tau_{k+1} = c \min \{\max\{T_f - \eta - t_k, 1/L\}, 1\}$ for k < Kand the associated discrete time scheme $\{t_k\}_{k=0}^K$ such that $t_0 = 0$ and $t_K = T_f - \eta$. Then, the following bound holds

$$\operatorname{KL}(\mu_{\eta}|\operatorname{Law}(\overleftarrow{X}_{T_{f}-\eta}^{\star})) \lesssim d\eta^{-1} \mathrm{e}^{-4\lambda(T_{f}-\eta)} + \lambda\epsilon(T_{f}-\eta) + \lambda cd[1 + \log(\eta^{-1})].$$
(33)

Theorem 2.6 is a consequence of Theorem 2.4 when the backward dynamic stops early at μ_{η} instead of μ^{\star} . We benefit from the structure of μ_{η} to obtain a bound growing linearly with dimension, which is advantageous for high-dimensional sampling. The full proof is deferred to Appendix F.6.1. It is worth noting that (16) ensures that μ_{η} is always positive for any $\eta \in (0, T_f)$, thus the Fisher-like

information $\beta_{\gamma^d}(\mu_\eta)$ is always finite. As a result, Assumption 2.2 is no longer required. To obtain then a complexity bound for DMPMs on its discretization error without Assumption 2.2, we bound in our next result, the total variation distance between μ^* and μ_η for $\eta > 0$.

Proposition 2.7. For any $\eta \in (0, \max\{T_f, \frac{1}{\lambda}\})$, the following holds

$$\|\mu_{\eta} - \mu^{\star}\|_{\mathrm{TV}} \leq 2 - 2(1 - \lambda \eta)^{d}$$
. (34)

The proof of Proposition 2.7 is provided in the supplement F.6.2. Combining Theorem 2.6 and Proposition 2.7, we deduce that

Corollary 2.8. Consider the sequence of step-size as in Theorem 2.6 and let Assumption 2.1 hold. Choosing

$$\eta = \frac{1 - (1 - \epsilon)^{1/d}}{\lambda} , \quad c = \frac{\epsilon^2}{\lambda d [1 + \log(\eta^{-1})]} ,$$

$$T_f = \eta + \frac{1}{4\lambda} \log \frac{d}{\eta \epsilon^2} ,$$
 (35)

implies that

$$K \lesssim d[1 + \log(\lambda d/\epsilon)][\log(d/\epsilon^2) + (\lambda + 1)\log(\lambda d/\epsilon)]/\epsilon^2 ,$$

and the following bound holds

$$\|\mu^{\star} - \operatorname{Law}(\overleftarrow{X}_{T_f-\eta}^{\star})\|_{\mathrm{TV}} \lesssim \epsilon + \sqrt{\lambda \epsilon (T_f-\eta)}$$

The proof of Corollary 2.8 is given in Appendix F.6.3.

3. Existing works on diffusion-based generative models for discrete data

We briefly review existing approaches to discrete generative modeling based on diffusion processes. Additional discussion is provided in Appendix A.

A first class of methods maps discrete variables into continuous spaces, enabling the use of classical diffusion machinery (Dieleman et al., 2022; Chen et al., 2022b; Richemond et al., 2022), but struggle to scale in dimensions and lacks theoretical guarantees.

Other methods, such as Argmax Flows and Multinomial Diffusion (Hoogeboom et al., 2021), operate directly in discrete spaces and use categorical noise models or argmax transformations to handle discrete tokens, but can impose considerable computational overhead.

More recently, CTMC-based frameworks were introduced (Campbell et al., 2022), upon which flow matching techniques were adapted to the discrete domains (Gat et al., 2024; Campbell et al., 2024), using conditionally constructed rate matrices built ad-hoc, contrasting with our

principled time-reversal derivation.

Holderrieth et al. (2024) proposed a general framework for generator matching over arbitrary Markov processes, assuming access to a conditional interpolating distribution, leaving the choice of interpolating process and loss function to problem-specific adaptation.

Alternative directions include masked diffusion models (Austin et al., 2021) and stochastic integral formulations (Ren et al., 2024), aiming to balance tractability and theoretical soundness. Shi et al. (2024); Sahoo et al. (2024) propose efficient training via absorbing-state kernels, optimizing model parameterization and loss. Meanwhile, Lou et al. (2024) model the score function as a density ratio rather than a discrete denoiser. This yields an entropic regularization loss equivalent to (22), missing the L^2 projection and cross-entropy terms applied to the discrete denoiser employed in our formulation.

Concurrently with our work, Bach & Saremi (2025) propose a discrete analogue of Gaussian smoothing, entirely forgoing the continuous-time framework. Their denoising-based method offers a static-noise alternative to CTMCs, yielding Langevin-type sampling dynamics on the hypercube.

Theoretical results for discrete generative models are much scarcer than for their continuous counterparts. To the best of our knowledge, only Campbell et al. (2022); Ren et al. (2024) provide theoretical guarantees, and these rely on significantly stronger assumptions than those used in our work.

Regarding Theorem 1 in Campbell et al. (2022), we note that our approach does not require an L^{∞} -bound on the score approximation error, but only an L²-bound. Moreover, we do not impose any assumptions on the marginal density of the forward process (cf. Assumption 2), but only on the data distribution itself. We also avoid placing assumptions on the backward transition rates (cf. Assumption 3). In contrast to Campbell et al. (2022), we impose no conditions on the rates or densities of the forward and backward processes. Additionally, their discretization error scales linearly with the time horizon, whereas our bounds do not incur such a cost.

Concerning Ren et al. (2024), and in particular Theorems 4.7 and 4.9, we highlight that we make no assumptions on the score function. In their Assumption 4.4, a time-dependent L^{∞} -bound of the form $|s_t(x)| \leq 1 \vee t^{-1}$ is imposed on the true score s_t , and a time-uniform L^{∞} -bound is assumed for the learned score s_t^{θ} . As in Campbell et al. (2022), these assumptions are arguably unnatural, as they are not placed directly on the data distribution but on more complex transformations of it. Furthermore, Assumption 4.5 in Ren et al. (2024) imposes a quantitative form of Lipschitz continuity on s_t , and as the authors themselves state, "Assumption 4.5 corresponds to the Lipschitz continuity of the score function." However, it is now well known that such an assumption is unnecessary in the continuous setting.

4. Experiments

The full experimental details are available in Appendix D. We evaluate our Discrete Markov Probabilistic Model (DMPM) on two datasets. The first is a low-dimensional synthetic *sawtooth* dataset, with dimension $4 \le d \le 16$. The second is binarized MNIST, with $d = 32 \times 32$. We explore various design choices, and compare DMPM against MD4 (masked diffusion) (Shi et al., 2024) and DFM (discrete flow matching) (Gat et al., 2024), two state-of-the-art discrete generative approaches.

4.1. Experiments on Small-Dimensional Bernoulli Data

We study a discrete data distribution p such that each component of $X = (X_i)_{i=1}^d \sim p$ is independently distributed as Bernoulli (p_i) . The map $i \mapsto p_i$ forms a sawtooth pattern (see Figure 6). We evaluate performance using a custom Sliced Wasserstein Distance (SWD) between the learned and true distributions (see Appendix D.3). Indeed, the state space size 2^d can get too big for traditional histogram-based metrics like KL divergence or Hellinger distance.



Figure 1. Comparison of time-schedules (*cosine*, *linear*, *quadratic*) and time horizon ($T_f = 3$ vs. $T_f = 10$).

Time horizon and time-schedule. We study the impact of the time horizon T_f and various time-schedules—uniform, quadratic, and cosine (see Table 2)—on model performance. Figure 1 presents results for a model trained with the basic \mathfrak{L}_{L^2} loss on data with d = 16, evaluated across multiple reverse step counts. The cosine schedule with $T_f = 3$ yields the best performance in terms of sliced Wasserstein distance (SWD), outperforming other schedules and longer horizons, while also requiring fewer reverse steps. This indicates that $T_f = 3$ is sufficient to reach near-uniformity during forward diffusion, without excessive transitions to uniform states. We adopt this configuration for all subsequent experiments.

Comparison with state-of-the-art methods. We compare

d	4	8	12	16
DFM	6.102	8.864	5.019	8.302
MD4	9.376	7.670	4.045	8.037
DMPM	3.174	3.308	2.342	2.515

Figure 2. SWD \downarrow , in 1e-3, for DMPM, MD4, and DFM across data dimension *d*. Selected the best result with #steps $2 \leq K \leq 200$ for each method.



Figure 3. FID \downarrow on MNIST, linear vs. constant flip-schedules scaled for *d* total bit flips, with various loss configurations.

DMPM (with cosine schedule, $T_f = 3$, and \mathcal{L}_{L^2} loss) against MD4 and Discrete Flow Matching (DFM). Figure 2 shows SWD scores across varying data dimensions *d*. DMPM consistently outperforms both baselines, achieving superior results with significantly fewer reverse steps (typically 30 vs 100), highlighting its sampling efficiency.

4.2. Experiments on Higher-Dimensional Binary MNIST

DMPM sampler and flip-schedule. We evaluate the DMPM sampler using constant and linear flip-schedules $\{M_{t_k}\}_{k=1}^{K}$. Empirically, performance is optimal when the total number of flipped bits matches the input dimension, i.e., $\sum_{k=1}^{K} M_{t_k} = d$. We scale each flip-schedule accordingly. Figure 3 illustrates that performance remains stable across different values of K as long as the total number of flips is held constant, allowing for significant speedups by reducing the number of reverse steps and thus network calls.

Using a model trained with the weighted loss $\mathfrak{L}^{w}_{1/3,1/3,1/3}$ and a linear flip-schedule, we achieve a best FID of 4.77 using only 25 network calls. The linear schedule consistently outperforms the constant variant; flipping fewer bits early helps guide the model toward more coherent samples, similarly as what is reported for MD4 (Shi et al., 2024).

Loss function configuration. Among the losses we tested, the balanced form $\mathfrak{L}_{1,,1,,1}^w$ consistently yields the best results. The weighting factor w helps normalize the scale of the ℓ_2 , cross-entropy, and KL components at each timestep, improving overall synergy. Figure 5 illustrates these effects. Nonetheless, simpler variants such as \mathfrak{L}_{L^2} and $\mathfrak{L}_{L^2}^w$ already achieve near-optimal performance in many settings.

Method		10	25	50	100	200	500
DFM	$\begin{array}{c} FID \\ F_1^{\rm dc} \end{array}$	227.55 0.00	156.26 0.00	88.93 0.01	39.62 0.14	16.26 0.41	7.34 0.68
MD4	$\begin{array}{c} FID \\ F_1^{\rm dc} \end{array}$	97.97 0.04	33.50 0.29	14.06 0.57	6.83 0.76	4.48 0.83	<i>3.43</i> 0.86
DMPM _{flips}	$\begin{array}{c} FID \\ F_1^{\rm dc} \end{array}$	16.30 0.64	9.98 0.92	11.07 0.93	9.07 0.93	7.80 0.93	10.84 0.70
DMPM _{denoise}	$\begin{array}{c} FID \\ F_1^{\rm dc} \end{array}$	78.20 0.13	20.94 0.67	8.62 0.87	<u>3.98</u> <u>0.96</u>	2.89 1.00	4.36 1.00

Table 1. FID \downarrow (first row of each method) and $F_1^{dc} \uparrow$ (second row) on MNIST for various total reverse steps. We highlight the best result in **bold**, the 2nd best in *italics*, and underline the 3rd best.

Denoise-renoise sampler and comparison with state-ofthe-art. We further exploit the discrete denoiser structure through a denoise-renoise sampler (Algorithm 5, Appendix C.4), which alternates single-step denoising and re-noising in a multistep loop. This approach leverages the model's learned transitions more effectively, leading to notable gains in sample quality.

We compare DMPM, trained with the balanced loss $\mathfrak{L}_{1,1,1}^w$, under two sampling strategies—denoise-renoise and linear flip-schedule—against state-of-the-art baselines: MD4 (masked diffusion) and DFM (discrete flow matching). Results are reported for varying numbers of reverse steps K, using both Fréchet Inception Distance (FID) and the F_1^{dc} metric, a harmonic mean of coverage and density (Naeem et al., 2020). While FID captures overall realism, F_1^{dc} reflects local distributional fidelity; full details are in Appendix D.3.

As shown in Table 1, both DMPM variants (rows 3 and 4) consistently outperform the baselines. At K = 200, DMPM (denoise-renoise) achieves the best FID (2.89) and perfect F_1^{dc} (1.00), surpassing MD4 (4.48) and DFM (16.26). Even with K = 50, it maintains strong results (FID 8.62, F_1^{dc} 0.87). Similarly, DMPM with flip-schedule achieves FID below 10 and F_1^{dc} above 0.90 at K = 25, demonstrating excellent efficiency with minimal network calls.

Sample grids illustrating visual quality are shown in Figure 7 and Figure 8.

4.3. Conclusions

Our experiments show that DMPM consistently matches or surpasses state-of-the-art performance on both low- and high-dimensional discrete datasets. On binarized MNIST, it achieves better FID and F_1^{dc} than competing methods, with fewer network calls. These gains stem from our principled reparameterization of the score function as a denoiser, and a stable, well-structured training objective. Together, they yield a scalable and theoretically sound framework for discrete generative modeling.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Acknowledgements

The work of L.T.N.P. was supported by the École Doctorale de Mathématiques Hadamard (EDMH). The work of A.O. was funded by the European Union (ERC-2022-SYG-OCEAN-101071601). The work of D.S was funded by the European Union (ERC, Dynasty, 101039676). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

This work was granted access to the HPC resources of IDRIS under the allocation 2025-AD011015323R1 made by GENCI.

References

- Austin, J., Johnson, D. D., Ho, J., Tarlow, D., and Van Den Berg, R. Structured denoising diffusion models in discrete state-spaces. *Advances in Neural Information Processing Systems*, 34:17981–17993, 2021.
- Bach, F. and Saremi, S. Sampling binary data by denoising through score functions, 2025. URL https://arxiv.org/abs/2502.00557.
- Bakry, D., Gentil, I., Ledoux, M., et al. *Analysis and geometry of Markov diffusion operators*, volume 103. Springer, 2014.
- Bar-Tal, O., Chefer, H., Tov, O., Herrmann, C., Paiss, R., Zada, S., Ephrat, A., Hur, J., Liu, G., Raj, A., et al. Lumiere: A space-time diffusion model for video generation. In *SIGGRAPH Asia 2024 Conference Papers*, pp. 1–11, 2024.
- Bertazzi, A., Oliviero-Durmus, A., Shariatian, D., Simsekli, U., and Moulines, E. Piecewise deterministic generative models. arXiv preprint arXiv:2407.19448, 2024.
- Campbell, A., Benton, J., De Bortoli, V., Rainforth, T., Deligiannidis, G., and Doucet, A. A continuous time framework for discrete denoising models. *Advances in Neural Information Processing Systems*, 35:28266–28279, 2022.
- Campbell, A., Yim, J., Barzilay, R., Rainforth, T., and Jaakkola, T. Generative flows on discrete state-spaces: Enabling multimodal flows with applications to protein

co-design, 2024. URL https://arxiv.org/abs/ 2402.04997.

- Chen, N., Zhang, Y., Zen, H., Weiss, R. J., Norouzi, M., and Chan, W. Wavegrad: Estimating gradients for waveform generation. *arXiv preprint arXiv:2009.00713*, 2020.
- Chen, S., Chewi, S., Li, J., Li, Y., Salim, A., and Zhang, A. R. Sampling is as easy as learning the score: theory for diffusion models with minimal data assumptions. arXiv preprint arXiv:2209.11215, 2022a.
- Chen, T., Zhang, R., and Hinton, G. Analog bits: Generating discrete data using diffusion models with selfconditioning. *arXiv preprint arXiv:2208.04202*, 2022b.
- Conforti, G. and Léonard, C. Time reversal of markov processes with jumps under a finite entropy condition. *Stochastic Processes and their Applications*, 144:85–124, 2022.
- Conforti, G., Durmus, A., and Silveri, M. G. Kl convergence guarantees for score diffusion models under minimal data assumptions. *SIAM Journal on Mathematics of Data Science*, 7(1):86–109, 2025.
- Delatte, P. Lecture 5. hahn-banach extension theorem. *Lecture note*, 2022.
- Dieleman, S., Sartran, L., Roshannai, A., Savinov, N., Ganin, Y., Richemond, P. H., Doucet, A., Strudel, R., Dyer, C., Durkan, C., et al. Continuous diffusion for categorical data. arXiv preprint arXiv:2211.15089, 2022.
- Gat, I., Remez, T., Shaul, N., Kreuk, F., Chen, R. T., Synnaeve, G., Adi, Y., and Lipman, Y. Discrete flow matching. Advances in Neural Information Processing Systems, 37:133345–133385, 2024.
- Gavrilyuk, I., Makarov, V., and Vasylyk, V. *Exponentially* convergent algorithms for abstract differential equations. Springer Science & Business Media, 2011.
- Gillespie, D. T. Stochastic simulation of chemical kinetics. Annual Review of Physical Chemistry, 58(Volume 58, 2007):35-55, 2007. ISSN 1545-1593. doi: https: //doi.org/10.1146/annurev.physchem.58.032806.104637. URL https://www.annualreviews.org/ content/journals/10.1146/annurev. physchem.58.032806.104637.
- Ho, J., Salimans, T., Gritsenko, A., Chan, W., Norouzi, M., and Fleet, D. J. Video diffusion models. *Advances in Neural Information Processing Systems*, 35:8633–8646, 2022.
- Holderrieth, P., Havasi, M., Yim, J., Shaul, N., Gat, I., Jaakkola, T., Karrer, B., Chen, R. T., and Lipman, Y.

Generator matching: Generative modeling with arbitrary markov processes. *arXiv preprint arXiv:2410.20587*, 2024.

- Hoogeboom, E., Nielsen, D., Jaini, P., Forré, P., and Welling, M. Argmax flows and multinomial diffusion: Learning categorical distributions. *Advances in Neural Information Processing Systems*, 34:12454–12465, 2021.
- Ikeda, N. and Watanabe, S. Stochastic differential equations and diffusion processes, volume 24. Elsevier, 2014.
- Karras, T., Aittala, M., Aila, T., and Laine, S. Elucidating the design space of diffusion-based generative models. *Advances in neural information processing systems*, 35: 26565–26577, 2022.
- Kong, Z., Ping, W., Huang, J., Zhao, K., and Catanzaro, B. Diffwave: A versatile diffusion model for audio synthesis. *arXiv preprint arXiv:2009.09761*, 2020.
- Kynkäänniemi, T., Karras, T., Laine, S., Lehtinen, J., and Aila, T. Improved precision and recall metric for assessing generative models. *Advances in neural information processing systems*, 32, 2019.
- Laszuk, D. Python implementation of empirical mode decomposition algorithm. https://github.com/ laszukdawid/PyEMD, 2017.
- Lee, H., Lu, J., and Tan, Y. Convergence of score-based generative modeling for general data distributions. In *International Conference on Algorithmic Learning Theory*, pp. 946–985. PMLR, 2023.
- Léonard, C. Orlicz spaces. Unpublished notes, available on the Internet, 2007.
- Léonard, C. Girsanov theory under a finite entropy condition. In Séminaire de Probabilités XLIV, pp. 429–465. Springer, 2012.
- Liggett, T. M. *Continuous time Markov processes: an introduction*, volume 113. American Mathematical Soc., 2010.
- Lou, A., Meng, C., and Ermon, S. Discrete diffusion language modeling by estimating the ratios of the data distribution, 2024. URL https://openreview.net/ forum?id=71mqtQdKB9.
- Méliot, P.-L. Convergence of random variables and large deviations. *Lecture note*, 2025.
- Mozumder, M. S. U. Some derivations on levy and jump processes. GANIT: Journal of Bangladesh Mathematical Society, 29:11–22, 2009.

- Naeem, M. F., Oh, S. J., Uh, Y., Choi, Y., and Yoo, J. Reliable fidelity and diversity metrics for generative models. In *International conference on machine learning*, pp. 7176–7185. PMLR, 2020.
- Nichol, A. Q. and Dhariwal, P. Improved denoising diffusion probabilistic models. In *International conference on machine learning*, pp. 8162–8171. PMLR, 2021.
- Nutz, M. Introduction to entropic optimal transport. *Lecture notes, Columbia University*, 2021.
- Owen, A. B. Monte carlo theory, methods and examples; 2013. statweb, 2021.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., and Chen, M. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- Rao, M. M. and Ren, Z. D. Theory of orlicz spaces. (*No Title*), 1991.
- Ren, Y., Chen, H., Rotskoff, G. M., and Ying, L. How discrete and continuous diffusion meet: Comprehensive analysis of discrete diffusion models via a stochastic integral framework. arXiv preprint arXiv:2410.03601, 2024.
- Richemond, P. H., Dieleman, S., and Doucet, A. Categorical sdes with simplex diffusion. *arXiv preprint arXiv:2210.14784*, 2022.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., and Ommer, B. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- Saharia, C., Chan, W., Saxena, S., Li, L., Whang, J., Denton, E. L., Ghasemipour, K., Gontijo Lopes, R., Karagol Ayan, B., Salimans, T., et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35: 36479–36494, 2022.
- Sahoo, S. S., Arriola, M., Schiff, Y., Gokaslan, A., Marroquin, E., Chiu, J. T., Rush, A., and Kuleshov, V. Simple and effective masked diffusion language models, 2024. URL https://arxiv.org/abs/2406.07524.
- Shariatian, D., Simsekli, U., and Durmus, A. Denoising Lévy probabilistic models. *arXiv preprint arXiv:2407.18609*, 2024.
- Shi, J., Han, K., Wang, Z., Doucet, A., and Titsias, M. Simplified and generalized masked diffusion for discrete data. Advances in neural information processing systems, 37:103131–103167, 2024.

- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., and Ganguli, S. Deep unsupervised learning using nonequilibrium thermodynamics. In *International conference on machine learning*, pp. 2256–2265. pmlr, 2015.
- Song, Y., Sohl-Dickstein, J., Kingma, D. P., Kumar, A., Ermon, S., and Poole, B. Score-based generative modeling through stochastic differential equations, 2021. URL https://arxiv.org/abs/2011.13456.
- Song, Y., Dhariwal, P., Chen, M., and Sutskever, I. Consistency models, 2023. URL https://arxiv.org/ abs/2303.01469.
- Touzi, N. Optimal stochastic control, stochastic target problems, and backward SDE, volume 29. Springer Science & Business Media, 2012.
- Villegas, R., Babaeizadeh, M., Kindermans, P.-J., Moraldo, H., Zhang, H., Saffar, M. T., Castro, S., Kunze, J., and Erhan, D. Phenaki: Variable length video generation from open domain textual descriptions. In *International Conference on Learning Representations*, 2022.
- Yoon, E. B., Park, K., Kim, S., and Lim, S. Score-based generative models with Lévy processes. Advances in Neural Information Processing Systems, 36:40694–40707, 2023.
- Zitkovic, G. Uniform integrability. *course notes: Theorem if Probability II, Fall 2015: Lecture*, 12, 2015.

A. Existing works on diffusion-based generative models for discrete data

This section provides details of the recent researches on discrete generative models.

Embedding discrete structure in the continuous space. To keep the benefits of continuous representations, Dieleman et al. (2022) and Chen et al. (2022b) mapped discrete structures into Euclidean space, while Richemond et al. (2022) placed them into the simplex, all while continuing to use forward continuous diffusion models. In particular, Dieleman et al. (2022) proposed a continuous diffusion model for categorical data, which has some advantages over autoregressive models, such as the ability to perform arbitrary infilling and a more flexible sampling process. However, this method comes with an expensive training cost and lacks of strong theoretical guarantees.

Argmax flows and Multinomial Diffusion. Hoogeboom et al. (2021) introduced two new generative models, Argmax Flows and Multinomial Diffusion, to handle categorical data like text and image segmentation. Argmax Flows connect discrete data with continuous models by using an argmax function combined with a probabilistic inverse, making categorical distributions easy-learning. Multinomial Diffusion process uses a categorical distribution to add noise to discrete data and then trains a model to reverse the process. However, both Argmax Flows and Multinomial Diffusion have some limitations: computational costs increase due to additional steps , and the theoretical guarantee is missing.

Designing the flow processes over the discrete state space. Campbell et al. (2022) introduced the first complete continuous-time framework for denoising diffusion models applied to discrete data. They used CTMCs to model the forward noising process and its time-reversal dynamics. While the core idea is similar to ours, their approach is more complex because their method consider generic CTMC and is not specialized to the noising process that we consider. As a result, their method essentially boils down learning density ratios which can be computationally demanding and fail to offer efficient approximation in high dimensions. They also added a correction step to bring the sample distribution closer to the desired one, which increased the practical training cost Gat et al. (2024). By focusing on the random-walk CTMC on X, we were able to provide a discrete counterpart to the score function that is learn in continuous diffusion models and also to establish strong convergence guarantees for our method. Campbell et al. (2024) further extended this line of work by adapting flow-matching techniques to discrete domains using conditionally defined rate matrices. However, their method does not derive the reverse process from a time-reversal principle, and thus relies on hand-crafted dynamics with limited theoretical justification.

Generator Matching. Another recent approach to handle discrete data is generative modeling with arbitrary Markov processes using generator matching, introduced by Holderrieth et al. (2024). In this approach, the authors design an appropriate Markov process that transforms a simple distribution into the desired data one using a generator, which can be efficiently trained with a neural network. This method is quite flexible and can be applied to different state spaces, especially in discrete settings. However, this method being very generic suffer from the same drawback as Campbell et al. (2022).

Masked diffusion models. One important step toward more advanced models is the "masked" diffusion process, a discrete diffusion approach first introduced by Austin et al. (2021). Recently, Shi et al. (2024) looked into this model further, simplifying its training objective by expressing it as a signal-to-noise ratio, which helps highlight some useful features. However, despite these improvements, the model still lacks theoretical guarantees. Sahoo et al. (2024) improved upon this direction by leveraging the structure of the absorbing kernel and refining the bridge-based reverse process, leading to more efficient optimization. The model's reliance on absorbing-state approximations and heuristic training objectives limits its theoretical grounding.

Direct score parameterization. Lou et al. (2024) propose to learn the score function directly as a density ratio rather than a denoising map. This formulation leads to a single entropic regularization loss equivalent to (22) in our paper. However, it misses our discrete denoiser decomposition we use and the associated L^2 projection and cross-entropy terms, which improve training stability.

Discrete Diffusion Models via a Stochastic Integral Framework. Ren et al. (2024) introduced a new way to analyze discrete diffusion models via Lévy-type stochastic integrals and expanded Poisson random measures. Specifically, they established the stochastic integral expressions of the noising and denoising processes for the categorical data. They provided a unified error analysis framework and showed the first error bound for their algorithms in KL divergence. However, their results rely on strong assumptions in contrast to our results. Besides, our bounds are simpler and better, in particular with respect to the time horizon.

Denoising without time dynamics. Concurrently with our work, Bach & Saremi (2025) propose a discrete denoising model

that avoids continuous-time formulations altogether. Their method treats Bernoulli corruption as an analogue of Gaussian smoothing, enabling a Langevin-style sampler on the hypercube. This approach is complementary to CTMC-based methods and their dynamic interpretability.

Our paper takes a step toward bridging these gaps. By clearly describing the forward Markov process, we can express the score function as a conditional expectation, which helps us avoid the costly signal-to-noise ratio training used in Shi et al. (2024). This way, we not only offer a simpler and more affordable training approach, but also provide solid theoretical guarantees for our models in practice.

B. Interpretation of DMPMs

B.1. The simple case $X = \{0, 1\}$

We start by explicitly constructing our forward process $(\vec{X}_t)_{t \in [0,T_f]}$ starting from $\vec{X}_0 \sim \mu^*$. Consider the fixed jump times $(T_i)_{i \in \{1,...,N\}} | N \stackrel{\text{iid}}{\sim} \text{Unif}([0,T_f])$ of a Poisson process over $[0,T_f]$ where $N \sim \text{Pn}(\lambda T_f)$ is the number of jumps, and $\lambda > 0$ is a prescribed jump rate. Without loss of generality, we assume that $0 = T_0 \leq T_1 < \ldots < T_N$. We define recursively $(\vec{X}_t)_{t \in [0,T_f]}$ over $(T_i, T_{i+1}]$. Suppose that \vec{X}_{T_i} has been defined we set $\vec{X}_t = \vec{X}_{T_i}$ for any $t \in (T_i, T_{i+1})$ and $\vec{X}_{T_{i+1}} = 1 - \vec{X}_{T_i}$. It is well known that $(\vec{X}_t)_{t \in [0,T_f]}$ is a Markov jump process (Owen, 2021, Section 6) with generator \vec{q}^1 defined for any $x, y \in X$ as

$$\overrightarrow{q}^{1}(x,y) := \begin{cases} \lambda , & \text{if } y \neq x ,\\ -\lambda , & \text{otherwise }. \end{cases}$$
(36)

The transition probability matrix $\mathbb{P}(\overrightarrow{X}_t = y | \overrightarrow{X}_0 = x) = \overrightarrow{p}_t^1(x, y)$, for $x, y \in \mathsf{X}, 0 \leq t \leq T_f$, is known to be

$$\overrightarrow{p}_{t}^{1}(x,y) = \begin{cases} \frac{1}{2} + \frac{1}{2}e^{-2\lambda t}, & \text{if } x = y, \\ \frac{1}{2} - \frac{1}{2}e^{-2\lambda t}, & \text{otherwise}. \end{cases}$$
(37)

Detailed calculation of the transition probability in (6). Based on the Kolmogorov equation, the transition matrix \overrightarrow{p}_t^1 for $0 \leq t \leq T_f$ admits the following formula

$$\overrightarrow{p}_t^1 = \mathrm{e}^{t \overrightarrow{q}^1}$$

where \vec{q}^{1} is define in (5). Clearly, the generator \vec{q}^{1} admits two eigenvalues 0 and -2λ associated with the eigenvectors $\begin{pmatrix} 1 & 1 \end{pmatrix}^{T}$ and $\begin{pmatrix} 1 & -1 \end{pmatrix}^{T}$ respectively. Thus we can diagonalize \vec{q}^{1} as

$$\overrightarrow{q}^{1} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} 0 & 0 \\ 0 & -2\lambda \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}^{-1} ,$$

and the transition matrix \overrightarrow{p}_t^1 follows

$$\vec{p}_{t}^{1} = e^{t\vec{q}^{1}} = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & e^{-2\lambda t} \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}^{-1} = \frac{1}{2} \begin{pmatrix} 1 + e^{-2\lambda t} & 1 - e^{-2\lambda t} \\ 1 - e^{-2\lambda t} & 1 + e^{-2\lambda t} \end{pmatrix}.$$

B.2. General state space $X = \{0, 1\}^d$

B.2.1. FORWARD PROCESS

We consider the jump times $(T_i)_{i \in \{1,...,N\}} | N \stackrel{\text{iid}}{\sim} \operatorname{Unif}([0, T_f])$ of a Poisson process over $[0, T_f]$ where $N \sim \operatorname{Pn}(\lambda T_f)$ is the number of jump. Without loss of generality, we suppose that $T_0 = 0 \leq T_1 < \ldots < T_N$. We define recursively $(\overrightarrow{X}_t)_{t \in [0, T_f]}$ over $(T_i, T_{i+1}]$ as follows. Suppose \overrightarrow{X}_{T_i} has been defined. We set $\overrightarrow{X}_t = \overrightarrow{X}_{T_i}$ for $t \in (T_i, T_{i+1})$, and finally, set $\overrightarrow{X}_{T_{i+1}}^{\ell_i} = 1 - \overrightarrow{X}_{T_i}^{\ell_i}$, where $\ell_i \sim \operatorname{Unif}(\{1, \ldots, d\})$, with ℓ_i independent from the past, and $\overrightarrow{X}_{T_{i+1}}^j = \overrightarrow{X}_{T_i}^j$ for $j \neq \ell_i$. The associated generator matrix is given in (14). We now seek to obtain the associated transition matrix.

Proof of (15). We start with a note that the generator matrix \vec{q} can be expressed as a sum of matrices \vec{q}^{ℓ} as follows

$$\overrightarrow{q} = \sum_{\ell=1}^{d} \overrightarrow{q}^{\ell}, \quad \text{with } \overrightarrow{q}^{\ell}(x, y) = \begin{cases} \lambda, & \text{if } x^{i} = y^{i} \text{ for } i \neq \ell \text{ and } x^{\ell} \neq y^{\ell}, \\ -\lambda, & \text{if } x = y, \\ 0, & \text{otherwise }. \end{cases}$$

Notice that \vec{q}^{ℓ} also admits the following formula with respect concerning the tensor product

$$\overrightarrow{q}^{\ell} = \underbrace{\mathbb{I} \otimes \mathbb{I} \otimes \dots \otimes \mathbb{I} \otimes \overrightarrow{A} \otimes \mathbb{I} \otimes \dots \otimes \mathbb{I}}_{d \text{ times}}, \qquad (38)$$

with \mathbb{I} the 2 × 2 identity matrix and $\overrightarrow{A} = \begin{pmatrix} -\lambda & \lambda \\ \lambda & -\lambda \end{pmatrix}$, which is the ℓ^{th} matrix in the previous product. Indeed, by the definition of tensor product, for any $x = (x^i)_{i=1}^d, y = (y^i)_{i=1}^d \in X$, we observe that

$$\begin{split} (\mathbb{I} \otimes \mathbb{I} \otimes ... \otimes \mathbb{I} \otimes \underbrace{\overrightarrow{A}}_{\ell^{th}} \otimes \mathbb{I} \otimes ... \otimes \mathbb{I})(x, y) &= \mathbb{I}(x^1, y^1) \mathbb{I}(x^2, y^2) ... \overrightarrow{A}(x^{\ell}, y^{\ell}) ... \mathbb{I}(x^d, y^d) \\ &= \begin{cases} \lambda \ , & \text{if } x^i = y^i \text{ for } i \neq \ell \text{ and } x^{\ell} \neq y^{\ell} \ , \\ -\lambda \ , & \text{if } x = y \ , \\ 0 \ , & \text{otherwise } . \end{cases} \end{split}$$

which is exactly the expression of $\vec{q}^{\ell}(x, y)$. We now use the Kolmogorov equation combined with the expression of \vec{q}_{t}^{ℓ} in (38), and apply the formula $e^{\mathbb{I}\otimes A+B\otimes\mathbb{I}} = e^{A} \otimes e^{B}$ for any matrix A, B (Gavrilyuk et al., 2011, Appendix) to get

$$\overrightarrow{p}_t = e^{t\overrightarrow{q}} = e^{\sum_{\ell=1}^d t\overrightarrow{q}^\ell} = \underbrace{e^{t\overrightarrow{A}} \otimes \dots \otimes e^{t\overrightarrow{A}}}_{d \text{ times}}$$

We are thus left with the computation of $e^{t\vec{A}}$. It is clear that the eigenvalues of \vec{A} are 0 and -2λ , with the corresponding eigenvectors $\begin{pmatrix} 1 & 1 \end{pmatrix}^T$ and $\begin{pmatrix} 1 & -1 \end{pmatrix}^T$ respectively. Consequently, we can compute $e^{t\vec{A}}$ as: for any $a, b \in \{0, 1\}$,

$$\overrightarrow{p}_t^1(a,b) := \mathrm{e}^{t\overrightarrow{A}}(a,b) = \begin{cases} \frac{1}{2} + \frac{1}{2}\mathrm{e}^{-2\lambda t} , & \text{if } a = b , \\ \frac{1}{2} - \frac{1}{2}\mathrm{e}^{-2\lambda t} , & \text{if } a \neq b , \end{cases}$$

and the formula of transition probability \overrightarrow{p}_t for $0 \leq t \leq T_f$ follows: for any $x = (x^i)_{i=1}^d$ and $y = (y^i)_{i=1}^d$ in X,

$$\overrightarrow{p}_t(x,y) = \prod_{i=1}^d \overrightarrow{p}_t^1(x^i, y^i), \quad \text{with } \overrightarrow{p}_t^1(x^i, y^i) = \begin{cases} \frac{1}{2} + \frac{1}{2}e^{-2\lambda t} , & \text{if } x^i = y^i , \\ \frac{1}{2} - \frac{1}{2}e^{-2\lambda t} , & \text{otherwise } . \end{cases}$$

B.2.2. CONDITIONAL EXPECTATION EXPRESSION OF THE SCORE FUNCTION

Proof of Proposition 1.1. Fix $x \in X$ and $\ell = 1, ..., d$. First note that by definition of μ_{T_f-t} as the marginal distribution of the noising process, we have

$$\mu_{T_f-t}(x) - \mu_{T_f-t}(\varphi^{(\ell)}(x)) = \sum_{z \in \mathsf{X}} \mu_0(z) (\overrightarrow{p}_{T_f-t}(z,x) - \overrightarrow{p}_{T_f-t}(z,\varphi^{(\ell)}(x))) .$$
(39)

The formula of transition probabilities $\overrightarrow{p}_{T_f-t}(z, \varphi^{(\ell)}(x))$ combined with the definition of $\varphi^{(\ell)}(x)$ lead to

$$\begin{split} \overrightarrow{p}_{T_f-t}(z,\varphi^{(\ell)}(x)) &= \prod_{i=1}^d \overrightarrow{p}_{T_f-t}^1(z^i,\varphi^i_\ell(x)) \\ &= \overrightarrow{p}_{T_f-t}^1(z^\ell,\varphi^\ell_\ell(x)) \prod_{\substack{i=1\\i\neq\ell}}^d \overrightarrow{p}_{T_f-t}^1(z^i,x^i) \\ &= \frac{\overrightarrow{p}_{T_f-t}^1(z^\ell,\varphi^\ell_\ell(x))}{\overrightarrow{p}_{T_f-t}^1(z^\ell,x^\ell)} \overrightarrow{p}_{T_f-t}(z,x) \,. \end{split}$$

Substituting this into (39) implies

$$\begin{split} \mu_{T_f-t}(x) &- \mu_{T_f-t}(\varphi^{(\ell)}(x)) = \sum_{z \in \mathsf{X}} \mu_0(z) \overrightarrow{p}_{T_f-t}^1(z, x) (1 - \frac{\overrightarrow{p}_{T_f-t}^1(z^{\ell}, \varphi^{(\ell), \ell}(x))}{\overrightarrow{p}_{T_f-t}^1(z^{\ell}, x^{\ell})}) \\ &= \sum_{z \in \mathsf{X}} \left[\frac{2\mathrm{e}^{-2\lambda(T_f-t)}}{1 + \mathrm{e}^{-2\lambda(T_f-t)}} - \frac{4\mathrm{e}^{-2\lambda(T_f-t)}(x^{\ell} - z^{\ell})^2}{1 - \mathrm{e}^{-4\lambda(T_f-t)}} \right] \mathbb{P} \left[\overrightarrow{X}_0 = z, \overrightarrow{X}_{T_f-t} = x \right] \;, \end{split}$$

where the last equality comes from the formula of $\overrightarrow{p}_{T_f-t}^1$ and the fact that if $z^{\ell} = \varphi^{(\ell),\ell}(x)$ then $z^{\ell} \neq x^{\ell}$. Therefore, the score function in components are

$$\begin{split} s_t^{\ell}(x) &= \frac{\mu_{T_f - t}(x) - \mu_{T_f - t}(\varphi^{(\ell)}(x))}{\mu_{T_f - t}(x)} \\ &= \sum_{z \in \mathsf{X}} \left[\frac{2\mathrm{e}^{-2\lambda(T_f - t)}}{1 + \mathrm{e}^{-2\lambda(T_f - t)}} - \frac{4\mathrm{e}^{-2\lambda(T_f - t)}(x^{\ell} - z^{\ell})^2}{1 - \mathrm{e}^{-4\lambda(T_f - t)}} \right] \mathbb{P} \left[\overrightarrow{X}_0 = z | \overrightarrow{X}_{T_f - t} = x \right] \\ &= \mathbb{E} \left[\frac{2\alpha_{T_f - t}}{1 + \alpha_{T_f - t}} - \frac{4\alpha_{T_f - t}(\overrightarrow{X}_{T_f - t}^{\ell} - \overrightarrow{X}_0^{\ell})^2}{1 - \alpha_{T_f - t}^2} \middle| \overrightarrow{X}_{T_f - t} = x \right] , \quad \text{for } t \in [0, T_f) , \end{split}$$

where $\alpha_t = e^{-2\lambda t}$, and we finish the proof of Proposition 1.1.

B.2.3. INVARIANT MEASURE OF THE FORWARD PROCESS

As we have a comprehensive understanding of the forward process, we observe that its invariant measure is the uniform distribution over X, denoted by γ^d . Indeed, for any $x \in X$ and $t \in [0, T_f]$,

$$(\gamma^d \overrightarrow{p}_t)(x) = \sum_{z \in \mathsf{X}} \gamma^d(z) \overrightarrow{p}_t(z, x) = \frac{1}{2^d} \sum_{z \in \mathsf{X}} \overrightarrow{p}_t(z, x) = \frac{1}{2^d} = \gamma^d(x) \; .$$

Furthermore, by formula of \overrightarrow{p} given in (15), we have $\overrightarrow{p}_t(x, y) \xrightarrow{t \to \infty} \frac{1}{2^d}$ for any $x, y \in X$. Consequently, the following holds for any $x \in X$,

$$\mu_t(x) = \sum_{z \in \mathsf{X}} \mu_0(z) \overrightarrow{p}_t(z, x) \xrightarrow{t \to \infty} \frac{1}{2^d} \sum_{z \in \mathsf{X}} \mu_0(z) = \frac{1}{2^d} = \gamma^d(x) ,$$

meaning that the forward dynamic $(\overrightarrow{X}_t)_{t \in [0,T_f]}$ converges geometrically fast to γ^d .

C. Implementation of DMPMs

C.1. Alternative ideal backward simulation

Besides the simulation of the backward process provided in Section 1.3, we can also use the following procedure to produce the time-reversal dynamic.

The second procedure to sample $(\overline{X}_t)_{t\in[0,T_f]}$ is to consider a sample \overline{X}_0 from μ_{T_f} and a sequence of i.i.d. random variables distributed according to the exponential distribution with parameter 1, $\{E_i^{\ell} : i \in \mathbb{N}, \ell \in \{1, \ldots, d\}\}$, we can define the jump times $(T_i)_{i\in\mathbb{N}}$ of the backward process and its transition by induction setting $T_0 = 0$. Given $(T_i, \overline{X}_{T_i})$, we define the next jump time as $T_{i+1}^j = T_i + \Delta T_{i+1}^j$, where $\Delta T_{i+1}^j = \inf\{t \ge 0 : \int_0^t \lambda(1 - s^j(\overline{X}_{T_i})) dr \ge E_i^j\}$. Then, set $T_{i+1} = T_{i+1}^{\ell_i}$, where $\ell_i = \arg\min_{j \in \{1, \ldots, d\}} T_{i+1}^j$, and $\overline{X}_t = \overline{X}_{T_i}$ for $t \in (T_i, T_{i+1} \wedge T_f)$, and finally if $T_{i+1} < T_f, \overline{X}_{T_{i+1}}^\ell = 1 - \overline{X}_{T_i}^{\ell_i}$ for $\ell_i \in \{1, \ldots, d\}$.

C.2. Perfect backward approximation

We provide here the pseudo-code of backward approximation sampling in continuous time scheme:

Algorithm 1 DMPMs Algorithm (Continuous time scheme)

Input: a time horizon $T_f \gg 1$ large enough, a prescribed jump rate λ , an approximate score function s^{θ^*} Backward process:

Set $T_0 = 0$ and initialize $\overleftarrow{X}_0 \sim \gamma^d$ $i \leftarrow 0$ while $T_i \leq T_f$ do Draw $E_i \sim \text{Exp}(1)$ Solve $\Delta T_{i+1} = \inf\{t \ge 0 : \int_0^t \lambda_{T_i+r}^{\theta^*}(\overleftarrow{X}_{T_i}) dr \ge E_i\}$, with $\lambda_t^{\theta^*}(x) = \lambda \sum_{\ell=1}^d (1 - s_t^{\theta^*,\ell}(x))$ Set $T_{i+1} = T_i + \Delta T_{i+1}$ if $T_i < t < \min(T_{i+1}, T_f)$ then Set $\overleftarrow{X}_t = \overleftarrow{X}_{T_i}$ end if if $T_{i+1} < T_f$ then Draw $\ell_i \in \{1, \dots, d\} \sim \text{Cate}(\{\lambda(1 - s_{T_{i+1}}^{\theta^*,\ell}(\overleftarrow{X}_{T_i}))/\lambda_{T_{i+1}}^{\theta^*}(\overleftarrow{X}_{T_i})\}_{\ell=1}^d)$ Set $\overleftarrow{X}_{T_{i+1}} = \varphi^{(\ell_i)}(\overleftarrow{X}_{T_i})$ end if $i \leftarrow i + 1$ end while Output: \overleftarrow{X}_{T_f}

C.3. Discrete denoiser and score reparameterization

Discrete-denoiser structure.

Recall from Proposition 1.1 that each score component admit the following conditional expectation:

$$s_t^{\ell}(x) = \mathbb{E}\left[f_t^{\ell}(\overrightarrow{X}_0, \overrightarrow{X}_{T_f-t}) | \overrightarrow{X}_{T_f-t} = x\right] , \qquad (40)$$

where

$$f_t^{\ell}(\vec{X}_0^{\ell}, \vec{X}_{T_f-t}) = \frac{2\alpha_{T_f-t}}{1 + \alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}(\vec{X}_{T_f-t}^{\ell} - \vec{X}_0^{\ell})^2}{1 - \alpha_{T_f-t}^2}$$
(41)

г ___

for $t \in [0, T_f)$, $x \in X$ and $\ell = 1, ..., d$.

Remark that

$$\mathbb{E}\left[f_t^{\ell}(\overrightarrow{X}_0^{\ell}, \overrightarrow{X}_{T_f-t}) | \overrightarrow{X}_{T_f-t} = x\right] = \frac{2\alpha_{T_f-t}}{1+\alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}\mathbb{E}\left[(\overrightarrow{X}_{T_f-t}^{\ell} - \overrightarrow{X}_0^{\ell})^2 | \overrightarrow{X}_{T_f-t} = x\right]}{1-\alpha_{T_f-t}^2} .$$
(42)

Thus we introduce the function d_t^{ℓ} defined as

$$d_t^{\ell}: x \mapsto \mathbb{E}\left[(\vec{X}_{T_f-t}^{\ell} - \vec{X}_0^{\ell})^2 | \vec{X}_{T_f-t} = x \right],$$
(43)

which can be further rewritten as

$$\begin{aligned} d_t^\ell(x) &= \mathbb{E}\left[\left(\overrightarrow{X}_{T_f-t}^\ell - \overrightarrow{X}_0^\ell \right)^2 \middle| \overrightarrow{X}_{T_f-t} = x \right] \\ &= \mathbb{E}\left[\mathbbm{1}_{\overrightarrow{X}_{T_f-t}^\ell \neq \overrightarrow{X}_0^\ell} \middle| \overrightarrow{X}_{T_f-t} = x \right] \\ &= \mathbb{P}\left(\left| \overrightarrow{X}_0^\ell \neq x^\ell \right| \middle| \overrightarrow{X}_{T_f-t} = x \right) \,. \end{aligned}$$

In some sense, this is the discrete version of the continuous denoiser $\mathbb{E}[\vec{X}_0|\vec{X}_t]$ approximated by classical diffusion models (Song et al., 2021), as obtained from the score by Tweedie's formula. Thus we call $d_t^{\ell}(x)$ the discrete denoiser.

Score reparameterization. Based on the previous derivations, each score component $s_t^{\ell}(x)$ can be written as a function of d_t^{ℓ} :

$$s_t^{\ell}(x) = \frac{2\alpha_{T_f-t}}{1+\alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}d_t^{\ell}(x)}{1-\alpha_{T_f-t}^2}, \qquad (44)$$

So we can reparameterize our score models s_t^{θ} as

$$s_t^{\theta,\ell}(x) = \frac{2\alpha_{T_f-t}}{1+\alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}d_t^{\theta,\ell}(x)}{1-\alpha_{T_f-t}^2} , \qquad (45)$$

where $d_t^{\theta,\ell}(x)$ aims to approximate $d_t^{\ell}(x)$.

C.4. Objective functions derived from the discrete denoiser structure

Inspired by the previous derivations, we modify our existing \mathfrak{L}_{L^2} loss function to replace by a denoising loss equivalent. We introduce a cross-entropy loss, and finally propose a scaling of the loss functions, based on the average output magnitude of the discrete denoiser, thus helping with the learning, and improving synergies between loss elements.

Score-matching objective \mathfrak{L}_{L^2} . We rewrite the objective function \mathfrak{L}_{L^2} to fit the *discrete denoiser*, considered as a conditional expectation:

$$\mathfrak{L}_{\mathrm{L}^{2}}^{\mathrm{den}}:\theta\mapsto\int_{0}^{T_{f}}\mathfrak{L}_{t,\mathrm{L}^{2}}(\theta)\mathrm{d}t\;,\quad\mathfrak{L}_{t,\mathrm{L}^{2}}(\theta)=\mathbb{E}\left[\|d_{T_{f}-t}^{\theta}(\overrightarrow{X}_{t})-(\overrightarrow{X}_{0}-\overrightarrow{X}_{t})\mathrm{d}(\overrightarrow{X}_{0}-\overrightarrow{X}_{t})\|^{2}\right]\;,\tag{46}$$

where d is the element-wise product.

Cross-entropy objective \mathcal{L}_{CE} . Instead of the \mathcal{L}_{L^2} loss suggested by the conditional expectation structure, we can consider a cross-entropy loss to fit our model to the correct distribution: classical derivations from the conditional log-likelihood $\sum_{\ell=1}^{d} \mathbb{E} \left[\log p_t^{\theta,\ell}(\vec{X}_{T_f-t}^{\ell} | \vec{X}_0^{\ell}) \right]$, where

$$p_t^{\theta,\ell}(x_{T_f-t}|x_0) = \begin{cases} d_t^{\theta,\ell}(x_{T_f-t}) & \text{if } x_{T_f-t} \neq x_0\\ 1 - d_t^{\theta,\ell}(x_{T_f-t}) & \text{else} \end{cases},$$
(47)

lead to the following cross entropy loss:

$$\mathfrak{L}_{CE}(\theta) = -\int_{0}^{T_{f}} \mathfrak{L}_{t,CE}(\theta) \mathrm{d}t , \qquad (48)$$

where

$$\mathfrak{L}_{t,\mathrm{CE}}(\theta) = \mathbb{E}\left[\sum_{l=1}^{d} Y_t^{\ell} \log d_t^{\theta,\ell}(\overrightarrow{X}_{T_f-t}) + (1-Y_t^{\ell}) \log\left(1-d_t^{\theta,\ell}(\overrightarrow{X}_{T_f-t})\right)\right], \quad Y_t^{\ell} = \begin{cases} 1 & \text{if } \overrightarrow{X}_0^{\ell} \neq \overrightarrow{X}_{T_f-t}^{\ell}, \\ 0 & \text{else}. \end{cases}$$

Further improvements. To address training efficiency, we inspect the average magnitude of the loss across the dataset, at each timestep. Indeed, the average value of d_t^{ℓ} is

$$w_{T_{f}-t} = \mathbb{E}\left[d_{t}^{\ell}(\vec{X}_{T_{f}-t})\right] = \mathbb{E}\left[\mathbb{E}\left[d_{t}^{\ell}(\vec{X}_{T_{f}-t})\middle|\vec{X}_{0}\right]\right]$$
$$= \mathbb{E}\left[\mathbb{P}\left(\vec{X}_{0}^{\ell} \neq \vec{X}_{T_{f}-t}^{\ell}\middle|\vec{X}_{0}\right)\right]$$
$$= \frac{1}{2}\left(1 - \alpha_{T_{f}-t}\right) , \qquad (49)$$

as given by the formulas for the transition kernels of the forward process. We can see that the value of w_t is close to zero for small values of t, which stalls the learning process. Empirically, we find that dividing the integrand of either loss terms \mathfrak{L}_{L^2} or \mathfrak{L}_{CE} by w_t yields improvements. As a result, we modify the losses to counterbalance their diminishing magnitude across timesteps: $\mathfrak{L}_{t,L^2}^w = \frac{\mathfrak{L}_{t,L^2}}{w_t}, \mathfrak{L}_{t,CE}^w = \frac{\mathfrak{L}_{t,CE}}{w_t}$, and define the associated losses

$$\mathfrak{L}_{\mathrm{L}^{2}}^{w}(\theta) = \int_{0}^{T_{f}} \mathfrak{L}_{t,\mathrm{L}^{2}}^{w}(\theta) \mathrm{d}t , \quad \mathfrak{L}_{\mathrm{CE}}^{w}(\theta) = -\int_{0}^{T_{f}} \mathfrak{L}_{t,\mathrm{CE}}^{w}(\theta) \mathrm{d}t .$$
(50)

Comparing $\mathfrak{L}_{L^2}, \mathfrak{L}_{CE}, \mathfrak{L}_{L^2}^w, \mathfrak{L}_{CE}^w$.

In Figure 4, we plot the average loss per timestep, for a trained model on MNIST (following the specifications given in Appendix D.2)). It shows that, on average, the \mathfrak{L}_L loss effectively becomes a scaled variant of the cross-entropy objective, which is reflected in similar performance results. This corroborates our derivation that L2 acts as an effective lower bound to the log-likelihood. This also supports its relevancy with respect to the underlying structure of this generative model. It must be noted that both losses still benefit from positive synergies when used together.

Importantly, dividing by $w_t = (1 - \alpha_t)/2$ particularly helps at smaller timesteps, and keeps the loss values at the same magnitude across timesteps, enhancing training dynamics. This is illustrated in Figure 5, where scaling the losses with the w scale factor consistently yields improvements.





Figure 4. Comparison of \mathfrak{L}_{L^2} , \mathfrak{L}_{CE} , $\mathfrak{L}_{L^2}^w$, \mathfrak{L}_{CE}^w average losses over timesteps. The two losses become scaled version of one another only when averaged over data, but otherwise benefit from positive synergies when mixed together.

Figure 5. FID \downarrow , on MNIST, for models trained with \mathfrak{L}_{ϖ} and $\mathfrak{L}_{\varpi}^{w}$ losses, evaluated using 200 reverse steps with the denoise-renoise sampler. Scaling with w yields consistent improvements, with the best loss configuration $\mathfrak{L}_{1/3,1/3,1/3}^{w}$ involving all the methodological improvements we discussed.

Final objective functions. We choose a linear combination of the previous loss objectives, weighted by positive coefficients $\varpi_1, \varpi_2, \varpi_3$:

$$\mathfrak{L}_{\varpi} = \varpi_1 \mathfrak{L}_{\mathrm{L}^2}^{\mathrm{den}} + \varpi_2 \mathfrak{L}_{\mathrm{e}} + \varpi_3 \mathfrak{L}_{\mathrm{CE}} , \qquad (51)$$

and, if we choose their version weighted by $1/w_t$:

$$\mathfrak{L}^w_{\varpi} = \varpi_1 \mathfrak{L}^w_{\mathrm{L}^2} + \varpi_2 \mathfrak{L}_{\mathrm{e}} + \varpi_3 \mathfrak{L}^w_{\mathrm{CE}} .$$
⁽⁵²⁾

Algorithm 2 Training Algorithm for DMPM (Reparameterized Score)

Require: Dataset \mathcal{D} of samples $X \in \{0, 1\}^d$; Time horizon $T_f > 0$ and rate $\lambda > 0$; Parameterized discrete denoiser model $\{d_t^{\theta,\ell}(x): \theta \in \Theta\}_{t,\ell,x};$ Derived score function $s_t^{\theta} := \frac{2\alpha_{T_f-t}}{1+\alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}d_t^{\theta}}{1-\alpha_{T_f-t}^2}$ (score reparameterization (24)); Define α_t as in (13), f_t as in (19) Loss coefficients $\varpi_1, \varpi_2, \varpi_3 \ge 0$; 1: while optimization has not converged do Sample a batch $\{X_i\}_{i=1}^B$ from \mathcal{D} . 2: Draw $t_1, \ldots, t_B \stackrel{\text{iid}}{\sim} \text{Unif}([0, T_f])$ 3: Forward sampling: fast simulation via $p_{t|0} = (p_{t|0}^1)^{\otimes d}$ 4: 5: for i = 1 to B do $X_{i.0} \leftarrow X_i$ 6: $p_{T_f-t_i} \leftarrow (1 - \alpha_{T_f-t_i})/2$ 7: Compute \vec{X}_{i,T_f-t_i} by flipping each bit of $\vec{X}_{i,0}$ independently with probability $p_{T_f-t_i}$ 8: if Scaling losses with average d_t magnitude then 9: $w_i \leftarrow (1 - \alpha_{T_f - t_i})/2$ 10: 11: else 12: $w_i \leftarrow 1$ 13: end if 14: end for $\mathfrak{L}_{L^2}(\theta) \leftarrow \frac{1}{B} \sum_{i=1}^{B} \frac{1}{w_i} \| d_t^{\theta}(\overrightarrow{X}_{i,T_f-t_i}) - (\overrightarrow{X}_{i,T_f-t_i} - \overrightarrow{X}_{i,0}) \mathrm{d}(\overrightarrow{X}_{i,t} - \overrightarrow{X}_{i,0}) \|^2$ 15: $\mathfrak{L}_{CE}(\theta) \leftarrow \tfrac{1}{Bd} \sum_{i=1}^{B} \tfrac{1}{w_i} \sum_{\ell=1}^{d} \left(\mathbb{1}_{\overrightarrow{X}_0^\ell \neq \overrightarrow{X}_{T_f - t_i}^\ell} \log d_{t_i}^{\theta, \ell} (\overrightarrow{X}_{T_f - t_i}^\ell) + (1 - \mathbb{1}_{\overrightarrow{X}_0^\ell \neq \overrightarrow{X}_{T_f - t_i}^\ell}) \log (1 - d_{t_i}^{\theta, \ell} (\overrightarrow{X}_{T_f - t_i}^\ell) \right)$ 16: $\mathfrak{L}_{\mathbf{e}}(\theta) \leftarrow \frac{1}{B} \sum_{i=1}^{B} \sum_{\ell=1}^{d} \left(-s_{T_f - t_i}^{\theta, \ell}(\overrightarrow{X}_{t_i}) + (f_{T_f - t_i}^{\ell}(\overrightarrow{X}_{t_i}) - 1) \log(1 - s_{T_f - t_i}^{\theta, \ell}(\overrightarrow{X}_{t_i})) \right)$ 17: $\mathfrak{L}_{\varpi}(\theta) \leftarrow \varpi_1 \, \mathfrak{L}_{\mathrm{L}^2}^{\mathrm{den}} + \varpi_2 \, \mathfrak{L}_{\mathrm{e}} + \varpi_3 \, \mathfrak{L}_{\mathrm{CE}}$ 18: Perform a gradient step on $\mathfrak{L}_{\varpi}(\theta)$ w.r.t. θ . 19: 20: end while 21: **Return** the final parameter θ^* .

C.5. Generative process and sampling procedures

Once we obtain our neural network d_t^{θ} approximating d_t , we use it to produce fresh samples that closely mimic the observed data. To do so, we first introduce a DMPM sampler based on the true reverse process. We then propose a slight modification, leveraging the distribution on indices available at each step, by flipping multiple bits instead of just one, using a flip-schedule. Finally, we derive a denoise-renoise sampler, solely based on the discrete denoiser structure of the problem, as inspired by similar lines of work in conitnuous diffusion.

DMPM sampler. A first sampling procedure is given in Algorithm 3. It is designed to be as close as possible to the true backward process, while enabling efficient parallelization when implemented. It consists in a piecewise-approximation of the functions of interest, parameterized by the choice of a time discretization grid $0 = t_0 < t_1 < \cdots < t_K = T_f$, which we call a **time-schedule**.

In Table 2, we give the different time-schedules we experiment with. We draw inspiration from numerous lines of work on continuous and discrete diffusion (Shi et al., 2024; Karras et al., 2022), in which these are common choices.

DMPM sampler with flip-schedule In Algorithm 4, we further take advantage of the specific structure of our backward process, by leveraging the distribution over indices given by the learned score model at each timestep t. Instead of flipping a single bit per timestep t_k , we flip a total of M_{t_k} bits sampled without replacements from the given distribution. We call the sequence $\{M_t\}_{0 \le t \le T_f}$ the flip-schedule. When a time-schedule $\{t_k\}_{k=1}^K$ has been chosen, we also call the corresponding discrete sequence $\{M_{t_k}\}_{k=1}^K$ a flip-schedule.

In Table 3, we give the two flip-schedules we explore in this paper. The choice for the linear schedule is inspired from the philosophy of the masking schedule introduced in the context of masked diffusion by Shi et al. (2024).

Time-schedule	Value of t_k
Linear	$T_f \frac{k}{K}$
Quadratic	$T_{f_{\ell}}\left(\frac{k}{K}\right)^2$
Cosine	$T_f \cos\left(\frac{(1-k/K)\pi}{2}\right)$

Flip-schedule	Value of M_t			
Constant	M			
Linear	$M \frac{t}{T_f}$			

Table 2. Different time schedules $(t_k)_{k=1}^K$ used in our experiments. T_f denotes the final time, and K is the number of reverse steps.

Table 3. Different flip schedules $(M_t)_{0 \le t \le T_f}$ used in our experiments. In both schedules, M is a constant to be fixed and controls the total number of bits flipped during generation.

Denoise-renoise sampler. In Algorithm 5, we introduce the following denoise/renoise cycle, interpreting the model output d_t^{θ} as the probability that each bit should be flipped at timestep t to reach timestep 0. After doing a full denoise pass (from time $T_f \to 0$), we noise the sample with the transition kernel of the forward process (from time $0 \to T_f - \Delta$). Then we can do another denoise pass from $(T_f - \Delta) \to 0$, etc.

Algorithm 3 Backward sampling of DMPM with piecewise-constant score

Require: Time horizon $T_f > 0$ and rate $\lambda > 0$; K > 0 number of reverse steps and time-schedule $0 = t_0 < t_1 < \cdots < t_K = T_f$; Flip-schedule, *i.e.*, sequence of positive integers $\{M_{t_k}\}_{k=1}^K$; Discrete denoiser model d^{θ} ; Derived score function $s_t^{\theta} := \frac{2\alpha_{T_f-t}}{1+\alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}d_t^{\theta}}{1-\alpha_{T_f-t}^2}$ (score reparameterization (24)); Define α_t as in (13); 1: $\overleftarrow{X}_0^{\star} \sim \text{Unif}(0,1)^{\otimes d}$ 2: $E \sim \mathcal{E}(1)$ 3: $\Lambda \leftarrow 0$ 4: for k = 0 to K - 1 do 5: $\overline{\lambda}_{t_k} \leftarrow \lambda \sum_{\ell=1}^d (1 - s_{t_k}^{\theta, \ell})$ 6: $\Delta t_k \leftarrow t_{k+1} - t_k$ $\Lambda \leftarrow \Lambda + \overline{\lambda}_{t_k} \,\Delta t_k$ 7: $\ell^{\star} \sim \operatorname{Cate}\left(\left\{\frac{\lambda\left(1-s_{t_{k}}^{\theta,\ell}\right)}{\overline{\lambda}_{t_{k}}}\right\}_{\ell=1}^{d}\right)$ $\overleftarrow{X}_{t_{k}}^{\star,\ell^{\star}} \leftarrow 1 - \overleftarrow{X}_{t_{k}}^{\star,\ell^{\star}}$ $\overrightarrow{\Lambda} \leftarrow 0$ $\text{ if } \Lambda > E \ \text{ then }$ 8: 9: 10: 11: $E \sim \mathcal{E}(1)$ 12: 13: end if 14: $\overleftarrow{X}_{t_{k+1}}^{\star} \leftarrow \overleftarrow{X}_{t_k}^{\star}$ 15: end for Output: $\overleftarrow{X}_{T_f}^{\star}$

Algorithm 4 Backward sampling of DMPM with piecewise-constant score and flip-schedule

Require: Time horizon $T_f > 0$ and rate $\lambda > 0$; K > 0 number of reverse steps and time-schedule $0 = t_0 < t_1 < \cdots < t_K = T_f$; Flip-schedule, *i.e.*, sequence of positive integers $\{M_{t_k}\}_{k=0}^K$; Discrete denoiser model d^{θ} ; Derived score function $s_t^{\theta} := \frac{2\alpha_{T_f-t}}{1+\alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}d_t^{\theta}}{1-\alpha_{T_f-t}^2}$ (score reparameterization (24)); Define α_t as in (13); 1: $\overleftarrow{X}_0^{\theta} \sim \text{Unif}(0,1)^{\otimes d}$ 2: $E \sim \mathcal{E}(1)$ 3: $\Lambda \leftarrow 0$ 4: for k = 0 to K - 1 do 5: $\overline{\lambda}_{t_k} \leftarrow \lambda \sum_{\ell=1}^d (1 - s_{t_k}^{\theta, \ell})$ 6: $\Delta t_k \leftarrow t_{k+1} - t_k$ 7: $\Lambda \leftarrow \Lambda + \overline{\lambda}_{t_k} \Delta t_k$ 8: **if** $\Lambda > E$ **then** $[\ell_1^{\star}, \dots, \ell_M^{\star}] \sim \text{Hypergeometric}\left(\left\{\frac{\lambda\left(1-s_{t_k}^{\theta,\ell}\right)}{\overline{\lambda}_{t_k}}\right\}_{\ell=1}^d, M_{t_k}\right)$ 9: $\begin{aligned} & \mathbf{for} \ i = 1 \ \mathrm{to} \ M_{t_k} \ \mathbf{do} \\ & \overleftarrow{X}_{t_k}^{\theta, \ell_i^*} \leftarrow 1 - \overleftarrow{X}_{t_k}^{\theta, \ell_i^*} \end{aligned}$ 10: 11: end for 12: $\Lambda \leftarrow 0$ 13: 14: $E \sim \mathcal{E}(1)$ 15: end if 16: $\overleftarrow{X}^{\theta}_{t_{k+1}} \leftarrow \overleftarrow{X}^{\theta}_{t_k}$ 17: end for \leftarrow **Output:** $\overleftarrow{X}^{\theta}_{T_{f}}$

Algorithm 5 Denoise–Noise Cycling with a Discrete Denoiser Model

Require: Time horizon T_f > 0 and rate λ > 0; K > 0 number of reverse steps and time-schedule 0 = t₀ < t₁ < ··· < t_K = T_f; Discrete denoiser model d^θ;
1: X₀^θ ~ Unif(0, 1)^{⊗d} {initial sample in {0, 1}^d}
2: for k = 0 to K − 1 do

- 3: **Denoise phase:**
- 4: $d_{t_k} \leftarrow d^{\theta}_{t_k}(\overline{X}^{\theta}_{t_k})$
- 5: Compute $X_{T_f}^{e_k}$ by flipping each component l of $X_{t_k}^{\theta}$ with probability $d_{t_k}^l$
- 6: Noise phase:
- 7: Sample $\overleftarrow{X}^{\theta}_{t_{k+1}} \sim p_{T_f t_{k+1}|0}(\cdot |\overleftarrow{X}^{\theta}_{T_f})$, as in Algorithm 2
- 8: end for
 - **Output:** $\overleftarrow{X}^{\theta}_{T_f}$

D. Experiments

All experiments are conducted using PyTorch. All the training and experiments are conducted on four NVIDIA RTX8000 GPU.

We use the score parameterization introduced in (45):

$$s_t^{\theta,\ell}(x) = \frac{2\alpha_{T_f-t}}{1+\alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}d_t^{\theta,\ell}(x)}{1-\alpha_{T_f-t}^2},$$
(53)

where the neural network $d_t^{\theta,\ell}(x)$ aims to approximate $d_t^{\ell}(x) = \mathbb{P}(\overrightarrow{X}_0^{\ell} \neq x^{\ell} | \overrightarrow{X}_{T_f-t} = x)$. Since the output of the neural network is $d_t^{\theta}(x) \in (0, 1)^d$, we add a sigmoid activation function at the last layer.

We consider various loss configurations \mathfrak{L}_{ϖ} , \mathfrak{L}_{ϖ}^w as introduced in (26), (52), with 6 choices of coefficients $(\varpi_1, \varpi_2, \varpi_3)$ normalized in the 2-simplex $\Delta_2 \subset \mathbb{R}^3$. We test all $2^3 - 1 = 7$ possible non-empty combinations, minus the single \mathfrak{L}_e loss combination $(\varpi_2 = 1)$, as the latter only acts as entropic regularization and does not perform well by itself. This lets us study the synergies between the different loss terms.

D.1. Small dimension data

We first conduct experiments on a discrete data distribution p supported on $\{0,1\}^d$. Each component of $X = (X_i)_{i=1}^d \sim p$ is independently distributed as a Bernoulli distribution with parameter p_i :

$$p(x) = \prod_{i=1}^{d} p_i(x_i) , \qquad (54)$$

where the map $i \mapsto p_i$ forms a sawtooth-like pattern, oscillating linearly between 0.05 and 0.95, as can be seen in Figure 6.



Figure 6. Sawtooth pattern used to define $i \mapsto p_i$, plotted with d = 16. Values oscillate linearly from 0.05 to 0.95, and back.

For training, we use 20 000 datapoints resampled at each epoch, and a batch size of 1024. We train each model for 300 epochs, using AdamW with a learning rate of 1e-3. We employ a network composed of multiple MLP blocks: 4 residual blocks, each consisting of two feed-forward layers of width 256; layer normalization and SiLU activations in each block; a feed-forward embedding for the timesteps, mapping \mathbb{R} to a hidden dimension of 256, whose output is then injected into each residual block by an additional MLP of dimension 256×256 .

For evaluation, we estimate each distribution with 20,000 samples, and draw 1000 vectors uniformly on the simplex Δ_d to compute our SWD metric (see Appendix D.3).

D.2. Image data

We work on the binarized MNIST dataset, which we scale from 28×28 to 32×32 in order to fit in the U-Net architecture. We set the pixel value to 0 if its intensity is below 0.5, and to 1 otherwise.

We compare DMPM to MD4 (masked diffusion, as in Shi et al. (2024)) and DFM (discrete flow matching, as in Gat et al. (2024)). We reimplement MD4 with the cosine schedule and the algorithms given in Appendix F of Shi et al. (2024). We implement DFM based on the Pytorch implementation in https://github.com/gle-bellier/discrete-fm,

and we use corrector sampling for better results.

For DMPM, we are using the cosine time-schedule and time horizon $T_f = 3$. For both MD4 and DFM, we set the mask value to the integer 2.

To establish a fair comparison, we use the same network model for every method. We use a U-Net following the implementation of (Nichol & Dhariwal, 2021) available in https://github.com/openai/improved-diffusion. We dimension the network as follows.

The first layer is an embedding layer of output dimension 32 and input dimension d_{input} , where $d_{input} = 2$ for DMPM (input values are either 0 and 1) and $d_{input} = 3$ for MD4 and DFM (input values are either 0, 1 or the mask value 2).

We set the hidden layers to [128, 256, 256, 256], fix the number of residual blocks to 2 at each level, and add self-attention block at resolution 16×16 , using 4 heads. We use an exponential moving average with a rate of 0.99. We use the silu activation function at every layer. Timestep t is fed to the model through the Transformer sinusoidal position embedding.

For DMPM and MD4, we set the number of output channels to 1 and add a sigmoid activation at the last layer. For DFM, we set the output channels to 3 and apply softmax channel-wise.

The optimizer is AdamW with learning rate 5e-4. We use the StepLR scheduler which scales the learning rate by $\gamma = .99$ every 400 steps. We train on MNIST for 120 000 steps with batch size 256. A single training run on MNIST takes approximately 6 hours per GPU, and requires about 6-12GB of VRAM for our settings.

To assess the quality of our generative models, we compute our metrics between $4\,000$ real images and $4\,000$ generated images. Generating $4\,000$ images with $1\,000$ reverse steps takes approximately 2 hours on one GPU.

D.3. Metrics

For low-dimensional data, we use a custom sliced Wasserstein metric. For image data, in addition to the classical FID metric, we use a F_1^{DC} summary score, based on the density and coverage metrics.

 $\mathbf{F}_{1}^{\text{DC.}}$ as summary metric of density-coverage The density and coverage metrics are introduced in the setting of generative models by (Naeem et al., 2020). They assess the overlap of sample distributions using local geometric structures. Density measures how much the generated distribution is contained in the original data distribution (measuring quality), and coverage measures how much of the original data distribution is covered by the generated distribution (diversity).

These metrics are improvements of the precision and recall metrics for generative models (Kynkäänniemi et al., 2019). They offer different measures to characterize the performance of generative models. For instance they can decorrelate the negative effect of mode collapse from the negative effect of noisy/blurry generations, each of them decreasing respectively coverage and density, and have been of importance in recent studies, e.g., in heavy-tailed generative modeling (Shariatian et al., 2024; Yoon et al., 2023).

We consider a single summary F_1^{DC} score, which we define as the harmonic mean of these two values:

$$F_1^{DC} = 2 \cdot \frac{\text{density} \cdot \text{coverage}}{\text{density} + \text{coverage}} .$$
(55)

Sliced Wasserstein metric SWD.

Since the state space of our dataset over $\{0, 1\}^d$ is of size 2^d , we cannot work with histogram-based metrics, which would require exponentially many samples when d increases.

We address this issue with our sliced Wasserstein metric SWD. This metric is defined between distributions μ , ν on $\{0,1\}^d$ as:

$$SWD(\mu,\nu) = \int_{\Delta_d} W\left(u_{\#}\mu, u_{\#}\nu\right) du , \qquad (56)$$

where, for $u \in \Delta_d$, the pushforward $u_{\#}$ is derived from the function

$$x \in \{0,1\}^d \mapsto \langle u, x \rangle \in [0,1] .$$
(57)

Simple Monte-Carlo averages are used to evaluate the integral with respect to the uniform distribution over the simplex Δ_d ,

and we compute the Wasserstein distance between the pushforward measures with the pyemd package (Laszuk, 2017).

E. Additional results

In this section, we give grid images of generated samples for DMPM models trained on binarized MNIST, with the loss $\mathfrak{L}^w_{1/3,1/3,1/3}$.



Figure 7. Default DMPM sampler, 25 reverse steps, cosine timeschedule, linear flip-schedule dimensioned for 1000 total bit flips.



Figure 8. Denoise-renoise sampler, 200 reverse steps, cosine time-schedule.

F. Convergence of DMPMs

The proof of DMPMs' convergence requires understanding the backward dynamic under the canonical process point of view equivalent to the transition matrix point of view we provided in Section 1. We provide first some essential preliminaries on Poisson measures and the corresponding Itô's formula.

F.1. Some basic facts on stochastic calculus for CTMCs

F.1.1. POINT PROCESSES AND POISSON POINT PROCESSES

Let (Y, \mathcal{Y}) be a measurable space. Let \mathcal{M}_Y be the set of non-negative (possibly infinite) integer-valued measures on (Y, \mathcal{Y}) . Let \mathcal{M}_Y be the smallest σ -field on \mathcal{M}_Y with respect to which the mappings $\mu \in \mathcal{M}_Y \mapsto \mu(B) \in \mathbb{N} \cup \{+\infty\}, B \in \mathcal{Y}$, are measurable.

Definition F.1 (Poisson random measure). An $(\mathcal{M}_{\mathsf{Y}}, \mathcal{M}_{\mathsf{Y}})$ -valued random variable μ (*i.e.*, a mapping $\mu : \Omega \to \mathcal{M}_{\mathsf{Y}}$ defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ which is $\mathcal{F}/\mathcal{M}_{\mathsf{Y}}$ -measurable) is called a *Poisson random measure* if

- for each $B \in \mathcal{Y}$, $\mu(B)$ is Poisson distributed; *i.e.*, $\mathbb{P}(\mu(B) = n) = \lambda_{\mu}(B)^n \exp[-\lambda_{\mu}(B)]/n!$ for any $n \in \mathbb{N}$ where $\lambda_{\mu}(B) = \mathbb{E}[\mu(B)], B \in \mathcal{Y}$. Note that if $\lambda_{\mu}(B) = \infty$, we understand that $\mu(B) = \infty$ a.s.
- If $B_1, B_2, \ldots, B_n \in \mathcal{Y}$ are disjoint, then $\mu(B_1), \mu(B_2), \ldots, \mu(B_n)$ are mutually independent.

Remark that it is easy to show that λ_{μ} is a non-negative measure on (Y, \mathcal{Y}) that uniquely determines the distribution of μ by the monotone class theorem; see (Ikeda & Watanabe, 2014, Chapter I, Section 9). It is called the *mean measure* or the *intensity measure* of the Poisson random measure μ .

Let (X, \mathcal{X}) be a measurable space. We define a *point function* \mathfrak{p} on X by a mapping $\mathfrak{p} : D_{\mathfrak{p}} \to X$, where its domain $D_{\mathfrak{p}}$ is a countable subset of $(0, \infty)$. \mathfrak{p} defines a non-negative (possibly infinite) integer-valued measure $N_{\mathfrak{p}}$ on $(0, \infty) \times X$ endowed with the product σ -field $\mathcal{B}((0, \infty)) \times \mathcal{X}$ by

$$N_{\mathfrak{p}}((0,t]\times\mathsf{U})=\mathrm{Card}\left\{s\in D_{\mathfrak{p}}\ :\ s\leqslant t,\ \mathfrak{p}(s)\in\mathsf{U}\right\},\quad t>0\ ,\quad\mathsf{U}\in\mathcal{X}\ .$$

A point process is obtained by randomizing the notion of point functions. Let Π_X be the set of point functions with values in X and $\mathcal{B}(\Pi_X)$ be the smallest σ -field on Π_X with respect to which the mappings $\mathfrak{p} \mapsto N_{\mathfrak{p}}((0,t] \times U), t > 0, U \in \mathcal{X}$, are measurable.

Definition F.2 (Point process). A *point process* \mathbf{p} on X is a $(\Pi_X, \mathcal{B}(\Pi_X))$ -valued random variable, that is, a mapping $\mathbf{p}: \Omega \to \Pi_X$ defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ which is $\mathcal{F}/\mathcal{B}(\Pi_X)$ -measurable.

A point process **p** is called *stationary* if for every t > 0, **p** and $\theta_t \mathbf{p}$ have the same probability law, where $\theta_t \mathbf{p}$ is defined by $D_{\theta_t \mathbf{p}} = \{s \in (0, \infty) : s + t \in D_{\mathbf{p}}\}$ and $(\theta_t \mathbf{p})(s) = \mathbf{p}(s + t)$.

Definition F.3 (Poisson point process). A point process **p** is called *Poisson* if $N_{\mathbf{p}}$ is a Poisson random measure on $(0, \infty) \times X$. A Poisson point process is stationary if and only if its intensity measure $\hat{N}_{\mathbf{p}}(dtdx) = \mathbb{E}[N_{\mathbf{p}}(dtdx)]$ is of the form

$$\hat{N}_{\mathbf{p}}(\mathrm{d}t\mathrm{d}x) = \mathrm{d}t\mathrm{n}(\mathrm{d}x)$$

for some measure n on (X, \mathcal{X}) .

F.1.2. STOCHASTIC INTEGRAL WITH RESPECT TO POINT PROCESS

Here, we review the construction and definition of stochastic integral with respect to Point Processes for completeness; see (Ikeda & Watanabe, 2014, Chapter II, Section 3).

Let $(\Omega, (\mathcal{F}_t)_{t \ge 0}, \mathbb{P})$ be a filtered probability space, where \mathbb{P} is defined on the smallest σ -algebra containing all $\mathcal{F}_t, t \ge 0$ and $(\mathcal{F}_t)_{t \ge 0}$ satisfies the usual conditions (*i.e.*, right-continuity and completeness). A point process $\mathbf{p} = (\mathbf{p}(t))$ on X defined on Ω is called (\mathcal{F}_t) -adapted if for every t > 0 and $U \in \mathcal{X}$, $N_{\mathbf{p}}(t, U) = \sum_{s \in D_{\mathbf{p}}, s \le t} \mathbb{1}_{U}(\mathbf{p}(s))$ is \mathcal{F}_t -measurable. \mathbf{p} is called σ -finite if there exist $U_n \in \mathcal{X}, n = 1, 2, \ldots$ such that $\bigcup_n U_n = X$ and $\mathbb{E}[N_{\mathbf{p}}(t, U_n)] < \infty$ for all t > 0 and $n = 1, 2, \ldots$ For a given (\mathcal{F}_t) -adapted, σ -finite point process \mathbf{p} , let $\Gamma_{\mathbf{p}} = \{U \in \mathcal{X}; \mathbb{E}[N_{\mathbf{p}}(t, U)] < \infty$ for all t > 0}. If $U \in \Gamma_{\mathbf{p}}$, then

 $t \in [0, \infty) \mapsto N_{\mathbf{p}}(t, \mathsf{U})$ is an adapted, integrable increasing process and hence there exists a natural integrable increasing process $\hat{N}_{\mathbf{p}}(t, \mathsf{U})$ such that $\tilde{N}_{\mathbf{p}}: t \in (0, \infty) \mapsto \tilde{N}_{\mathbf{p}}(t, \mathsf{U}) = N_{\mathbf{p}}(t, \mathsf{U}) - \hat{N}_{\mathbf{p}}(t, \mathsf{U})$ is a martingale.

Definition F.4. An (\mathcal{F}_t) -adapted point process \mathbf{p} on $(\Omega, \mathcal{F}, \mathbb{P})$ is said to be *of the class* (QL) (quasi left-continuous) (w.r.t. (\mathcal{F}_t)) if it is σ -finite and there exists $\hat{N}_{\mathbf{p}}$ such that

- 1. for $U \in \Gamma_{\mathbf{p}}, t \in (0, \infty) \mapsto \hat{N}_{\mathbf{p}}(t, U)$ is a continuous (\mathcal{F}_t) -adapted increasing process,
- 2. for each $t \in (0,\infty)$ and a.a. $\omega \in \Omega$, $U \mapsto \hat{N}_{\mathbf{p}}(t, U)$ is a σ -finite measure on $(\mathsf{X}, \mathcal{X})$,
- 3. for $\mathsf{U} \in \Gamma_{\mathbf{p}}, t \in (0,\infty) \mapsto \tilde{N}_{\mathbf{p}}(t,\mathsf{U}) = N_{\mathbf{p}}(t,\mathsf{U}) \hat{N}_{\mathbf{p}}(t,\mathsf{U})$ is an (\mathcal{F}_t) -martingale.

The random measure $\hat{N}_{\mathbf{p}}$ is called the *compensator* of the point process \mathbf{p} (or $N_{\mathbf{p}}$).

Definition F.5. A point process **p** is called an (\mathcal{F}_t) -Poisson point process if it is an (\mathcal{F}_t) -adapted, σ -finite Poisson point process such that $\{N_{\mathbf{p}}(t+h, \mathsf{U}) - N_{\mathbf{p}}(t, \mathsf{U})\}_{h>0,\mathsf{U}\in\mathcal{X}}$ is independent of \mathcal{F}_t for any $t \in [0, \infty)$.

An (\mathcal{F}_t) -Poisson point process is of class (QL) if and only if $t \mapsto \mathbb{E}[N_{\mathbf{p}}(t, \mathsf{U})]$ is continuous for $\mathsf{U} \in \Gamma_{\mathbf{p}}$; in this case, the compensator $\hat{N}_{\mathbf{p}}$ is given by $\hat{N}_{\mathbf{p}}(t, \mathsf{U}) = \mathbb{E}[N_{\mathbf{p}}(t, \mathsf{U})]$.

Theorem F.6. Let \mathbf{p} be a point process of class (QL) w.r.t. (\mathcal{F}_t) on some state space (X, \mathcal{X}) such that its compensator $\hat{N}_{\mathbf{p}}(dtdx)$ is a non-random σ -finite measure on $[0, \infty) \times X$. Then \mathbf{p} is an (\mathcal{F}_t)-Poisson point process. If, in particular, $\hat{N}_{\mathbf{p}}(dtdx) = dtn(dx)$ where n(dx) is a non-random σ -finite measure on X, \mathbf{p} is a stationary (\mathcal{F}_t)-Poisson point process with n as its characteristic measure.

We are now going to discuss stochastic integrals w.r.t. a given point process of the class (QL). For this it is convenient to generalize the notion of predictable processes.

Definition F.7 (Predictable functions). A real function $f(t, x, \omega)$ defined on $[0, \infty) \times X \times \Omega$ is called (\mathcal{F}_t) -predictable if the mapping $(t, x, \omega) \mapsto f(t, x, \omega)$ is $S/\mathcal{B}(\mathbb{R})$ -measurable where S is the smallest σ -field on $[0, \infty) \times X \times \Omega$ w.r.t. which all g having the following properties are measurable:

- 1. for each $t > 0, (x, \omega) \mapsto g(t, x, \omega)$ is $\mathcal{X} \times \mathcal{F}_t$ -measurable;
- 2. for each $(x, \omega), t \in [0, \infty) \mapsto g(t, x, \omega)$ is left continuous.

We introduce the following classes

$$F_{\mathbf{p}} = \left\{ f(t, x, \omega); f \text{ is } (\mathcal{F}_t) \text{-predictable and for each } t > 0, \ \int_0^{t+} \int_{\mathsf{X}} |f(s, x, \omega)| N_{\mathbf{p}}(\mathrm{dsd}x) < \infty \text{ a.s.} \right\},$$

$$F_{\mathbf{p}}^1 = \left\{ f(t, x, \omega); f \text{ is } (\mathcal{F}_t) \text{-predictable and for each } t > 0, \ \mathbb{E} \left[\int_0^t \int_{\mathsf{X}} |f(s, x, \cdot)| \hat{N}_{\mathbf{p}}(\mathrm{dsd}x) \right] < \infty \right\}.$$
(58)

Theorem F.8 (Stochastic integral). For $f \in F_{\mathbf{p}}$, the stochastic integral $\int_0^{t+} \int_X f(s, x, \cdot) N_{\mathbf{p}}(\mathrm{d}s\mathrm{d}x)$ is well-defined a.s. and equals the absolutely convergent sum,

$$\int_0^{t+} \int_{\mathsf{X}} f(s, x, \cdot) N_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x) = \sum_{s \leqslant t, s \in D_{\mathbf{p}}} f(s, \mathbf{p}(s), \cdot) \; .$$

We then denote the stochastic integral w.r.t. $N_{\mathbf{p}}$ as

$$\int_0^{t+} \int_{\mathsf{X}} f(s, x, \cdot) \tilde{N}_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x) = \int_0^{t+} \int_{\mathsf{X}} f(s, x, \cdot) N_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x) - \int_0^t \int_{\mathsf{X}} f(s, x, \cdot) \hat{N}_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x) ,$$

and the process $t \mapsto \int_0^{t+} \int_X |f(s, x, \cdot)| \tilde{N}_{\mathbf{p}}(\mathrm{dsd}x)$ is an (\mathcal{F}_t) -local martingale.

Now let $f \in F_{\mathbf{p}}^1$, the stochastic integral $\int_0^t \int_{\mathsf{X}} f(s, x, \cdot) \hat{N}_{\mathbf{p}}(\mathrm{dsd}x)$ satisfies

$$\mathbb{E}\left[\int_0^{t+} \int_{\mathsf{X}} |f(s,x,\cdot)| N_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x)\right] = \mathbb{E}\left[\int_0^t \int_{\mathsf{X}} |f(s,x,\cdot)| \hat{N}_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x)\right] \ .$$

This implies, in particular, that $F_{\mathbf{p}}^1 \subset F_{\mathbf{p}}$. Then the stochastic integral $t \mapsto \int_0^{t+} \int_{\mathsf{X}} f(s, x, \cdot) \tilde{N}_{\mathbf{p}}(\mathrm{d}s\mathrm{d}x)$ with $f \in F_{\mathbf{p}}^1$ is an (\mathcal{F}_t) -true martingale.

F.1.3. ITÔ'S FORMULA FOR POINT PROCESS

Itô's formula is one of the most important tools in the study of semi-martingales. It provides us with the differential-integral calculus for sample functions of stochastic processes.

Let $(\Omega, (\mathcal{F}_t)_{t \ge 0}, \mathbb{P})$ be a filtered probability space given as above and let (X, \mathcal{X}) be a measurable space. Define a *d*-dimensional semi-martingale $\mathbf{X}_t = (\mathbf{X}_t^1, \mathbf{X}_t^2, \dots, \mathbf{X}_t^d)$ adapted to $(\mathcal{F}_t)_{t \ge 0}$, taking values in X by

$$\mathbf{X}_t^i = \mathbf{X}_0^i + M_t^i + A_t^i + \int_0^{t+} \int_{\mathsf{X}} f^i(s, x, \cdot) N_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x) + \int_0^{t+} \int_{\mathsf{X}} g^i(s, x, \cdot) \tilde{N}_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x) ,$$

for i = 1, 2, ..., d, where

• each $(M_t^i)_{t \ge 0} \in \mathcal{M}_2^{c, \text{loc}}$ with

 $\mathcal{M}_2^{\text{c,loc}} = \{(M_t)_{t \ge 0}; (M_t)_{t \ge 0} \text{ is a locally square integrable } (\mathcal{F}_t) \text{-martingale}, M_0 = 0 \text{ a.s. } ; t \mapsto X_t \text{ is continuous a.s.} \};$

- each $(A_t^i)_{t \ge 0}$ is a continuous (\mathcal{F}_t) -adapted process satisfying $A_0^i = 0$ and for almost every $\omega \in \Omega$, the sample path $t \mapsto A_t^i(\omega)$ has bounded variation on each finite interval;
- **p** is a point process of the class (QL) w.r.t. (\mathcal{F}_t) on some state such that $f^i(t, x, \omega)g^j(t, x, \omega) = 0$ for any $(i, j, t, x, \omega) \in \{1, \dots, d\}^2 \times [0, \infty) \times \mathsf{X} \times \Omega$; furthermore, we assume that a constant C > 0 exists such that

$$|g^i(t, x, \omega)| \leq C$$
 for all $(i, t, x, \omega) \in \{1, \dots, d\} \times [0, \infty) \times \mathsf{X} \times \Omega;$

• each \mathbf{X}_0^i is an \mathcal{F}_0 -measurable random variable.

Denote also $f = (f^1, \ldots, f^d)$ and $g = (g^1, \ldots, g^d)$.

Theorem F.9 (Itô's formula, (Ikeda & Watanabe, 2014, Chapter II, Section 5)). Let F be a function of class C^2 on \mathbb{R}^d and let $(\mathbf{X}_t)_{t \ge 0}$ be a d-dimensional semi-martingale given above. Then the stochastic process $(F(\mathbf{X}_t))_{t \ge 0}$ is also a semi-martingale w.r.t. $(\mathcal{F}_t)_{t \ge 0}$ and the following formula holds:

$$\begin{split} F(\mathbf{X}_{t}) - F(\mathbf{X}_{0}) &= \sum_{i=1}^{d} \int_{0}^{t} \frac{\partial F}{\partial x_{i}} (\mathbf{X}_{s}) \mathrm{d}M_{s}^{i} + \sum_{i=1}^{d} \int_{0}^{t} \frac{\partial F}{\partial x_{i}} (\mathbf{X}_{s}) \mathrm{d}A_{s}^{i} \\ &+ \frac{1}{2} \sum_{i,j=1}^{d} \int_{0}^{t} \frac{\partial F^{2}}{\partial x_{i} \partial x_{j}} (\mathbf{X}_{s}) \mathrm{d}\langle M^{i}, M^{j} \rangle_{s} \\ &+ \int_{0}^{t+} \int_{\mathbf{X}} \left\{ F(\mathbf{X}_{s-} + f(s, x, \cdot)) - F(\mathbf{X}_{s-}) \right\} N_{\mathbf{p}}(\mathrm{d}s\mathrm{d}x) \\ &+ \int_{0}^{t+} \int_{\mathbf{X}} \left\{ F(\mathbf{X}_{s-} + g(s, x, \cdot)) - F(\mathbf{X}_{s-}) \right\} \tilde{N}_{\mathbf{p}}(\mathrm{d}s\mathrm{d}x) \\ &+ \int_{0}^{t} \int_{\mathbf{X}} \left\{ F(\mathbf{X}_{s} + g(s, x, \cdot)) - F(\mathbf{X}_{s}) - \sum_{i=1}^{d} g^{i}(s, x, \cdot) \frac{\partial F}{\partial x_{i}} (\mathbf{X}_{s}) \right\} \hat{N}_{\mathbf{p}}(\mathrm{d}s\mathrm{d}x) \; . \end{split}$$

Corollary F.10 (Itô's isometry). Let $N_{\mathbf{p}}$, $\tilde{N}_{\mathbf{p}}$ and $\hat{N}_{\mathbf{p}}$ given as above. Then the following holds for real-valued functions $f \in F_{\mathbf{p}}^1$,

$$\mathbb{E}\left[\left(\int_0^{t+} \int_{\mathsf{X}} f(s, x, \omega) \tilde{N}_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x)\right)^2\right] = \mathbb{E}\left[\int_0^t \int_{\mathsf{X}} f^2(s, x, \omega) \hat{N}_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x)\right] \,.$$

Proof of Corollary F.10. Denote the process $(\mathbf{X}_t)_{t \ge 0}$ by

$$\mathbf{X}_t = \int_0^{t+} \int_{\mathsf{X}} f(s, x, \omega) \tilde{N}_{\mathbf{p}}(\mathrm{d}s\mathrm{d}x) ,$$

and apply Itô's formula (see Theorem F.9) for function $F(x) = x^2$, we get that

$$\begin{split} \mathbf{X}_t^2 - \mathbf{X}_0^2 &= \int_0^{t+} \int_{\mathbf{X}} \left\{ (\mathbf{X}_{s-} + f(s, x, \omega))^2 - \mathbf{X}_{s-}^2 \right\} \tilde{N}_{\mathbf{p}}(\mathrm{d}s\mathrm{d}x) \\ &+ \int_0^t \int_{\mathbf{X}} \left\{ (\mathbf{X}_s + f(s, x, \omega))^2 - \mathbf{X}_s^2 - 2\mathbf{X}_s f(s, x, \omega) \right\} \hat{N}_{\mathbf{p}}(\mathrm{d}s\mathrm{d}x) , \quad \text{for } t \ge 0 . \end{split}$$

Taking expectations on both sides and using the martingale property of the stochastic integral with respect to $\tilde{N}_{\mathbf{p}}$ implies that

$$\mathbb{E}\left[\mathbf{X}_{t}^{2} - \mathbf{X}_{0}^{2}\right] = \mathbb{E}\left[\int_{0}^{t} \int_{\mathsf{X}} f^{2}(s, x, \omega) \hat{N}_{\mathbf{p}}(\mathrm{d}s\mathrm{d}x)\right] \,.$$

Plugging the formula of X_t into the above and noting that $X_0 = 0$ yields the desired equation,

$$\mathbb{E}\left[\left(\int_0^{t+} \int_{\mathsf{X}} f(s, x, \omega) \tilde{N}_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x)\right)^2\right] = \mathbb{E}\left[\int_0^t \int_{\mathsf{X}} f^2(s, x, \omega) \hat{N}_{\mathbf{p}}(\mathrm{d} s \mathrm{d} x)\right],$$
 poof.

and concludes the proof.

F.1.4. APPLICATION TO CTMCs

Let (X, \mathcal{X}) be a measurable space with $X \subset \mathbb{R}^d$ and let $(X_t)_{t \in [0, T_f]}$ be a CTMC on X associated with the jump rate function $r: (0, T_f] \times X \to \mathbb{R}_+$ and the kernel $k_t(X_{t-}, x)$ determining the probability of jumping into $x \in X$ at time t, given a jump occurs from a current state X_{t-} . The generator \mathfrak{q}_t is defined by $\mathfrak{q}_t(X_{t-}, x) = r_t(X_{t-})k_t(X_{t-}, x)$ represents the rate of jumping from the current state X_{t-} to the new state $x \in X$ at time $t \in (0, T_f]$. The CTMC $(X_t)_{t \in [0, T_f]}$ defines a Poisson point process $\mathbf{p}_X = (\mathbf{p}_X(t))_{t \in [0, T_f]}$ on (X, \mathcal{X}) , where

$$\mathbf{p}_X : D_{\mathbf{p}_X} \subset (0, T_f] \to \mathsf{X} , \quad \mathbf{p}_X(t) = X_t , t \in D_{\mathbf{p}_X}$$

with $D_{\mathbf{p}_X}$ is the set of jump times of $(X_t)_{t \in [0, T_f]}$. We observe that \mathbf{p}_X describes the new state after jumping at time $t \in (0, T_f]$ and it constructs a corresponding Poisson random measure $N_{\mathbf{p}_X}(\mathrm{d}t\mathrm{d}x)$ on $(0, T_f] \times \mathsf{X}$ by

$$\begin{split} N_{\mathbf{p}_X}((0,t]\times\mathsf{U}) &= \operatorname{Card}\left\{s\in D_{\mathbf{p}_X}: s\leqslant t, \mathbf{p}_X(s)\in\mathsf{U}\right\} \\ &= \sum_{s\in D_{\mathbf{p}_X}} \delta_{(s,X_s)}((0,t]\times\mathsf{U}) \qquad \text{ for } t\in(0,T_f], \quad \mathsf{U}\in\mathcal{X} \;, \end{split}$$

that counts the total jumps into $U \in \mathcal{X}$ occurring during (0, t]. Then the random compensator \bar{n}^q of $N_{\mathbf{p}_X}$ is given by

$$\bar{\mathbf{n}}^{\mathfrak{q}}(\mathrm{d} t \mathrm{d} x) = \sum_{y \in \mathsf{X}} \mathfrak{q}(X_{t-}, y) \mathbb{1}_{X_{t-} \neq y} \delta_y(\mathrm{d} x) \mathrm{d} t ,$$

since the corresponding compensated measure

$$\tilde{N}_{\mathbf{p}_X}^{\mathfrak{q}}(\mathrm{d}t\mathrm{d}x) = N_{\mathbf{p}_X}(\mathrm{d}t\mathrm{d}x) - \bar{\mathbf{n}}^{\mathfrak{q}}(\mathrm{d}t\mathrm{d}x)$$

is an (\mathcal{F}_t) -martingale, where $(\mathcal{F}_t)_{t \in [0, T_f]}$ denotes the right-continuous and complete natural filtration generated by the process $(X_t)_{t \in [0, T_f]}$. Indeed, we can show the martingale property of $\tilde{N}_{\mathbf{p}_X}^{\mathfrak{q}}$ as follows. For any function $f \in F_{\mathbf{p}_X}^1$, where the class $F_{\mathbf{p}_X}^1$ is given in (58), define the following stochastic integrals:

$$\begin{split} \int_0^{t+} \int_{\mathsf{X}} f(s, x, \cdot) N_{\mathbf{p}_X}(\mathrm{d}s\mathrm{d}x) &= \sum_{\substack{0 < s \leq t \\ s \in D_{\mathbf{p}_X}}} f(s, \mathbf{p}_X(s), \cdot) \ ,\\ \int_0^{t+} \int_{\mathsf{X}} f(s, x, \cdot) \tilde{N}_{\mathbf{p}_X}^{\mathfrak{q}}(\mathrm{d}s\mathrm{d}x) &= \int_0^{t+} \int_{\mathsf{X}} f(s, x, \cdot) N_{\mathbf{p}_X}(\mathrm{d}s\mathrm{d}x) - \int_0^t \int_{\mathsf{X}} f(s, x, \cdot) \bar{n}^{\mathfrak{q}}(\mathrm{d}s\mathrm{d}x) \ . \end{split}$$

Then for $0 \leqslant s < t \leqslant T_f$, we have

$$\begin{split} \mathbb{E}\left[\int_{s}^{t+}\int_{\mathsf{X}}f(z,x,\cdot)N_{\mathbf{p}_{X}}(\mathrm{d}z\mathrm{d}x)\Big|\mathcal{F}_{s}\right] &= \mathbb{E}\left[\sum_{\substack{s$$

meaning that \bar{n}^q is indeed the compensator of the Poisson random measure $N_{\mathbf{p}_X}$. With those notations in hand, we can decompose the CTMC $(X_t)_{t \in [0,T_f]}$ as

$$X_t = X_0 + \sum_{\substack{0 < s \le t \\ s \in D_{\mathbf{p}_X}}} (X_s - X_{s-}) = X_0 + \int_0^{t+} \int_{\mathbf{X}} (x - X_{s-}) N_{\mathbf{p}_X}(\mathrm{d}s\mathrm{d}x) , \quad \text{for } t \in [0, T_f] ,$$

under the assumption $f(s, x, \omega) := x - \omega_{s-} \in F_{\mathbf{p}_X}$, where $F_{\mathbf{p}_X}$ defined in (58). Applying Itô's formula to this process, for any function F in $C_b^2(X)$, we get that

$$F(X_t) - F(X_0) = \int_0^{t+} \int_X \{F(X_{s-} + x - X_{s-}) - F(X_{s-})\} N_{\mathbf{p}_X}(\mathrm{d}s\mathrm{d}x)$$
$$= \int_0^{t+} \int_X \{F(x) - F(X_{s-})\} N_{\mathbf{p}_X}(\mathrm{d}s\mathrm{d}x)$$

Expressing the Poisson random measure as $N_{\mathbf{p}_X} = \tilde{N}_{\mathbf{p}_X}^{\mathfrak{q}} + \bar{\mathbf{n}}^{\mathfrak{q}}$ and plugging it into the formula above yield

$$F(X_t) - F(X_0) - \int_0^t \int_{\mathbf{X}} \{F(x) - F(X_{s-})\} \,\bar{\mathbf{n}}^{\mathfrak{q}}(\mathrm{d}s\mathrm{d}x) = \int_0^{t+} \int_{\mathbf{X}} \{F(x) - F(X_{s-})\} \,\tilde{N}_{\mathbf{p}_X}^{\mathfrak{q}}(\mathrm{d}s\mathrm{d}x) \,.$$

In other words, the process

$$\left(F(X_t) - F(X_0) - \int_0^t \int_{\mathsf{X}} \{F(x) - F(X_{s-})\} \mathbb{1}_{X_{s-} \neq x} \mathfrak{q}(X_{s-}, \mathrm{d}x) \mathrm{d}s\right)_{t \in [0, T_f]}$$

is an (\mathcal{F}_t) -local martingale as the compensated measure $\tilde{N}_{\mathbf{p}_X}^{\mathfrak{q}}$ was shown to be an (\mathcal{F}_t) -martingale in the previous computation. It follows that for the CTMC $(X_t)_{t \in [0,T_f]}$ with generator \mathfrak{q} , Itô's formula asserts that the process

$$\left(F(X_t) - F(X_0) - \int_0^t \mathfrak{q} F(X_{s-}) \mathrm{d}s\right)_{t \in [0, T_f]}$$

is an (\mathcal{F}_t) -local martingale for any function $F \in C_b^2(X)$. This result aligns with Dynkin's formula.

F.2. Canonical process point of view

We now want to give a description of the time reversal process as a controlled process like in the continuous setting in Conforti et al. (2025). To this purpose, we consider the following canonical setting. Let $\mathbb{D}_{T_f} = \mathbb{D}([0, T_f]; X)$ be the canonical space of all càdlàg (right-continuous with left limits) paths from $[0, T_f]$ to $X = \{0, 1\}^d$, and let $(\mathbf{X}_t)_{t \in [0, T_f]}$ be the canonical process defined by

$$\mathbf{X}_t(\omega) = \omega_t, \quad \text{for } t \in [0, T_f] , \ (\omega_t)_{t \in [0, T_f]} \in \mathbb{D}_{T_f} .$$

We endow this space with the canonical filtration, that is, the right-continuous and complete augmentation of the filtration generated by $(\mathbf{X}_t)_{t \in [0,T_f]}$, denoted by $(\mathcal{F}_t)_{t \in [0,T_f]}$.

For any $\mathbb{P} \in \mathcal{P}(\mathbb{D}_{T_f})$ we denote by $\mathbb{E}_{\mathbb{P}}$ the corresponding expectation. For any $t \in [0, T_f]$, we denote by \mathbb{P}_t , the distribution of \mathbf{X}_t under \mathbb{P} .

As usual convention, we simply denote the random variable $\omega \mapsto U(\omega, t, x)$ as U(t, x) for a predictable process $U : \mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X} \to \mathbb{R}$.

For any random generator $\mathbf{q}: \mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X}^2 \to \mathbb{R}$ and predictable process $U: \mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X} \to \mathbb{R}$, the new random generator $U\mathbf{q}: \mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X}^2 \to \mathbb{R}$ is defined as $(U\mathbf{q})(\omega, t, x, y) := U(\omega, t, y)\mathbf{q}(\omega, t, x, y)$ for $x \neq y$.

For convenience, we denote by $\mathbb{F}(X)$ the set of functions from X to \mathbb{R} .

We now follow the approach for time reversal used in Léonard (2012) to characterize the distribution of $(\mathbf{X}_t)_{t \in [0,T_f]}$ as the solution of a martingale problem.

Definition F.11 (Martingale problem). Let $\mathbf{q} : \mathbb{D}_{T_f} \times [0, T_f] \times X^2$ be a non-homogeneous predictable random generator. We say that $\mathbb{P} \in \mathcal{P}(\mathbb{D}_{T_f})$ solves *the Martingale problem* MP(\mathbf{q}) with initial condition μ_0 , and write $\mathbb{P} \in MP(\mathbf{q})$, if under \mathbb{P} , \mathbf{X}_0 has distribution μ_0 and the process

$$\left(f(\mathbf{X}_t) - f(\mathbf{X}_0) - \int_0^t \mathbf{q}(s)f(\mathbf{X}_{s-}) \mathrm{d}s\right)_{t \in [0, T_f]}$$

is an $(\mathcal{F}_t)_{t \in [0,T_f]}$ -local martingale for any function $f \in \mathbb{F}(\mathsf{X})$, where we denote $\mathbf{q}(t)f(x) = \sum_{y \in \mathsf{X}} \mathbf{q}(\omega, t, x, y)f(y)$.

Note that by Itô's formula and Appendix F.1.4, if under \mathbb{P} , $(\mathbf{X}_t)_{t \in [0,T_f]}$ is a CTMC with generator \mathbf{q} , then \mathbb{P} solves MP(\mathbf{q}). Recall that for the forward process under consideration, the generator \overrightarrow{q} is defined as

$$\overrightarrow{q}(x,y) := \begin{cases} \lambda , & \text{if } y = \varphi^{(\ell)}(x) \text{ for some } \ell \in \{1,\dots,d\} ,\\ -\lambda d , & \text{if } y = x ,\\ 0 , & \text{otherwise }, \end{cases}$$
(59)

where $\varphi^{(\ell)} : \mathsf{X} \to \mathsf{X}$ is the function which flips the ℓ -th component for $\ell \in \{1, \ldots, d\}$. Note that by (15), $\gamma^d = \text{Unif}(\mathsf{X})$ is an invariant distribution for the CTMC with generator \overrightarrow{q} , *i.e.*, for any measurable function f, $\sum_{x \in \mathsf{X}} \overrightarrow{p}_t f(x) \gamma^d(x) = \sum_{x \in \mathsf{X}} f(x) \gamma^d(x)$. In fact the transition density \overrightarrow{p}_t is reversible with respect to γ^d for any $t \in [0, T_f]$ using (15), *i.e.*, for any $x, y \in \mathsf{X}$ and $t \in [0, T_f]$,

$$\gamma^d(x)\overrightarrow{p}_t(x,y) = \gamma^d(y)\overrightarrow{p}_t(y,x)$$
.

As a result, we get that for any $0 \le t_1 < \ldots < t_n \le T_f$, under \overrightarrow{R} , where \overrightarrow{R} denoted the distribution of the CTMC with

generator \overrightarrow{q} started at stationarity γ^d , $(\mathbf{X}_{t_1}, \ldots, \mathbf{X}_{t_n})$ has the same distribution as $(\mathbf{X}_{T_f-t_1}, \ldots, \mathbf{X}_{T_f-t_n})$ and therefore the reference path measure \overrightarrow{R} is reversible, *i.e.*, $\overrightarrow{R} = \overleftarrow{R}$, where \overleftarrow{R} is the distribution of $(\mathbf{X}_{T_f-t_1})_{t \in [0,T_f]}$ under \overrightarrow{R} .

Following Appendix F.1.4, for any $\mathbb{P} \in \mathcal{P}(\mathbb{D}_{T_f})$ such that $(\mathbf{X}_t)_{t \in [0, T_f]}$ is a CTMC with generator $\mathfrak{q} : [0, T_f] \times X^2 \to \mathbb{R}$ under \mathbb{P} , we denote by $\mathbf{p}_{\mathbf{X}}$ the point process and by $N_{\mathbf{X}}$ the Poisson random measure associated with $(\mathbf{X}_t)_{t \in [0, T_f]}$.

We also define for any $(\omega_t)_{t \in [0, T_f]} \in \mathbb{D}_{T_f}$:

$$\bar{\mathbf{n}}^{\mathfrak{q}}((\omega_t)_{t\in[0,T_f]}, \mathrm{d} t \mathrm{d} x) = \mathbf{n}^{\mathfrak{q}}(t, \mathrm{d} x) \mathrm{d} t , \quad \mathbf{n}^{\mathfrak{q}}(t, \mathrm{d} x) = \sum_{y\in\mathsf{X}} \mathbb{1}_{\omega_{t-}\neq y} \mathfrak{q}_t(\omega_{t-}, y) \delta_y(\mathrm{d} x) .$$
(60)

By convention, we denote $\bar{n}^{\mathfrak{q}}((\mathbf{X}_t)_{t\in[0,T_f]}, dtdx)$ by $\bar{n}^{\mathfrak{q}}(dtdx)$ which corresponds to the compensator of $(\mathbf{X}_t)_{t\in[0,T_f]}$ under \mathbb{P} , if under this distribution $(\mathbf{X}_t)_{t\in[0,T_f]}$ is a CTMC with generator $\mathfrak{q}: \mathsf{X}^2 \to \mathbb{R}$. Consequently, the compensated sum of jumps $\tilde{N}_{\mathbf{X}}^{\mathfrak{q}} = N_{\mathbf{X}} - \bar{n}^{\mathfrak{q}}$ forms a martingale under \mathbb{P} .

F.2.1. GIRSANOV'S THEOREM

As proven in Léonard (2012, Theorem 2.6-2.9), the relative entropy of two path measures associated to two jump processes can be expressed using the Young function

$$\varrho(a) := e^a - a - 1$$
, for $a \in \mathbb{R}$.

Note that the convex conjugate ρ is equal to

$$\varrho^*(b) = (b+1)\log(b+1) - b$$
, for $b > -1$,

with the convention $\varrho^*(-1) = 1$ and $\varrho^*(b) = \infty$, for b < -1. We recall that the functions ϱ and ϱ^* are respectively equivalent to $a^2/2$ and $b^2/2$ near zero. Finally, we recall the definition of the function

$$h(a) = \varrho^*(a-1), \qquad a \in \mathbb{R}$$

Theorem F.12 (Girsanov's theorem). (Léonard, 2012, Theorem 2.6-2.9) Let $\mathbb{P} \in \mathcal{P}(\mathbb{D}_{T_f})$ verifying $\mathrm{KL}(\mathbb{P}|\overline{R}) < \infty$. Then, there exists a unique predictable non-negative process $U : \mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X} \to [0, \infty)$ satisfying the integrability condition

$$\mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}} \varrho^*(|U(t,x)-1|)\bar{\mathbf{n}}^{\overrightarrow{q}}(\mathrm{d} t\mathrm{d} x)\right] < \infty , \qquad (61)$$

and $\mathbb{P} \in MP(U\overrightarrow{q})$ where $(U\overrightarrow{q})(\omega, t, x, y) = U(\omega, t, y)\overrightarrow{q}(x, y)$ for $x \neq y$. Moreover, we have that

$$\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\vec{R}}((\mathbf{X}_t)_{t\in[0,T_f]}) = \frac{\mathrm{d}\mathbb{P}_0}{\mathrm{d}\vec{R}_0}(\mathbf{X}_0)\exp\left(\int_{[0,T_f]\times\mathsf{X}}\log U(t,x)\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) - \int_{[0,T_f]\times\mathsf{X}}\varrho(\log U(t,x))\bar{n}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x)\right)$$

and the KL divergence reads as

$$\mathrm{KL}(\mathbb{P}|\overrightarrow{R}) = \mathrm{KL}(\mathbb{P}_0|\overrightarrow{R}_0) + \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}} h(U(t,x))\bar{\mathbf{n}}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x)\right]$$

The proof of Theorem F.12 is given for completeness and is based on several technical lemmas, which we introduce in the following framework. Let $\mathbb{P}^q \in \mathcal{P}(\mathbb{D}_{T_f})$ such that $(\mathbf{X}_t)_{t \in [0,T_f]}$ is a CTMC with generator $q : [0, T_f] \times X^2 \to \mathbb{R}$, *i.e.*, $\sum_{y \in \mathbf{X}} q(t, x, y) = 0$ for any $(t, x) \in [0, T_f] \times X$, and denote $\mathbf{p}_{\mathbf{X}}$, $N_{\mathbf{X}}$, \bar{n}^q , n^q and $\tilde{N}_{\mathbf{X}}^q$ as in previous Section.

Let χ be a \mathbb{R} -valued predictable process on $\mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X}$ such that $\int_{[0, T_f] \times \mathsf{X}} \varrho(\chi_t(\omega, x)) \bar{\mathrm{n}}^{\mathfrak{q}}(\mathrm{d}t \mathrm{d}x) < \infty$, $\mathbb{P}^{\mathfrak{q}}$ -a.s. Define

$$Z_t^{\chi} := \exp\left(\int_{[0,t]\times\mathsf{X}} \chi_s(\omega, x) \tilde{N}_{\mathbf{X}}^{\mathfrak{q}}(\mathrm{d} s \mathrm{d} x) - \int_{[0,t]\times\mathsf{X}} \varrho(\chi_s(\omega, x)) \bar{\mathrm{n}}^{\mathfrak{q}}(\mathrm{d} s \mathrm{d} x)\right) , \quad \text{for } t \in [0, T_f] .$$
(62)

Lemma F.13. Let $\mathbb{P}^q \in \mathcal{P}(\mathbb{D}_{T_f})$ such that $(\mathbf{X}_t)_{t \in [0, T_f]}$ is a CTMC with generator $\mathfrak{q} : [0, T_f] \times X^2 \to \mathbb{R}$. Assume that χ is a \mathbb{R} -valued predictable process on $\mathbb{D}_{T_f} \times [0, T_f] \times X$ satisfying the integrability condition

$$\mathbb{E}_{\mathbb{P}^{q}} \int_{[0,T_{f}] \times \mathsf{X}} \varrho(\chi_{s}(\omega, x)) \bar{\mathbf{n}}^{q}(\mathrm{d} s \mathrm{d} x) < \infty .$$
(63)

Then $\int_{[0,t]\times \mathbf{X}} \chi_s(\omega, x) \tilde{N}^{\mathfrak{q}}_{\mathbf{X}}(\mathrm{dsd}x)$ is a local $\mathbb{P}^{\mathfrak{q}}$ -martingale. Moreover, the process $(\mathbf{Z}^{\chi}_t)_{t\in[0,T_f]}$ defined in (62) is a local $\mathbb{P}^{\mathfrak{q}}$ -martingale and a positive $\mathbb{P}^{\mathfrak{q}}$ -supermartingale, which satisfies

$$\mathrm{d} \mathbf{Z}_t^{\chi} = \mathbf{Z}_{t-}^{\chi} \int_{\mathsf{X}} (\mathrm{e}^{\chi_t(\omega, x)} - 1) \tilde{N}_{\mathbf{X}}^{\mathfrak{q}}(\mathrm{d} t \mathrm{d} x) \; .$$

Proof of Lemma F.13. By Theorem F.8, the process

$$M_t^{\chi} := \int_{[0,t] \times \mathbf{X}} \chi_s(\omega, x) \tilde{N}_{\mathbf{X}}^{\mathfrak{q}}(\mathrm{d} s \mathrm{d} x)$$

is a local \mathbb{P}^{q} -martingale. Denote $Y_{t}^{\chi} := M_{t}^{\chi} - \int_{[0,t]} \beta_{s} ds$ with $\beta_{s} := \int_{\mathsf{X}} \varrho(\chi_{s}(\omega, x)) n^{q}(s, dx)$. Applying Itô's formula provided in Theorem F.9 for the jump process $(Y_{t}^{\chi})_{t \in [0,T_{f}]}$ and for a function f of class $C^{2}(\mathbb{R})$ implies

$$\begin{split} \mathrm{d}f(\mathbf{Y}_t^{\chi}) &= \left[\int_{\mathbf{X}} \left[f(\mathbf{Y}_{t-}^{\chi} + \chi_t(\omega, x)) - f(\mathbf{Y}_{t-}^{\chi}) - f'(\mathbf{Y}_{t-}^{\chi}) \cdot \chi_t(\omega, x) \right] \mathbf{n}^{\mathfrak{q}}(t, \mathrm{d}x) \right] \mathrm{d}t \\ &+ f'(\mathbf{Y}_{t-}^{\chi}) \cdot \beta_t \mathrm{d}t + \mathrm{d}M_t^f \ , \quad \mathbb{P}^{\mathfrak{q}}\text{-a.s.} \ , \end{split}$$

where M_t is given by

$$M_t^f = \int_{[0,t]\times\mathsf{X}} \left[f(Y_{s-}^{\chi} + \chi_s(\omega, x)) - f(Y_{s-}^{\chi}) \right] \tilde{N}_{\mathbf{X}}^{\mathfrak{q}}(\mathrm{d} s \mathrm{d} x)$$

is a local \mathbb{P}^q -martingale, since the integrand is \mathbb{R} -valued predictable process and $\tilde{N}^q_{\mathbf{X}}$ forms a martingale under \mathbb{P}^q . Using this formula for $f(y) = e^y$, we obtain

$$\begin{split} \mathrm{d}\mathrm{e}^{\mathrm{Y}_{t}^{\chi}} &= \left[\int_{\mathsf{X}} (\mathrm{e}^{\mathrm{Y}_{t-}^{\chi} + \chi_{t}(\omega, x)} - \mathrm{e}^{\mathrm{Y}_{t-}^{\chi}} - \mathrm{e}^{\mathrm{Y}_{t-}^{\chi}} \cdot \chi_{t}(\omega, x)) \mathrm{n}^{\mathfrak{q}}(t, \mathrm{d}x) \right] \mathrm{d}t - \mathrm{e}^{\mathrm{Y}_{t-}^{\chi}} \beta_{t} \mathrm{d}t + \mathrm{d}M_{t}^{\exp} \\ &= \mathrm{e}^{\mathrm{Y}_{t-}^{\chi}} \beta_{t} \mathrm{d}t - \mathrm{e}^{\mathrm{Y}_{t-}^{\chi}} \beta_{t} \mathrm{d}t + \mathrm{d}M_{t}^{\exp} = \mathrm{d}M_{t}^{\exp} \ , \quad \mathbb{P}^{\mathfrak{q}}\text{-a.s.} \ . \end{split}$$

This implies $Z_t^{\chi} = e^{Y_t^{\chi}}$ is a local \mathbb{P}^q -martingale and, since Z_t^{χ} is positive, we can conclude that Z_t^{χ} is a \mathbb{P}^q -supermartingale thanks to Fatou's lemma. In addition, we have

$$\mathrm{d}M_t^{\mathrm{exp}} = \int_{\mathsf{X}} \left(\mathrm{e}^{Y_{t-}^{\chi} + \chi_t(\omega, x)} - \mathrm{e}^{Y_{t-}^{\chi}} \right) \tilde{N}_{\mathbf{X}}^{\mathfrak{q}}(\mathrm{d}t\mathrm{d}x) = \mathrm{e}^{Y_{t-}^{\chi}} \int_{\mathsf{X}} (\mathrm{e}^{\chi_t(\omega, x)} - 1) \tilde{N}_{\mathbf{X}}^{\mathfrak{q}}(\mathrm{d}t\mathrm{d}x) ,$$

i.e., $dZ_t^{\chi} = Z_{t-}^{\chi} \int_{\mathbf{X}} (e^{\chi_t(\omega, x)} - 1) \tilde{N}_{\mathbf{X}}^{\mathfrak{q}}(dt dx)$ and we conclude the proof of Lemma F.13.

We now define the stopping time for $k, j \ge 1$,

$$\sigma_j^k := \inf \left\{ t \in [0, T_f] : \int_{[0,t] \times \mathsf{X}} \varrho(\chi_s(\omega, x)) \bar{\mathbf{n}}^{\mathfrak{q}}(\mathrm{d}s\mathrm{d}x) \ge k \text{ or } \chi_t(\omega, \omega_t) \notin [-j,k] \right\} .$$
(64)

For $\mathbb{P} \in \mathcal{P}(\mathbb{D}_{T_f})$, let us denote $\mathbb{P}^{\sigma_j^k} = \mathbf{X}_{\#}^{\sigma_j^k} \mathbb{P}$ the law under \mathbb{P} of the process $\mathbf{X}^{\sigma_j^k}$ which is stopped at the stopping time σ_j^k . **Lemma F.14.** Let $\mathbb{P}^q \in \mathcal{P}(\mathbb{D}_{T_f})$ such that $(\mathbf{X}_t)_{t \in [0, T_f]}$ is a CTMC with generator $\mathfrak{q} : [0, T_f] \times X^2 \to \mathbb{R}$ and let χ be a \mathbb{R} -valued predictable process on $\mathbb{D}_{T_f} \times [0, T_f] \times X$ satisfying the integrability condition (63). Let $(\mathbf{Z}_t^{\chi})_{t \in [0, T_f]}$ be

defined in (62) and σ_j^k be defined in (64). For all $j, k \ge 1$, the process $(Z_t^{\sigma_j^k})_{t \in [0,T_f]}$ defined as $Z_t^{\sigma_j^k} := Z_t^{\chi_j^k}$, is a genuine \mathbb{P}^q -martingale with $\chi_j^k = \mathbb{1}_{[0,\sigma_j^k]}\chi$, and the measure \mathbb{Q}_j^k defined for any measurable function $F : \mathbb{D}_{T_f} \to \mathbb{R}_+$ by

$$\mathbb{E}_{\mathbb{Q}_j^k}[F((\mathbf{X}_t)_{t\in[0,T_f]})] = \mathbb{E}_{\mathbb{P}^q}[F((\mathbf{X}_{t\wedge\sigma_j^k})_{t\in[0,T_f]})\mathbf{Z}_{T_f}^{\sigma_j^k}], \quad i.e. \ , \quad \mathbb{Q}_j^k = \mathbf{Z}_{T_f}^{\sigma_j^k}(\mathbb{P}^q)^{\sigma_j^k}$$

is a probability measure on \mathbb{D}_{T_f} which satisfies

$$\mathbb{Q}_{j}^{k} \in \mathrm{MP}(\mathbb{1}_{\left[0,\sigma_{j}^{k}\right]} \mathrm{e}^{\chi} \mathfrak{q}) \ .$$

Proof of Lemma F.14. Fix $j, k \ge 1$. We have

$$\mathbf{Z}_{t}^{\sigma_{j}^{k}} = \exp\left(\int_{[0,t]\times\mathbf{X}} \chi_{j}^{k} \mathrm{d}\tilde{N}_{\mathbf{X}}^{\mathfrak{q}} - \int_{[0,t]\times\mathbf{X}} \varrho(\chi_{j}^{k}) \mathrm{d}\bar{\mathbf{n}}^{\mathfrak{q}}\right) \;,$$

where $\chi_j^k = \mathbb{1}_{[0,\sigma_j^k]}\chi$ is predictable since χ is predictable and $\mathbb{1}_{[0,\sigma_j^k]}$ is left continuous. For simplicity, we drop the subscripts and superscripts and write $\tilde{\chi} = \chi_j^k$ and $\tilde{Z}_t = Z_t^{\sigma_j^k}$ for the rest of the proof. From the definition of σ_j^k , we obtain that \mathbb{P}^q a.s.,

$$\int_{[0,t]\times\mathsf{X}} \varrho(\tilde{\chi}_s) \mathrm{d}\bar{\mathrm{n}}^{\mathfrak{q}} \leqslant k , \quad \text{and } \tilde{\chi}_t(\omega, \mathbf{X}_t) \in [-j, k] , \text{ for any } t \in [0, T_f] .$$
(65)

First, we prove that (\tilde{Z}_t) is a \mathbb{P}^q -martingale. From Lemma F.13, (\tilde{Z}_t) is a local martingale. Therefore, it is enough to show that for $t \in [0, T_f]$,

$$\mathbb{E}_{\mathbb{P}^q}[\tilde{\mathbf{Z}}_t^p] < \infty$$
, for some $p > 1$.

For p > 1, we have

$$\tilde{\mathbf{Z}}_{t}^{p} = \exp\left(p\int_{[0,t]\times\mathsf{X}}\tilde{\chi}_{s}\mathrm{d}\tilde{N}_{\mathbf{X}}^{\mathsf{q}} - p\int_{[0,t]\times\mathsf{X}}\varrho(\tilde{\chi}_{s})\mathrm{d}\bar{\mathbf{n}}^{\mathsf{q}}\right) \leqslant \exp\left(p\int_{[0,t]\times\mathsf{X}}\tilde{\chi}_{s}\mathrm{d}\tilde{N}_{\mathbf{X}}^{\mathsf{q}}\right)$$

and

$$\exp\left(p\int_{[0,t]\times\mathbf{X}}\tilde{\chi}_{s}\mathrm{d}\tilde{N}_{\mathbf{X}}^{\mathfrak{q}} - \int_{[0,t]\times\mathbf{X}}\varrho(p\tilde{\chi}_{s})\mathrm{d}\bar{n}^{\mathfrak{q}}\right) \ge \exp\left(p\int_{[0,t]\times\mathbf{X}}\tilde{\chi}_{s}\mathrm{d}\tilde{N}_{\mathbf{X}}^{\mathfrak{q}}\right) \middle/ C(k,p,t) ,$$

for some finite deterministic constant $0 < C(k, p, t) < \infty$. Indeed, \mathbb{P}^{q} -a.s., for any $s \in [0, T_{f}]$, it holds that $\varrho(p\tilde{\chi}_{s}) \leq e^{k(p-1)}(\varrho(\tilde{\chi}_{s}) + k + 1)$ since $\tilde{\chi}_{s} \leq k$ and p > 1. It yields that \mathbb{P}^{q} -a.s., it holds

$$\exp\left(\int_{[0,t]\times\mathsf{X}} \varrho(p\tilde{\chi}_s)\mathrm{d}\bar{\mathrm{n}}^{\mathfrak{q}}\right) \leqslant \exp\left(\mathrm{e}^{k(p-1)}\int_{[0,t]\times\mathsf{X}} (\varrho(\tilde{\chi}_s)+k+1)\mathrm{d}\bar{\mathrm{n}}^{\mathfrak{q}}\right)$$

$$\stackrel{(65)}{\leqslant} \exp\left(k\mathrm{e}^{k(p-1)}+(k+1)\mathrm{e}^{k(p-1)}\int_{[0,t]\times\mathsf{X}} 1\mathrm{d}\bar{\mathrm{n}}^{\mathfrak{q}}\right) \leqslant C(k,p,t) < \infty ,$$

where the last inequality follows from the formula of \bar{n}^q given in (60) and the fact that X is finite. This implies \mathbb{P}^{q} -a.s.,

$$\tilde{Z}_{t}^{p} \leqslant \exp\left(p\int_{[0,t]\times\mathbf{X}}\tilde{\chi}_{s}\mathrm{d}\tilde{N}_{\mathbf{X}}^{\mathfrak{q}}\right) \leqslant C(k,p,t)\exp\left(p\int_{[0,t]\times\mathbf{X}}\tilde{\chi}_{s}\mathrm{d}\tilde{N}_{\mathbf{X}}^{\mathfrak{q}} - \int_{[0,t]\times\mathbf{X}}\varrho(p\tilde{\chi}_{s})\mathrm{d}\bar{\mathbf{n}}^{\mathfrak{q}}\right) \,. \tag{66}$$

On the other hand, applying Lemma F.13 for $p\tilde{\chi}$ yields that $\exp\left(p\int_{[0,t]\times X} \tilde{\chi}_s d\tilde{N}^{\mathfrak{q}}_{\mathbf{X}} - \int_{[0,t]\times X} \varrho(p\tilde{\chi}_s) d\bar{\mathbb{n}}^{\mathfrak{q}}\right)$ is a $\mathbb{P}^{\mathfrak{q}}$ -

supermartingale, and we get

$$\mathbb{E}_{\mathbb{P}^{\mathfrak{q}}}\left[\exp\left(p\int_{[0,t]\times\mathsf{X}}\tilde{\chi}_{s}\mathrm{d}\tilde{N}_{\mathbf{X}}^{\mathfrak{q}}-\int_{[0,t]\times\mathsf{X}}\varrho(p\tilde{\chi}_{s})\mathrm{d}\bar{n}^{\mathfrak{q}}\right)\right] \leqslant \mathbb{E}_{\mathbb{P}^{\mathfrak{q}}}\left[\exp\left(p\int_{[0,0]\times\mathsf{X}}\tilde{\chi}_{s}\mathrm{d}\tilde{N}_{\mathbf{X}}^{\mathfrak{q}}-\int_{[0,0]\times\mathsf{X}}\varrho(p\tilde{\chi}_{s})\mathrm{d}\bar{n}^{\mathfrak{q}}\right)\right]=1.$$

Plugging this estimate into (66) gives

$$\mathbb{E}_{\mathbb{P}^{\mathsf{q}}}[\tilde{\mathbf{Z}}_t^p] \leqslant C(k,p,t) < \infty \;, \quad \text{for any } t \in [0,T_f] \;,$$

which allow us to conclude that (\tilde{Z}_{T_f}) is a \mathbb{P}^q -martingale (see, *e.g.*, Zitkovic, 2015). Thereby $\mathbb{E}_{\mathbb{P}^q}[\tilde{Z}_t] = \mathbb{E}_{\mathbb{P}^q}[\tilde{Z}_0] = 1$ for any $t \in [0, T_f]$ and it follows \mathbb{Q}_j^k is a probability measure on \mathbb{D}_{T_f} .

Now, we show the second claim of Lemma F.14:

$$\mathbb{Q}_{j}^{k} \in \mathrm{MP}(\mathbb{1}_{[0,\sigma_{j}^{k}]}\mathrm{e}^{\chi}\mathfrak{q}) \ .$$

Let τ be a finitely valued stopping time which will be specified later, and for any function $f \in \mathbb{F}(X)$, we denote

$$F_t = \sum_{0 \le s \le t} \{ f(\mathbf{X}_s) - f(\mathbf{X}_{s-}) \} = \int_0^{t+} \int_{\mathsf{X}} \{ f(x) - f(\mathbf{X}_{s-}) \} N_{\mathbf{X}}(\mathrm{d}s\mathrm{d}x) , \quad \text{for } t \in [0, T_f] .$$

Recall that by Lemma F.13, the martingale (\tilde{Z}_t) satisfies the followings for \mathbb{P}^{q} -a.s.,

$$\mathrm{d}\tilde{\mathbf{Z}}_{t} = \mathbb{1}_{\left[0,\sigma_{j}^{k}\right]}(t)\tilde{\mathbf{Z}}_{t-}\int_{\mathsf{X}}(\mathrm{e}^{\tilde{\chi}_{t}(\omega,x)}-1)\tilde{N}_{\mathbf{X}}^{\mathfrak{q}}(\mathrm{d}t\mathrm{d}x)$$

We have

$$\mathbb{E}_{\mathbb{Q}_{j}^{k}}\left[\sum_{0\leqslant s\leqslant t\wedge\tau}\left\{f(\mathbf{X}_{s})-f(\mathbf{X}_{s-})\right\}\middle|\mathcal{F}_{0}\right]=\mathbb{E}_{\mathbb{Q}_{j}^{k}}\left[\mathbf{F}_{t\wedge\tau}\middle|\mathcal{F}_{0}\right]=\mathbb{E}_{\mathbb{P}^{\mathsf{q}}}\left[\tilde{\mathbf{Z}}_{t\wedge\tau\wedge\sigma_{j}^{k}}\mathbf{F}_{t\wedge\tau\wedge\sigma_{j}^{k}}-\tilde{\mathbf{Z}}_{0}\mathbf{F}_{0}\middle|\mathcal{F}_{0}\right].$$
(67)

Let us denote the two-dimensional process $(I_t)_{t \in [0,T_f]} = (I_t^1, I_t^2)_{t \in [0,T_f]}$, where

$$\mathbf{I}_{t}^{1} := \mathbf{F}_{t} = \int_{[0,t]\times\mathsf{X}} \underbrace{\left[f(x) - f(\mathbf{X}_{s-})\right]}_{=:v^{1}(s,x,\omega)} \tilde{N}_{\mathbf{X}}^{\mathfrak{q}}(\mathrm{d}s\mathrm{d}x) + \underbrace{\int_{[0,t]\times\mathsf{X}} [f(x) - f(\mathbf{X}_{s-})]\bar{\mathbf{n}}^{\mathfrak{q}}(\mathrm{d}s\mathrm{d}x)}_{=:A_{t}^{1}} + \underbrace{\int_{[0,t]\times\mathsf{X}} [f(x) - f(\mathbf{X}_{s-})]\bar$$

and

$$\mathbf{I}_t^2 = \tilde{\mathbf{Z}}_t = \int_{[0,t]\times\mathbf{X}} \underbrace{\tilde{\mathbf{Z}}_{s-}(\mathbf{e}^{\tilde{\chi}_s(\omega,x)} - 1)}_{::=v^2(s,x,\omega)} \tilde{N}_{\mathbf{X}}^{\mathfrak{q}}(\mathrm{d}s\mathrm{d}x) \; .$$

Let $v = (v^1, v^2)$ and apply Itô's formula (see Theorem F.9) to the process $(I_t)_{t \in [0, T_f]}$ using the function given by the product of the coordinates, treating $(A_t^1)_{t \in [0, T_f]}$ as a continuous, finite variation process adapted to the filtration (\mathcal{F}_t) ,

$$\begin{split} \mathbb{E}_{\mathbb{P}^{q}} \left[\mathbb{F}_{t\wedge\tau\wedge\sigma_{j}^{k}}\tilde{Z}_{t\wedge\tau\wedge\sigma_{j}^{k}} - \mathbb{F}_{0}\tilde{Z}_{0} \Big| \mathcal{F}_{0} \right] &= \mathbb{E}_{\mathbb{P}^{q}} \left[I_{t\wedge\tau\wedge\sigma_{j}^{k}}^{1} I_{t\wedge\tau\wedge\sigma_{j}^{k}}^{2} - I_{0}^{1} I_{0}^{2} \Big| \mathcal{F}_{0} \right] \\ &= \mathbb{E}_{\mathbb{P}^{q}} \left[\int_{[0,t\wedge\tau\wedge\sigma_{j}^{k}]\times\mathbf{X}} \left\{ (\mathbf{I}_{s-}^{1} + v^{1}(s,x,\omega))(\mathbf{I}_{s-}^{2} + v^{2}(s,x,\omega)) - \mathbf{I}_{s-}^{1} \mathbf{I}_{s-}^{2} \right\} \tilde{N}_{\mathbf{X}}^{\mathbf{q}}(\mathrm{d}s\mathrm{d}x) \Big| \mathcal{F}_{0} \right] \\ &+ \mathbb{E}_{\mathbb{P}^{q}} \left[\int_{[0,t\wedge\tau\wedge\sigma_{j}^{k}]\times\mathbf{X}} \left\{ (\mathbf{I}_{s}^{1} + v^{1}(s,x,\omega))(\mathbf{I}_{s}^{2} + v^{2}(s,x,\omega)) - \mathbf{I}_{s}^{1} \mathbf{I}_{s}^{2} - \mathbf{I}_{s}^{2} v^{1}(s,x,\omega) - \mathbf{I}_{s}^{1} v^{2}(s,x,\omega) \right\} \bar{\mathbf{n}}^{\mathbf{q}}(\mathrm{d}s\mathrm{d}x) \Big| \mathcal{F}_{0} \right] \\ &+ \mathbb{E}_{\mathbb{P}^{q}} \left[\int_{[0,t\wedge\tau\wedge\sigma_{j}^{k}]\times\mathbf{X}} \left\{ \mathbf{I}_{s}^{2} v^{1}(s,x,\omega) + v^{1}(s,x,\omega) v^{2}(s,x,\omega) \right\} \bar{\mathbf{n}}^{\mathbf{q}}(\mathrm{d}s\mathrm{d}x) \Big| \mathcal{F}_{0} \right] \\ &= \mathbb{E}_{\mathbb{P}^{q}} \left[\int_{[0,t\wedge\tau\wedge\sigma_{j}^{k}]\times\mathbf{X}} \left\{ \mathbf{I}_{s}^{2} v^{1}(s,x,\omega) + v^{1}(s,x,\omega) v^{2}(s,x,\omega) \right\} \bar{\mathbf{n}}^{\mathbf{q}}(\mathrm{d}s\mathrm{d}x) \Big| \mathcal{F}_{0} \right] \quad (\mathrm{as} \ \tilde{N}_{\mathbf{X}}^{\mathbf{q}} \ \mathrm{is a} \ \mathbb{P}^{\mathbf{q}} - \mathrm{martingale}) \\ &= \mathbb{E}_{\mathbb{P}^{q}} \left[\int_{[0,t\wedge\tau\wedge\sigma_{j}^{k}]\times\mathbf{X}} \left\{ \tilde{\mathbf{Z}}_{s} v^{1}(s,x,\omega) + v^{1}(s,x,\omega) \tilde{\mathbf{Z}}_{s-} \left(\mathrm{e}^{\tilde{\chi}_{s}(\omega,x)} - \mathbf{1} \right) \right\} \bar{\mathbf{n}}^{\mathbf{q}}(\mathrm{d}s\mathrm{d}x) \Big| \mathcal{F}_{0} \right] , \end{split}$$

where we reduce the stochastic integral w.r.t. $\tilde{N}_{\mathbf{X}}^{\mathfrak{q}}$ as it is a local $\mathbb{P}^{\mathfrak{q}}$ -martingale, since the integrand is \mathbb{R} -valued predictable process and $\tilde{N}_{\mathbf{X}}^{\mathfrak{q}}$ forms a martingale under $\mathbb{P}^{\mathfrak{q}}$. Since $\tilde{Z}_{s} = \tilde{Z}_{s-}$ for Lebesgue almost all $s \in [0, t \wedge \tau \wedge \sigma_{j}^{k}]$ (Mozumder, 2009, Proposition 2.1) and $\bar{n}^{\mathfrak{q}}$ is atomless in time, the calculation follows

$$\mathbb{E}_{\mathbb{P}^{\mathfrak{q}}}\left[F_{t\wedge\tau\wedge\sigma_{j}^{k}}\tilde{Z}_{t\wedge\tau\wedge\sigma_{j}^{k}}-F_{0}\tilde{Z}_{0}\middle|\mathcal{F}_{0}\right] = \mathbb{E}_{\mathbb{P}^{\mathfrak{q}}}\left[\int_{[0,t\wedge\tau\wedge\sigma_{j}^{k}]\times\mathsf{X}}\tilde{Z}_{s}v^{1}(s,x,\omega)\left(1+\mathrm{e}^{\tilde{\chi}_{s}(\omega,x)}-1\right)\bar{n}^{\mathfrak{q}}(\mathrm{d}s\mathrm{d}x)\middle|\mathcal{F}_{0}\right] \\
= \mathbb{E}_{\mathbb{P}^{\mathfrak{q}}}\left[\int_{[0,t\wedge\tau\wedge\sigma_{j}^{k}]\times\mathsf{X}}\tilde{Z}_{s}v^{1}(s,x,\omega)\mathrm{e}^{\tilde{\chi}_{s}(\omega,x)}\bar{n}^{\mathfrak{q}}(\mathrm{d}s\mathrm{d}x)\middle|\mathcal{F}_{0}\right].$$
(68)

Denote $G_t := \int_{[0,t] \times X} v^1(s, x, \omega) e^{\tilde{\chi}_s(\omega, x)} \bar{n}^{\mathfrak{q}}(dsdx)$. Applying Itô's formula for the process (G_t, \tilde{Z}_t) analogously as argued before, we obtain that

$$\mathbb{E}_{\mathbb{P}^{q}} \left[\tilde{Z}_{t\wedge\tau\wedge\sigma_{j}^{k}} \mathbf{G}_{t\wedge\tau\wedge\sigma_{j}^{k}} - \tilde{Z}_{0}\mathbf{G}_{0} \middle| \mathcal{F}_{0} \right] \\
= \mathbb{E}_{\mathbb{P}^{q}} \left[\int_{[0,t\wedge\tau\wedge\sigma_{j}^{k}]\times\mathbf{X}} \left\{ \mathbf{G}_{s-}(\tilde{Z}_{s-} + v^{2}(s,x,\omega)) - \mathbf{G}_{s-}\tilde{Z}_{s-} \right\} \tilde{N}_{\mathbf{X}}^{\mathbf{q}}(\mathrm{d}s\mathrm{d}x) \middle| \mathcal{F}_{0} \right] \\
+ \mathbb{E}_{\mathbb{P}^{q}} \left[\int_{[0,t\wedge\tau\wedge\sigma_{j}^{k}]\times\mathbf{X}} \left\{ (\mathbf{G}_{s}(\tilde{Z}_{s} + v^{2}(s,x,\omega)) - \mathbf{G}_{s}\tilde{Z}_{s} - \mathbf{G}_{s}v^{2}(s,x,\omega) \right\} \bar{\mathbf{n}}^{\mathbf{q}}(\mathrm{d}s\mathrm{d}x) \middle| \mathcal{F}_{0} \right] \\
+ \mathbb{E}_{\mathbb{P}^{q}} \left[\int_{[0,t\wedge\tau\wedge\sigma_{j}^{k}]\times\mathbf{X}} \tilde{Z}_{s}v^{1}(s,x,\omega) \mathrm{e}^{\tilde{\chi}_{s}(\omega,x)} \bar{\mathbf{n}}^{\mathbf{q}}(\mathrm{d}s\mathrm{d}x) \middle| \mathcal{F}_{0} \right] \\
= \mathbb{E}_{\mathbb{P}^{q}} \left[\int_{[0,t\wedge\tau\wedge\sigma_{j}^{k}]\times\mathbf{X}} \tilde{Z}_{s}v^{1}(s,x,\omega) \mathrm{e}^{\tilde{\chi}_{s}(\omega,x)} \bar{\mathbf{n}}^{\mathbf{q}}(\mathrm{d}s\mathrm{d}x) \middle| \mathcal{F}_{0} \right],$$
(69)

as the stochastic integral w.r.t. $\tilde{N}_{\mathbf{X}}^{\mathfrak{q}}$ is a local $\mathbb{P}^{\mathfrak{q}}$ -martingale, since the integrand is \mathbb{R} -valued predictable process and $\tilde{N}_{\mathbf{X}}^{\mathfrak{q}}$ forms a martingale under $\mathbb{P}^{\mathfrak{q}}$. Combining (67), (68) and (69) implies

$$\begin{split} \mathbb{E}_{\mathbb{Q}_{j}^{k}} \left[\sum_{0 \leqslant s \leqslant t \wedge \tau} \{f(\mathbf{X}_{s}) - f(\mathbf{X}_{s-})\} \middle| \mathcal{F}_{0} \right] &= \mathbb{E}_{\mathbb{P}^{q}} \left[\tilde{Z}_{t \wedge \tau \wedge \sigma_{j}^{k}} \mathbf{G}_{t \wedge \tau \wedge \sigma_{j}^{k}} \middle| \mathcal{F}_{0} \right] \quad (\text{since } \mathbf{G}_{0} = 0) \\ &= \mathbb{E}_{\mathbb{Q}_{j}^{k}} \left[\mathbf{G}_{t \wedge \tau} \middle| \mathcal{F}_{0} \right] \\ &= \mathbb{E}_{\mathbb{Q}_{j}^{k}} \left[\int_{[0, t \wedge \tau] \times \mathsf{X}} \{f(x) - f(\mathbf{X}_{s-})\} \mathbf{e}^{\tilde{\chi}_{s}(\omega, x)} \bar{\mathbf{n}}^{\mathfrak{q}}(\mathrm{d}s\mathrm{d}x) \middle| \mathcal{F}_{0} \right] \\ &= \mathbb{E}_{\mathbb{Q}_{j}^{k}} \left[\int_{[0, t \wedge \tau] \times \mathsf{X}} \{f(x) - f(\mathbf{X}_{s-})\} \mathbb{1}_{\mathbf{X}_{s-} \neq x} \mathbf{e}^{\tilde{\chi}_{s}(\omega, x)} \mathfrak{q}_{s}(\mathbf{X}_{s-}, \mathrm{d}x) \mathrm{d}s \middle| \mathcal{F}_{0} \right] \,. \end{split}$$

Recall that the random generator $e^{\tilde{\chi}}\mathfrak{q}: \mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X}^2 \to \mathbb{R}$ is defined by $(e^{\tilde{\chi}}\mathfrak{q})(\omega, t, x, y) := e^{\tilde{\chi}(\omega, t, y)}\mathfrak{q}(w, t, x, y) = e^{\tilde{\chi}_t(\omega, y)}\mathfrak{q}(t, x, y)$ for $y \neq x$, since the generator \mathfrak{q} under consideration is deterministic. Denote by $\bar{n}^{e^{\tilde{\chi}}\mathfrak{q}}$ the corresponding jump kernel for $\omega = (\omega_t)_{t \in [0, T_f]} \in \mathbb{D}_{T_f}$ and $(t, x) \in [0, T_f] \times \mathsf{X}$,

$$\begin{split} \bar{\mathbf{n}}^{\mathrm{e}^{\chi}\mathfrak{q}}(\omega, \mathrm{d}t\mathrm{d}x) &:= \mathbb{1}_{\omega_{t-} \neq x}(\mathrm{e}^{\tilde{\chi}}\mathfrak{q})(\omega, t, \omega_{t-}, \mathrm{d}x)\mathrm{d}t \\ &= \mathbb{1}_{\omega_{t-} \neq x}\mathrm{e}^{\tilde{\chi}_{t}(\omega, x)}\mathfrak{q}_{t}(\omega_{t-}, \mathrm{d}x)\mathrm{d}t \;, \end{split}$$

then the previous equation rewrites

$$\mathbb{E}_{\mathbb{Q}_{j}^{k}}\left[f(\mathbf{X}_{t\wedge\tau}) - f(\mathbf{X}_{0})|\mathcal{F}_{0}\right] = \mathbb{E}_{\mathbb{Q}_{j}^{k}}\left[\int_{[0,t\wedge\tau]\times\mathsf{X}} \{f(x) - f(\mathbf{X}_{s-})\}\bar{\mathbf{n}}^{e^{\tilde{\chi}}\mathfrak{q}}(\mathrm{d}s\mathrm{d}x)\middle|\mathcal{F}_{0}\right]$$
$$= \mathbb{E}_{\mathbb{Q}_{j}^{k}}\left[\int_{[0,t\wedge\tau]} (e^{\tilde{\chi}}\mathfrak{q})(s)f(\mathbf{X}_{s-})\mathrm{d}s\middle|\mathcal{F}_{0}\right].$$

Choosing τ such that the above terms are meaningful, we conclude that $\mathbb{Q}_j^k \in MP(e^{\chi_j^k}\mathfrak{q})$ and finish the proof. \Box

Proof of Theorem F.12. This proof is an adaptation of Theorem 2.6 in Léonard (2012) based on technical lemmas provided above applying on the reference measure $\vec{R} \in MP(\vec{q})$ defined at the beginning of Section F. Consider $\|.\|_{\varrho}$ defined as

$$\|\phi\|_{\varrho} := \inf \left\{ a > 0 : \mathbb{E}_{\mathbb{P}} \left[\int_{[0,T_f] \times \mathsf{X}} \varrho(\phi/a) \mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}} \right] \leqslant 1 \right\}$$

for any $\phi \in S_{\varrho}$, where S_{ϱ} is the Orlicz space defined by

$$S_{\varrho} := \left\{ \phi : \mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X} \to \mathbb{R} \text{ measurable s.t. } \mathbb{E}_{\mathbb{P}} \left[\int_{[0, T_f] \times \mathsf{X}} \varrho(b|\phi|) \mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}} \right] < \infty, \text{ for any } b \geqslant 0 \right\}$$

It is well-known that it is a norm called the Luxemburg norm of S_{ϱ} .

Furthermore, for any $\chi \in S_{\varrho}$, by Theorem F.8, the process

$$M_t^{\chi,\overrightarrow{q}} := \int_{[0,t]\times\mathsf{X}} \chi_s(\omega, x) \tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}(\mathrm{d} s \mathrm{d} x)$$

is a local \overrightarrow{R} -martingale. Then, we define the two stochastic processes $(\mathbf{Y}_t^{\chi, \overrightarrow{q}})_{t \in [0, T_f]}$ and $(\mathbf{Z}_t^{\chi, \overrightarrow{q}})_{t \in [0, T_f]}$ as $\mathbf{Y}_t^{\chi, \overrightarrow{q}} := M_t^{\chi, \overrightarrow{q}} - \int_{[0,t]} \beta_s^{\chi, \overrightarrow{q}} \, \mathrm{d}s$ with $\beta_s^{\chi, \overrightarrow{q}} := \int_{\mathsf{X}} \varrho(\chi_s(\omega, x)) \mathbf{n}^{\overrightarrow{q}}(s, \mathrm{d}x)$, and $\mathbf{Z}_t^{\chi, \overrightarrow{q}} = \mathbf{e}^{\mathbf{Y}_t^{\chi, \overrightarrow{q}}}$. By Lemma F.13, the process $\mathbf{Z}_{T_f}^{\chi, \overrightarrow{q}}$ is

a \overrightarrow{R} -supermartingale, thus $0 < \mathbb{E}_{\overrightarrow{R}} \left[\mathbb{Z}_{T_f}^{\chi, \overrightarrow{q}} \right] \leqslant 1$.

For $\mathbb{P} \in \mathcal{P}(\mathbb{D}_{T_f})$ such that $\mathrm{KL}(\mathbb{P}|\vec{R}) < \infty$, the Donsker-Varadhan variational formulation of the KL implies that

$$\operatorname{KL}(\mathbb{P}|\overrightarrow{R}) = \sup\left\{\int u \mathrm{d}\mathbb{P} - \log\int \mathrm{e}^{u} \mathrm{d}\overrightarrow{R} : u \text{ measurable and such that } \int \mathrm{e}^{u} \mathrm{d}\overrightarrow{R} < \infty\right\}.$$
 (70)

For any $\chi \in S_{\varrho}$, choosing $u = Y_{T_f}^{\chi, \vec{q}}$ and noting that $\log \mathbb{E}_{\vec{R}} \left[Z_{T_f}^{\chi, \vec{q}} \right] \leq \log 1 = 0$, we derive

$$\mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}}\chi_t(\omega,x)\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}(\mathrm{d} t\mathrm{d} x)-\int_{[0,T_f]\times\mathsf{X}}\varrho(\chi_t(\omega,x))\bar{\mathrm{n}}^{\overrightarrow{q}}(\mathrm{d} t\mathrm{d} x)\right]\leqslant \mathrm{KL}(\mathbb{P}|\overrightarrow{R})\;.$$

Therefore, for any $\chi \in S_{\varrho}$,

$$\mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}} \chi \mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}\right] \leqslant \mathrm{KL}(\mathbb{P}|\overrightarrow{R}) + \int_{[0,T_f]\times\mathsf{X}} \varrho(\chi) \mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}} .$$
(71)

For any function $\phi \in S_{\varrho}$, taking $\chi := \phi/\|\phi\|_{\varrho}$ in (71) implies that for any $\phi \in S_{\varrho}$,

$$\mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}}\phi\mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}\right] \leqslant [\mathrm{KL}(\mathbb{P}|\overrightarrow{R})+1]\|\phi\|_{\varrho} .$$
(72)

Consider now the sub-space $\mathcal{H} \subset S_{\varrho}$,

$$\mathcal{H} := \left\{ \phi : \mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X} \to \mathbb{R} \text{ predictable and bounded s.t. } \mathbb{E}_{\mathbb{P}} \left[\int_{[0, T_f] \times \mathsf{X}} \varrho(b|\phi|) \mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}} \right] < \infty, \text{ for any } b \ge 0 \right\} .$$

Since any $\phi \in \mathcal{H}$ satisfies (63), Lemma F.13 entails (72) for all $\phi \in \mathcal{H}$, as $\operatorname{KL}(\mathbb{P}|\vec{R}) < \infty$. This implies the linear mapping $\phi \mapsto \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f] \times \mathbf{X}} \phi d\tilde{N}_{\mathbf{X}}^{\vec{q}}\right]$ is continuous on \mathcal{H} equipped with the norm $\|\cdot\|_{\varrho}$. It is worth noting that $(S_{\varrho}, \|\cdot\|_{\varrho})$ forms a Banach space (see Léonard, 2007, Proposition 1.18), and \mathcal{H} is a linear subspace of S_{ϱ} (see Léonard, 2007, Proposition 1.18), hence the Hahn-Banach extension theorem (Delatte (2022, Theorem 1)) ensures that the previously defined linear functional on \mathcal{H} can be extended to the entire space S_{ϱ} without loss of continuity since the Luxemburg norm $\|\cdot\|_{\varrho}$ is a convex function on S_{ϱ} , see Rao & Ren (1991, Section 2.2). Note that the convex conjugate of the Young function $\varrho(|a|)$ is $\varrho^*(|b|)$. Therefore, as showed in Rao & Ren (1991, Theorem 3.1.9), the dual space of $(S_{\varrho}, \|\cdot\|_{\varrho})$ is isomorphic to the space

$$L_{\varrho^*} := \left\{ K : \mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X} \to \mathbb{R} \text{ measurable s.t. } \mathbb{E}_{\mathbb{P}} \left[\int_{[0, T_f] \times \mathsf{X}} \varrho^*(|K|) \mathrm{d}\bar{\mathrm{n}}^{\overrightarrow{q}} \right] < \infty \right\} \ ,$$

that means there exists some $K \in L_{\rho^*}$ such that

$$\mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}}\phi\mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}\right] = \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}}K\phi\mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}}\right], \quad \text{for any } \phi \in \mathcal{H}.$$
(73)

We now prove the uniqueness and predictability of K. Introduce the predictable projection of $K \in L_{\varrho^*}$ as $K^{pr} := \mathbb{E}_{\mathbb{P}}(K|\mathbf{X}_{[0,t)})$, for $t \in [0, T_f]$. Since \mathcal{B} is dense in S_{ϱ} , \mathcal{H} is dense in the subspace of all the predictable processes in S_{ϱ} . Then, any two functions $K_1, K_2 \in L_{\varrho^*}$ satisfying (73) must share the same projection, *i.e.*, $K_1^{pr} = K_2^{pr}$. It follows that there exists a unique predictable process K in the space

$$\mathcal{K}(P) := \left\{ K : \mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X} \to \mathbb{R} \text{ predictable s.t. } \mathbb{E}_{\mathbb{P}} \left[\int_{[0, T_f] \times \mathsf{X}} \varrho^*(|K|) \mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}} \right] < \infty \right\} \ ,$$

which satisfies (73). Moreover, for any function $\phi \in \mathcal{H}$,

$$\begin{split} \int_{[0,T_f]\times\mathsf{X}} \phi \mathrm{d}(\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}} - K\bar{\mathbf{n}}^{\overrightarrow{q}}) &= \int_{[0,T_f]\times\mathsf{X}} \phi \mathrm{d}(N_{\mathbf{X}} - \bar{\mathbf{n}}^{\overrightarrow{q}} - K\bar{\mathbf{n}}^{\overrightarrow{q}}) \\ &= \int_{[0,T_f]\times\mathsf{X}} \phi \mathrm{d}(N_{\mathbf{X}} - (K+1)\bar{\mathbf{n}}^{\overrightarrow{q}}) \\ &= \int_{[0,T_f]\times\mathsf{X}} \phi \mathrm{d}(N_{\mathbf{X}} - U\bar{\mathbf{n}}^{\overrightarrow{q}}) \;, \end{split}$$

with U = K + 1, and the equation (73) is thus equivalent to

$$\mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}}\phi\mathrm{d}(N_{\mathbf{X}}-U\bar{\mathbf{n}}^{\overrightarrow{q}})\right] = 0, \quad \text{for any } \phi \in \mathcal{H} .$$
(74)

Thus, $U\bar{n}^{\vec{q}}$ is a positive measure and U is non-negative. Furthermore, we can argue analogously to obtain equation (74) on the interval [s,t] for $0 \le s \le t \le T_f$ then choose $\phi_s(\omega, x) = f(x) - f(\omega_{s-})$ to deduce

$$\mathbb{E}_{\mathbb{P}}\left[f(\mathbf{X}_{t}) - f(\mathbf{X}_{s})|\mathcal{F}_{s}\right] = \mathbb{E}_{\mathbb{P}}\left[\int_{[s,t]\times\mathsf{X}} (f(x) - f(\mathbf{X}_{z-}))(U\bar{\mathbf{n}}^{\overrightarrow{q}})(\mathrm{d}z\mathrm{d}x)\Big|\mathcal{F}_{s}\right], \quad \text{for any } f \in \mathbb{F}(\mathsf{X}).$$
(75)

Define the random generator $U\overrightarrow{q}$ on $\mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X}^2$ by $(U\overrightarrow{q})(\omega, t, x, y) := U(\omega, t, y)\overrightarrow{q}(x, y)$ for $y \neq x$ and use the convention

$$(U\overrightarrow{q})(\omega,t,x,x) := U(\omega,t,x)\overrightarrow{q}(x,x) = -\sum_{y \neq x} U(\omega,t,y)\overrightarrow{q}(x,y) \; .$$

Then $U\overrightarrow{q}$ forms a generator and (75) rewrites

$$\mathbb{E}_{\mathbb{P}}\left[f(\mathbf{X}_t) - f(\mathbf{X}_s) | \mathcal{F}_s\right] = \mathbb{E}_{\mathbb{P}}\left[\int_{[s,t]} (U \overrightarrow{q})(z) f(\mathbf{X}_{z-}) dz \middle| \mathcal{F}_s\right], \quad \text{for any } f \in \mathbb{F}(\mathsf{X}) \;.$$

As a result, we conclude that $\mathbb{P} \in MP(U\overrightarrow{q})$. We now show the formulation of the Radon-Nikodym density $d\mathbb{P}/d\overrightarrow{R}$. When $\mathbb{P} \sim \overrightarrow{R}$, define the stopping time τ_j^k as

$$\tau_j^k := \inf\left\{t \in [0, T_f]; \int_{[0,t] \times \mathsf{X}} \varrho(\log U_s(\omega, x)\bar{\mathbf{n}}^{\overrightarrow{q}}(\mathrm{d}x\mathrm{d}s)) \ge k \text{ or } \log U_t(\omega, \omega_t) \notin [-j, k]\right\},\$$

which coincides with the stopping time σ_j^k when $\chi = \log U$. Denote $U^{\sigma_j^k} := \mathbb{1}_{[0,\sigma_j^k]}U$ and for simplicity, we write $U = U^{\sigma_j^k}$. By conditioning w.r.t. \mathbf{X}_0 , we can assume without loss of generality that $\overrightarrow{R}_0 = \mathbb{P}_0$, *i.e.*, $\frac{d\mathbb{P}_0}{d\overrightarrow{R}_0}(\mathbf{X}_0) = 1$.

Applying Lemma F.14 for $\mathbb{P} \in MP(U\overrightarrow{q})$ and $\chi = -\log U$, we have

$$\mathbb{Q}^{\tau_j^k} := \mathbb{Z}_{T_f}^{-\log U, U\overrightarrow{q}} \mathbb{P}^{\tau_j^k} \in \mathrm{MP}(\mathbb{1}_{[0, \tau_j^k]} \mathrm{e}^{-\log U} U\overrightarrow{q}) = \mathrm{MP}(\mathbb{1}_{[0, \tau_j^k]} \overrightarrow{q}) \,. \tag{76}$$

Furthermore, $\mathbb{Q}_{0}^{\tau_{j}^{k}} = \mathbb{P}_{0}^{\tau_{j}^{k}} = \overrightarrow{R}_{0}^{\tau_{j}^{k}} = \gamma^{d}$, where γ^{d} is the invariant distribution of $(\mathbf{X}_{t})_{t \in [0, T_{f}]}$, which combined with the equation (76) imply $\mathbb{Q}^{\tau_{j}^{k}}$ is the invariant path measure, *i.e.*,

$$\mathbb{Q}^{ au_j^k} = \overrightarrow{R}^{ au_j^k}$$
 .

Now, applying first Lemma F.14 for $\overrightarrow{R} \in MP(\overrightarrow{q})$ and $\chi = \log U$ yields

$$\tilde{\mathbb{P}}^{\tau_j^k} := \mathbf{Z}_{T_f}^{\log U, \overrightarrow{q}} \overrightarrow{R}^{\tau_j^k} \in \mathrm{MP}(\mathbb{1}_{[0, \tau_j^k]} \mathrm{e}^{\log U} \overrightarrow{q}) = \mathrm{MP}(\mathbb{1}_{[0, \tau_j^k]} U \overrightarrow{q}) \ .$$

Secondly, applying Lemma F.14 with $\tilde{\mathbb{P}}^{\tau_j^k} \in \mathrm{MP}(\mathbb{1}_{[0,\tau_i^k]}U\overrightarrow{q})$ and $\chi = -\log U$ implies

$$\tilde{\mathbb{Q}}^{\tau_j^k} := \mathbf{Z}_{T_f}^{-\log U, U \overrightarrow{q}} \, \tilde{\mathbb{P}}^{\tau_j^k} \in \mathrm{MP}(\mathbb{1}_{[0, \tau_j^k]} \mathrm{e}^{-\log U} U \overrightarrow{q}) = \mathrm{MP}(\mathbb{1}_{[0, \tau_j^k]} \overrightarrow{q}) \,.$$

Argue as before, the previous equation together with the initial condition $\tilde{\mathbb{Q}}_{0}^{\tau_{j}^{k}} = \tilde{\mathbb{P}}_{0}^{\tau_{j}^{k}} = \vec{R}_{0}^{\tau_{j}^{k}} = \gamma^{d}$ yield that $\tilde{\mathbb{Q}}_{j}^{\tau_{j}^{k}} = \vec{R}_{0}^{\tau_{j}^{k}}$. Combining it with $\mathbb{Q}_{0}^{\tau_{j}^{k}} = \vec{R}_{0}^{\tau_{j}^{k}}$ implies

 $\mathbb{Q}^{\tau_j^k} = \tilde{\mathbb{Q}}^{\tau_j^k}$,

which means that

$$\mathbf{Z}_{T_f}^{-\log U, U \overrightarrow{q}} \mathbb{P}^{\tau_j^k} = \mathbf{Z}_{T_f}^{-\log U, U \overrightarrow{q}} \widetilde{\mathbb{P}}^{\tau_j^k} .$$

$$\tag{77}$$

Now observe that $Z_{T_f}^{-\log U, U \overrightarrow{q}} > 0$, therefore, equation (77) implies $\mathbb{P}^{\tau_j^k} = \tilde{\mathbb{P}}^{\tau_j^k}$. Hence $\mathbb{P}^{\tau_j^k} = Z_{T_f}^{\log U, \overrightarrow{q}} \overrightarrow{R}^{\tau_j^k}$, *i.e.*,

$$\begin{split} \mathbb{1}_{[0,\tau_j^k \wedge T_f]} \frac{d\mathbb{P}}{d\overrightarrow{R}}((\mathbf{X}_t)_{t \in [0,T_f]}) &= \mathbb{1}_{[0,\tau_j^k \wedge T_f]} \frac{d\mathbb{P}_0}{d\overrightarrow{R}_0}(\mathbf{X}_0) \\ & \exp\left(\int_{[0,\tau_j^k \wedge T_f] \times \mathsf{X}} (\mathbb{1}_{[0,\tau_j^k \wedge T_f]} \log U) \mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}} - \int_{[0,\tau_j^k \wedge T_f] \times \mathsf{X}} \varrho(\log U) \mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}}\right) \,. \end{split}$$

Letting k and j tend to infinity, since $\tau := \lim_{k,j\to\infty} \tau_j^k = \infty$, we get

$$\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\vec{R}}((\mathbf{X}_t)_{t\in[0,T_f]}) = \frac{\mathrm{d}\mathbb{P}_0}{\mathrm{d}\vec{R}_0}(\mathbf{X}_0)\exp\left(\int_{[0,T_f]\times\mathsf{X}}\log U\mathrm{d}\tilde{N}_{\mathbf{X}}^{\vec{q}} - \int_{[0,T_f]\times\mathsf{X}}\varrho(\log U)\mathrm{d}\bar{n}^{\vec{q}}\right) \,.$$

We now extend the result above to the case when \mathbb{P} might not be equivalent to \overrightarrow{R} . The idea is to approximate \mathbb{P} by a sequence (\mathbb{P}_n) , which satisfies $\mathbb{P}_n \sim \overrightarrow{R}$ for all $n \ge 1$. Denoting

$$\mathbb{P}_n = \left(1 - \frac{1}{n}\right) \mathbb{P} + \frac{\overrightarrow{R}}{n} \quad \text{for } n \ge 1 , \qquad (78)$$

we have $\mathbb{P}_n \sim \overrightarrow{R}$ and $\lim_{n\to\infty} \mathrm{KL}(\mathbb{P}|\mathbb{P}_n) = 0$. For simplicity, we write $\chi = \log U$ and $\chi^n = \log U^n$, which are welldefined \mathbb{P} -a.s. From the variational representation given in (70) and using $\mathbb{P} \in \mathrm{MP}(U\overrightarrow{q})$ combined with Lemma F.13, we obtain

$$\mathrm{KL}(\mathbb{P}|\mathbb{P}_n) \ge \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}} (\chi-\chi^n)\mathrm{d}\tilde{N}_{\mathbf{X}}^{U^n\overrightarrow{q}} - \int_{[0,T_f]\times\mathsf{X}} \varrho(\chi-\chi^n)\mathrm{d}(U^n\overline{\mathbf{n}}^{\overrightarrow{q}})\right]$$

By definition, we have

$$\tilde{N}_{\mathbf{X}}^{U^n \overrightarrow{q}} = N_{\mathbf{X}} - U^n \bar{\mathbf{n}}^{\overrightarrow{q}} = N_{\mathbf{X}} - U \bar{\mathbf{n}}^{\overrightarrow{q}} + (U - U^n) \bar{\mathbf{n}}^{\overrightarrow{q}} = \tilde{N}_{\mathbf{X}}^{U \overrightarrow{q}} + (U - U^n) \bar{\mathbf{n}}^{\overrightarrow{q}} ,$$

which yields

$$\operatorname{KL}(\mathbb{P}|\mathbb{P}_n) \geq \mathbb{E}_{\mathbb{P}} \left[\int_{[0,T_f] \times \mathbf{X}} (\chi - \chi^n) \mathrm{d}(\tilde{N}_{\mathbf{X}}^{U\overrightarrow{q}} + \overline{\mathbf{n}}^{\overrightarrow{q}}(U - U^n)) - \int_{[0,T_f] \times \mathbf{X}} \left(\frac{U}{U^n} - \log \frac{U}{U^n} - 1 \right) U^n \mathrm{d}\overline{\mathbf{n}}^{\overrightarrow{q}} \right]$$
$$= \mathbb{E}_{\mathbb{P}} \left[\int_{[0,T_f] \times \mathbf{X}} (\chi - \chi^n) \mathrm{d}\tilde{N}_{\mathbf{X}}^{U\overrightarrow{q}} + \int_{[0,T_f] \times \mathbf{X}} U \log \frac{U}{U^n} \mathrm{d}\overline{\mathbf{n}}^{\overrightarrow{q}} - \int_{[0,T_f] \times \mathbf{X}} \left(\frac{U}{U^n} - 1 \right) U^n \mathrm{d}\overline{\mathbf{n}}^{\overrightarrow{q}} \right] .$$

Since $\mathbb{P} \in MP(U\overrightarrow{q})$, we deduce that the stochastic integral $\int_{[0,T_f]\times \mathbf{X}} (\chi - \chi^n) d\widetilde{N}_{\mathbf{X}}^{U\overrightarrow{q}}$ is a local \mathbb{P} -martingale. Therefore,

$$\begin{aligned} \operatorname{KL}(\mathbb{P}|\mathbb{P}_n) &\geq \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}} \left(U^n - U - U\log\frac{U^n}{U}\right) \mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}}\right] \\ &= \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}} \left(\frac{U^n}{U} - \log\frac{U^n}{U} - 1\right) U \mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}}\right] \\ &= \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}} \varrho(\chi^n - \chi) \mathrm{d}(U\bar{\mathbf{n}}^{\overrightarrow{q}})\right].\end{aligned}$$

Since $\lim_{n\to\infty} \operatorname{KL}(\mathbb{P}|\mathbb{P}_n) = 0$, we obtain

$$\lim_{n \to \infty} \mathbb{E}_{\mathbb{P}} \left[\int_{[0, T_f] \times \mathsf{X}} \varrho(\chi^n - \chi) \mathrm{d}(U \bar{\mathbf{n}}^{\overrightarrow{q}}) \right] = 0 .$$
⁽⁷⁹⁾

On the other hand, the fact that $\mathbb{P}_n \sim \overrightarrow{R}$ yields

$$\frac{\mathrm{d}\mathbb{P}_{n}}{\mathrm{d}\vec{R}}((\mathbf{X}_{t})_{t\in[0,T_{f}]}) = \frac{\mathrm{d}\mathbb{P}_{n,0}}{\mathrm{d}\vec{R}_{0}}(\mathbf{X}_{0})\exp\left(\int_{[0,T_{f}]\times\mathsf{X}}\chi^{n}\mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}} - \int_{[0,T_{f}]\times\mathsf{X}}\varrho(\chi^{n})\mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}}\right) \,. \tag{80}$$

To obtain the desired expression for the Radon–Nikodym density $\frac{d\mathbb{P}}{d\vec{R}}$, we represent it as

$$\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\vec{R}}((\mathbf{X}_{t})_{t\in[0,T_{f}]}) = \left(\frac{\mathrm{d}\mathbb{P}_{n}}{\mathrm{d}\mathbb{P}_{n}}\cdot\frac{\mathrm{d}\mathbb{P}_{n}}{\mathrm{d}\vec{R}}\right)((\mathbf{X}_{t})_{t\in[0,T_{f}]})$$

$$\stackrel{(80)}{=} \left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{P}_{n}}\cdot\frac{\mathrm{d}\mathbb{P}_{n,0}}{\mathrm{d}\vec{R}_{0}}\right)(\mathbf{X}_{0})\exp\left(\int_{[0,T_{f}]\times\mathbf{X}}\chi^{n}\mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}} - \int_{[0,T_{f}]\times\mathbf{X}}\varrho(\chi^{n})\mathrm{d}\bar{n}^{\overrightarrow{q}}\right)$$

$$= \left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{P}_{n}}\cdot\frac{\mathrm{d}\mathbb{P}_{n,0}}{\mathrm{d}\mathbb{P}_{0}}\right)(\mathbf{X}_{0})\frac{\mathrm{d}\mathbb{P}_{0}}{\mathrm{d}\vec{R}_{0}}(\mathbf{X}_{0})\exp\left(\int_{[0,T_{f}]\times\mathbf{X}}\chi\mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}} - \int_{[0,T_{f}]\times\mathbf{X}}\varrho(\chi)\mathrm{d}\bar{n}^{\overrightarrow{q}}\right)$$

$$\exp\left(\int_{[0,T_{f}]\times\mathbf{X}}(\chi^{n}-\chi)\mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}} - \int_{[0,T_{f}]\times\mathbf{X}}(\varrho(\chi^{n}) - \varrho(\chi))\mathrm{d}\bar{n}^{\overrightarrow{q}}\right),\mathbb{P}\text{-a.s.} \tag{81}$$

The last part can be rewritten as follows using equation (73),

$$\exp\left(\int_{[0,T_f]\times\mathsf{X}} (\chi^n-\chi)\mathrm{d}\tilde{N}^{U\overrightarrow{q}}_{\mathbf{X}} + \int_{[0,T_f]\times\mathsf{X}} (\chi^n-\chi)(\mathrm{e}^{\chi}-1)\mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}} - \int_{[0,T_f]\times\mathsf{X}} (\mathrm{e}^{\chi^n}-\chi^n-\mathrm{e}^{\chi}+\chi)\mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}}\right)$$
$$= \exp\left(\int_{[0,T_f]\times\mathsf{X}} (\chi^n-\chi)\mathrm{d}\tilde{N}^{U\overrightarrow{q}}_{\mathbf{X}} - \int_{[0,T_f]\times\mathsf{X}} \varrho(\chi-\chi^n)\mathrm{d}(U\bar{\mathbf{n}}^{\overrightarrow{q}})\right)$$

We first handle the second integral above using (79) and the fact that ρ is a non-negative function,

$$\mathbb{E}_{\mathbb{P}}\left[\left|\int_{[0,T_f]\times\mathsf{X}}\varrho(\chi-\chi^n)\mathrm{d}(U\bar{\mathrm{n}}^{\overrightarrow{q}})\right|\right]\xrightarrow{n\to\infty}0.$$

This together with Markov's inequality lead to

$$\int_{[0,T_f]\times\mathsf{X}} \varrho(\chi-\chi^n) \mathrm{d}(U\bar{\mathrm{n}}^{\overrightarrow{q}}) \xrightarrow{\mathbb{P}} 0 , \qquad (82)$$

and therefore, from Méliot (2025, Proposition 1.4), there is a subsequence χ^{n_k} such that

$$\int_{[0,T_f]\times\mathsf{X}} \varrho(\chi-\chi^{n_k}) \mathrm{d}(U\bar{\mathrm{n}}^{\overrightarrow{q}}) \xrightarrow{n\to\infty} 0, \quad \mathbb{P}\text{-a.s.}$$
(83)

Furthermore, recall that $\rho(a) = e^a - a - 1 \ge a^2/2$, hence (79) can be used to control the stochastic integral w.r.t. the \mathbb{P} -martingale $\tilde{N}_{\mathbf{X}}^{U\overrightarrow{q}}$ as follows

$$\begin{split} \mathbb{E}_{\mathbb{P}}\left[\left(\int_{[0,T_f]\times\mathsf{X}} (\chi^n-\chi)\mathrm{d}\tilde{N}^{U\overrightarrow{q}}_{\mathbf{X}}\right)^2\right] \stackrel{Corollary\ F.10}{=} \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}} (\chi^n-\chi)^2\mathrm{d}(U\overline{\mathbf{n}}^{\overrightarrow{q}})\right] \\ \leqslant 2\mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathsf{X}} \varrho(\chi-\chi^n)\mathrm{d}(U\overline{\mathbf{n}}^{\overrightarrow{q}})\right] \xrightarrow{n\to\infty} 0 \;. \end{split}$$

This, along with Markov's inequality, results in

$$\int_{[0,T_f]\times\mathsf{X}} (\chi^n - \chi) \mathrm{d}\tilde{N}^{U\overrightarrow{q}}_{\mathbf{X}} \xrightarrow{\mathbb{P}} 0$$

Combining this with (82) implies that

$$\int_{[0,T_f]\times\mathsf{X}} (\chi^n-\chi) \mathrm{d}\tilde{N}^{U\overrightarrow{q}}_{\mathbf{X}} - \int_{[0,T_f]\times\mathsf{X}} \varrho(\chi-\chi^n) \mathrm{d}(U\overline{\mathbf{n}}^{\overrightarrow{q}}) \xrightarrow{\mathbb{P}} 0 \; .$$

As a consequence, Méliot (2025, Proposition 1.4) asserts that there is a subsequence (χ^{n_k}) such that

$$\left[\int_{[0,T_f]\times\mathsf{X}} (\chi^{n_k} - \chi) \mathrm{d}\tilde{N}^{U\overrightarrow{q}}_{\mathbf{X}} - \int_{[0,T_f]\times\mathsf{X}} \varrho(\chi - \chi^{n_k}) \mathrm{d}(U\overline{n}^{\overrightarrow{q}})\right] \xrightarrow{k\to\infty} 0 , \quad \mathbb{P}\text{-a.s.}$$

It helps interpreting (81) as

$$\begin{split} \frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\vec{R}}((\mathbf{X}_t)_{t\in[0,T_f]}) &= \left(\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\mathbb{P}_{n_k}}, \frac{\mathrm{d}\mathbb{P}_{n_k,0}}{\mathrm{d}\mathbb{P}_0}\right)(\mathbf{X}_0) \frac{\mathrm{d}\mathbb{P}_0}{\mathrm{d}\vec{R}_0}(\mathbf{X}_0) \exp\left(\int_{[0,T_f]\times\mathbf{X}} \chi \mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}} - \int_{[0,T_f]\times\mathbf{X}} \varrho(\chi) \mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}}\right) \\ &\qquad \exp\left(\int_{[0,T_f]\times\mathbf{X}} (\chi^{n_k} - \chi) \mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}} - \int_{[0,T_f]\times\mathbf{X}} (\varrho(\chi^{n_k}) - \varrho(\chi)) \mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}}\right) \\ &\qquad \frac{k \to \infty}{(78)} \frac{\mathrm{d}\mathbb{P}_0}{\mathrm{d}\vec{R}_0}(\mathbf{X}_0) \exp\left(\int_{[0,T_f]\times\mathbf{X}} \chi \mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}} - \int_{[0,T_f]\times\mathbf{X}} \varrho(\chi) \mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}}\right), \quad \mathbb{P}\text{-a.s.} \end{split}$$

Replacing $\chi = \log U$, we arrive at our desired claim

$$\frac{\mathrm{d}\mathbb{P}}{\mathrm{d}\vec{R}}((\mathbf{X}_t)_{t\in[0,T_f]}) = \frac{\mathrm{d}\mathbb{P}_0}{\mathrm{d}\vec{R}_0}(\mathbf{X}_0)\exp\left(\int_{[0,T_f]\times\mathsf{X}}\log U\mathrm{d}\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}} - \int_{[0,T_f]\times\mathsf{X}}\varrho(\log U)\mathrm{d}\bar{\mathbf{n}}^{\overrightarrow{q}}\right), \quad \mathbb{P}\text{-a.s.}$$

Consequently, the KL divergence reads as

$$\mathrm{KL}(\mathbb{P}|\vec{R}) = \mathrm{KL}(\mathbb{P}_0|\vec{R}_0) + \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]\times\mathbf{X}} \log U \mathrm{d}\tilde{N}_{\mathbf{X}}^{\vec{q}} - \int_{[0,T_f]\times\mathbf{X}} \rho(\log U) \mathrm{d}\bar{\mathbf{n}}^{\vec{q}}\right] \,.$$

Applying (74) to the function $\phi = \log U$, we get

$$\begin{split} \operatorname{KL}(\mathbb{P}|\overrightarrow{R}) &= \operatorname{KL}(\mathbb{P}_{0}|\overrightarrow{R}_{0}) + \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_{f}]\times\mathsf{X}} (U-1)\log U\mathrm{d}\overline{\mathbf{n}}^{\overrightarrow{q}} - \int_{[0,T_{f}]\times\mathsf{X}} \varrho(\log U)d\overline{\mathbf{n}}^{\overrightarrow{q}}\right] \\ &= \operatorname{KL}(\mathbb{P}_{0}|\overrightarrow{R}_{0}) + \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_{f}]\times\mathsf{X}} [(U-1)\log U - U + \log U + 1]\mathrm{d}\overline{\mathbf{n}}^{\overrightarrow{q}}\right] \\ &= \operatorname{KL}(\mathbb{P}_{0}|\overrightarrow{R}_{0}) + \mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_{f}]\times\mathsf{X}} h(U(t,x))\overline{\mathbf{n}}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x)\right] \,, \end{split}$$

with $h(a) := \varrho^*(a-1) = a \log a - a + 1$ for a > 0. The proof of Theorem F.12 is then finished.

In fact, we can simplify the KL expression above by replacing $\bar{n}^{\vec{q}}(dtdx) = \sum_{y \in X} \mathbb{1}_{\mathbf{X}_{t-}\neq y} \vec{q}(\mathbf{X}_{t-}, y) \delta_y(dx) dt$ and \vec{q} by the formula given in (59) to arrive at

$$\begin{aligned} \operatorname{KL}(\mathbb{P}|\overrightarrow{R}) &= \operatorname{KL}(\mathbb{P}_0|\overrightarrow{R}_0) + \mathbb{E}_{\mathbb{P}} \int_{[0,T_f] \times \mathsf{X}} h(U(t,x)) \sum_{y \in \mathsf{X}} \mathbb{1}_{\mathbf{X}_{t-} \neq y} \overrightarrow{q}(\mathbf{X}_{t-},y) \delta_y(\mathrm{d}x) \mathrm{d}t \\ &= \operatorname{KL}(\mathbb{P}_0|\overrightarrow{R}_0) + \lambda \mathbb{E}_{\mathbb{P}} \int_{[0,T_f]} \sum_{\ell=1}^d h(U(t,\varphi^{(\ell)}(\mathbf{X}_{t-}))) \mathrm{d}t \;. \end{aligned}$$

F.2.2. INTERPRETING THE TIME-REVERSED DYNAMIC AS A CONTROL-DRIVEN PROCESS

Let $\overrightarrow{R} \in MP(\overrightarrow{q})$ be the stationary measure on the path space \mathbb{D}_{T_f} introduced in the previous section. Then, the process \overrightarrow{R} is showed to be reversible, meaning $\overleftarrow{R} = \overrightarrow{R}$, and corresponds to the invariant measure $\gamma^d = \text{Uniform}(X)$. Let $\overrightarrow{\mathbb{P}}^{\mu^*} = \text{Law}((\overrightarrow{X}_t)_{t \in [0, T_f]})$ represent the forward probability measure on the interval $[0, T_f]$ starting from μ^* and governed by the forward generator \overrightarrow{q} given in (59). We denote the corresponding time-reversed probability measure ending at μ^* by $\overleftarrow{\mathbb{P}}^{\mu^*} = \text{Law}((\overleftarrow{X}_t)_{t \in [0, T_f]})$. Let μ_t be the marginal density of the forward dynamic at time $t \in [0, T_f]$ and denote by $\widetilde{\mu}_t := \mu_t/\gamma^d$ the corresponding relative density.

Proposition F.15. The time reversal process $\overleftarrow{\mathbb{P}}^{\mu^*}$ solves the Martingale Problem $MP(u\overrightarrow{q})$ with the control u given by: for $(t, x) \in [0, T_f] \times X$,

$$u_{t}(x,y) = \begin{cases} \tilde{\mu}_{T_{f}-t}(\varphi^{(\ell)}(x))/\tilde{\mu}_{T_{f}-t}(x), & \text{if } y = \varphi^{(\ell)}(x) \text{ for some } \ell = 1, \dots, d, \\ \sum_{\ell=1}^{d} u_{t}(x, \varphi^{(\ell)}(x))/d, & \text{if } y = x, \\ 1, & \text{otherwise}. \end{cases}$$
(84)

Proof of Proposition F.15. Recall that the generator of the backward process \overleftarrow{q} satisfies the following equation for $t \in [0, T_f]$ and $x \neq y$,

$$\overleftarrow{q}_t(x,y) = \frac{\mu_{T_f-t}(y)}{\mu_{T_f-t}(x)} \overrightarrow{q}(y,x) \; .$$

Using the fact that the stationary distribution γ^d satisfies the following balance equation for $x, y \in X, x \neq y$,

$$\gamma^d(x)\overrightarrow{q}(x,y) = \gamma^d(y)\overrightarrow{q}(y,x) \; ,$$

we can express the backward generator as the perturbation of the forward one as follows

$$\overleftarrow{q}_{t}(x,y) = \frac{\mu_{T_{f}-t}(y)}{\mu_{T_{f}-t}(x)}\overrightarrow{q}(y,x) = \frac{\mu_{T_{f}-t}(y)\gamma^{d}(x)}{\mu_{T_{f}-t}(x)\gamma^{d}(y)}\overrightarrow{q}(x,y) = \frac{\widetilde{\mu}_{T_{f}-t}(y)}{\widetilde{\mu}_{T_{f}-t}(x)}\overrightarrow{q}(x,y)$$

Note that $\overrightarrow{q}(x,y) = 0$ for $y \notin \{x, \varphi^{(\ell)}(x); \ell = 1, \dots, d\}$, thus we can define the control u as: for $(t,x) \in [0,T_f] \times X$,

$$u_t(x,\varphi^{(\ell)}(x)) = \frac{\tilde{\mu}_{T_f-t}(\varphi^{(\ell)}(x))}{\tilde{\mu}_{T_f-t}(x)} \quad \text{for } \ell = 1, \dots, d , \quad \text{and} \quad u_t(x,y) = 1 \quad \text{for } y \notin \{x,\varphi^{(\ell)}(x); \ \ell = 1, \dots, d\} ,$$

to obtain the relation

$$\overleftarrow{q}_t(x,y) = u_t(x,y)\overrightarrow{q}(x,y)$$
, for $(t,x,y) \in [0,T_f] \times \mathsf{X}^2$ and $y \neq x$.

Furthermore, under the convention

$$u_t(x,x) = \frac{-\sum_{y \neq x} u_t(x,y) \overrightarrow{q}(x,y)}{\overrightarrow{q}(x,x)} = \frac{1}{d} \sum_{\ell=1}^d u_t(x,\varphi^{(\ell)}(x)) , \quad \text{for } (t,x) \in [0,T_f] \times \mathsf{X} ,$$

 $u \overrightarrow{q}$ in fact forms a generator and satisfies $\overleftarrow{q} = u \overrightarrow{q}$, which implies that $\overleftarrow{\mathbb{P}}^{\mu^*} \in MP(u \overrightarrow{q})$ and we conclude the proof. \Box

Proposition F.15 shows that the unique control u associated with the backward dynamic $(\overleftarrow{X}_t)_{t \in [0, T_f]}$ in Girsanov's thereom F.12 is in fact Markovian. This enables expressing the Radon-Nikodym density $d\overleftarrow{\mathbb{P}}^{\mu^*}/d\overrightarrow{R}$ as

$$\frac{\mathrm{d}\overleftarrow{\mathbb{P}}^{\mu^{\star}}}{\mathrm{d}\overrightarrow{R}}((\overleftarrow{X}_{t})_{t\in[0,T_{f}]}) = \frac{\mathrm{d}\overleftarrow{\mathbb{P}}_{0}^{\mu^{\star}}}{\mathrm{d}\overrightarrow{R}_{0}}(\overleftarrow{X}_{0})\exp\left(\int_{[0,T_{f}]\times\mathsf{X}}\log u_{t}(\overleftarrow{X}_{t-},x)\mathrm{d}\widetilde{N}_{\mathbf{X}}^{\overrightarrow{q}} - \int_{[0,T_{f}]\times\mathsf{X}}\varrho(\log u_{t}(\overleftarrow{X}_{t-},x))\mathrm{d}\overline{n}^{\overrightarrow{q}}\right) ,$$

where u is explicitly given in (84).

F.2.3. EVOLUTION OF THE CONTROL IN THE REVERSED-TIME SYSTEM

This section aims to characterize the control corresponding to the backward dynamics through the Hamilton–Jacobi–Bellman (HJB) equation. This characterization serves as the key ingredient for applying Itô's formula to analyze the evolution of the time-reversed process.

Proposition F.16 (Hamilton–Jacobi–Bellman equation). The control u given in (84) admits the following formula

$$u_t(x,y) = e^{V(t,x) - V(t,y)}$$
, for $(t,x,y) \in [0,T_f] \times \mathsf{X}^2$, $x \neq y$,

where $V(t, x) = -\log \tilde{\mu}_{T_f-t}(x)$, which satisfies the following HJB equation

$$\begin{cases} \partial_t V(t,x) = \lambda \sum_{\ell=1}^d [e^{V(t,x) - V(t,\varphi^{(\ell)}(x))} - 1], \\ V(T_f,x) = g(x) = -\log \frac{d\mu^*}{d\gamma^d}(x), \end{cases} \quad \text{for } (t,x) \in [0,T_f) \times \mathsf{X}.$$
(85)

Proof of Proposition F.16. Denote $V(t, x) := -\log \tilde{\mu}_{T_f-t}(x)$, then the optimal control can be interpreted as

$$u_t(x,y) = \frac{\tilde{\mu}_{T_f-t}(y)}{\tilde{\mu}_{T_f-t}(x)} = \frac{\mathrm{e}^{-V(t,y)}}{\mathrm{e}^{-V(t,x)}} = \mathrm{e}^{V(t,x)-V(t,y)} , \quad \text{for } (t,x,y) \in [0,T_f] \times \mathsf{X}^2 \text{ and } x \neq y .$$

In addition, the function V fulfills the following equation: for $(t, x) \in [0, T) \times X$,

$$\begin{split} \partial_t V(t,x) &= \frac{\partial_t \mu_{T_f - t}(x)}{\mu_{T_f - t}(x)} \\ &= \frac{\sum_{y \in \mathbf{X}} \mu_{T_f - t}(y) \overrightarrow{q}(y,x)}{\mu_{T_f - t}(x)} \quad \text{(by forward Kolmogorov equation)} \\ &= \frac{\sum_{y \in \mathbf{X}} \mu_{T_f - t}(y) \overrightarrow{q}(x,y) \gamma^d(x) / \gamma^d(y)}{\mu_{T_f - t}(x)} \quad \text{(by balance equation)} \\ &= \sum_{y \in \mathbf{X}} \frac{\widetilde{\mu}_{T_f - t}(y)}{\widetilde{\mu}_{T_f - t}(x)} \overrightarrow{q}(x,y) \\ &= \lambda \sum_{\ell = 1}^d [\mathrm{e}^{V(t,x) - V(t,\varphi^{(\ell)}(x))} - 1] \,. \end{split}$$

Moreover, V also satisfies the final condition

$$V_{T_f}(x) = -\log \tilde{\mu}_0(x) = -\log \tilde{\mu}^{\star}(x) = g(x) , \quad \text{for } x \in \mathsf{X} ,$$

therefore V solves the HJB equation (85) and thus concludes the proof of Proposition F.16.

To derive the convergence bound, it is essential to interpret the evolution of the control using the HJB equation first. In fact, the control above satisfies the following martingale and monotone property due to its characterization given in Proposition F.16.

Proposition F.17. With all the notations above, $u_t(\overleftarrow{X}_t, \varphi^{(\ell)}(\overleftarrow{X}_t))$ is a $\overleftarrow{\mathbb{P}}^{\mu^*}$ -martingale for fixed $\ell = 1, \ldots, d$. Consequently, $h(u_t(\overleftarrow{X}_t, \varphi^{(\ell)}(\overleftarrow{X}_t)))$ is a $\overleftarrow{\mathbb{P}}^{\mu^*}$ -submartingale and the monotonicity follows:

$$\mathbb{E}_{\overline{\mathbb{P}}^{\mu\star}}[h(u_s(\overleftarrow{X}_s,\varphi^{(\ell)}(\overleftarrow{X}_s)))] \leqslant \mathbb{E}_{\overline{\mathbb{P}}^{\mu\star}}[h(u_t(\overleftarrow{X}_t,\varphi^{(\ell)}(\overleftarrow{X}_t)))], \quad for \quad 0 \leqslant s \leqslant t \leqslant T_f.$$

Proof of Proposition F.17. Fix $t \in [0, T_f]$ and $\ell = 1, \ldots, d$, applying Itô's formula on

$$f^{\ell}(t, \overline{X}_t) := u_t(\overline{X}_t, \varphi^{(\ell)}(\overline{X}_t)) = e^{V(t, \overline{X}_t) - V(t, \varphi^{(\ell)}(\overline{X}_t))} ,$$

and note that $\text{Law}(\overleftarrow{X}_{\cdot}) = \overleftarrow{\mathbb{P}}^{\mu^{\star}}$ solves $\text{MP}(u\overrightarrow{q})$ as well as $\overleftarrow{X}_{t} = \overleftarrow{X}_{t-}$ for Lebesgue almost all $t \in (0, T_{f}]$ (see Mozumder,

2009, Proposition 2.1), we obtain that the process

$$f^{\ell}(t,\overleftarrow{X}_{t}) - f^{\ell}(0,\overleftarrow{X}_{0}) - \int_{0}^{t} \left[\partial_{s}f^{\ell}(s,\overleftarrow{X}_{s}) + (u\overrightarrow{q})f^{\ell}_{s}(\overleftarrow{X}_{s})\right] \mathrm{d}s$$

is a $\overleftarrow{\mathbb{P}}^{\mu^*}$ -martingale. Denote

$$b_s^\ell := \partial_s f^\ell(s, \overleftarrow{X}_s) + (u \overrightarrow{q}) f_s^\ell(\overleftarrow{X}_s), \quad \text{for } s \in [0, t] \ .$$

We aim to prove that $b_s^{\ell} = 0$. Indeed, by the definition of f^{ℓ} , \vec{q} and the HJB equation (85), we get that

$$\begin{split} b_s^\ell &= u_s(\overleftarrow{X}_s, \varphi^{(\ell)}(\overleftarrow{X}_s)) \left[\partial_s V(s, \overleftarrow{X}_s) - \partial_s V(s, \varphi^{(\ell)}(\overleftarrow{X}_s)) \right] \\ &+ \sum_{i=1}^d \left[u_s(\varphi^{(i)}(\overleftarrow{X}_s), \varphi^{(\ell)}(\varphi^{(i)}(\overleftarrow{X}_s))) - u_s(\overleftarrow{X}_s, \varphi^{(\ell)}(\overleftarrow{X}_s)) \right] \lambda u_s(\overleftarrow{X}_s, \varphi^{(i)}(\overleftarrow{X}_s)) \\ &= \lambda u_s(\overleftarrow{X}_s, \varphi^{(\ell)}(\overleftarrow{X}_s)) \left[\sum_{i=1}^d u_s(\overleftarrow{X}_s, \varphi^{(i)}(\overleftarrow{X}_s)) - \sum_{i=1}^d u_s(\varphi^{(\ell)}(\overleftarrow{X}_s), \varphi^{(i)}(\varphi^{(\ell)}(\overleftarrow{X}_s))) \right] \\ &+ \sum_{i=1}^d \left[u_s(\varphi^{(i)}(\overleftarrow{X}_s), \varphi^{(\ell)}(\varphi^{(i)}(\overleftarrow{X}_s))) - u_s(\overleftarrow{X}_s, \varphi^{(\ell)}(\overleftarrow{X}_s)) \right] \lambda u_s(\overleftarrow{X}_s, \varphi^{(i)}(\overleftarrow{X}_s)) \\ &= \lambda \sum_{i=1}^d \left[u_s(\varphi^{(i)}(\overleftarrow{X}_s), \varphi^{(\ell)}(\varphi^{(i)}(\overleftarrow{X}_s))) u_s(\overleftarrow{X}_s, \varphi^{(i)}(\overleftarrow{X}_s)) \right] \\ &- u_s(\varphi^{(\ell)}(\overleftarrow{X}_s)) u_s(\overleftarrow{X}_s, \varphi^{(\ell)}(\overleftarrow{X}_s)) u_s(\overleftarrow{X}_s, \varphi^{(\ell)}(\overleftarrow{X}_s)) \right] . \end{split}$$

Using the identity $u_s(x, \varphi^{(i)}(x)) = e^{V(s,x) - V(s,\varphi^{(i)}(x))}$ for i = 1, 2, ..., d in Proposition F.22 yields

$$\begin{split} b_s^\ell &= \lambda \sum_{i=1}^d \left[\mathrm{e}^{V(s,\varphi^{(i)}(\overleftarrow{X}_s)) - V(s,\varphi^{(\ell)}(\varphi^{(i)}(\overleftarrow{X}_s))) + V(s,\overleftarrow{X}_s) - V(s,\varphi^{(i)}(\overleftarrow{X}_s))} \\ &\quad - \mathrm{e}^{V(s,\varphi^{(\ell)}(\overleftarrow{X}_s)) - V(s,\varphi^{(i)}(\varphi^{(\ell)}(\overleftarrow{X}_s))) + V(s,\overleftarrow{X}_s) - V(s,\varphi^{(\ell)}(\overleftarrow{X}_s))} \right] \\ &= \lambda \sum_{i=1}^d \left[\mathrm{e}^{-V(s,\varphi^{(\ell)}(\varphi^{(i)}(\overleftarrow{X}_s))) + V(s,\overleftarrow{X}_s)} - \mathrm{e}^{-V(s,\varphi^{(i)}(\varphi^{(\ell)}(\overleftarrow{X}_s))) + V(s,\overleftarrow{X}_s)} \right] = 0 \;, \end{split}$$

as $\varphi^{(\ell)}(\varphi^{(i)}(\overleftarrow{X}_s)) = \varphi^{(i)}(\varphi^{(\ell)}(\overleftarrow{X}_s))$ for any $\ell, i = 1, ..., d$. We thus conclude that $u_t(\overleftarrow{X}_t, \varphi^{(\ell)}(\overleftarrow{X}_t))$ is a $\overleftarrow{\mathbb{P}}^{\mu^*}$ -martingale. Furthermore, since h is convex, it follows that $h(u_t(\overleftarrow{X}_t, \varphi^{(\ell)}(\overleftarrow{X}_t)))$ is a $\overleftarrow{\mathbb{P}}^{\mu^*}$ -submartingale, which implies the desired monotonicity for $\ell = 1, 2, ..., d$ and concludes the proof.

F.3. Connection between the transition matrix and canonical process point of view

As we see in previous sections, the time reversal process can be understood not only via the backward transition matrix but also via the process driven by the control. The transition matrix point of view provides an approximation of the score to simulate the backward process, which is very useful in practice. In parallel, the canonical process point of view gives us a better understanding of the evolution of the time reversal process, which allows us to show a theoretical guarantee on our algorithm. These two points of view in fact have a strong relation, which will be specified in the following Proposition.

Proposition F.18. The control u driving the backward process satisfies the following relation w.r.t. the score function

defined in (17) as

$$u_t(x,\varphi^{(\ell)}(x)) = 1 - s_t^\ell(x), \quad \text{with } \ell = 1, \dots, d \text{ and } (t,x) \in [0,T_f) \times \mathsf{X}.$$
 (86)

Proof of Proposition F.18. The result follows directly from the definition of the score function s in (17) and the form of the control u in (84). This identity confirms the equivalence between the transition matrix viewpoint and the canonical process formulation.

F.4. Optimal control perspective on the time-reversed process

This section presents an alternative perspective on deriving the HJB equation, viewing it through the lens of optimal control and leveraging the Dynamic Programming Principle. In particular, the predictable process u defined in Theorem F.12 can be characterized not only via the backward generator but also as the solution to an optimal control problem involving the expression of relative entropy.

In the continuous case, Conforti et al. (2025) demonstrated that the time reversal process can be formulated as a solution to an optimal control problem. This characterization not only describes the dynamics of the process but also serves as a powerful framework for deriving the HJB equation, which can be obtained by leveraging Girsanov's theorem (Theorem F.12)

Let \overrightarrow{R} be the stationary measure on the path space \mathbb{D}_{T_f} introduced in the previous section. Then, the process \overrightarrow{R} is reversible, meaning $\overleftarrow{R} = \overrightarrow{R}$, and corresponds to the invariant measure $\gamma^d = \text{Uniform}(X)$. Let $\overrightarrow{\mathbb{P}}^{\mu^*}$ represent the forward probability measure on the interval $[0, T_f]$ starting from μ^* and governed by the forward generator \overrightarrow{q} given in (59). We denote the corresponding backward probability measure ending at μ^* by $\overleftarrow{\mathbb{P}}^{\mu^*}$.

Proposition F.19. The time reversal process $\overleftarrow{\mathbb{P}}^{\mu^*}$ satisfies the following optimization problem

$$\overleftarrow{\mathbb{P}}^{\mu^{\star}} = \operatorname*{arg\,min}_{\mathbb{P}\in\mathcal{P}(\mathbb{D}_{T_{f}}):\;\mathrm{KL}(\mathbb{P}|\overrightarrow{R}) < \infty} \left(\mathrm{KL}(\mathbb{P}|\overrightarrow{R}) + \int g \mathrm{d}\mathbb{P}_{T_{f}}\right), \quad \textit{with } g = -\log \frac{\mathrm{d}\mu^{\star}}{\mathrm{d}\gamma^{d}} \;.$$

Proof of Proposition F.19. For $\mathbb{P} = \text{Law}((\mathbf{X}_t)_{t \in [0, T_f]}) \in \mathcal{P}(\mathbb{D}_{T_f})$ such that $\text{KL}(\mathbb{P}|\vec{R}) < \infty$, the Donsker-Varadhan variational formulation of the KL implies that

$$\mathrm{KL}(\mathbb{P}|\overrightarrow{R}) = \sup_{f \in L^1(\mathbb{P}) \text{ s.t. } \int \mathrm{e}^f \mathrm{d}\overrightarrow{R} < \infty} \left(\int f \mathrm{d}\mathbb{P} - \log \int \mathrm{e}^f \mathrm{d}\overrightarrow{R} \right) \ .$$

Taking $f((\mathbf{X}_t)_{t \in [0, T_f]}) = -g(\mathbf{X}_{T_f})$ yields

$$\mathrm{KL}(\mathbb{P}|\overrightarrow{R}) \ge \int -g\mathrm{d}\mathbb{P}_{T_f} - \log \int \mathrm{e}^{-g}\mathrm{d}\overrightarrow{R}_{T_f} = -\int g\mathrm{d}\mathbb{P}_{T_f} - \log \int \mathrm{d}\mu^* = -\int g\mathrm{d}\mathbb{P}_{T_f} \,,$$

since $\int d\mu^* = 1$. On the other hand, since $\overleftarrow{\mathbb{P}}^{\mu^*}$ is the backward process ended at μ^* and \overrightarrow{R} is a reversible path probability measure on $[0, T_f]$, *i.e.*, $\overrightarrow{R} = \overleftarrow{R}$, we have

$$\frac{\mathrm{d}\widetilde{\mathbb{P}}^{\mu^{\star}}}{\mathrm{d}\overrightarrow{R}}((\overleftarrow{X}_{t})_{t\in[0,T_{f}]}) = \frac{\mathrm{d}\mu^{\star}}{\mathrm{d}\gamma^{d}}(\overleftarrow{X}_{T_{f}}) = \mathrm{e}^{-g(\overleftarrow{X}_{T_{f}})}.$$

This implies

$$\mathrm{KL}(\overleftarrow{\mathbb{P}}^{\mu^{\star}}|\overrightarrow{R}) = \mathrm{KL}(\mu^{\star}|\gamma^{d}) = \int \log \frac{\mathrm{d}\mu^{\star}}{\mathrm{d}\gamma^{d}} \mathrm{d}\mu^{\star} = -\int g \mathrm{d}\mu^{\star} = -\int g \mathrm{d}\overleftarrow{\mathbb{P}}_{T_{f}}^{\mu^{\star}}.$$

Combining the previous results, we obtain that the time reversal $\overleftarrow{\mathbb{P}}^{\mu^*}$ is the optimal solution to the following problem

$$\overleftarrow{\mathbb{P}}^{\mu^{\star}} = \operatorname*{arg\,min}_{\mathbb{P}\in\mathcal{P}(\mathbb{D}_{T_{f}}):\;\mathrm{KL}(\mathbb{P}|\overrightarrow{R})<\infty} \left(\mathrm{KL}(\mathbb{P}|\overrightarrow{R}) + \int g \mathrm{d}\mathbb{P}_{T_{f}}\right) \;,$$

which is the desired conclusion.

Utilizing the expression for $KL(\mathbb{P}|\vec{R})$ given by Girsanov's Theorem F.12, we can now frame the corresponding Optimal Control problem.

Theorem F.20. Denote by \mathcal{D} the set of all $U : \mathbb{D}_{T_f} \times [0, T_f] \times \mathsf{X} \to [0, \infty)$ satisfying the integrability condition

$$\mathbb{E}_{\mathbb{P}}\left[\int_{[0,T_f]} \sum_{\ell=1}^d \varrho^*(|U_t(\overleftarrow{X}^U_{[0,t)},\varphi^{(\ell)}(\overleftarrow{X}^U_{t-}))-1|) \mathrm{d}t\right] < \infty ,$$

which is indeed equivalent to condition (61). Then $\overleftarrow{\mathbb{P}}^{\mu^*}$ is the law of \overleftarrow{X}^{u^*} with u^* is the optimal solution to

$$\inf_{U \in \mathcal{D}} \mathbb{E} \left[\lambda \int_{[0,T_f]} \sum_{\ell=1}^d h(U_t(\overleftarrow{X}^U_{[0,t)}, \varphi^{(\ell)}(\overleftarrow{X}^U_{t-}))) dt + g(\overleftarrow{X}^U_{T_f}) \right],$$
s.t. Law($(\overleftarrow{X}^U_t)_{t \in [0,T_f]}$) $\in MP(U\overrightarrow{q})$, with $(U\overrightarrow{q})(\omega, t, x, y) := U_t(\omega_{[0,t)}, y)\overrightarrow{q}(x, y)$ for $x \neq y$.
$$(87)$$

Proof of Theorem F.20. Theorem F.20 is a consequence of Theorem F.12 and Proposition F.19.

F.4.1. HAMILTON-JACOBI-BELLMAN EQUATION

The goal of this section is to derive the Hamilton–Jacobi–Bellman (HJB) equation as in Proposition F.16 using the optimal control viewpoint. To this purpose, we first consider the generalization of the previous control problem. Let J be the following cost

$$\begin{split} J(t,x,U) &:= \mathbb{E}\left[\lambda \int_{[t,T_f]} \sum_{\ell=1}^d h(U_s(\overleftarrow{X}_{[t,s)}^{t,x,U}, \varphi^{(\ell)}(\overleftarrow{X}_{s-}^{t,x,U}))) \mathrm{d}s + g(\overleftarrow{X}_{T_f}^{t,x,U})\right] \,,\\ \text{s.t.} & \left\{ \begin{aligned} &\mathrm{Law}((\overleftarrow{X}_t^{t,x,U})_{t \in [0,T_f]}) \in \mathrm{MP}(U\overrightarrow{q}) \,, \\ & \overleftarrow{X}_{t-}^{t,x,U} = x \,, \end{aligned} \right. \quad \text{for } (x,t,U) \in \mathsf{X} \times [0,T_f] \times \mathcal{D} \,. \end{split}$$

Consider V(t, x) to be the value function of the previous cost function, *i.e.*,

$$V(t,x) := \inf_{U \in \mathcal{D}} J(t,x,U) \; .$$

The following Dynamic Programming Principle is the main tool to derive the HJB equation.

Lemma F.21. (*Touzi, 2012, Theorem 3.3*) For any stopping time $\kappa \in [t, T_f]$, the Dynamic Programming Principle (DPP) implies

$$V(t,x) = \inf_{U \in \mathcal{D}} \mathbb{E} \left[\lambda \int_{[t,\kappa]} \sum_{\ell=1}^{d} h(U_s(\overleftarrow{X}_{[t,s)}^{t,x,U}, \varphi^{(\ell)}(\overleftarrow{X}_{s-}^{t,x,U}))) \mathrm{d}s + V(\kappa, \overleftarrow{X}_{\kappa}^{t,x,U}) \right] \,. \tag{88}$$

Proof of Lemma F.21. Refer to Touzi (2012, Section 3.2).

The expression of V given in Lemma F.21 leads us to the following HJB equation, that in fact coincides with the one derived in Proposition F.16, and is a characterization of the optimal control to the problem (87).

Proposition F.22. Assume that V is continuously differentiable in time. Then, the optimal control u^* to the problem (87) is a Markovian, i.e., $u_t^*(\omega_{[0,t)}, x) = u_t^*(\omega_{t-}, x)$ for $t \in (0, T_f]$, and admits the following formula

$$u_t^*(x, \varphi^{(\ell)}(x)) = e^{V(t,x) - V(t, \varphi^{(\ell)}(x))}$$
 for $\ell = 1, 2, \dots, d$,

with V satisfies the following HJB equation

$$\begin{cases} \partial_t V(t,x) - \lambda \sum_{\ell=1}^d e^{V(t,x) - V(t,\varphi^{(\ell)}(x))} = -\lambda d ,\\ V(T_f,x) = g(x) = -\log \frac{\mathrm{d}\mu^*}{\mathrm{d}\gamma^d}(x) , \end{cases} \quad \text{for } (t,x) \in [0,T_f) \times \mathsf{X} . \end{cases}$$
(89)

Proof of Proposition F.22. The proof is an adaptation of Proposition 3.5 in Touzi (2012). First, the DPP formula (88) for $t \in [0, T_f)$ and $\kappa = t + \alpha$ with $\alpha > 0$ leads to

$$\mathbb{E}\left[\lambda \int_{[t,t+\alpha]} \sum_{\ell=1}^d h(U_s(\overleftarrow{X}_{[t,s)}^{t,x,U}, \varphi^{(\ell)}(\overleftarrow{X}_{s-}^{t,x,U}))) \mathrm{d}s + V(t+\alpha, \overleftarrow{X}_{t+\alpha}^{t,x,U}) - V(t,x)\right] \ge 0 ,$$

for any admissible control $U \in \mathcal{D}$. Using Itô's formula on the process $X^{t,x,U}$ with the law $\stackrel{\leftarrow}{\mathbb{P}} \in MP(U\overrightarrow{q})$, we get

$$\mathbb{E}\left[\int_{[t,t+\alpha]} \lambda \sum_{\ell=1}^{d} h(U_s(\overleftarrow{X}_{[t,s)}^{t,x,U}, \varphi^{(\ell)}(\overleftarrow{X}_{s-}^{t,x,U}))) \mathrm{d}s + \int_{[t,t+\alpha]} (\partial_t V(s, \overleftarrow{X}_s^{t,x,U}) + (U_s \overrightarrow{q}) V_s(\overleftarrow{X}_s^{t,x,U})) \mathrm{d}s\right] \ge 0.$$

Using the formula of \overrightarrow{q} and multiplying the both hand sides by $\frac{1}{\alpha}$ and pushing $\alpha \to 0$, we arrive at

$$\lambda \sum_{\ell=1}^{d} h(U_t(x,\varphi^{(\ell)}(x))) + \partial_t V(t,x) + \lambda \sum_{\ell=1}^{d} \left[V(t,\varphi^{(\ell)}(x)) - V(t,x) \right] U_t(x,\varphi^{(\ell)}(x)) \ge 0 ,$$

for any $U \in \mathcal{D}$. Taking the infimum w.r.t. U, we get

$$\partial_t V(t,x) + \lambda \inf_{U \in \mathcal{D}} \sum_{\ell=1}^d \left[h(U_t(x,\varphi^{(\ell)}(x))) + [V(t,\varphi^{(\ell)}(x)) - V(t,x)]U_t(x,\varphi^{(\ell)}(x)) \right] \ge 0, \quad \text{for } (t,x) \in [0,T_f) \times \mathsf{X}.$$

We prove next the equality by contradiction. Assume that there exists $(t_0, x_0) \in [0, T_f] \times X$ such that

$$\partial_t V(t_0, x_0) + \lambda \inf_{U \in \mathcal{D}} \sum_{\ell=1}^d \left[h(U_{t_0}(x_0, \varphi^{(\ell)}(x_0))) + [V(t_0, \varphi^{(\ell)}(x_0)) - V(t_0, x_0)] U_{t_0}(x_0, \varphi^{(\ell)}(x_0)) \right] > 0 .$$

Denote $\Delta V(t_0, x_0, \varphi^{(\ell)}(x_0)) := V(t_0, \varphi^{(\ell)}(x_0)) - V(t_0, x_0)$. The previous inequality implies that there exists $\varepsilon > 0$ such that

$$\partial_t V(t_0, x_0) + \lambda \inf_{U \in \mathcal{D}} \sum_{\ell=1}^d \left[h(U) + U\Delta V \right](t_0, x_0, \varphi^{(\ell)}(x_0)) \ge \varepsilon > 0 .$$
(90)

Take $\xi > 0$ small enough such that

$$\lambda \sum_{\ell=1}^{d} (\mathrm{e}^{-\Delta V + \xi} - \mathrm{e}^{-\Delta V})(t_0, x_0, \varphi^{(\ell)}(x_0)) < \frac{\varepsilon}{2} , \qquad (91)$$

and define the function $f \leq V$ as

$$f(t,x) := V(t,x) - \xi \left[|t - t_0|^2 + \delta_{\{x_0\}}(x) \right] , \quad \text{for } (t,x) \in [0,T_f] \times \mathsf{X} .$$

It is clear that

$$f(t_0, x_0) = V(t_0, x_0) , \quad \partial_t f(t_0, x_0) = \partial_t V(t_0, x_0) , \quad \text{and} \quad f(t_0, x) - V(t_0, x) = -\xi \text{ for } x \neq x_0 .$$

Therefore,

$$\begin{split} \partial_t f(t_0, x_0) &+ \lambda \inf_{U \in \mathcal{D}} \sum_{\ell=1}^d \left[h(U_{t_0}(x_0, \varphi^{(\ell)}(x_0))) + \left[f(t_0, \varphi^{(\ell)}(x_0)) - f(t_0, x_0) \right] U_{t_0}(x_0, \varphi^{(\ell)}(x_0)) \right] \\ &= \partial_t V(t_0, x_0) + \lambda \inf_{U \in \mathcal{D}} \sum_{\ell=1}^d \left[h(U_{t_0}(x_0, \varphi^{(\ell)}(x_0))) + \left[V(t_0, \varphi^{(\ell)}(x_0)) - V(t_0, x_0) - \xi \right] U_{t_0}(x_0, \varphi^{(\ell)}(x_0)) \right] \\ &= \partial_t V(t_0, x_0) + \lambda \inf_{U \in \mathcal{D}} \sum_{\ell=1}^d \left[h(U) + (\Delta V - \xi) U \right] (t_0, x_0, \varphi^{(\ell)}(x_0)) \,. \end{split}$$

The minimum above is attained at U such that $U(t_0, x_0, \varphi^{(\ell)}(x_0)) = e^{-\Delta V + \xi}(t_0, x_0, \varphi^{(\ell)}(x_0))$, thus

$$\begin{split} \partial_t f(t_0, x_0) &+ \lambda \inf_{U \in \mathcal{D}} \sum_{\ell=1}^d \left[h(U_{t_0}(x_0, \varphi^{(\ell)}(x_0))) + [f(t_0, \varphi^{(\ell)}(x_0)) - f(t_0, x_0)] U_{t_0}(x_0, \varphi^{(\ell)}(x_0)) \right] \\ &= \partial_t V(t_0, x_0) + \lambda \sum_{\ell=1}^d \left[h(\mathrm{e}^{-\Delta V + \xi}) + (\Delta V - \xi) \mathrm{e}^{-\Delta V + \xi} \right] (t_0, x_0, \varphi^{(\ell)}(x_0)) \\ &= \partial_t V(t_0, x_0) + \lambda \sum_{\ell=1}^d (1 - \mathrm{e}^{-\Delta V + \xi}) (t_0, x_0, \varphi^{(\ell)}(x_0)) \\ &= \partial_t V(t_0, x_0) + \lambda \sum_{\ell=1}^d (1 - \mathrm{e}^{-\Delta V}) (t_0, x_0, \varphi^{(\ell)}(x_0)) + \lambda \sum_{\ell=1}^d (\mathrm{e}^{-\Delta V} - \mathrm{e}^{-\Delta V + \xi}) (t_0, x_0, \varphi^{(\ell)}(x_0)) \\ &> \varepsilon - \frac{\varepsilon}{2} = \frac{\varepsilon}{2} > 0 \;, \end{split}$$

where the last inequality relies on (91) and (90) with $U = e^{-\Delta V}$. Therefore, we obtain

$$\partial_t f(t_0, x_0) + \lambda \inf_{U \in \mathcal{D}} \sum_{\ell=1}^d \left[h(U_{t_0}(x_0, \varphi^{(\ell)}(x_0))) + [f(t, \varphi^{(\ell)}(x_0)) - f(t_0, x_0)] U_{t_0}(x_0, \varphi^{(\ell)}(x_0)) \right] > 0 \; .$$

From the continuity in time of the Hamiltonian, the previous inequality yields that

$$\partial_t f(t,x) + \lambda \inf_{U \in \mathcal{D}} \sum_{\ell=1}^d \left[h(U_t(x,\varphi^{(\ell)}(x))) + [f(t,\varphi^{(\ell)}(x)) - f(t,x)] U_t(x,\varphi^{(\ell)}(x)) \right] \ge 0 ,$$

for $(t,x) \in (t_0 - r, t_0 + r) \times \{x_0\} ,$ (92)

for some 0 < r < 1. Defining the stopping time κ^U as

$$\kappa^U := \inf \left\{ t \in (t_0, T_f] : \overleftarrow{X}_{t-}^{t_0, x_0, U} \neq x_0 \right\} \wedge (t_0 + r) ,$$

for an arbitrary control U, we have

$$f(\kappa^{U}, \overleftarrow{X}_{\kappa^{U}}^{t_{0}, x_{0}, u}) = \begin{cases} f(t_{0} + r, \overleftarrow{X}_{t_{0} + r}^{t_{0}, x_{0}, U}) , & \text{if} & \overleftarrow{X}_{\kappa^{U} -}^{t_{0}, x_{0}, U} = x_{0} , \\ f(\kappa^{U}, \overleftarrow{X}_{\kappa^{U}}^{t_{0}, x_{0}, U}) , & \text{if} & \overleftarrow{X}_{\kappa^{U} -}^{t_{0}, x_{0}, U} \neq x_{0} . \end{cases}$$

This implies that

$$f(\kappa^{U}, \overleftarrow{X}_{\kappa^{U}}^{t_{0}, x_{0}, u}) - V(\kappa^{U}, \overleftarrow{X}_{\kappa^{U}}^{t_{0}, x_{0}, u}) = \begin{cases} -\xi r^{2} , & \text{if } \overleftarrow{X}_{\kappa^{U}-}^{t_{0}, x_{0}, u} = x_{0} , \\ -\xi(|\kappa^{U} - t_{0}|^{2} + 1) , & \text{if } \overleftarrow{X}_{\kappa^{U}-}^{t_{0}, x_{0}, u} \neq x_{0} , \end{cases}$$
$$\leqslant -\xi r^{2} .$$

Note that for any $s \in [t_0, \kappa^U]$, not only $\overleftarrow{X}_{s-}^{t_0, x_0, U} = x_0$, thus $\overleftarrow{X}_{[t_0, s)}^{t_0, x_0, U} = x_0$. Therefore,

$$\mathbb{E}\left[\int_{[t_{0},\kappa^{U}]} \lambda \sum_{\ell=1}^{d} h(U_{s}(\overleftarrow{X}_{[t_{0},s)}^{t_{0},x_{0},U},\varphi^{(\ell)}(\overleftarrow{X}_{s-}^{t_{0},x_{0},U}))) ds + V(\kappa^{U},\overleftarrow{X}_{\kappa^{U}}^{t_{0},x_{0},U})\right]$$

$$\geq \mathbb{E}\left[\int_{[t_{0},\kappa^{U}]} \lambda \sum_{\ell=1}^{d} h(U_{s}(x_{0},\varphi^{(\ell)}(x_{0}))) ds + f(\kappa^{U},\overleftarrow{X}_{\kappa^{U}}^{t_{0},x_{0},U}) + \xi r^{2}\right]$$

$$= \mathbb{E}\left[\int_{[t_{0},\kappa^{U}]} \lambda \sum_{\ell=1}^{d} h(U_{s}(x_{0},\varphi^{(\ell)}(x_{0}))) ds + f(\kappa^{U},\overleftarrow{X}_{\kappa^{U}}^{t_{0},x_{0},U}) - f(t_{0},x_{0})\right] + f(t_{0},x_{0}) + \xi r^{2}$$

Using Itô's formula to compute the difference $f(\kappa^U, \overleftarrow{X}_{\kappa^U}^{t_0, x_0, U}) - f(t_0, x_0)$, and relies on the fact that $V(t_0, x_0) = f(t_0, x_0)$, the calculation follows

$$\mathbb{E}\left[\int_{[t_0,\kappa^U]} \lambda \sum_{\ell=1}^d h(U_s(\overleftarrow{X}_{[t_0,s_0]}^{t_0,x_0,U},\varphi^{(\ell)}(\overleftarrow{X}_{s_-}^{t_0,x_0,U}))) ds + V(\kappa^U,\overleftarrow{X}_{\kappa^U}^{t_0,x_0,U})\right] \\
\geqslant \mathbb{E}\int_{[t_0,\kappa^U]} \left[\partial_t f(s,x_0) + \lambda \sum_{\ell=1}^d h(U_s(x_0,\varphi^{(\ell)}(x_0))) \\
+ (f(s,\varphi^{(\ell)}(x_0)) - f(s,x_0))U_s(x_0,\varphi^{(\ell)}(x_0))\right] ds + V(t_0,x_0) + \xi r^2 \\
\stackrel{(92)}{\geqslant} V(t_0,x_0) + \xi r^2 .$$

Since the above control U is arbitrary, the previous inequality is indeed a contradiction to DPP formula (88). Consequently, we can deduce the following HJB equation satisfied by the value function for $(t, x) \in [0, T_f) \times X$,

$$\begin{cases} \partial_t V(t,x) + \lambda \inf_{U \in \mathcal{D}} \sum_{\ell=1}^d \left[h(U_t(x,\varphi^{(\ell)}(x))) + [V(t,\varphi^{(\ell)}(x)) - V(t,x)] U_t(x,\varphi^{(\ell)}(x)) \right] = 0, \\ V(T_f,x) = g(x). \end{cases}$$
(93)

Proceeding as before, we minimize (93) and obtain the optimal control

$$u_t^*(\overleftarrow{X}_{[0,t]}^{t,x,u^*},\varphi^{(\ell)}(\overleftarrow{X}_{t-}^{t,x,u^*})) = u_t^*(\overleftarrow{X}_{t-}^{t,x,u^*},\varphi^{(\ell)}(\overleftarrow{X}_{t-}^{t,x,u^*})) = u_t^*(x,\varphi^{(\ell)}(x)) = e^{V(t,x) - V(t,\varphi^{(\ell)}(x))}, \quad \text{for } \ell = 1, \dots, d,$$

that is Markovian. Replacing the formulation of u^* into (93) boils down to

$$\begin{cases} \partial_t V(t,x) - \lambda \sum_{\ell=1}^d e^{V(t,x) - V(t,\varphi^{(\ell)}(x))} = -\lambda d ,\\ V(T_f,x) = g(x) = -\log \frac{\mathrm{d}\mu^*}{\mathrm{d}\gamma^d}(x) , \end{cases} \quad \text{for } (t,x) \in [0,T_f) \times \mathsf{X} \end{cases}$$

which in fact coincides with the equation in Proposition F.16 and concludes the proof of Theorem F.22.

F.5. Convergence of DMPMs

Based on the perspective of the canonical process, the backward evolution can be described as a control-driven dynamic. These tools enable us to establish the error bounds presented in Theorem 2.3 and Theorem 2.4.

F.5.1. PROOF OF THEOREM 2.3

We first prove the curvature–dimension inequality satisfied by our forward dynamics, which is associated with the stationary distribution $\gamma^d = \text{Uniform}(X)$. This serves as the key estimate for deriving the entropy decay results later.

Lemma F.23 (Curvature–dimension inequality). *The forward dynamic described above satisfies* curvature–dimension inequality $CD(2\lambda, \infty)$, *i.e.*,

$$\Gamma_2(f) \ge 2\lambda\Gamma(f)$$
 for any function f ,

where Γ is the carré du champ operator and Γ_2 is the iterated carré du champ operator.

Proof of Lemma F.23. Recall the formulations of Γ and Γ_2 for functions f and g, which are typically defined as follows:

$$\Gamma(f,g) = \frac{1}{2} \left[\overrightarrow{q} (fg) - f(\overrightarrow{q}g) - (\overrightarrow{q}f)g \right] \quad \text{and} \quad \Gamma(f,f) = \Gamma(f) = \frac{1}{2} \left[\overrightarrow{q} (f^2) - 2f \overrightarrow{q}f \right]$$
$$\Gamma_2(f) = \frac{1}{2} \left[\overrightarrow{q} \Gamma(f) - 2\Gamma(\overrightarrow{q}f,f) \right] ,$$

where \overrightarrow{q} is the forward generator defined in (14). These quantities capture the interaction between the functions f and g under the generator \overrightarrow{q} and play a crucial role in establishing results related to curvature–dimension inequalities and entropy decay. We now compute explicitly $\Gamma(f)(x)$ for $x \in X$ as

$$\begin{split} \Gamma(f)(x) &= \frac{1}{2} \left[\overrightarrow{q}(f^2) - 2f \overrightarrow{q} f \right] \\ &= \frac{\lambda}{2} \left[\sum_{\ell=1}^d \left(f^2(\varphi^{(\ell)}(x)) - f^2(x) \right) - 2f(x) \sum_{\ell=1}^d \left(f(\varphi^{(\ell)}(x)) - f(x) \right) \right] \\ &= \frac{\lambda}{2} \sum_{\ell=1}^d \left[f(\varphi^{(\ell)}(x)) - f(x) \right]^2 \,. \end{split}$$

Regarding the iterated carré du champ $\Gamma_2(f)$, the first term can be calculated as

$$\vec{q} \Gamma(f)(x) = \lambda \sum_{i=1}^{d} \left[\Gamma(f)(\varphi^{(i)}(x)) - \Gamma(f)(x) \right]$$

$$= \lambda \sum_{i=1}^{d} \left[\frac{\lambda}{2} \sum_{\ell=1}^{d} \left(f(\varphi^{(\ell)}(\varphi^{(i)}(x))) - f(\varphi^{(i)}(x)) \right)^2 - \frac{\lambda}{2} \sum_{\ell=1}^{d} \left(f(\varphi^{(\ell)}(x)) - f(x) \right)^2 \right]$$

$$= \frac{\lambda^2}{2} \sum_{i=1}^{d} \sum_{\ell=1}^{d} \left[\left(f(\varphi^{(\ell)}(\varphi^{(i)}(x))) - f(\varphi^{(i)}(x)) \right)^2 - \left(f(\varphi^{(\ell)}(x)) - f(x) \right)^2 \right].$$
(94)

To simplify the second term of Γ_2 , we note that

$$\begin{split} 2\Gamma(f,g)(x) &= \overrightarrow{q}\,(fg) - f(\overrightarrow{q}\,g) - g(\overrightarrow{q}\,f) \\ &= \lambda \sum_{\ell=1}^d \left[f(\varphi^{(\ell)}(x))g(\varphi^{(\ell)}(x)) - f(x)g(\varphi^{(\ell)}(x)) + f(x)g(x) - g(x)f(\varphi^{(\ell)}(x)) \right] \\ &= \lambda \sum_{\ell=1}^d \left[f(x) - f(\varphi^{(\ell)}(x)) \right] \left[g(x) - g(\varphi^{(\ell)}(x)) \right] \,. \end{split}$$

Applying this to $g = \overrightarrow{q} f$ yields

$$2\Gamma(f, \overrightarrow{q}f)(x) = \lambda \sum_{\ell=1}^{d} \left[f(x) - f(\varphi^{(\ell)}(x)) \right] \left[\lambda \sum_{i=1}^{d} \left(f(\varphi^{(i)}(x)) - f(x) \right) - \lambda \sum_{i=1}^{d} \left(f(\varphi^{(\ell)}(\varphi^{(i)}(x))) - f(\varphi^{(\ell)}(x)) \right) \right]$$
$$= \lambda^{2} \sum_{\ell=1}^{d} \sum_{i=1}^{d} \left[f(x) - f(\varphi^{(\ell)}(x)) \right] \left[f(\varphi^{(i)}(x)) - f(x) - f(\varphi^{(\ell)}(\varphi^{(i)}(x))) + f(\varphi^{(\ell)}(x)) \right] . \tag{95}$$

Plugging (94) and (95) into Γ_2 yields

$$\begin{split} \Gamma_2(f)(x) &= \frac{1}{2} \left[\overrightarrow{q} \, \Gamma(f) - 2\Gamma(\overrightarrow{q} \, f, f) \right] \\ &= \frac{\lambda^2}{4} \sum_{\ell=1}^d \sum_{i=1}^d \left[\left(f(\varphi^{(\ell)}(\varphi^{(i)}(x))) - f(\varphi^{(i)}(x)) \right)^2 - \left(f(\varphi^{(\ell)}(x)) - f(x) \right)^2 \right. \\ &\quad + 2 \left(f(x) - f(\varphi^{(\ell)}(x)) \right)^2 - 2 \left(f(x) - f(\varphi^{(\ell)}(x)) \right) \left(f(\varphi^{(i)}(x)) - f(\varphi^{(\ell)}(\varphi^{(i)}(x))) \right) \right] \\ &= \frac{\lambda^2}{4} \sum_{\ell=1}^d \sum_{i=1}^d \left[f(\varphi^{(\ell)}(\varphi^{(i)}(x))) - f(\varphi^{(i)}(x)) - f(\varphi^{(\ell)}(x)) + f(x) \right]^2 \\ &\geq \frac{\lambda^2}{4} \sum_{\ell=1}^d \left[f(x) - f(\varphi^{(\ell)}(x)) - f(\varphi^{(\ell)}(x)) + f(x) \right]^2 \\ &= \lambda^2 \sum_{\ell=1}^d \left[f(\varphi^{(\ell)}(x)) - f(x) \right]^2 = 2\lambda \Gamma(f)(x) \,, \quad \text{for any } x \in \mathsf{X} \,. \end{split}$$

Thus the desired inequality holds and we conclude the proof.

We are now prepared to analyze the key distinguishing result of this paper.

Proof of Theorem 2.3. We begin by establishing a bound on the "distance" between the backward path measure in continuous time, $\overleftarrow{\mathbb{P}}^{\mu^*}$, associated with the controlled process $(\overleftarrow{X}_t)_{t\in[0,T_f]}$, and the path measure $\overleftarrow{\mathbb{P}}^*$ corresponding to the simulated backward process $(\overleftarrow{X}_t^*)_{t\in[0,T_f]}$ generated in Algorithm 3. For brevity, we denote the backward path measure $\overleftarrow{\mathbb{P}}^{\mu^*}$ by $\overleftarrow{\mathbb{P}}$ throughout the remainder of the paper. By taking the stationary forward path measure $\overrightarrow{R} \in \mathrm{MP}(\overrightarrow{q})$ as the reference in Girsanov's Theorem F.12, we derive the Radon–Nikodym density of $\overleftarrow{\mathbb{P}} \in \mathrm{MP}(u\overrightarrow{q})$ with respect to \overrightarrow{R} , where the control u is specified in (84), as follows

$$\frac{\mathrm{d}\widetilde{\mathbb{P}}}{\mathrm{d}\vec{R}}((\mathbf{X}_t)_{t\in[0,T_f]}) = \frac{\mathrm{d}\widetilde{\mathbb{P}}_0}{\mathrm{d}\vec{R}_0}(\mathbf{X}_0)\exp\left(\int_{[0,T_f]\times\mathsf{X}}\log u_t(\mathbf{X}_{t-},x)\tilde{N}_{\mathbf{X}}^{\vec{q}}(\mathrm{d}t\mathrm{d}x) - \int_{[0,T_f]\times\mathsf{X}}\varrho(\log u_t(\mathbf{X}_{t-},x))\bar{n}^{\vec{q}}(\mathrm{d}t\mathrm{d}x)\right) \,.$$

With a partition $0 = t_0 < ... < t_K = T_f$ for $K \ge 1$ of $[0, T_f]$ associated with the sequence of step-size $\{\tau_k\}_{k=1}^K : \tau_{k+1} = t_{k+1} - t_k$, the previous expression rewrites

$$\begin{split} \frac{\mathrm{d}\widetilde{\mathbb{P}}}{\mathrm{d}\vec{R}}((\mathbf{X}_{t})_{t\in[0,T_{f}]}) &= \frac{\mathrm{d}\widetilde{\mathbb{P}}_{0}}{\mathrm{d}\vec{R}_{0}}(\mathbf{X}_{0}) \exp\sum_{k=0}^{K-1} \left(\int_{[t_{k},t_{k+1})\times\mathbf{X}} \log u_{t}(\mathbf{X}_{t-,x})\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) \\ &\quad - \int_{[t_{k},t_{k+1})\times\mathbf{X}} \varrho(\log u_{t}(\mathbf{X}_{t-,x}))\bar{n}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) \right) \\ &= \frac{\mathrm{d}\widetilde{\mathbb{P}}_{0}}{\mathrm{d}\vec{R}_{0}}(\mathbf{X}_{0}) \exp\sum_{k=0}^{K-1} \left(\int_{[t_{k},t_{k+1})\times\mathbf{X}} \log u_{t}(\mathbf{X}_{t-,x})\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) \\ &\quad + \int_{[t_{k},t_{k+1})\times\mathbf{X}} \log u_{t}(\mathbf{X}_{t-,x})\bar{n}^{(u-1)\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) - \int_{[t_{k},t_{k+1})\times\mathbf{X}} \varrho(\log u_{t}(\mathbf{X}_{t-,x}))\bar{n}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) \right) \\ &= \frac{\mathrm{d}\widetilde{\mathbb{P}}_{0}}{\mathrm{d}\vec{R}_{0}}(\mathbf{X}_{0}) \exp\sum_{k=0}^{K-1} \left(\int_{[t_{k},t_{k+1})\times\mathbf{X}} \log u_{t}(\mathbf{X}_{t-,x})\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) \\ &\quad + \lambda \int_{[t_{k},t_{k+1})} \sum_{\ell=1}^{d} (u_{t}-1) \log u_{t}(\mathbf{X}_{t-,\varphi}^{(\ell)}(\mathbf{X}_{t-})) \mathrm{d}t \\ &\quad - \lambda \int_{[t_{k},t_{k+1})} \sum_{\ell=1}^{d} (u_{t}-\log u_{t}-1)(\mathbf{X}_{t-,\varphi}^{(\ell)}(\mathbf{X}_{t-})) \mathrm{d}t \right) \\ &= \frac{\mathrm{d}\widetilde{\mathbb{P}}_{0}}{\mathrm{d}\vec{R}_{0}}(\mathbf{X}_{0}) \exp\sum_{k=0}^{K-1} \left(\int_{[t_{k},t_{k+1})\times\mathbf{X}} \log u_{t}(\mathbf{X}_{t-,x})\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) \\ &\quad + \lambda \int_{[t_{k},t_{k+1})} \sum_{\ell=1}^{d} (u_{t}-\log u_{t}-1)(\mathbf{X}_{t-,\varphi}^{(\ell)}(\mathbf{X}_{t-})) \mathrm{d}t \right) \\ &= \frac{\mathrm{d}\widetilde{\mathbb{P}}_{0}}{\mathrm{d}\vec{R}_{0}}(\mathbf{X}_{0}) \exp\sum_{k=0}^{K-1} \left(\int_{[t_{k},t_{k+1})\times\mathbf{X}} \log u_{t}(\mathbf{X}_{t-,x})\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) \\ &\quad + \lambda \int_{[t_{k},t_{k+1})} \sum_{\ell=1}^{d} (u_{t}\log u_{t}-u_{t}+1)(\mathbf{X}_{t},\varphi^{(\ell)}(\mathbf{X}_{t})) \mathrm{d}t \right), \end{split}$$

as $\mathbf{X}_{t-} = \mathbf{X}_t$ for Lebesgue almost all $t \in [0, T_f]$. Applying Girsanov's theorem F.12 once more to the path measure $\overleftarrow{\mathbb{P}}^{\star} \in \mathrm{MP}(\widehat{u}^{\theta^{\star}}\overrightarrow{q})$ associated with the process $(\overleftarrow{X}_t^{\star})_{t\in[0,T_f]}$ generated by Algorithm 3, we obtain the following expression:

$$\begin{split} \frac{\mathrm{d}\widetilde{\mathbb{P}^{\star}}}{\mathrm{d}\overrightarrow{R}}((\mathbf{X}_{t})_{t\in[0,T_{f}]}) &= \frac{\mathrm{d}\widetilde{\mathbb{P}}_{0}^{\star}}{\mathrm{d}\overrightarrow{R}_{0}}(\mathbf{X}_{0})\exp\sum_{k=0}^{K-1}\left(\int_{[t_{k},t_{k+1})\times\mathbf{X}}\log\hat{u}_{t}^{\theta^{\star}}(\mathbf{X}_{t-},x)\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x)\right) \\ &\quad -\int_{[t_{k},t_{k+1})\times\mathbf{X}}\varrho(\log\hat{u}_{t}^{\theta^{\star}}(\mathbf{X}_{t-},x))\bar{n}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x)\right) \\ &= \frac{\mathrm{d}\widetilde{\mathbb{P}}_{0}^{\star}}{\mathrm{d}\overrightarrow{R}_{0}}(\mathbf{X}_{0})\exp\sum_{k=0}^{K-1}\left(\int_{[t_{k},t_{k+1})\times\mathbf{X}}\log\hat{u}_{t}^{\theta^{\star}}(\mathbf{X}_{t-},x)\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) \\ &\quad +\int_{[t_{k},t_{k+1})\times\mathbf{X}}\log\hat{u}_{t}^{\theta^{\star}}(\mathbf{X}_{t-},x)\bar{n}^{(u-1)\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) - \int_{[t_{k},t_{k+1})\times\mathbf{X}}\varrho(\log\hat{u}_{t}^{\theta^{\star}}(\mathbf{X}_{t-},x))\bar{n}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x)\right) \\ &= \frac{\mathrm{d}\widetilde{\mathbb{P}}_{0}^{\star}}{\mathrm{d}\overrightarrow{R}_{0}}(\mathbf{X}_{0})\exp\sum_{k=0}^{K-1}\left(\int_{[t_{k},t_{k+1})\times\mathbf{X}}\log\hat{u}_{t}^{\theta^{\star}}(\mathbf{X}_{t-},x)\tilde{N}_{\mathbf{X}}^{\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) \\ &\quad +\lambda\int_{[t_{k},t_{k+1})}\sum_{\ell=1}^{d}(u_{t}\log\hat{u}_{t}^{\theta^{\star}}-\hat{u}_{t}^{\theta^{\star}}+1)(\mathbf{X}_{t-},\varphi^{(\ell)}(\mathbf{X}_{t-}))\mathrm{d}t\right). \end{split}$$

Since $\mathbf{X}_t = \mathbf{X}_{t-}$ for Lebesgue almost all $t \in [0, T_f]$ and using the fact that $\hat{u}_t^{\theta^*}(\mathbf{X}_t, \varphi^{(\ell)}(\mathbf{X}_t)) = u_{t_k}^{\theta^*}(\mathbf{X}_{t_k}, \varphi^{(\ell)}(\mathbf{X}_{t_k}))$ for any $t \in [t_k, t_{k+1}), \ell = 1, \ldots, d$ and $(\mathbf{X}_t)_{t \in [0, T_f]}$, the calculation follows

$$\frac{\mathrm{d}\widetilde{\mathbb{P}}^{\star}}{\mathrm{d}\vec{R}}((\mathbf{X}_{t})_{t\in[0,T_{f}]}) = \frac{\mathrm{d}\widetilde{\mathbb{P}}_{0}^{\star}}{\mathrm{d}\vec{R}_{0}}(\mathbf{X}_{0})\exp\sum_{k=0}^{K-1} \left(\int_{[t_{k},t_{k+1})\times\mathbf{X}}\log\hat{u}_{t}^{\theta^{\star}}(\mathbf{X}_{t-},x)\widetilde{N}_{\mathbf{X}}^{u\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) + \lambda \int_{[t_{k},t_{k+1})}\sum_{\ell=1}^{d}(u_{t}(\mathbf{X}_{t},\varphi^{(\ell)}(\mathbf{X}_{t}))\log u_{t_{k}}^{\theta^{\star}}(\mathbf{X}_{t_{k}},\varphi^{(\ell)}(\mathbf{X}_{t_{k}})) - u_{t_{k}}^{\theta^{\star}}(\mathbf{X}_{t_{k}},\varphi^{(\ell)}(\mathbf{X}_{t_{k}})) + 1)\mathrm{d}t\right).$$

Combining two Radon-Nikodym densities $d\overleftarrow{\mathbb{P}}/d\overrightarrow{R}$ and $d\overleftarrow{\mathbb{P}}^*/d\overrightarrow{R}$, we deduce that

$$\begin{aligned} \frac{\mathrm{d}\widehat{\mathbb{P}}}{\mathrm{d}\widehat{\mathbb{P}}^{\star}}((\mathbf{X}_{t})_{t\in[0,T_{f}]}) &= \frac{\mathrm{d}\widehat{\mathbb{P}}_{0}}{\mathrm{d}\widehat{\mathbb{P}}_{0}^{\star}}(\mathbf{X}_{0}) \exp\sum_{k=0}^{K-1} \left(\int_{[t_{k},t_{k+1})\times\mathbf{X}} (\log u_{t}(\mathbf{X}_{t-},x) - \log \hat{u}_{t}^{\theta^{\star}}(\mathbf{X}_{t-},x)) \tilde{N}_{\mathbf{X}}^{u\overrightarrow{q}}(\mathrm{d}t\mathrm{d}x) \right. \\ &+ \lambda \int_{[t_{k},t_{k+1})} \sum_{\ell=1}^{d} \left(u_{t}(\mathbf{X}_{t},\varphi^{(\ell)}(\mathbf{X}_{t})) \log \frac{u_{t}(\mathbf{X}_{t},\varphi^{(\ell)}(\mathbf{X}_{t}))}{u_{t_{k}}^{\theta^{\star}}(\mathbf{X}_{t_{k}},\varphi^{(\ell)}(\mathbf{X}_{t_{k}}))} \right. \\ &- u_{t}(\mathbf{X}_{t},\varphi^{(\ell)}(\mathbf{X}_{t})) + u_{t_{k}}^{\theta^{\star}}(\mathbf{X}_{t_{k}},\varphi^{(\ell)}(\mathbf{X}_{t_{k}})) \right) \mathrm{d}t \right). \end{aligned}$$

This leads to the following expression of the KL divergence

$$\begin{split} \operatorname{KL}(\overleftarrow{\mathbb{P}} | \overleftarrow{\mathbb{P}}^{\star}) &= \operatorname{KL}(\mu_{T_{f}} | \gamma^{d}) + \sum_{k=0}^{K-1} \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[\int_{[t_{k}, t_{k+1}) \times \mathsf{X}} (\log u_{t}(\overleftarrow{X}_{t-}, x) - \log \hat{u}_{t}^{\theta^{\star}}(\overleftarrow{X}_{t-}, x)) \tilde{N}_{\overleftarrow{X}}^{u \overrightarrow{q}}(\mathrm{d}t \mathrm{d}x) \right. \\ &+ \lambda \int_{[t_{k}, t_{k+1})} \sum_{\ell=1}^{d} \left(u_{t}(\overleftarrow{X}_{t}, \varphi^{(\ell)}(\overleftarrow{X}_{t})) \log \frac{u_{t}(\overleftarrow{X}_{t}, \varphi^{(\ell)}(\overleftarrow{X}_{t}))}{u_{t_{k}}^{\theta^{\star}}(\overleftarrow{X}_{t_{k}}, \varphi^{(\ell)}(\overleftarrow{X}_{t_{k}}))} \right. \\ &- u_{t}(\overleftarrow{X}_{t}, \varphi^{(\ell)}(\overleftarrow{X}_{t})) + u_{t_{k}}^{\theta^{\star}}(\overleftarrow{X}_{t_{k}}, \varphi^{(\ell)}(\overleftarrow{X}_{t_{k}})) \right] \mathrm{d}t \right]. \end{split}$$

Note that $\overleftarrow{\mathbb{P}} \in MP(u\overrightarrow{q})$, thus $\widetilde{N}_{\mathbf{X}}^{u\overrightarrow{q}}$ is a $\overleftarrow{\mathbb{P}}$ -martingale, which in turn allow us to reduce the first integral above. Hence we can simplify the expression for the KL divergence as follows

$$\operatorname{KL}(\overleftarrow{\mathbb{P}}|\overleftarrow{\mathbb{P}}^{\star}) = \operatorname{KL}(\mu_{T_{f}}|\gamma^{d}) + \sum_{k=0}^{K-1} \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[\lambda \int_{[t_{k},t_{k+1})} \sum_{\ell=1}^{d} \left(u_{t}(\overleftarrow{X}_{t},\varphi^{(\ell)}(\overleftarrow{X}_{t})) \log \frac{u_{t}(\overleftarrow{X}_{t},\varphi^{(\ell)}(\overleftarrow{X}_{t}))}{u_{t_{k}}^{\theta^{\star}}(\overleftarrow{X}_{t_{k}},\varphi^{(\ell)}(\overleftarrow{X}_{t_{k}}))} - u_{t}(\overleftarrow{X}_{t},\varphi^{(\ell)}(\overleftarrow{X}_{t})) + u_{t_{k}}^{\theta^{\star}}(\overleftarrow{X}_{t_{k}},\varphi^{(\ell)}(\overleftarrow{X}_{t_{k}})) \right) \mathrm{d}t \right].$$

For brevity, we denote $u_t(\overleftarrow{X}_t, \varphi^{(\ell)}(\overleftarrow{X}_t))$ as u_t^ℓ throughout the remainder of the proof. We rewrite $\mathrm{KL}(\overleftarrow{\mathbb{P}} | \overleftarrow{\mathbb{P}}^{\star})$ as

$$\mathrm{KL}(\overleftarrow{\mathbb{P}}|\overleftarrow{\mathbb{P}}^{\star}) = \mathrm{KL}(\mu_{T_f}|\gamma^d) + \lambda \sum_{k=0}^{K-1} \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[\int_{[t_k, t_{k+1})} \sum_{\ell=1}^d \left(u_t^{\ell} \log u_t^{\ell} - u_t^{\ell} \log u_{t_k}^{\theta^{\star}, \ell} - u_t^{\ell} + u_{t_k}^{\theta^{\star}, \ell} \right) \mathrm{d}t \right]$$

We separate this expression into three terms to control:

$$\operatorname{KL}(\overleftarrow{\mathbb{P}} | \overleftarrow{\mathbb{P}}^{\star}) = \underbrace{\operatorname{KL}(\mu_{T_{f}} | \gamma^{d})}_{E_{1}} + \lambda \underbrace{\sum_{k=0}^{K-1} \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[\int_{[t_{k}, t_{k+1})} \sum_{\ell=1}^{d} \left(u_{t_{k}}^{\ell} \log u_{t_{k}}^{\ell} - u_{t}^{\ell} \log u_{t_{k}}^{\theta^{\star}, \ell} - u_{t_{k}}^{\ell} + u_{t_{k}}^{\theta^{\star}, \ell} \right) \mathrm{d}t \right]}_{E_{2}}_{E_{2}} + \lambda \underbrace{\sum_{k=0}^{K-1} \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[\int_{[t_{k}, t_{k+1})} \sum_{\ell=1}^{d} \left(u_{t}^{\ell} \log u_{t}^{\ell} - u_{t}^{\ell} - u_{t_{k}}^{\ell} \log u_{t_{k}}^{\ell} + u_{t_{k}}^{\ell} \right) \mathrm{d}t \right]}_{E_{3}}}_{E_{3}}.$$
(96)

We begin by observing that the uniform distribution γ^d is the invariant measure of the forward process, which satisfies the curvature–dimension condition $CD(2\lambda, \infty)$ (see Lemma F.23). As a consequence, it satisfies a logarithmic Sobolev inequality by Bakry et al. (2014, Theorem 5.10). This, in turn, implies exponential decay of entropy over time by Bakry et al. (2014, Theorem 5.12), and thus we obtain:

$$E_1 = \mathrm{KL}(\mu_{T_f} | \gamma^d) \leqslant \mathrm{e}^{-4\lambda T_f} \mathrm{KL}(\mu^* | \gamma^d) \,. \tag{97}$$

Next, we bound E_2 using the tower property, the martingale property established in Proposition F.17, and the approximation error assumption in Assumption 2.1,

$$\begin{split} E_{2} &= \sum_{k=0}^{K-1} \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[\int_{[t_{k}, t_{k+1})} \sum_{\ell=1}^{d} \left(u_{t_{k}}^{\ell} \log u_{t_{k}}^{\ell} - \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[u_{t}^{\ell} \big| \mathcal{F}_{t_{k}} \right] \log u_{t_{k}}^{\theta^{\star}, \ell} - u_{t_{k}}^{\ell} + u_{t_{k}}^{\theta^{\star}, \ell} \right) \mathrm{d}t \right] \\ \stackrel{Proposition F.17}{=} \sum_{k=0}^{K-1} \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[\int_{[t_{k}, t_{k+1})} \sum_{\ell=1}^{d} \left(u_{t_{k}}^{\ell} \log u_{t_{k}}^{\ell} - u_{t_{k}}^{\ell} \log u_{t_{k}}^{\theta^{\star}, \ell} - u_{t_{k}}^{\ell} + u_{t_{k}}^{\theta^{\star}, \ell} \right) \mathrm{d}t \right] \\ &= \sum_{k=0}^{K-1} (t_{k+1} - t_{k}) \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[\sum_{\ell=1}^{d} \left(u_{t_{k}}^{\ell} \log \frac{u_{t_{k}}^{\ell}}{u_{t_{k}}^{\theta^{\star}, \ell}} - u_{t_{k}}^{\ell} + u_{t_{k}}^{\theta^{\star}, \ell} \right) \right] \\ &= \sum_{k=0}^{K-1} (t_{k+1} - t_{k}) \mathbb{E}_{\overleftarrow{\mathbb{P}}} \left[\sum_{\ell=1}^{d} u_{t_{k}}^{\theta^{\star}, \ell} h \left(\frac{u_{t_{k}}^{\ell}}{u_{t_{k}}^{\theta^{\star}, \ell}} \right) \right] , \end{split}$$

where $h(a) = a \log a - a + 1$ for a > 0 and $(\mathcal{F}_t)_{t \in [0, T_f]}$ denotes the right-continuous and complete natural filtration generated by the process $(\overleftarrow{X}_t)_{t \in [0, T_f]}$. Using Assumption 2.1, we can now bound the term E_2 as

$$E_2 \leqslant \epsilon \sum_{k=0}^{K-1} (t_{k+1} - t_k) = \epsilon (t_K - t_0) = \epsilon T_f , \qquad (98)$$

since $\sum_{k=0}^{K-1} (t_{k+1} - t_k)$ is a telescoping sum. It remains to control E_3 . To do so, we leverage the monotonicity of the function h established in Proposition F.17,

$$\begin{split} E_3 &= \sum_{k=0}^{K-1} \mathbb{E}_{\widehat{\mathbb{P}}} \left[\int_{[t_k, t_{k+1})} \sum_{\ell=1}^d \left(\mathbb{E}_{\widehat{\mathbb{P}}} \left[h(u_t^\ell) \big| \mathcal{F}_{t_k} \right] - h(u_{t_k}^\ell) \right) \mathrm{d}t \right] \\ &\leqslant \sum_{k=0}^{K-1} \mathbb{E}_{\widehat{\mathbb{P}}} \left[\int_{[t_k, t_{k+1})} \sum_{\ell=1}^d \left(\mathbb{E}_{\widehat{\mathbb{P}}} \left[h(u_{t_{k+1}}^\ell) \big| \mathcal{F}_{t_k} \right] - h(u_{t_k}^\ell) \right) \mathrm{d}t \right] \\ &= \sum_{k=0}^{K-1} (t_{k+1} - t_k) \left(\mathbb{E}_{\widehat{\mathbb{P}}} \left[\sum_{\ell=1}^d h(u_{t_{k+1}}^\ell) \right] - \mathbb{E}_{\widehat{\mathbb{P}}} \left[\sum_{\ell=1}^d h(u_{t_k}^\ell) \right] \right) \end{split}$$

Define $\tau = \max\{t_{k+1} - t_k : k = 0, \dots, K-1\}$, the calculation follows

$$E_{3} \leqslant \tau \sum_{k=0}^{K-1} \left(\mathbb{E}_{\widetilde{\mathbb{P}}} \left[\sum_{\ell=1}^{d} h(u_{t_{k+1}}^{\ell}) \right] - \mathbb{E}_{\widetilde{\mathbb{P}}} \left[\sum_{\ell=1}^{d} h(u_{t_{k}}^{\ell}) \right] \right)$$
$$= \tau \left(\mathbb{E}_{\widetilde{\mathbb{P}}} \left[\sum_{\ell=1}^{d} h(u_{T_{f}}^{\ell}) \right] - \mathbb{E}_{\widetilde{\mathbb{P}}} \left[\sum_{\ell=1}^{d} h(u_{0}^{\ell}) \right] \right) \leqslant \tau \beta_{\gamma^{d}}(\mu^{\star}) ,$$
(99)

since h is non-negative. Here, the Fisher-like information functional β_{γ^d} of the data distribution μ^* is defined as $\beta_{\gamma^d}(\mu^*) := \mathbb{E}_{\mathbb{P}}\left[\sum_{\ell=1}^d h(u_{T_f}^\ell)\right]$. Combining (97), (98) and (99), we arrive at

$$\mathrm{KL}(\overleftarrow{\mathbb{P}}|\overleftarrow{\mathbb{P}}^{\star}) \leqslant \mathrm{e}^{-4\lambda T_f} \mathrm{KL}(\mu^{\star}|\gamma^d) + \lambda \tau \beta_{\gamma^d}(\mu^{\star}) + \lambda \epsilon T_f .$$
(100)

Finally, noting that $\mu^* = \text{Law}(\overleftarrow{X}_{T_f}^{u^*})$, we obtain the relation

$$\mathrm{KL}(\mu^{\star}|\mathrm{Law}(\overleftarrow{X}_{T_{f}}^{\star})) = \mathrm{KL}(\mathrm{Law}(\overleftarrow{X}_{T_{f}})|\mathrm{Law}(\overleftarrow{X}_{T_{f}}^{\star})) \leqslant \mathrm{KL}(\mathrm{Law}((\overleftarrow{X}_{t})_{t \in [0, T_{f}]})|\mathrm{Law}((\overleftarrow{X}_{t}^{\star})_{t \in [0, T_{f}]})) = \mathrm{KL}(\overleftarrow{\mathbb{P}}|\overleftarrow{\mathbb{P}}^{\star}),$$

where the inequality is known as *Data processing* inequality for KL divergence (Nutz, 2021, Lemma 1.6). Combining this with (100), we conclude that

$$\mathrm{KL}(\mu^{\star}|\mathrm{Law}(\overleftarrow{X}_{T_{f}}^{\star})) \leqslant \mathrm{e}^{-4\lambda T_{f}} \mathrm{KL}(\mu^{\star}|\gamma^{d}) + \lambda \tau \beta_{\gamma^{d}}(\mu^{\star}) + \lambda \epsilon T_{f} ,$$

and successfully provide a theoretical guarantee for our generative models.

F.5.2. PROOF OF THEOREM 2.4

We demonstrate Theorem 2.4 through the following result:

Lemma F.24. Denoting $y_t = \mathbb{E}_{\mathbb{F}}\left[\sum_{\ell=1}^d h(u_t(\overleftarrow{X}_t, \varphi^{(\ell)}(\overleftarrow{X}_t)))\right]$ then it holds for $t \in [0, T_f)$,

$$y_t \lesssim \frac{d}{T_f - t} \ . \tag{101}$$

Proof of Lemma F.24. Recall the definition of $h(a) = a \log a - a + 1$ for a > 0, and the connection between the optimal control u_t and the score function s_t , as provided by Proposition F.18, is given by

$$u_t(x, \varphi^{(\ell)}(x)) = 1 - s_t^{\ell}(x) \quad \text{for } \ell = 1, \dots, d \text{ and } (t, x) \in [0, T_f) \times \mathsf{X} ,$$

with the score function admitting a conditional expectation expression as given in (18),

$$s_t^{\ell}(x) = \mathbb{E}\left[\frac{2\alpha_{T_f-t}}{1+\alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}(\overrightarrow{X}_{T_f-t}^{\ell} - \overrightarrow{X}_0^{\ell})^2}{1-\alpha_{T_f-t}^2}\middle|\overrightarrow{X}_{T_f-t} = x\right], \quad \text{with } \alpha_t = e^{-2t}.$$

Using the above formulation of the score function, we estimate y_t as follows

$$\begin{split} y_t &= \mathbb{E}_{\widehat{\mathbb{T}}} \left[\sum_{\ell=1}^d (1 - s_t^\ell(\overleftarrow{X}_t)) \log(1 - s_t^\ell(\overleftarrow{X}_t)) + s_t^\ell(\overleftarrow{X}_t) \right] \\ &\leqslant \mathbb{E}_{\widehat{\mathbb{T}}} \left[\sum_{\ell=1}^d (1 - s_t^\ell(\overleftarrow{X}_t))(1 - s_t^\ell(\overleftarrow{X}_t) - 1) + s_t^\ell(\overleftarrow{X}_t) \right] \quad (\text{since } \log a \leqslant a - 1) \\ &= \mathbb{E}_{\widehat{\mathbb{T}}} \left[\sum_{\ell=1}^d (s_t^\ell(\overleftarrow{X}_t))^2 \right] \\ &= \sum_{\ell=1}^d \mathbb{E} \left[\left(\mathbb{E} \left[\frac{2\alpha_{T_f-t}}{1 + \alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}(\overrightarrow{X}_{T_f-t}^\ell - \overrightarrow{X}_0^\ell)^2}{1 - \alpha_{T_f-t}^2} \right| \overrightarrow{X}_{T_f-t} = x \right] \right)^2 \right] \\ &\leqslant \sum_{\ell=1}^d \mathbb{E} \left[\mathbb{E} \left[\left(\frac{2\alpha_{T_f-t}}{1 + \alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}(\overrightarrow{X}_{T_f-t}^\ell - \overrightarrow{X}_0^\ell)^2}{1 - \alpha_{T_f-t}^2} \right)^2 \right| \overrightarrow{X}_{T_f-t} = x \right] \right] \\ &= \sum_{\ell=1}^d \mathbb{E} \left[\left(\frac{2\alpha_{T_f-t}}{1 + \alpha_{T_f-t}} - \frac{4\alpha_{T_f-t}(\overrightarrow{X}_{T_f-t}^\ell - \overrightarrow{X}_0^\ell)^2}{1 - \alpha_{T_f-t}^2} \right)^2 \right] . \end{split}$$

Expanding the last quantity and noting that $(\overrightarrow{X}_{T_f-t}^{\ell} - \overrightarrow{X}_0^{\ell})^2 = (\overrightarrow{X}_{T_f-t}^{\ell} - \overrightarrow{X}_0^{\ell})^4$ since $(\overrightarrow{X}_{T_f-t}^{\ell} - \overrightarrow{X}_0^{\ell}) \in \{0, \pm 1\}$, we obtain that

$$\begin{split} y_t &\leqslant \sum_{\ell=1}^d \mathbb{E} \left[\frac{4\alpha_{T_f-t}^2}{(1+\alpha_{T_f-t})^2} + \frac{16\alpha_{T_f-t}^2(\vec{X}_{T_f-t}^\ell - \vec{X}_0^\ell)^2}{(1+\alpha_{T_f-t})(1-\alpha_{T_f-t}^2)} \left(-1 + \frac{1}{1-\alpha_{T_f-t}} \right) \right] \\ &= \sum_{\ell=1}^d \mathbb{E} \left[\frac{4\alpha_{T_f-t}^2}{(1+\alpha_{T_f-t})^2} + \frac{16\alpha_{T_f-t}^3(\vec{X}_{T_f-t}^\ell - \vec{X}_0^\ell)^2}{(1-\alpha_{T_f-t}^2)^2} \right] \\ &\leqslant \sum_{\ell=1}^d \left(\frac{4\alpha_{T_f-t}^2}{(1+\alpha_{T_f-t})^2} + \frac{16\alpha_{T_f-t}^3 \mathbb{E} \left[(\vec{X}_{T_f-t}^\ell - \vec{X}_0^\ell)^2 \right]}{(1-\alpha_{T_f-t}^2)^2} \right) \,. \end{split}$$

Note that

$$\mathbb{E}\left[(\overrightarrow{X}_{T_f-t}^{\ell} - \overrightarrow{X}_0^{\ell})^2\right] = \mathbb{P}\left(\overrightarrow{X}_{T_f-t}^{\ell} \neq \overrightarrow{X}_0^{\ell}\right) = \frac{1}{2}(1 - \alpha_{T_f-t}).$$

Thus the upper bound of y_t is

$$\begin{split} y_t &\leqslant \sum_{\ell=1}^d \left(\frac{4\alpha_{T_f-t}^2}{(1+\alpha_{T_f-t})^2} + \frac{8\alpha_{T_f-t}^3(1-\alpha_{T_f-t})}{(1-\alpha_{T_f-t}^2)^2} \right) \\ &= \frac{4d\alpha_{T_f-t}^2}{(1+\alpha_{T_f-t})^2} \left[1 + \frac{2\alpha_{T_f-t}}{(1-\alpha_{T_f-t})} \right] \\ &= \frac{4d\alpha_{T_f-t}^2}{1-\alpha_{T_f-t}^2} \lesssim \frac{d}{e^{4(T_f-t)}-1} \lesssim \frac{d}{T_f-t} \,, \quad \text{for } t \in [0,T_f) \;, \end{split}$$

where the last estimate follows from the elementary inequality $e^a \ge a + 1$ for all $a \in \mathbb{R}$. Therefore, the bound in (101) holds for all $t \in [0, T_f)$.

Proof of Theorem 2.4. We follow the same strategy as in the proof of Theorem 2.3, the only difference being the way we handle the following term in (99)

$$E_3 := \sum_{k=0}^{K-1} (t_{k+1} - t_k) \left(\mathbb{E}_{\widehat{\mathbb{P}}} \left[\sum_{\ell=1}^d h(u_{t_{k+1}}) \right] - \mathbb{E}_{\widehat{\mathbb{P}}} \left[\sum_{\ell=1}^d h(u_{t_k}) \right] \right) = \sum_{k=0}^{K-1} \tau_{k+1} \left(y_{t_{k+1}} - y_{t_k} \right)$$

Following precisely the argument structure in the proof of Theorem 3 from Conforti et al. (2025), and fixing T_f , a, and c, we choose the sequence of step-size as

$$\tau_{k+1} = \begin{cases} T_f - t_{N-1} , & k = N-1 ,\\ ca , & k_0 + k_1 \le k \le k_0 + k_1 + k_2 - 1 ,\\ c(T_f - t_k) , & k_0 \le k \le k_0 + k_1 - 1 ,\\ c , & 0 \le k \le k_0 - 1 , \end{cases}$$
(102)

and set the number of iterations $K = k_0 + k_1 + k_2 + 1$, with

$$k_0 = \max\{k \ge 0 : T_f - t_k \ge 1\}, k_1 = \max\{k \ge 0 : T_f - t_{k_0 + k} \ge a\} \text{ and } k_2 = \max\{k \ge 0 : T_f - t_{k_0 + k_1 + k} \ge 0\}.$$
(103)

It is shown in Conforti et al. (2025) that

$$k_0 = \lfloor c^{-1}(T_f - 1) \rfloor, \quad k_1 = \lfloor \log(a/(T_f - t_{k_0})) / \log(1 - c) \rfloor \lesssim \log(1/a) / c ,$$

$$K - k_0 - k_1 = k_2 + 1 \lesssim 1/c \quad \text{and} \quad \tau_{k+1} = c(1 - c)^{k-k_0}(T_f - t_{k_0}) \quad \text{for } k \in \{k_0, \dots, k_0 + k_1 - 1\} .$$
(104)

Using (102) and the monotonicity of y_t established in Proposition F.17, we can bound E_3 as

$$\begin{split} E_{3} &\leqslant \sum_{k=0}^{K-1} \tau_{k+1} \left(y_{t_{k+1}} - y_{t_{k}} \right) \\ &= \tau_{K} y_{t_{K}} + \sum_{k=1}^{K-1} y_{t_{k}} (\tau_{k} - \tau_{k+1}) \\ &= \tau_{K} y_{t_{K}} + \sum_{k=1}^{k_{0}} y_{t_{k}} (\tau_{k} - \tau_{k+1}) + \sum_{k=k_{0}+1}^{k_{0}+k_{1}} y_{t_{k}} (\tau_{k} - \tau_{k+1}) \\ &+ \sum_{k=k_{0}+k_{1}+1}^{k_{0}+k_{1}+k_{2}-1} y_{t_{k}} (\tau_{k} - \tau_{k+1}) + y_{t_{k_{0}+k_{1}+k_{2}}} (\tau_{K-1} - \tau_{K}) \\ &\lesssim \underbrace{y_{t_{k_{0}}} [c - c(T_{f} - t_{k_{0}})]}_{(1)} + \underbrace{c \sum_{k=k_{0}+1}^{k_{0}+k_{1}-1} y_{t_{k}} \tau_{k}}_{(2)} \\ &+ \underbrace{cy_{t_{k_{0}+k_{1}}} (T_{f} - t_{k_{0}+k_{1}-1} - a)}_{(3)} + \underbrace{y_{t_{K}} \tau_{K-1}}_{(4)} \,. \end{split}$$

We now bound (1)-(2)-(3)-(4) one-by-one. We start with

$$(1): y_{t_{k_0}}[c - c(T_f - t_{k_0})] \leqslant c y_{t_{k_0}} \overset{Lemma \ F.24}{\lesssim} \frac{cd}{T_f - t_{k_0}} \overset{(103)}{\leqslant} cd \ .$$

Next, we bound the second term

$$(2): c \sum_{k=k_0+1}^{k_0+k_1-1} y_{t_k} \tau_k \overset{Lemma}{\lesssim} \overset{F.24}{c} c \sum_{k=k_0+1}^{k_0+k_1-1} \frac{d\tau_k}{T_f - t_k} = c^2 d \sum_{k=k_0+1}^{k_0+k_1-1} \frac{\tau_k}{\tau_{k+1}}$$
$$\stackrel{(104)}{=} c^2 d \sum_{k=k_0+1}^{k_0+k_1-1} \frac{c(1-c)^{k-k_0-1}(T_f - t_{k_0})}{c(1-c)^{k-k_0}(T_f - t_{k_0})}$$
$$\leqslant c^2 d \sum_{k=k_0+1}^{k_0+k_1-1} \frac{1}{1-c} \overset{c\leqslant 1/2}{\lesssim} c^2 dk_1 \lesssim c d \log(1/a) .$$

Recall that $T_f - t_{k_0+k_1-1} \leq 1$. As a result, we have

$$(3): cy_{t_{k_0+k_1}}(T_f - t_{k_0+k_1-1} - a) \overset{Lemma \ F.24}{\lesssim} \frac{cd}{T_f - t_{k_0+k_1}}(T_f - t_{k_0+k_1-1}) = \frac{cd}{T_f - t_{k_0+k_1-1} - \tau_{k_0+k_1}}(T_f - t_{k_0+k_1-1}) \\ \overset{(102)}{\leqslant} \frac{cd}{1 - c} \overset{c\leqslant 1/2}{\lesssim} cd .$$

Finally, for the last term, we have by definition of $L = y_{T_f}/d$,

$$(4): y_{t_K}h_{K-1} = y_{T_f}ca = cadL$$
.

Plugging all the bounds of (1) - (2) - (3) - (4) into E_3 gives

$$E_3 \lesssim cd[1 + \log(1/a) + aL] .$$

Moreover, choosing a = 1/L yields the following bound on the sampling error

$$\mathrm{KL}(\mu^*|\mathrm{Law}(\overline{X}_{T_f}^*)) \lesssim \mathrm{e}^{-4\lambda T_f} \mathrm{KL}(\mu^*|\gamma^d) + \lambda cd[1 + \log(L)] + \lambda \epsilon T_f ,$$

and this concludes the proof of Theorem 2.4.

F.5.3. PROOF OF COROLLARY 2.5

Proof of Corollary 2.5. Applying Theorem 2.4 with c, T_f chosen in (32) implies

$$\mathrm{KL}(\mu^{\star}|\mathrm{Law}(\overleftarrow{X}_{T_{f}}^{\star})) \lesssim \epsilon + \epsilon \log \frac{\mathrm{KL}(\mu^{\star}|\gamma^{d})}{\epsilon} .$$

Therefore we obtain the approximation error $\tilde{O}(\epsilon \log(\mathrm{KL}(\mu^*|\gamma^d)))$. In addition, the number of iterations is given by

$$\begin{split} K &= k_0 + k_1 + K - k_0 - k_1 \overset{(104)}{\lesssim} \frac{T_f - 1}{c} + \frac{\log(1/a)}{c} + \frac{1}{c} = \frac{T_f + \log(L)}{c} \\ \overset{(32)}{=} \frac{\lambda d[1 + \log(L)][\log(\mathrm{KL}(\mu^*|\gamma^d)/\epsilon) / \lambda + \log(L)]}{\epsilon} \\ &= \frac{d[1 + \log(L)][\log(\mathrm{KL}(\mu^*|\gamma^d)/\epsilon) + \lambda \log(L)]}{\epsilon} \,, \end{split}$$

which grows logarithmically rather than linearly with respect to the discrete Fisher information $\beta_{\gamma^d}(\mu^*)$. We thus conclude that this step-size sequence offers improved performance compared to the constant step-size.

F.6. Convergence of DMPMs with early stopping strategy

F.6.1. PROOF OF THEOREM 2.6

Proof of Theorem 2.6. Applying the proof of Theorem 2.4 with the target distribution μ_{η} in place of μ^* leads to the following analogous result:

$$\operatorname{KL}(\mu_{\eta}|\operatorname{Law}(\overleftarrow{X}_{T_{f}-\eta}^{\star})) \leq e^{-4\lambda(T_{f}-\eta)}\operatorname{KL}(\mu_{\eta}|\gamma^{d}) + \lambda cd[1+\log(L)] + \lambda\epsilon(T_{f}-\eta) .$$
(105)

Recall that the forward dynamic satisfies the curvature–dimension condition $CD(2\lambda, \infty)$ (see Lemma F.23) and consequently, the logarithm Sobolev inequality holds (Bakry et al., 2014, Theorem 5.10):

$$\mathrm{KL}(\mu_{\eta}|\gamma^{d}) \leqslant \frac{1}{4\lambda} \mathcal{E}\left(\frac{\mu_{\eta}}{\gamma^{d}}, \log\frac{\mu_{\eta}}{\gamma^{d}}\right) .$$
(106)

As computed in Lemma F.23, we have

$$\mathcal{E}\left(\frac{\mu_{\eta}}{\gamma^{d}},\log\frac{\mu_{\eta}}{\gamma^{d}}\right) = \lambda \sum_{x \in \mathbf{X}} \sum_{\ell=1}^{d} \frac{\mu_{\eta}}{\gamma^{d}}(x) \left(\log\frac{\mu_{\eta}}{\gamma^{d}}(x) - \log\frac{\mu_{\eta}}{\gamma^{d}}(\varphi^{(\ell)}(x))\right) \gamma^{d}(x)$$
$$= \lambda \mathbb{E}\left[\sum_{\ell=1}^{d} \left(\log\frac{\mu_{\eta}(\overrightarrow{X}_{\eta})}{\mu_{\eta}(\varphi^{(\ell)}(\overrightarrow{X}_{\eta}))}\right)\right].$$
(107)

On the other hand, the discrete Fisher information of μ_{η} is given by

$$\begin{split} \beta_{\gamma^d}(\mu_\eta) &= \mathbb{E}\left[\sum_{\ell=1}^d h\left(\mathrm{e}^{-\log\left(\frac{\mu\eta}{\gamma^d}(\vec{X}_\eta)\right) + \log\left(\frac{\mu\eta}{\gamma^d}(\varphi^{(\ell)}(\vec{X}_\eta))\right)}\right)\right] \\ &= \mathbb{E}\left[\sum_{\ell=1}^d h\left(\mathrm{e}^{\log\left(\frac{\mu\eta(\varphi^{(\ell)}(\vec{X}_\eta))}{\mu\eta(\vec{X}_\eta)}\right)}\right)\right] \\ &= \mathbb{E}\left[\sum_{\ell=1}^d h\left(\frac{\mu\eta(\varphi^{(\ell)}(\vec{X}_\eta))}{\mu\eta(\vec{X}_\eta)}\log\frac{\mu\eta(\varphi^{(\ell)}(\vec{X}_\eta))}{\mu\eta(\vec{X}_\eta)} - \frac{\mu\eta(\varphi^{(\ell)}(\vec{X}_\eta))}{\mu\eta(\vec{X}_\eta)} + 1\right)\right] \\ &= \sum_{x\in\mathsf{X}}\sum_{\ell=1}^d \left(\mu\eta(\varphi^{(\ell)}(x))\log\frac{\mu\eta(\varphi^{(\ell)}(x))}{\mu\eta(x)} - \mu\eta(\varphi^{(\ell)}(x)) + \mu\eta(x)\right) \\ &= \sum_{x\in\mathsf{X}}\sum_{\ell=1}^d \left(\mu\eta(\varphi^{(\ell)}(x))\log\frac{\mu\eta(\varphi^{(\ell)}(x))}{\mu\eta(x)}\right) \\ &= \sum_{x\in\mathsf{X}}\sum_{\ell=1}^d \left(\mu\eta(x)\log\frac{\mu\eta(x)}{\mu\eta(\varphi^{(\ell)}(x))}\right) \\ &= \mathbb{E}\left[\sum_{\ell=1}^d \left(\log\frac{\mu\eta(\vec{X}_\eta)}{\mu\eta(\varphi^{(\ell)}(\vec{X}_\eta))}\right)\right] \stackrel{(107)}{=} \frac{1}{\lambda}\mathcal{E}\left(\frac{\mu\eta}{\gamma^d},\log\frac{\mu\eta}{\gamma^d}\right) \,. \end{split}$$

Plugging it into (106) implies that the discrete Fisher information dominates the KL divergence as

$$\operatorname{KL}(\mu_{\eta}|\gamma^{d}) \leqslant \frac{1}{4} \beta_{\gamma^{d}}(\mu_{\eta}) \;.$$

To complete the proof, it remains to bound the discrete Fisher information. To this end, we employ the elementary inequality $\log a \le a - 1$ for a > 0, and proceed using the same reasoning as in Lemma F.24, as detailed below,

$$\begin{split} \beta_{\gamma^d}(\mu_\eta) &= \mathbb{E}\left[\sum_{\ell=1}^d \left(\frac{\mu_\eta(\varphi^{(\ell)}(\overrightarrow{X}_\eta))}{\mu_\eta(\overrightarrow{X}_\eta)}\log\frac{\mu_\eta(\varphi^{(\ell)}(\overrightarrow{X}_\eta))}{\mu_\eta(\overrightarrow{X}_\eta)} - \frac{\mu_\eta(\varphi^{(\ell)}(\overrightarrow{X}_\eta))}{\mu_\eta(\overrightarrow{X}_\eta)} + 1\right)\right] \\ &\leqslant \mathbb{E}\left[\sum_{\ell=1}^d \left(\frac{\mu_\eta(\varphi^{(\ell)}(\overrightarrow{X}_\eta))}{\mu_\eta(\overrightarrow{X}_\eta)}\left(\frac{\mu_\eta(\varphi^{(\ell)}(\overrightarrow{X}_\eta))}{\mu_\eta(\overrightarrow{X}_\eta)} - 1\right) - \frac{\mu_\eta(\varphi^{(\ell)}(\overrightarrow{X}_\eta))}{\mu_\eta(\overrightarrow{X}_\eta)} + 1\right)\right] \\ &= \mathbb{E}\left[\sum_{\ell=1}^d \left(\frac{\mu_\eta(\varphi^{(\ell)}(\overrightarrow{X}_\eta))}{\mu_\eta(\overrightarrow{X}_\eta)} - 1\right)^2\right] \\ &= \mathbb{E}\left[\sum_{\ell=1}^d \left(s_{T_f-\eta}^\ell(\overrightarrow{X}_\eta)\right)^2\right] \lesssim \frac{d}{T_f - (T_f - \eta)} = \frac{d}{\eta} \,. \end{split}$$

It follows that

$$\operatorname{KL}(\mu_\eta|\gamma^d) \lesssim rac{d}{\eta} \quad ext{and} \quad L = d^{-1} \beta_{\gamma^d}(\mu_\eta) \lesssim \eta^{-1} \; .$$

Combined with (105), this leads to the desired conclusion:

$$\operatorname{KL}(\mu_{\eta}|\operatorname{Law}(\overleftarrow{X}_{T_{f}-\eta})) \lesssim d\eta^{-1} \mathrm{e}^{-4\lambda(T_{f}-\eta)} + \lambda cd[1+\log(\eta^{-1})] + \lambda \epsilon(T_{f}-\eta) .$$

F.6.2. PROOF OF PROPOSITION 2.7

Proof of Proposition 2.7. Recall that the total variation distance between μ_{η} and μ^* for $\eta \in (0, \max\left\{T_f, \frac{1}{\lambda}\right\})$ is defined as

$$\|\mu_{\eta} - \mu^{\star}\|_{\mathrm{TV}} = \sum_{x \in \mathsf{X}} |\mu_{\eta}(x) - \mu^{\star}(x)|.$$

By the triangle inequality, we obtain

$$\|\mu_{\eta} - \mu^{\star}\|_{\mathrm{TV}} \leqslant \sum_{x \in \mathsf{X}} |\mu_{\eta}(x) - \mu^{\star}(x)\overrightarrow{p}_{\eta}(x,x)| + |\mu^{\star}(x) - \mu^{\star}(x)\overrightarrow{p}_{\eta}(x,x)| ,$$

where the transition probability $\overrightarrow{p}_{\eta}$ is defined in (15). The two terms above are non-negative as

$$\mu_{\eta}(x) = \sum_{z \in \mathsf{X}} \mu^{\star}(z) \overrightarrow{p}_{\eta}(z, x) \ge \mu^{\star}(x) \overrightarrow{p}_{\eta}(x, x) ,$$

and the transition probability $\overrightarrow{p}_{\eta}(x,x) \leq 1$ for any $x \in X$. This together with the formula of $\overrightarrow{p}_{\eta}$ in (15) yield

$$\|\mu_{\eta} - \mu^{\star}\|_{\mathrm{TV}} \leqslant \sum_{x \in \mathsf{X}} \left[\mu_{\eta}(x) + \mu^{\star}(x) - 2\mu^{\star}(x) \left(\frac{1}{2} + \frac{1}{2} \mathrm{e}^{-2\lambda\eta}\right)^{d} \right]$$
$$\leqslant 2 - 2 \left(\frac{1}{2} + \frac{1}{2} \mathrm{e}^{-2\lambda\eta}\right)^{d}.$$

To simplify this upper bound, we apply the exponential inequality $e^a \ge a + 1$ to $e^{-2\lambda\eta}$ and note that $1 - \lambda\eta > 0$:

$$\|\mu_{\eta} - \mu^{\star}\|_{\mathrm{TV}} \leq 2 - 2\left(\frac{1}{2} + \frac{1}{2}(-2\lambda\eta + 1)\right)^{d} = 2 - 2(1 - \lambda\eta)^{d},$$

and thus the proof is complete.

F.6.3. PROOF OF COROLLARY 2.8

Proof of Corollary 2.8. We observe the following by applying the triangle inequality and Pinsker's inequality,

$$\begin{aligned} \|\mu^{\star} - \operatorname{Law}(\overleftarrow{X}_{T_{f}-\eta}^{\star})\|_{\mathrm{TV}} &\leq \|\mu^{\star} - \mu_{\eta}\|_{\mathrm{TV}} + \|\mu_{\eta} - \operatorname{Law}(\overleftarrow{X}_{T_{f}-\eta}^{\star})\|_{\mathrm{TV}} \\ &\leq \|\mu^{\star} - \mu_{\eta}\|_{\mathrm{TV}} + \sqrt{2\operatorname{KL}(\mu_{\eta}|\operatorname{Law}(\overleftarrow{X}_{T_{f}-\eta}^{\star})))} \,. \end{aligned}$$

Then Theorem 2.6 and Proposition 2.7 together imply

$$\|\mu^{\star} - \operatorname{Law}(\overleftarrow{X}_{T_{f}-\eta}^{\star})\|_{\mathrm{TV}} \lesssim 1 - (1 - \lambda\eta)^{d} + \sqrt{d\eta^{-1} \mathrm{e}^{-4\lambda(T_{f}-\eta)}} + \sqrt{\lambda cd[1 + \log(\eta^{-1})]} + \sqrt{\lambda \epsilon(T_{f}-\eta)} \,. \tag{108}$$

The choices of η , c and T_f in (35) lead to

$$1 - (1 - \lambda \eta)^d = \epsilon \quad \text{and} \quad d\eta^{-1} e^{-4\lambda(T_f - \eta)} = \epsilon^2 \quad \text{and} \quad \lambda cd[1 + \log(\eta^{-1})] = \epsilon^2 .$$
(109)

Substituting (109) into (108) gives the desired upper bound

$$\|\mu^{\star} - \operatorname{Law}(\overleftarrow{X}_{T_f-\eta}^{\star})\|_{\mathrm{TV}} \lesssim \epsilon + \sqrt{\lambda \epsilon (T_f-\eta)} ,$$

with the number of iterations is

$$K \lesssim \frac{T_f - \eta + \log(\eta^{-1})}{c} \stackrel{(35)}{=} \frac{\lambda d[1 + \log(\eta^{-1})][\log(d/\eta\epsilon^2) / \lambda + \log(\eta^{-1})]}{\epsilon^2} \\ = \frac{d[1 + \log(\eta^{-1})][\log(d/\eta\epsilon^2) + \lambda\log(\eta^{-1})]}{\epsilon^2} \\ = \frac{d[1 + \log(\eta^{-1})][\log(d/\epsilon^2) + (\lambda + 1)\log(\eta^{-1})]}{\epsilon^2} .$$
(110)

Finally, the term $\log \eta^{-1}$ can be bounded from above by Bernoulli's inequality $(1-\epsilon)^{1/d} \leq 1 - \frac{\epsilon}{d}$ for $\epsilon \in (0,1)$ and $d \ge 1$,

$$\log(\eta^{-1}) = \log\left(\frac{\lambda}{1 - (1 - \epsilon)^{1/d}}\right) \leqslant \log\frac{\lambda d}{\epsilon} \; .$$

Plugging it into (110) concludes the proof of Corollary 2.8.