A Extended background: KL divergence with the Wiener process plan

This section illustrates that (7) holds. Consider a process $T \in \mathcal{F}(\mathbb{P}_0)$, i.e., T is a probability distribution on Ω with the marginal $\mathbb{P}_0$ at $t = 0$.

Let W^ϵ be the Wiener process with variance ϵ starting at $\mathbb{P}_0$, i.e., it satisfies $dX_t = \sqrt{\epsilon}dW_t$ with $X_0 \sim \mathbb{P}_0$. Hence, $\pi^{W^\epsilon}(y|x)$ is the normal distribution $\frac{d\pi^{W^\epsilon}(y|x)}{dy} = \mathcal{N}(y|x, \epsilon I)$. Then $\mathrm{KL}(\pi^T || \pi^{W^\epsilon})$ between joint distributions at times $t = 0$ and $t = 1$ of these processes is given by:

$$\mathrm{KL}(\pi^T || \pi^{W^\epsilon}) = -\int_{\mathcal{X} \times \mathcal{Y}} \log \frac{d\pi^{W^\epsilon}(x,y)}{d[x,y]} d\pi^T(x,y) + \underbrace{\int_{\mathcal{X} \times \mathcal{Y}} \log \frac{d\pi^T(x,y)}{d[x,y]} d\pi^T(x,y)}_{=-H(\pi^T)}, \quad (18)$$

where $\frac{d\pi(x,y)}{d[x,y]}$ denotes the joint density of distribution π. We derive

$$-\int_{\mathcal{X} \times \mathcal{Y}} \log \frac{d\pi^{W^\epsilon}(x,y)}{d[x,y]} d\pi^T(x,y) = -\int_{\mathcal{X} \times \mathcal{Y}} \log \frac{d\pi^{W^\epsilon}(y|x)}{dy} \frac{d\pi^{W^\epsilon}(x)}{dx} d\pi^T(x,y) =$$

$$-\int_{\mathcal{X} \times \mathcal{Y}} \log \frac{d\pi^{W^\epsilon}(y|x)}{dy} d\pi^T(x,y) - \int_{\mathcal{X}} \int_{\mathcal{Y}} \log \frac{d\pi^{W^\epsilon}(x)}{dx} d\pi^T(y|x) \overbrace{d\mathbb{P}_0(x)}^{d\pi_0^T(x)} =$$

$$-\int_{\mathcal{X} \times \mathcal{Y}} \log \frac{d\pi^{W^\epsilon}(y|x)}{dy} d\pi^T(x,y) - \int_{\mathcal{X}} \log \frac{d\pi^{W^\epsilon}(x)}{dx} \Big[\int_{\mathcal{Y}} 1 d\pi^T(y|x) \Big] d\mathbb{P}_0(x) =$$

$$-\int_{\mathcal{X}} \int_{\mathcal{Y}} \log \frac{d\pi^{W^\epsilon}(y|x)}{dy} d\pi^T(x,y) - \int_{\mathcal{X}} \log \frac{d\pi^{W^\epsilon}(x)}{dx} d\mathbb{P}_0(x) =$$

$$-\int_{\mathcal{X}} \int_{\mathcal{Y}} \log \frac{d\pi^{W^\epsilon}(y|x)}{dy} d\pi^T(x,y) - \int_{\mathcal{X}} \log \frac{d\mathbb{P}_0(x)}{dx} d\mathbb{P}_0(x) =$$

$$-\int_{\mathcal{X} \times \mathcal{Y}} \log \frac{d\pi^{W^\epsilon}(y|x)}{dy} d\pi^T(x,y) + H(\mathbb{P}_0) =$$

$$-\int_{\mathcal{X} \times \mathcal{Y}} \log \Big((2\pi\epsilon)^{\frac{-D}{2}} \exp\Big(-\frac{||x-y||^2}{2\epsilon} \Big) \Big) d\pi^T(x,y) + H(\mathbb{P}_0) =$$

$$+\frac{D}{2} \log(2\pi\epsilon) + \int_{\mathcal{X} \times \mathcal{Y}} \frac{||x-y||^2}{2\epsilon} d\pi^T(x,y) + H(\mathbb{P}_0).$$

After substituting this result into (18), one obtains:

$$\mathrm{KL}(\pi^T || \pi^{W^\epsilon}) = \int_{\mathcal{X} \times \mathcal{Y}} \frac{||x-y||^2}{2\epsilon} d\pi^T(x,y) - H(\pi^T) + \underbrace{\frac{D}{2} \log(2\pi\epsilon) + H(\mathbb{P}_0)}_{=C \text{ in } (7)}. \quad (19)$$

B Proofs

In this section, we provide the proof for our main theoretical results (Theorems 4.1 and 4.3). The proofs require several auxiliary results which we formulate and prove in §B.1 and §B.2.

In §B.1, we show that entropic OT can be reformulated as a maximin problem. This is a technical intermediate result needed to derive our main maximin reformulation of SB (Theorem 4.1). More precisely, in §B.2, we show that these maximin problems for entropic OT and SB are actually equivalent. By using this observation and related facts, in §B.3, we prove our Theorems 4.1 and 4.3.

B.1 Relaxation of entropic OT

To begin with, we recall some facts regarding EOT and SB. Recall the definition of EOT (2):

$$\inf_{\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)} \int_{\mathcal{X} \times \mathcal{Y}} \frac{||x-y||^2}{2} d\pi(x,y) - \epsilon H(\pi). \quad (20)$$

Henceforth, we assume that $\mathbb{P}_0$ and $\mathbb{P}_1$ are *absolutely continuous*. The situation when $\mathbb{P}_0$ or $\mathbb{P}_1$ is not absolutely continuous is not of any practical interest: there is no $\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)$ for which the differential entropy $H(\pi)$ is finite which means that (20) equals to $+\infty$ for every $\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)$, i.e., every plan is optimal. In turn, when $\mathbb{P}_0$ and $\mathbb{P}_1$ are absolutely continuous, the OT plan is unique thanks to the strict convexity of entropy (on the set of absolutely continuous plans).

Recall equation (19) for $\mathrm{KL}(\pi||\pi^{W^\epsilon})$:

$$\mathrm{KL}(\pi||\pi^{W^\epsilon}) = \int_{\mathcal{X} \times \mathcal{Y}} \frac{||x - y||^2}{2\epsilon} d\pi(x, y) - H(\pi) + C. \tag{21}$$

We again note that

$$\inf_{\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)} \mathrm{KL}(\pi||\pi^{W^\epsilon}) = \frac{1}{\epsilon} \inf_{\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)} \left\{ \int_{\mathcal{X} \times \mathcal{Y}} \frac{||x - y||^2}{2} d\pi(x, y) - \epsilon H(\pi) \right\} + C,$$

i.e., problems (20) and (21) can be viewed as equivalent as their minimizers are the same. For convenience, we proceed with $\inf_{\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)} \mathrm{KL}(\pi||\pi^{W^\epsilon})$ and denote its optimal value by $\mathcal{L}^*$, i.e.,

$$\mathcal{L}^* \stackrel{\text{def}}{=} \inf_{\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)} \mathrm{KL}(\pi||\pi^{W^\epsilon}).$$

For a given $\beta \in \mathcal{C}_{b,2}(\mathcal{Y})$, we define an auxiliary joint distribution $d\pi^\beta(x, y) = d\pi^\beta(y|x)d\mathbb{P}_0(x)$, where $d\pi^\beta(y|x)$ is given by

$$d\pi^\beta(y|x) = \frac{1}{C_\beta^x} \exp(\beta(y))d\pi^{W^\epsilon}(y|x),$$

where $C_\beta^x(x) \stackrel{\text{def}}{=} \int_{\mathcal{Y}} \exp(\beta(y))d\pi^{W^\epsilon}(y|x)$. Note that $C_\beta^x < \infty$ since $\beta \in \mathcal{C}_{b,2}(\mathcal{Y})$ is upper bounded.

Before going further, we need to introduce several technical auxiliary results.

Proposition B.1. *For $\nu \in \mathcal{P}_2(\mathcal{Y})$ and $x \in \mathcal{X}$ it holds that*

$$\mathrm{KL}(\nu||\pi^{W^\epsilon}(\cdot|x)) - \int_{\mathcal{Y}} \beta(y)d\nu(y) = \mathrm{KL}(\nu||\pi^\beta(\cdot|x)) - \log C_\beta^x. \tag{22}$$

Proof of Proposition B.1. We derive

$$\mathrm{KL}(\nu||\pi^{W^\epsilon}(\cdot|x)) - \int_{\mathcal{Y}} \beta(y)d\nu(y) = \int_{\mathcal{Y}} \log \frac{d\nu(y)}{d\pi^{W^\epsilon}(y|x)} d\nu(y) - \int_{\mathcal{Y}} \beta(y)d\nu(y) =$$

$$\int_{\mathcal{Y}} \log \frac{d\nu(y)}{\exp(\beta(y))d\pi^{W^\epsilon}(y|x)} d\nu(y) = \int_{\mathcal{Y}} \log \frac{C_\beta^x d\nu(y)}{C_\beta^x \exp(\beta(y))d\pi^{W^\epsilon}(y|x)} d\nu(y) =$$

$$\int_{\mathcal{Y}} \log \frac{d\nu(y)}{d\pi^\beta(y|x)} d\nu(y) - \log C_\beta^x = \mathrm{KL}(\nu||\pi^\beta(\cdot|x)) - \log C_\beta^x.$$

$\square$

Lemma B.2. *For $\pi \in \Pi(\mathbb{P}_0)$, i.e., probability distributions $\pi \in \mathcal{P}_2(\mathcal{X} \times \mathcal{Y})$ whose projection to $\mathcal{X}$ equals $\mathbb{P}_0$, we have*

$$\mathrm{KL}(\pi||\pi^{W^\epsilon}) - \int_{\mathcal{Y}} \beta(y)d\pi(y) = \mathrm{KL}(\pi||\pi^\beta) - \int_{\mathcal{X}} \log C_\beta^x d\mathbb{P}_0(x). \tag{23}$$

Proof of Lemma B.2. For each $x \in \mathcal{X}$, we substitute $\nu = \pi(\cdot|x)$ to (22) and integrate over $x \sim \mathbb{P}_0$. For the left part, we obtain the following:

$$\int_{\mathcal{X}} \left(\mathrm{KL}(\pi(\cdot|x)||\pi^{W^\epsilon}(\cdot|x)) - \int_{\mathcal{Y}} \beta(y)d\pi(y|x) \right) d\mathbb{P}_0(x) =$$

$$\int_{\mathcal{X}} \mathrm{KL}(\pi(\cdot|x)||\pi^{W^\epsilon}(\cdot|x))d\mathbb{P}_0(x) - \int_{\mathcal{X} \times \mathcal{Y}} \beta(y)d\pi(y|x)d\mathbb{P}_0(x) =$$

$$\int_{\mathcal{X}} \int_{\mathcal{Y}} \log \frac{d\pi(y|x)}{d\pi^{W^\epsilon}(y|x)} d\pi(y|x) d\mathbb{P}_0(x) - \int_{\mathcal{Y}} \beta(y) d\pi_1(y) =$$

$$\int_{\mathcal{X}} \int_{\mathcal{Y}} \log \frac{d\pi(y|x) d\mathbb{P}_0(x)}{d\pi^{W^\epsilon}(y|x) d\mathbb{P}_0(x)} d\pi(y|x) d\mathbb{P}_0(x) - \int_{\mathcal{Y}} \beta(y) d\pi_1(y) =$$

$$\int_{\mathcal{X} \times \mathcal{Y}} \log \frac{d\pi(x,y)}{d\pi^{W^\epsilon}(x,y)} d\pi(x,y) - \int_{\mathcal{Y}} \beta(y) d\pi_1(y) =$$

$$\mathrm{KL}(\pi||\pi^{W^\epsilon}) - \int_{\mathcal{Y}} \beta(y) d\pi_1(y).$$

For the right part, we obtain:

$$\int_{\mathcal{X}} \left\{ \mathrm{KL}(\pi(\cdot|x)||\pi^\beta(\cdot|x)) - \log C_\beta^x \right\} d\mathbb{P}_0(x) =$$

$$\int_{\mathcal{X}} \mathrm{KL}(\pi(\cdot|x)||\pi^\beta(\cdot|x)) d\mathbb{P}_0(x) - \int_{\mathcal{X}} \log C_\beta^x d\mathbb{P}_0(x) =$$

$$\int_{\mathcal{X}} \int_{\mathcal{Y}} \log \frac{d\pi(y|x)}{d\pi^\beta(y|x)} d\pi(y|x) d\mathbb{P}_0(x) - \int_{\mathcal{X}} \log C_\beta^x d\mathbb{P}_0(x) =$$

$$\int_{\mathcal{X}} \int_{\mathcal{Y}} \log \frac{d\pi(y|x) d\mathbb{P}_0(x)}{d\pi^\beta(y|x) d\mathbb{P}_0(x)} d\pi(y|x) d\mathbb{P}_0(x) - \int_{\mathcal{X}} \log C_\beta^x d\mathbb{P}_0(x) =$$

$$\int_{\mathcal{X} \times \mathcal{Y}} \log \frac{d\pi(x,y)}{d\pi^\beta(x,y)} d\pi(x,y) - \int_{\mathcal{X}} \log C_\beta^x d\mathbb{P}_0(x) =$$

$$\mathrm{KL}(\pi||\pi^\beta) - \int_{\mathcal{X}} \log C_\beta^x d\mathbb{P}_0(x).$$

Hence, the equality (23) holds. $\qquad\square$

Now we introduce the following auxiliary functional $\widetilde{\mathcal{L}}$:

$$\widetilde{\mathcal{L}}(\beta, \pi) \stackrel{\text{def}}{=} \mathrm{KL}(\pi||\pi^{W^\epsilon}) - \int_{\mathcal{Y}} \beta(y) d\pi_1(y) + \int_{\mathcal{Y}} \beta(y) d\mathbb{P}_1(y).$$

Recall that π_1 denotes the second marginal distribution of π. We use this functional to derive the saddle point reformulation of EOT.

Lemma B.3 (Relaxation of entropic optimal transport). *It holds that*

$$\mathcal{L}^* = \inf_{\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)} \mathrm{KL}(\pi||\pi^{W^\epsilon}) = \sup_\beta \inf_{\pi \in \Pi(\mathbb{P}_0)} \widetilde{\mathcal{L}}(\beta, \pi), \qquad (24)$$

where sup *is taken over potentials* $\beta \in \mathcal{C}_{b,2}(\mathcal{Y})$ *and* inf *over* $\pi \in \Pi(\mathbb{P}_0)$.

Proof of Lemma B.3. We obtain

$$\inf_{\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)} \mathrm{KL}(\pi||\pi^{W^\epsilon}) = \inf_{\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)} \left\{ \int_{\mathcal{X}} \mathrm{KL}(\pi(y|x)||\pi^{W^\epsilon}(y|x)) d\mathbb{P}_0(x) \right\} =$$

$$\inf_{\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)} \int_{\mathcal{X}} C(x, \pi(y|x)) d\mathbb{P}_0(x), \qquad (25)$$

where $C(x, \nu) \stackrel{\text{def}}{=} \mathrm{KL}(\nu||\pi^{W^\epsilon}(y|x))$. The last problem in (25) is known as weak OT [7, 20] with a weak OT cost C. For a given $\beta \in \mathcal{C}_{b,2}(\mathcal{Y})$, consider its weak C-transform given by:

$$\beta^C(x) \stackrel{\text{def}}{=} \inf_{\nu \in \mathcal{P}_2(\mathcal{Y})} \left\{ C(x, \nu) - \int_{\mathcal{Y}} \beta(y) d\nu(y) \right\}. \qquad (26)$$

Since $C : \mathcal{X} \times \mathcal{P}_2(\mathcal{Y}) \to \mathbb{R}$ is lower bounded (by zero), convex in the second argument and jointly lower semi-continuous, the following equality holds [7, Theorem 1.3]:

$$\mathcal{L}^* = \inf_{\pi \in \Pi(\mathbb{P}_0, \mathbb{P}_1)} \int_{\mathcal{X}} C(x, \pi(y|x)) d\mathbb{P}_0(x) = \sup_\beta \left\{ \int_{\mathcal{X}} \beta^C(x) d\mathbb{P}_0(x) + \int_{\mathcal{Y}} \beta(y) d\mathbb{P}_1(y) \right\}, \qquad (27)$$

where $\sup$ is taken over $\beta \in \mathcal{C}_{b,2}(\mathcal{Y})$. We use our Proposition B.1 to note that

$$\beta^C(x) = \inf_{\nu \in \mathcal{P}_2(\mathcal{Y})} \left\{ \mathrm{KL}\big(\nu || \pi^{W^\epsilon}(y|x)\big) - \int_{\mathcal{Y}} \beta(y) d\nu(y) \right\} =$$
$$\inf_{\nu \in \mathcal{P}_2(\mathcal{Y})} \left\{ \mathrm{KL}(\nu || \pi^\beta(\cdot|x)) - \log C_\beta^x \right\} = - \log C_\beta^x.$$

This allows us to derive

$$\int_{\mathcal{X}} \beta^C(x) d\mathbb{P}_0(x) + \int_{\mathcal{Y}} \beta(y) d\mathbb{P}_1(y) = \tag{28}$$

$$-\int_{\mathcal{X}} \log C_\beta^x d\mathbb{P}_0(x) + \int_{\mathcal{Y}} \beta(y) d\mathbb{P}_1(y) =$$

$$\overbrace{\left\{ \inf_{\pi \in \Pi(\mathbb{P}_0)} \mathrm{KL}(\pi || \pi^\beta) \right\}}^{=0} - \int_{\mathcal{X}} \log C_\beta^x d\mathbb{P}_0(x) + \int_{\mathcal{Y}} \beta(y) d\mathbb{P}_1(y) =$$

$$\inf_{\pi \in \Pi(\mathbb{P}_0)} \left\{ \mathrm{KL}(\pi || \pi^\beta) - \overbrace{\int_{\mathcal{X}} \log C_\beta^x d\mathbb{P}_0(x) + \int_{\mathcal{Y}} \beta(y) d\mathbb{P}_1(y)}^{\text{Do not depend on } \pi} \right\} =$$

$$\inf_{\pi \in \Pi(\mathbb{P}_0)} \left\{ \mathrm{KL}(\pi || \pi^{W^\epsilon}) - \int_{\mathcal{Y}} \beta(y) d\pi(y) + \int_{\mathcal{Y}} \beta(y) d\mathbb{P}_1(y) \right\} = \inf_{\pi \in \Pi(\mathbb{P}_0)} \widetilde{\mathcal{L}}(\beta, \pi). \tag{29}$$

Here in transition to line (29), we use our Lemma B.2. It remains to take $\sup_\beta$ in equality between (28) and (29) and then recall (27) to finish the proof and obtain desired (24). $\qquad\square$

Thus, we can obtain the value $\mathcal{L}^*$ (8) by solving maximin problem (24) with only one constraint $\pi \in \Pi(\mathbb{P}_0)$. Moreover, our following lemma shows that in all optimal pairs (β^*, π^*) which solve maximin problem (24), π^* is necessary the unique entropic OT plan between $\mathbb{P}_0$ and $\mathbb{P}_1$.

Lemma B.4 (Entropic OT plan solves the relaxed entropic OT problem)**.** *Let π^* be the entropic OT plan between $\mathbb{P}_0$ and $\mathbb{P}_1$. For every optimal $\beta^* \in \mathrm{argsup}_\beta \inf_{\pi \in \Pi(\mathbb{P}_0)} \mathcal{L}(\beta, \pi)$, we have*

$$\pi^* = \underset{\pi \in \Pi(\mathbb{P}_0)}{\mathrm{arginf}} \, \widetilde{\mathcal{L}}(\beta^*, \pi). \tag{30}$$

Proof of Lemma B.4. Since β^* is optimal, we know from Lemma B.3 that $\inf_{\pi \in \Pi(\mathbb{P}_0)} \mathcal{L}(\beta^*, \pi) = \mathcal{L}^*$. Thanks to $\pi^* \in \Pi(\mathbb{P}_0, \mathbb{P}_1)$, we have $\pi_1^* = \mathbb{P}_1$. We substitute π^* to $\mathcal{L}(\beta^*, \pi)$ and obtain

$$\mathcal{L}(\beta^*, \pi^*) = \mathrm{KL}(\pi^* || \pi^{W^\epsilon}) + \int_{\mathcal{Y}} \beta(y) d\mathbb{P}_1(y) - \int_{\mathcal{Y}} \beta(y) \overbrace{d\pi_1^*(y)}^{=d\mathbb{P}_1(y)} = \mathrm{KL}(\pi^* || \pi^{W^\epsilon}) = \mathcal{L}^*. \tag{31}$$

The functional $\pi \mapsto \mathcal{L}(\beta^*, \pi)$ is strictly convex (in the convex subset of $\Pi(\mathbb{P}_0)$ of distributions π for which $\mathrm{KL}(\pi || \pi^{W^\epsilon}) < \infty$). Thus, it has a unique minimizer, which is π^*. $\qquad\square$

From our Lemmas B.3 and B.4 it follows that to get the OT plan π^* one may solve the maximin problem (24) to obtain an optimal saddle point (β^*, π^*). Unfortunately, it is challenging to estimate $\mathrm{KL}(\pi || \pi^{W^\epsilon})$ from samples, which limits the usage of this objective in practice.

B.2 Equivalence of EOT and DSB relaxed problems

Below we show how to relax SB problem (11) and link its solution to the relaxed entropic OT (24).

For a given $\beta \in \mathcal{C}_{b,2}(\mathcal{Y})$, we define an auxiliary process T^β such that its conditional distributions are $T_{|x,y}^\beta = W_{|x,y}^\epsilon$ and its joint distribution π^{T^β} at $t = 0, 1$ is given by π^β.

To simplify many of upcoming formulas, we introduce $C_\beta \stackrel{\text{def}}{=} \int_{\mathcal{X}} \log C_\beta^x d\mathbb{P}_0(x)$. Also, we introduce $\mathcal{F}(\mathbb{P}_0)$ to denote the set of processes starting at $\mathbb{P}_0$ at time $t = 0$.

Lemma B.5 (Inner objectives of relaxed EOT and SB are KL with T^β and π^{T^β}). *For $\pi \in \Pi(\mathbb{P}_0)$ and $T \in \mathcal{F}(\mathbb{P}_0)$, the following equations hold:*

$$\widetilde{\mathcal{L}}(\beta, \pi) = \mathrm{KL}(\pi || \pi^{T^\beta}) - C_\beta + \int_\mathcal{Y} \beta(y) d\mathbb{P}_1(y), \tag{32}$$

$$\mathcal{L}(\beta, T) = \mathrm{KL}(T || T^\beta) - C_\beta + \int_\mathcal{Y} \beta(y) d\mathbb{P}_1(y). \tag{33}$$

Note that the last two terms in each line depend only on β but not on π or T.

Proof of Lemma B.5. The first equation (32) directly follows from Lemma B.2. Now we prove (33):

$$\mathcal{L}(\beta, T) - \int_\mathcal{Y} \beta(y) d\mathbb{P}_1(y) = \mathrm{KL}(T || W^\epsilon) - \int \beta(y) d\pi_1^T(y) =$$

$$\mathrm{KL}(\pi^T || \pi^{W^\epsilon}) + \int_{\mathcal{X} \times \mathcal{Y}} \mathrm{KL}(T_{|x,y} || W^\epsilon_{|x,y}) d\pi^T(x,y) - \int \beta(y) d\pi_1^T(y) = \tag{34}$$

$$\mathrm{KL}(\pi^T || \pi^{T^\beta}) - C_\beta + \int_{\mathcal{X} \times \mathcal{Y}} \mathrm{KL}(T_{|x,y} || W^\epsilon_{|x,y}) d\pi^T(x,y) =$$

$$\mathrm{KL}(\pi^T || \pi^{T^\beta}) - C_\beta + \int_{\mathcal{X} \times \mathcal{Y}} \mathrm{KL}(T_{|x,y} || T^\beta_{|x,y}) d\pi^T(x,y) = \mathrm{KL}(T || T^\beta) - C_\beta. \tag{35}$$

In the transition to line (34), we use the disintegration formula (6). In line (35), we use the definition of T^β, i.e., we exploit the fact that $T^\beta_{|x,y} = W^\epsilon_{|x,y}$ and again use (6). $\qquad \square$

As a result of Lemma B.5, we obtain the following important corollary.

Corollary B.6 (The solution to the inner problem of relaxed SB is a diffusion). *Consider the problem*

$$\inf_{T \in \mathcal{F}(\mathbb{P}_0)} \mathcal{L}(\beta, T). \tag{36}$$

Then T^β is the unique optimizer of (36) and it holds that $T^\beta \in \mathcal{D}(\mathbb{P}_0)$, i.e., it is a diffusion process:

$$T^\beta = \operatorname*{arginf}_{T \in \mathcal{F}(\mathbb{P}_0)} \mathcal{L}(\beta, T) = \operatorname*{arginf}_{T_f \in \mathcal{D}(\mathbb{P}_0)} \mathcal{L}(\beta, T_f). \tag{37}$$

Proof. Thanks to (33), we see that T^β is the unique minimizer of (36). Now let $\mathbb{Q} \overset{\text{def}}{=} \pi_1^{T^\beta}$. Then

$$T^\beta = \operatorname*{arginf}_{T \in \mathcal{F}(\mathbb{P}_0)} \mathcal{L}(\beta, T) = \operatorname*{arginf}_{T \in \mathcal{F}(\mathbb{P}_0)} \left[\mathrm{KL}(T || W^\epsilon) - \int_\mathcal{Y} \beta(y) d\pi_1^T(y) \right] =$$

$$\operatorname*{arginf}_{T \in \mathcal{F}(\mathbb{P}_0, \mathbb{Q})} \left[\mathrm{KL}(T || W^\epsilon) - \underbrace{\int_\mathcal{Y} \beta(y) d\pi_1^T(y)}_{=\text{Const, since } \pi_1^T = \pi_1^{T^\beta} = \mathbb{Q}} \right] = \operatorname*{arginf}_{T \in \mathcal{F}(\mathbb{P}_0, \mathbb{Q})} \mathrm{KL}(T || W^\epsilon) =$$

$$\operatorname*{arginf}_{T_f \in \mathcal{D}(\mathbb{P}_0, \mathbb{Q})} \mathrm{KL}(T_f || W^\epsilon) = \operatorname*{arginf}_{T_f \in \mathcal{D}(\mathbb{P}_0, \mathbb{Q})} \frac{1}{2\epsilon} \mathbb{E}_{T_f} \left[\int_0^1 ||f(X_t, t)||^2 dt \right]. \tag{38}$$

In transition to (38), we use the fact that the process solving the Schrödinger Bridge (this time between $\mathbb{P}_0$ and $\mathbb{Q}$) with the Wiener Prior is a diffusion process (see Dynamic SB problem in §2.2 for details). As a result, we obtain $T^\beta \in \mathcal{D}(\mathbb{P}_0, \mathbb{Q}) \subset \mathcal{D}(\mathbb{P}_0)$ and finish the proof. $\qquad \square$

Below we show that for a given β, minimization of the SB relaxed functional $\mathcal{L}(\beta, T_f)$ over T_f is equivalent to the minimization of relaxed EOT functional $\widetilde{\mathcal{L}}(\beta, \pi)$ (24) with the same β.

Lemma B.7 (Equivalence of the inf values of the relaxed functionals). *It holds that*

$$\inf_{T_f \in \mathcal{D}(\mathbb{P}_0)} \mathcal{L}(\beta, T_f) = \inf_{\pi \in \Pi(\mathbb{P}_0)} \widetilde{\mathcal{L}}(\beta, \pi) = -C_\beta + \int \beta(y) d\mathbb{P}_1(y). \tag{39}$$

Moreover, the unique minimizers are given by $T^\beta \in \mathcal{D}(\mathbb{P}_0)$ and $\pi^{T^\beta} \in \Pi(\mathbb{P}_0)$, respectively.

Proof of Lemma B.7. Follows from Lemma B.5 and Corollary B.6. □

Finally, we see that both the maximin problems are equivalent.

Corollary B.8 (Equivalence of EOT and DSB maximin problems)**.** *It holds that*

$$\mathcal{L}^* = \sup_\beta \inf_{T_f \in \mathcal{D}(\mathbb{P}_0)} \mathcal{L}(\beta, T_f) = \sup_\beta \inf_{\pi \in \Pi(\mathbb{P}_0)} \widetilde{\mathcal{L}}(\beta, \pi) \qquad (40)$$

Proof of Corollary B.8. We take $\sup_\beta$ of both parts in equation (39). □

Also, it follows that the maximization of $\inf_{T_f \in \mathcal{D}(\mathbb{P}_0)} \mathcal{L}(\beta, T_f)$ over β allows to solve entropic OT.

B.3 Proofs of main results

Finally, after long preparations, we prove our main Theorem 4.1.

Proof of Theorem 4.1 and Corollary 4.2. From our Lemma B.7 and Corollary B.8 it follows that

$$\beta^* \in \operatorname*{argsup}_\beta \inf_{T_f \in \mathcal{D}(\mathbb{P}_0)} \mathcal{L}(\beta, T_f) \Leftrightarrow \beta^* \in \operatorname*{argsup}_\beta \inf_{\pi \in \Pi(\mathbb{P}_0)} \widetilde{\mathcal{L}}(\beta, \pi),$$

i.e., both maximin problems share the same optimal β^*. Thanks to our Lemma B.7, we already know that the process $T^{\beta^*} \in \mathcal{D}(\mathbb{P}_0)$ and the plan $\pi^{T^{\beta^*}} \in \Pi(\mathbb{P}_0)$ are the unique minimizers of problems

$$\inf_{T_f \in \mathcal{D}(\mathbb{P}_0)} \mathcal{L}(\beta^*, T_f) = \inf_{\pi \in \Pi(\mathbb{P}_0)} \widetilde{\mathcal{L}}(\beta^*, \pi),$$

respectively. Therefore, $T_{f^*} = T^{\beta^*}$ and, in particular, $\pi^{T_{f^*}} = \pi^{T^{\beta^*}}$. Moreover, since $(\beta^*, \pi^{T_{f^*}})$ is an optimal saddle point for $\widetilde{\mathcal{L}}$, from Lemma B.4 we conclude that $\pi^{T_{f^*}} = \pi^*$, i.e., $\pi^{T_{f^*}}$ is the **EOT plan** between $\mathbb{P}_0$ and $\mathbb{P}_1$. In particular, $\pi^{T_{f^*}} \in \Pi(\mathbb{P}_0, \mathbb{P}_1)$ which also implies that $T_{f^*} \in \mathcal{D}(\mathbb{P}_0, \mathbb{P}_1)$. The last step is to derive

$$\mathcal{L}^* = \mathcal{L}(\beta^*, T_{f^*}) = \mathrm{KL}(T_{f^*} || W^\epsilon) + \underbrace{\int_{\mathcal{Y}} \beta^*(y) d\mathbb{P}_1(y) - \int_{\mathcal{Y}} \beta^*(y) \overbrace{d\pi_1^{T_{f^*}}(y)}^{=d\mathbb{P}_1(y)}}_{=0 \text{ since } T_{f^*} \in \mathcal{D}(\mathbb{P}_0, \mathbb{P}_1)} = \mathrm{KL}(T_{f^*} || W^\epsilon).$$

which concludes that T_{f^*} is the **solution to SB** (5). □

Proof of Theorem 4.3. <u>Part 1.</u> From Lemma B.5 and Corollary B.6 it follows that that $\inf_{T_f} \mathcal{L}(\hat{\beta}, T_f)$ has the unique minimizer $T^{\widehat{\beta}}$ whose conditional distributions are $T_{|x,y}^{\widehat{\beta}} = W_{|x,y}^\epsilon$. Therefore,

$$\epsilon_1 = \mathcal{L}(\hat{\beta}, T_{\hat{f}}) - \inf_{T_f} \mathcal{L}(\hat{\beta}, T_f) =$$

$$\left[\mathrm{KL}(T_{\hat{f}} || T^{\widehat{\beta}}) - C_{\hat{\beta}} + \int_{\mathcal{Y}} \widehat{\beta}(y) d\mathbb{P}_1(y) \right] - \left[-C_{\widehat{\beta}} + \int_{\mathcal{Y}} \widehat{\beta}(y) d\mathbb{P}_1(y) \right] = \mathrm{KL}(T_{\hat{f}} || T^{\widehat{\beta}}). \qquad (41)$$

<u>Part 2.</u> Now we consider ϵ_2. We know that

$$\mathcal{L}^* = \mathrm{KL}(T_{f^*} || W^\epsilon) =$$

$$\mathrm{KL}(\pi^{T_{f^*}} || \pi^{W^\epsilon}) + \int_{\mathcal{X} \times \mathcal{Y}} \mathrm{KL}(T_{f^*|x,y} || W_{|x,y}^\epsilon) d\pi^{T_{f^*}}(x,y) = \mathrm{KL}(\pi^{T_{f^*}} || \pi^{W^\epsilon}).$$

From Lemma B.5 and Corollary B.6, we also know that

$$\inf_{T_f} \mathcal{L}(\hat{\beta}, T_f) = -C_{\hat{\beta}} + \int \hat{\beta}(y) d\mathbb{P}_1(y).$$

Therefore:

$$\epsilon_2 = \mathcal{L}^* - \inf_{T_f} \mathcal{L}(\hat{\beta}, T_{f^*}) = \mathrm{KL}(\pi^{T_{f^*}} || \pi^{W^\epsilon}) + C_{\hat{\beta}} - \int \hat{\beta}(y) d\mathbb{P}_1(y) =$$

19

$$\mathrm{KL}(\pi^{T_{f^*}}||\pi^{W_\epsilon}) + \int_\mathcal{X} \log C^x_{\hat\beta} d\mathbb{P}_0(x) - \int \hat\beta(y) d\mathbb{P}_1(y) =$$

$$\int_\mathcal{X} \mathrm{KL}(\pi^{T_{f^*}}(\cdot|x)||\pi^{W^\epsilon}(\cdot|x)) d\mathbb{P}_0(x) + \int_\mathcal{X} \log C^x_{\hat\beta} d\mathbb{P}_0(x) - \int \hat\beta(y) d\mathbb{P}_1(y) =$$

$$\int_\mathcal{X} \mathrm{KL}(\pi^{T_{f^*}}(\cdot|x)||\pi^{W^\epsilon}(\cdot|x)) d\mathbb{P}_0(x) + \int_\mathcal{X} \log C^x_{\hat\beta} d\mathbb{P}_0(x) - \int \hat\beta(y) d\pi^{T_{f^*}}(y|x) d\mathbb{P}_0(x) =$$

$$\int_\mathcal{X} \left\{ \mathrm{KL}(\pi^{T_{f^*}}(\cdot|x)||\pi^{W^\epsilon}(\cdot|x)) + \log C^x_{\hat\beta} - \int_\mathcal{Y} \hat\beta(y) d\pi^{T_{f^*}}(y|x) \right\} d\mathbb{P}_0(x) =$$

$$\int_\mathcal{X} \left\{ \mathrm{KL}(\pi^{T_{f^*}}(\cdot|x)||\pi^{T^{\hat\beta}}(\cdot|x)) - \log C^x_{\hat\beta} + \log C^x_{\hat\beta} \right\} d\mathbb{P}_0(x) =$$

$$\int_\mathcal{X} \mathrm{KL}(\pi^{T_{f^*}}(\cdot|x)||\pi^{T^{\hat\beta}}(\cdot|x)) d\mathbb{P}_0(x) = \mathrm{KL}(\pi^{T_{f^*}}||\pi^{T^{\hat\beta}}) =$$

$$\mathrm{KL}(\pi^{T_{f^*}}||\pi^{T^{\hat\beta}}) + \underbrace{\int_{\mathcal{X}\times\mathcal{Y}} \mathrm{KL}(T_{f^*|x,y}||T^{\hat\beta}_{|x,y}) d\pi^{T_{f^*}}(x,y)}_{=0,\ \text{since } T_{f^*|x,y}=T^{\hat\beta}_{|x,y}=W^\epsilon_{|x,y}} = \mathrm{KL}(T_{f^*}||T^{\hat\beta}). \quad (42)$$

Thus, we obtain $\epsilon_2 = \mathrm{KL}(T_{f^*}||T^{\hat\beta})$.

Part 3. By summing (41) and (42) and using the Pinsker inequality, we obtain

$$\epsilon_1 + \epsilon_2 = \mathrm{KL}(T_{\hat f}||T^{\hat\beta}) + \mathrm{KL}(T_{f^*}||T^{\hat\beta}) \geq 2\rho^2_{\mathrm{TV}}(T_{\hat f}, T^{\hat\beta}) + 2\rho^2_{\mathrm{TV}}(T_{f^*}, T^{\hat\beta}) \geq$$
$$\left[\rho_{\mathrm{TV}}(T_{\hat f}, T^{\hat\beta}) + \rho_{\mathrm{TV}}(T_{f^*}, T^{\hat\beta})\right]^2 \geq \rho^2_{\mathrm{TV}}(T_{\hat f}, T_{f^*}). \quad (43)$$

Here we use the triangle inequality in line (43). Therefore, $\rho_{\mathrm{TV}}(T_{\hat f}, T_{f^*}) \leq \sqrt{\epsilon_1 + \epsilon_2}$.

Part 4. By summing (41) and (42) and using the Pinsker inequality, we obtain

$$\epsilon_1 + \epsilon_2 = \mathrm{KL}(T_{\hat f}||T^{\hat\beta}) + \mathrm{KL}(T_{f^*}||T^{\hat\beta}) =$$

$$\mathrm{KL}(\pi^{T_{\hat f}}||\pi^{T^{\hat\beta}}) + \int_{\mathcal{X}\times\mathcal{Y}} \mathrm{KL}(T_{\hat f|x,y}||T^{\hat\beta}_{|x,y}) d\pi^{T_{\hat f}}(x,y) +$$

$$\mathrm{KL}(\pi^{T_{f^*}}||\pi^{T^{\hat\beta}}) + \int_{\mathcal{X}\times\mathcal{Y}} \mathrm{KL}(T_{f^*|x,y}||T^{\hat\beta}_{|x,y}) d\pi^{T_{f^*}}(x,y) \geq$$

$$\mathrm{KL}(\pi^{T_{\hat f}}||\pi^{T^{\hat\beta}}) + \mathrm{KL}(\pi^{T_{f^*}}||\pi^{T^{\hat\beta}}) \geq 2\rho^2_{\mathrm{TV}}(\pi^{T_{\hat f}}, \pi^{T^{\hat\beta}}) + 2\rho^2_{\mathrm{TV}}(\pi^{T_{f^*}}, \pi^{T^{\hat\beta}}) \geq$$
$$\left[\rho_{\mathrm{TV}}(\pi^{T_{\hat f}}, \pi^{T^{\hat\beta}}) + \rho_{\mathrm{TV}}(\pi^{T_{f^*}}, \pi^{T^{\hat\beta}})\right]^2 \geq \rho^2_{\mathrm{TV}}(\pi^{T_{\hat f}}, \pi^{T_{f^*}}).$$

Thus, $\rho_{\mathrm{TV}}(\pi^{T_{f^*}}, \pi^{T^{\hat\beta}}) \leq \sqrt{\epsilon_1 + \epsilon_2}$. $\qquad\square$

C Euler-Maruyama

In our Algorithm 1, at both the training and the inference stages, we use the Euler-Maruyama Algorithm 2 to solve SDE.

D Drift Norm Constant Multiplication Invariance

Our Algorithm 1 aims to solve the following optimization problem:

$$\sup_\beta \inf_{T_f \in \mathcal{D}(\mathbb{P}_0)} \underbrace{\left\{ \mathbb{E}_{T_f}\left[\int_0^1 C||f(X_t,t)||^2 dt\right] + \int_\mathcal{Y} \beta(y) d\mathbb{P}_1(y) - \int_\mathcal{Y} \beta(y) d\mathbb{P}_1^{T_f}(y) \right\}}_{\overset{\text{def}}{=} \mathcal{L}^C(\beta, T_f)},$$

with $C = 1$. At the same time, we use $C = \frac{1}{2\epsilon}$ in our theoretical derivations (12). We emphasize that the actual value of $C > 0$ **does not affect** the optimal solution T_{f^*} to this problem. Specifically, if

20

Algorithm 2: Euler-Maruyama algorithm

Input : batch of initial states X_0 at time moment $t = 0$;
SDE drift network $f_\theta : \mathbb{R}^D \times [0,1] \to \mathbb{R}^D$;
number of steps for the SDE solver $N \geq 1$;
noise variance $\epsilon \geq 0$.

Output : batches $\{X_n\}_{n=0}^N$ of intermediate states at $t = \frac{n}{N}$ simulating the proccess
$dX_t = f(X_t, t)dt + \sqrt{\epsilon}dW_t$;
batches $\{f_n\}_{n=0}^N$ of drift values $f(X_n, t_n)$ at $t = \frac{n-1}{N}$ simulating the process;

$\Delta t \leftarrow \frac{1}{N}$;
for $t = 1, 2, \ldots, N$ **do**
    for $i = 1, 2, \ldots, |X_0|$ **do**
        Sample noise W from $\mathcal{N}(0, I)$;
        $f_{t-1,i} \leftarrow f(X_{t-1}, t-1)$;
        $X_{t,i} \leftarrow X_{t-1,i} + f_{t-1,i}\Delta t + \sqrt{\epsilon \Delta t}W$;

(β^*, T_{f^*}) is the optimal point for the problem with $C = 1$, then $(\widetilde{C}\beta^*, T_{f^*})$ is the optimal point for $C = \widetilde{C}$. Indeed, for a pair (β, T_f) it holds that

$$\mathcal{L}^1(\beta, T_f) = \mathbb{E}_{T_f}\Big[\int_0^1 ||f(X_t, t)||^2 dt\Big] + \int_{\mathcal{Y}} \beta(y)d\mathbb{P}_1(y) - \int_{\mathcal{Y}} \beta(y)d\mathbb{P}_1^{T_f}(y) =$$

$$\frac{1}{\widetilde{C}}\Big\{\mathbb{E}_{T_f}\Big[\int_0^1 \widetilde{C}||f(X_t, t)||^2 dt\Big] + \int_{\mathcal{Y}} \widetilde{C}\beta(y)d\mathbb{P}_1(y) - \int_{\mathcal{Y}} \widetilde{C}\beta(y)d\mathbb{P}_1^{T_f}(y)\Big\} =$$

$$\frac{1}{\widetilde{C}}\Big\{\mathbb{E}_{T_f}\Big[\int_0^1 \widetilde{C}||f(X_t, t)||^2 dt\Big] + \int_{\mathcal{Y}} \widetilde{\beta}(y)d\mathbb{P}_1(y) - \int_{\mathcal{Y}} \widetilde{\beta}(y)d\mathbb{P}_1^{T_f}(y)\Big\} = \frac{1}{\widetilde{C}}\mathcal{L}^{\widetilde{C}}(\widetilde{\beta}, T_f), \quad (44)$$

where we use $\widetilde{\beta} \overset{\text{def}}{=} \widetilde{C}\beta$. Hence problems $\sup_\beta \inf_{T_f} \mathcal{L}^1(\beta, T_f)$ and $\sup_{\widetilde{\beta}} \inf_{T_f} \mathcal{L}^{\widetilde{C}}(\widetilde{\beta}, T_f)$ can be viewed as **equivalent** in the sense that one can be derived one from the other via the change of variables and multiplication by $\widetilde{C} > 0$. For completeness, we also note that the change of variables $\beta \leftrightarrow \widetilde{\beta}$ actually preserves the functional class of β, i.e., $\beta \in \mathcal{C}_{b,2}(\mathcal{Y}) \Longleftrightarrow \widetilde{\beta} \in \mathcal{C}_{b,2}(\mathcal{Y})$.

For convenience, we get rid of dependence on ϵ in the objective (12) and consider $\mathcal{L}^1$ for optimization, i.e., use $C = 1$ in Algorithm 1. Still the dependence on ϵ remains in $\sup_\beta \inf_{T_f} \mathcal{L}^1(\beta, T_f)$ as $T_f \in \mathcal{D}(\mathbb{P}_0)$ is a diffusion process with volatility ϵ. Interestingly, this point of view (optimizing $\mathcal{L}^1$ instead of $\mathcal{L}^{\frac{1}{2\epsilon}}$) **technically** allows to consider even $\epsilon = 0$. In this case, the optimization is performed over *deterministic* trajectories T_f determined by the velocity field $f(X_t, t)$. The problem $\sup_\beta \inf_{T_f} \mathcal{L}^1(\beta, T_f)$ may be viewed as a saddle point reformulation of the **unregularized** OT with the quadratic cost in the *dynamic* form, also known as the Benamou-Brenier formula [43, §6.1]. This particular case is out of scope of our paper (it is not EOT/SB) and we do not study the properties of $\mathcal{L}^1$ in this case. However, for completeness, we provide experimental results for $\epsilon = 0$.

E ENOT for Toy Experiments and High-dimensional Gaussians

In 2D toy experiments, we consider 2 tasks: *Gaussian → 8 gaussians* and *Gaussian → Swiss roll*. Results for the last one (Figure 5) are qualitatively similar to results of the first one (Figure 2), which we discussed earlier (§5.1). For both tasks, we parametrize the SDE drift function in Algorithm 1 by a feedforward neural network f_θ with 3 inputs, 3 linear layers (100 hidden neurons and ReLU activations) and 2 outputs. As inputs, we use 2 coordinates and time value t (as is). Analogically, we parametrize the potential by a feedforward neural network β_ϕ with 2 inputs, 3 linear layers (100 hidden neural and ReLU activations) and 2 outputs. In all the cases, we use $N = 10$ discretization steps for solving SDE by Euler-Maruyama Algorithm 2, Adam with lr $= 10^{-4}$, batch size 512. We train the model for 20000 total iterations of β_ϕ, and on each of them, we do $K_f = 10$ updates for the SDE drift function f_θ.

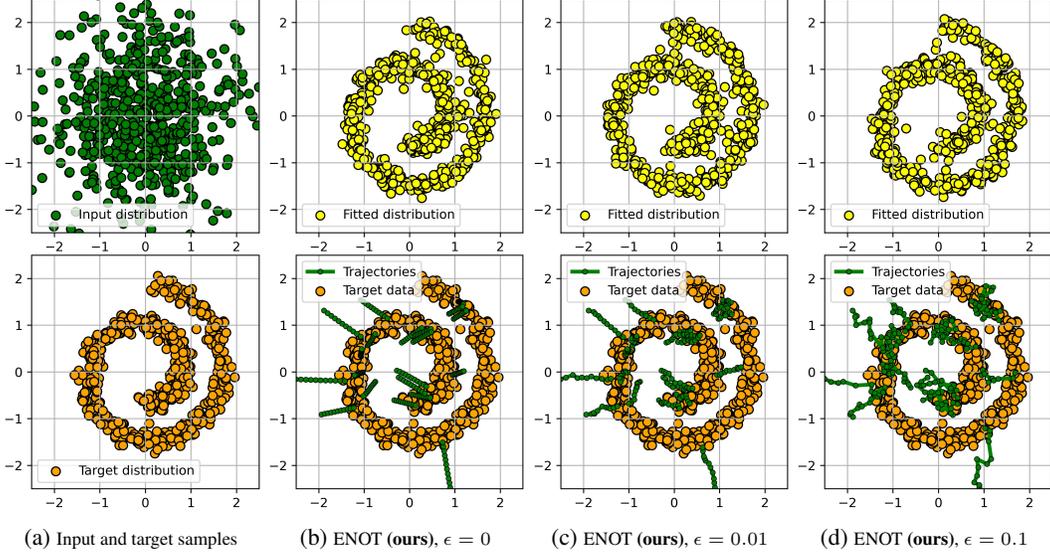Figure 5: Gaussian → Swiss roll, learned stochastic process with ENOT **(ours)**.

In the experiments with high-dimensional Gaussians, we use exactly the same setup as for toy 2D experiments but chose $N = 200$ discretization steps for SDE, all hidden sizes in neural networks are 512, and we train our model for 10000 iterations. To illustrate the stability of the algorithm, we provide the plot of BW_2^2-UVP (%) between the ground truth EOT plan π^* and the learned plan π of ENOT during training for $DIM = 128$ in Figure 6.
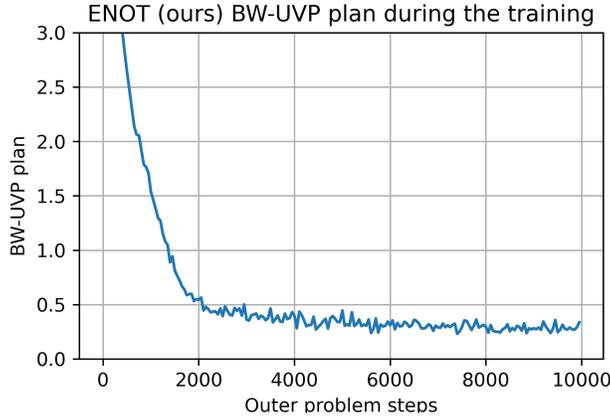

Figure 6: BW_2^2-UVP $\downarrow$ (%) between the the EOT plan π^* and the learned plan π of ENOT and MLE-SB during the training (DIM = 128).

F   ENOT for Colored MNIST and Unpaired Super-resolution of Celeba Faces

For the image tasks (§5.3, §5.4), we find out that using the following reparametrization of Euler-Maruyama Algorithm 2 considerably improves the quality of our Algorithm 1. In the Euler-Maruyama Algorithm 2, instead of using a neural network to parametrize drift function $f(X_t, t)$ and calculating the next state as $X_{t+1} = X_t + f(X_t, t)\Delta t + \sqrt{\epsilon \Delta t}$, we parametrize $g(X_t, t) = X_t + f(X_t, t)\Delta t$ by a neural network g_θ, and calculate the next state as $X_{t+1} = g_\theta(X_t, t) + \sqrt{\epsilon \Delta t}$. In turn, the drift function is given by $f(X_t, t) = \frac{1}{\Delta t} g(X_t, t) - X_t$. Also, we do not add noise at the last step of the Euler-Maruyama simulation because we find out that it provides better empirical performance.
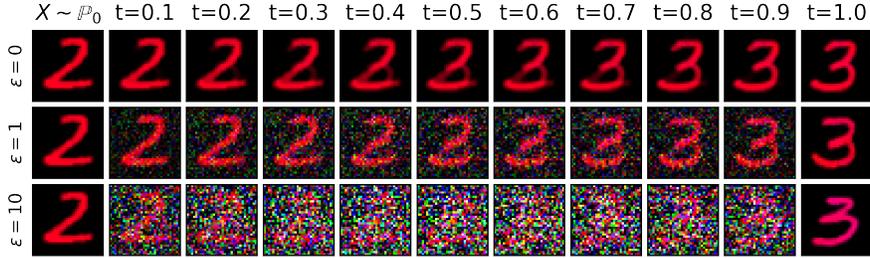
22

Figure 7: Trajectories from our learned ENOT **(ours)** for colored MNIST for different ϵ.

We use WGAN-QC discriminator's ResNet architecture [5] for the potential β. We use UNet [6] as $g_\theta(X_t, t)$ of SDE in our model. To condition it on t, we first obtain the embedding of t by using the positional embedding [7]. Then we add conditional instance normalization (CondIN) layers after each UNet's upscaling block [8]. We use Adam with lr $= 10^{-4}$, batch size 64 and 10:1 update ratio for f_θ/β_ϕ. For $\epsilon = 0$ and $\epsilon = 1$ our model converges in ≈ 20000 iterations, while for $\epsilon = 10$ it takes ≈ 70000 iteration to convergence. The last setup takes more iterations to converge because adding noise with higher variance during solving SDE by Euler-Maruyama Algorithm 2 increases the variance of stochastic gradients.

In the unpaired super-resolution of Celeba faces, we use the same experimental setup as for the colored MNIST experiment. It takes ≈ 40000 iterations for $\epsilon = 0$ and ≈ 70000 iterations for $\epsilon = 1$ and $\epsilon = 10$ to converge. In Figures 7, 1 we present trajectories provided by our algorithm for Colored MNIST and Celeba experiments.

Computational complexity. In the most challenging task (§5.4), ENOT converges in one week on $2\times$ A100 GPUs.

G Details of the baseline methods

In this section, we discuss details of the baseline methods with which we compare our method.

G.1 Gaussian case (§5.2).

SCONES [14]. We use the code from the authors' repository

$$\texttt{https://github.com/mdnls/scones-synthetic}$$

for their evaluation in the Gaussian case. We employ their configuration `blob/main/config.py`.

LSOT [45]. We use the part of the code of SCONES corresponding to learning dual OT potentials `blob/main/cpat.py` and the barycentric projection `blob/main/bproj.py` in the Gaussian case with configuration `blob/main/config.py`.

FB-SDE-J [10]. We utilize the official code from

$$\texttt{https://github.com/ghliu/SB-FBSDE}$$

with their configuration `blob/main/configs/default_checkerboard_config.py` for the checkerboard-to-noise toy experiment, changing the number of steps of dynamics from 100 to 200 steps. Since their hyper-parameters are developed for their 2-dimensional experiments, we increase the number of iterations for dimensions 16, 64 and 128 to 15 000.

FB-SDE-A [10]. We also take the code from the same repository as above. We base our configuration on the authors' one (`blob/main/configs/default_moon_to_spiral_config.py`) for the moon-to-spiral experiment. As earlier, we increase the number of steps of dynamics up to 200. Also, we change the number of training epochs for dimensions 16, 64 and 128 to 2,4 and 8 correspondingly.

---

[5]github.com/harryliew/WGAN-QC
[6]github.com/milesial/Pytorch-UNet
[7]github.com/rosinality/denoising-diffusion-pytorch
[8]github.com/kgkgzrtk/cUNet-Pytorch

DiffSB [15]. We utilize the official code from

<div align="center">

`https://github.com/JTT94/diffusion_schrodinger_bridge`

</div>

with their configuration `blob/main/conf/dataset/2d.yaml` for toy problems. We increase the amount of steps of dynamics to 200 and the number of steps of IPF procedure for dimensions 16, 64 and 128 to 30, 40 and 60, respectively.

MLE-SB [48]. We use the official code from

<div align="center">

`https://github.com/franciscovargas/GP_Sinkhorn`

</div>

with hyper-parameters from `blob/main/notebooks/2D Toy Data/2d_examples.ipynb`. We set the number of steps to 200. As earlier, we increase the number of steps of IPF procedure for dimensions 16, 64 and 128 to 1000, 3500 and 5000, respectively.

G.2 Colored MNIST (§5.3)

SCONES [14]. In order to prepare a score-based model, we use the code from

<div align="center">

`https://github.com/ermongroup/ncsnv2`

</div>

with their configuration `blob/master/configs/cifar10.yml`. Next, we utilize the code of SCONES from the official repository for their unpaired Celeba super-resolution experiment (`blob/main/scones/configs/superres_KL_0.005.yml`). We adapt it for 32×32 ColorMNIST images instead of 64×64 celebrity faces.

DiffSB [15]. We use the official code with their configuration `blob/main/conf/mnist.yaml` adopting it for three-channel ColorMNIST images instead of one-channel MNIST digits.

G.3 CelebA (§5.4)

SCONES [14].   For the SCONES, we use their exact code and configuration from `blob/main/scones/configs/superres_KL_0.005.yml`. As for the score-based model for celebrity faces, we pick the pre-trained model from

<div align="center">

`https://github.com/ermongroup/ncsnv2`

</div>

It is the one used by the authors of SCONES in their paper.

Augmented Cycle GAN [2]. We use the official code from

<div align="center">

`https://github.com/NathanDeMaria/AugmentedCycleGAN`

</div>

with their default hyper-parameters.

ICNN [38]. We utilize the reworked implementation by

<div align="center">

`https://github.com/iamalexkorotin/Wasserstein2Benchmark.`

</div>

which is a non-minimax version [26] of ICNN-based approach [38].   That is, we use `blob/main/notebooks/W2_test_images_benchmark.ipynb` and only change the dataloaders.

H Mean-Field Games

This appendix discusses the relation between the Mean-Field Game problem and Schrödinger Brdiges.

H.1 Intro to the Mean-Field game.

Consider a game with infinitely many small players. At time moment $t = 0$, they are distributed according to $X_0 \sim \rho_0$. Every player controls its behavior through drift α of the SDE:

$$dX_t = \alpha(X_t, t, \rho_t)dt + \sqrt{2\nu}dW_t$$

Here ρ_t is the distribution of all the players at the time moment t. When we consider a specific player, we consider ρ_t as a parameter. Each player aims to minimize the quantity:

$$\mathbb{E}[\int_0^T (L(X_t, \alpha_t, \rho_t) + f(X_t, \rho_t))dt + g(X_T, \rho_T)].$$

Here $L(x, \alpha, \rho)$ is similar to the Lagrange function in physics and describes the cost of moving in some direction given the current position and the other players' distribution. The additional function $f(X_t, \rho_t)$ is interpreted as the cost of the player's interaction at coordinate x with all the others. Now we can introduce the value function $\phi(x, t)$, which for position x and start time t returns the cost in case of the optimal control:

$$\phi(x, t) \stackrel{\text{def}}{=} \inf_\alpha \mathbb{E}[\int_t^T (L(X_t, \alpha_t, \rho_t) + f(X_t, \rho_t))dt + g(X_T, \rho_T)].$$

Before considering the Mean-Field game, we need to define an additional function $H(x, p, \rho)$. It is similar to the Hamilton function and is defined as the Legendre transform of Lagrange function L:

$$H(x, p, \rho) \stackrel{\text{def}}{=} \sup_\alpha [-\alpha p - L(x, \alpha, \rho)].$$

Mean-Field game implies finding the Nash equilibrium for all players of such the game. It is known [1] that the Nash equilibrium is the solution of the system of Hamilton-Jacobi-Bellman (HJB) and Fokker-Planck (FP) PDE equations. For two functions $H(x, p, \rho)$ and $f(x, \rho)$, Mean-Field game formulates as a system of two PDE with two constraints:

$$-\partial_t \phi - \nu \Delta \phi + H(x, \nabla \phi, \rho) = f(x, \rho) \text{ (HJB)}$$

$$-\partial_t \rho - \nu \Delta \rho - \mathbf{div}(\rho \nabla_p H(x, \nabla \phi)) = 0 \text{ (FP)}$$

$$\text{s.t. } \rho(x, 0) = \rho_0 , \phi(x, T) = g(x, \rho(\cdot, T))$$

The solution of this system is two functions $\rho(x, t)$ and $\phi(x, t)$, which describe all players' dynamics. Also, in Nash equilibrium, the specific player's behavior is described by the following SDE:

$$dX_t = -\nabla_p H(X_t, \nabla \phi(X_t, t), \rho)dt + \sqrt{2\nu}dW_t.$$

H.2 Relation to our work.

In recent work [34], the authors show that the Schrodinger Bridger problem could be formulated as a Mean-Field game with hard constraints on distribution $\rho(\cdot, T) = \rho_{target}(\cdot, T)$ via choosing proper function $g(x, \rho(\cdot, T))$ such as:

$$g(x, \rho(\cdot, T)) = \begin{cases} \infty, & \text{if } \rho(\cdot, T) \neq \rho_{target}(\cdot, T) \\ 0, & \rho(\cdot, T) = \rho_{target}(\cdot, T) \end{cases}$$

Also, the authors proposed an extension of DiffSB [34] algorithm for the Mean-Field game problem.

In [33], the authors in their experiments **consider only soft constraints** on the target density. More precisely, they consider only simple constraints such as $g(x, \rho) = ||x - x_{target}||_2$, where x_{target} is a given shared target point for every player, and every player is penalized for being far from this. Such soft constraint force players to have delta distribution at point x_{target}.

To solve the Mean-Field problem, the authors parameterize value function $\phi(x, t)$ by a neural network and use different neural network N_θ to sample from ρ_t. The authors penalize the violation of Mean-Field game PDEs for optimizing these networks. After the convergence, one can sample from the distribution ρ_t by using neural network N_θ. *Approach [33] has the advantage that authors do not need to use SDE solvers, which require more steps with growing parameter ν of diffusion operator.* However, computation of Laplacian and divergence for high-dimensional spaces (e.g., space 12228-dimensional space of 3x64x64 images) at each iteration of the training step may be computationally hard, restricting the applicability of their method to large-scale setups.

In our approach, we initially work with the SDE:

$$dX_t = \alpha(X_t, t, \rho_t)dt + \sqrt{2\nu}dW_t,$$

which describes the player's behavior and use a neural network to parametrize the drift α. We consider only **hard constraints** on the target distribution, $f(X_t, \rho_t) = 0$ and $L(X_t, \alpha_t, \rho_t) = \frac{1}{2}||\alpha_t||^2$ since this variant of Mean-Field game is also the particular case of Schrodinger Bridge problem and is equivalent to the entropic optimal transport. *Since we do not need to compute Laplacian or divergence, our approach scales better with the dimension.* However, for high values of diffusion parameter ν (which is equal to the $\frac{1}{2}\epsilon$ in our notation, where ϵ is the entropic regularization strength), our approach needs more steps for accurate solving of the SDE to provide samples, as we mentioned in limitations.

I   Extending ENOT to other costs

In the main text, we focus only on EOT with the quadratic cost $c(x, y) = \frac{1}{2}||x - y||^2$ which coincides with SB with the Wiener prior W^ϵ. However, one could use a different prior Q_v instead of W^ϵ in (5):

$$Q_v : dX_t = v(X_t, t)dt + \sqrt{\epsilon}dW_t,$$

and solve the problem

$$\inf_{T_f \in \mathcal{D}(\mathbb{P}_0, \mathbb{P}_1)} \text{KL}(T_f || Q_v) = \inf_{T_f \in \mathcal{D}(\mathbb{P}_0, \mathbb{P}_1)} \frac{1}{2\epsilon} \mathbb{E}_{T_f}[\int_0^1 ||f(X_t, t) - v(X_t, t)||^2 dt].$$

Here we just use the known expression (6) for $\text{KL}(T_f || Q_v)$ between two diffusion processes through their drift functions. Using the same derivation as in the main text §2.2, it can be shown that this new problem is equivalent to solving the EOT with cost $c(x, y) = -\log \pi^{Q_v}(y|x)$, where $\pi^{Q_v}(y|x)$ is a conditional distribution of the stochastic process Q_v at time $t = 1$ given the starting point x at time $t = 0$. For example, for W^ϵ (which we consider in the main text) we have

$$c(x, y) = -\log \pi^{W^\epsilon}(y|x) = \frac{1}{2\epsilon}(y - x)^T(y - x) + \text{Const},$$

i.e., we get the quadratic cost. Thus, using different priors for the Schrodinger bridge problem makes it possible to solve Entropic OT for other costs. We conjecture that most of our proofs and derivations can be extended to arbitrary prior process Q_v just by slightly changing the minimax functional (12):

$$\sup_\beta \inf_{T_f} \left(\frac{1}{2\epsilon} \mathbb{E}_{T_f}[\int_0^1 ||f(X_t, t) - v(X_t, t)||^2 dt] + \int_\mathcal{Y} \beta_\phi(y)d\mathbb{P}_1(y) - \int_\mathcal{Y} \beta_\phi(y)d\pi_1^{T_f}(y) \right).$$

We conduct a toy experiment to support this claim and consider Q_v with $\epsilon = 0.01$ and $v(x, t) = \nabla \log p(x)$, where $\log p(x)$ is a 2D distribution with a wave shape, see Figure 8. Intuitively, it means that trajectories should be concentrated in the regions with a high density of p. In Figure e 8, there the grey-scale color map represents the density of p, start points ($\mathbb{P}_0$) are green, target points ($\mathbb{P}_1$) are red, obtained trajectories are pink and mapped points are blue.

J   ENOT for the unregularized OT ($\epsilon = 0$)

Our proposed algorithm is designed to solve entropic OT and the equivalent SB problem. This implies that $\epsilon > 0$. Nevertheless, our algorithm *technically* allows using even $\epsilon = 0$, in which case it presumably computes the unegularized OT map for the quadratic cost. Here we present some empirical evidence supporting this claim as well some theoretical insights.

EMPIRICAL EVIDENCE. We consider the experimental setup with images from the continuous Wasserstein-2 benchmark [28, §4.4]. The images benchmark provides 3 pairs of distributions (Early, Mid, Late) for which the ground truth unregularized OT map for the quadratic cost is known by the construction. Hence, we may compare the map learned with our method ($\epsilon = 0$) with the true one.

We train our method with $\epsilon = 0$ on each of 3 benchmark pairs and present the quantitative results in Table 6. We use the same $\mathcal{L}^2$-UVP metric [28, §4.2] as the authors of the benchmark. As the baselines, we include the results of MM:R method from [28] and the method from [3]. Both methods are minimax and have some similarities with our approach. As we can see, ENOT with $\epsilon = 0$
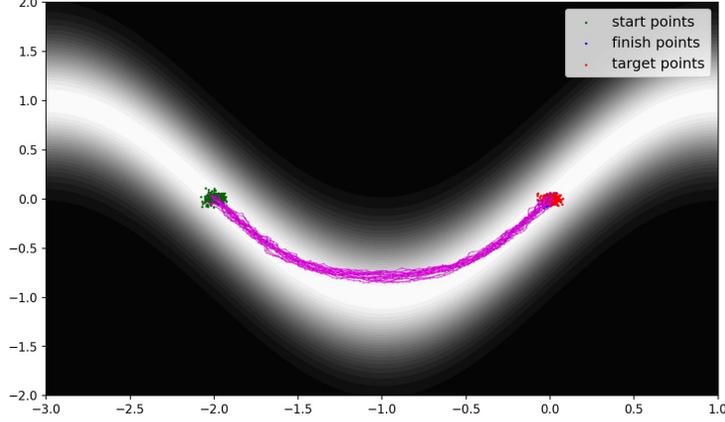
Figure 8: Toy example with ENOT (**ours**) for the complex prior $Q_v : dX_t = v(X_t, t)dt + \sqrt{\epsilon}dW_t$.

Benchmark	Early	Mid	Late
[28]*	1.4	0.4	0.22
[3]*	0.61	0.20	0.09
ENOT (**ours**)	0.77	0.21	0.09

Table 6: Comparison on W2 benchmark. *Results are taken from [3, Table 2].

works better than the MM:R solver but slightly underperforms compared to [3]. This evaluation demonstrates that our method recovers the unregularized OT map for the quadratic cost with the comparable quality to the existing saddle point OT methods.

THEORETICAL INSIGHTS. We see that empirically our method with $\epsilon = 0$ recovers the unregularized OT map. At the same time, this is not supported by our theoretical results as they work exclusively for $\epsilon > 0$ and rely on the properties of the KL divergence.

Overall, it seems like for $\epsilon = 0$ our method yields a saddle point reformulation of the Benamou-Brenier (BB) [8] problem which is also known as the dynamic version of the unregularized OT ($\epsilon = 0$) with the quadratic cost. This problem can be formulated as follows:

$$\inf_{T_f} \left\{ \frac{1}{2}\mathbb{E}_{T_f}[\int_0^1 ||f(X_t, t)||^2 dt] \right\} \quad \text{s.t.} \quad T_f : dX_t = f(X_t, t)dt, \quad X_0 \sim \mathbb{P}_0, X_1 \sim \mathbb{P}_1, \quad (45)$$

i.e., the goal is to find the process T_f of the minimal energy which moves the probability mass of $\mathbb{P}_0$ to $\mathbb{P}_1$. BB (45) is very similar to DSB (11) but there is no multiplier $\frac{1}{\epsilon}$, and the stochastic process T_f is restricted to be deterministic ($\epsilon = 0$). It is governed by a vector field f. Just like the DSB (11) is equivalent to EOT (2), it is known that BB (45) is **equivalent** to unregularized OT with the quadratic cost ($\epsilon = 0$). Namely, the distribution $\pi^{T_{f*}}$ is the unregularized OT plan between $\mathbb{P}_0$ and $\mathbb{P}_1$.

In turn, our Algorithm 1 for $\epsilon = 0$ optimizes the following saddle point objective:

$$\sup_{\beta} \inf_{T_f} \mathcal{L}(\beta, T_f) \stackrel{\text{def}}{=} \sup_{\beta} \inf_{T_f} \left\{ \frac{1}{2}\mathbb{E}_{T_f}[\int_0^1 ||f(X_t, t)||^2 dt] + \int_{\mathcal{Y}} \beta(y)d\mathbb{P}_1(y) - \int_{\mathcal{Y}} \beta(y)d\mathbb{P}_1^{T_f}(y) \right\}, \quad (46)$$

where $T_f : dX_t = f(X_t, t)dt$ with $X_0 \sim \mathbb{P}_0$ (the constraint $X_1 \sim \mathbb{P}_1$ here is lifted) and $\beta \in \mathcal{C}_{2,b}(\mathcal{Y})$. Just like in the Entropic case, functional $\mathcal{L}$ can be viewed as the Lagrangian for BB (45) with β playing the role of the Lagrange multiplier for the constraint $d\pi_1^{T_f}(y) = d\mathbb{P}_1(y)$. Naturally, it is expected that the value (45) coincides with (46), and we provide a *sketch of the proof* of this fact.

Overall, the proof logic is analogous to the Entropic case but the actual proof is much more technical as we can not use the KL-divergence machinery which helps to avoid non-uniqueness, etc.

Step 1 (Auxiliary functional, analog of Lemma B.3). We introduce an auxiliary functional

$$\widetilde{\mathcal{L}}(\beta, H) \stackrel{\text{def}}{=} \int_{\mathcal{X}} \frac{1}{2}||x - H(x)||^2 d\mathbb{P}_0(x) - \int_{\mathcal{X}} \beta(H(x))d\mathbb{P}_0(x) + \int_{\mathcal{Y}} \beta(y)d\mathbb{P}_1(y),$$

27

where β is a potential and $H : \mathbb{R}^D \to \mathbb{R}^D$ is a measurable map. This functional is nothing but the well-known max-min reformulation of static OT problem (in Monge's form) with the quadratic cost [3, Eq. 4], [28, Eq.9]. Hence,

$$\sup_\beta \inf_H \widetilde{\mathcal{L}}(\beta, H) = \underbrace{\inf_{H \sharp \mathbb{P}_0 = \mathbb{P}_1} \int_\mathcal{X} \frac{1}{2} \|x - H(x)\|^2 d\mathbb{P}_0(x)}_{\overset{\text{def}}{=} \mathcal{L}^*}.$$

Step 2 (Solution of the inner problem is always an OT map). An existence of some minimizer $H = H^\beta$ in $\inf_H \widetilde{L}(\beta, H)$ can be deduced from the measurable argmin selection theorem, e.g., [21, Theorem 18.19]. For this H^β we consider $\mathbb{P}' \overset{def}{=} H^\beta \sharp \mathbb{P}_0$. Recall that

$$H^\beta \in \underset{H}{\operatorname{arginf}} \, \widetilde{L}(\beta, H) = \underset{H}{\operatorname{arginf}} \int_\mathcal{X} \Big\{ \frac{\|x - H(x)\|^2}{2} - \beta(H(x)) \Big\} d\mathbb{P}_0(x).$$

Here we may add the fictive constraint $H \sharp \mathbb{P}_0 = \mathbb{P}'$ which is anyway satisfied by H^β and get

$$H^\beta \in \underset{H \sharp \mathbb{P}_0 = \mathbb{P}'}{\operatorname{arginf}} \int_\mathcal{X} \Big\{ \frac{\|x - H(x)\|^2}{2} - \beta(H(x)) \Big\} d\mathbb{P}_0(x) = \underset{H \sharp \mathbb{P}_0 = \mathbb{P}'}{\operatorname{arginf}} \int_\mathcal{X} \frac{\|x - H(x)\|^2}{2} d\mathbb{P}_0(x).$$

The last equality holds since $\int \beta(H(x)) d\mathbb{P}_0(x) = \int \beta(y) d\mathbb{P}'(y)$ does not depend on the choice of H due to the constraint $H \sharp \mathbb{P}_0 = \mathbb{P}'$. The latter is the OT problem between $\mathbb{P}$ and $\mathbb{P}'$ and we see that H^β is its solution.

Step 3 (Equivalence for inner objective values). Since H^β is the OT map between $\mathbb{P}_0, \mathbb{P}'$ (it is unique as $\mathbb{P}_0$ is absolutely continuous [43]), it can be represented as an ODE solution T_{f^β} to the Benamour Brenier problem between $\mathbb{P}_0, \mathbb{P}'$, i.e., $T_{f^\beta} : dX_t = f^\beta(X_t, t)dt$ and $H^\beta(X_0) = X_0 + \int_0^1 f^\beta(X_t, t)dt$. Furthermore, in this case, $\|X_0 - H^\beta(X_0)\|^2 = \int_0^1 \|f^\beta(X_t, t)\|^2 dt$. Hence, it can be derived that

$$\inf_H \widetilde{\mathcal{L}}(\beta, H) = \inf_{T_f} \mathcal{L}(\beta, T_f).$$

Step 4 (Equivalence of the saddle point objective). Take sup over $\beta \in \mathcal{C}_{b,2}(\mathcal{Y})$ and get the final equivalence:

$$\sup_\beta \inf_H \widetilde{\mathcal{L}}(\beta, H) = \sup_\beta \inf_H \mathcal{L}(\beta, T_f) = \mathcal{L}^*.$$

Step 5 (Dynamic OT solutions are contained in optimal saddle points). Pick any optimal $\beta^* \in \operatorname{argsup}_\beta \inf_H \mathcal{L}(\beta, T_{f^*})$ and let T_{f^*} be any solution to the Benamou-Brenier problem. Checking that $T^* \in \inf_H \mathcal{L}(\beta^*, T_f)$ can be done analogously to [31, Lemma 4], [42, Lemma 4.1]. $\qquad\square$

The derivation above shows the equivalence of objective values of dynamic unregularized OT (45) and our saddle point reformulation of BB (45). Additionally, it shows that solutions T_{f^*} can be recovered from *some* optimal saddle points (β^*, T_{f^*}) of our problem. At the same time, *unlike the EOT case* ($\epsilon > 0$), it is not guaranteed that for all the optimal saddle points (β^*, T_{f^*}) it holds that T_{f^*} is the solution to the BB problem. This aspect seems to be closely related to the *fake solutions* issue in the saddle point methods of OT [30] and may require further studies.