On the geometry of geodesics in discrete optimal transport

Matthias Erbar¹ · Jan Maas² · Melchior Wirth³

Received: 31 May 2018 / Accepted: 26 October 2018 / Published online: 11 December 2018
© The Author(s) 2018

Abstract

We consider the space of probability measures on a discrete set $\mathcal{X}$, endowed with a dynamical optimal transport metric. Given two probability measures supported in a subset $\mathcal{Y} \subseteq \mathcal{X}$, it is natural to ask whether they can be connected by a constant speed geodesic with support in $\mathcal{Y}$ at all times. Our main result answers this question affirmatively, under a suitable geometric condition on $\mathcal{Y}$ introduced in this paper. The proof relies on an extension result for subsolutions to discrete Hamilton–Jacobi equations, which is of independent interest.

Mathematics Subject Classification 49Q20 · 53C21

Contents

1 Introduction	2
2 The discrete transport distance	3
3 Duality for discrete optimal transport	5
4 Locality of optimal curves	10
5 Optimal transport avoids dead ends	15
6 Nonlocality of optimal transport on the triangle	17
References	18

Communicated by L. Ambrosio.

✉ Jan Maas
jan.maas@ist.ac.at

Matthias Erbar
erbar@iam.uni-bonn.de

Melchior Wirth
melchior.wirth@uni-jena.de

¹ Institut für angewandte Mathematik, Universität Bonn, Endenicher Allee 60, 53115 Bonn, Germany

² Institute of Science and Technology Austria (IST Austria), Am Campus 1, 3400 Klosterneuburg, Austria

³ Institute of Mathematics, Friedrich Schiller University Jena, 07737 Jena, Germany

1 Introduction

Optimal transport continues to be a very active field of research, both in mathematics and in applications. One of the central objects is the L^p -Kantorovich metric W_p , defined by

$$W_p(\mu, \nu) = \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{X}} d(x, y)^p d\pi(x, y) \right)^{1/p},$$

where μ and ν are Borel probability measures on the metric space $(\mathcal{X}, d)$, and $\Pi(\mu, \nu)$ is the set of all couplings of μ and ν .

The metric W_2 plays a special role in the theory, as it is the crucial object in the gradient flow formulation of dissipative PDE (starting from [11,20]) and in the synthetic theory of Ricci curvature [14,22], which builds on McCann's discovery that several important functionals enjoy convexity properties along W_2 -geodesics [16].

In spite of the robustness of the optimal transport theory, it is well known that if the underlying space is discrete, W_2 has several undesirable properties that hamper its usefulness. In particular, if $\mathcal{X}$ is discrete, the metric space $(\mathcal{P}(\mathcal{X}), W_2)$ does not contain any non-trivial geodesics.

To circumvent this problem, several authors introduced discrete dynamical transport metrics $\mathcal{W}$, based on discrete versions of the Benamou–Brenier formulation of optimal transport [2,15,17]. These metrics have been intensively studied in recent years; in particular, gradient flow formulations have been obtained for nonlinear evolution equations [6,19], and a discrete theory of Ricci curvature has been developed based on geodesic convexity of entropy functionals along discrete optimal transport [5,18]. Such Ricci curvature bounds have subsequently been obtained in various discrete probabilistic models [4,7,8].

In spite of the relevance of the notion of geodesic convexity, geometric properties of $\mathcal{W}$ -geodesics are currently poorly understood. The aim of this paper is to obtain results of this type. We focus on the issue of *locality* of geodesics in the space of probability measures.

More precisely, let $(\mathcal{X}, d)$ be a metric space, and consider a geodesic metric D on (a subset of) the space of Borel probability measures $\mathcal{P}(\mathcal{X})$. We say that a subset $\mathcal{Y} \subseteq \mathcal{X}$ has the *weak locality* property if any pair of probability measures $\mu_0, \mu_1 \in \mathcal{P}(\mathcal{X})$ supported in $\mathcal{Y}$ can be connected by a geodesic that is supported in $\mathcal{Y}$ at all times. The notion of *strong locality* is defined by requiring this property to hold for *any* geodesic connecting μ_0 and μ_1 . If any pair of measures can be connected by a unique geodesic, the notions of weak and strong locality coincide, but this property is currently unknown for discrete dynamical transport metrics.

If $(\mathcal{X}, d)$ is a geodesic metric space, and D is the Kantorovich metric W_p for some $1 \leq p < \infty$, it is well known that a subset $\mathcal{Y}$ has the weak (resp. strong) locality property if and only if it is weakly (resp. strongly) geodesically convex. This follows from the fact that geodesics in $(\mathcal{P}_p(\mathcal{X}), W_p)$ are supported on geodesics in $(\mathcal{X}, d)$; cf. [12] for a precise formulation of this result in a general setting.

Interestingly, the issue of locality in the discrete setting [with a discrete dynamical transport metric $\mathcal{W}$ on $\mathcal{P}(\mathcal{X})$ instead of W_p] turns out to be much more delicate. For example, if one considers the complete graph on a three-point set K_3 , then any geodesic connecting two Dirac masses transports a nontrivial part of the mass via the third point. Hence, two-point subsets of K_3 do not have the locality property. This is shown in Sect. 6 of this paper.

Based on this observation one may conjecture that any nontrivial $\mathcal{W}$ -geodesic has support on the whole graph. However, we show that this is not the case. In fact, the main contribution of this paper is the introduction of a geometric condition for subsets $\mathcal{Y} \subseteq \mathcal{X}$ (the *retraction*

property), that is shown to be sufficient for locality; see Theorem 4.11. The retraction property is easy to check in concrete examples, as is shown in Sect. 4.

As an application of our main result, we show that if $\mathcal{X}$ is any subset of the grid $\mathbb{Z}^d$ with the usual graph structure, and $\mathcal{Y} \subseteq \mathcal{X}$ is a hyperrectangle, then any pair of measures supported in $\mathcal{Y}$ can be connected by a geodesic supported in $\mathcal{Y}$. In particular, this property holds for measures supported on subsets of lines, or k -dimensional hyperplanes of dimension less than d . Let us also mention that discrete Ricci curvature bounds in the sense of [5, 18] are inherited by subsets with the retraction property; see Corollary 4.12.

A key ingredient in the proof of our main result is a duality result for the discrete transport metric $\mathcal{W}$, which was recently obtained by Gangbo, Li, and Mou (under slightly more restrictive conditions on the transition rates) [9]. We interpret this result (Theorem 3.4 below) in terms of subsolutions of a discrete Hamilton–Jacobi equation and present a different proof based on Fenchel–Rockafellar duality. We then show that subsolutions of the Hamilton–Jacobi equation on a subset $\mathcal{Y} \subseteq \mathcal{X}$ can be extended to the full space $\mathcal{X}$, provided that $\mathcal{Y}$ has the retraction property; cf. Theorem 4.10. Our main theorem is then a straightforward consequence of this result.

Structure of the paper In Sect. 2 we collect the necessary preliminaries on discrete transport metrics. Section 3 contains the dual formulation of the transport problem in terms of Hamilton–Jacobi subsolutions. In Sect. 4 we introduce the retraction property, we show the extension result for subsolutions to the Hamilton–Jacobi equation (Theorem 4.10), and we prove the main result on weak locality of subsets with the retraction property (Theorem 4.11). In Sect. 5 we show that the strong locality property holds for Markov chains with “dead ends”. Finally, it is shown in Sect. 6 that geodesics between Dirac measures on the triangle have full support.

2 The discrete transport distance

In this section we briefly recall the definition and basic properties of the discrete transport distance constructed in [2, 15, 17].

Let $\mathcal{X}$ be a finite set, and let $Q : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}_+$ denote the transition rates for a Markov chain on $\mathcal{X}$. Without loss of generality, we use the convention that $Q(x, x) = 0$ for all $x \in \mathcal{X}$. The corresponding generator $\mathcal{L}$ acts on functions $\phi : \mathcal{X} \rightarrow \mathbb{R}$ by

$$\mathcal{L}\phi(x) = \sum_{y \in \mathcal{X}} Q(x, y)(\phi(y) - \phi(x)) .$$

We assume that Q is *irreducible*, i.e., each pair $(x, y) \in \mathcal{X} \times \mathcal{X}$ can be connected, for some $n \in \mathbb{N}$, by a path $\{x_i\}_{i=0}^n$ satisfying $x_0 = x$, $x_n = y$, and $Q(x_{i-1}, x_i) > 0$ for $i = 1, \dots, n$. This assumption implies the existence of a unique stationary probability measure π on $\mathcal{X}$. Moreover, π is strictly positive. We will furthermore assume that Q is *reversible* with respect to π , i.e., the *detailed balance condition* holds:

$$\pi(x)Q(x, y) = \pi(y)Q(y, x) \quad \text{for all } x, y \in \mathcal{X} .$$

The triple $(\mathcal{X}, Q, \pi)$ will be referred to as a *Markov triple*.

A Markov chain induces a graph on the vertex set $\mathcal{X}$, whose edge set is given by $\mathcal{E} = \{(x, y) \in \mathcal{X} \times \mathcal{X} : Q(x, y) > 0\}$. We write $x \sim y$ iff $Q(x, y) > 0$. The assumption that Q is

irreducible corresponds to the graph $(\mathcal{X}, \mathcal{E})$ being connected. The detailed balance condition implies that the graph is undirected.

In order to define the discrete transport distance on the set $\mathcal{P}(\mathcal{X})$ of probability measures on $\mathcal{X}$, we introduce the following objects.

Definition 2.1 (*Continuity equation*) A pair (μ, V) is said to satisfy the *continuity equation* if

- (i) $\mu : [0, T] \rightarrow \mathbb{R}^{\mathcal{X}}$ is continuous;
- (ii) $V : [0, T] \rightarrow \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ is locally integrable;
- (iii) $\mu_t \in \mathcal{P}(\mathcal{X})$ for all $t \in [0, T]$;
- (iv) the continuity equation holds in the sense of distributions:

$$\frac{d}{dt} \mu_t(x) + \frac{1}{2} \sum_{y \in \mathcal{X}} (V_t(x, y) - V_t(y, x)) = 0 \quad \text{for all } x \in \mathcal{X}. \quad (2.1)$$

In this case, we write $(\mu, V) \in \text{CE}_T$. Furthermore, $\text{CE}_T(\mu^0, \mu^1)$ denotes the collection of pairs $(\mu, V) \in \text{CE}_T$ satisfying $\mu|_{t=0} = \mu^0$ and $\mu|_{t=T} = \mu^1$.

Definition 2.2 (*Admissible mean*) An *admissible mean* is a continuous function $\Lambda : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ that is C^∞ on $(0, \infty) \times (0, \infty)$, symmetric, positively 1-homogeneous, non-decreasing in each of its variables, jointly concave, and normalised, i.e., $\Lambda(1, 1) = 1$.

Of particular interest to us is the *logarithmic mean* given by

$$\Lambda_{\log}(s, t) := \int_0^1 s^\alpha t^{1-\alpha} d\alpha,$$

since it arises in the entropic gradient flow structure for the master equation $\partial_t \mu = \mathcal{L}^* \mu$. Other relevant examples of admissible means are the harmonic mean $\Lambda_{\text{har}}(s, t) = \frac{2st}{s+t}$, the geometric mean $\Lambda_{\text{geo}}(s, t) = \sqrt{st}$, and the arithmetic mean $\Lambda_{\text{ari}}(s, t) = \frac{s+t}{2}$. Some of these means arise in gradient structures for porous medium equations; cf. [6]. From now on, we will fix an admissible mean Λ .

The action functional for the discrete transport distance is defined using the convex and lower semicontinuous function $A : \mathbb{R}^3 \rightarrow [0, \infty]$ given by

$$A(s, t, w) := \begin{cases} \frac{w^2}{\Lambda(s, t)}, & w \in \mathbb{R}, s, t > 0, \\ 0, & w = 0, s, t \geq 0, \\ +\infty, & \text{otherwise.} \end{cases} \quad (2.2)$$

For $\mu \in \mathcal{P}(\mathcal{X})$ and $V : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ we define the action by

$$\mathcal{A}(\mu, V) = \frac{1}{2} \sum_{x, y \in \mathcal{X}} A(\mu(x)Q(x, y), \mu(y)Q(y, x), V(x, y)).$$

For brevity we sometimes write

$$\widehat{\mu}(x, y) := \Lambda(\mu(x)Q(x, y), \mu(y)Q(y, x)).$$

Definition 2.3 (*Discrete transport distance*) For a Markov triple $(\mathcal{X}, Q, \pi)$ and an admissible mean Λ , the *discrete transport distance* $\mathcal{W}$ is defined for $\mu_0, \mu_1 \in \mathcal{P}(\mathcal{X})$ by

$$\mathcal{W}(\mu_0, \mu_1) := \inf \left\{ \sqrt{\int_0^1 \mathcal{A}(\mu_t, V_t) dt} : (\mu, V) \in \text{CE}_1(\mu_0, \mu_1) \right\}. \quad (2.3)$$

It has been shown in [5] that minimisers exist in the minimisation problem above. Any minimal curve $(\mu_t)_{t \in [0,1]}$ is a constant speed geodesic, i.e., it satisfies $\mathcal{W}(\mu_s, \mu_t) = |t - s| \mathcal{W}(\mu_0, \mu_1)$ for all $s, t \in [0, 1]$.

Remark 2.4 Without loss of generality we may assume in the minimisation (2.3) that V is anti-symmetric, i.e., $V_t(x, y) = -V_t(y, x)$. In fact, for each $U \in \mathbb{R}$, the quantity $|V(x, y)|^2 + |V(y, x)|^2$ is minimised among all choices of $V(x, y), V(y, x)$ such that $V(x, y) - V(y, x) = U$ by choosing $V(y, x) = -V(x, y) = U/2$.

Finally, let us introduce some convenient notation to be used in the sequel. We denote the Euclidean inner products on $\mathbb{R}^{\mathcal{X}}$ and $\mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ by

$$\langle \phi, \psi \rangle = \sum_{x \in \mathcal{X}} \phi(x) \psi(x) \quad \text{and} \quad \langle \Phi, \Psi \rangle = \frac{1}{2} \sum_{x, y \in \mathcal{X}} \Phi(x, y) \Psi(x, y).$$

The discrete gradient of a function $\phi \in \mathbb{R}^{\mathcal{X}}$ will be denoted by $\nabla \phi(x, y) = \phi(y) - \phi(x)$, and the discrete divergence of $\Phi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ is given by

$$\nabla \cdot \Phi(x) = \frac{1}{2} \sum_{y \in \mathcal{X}} (\Phi(x, y) - \Phi(y, x)).$$

Furthermore, for $\mu \in \mathcal{P}(\mathcal{X})$ and $\Phi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ we write

$$\|\Phi\|_{\mu} := \sqrt{\langle \Phi, \Phi \cdot \hat{\mu} \rangle}.$$

where the multiplication of $\Phi \cdot \hat{\mu}$ is understood componentwise. For all $\Phi, V \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ and $\mu \in \mathcal{P}(\mathcal{X})$, Young's inequality yields

$$\langle \Phi, V \rangle \leq \frac{1}{2} \|\Phi\|_{\mu}^2 + \frac{1}{2} \mathcal{A}(\mu, V). \quad (2.4)$$

3 Duality for discrete optimal transport

We present a dual formulation for the discrete transport distance which can be seen as a discrete analogue of the Kantorovich duality. This result has recently been proved in [9] using different methods; cf. Proposition 3.10 and Theorems 5.10 and 7.4 in that paper. Note that the result in [9] is stated under slightly stronger assumptions on the transition rates. In our notation, it is assumed there that $Q(x, y) = Q(y, x)$ and π is constant. The slightly greater generality here does not cause additional difficulties.

Definition 3.1 (*Hamilton–Jacobi subsolution*) A function $\phi \in H^1((0, T); \mathbb{R}^{\mathcal{X}})$ is said to be a *Hamilton–Jacobi subsolution* if for a.e. t in $(0, T)$, we have

$$\langle \dot{\phi}_t, \mu \rangle + \frac{1}{2} \|\nabla \phi_t\|_{\mu}^2 \leq 0 \quad \text{for all } \mu \in \mathcal{P}(\mathcal{X}). \quad (3.1)$$

The collection of all Hamilton–Jacobi subsolutions is denoted $\text{HJ}_{\mathcal{X}}^T$.

Remark 3.2 Hamilton–Jacobi subsolutions obey a simple scaling relation: given $\phi \in \text{HJ}_{\mathcal{X}}^T$ and $\lambda > 0$, set $\phi_t^{\lambda} := \lambda \phi_{\lambda t}$. It is immediate to check that $\phi^{\lambda} \in \text{HJ}_{\mathcal{X}}^{\lambda T}$.

Remark 3.3 Informally, (3.1) may be seen as a one-sided discrete version of the Hamilton–Jacobi equation $\partial_t \phi + \frac{1}{2} |\nabla \phi|^2 = 0$. Note however that the dependence on μ in (3.1) is nonlinear, which prevents us from formulating the inequality pointwise in terms of ϕ only. This is a crucial difference between the discrete and the continuous setting, and a source of several difficulties.

In the continuous setting, a full treatment of Hamilton–Jacobi equations relies on the theory of viscosity solutions [3], but this concept will not play any role in our discrete setting. Let us also mention that Hamilton–Jacobi equations have been studied in the setting of metric length spaces [10, 13] as well as on graphs [21]. Our discrete notion of Hamilton–Jacobi subsolution is different from the one studied in [21].

Theorem 3.4 (Duality formula) *For $\mu_0, \mu_1 \in \mathcal{P}(\mathcal{X})$ we have*

$$\frac{1}{2} \mathcal{W}^2(\mu_0, \mu_1) = \sup \{ \langle \phi_1, \mu_1 \rangle - \langle \phi_0, \mu_0 \rangle : \phi \in HJ_{\mathcal{X}}^1 \}. \quad (3.2)$$

This representation remains true if the supremum is restricted to the class of functions $\phi \in C^1([0, 1], \mathbb{R}^{\mathcal{X}})$ satisfying (3.1).

Let us first give a heuristic argument for the duality result above. We start by introducing a Lagrange multiplier for the continuity equation constraint and write

$$\frac{1}{2} \mathcal{W}^2(\mu_0, \mu_1) = \inf_{\mu, V} \sup_{\phi} \left\{ \int_0^1 \frac{1}{2} \mathcal{A}(\mu_t, V_t) dt + \int_0^1 \langle \phi_t, \dot{\mu}_t + \nabla \cdot V_t \rangle dt \right\}, \quad (3.3)$$

where the supremum is taken over all (sufficiently smooth) functions $\phi : [0, 1] \rightarrow \mathbb{R}^{\mathcal{X}}$ and the infimum is taken over all (sufficiently smooth) curves $\mu : [0, 1] \rightarrow \mathbb{R}_+$ connecting μ_0 and μ_1 , and over all $V : [0, 1] \rightarrow \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$. Here we do not require that (μ, V) satisfies the continuity equation, but the inner supremum takes the value $+\infty$ if (μ, V) does not belong to $\text{CE}_1(\mu_0, \mu_1)$. We also do not require that μ takes values in $\mathcal{P}(\mathcal{X})$, but this is automatically enforced by the continuity equation.

Integrating by parts and using the min–max principle we obtain

$$\begin{aligned} \frac{1}{2} \mathcal{W}^2(\mu_0, \mu_1) &= \inf_{\mu, V} \sup_{\phi} \left\{ \langle \phi_1, \mu_1 \rangle - \langle \phi_0, \mu_0 \rangle + \int_0^1 \frac{1}{2} \mathcal{A}(\mu_t, V_t) - \langle \dot{\phi}_t, \mu_t \rangle - \langle \nabla \phi_t, V_t \rangle dt \right\} \\ &= \sup_{\phi} \left\{ \langle \phi_1, \mu_1 \rangle - \langle \phi_0, \mu_0 \rangle + \inf_{\mu, V} \int_0^1 \frac{1}{2} \mathcal{A}(\mu_t, V_t) - \langle \dot{\phi}_t, \mu_t \rangle - \langle \nabla \phi_t, V_t \rangle dt \right\}. \end{aligned}$$

As the quantity to be minimised is positively 1-homogeneous in (μ, V) , the infimum takes the value $-\infty$ if ϕ does not belong to $\mathcal{H}$, the set of C^1 functions $\phi : [0, 1] \rightarrow \mathbb{R}^{\mathcal{X}}$ satisfying

$$\int_0^1 \frac{1}{2} \mathcal{A}(\mu_t, V_t) - \langle \dot{\phi}_t, \mu_t \rangle - \langle \nabla \phi_t, V_t \rangle dt \geq 0$$

for all $\mu : [0, 1] \rightarrow \mathbb{R}_+^{\mathcal{X}}$ and all $V : [0, 1] \rightarrow \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$. Consequently,

$$\frac{1}{2} \mathcal{W}(\mu_0, \mu_1)^2 = \sup \{ \langle \phi_1, \mu_1 \rangle - \langle \phi_0, \mu_0 \rangle : \phi \in \mathcal{H} \}.$$

A simple localisation argument in t shows that $\phi \in \mathcal{H}$ iff for all $t \in [0, 1]$ and $(\mu, V) \in \mathbb{R}_+^{\mathcal{X}} \times \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$:

$$\frac{1}{2} \mathcal{A}(\mu, V) - \langle \dot{\phi}_t, \mu \rangle - \langle \nabla \phi_t, V \rangle \geq 0.$$

We may write

$$\frac{1}{2} \mathcal{A}(\mu, V) - \langle \nabla \phi_t, V \rangle = \frac{1}{4} \sum_{x,y} \frac{(V(x, y) - \hat{\mu}(x, y) \nabla \phi_t(x, y))^2}{\hat{\mu}(x, y)} - \frac{1}{2} \|\nabla \phi_t\|_\mu^2.$$

Minimising over V we conclude that $\phi \in \mathcal{H}$ iff the inequality

$$\langle \dot{\phi}_t, \mu \rangle + \frac{1}{2} \|\nabla \phi_t\|_\mu^2 \leq 0$$

holds for all $\mu \in \mathbb{R}_+^{\mathcal{X}}$ and $t \in [0, 1]$, which means that $\phi \in \text{HJ}_{\mathcal{X}}^1$.

We present a proof of Theorem 3.4 using the Fenchel–Rockafellar duality theorem; see, e.g., [23, Theorem 1.9]. Recall that given a normed vector space E with topological dual space E^* and a proper convex function $F : E \rightarrow \mathbb{R} \cup \{+\infty\}$, its Legendre–Fenchel transform $F^* : E^* \rightarrow \mathbb{R} \cup \{+\infty\}$ is defined by

$$F^* : E^* \rightarrow \mathbb{R} \cup \{+\infty\}, \quad F^*(x^*) := \sup_{x \in E} \{ \langle x, x^* \rangle - F(x) \}.$$

Theorem 3.5 (Fenchel–Rockafellar duality) *Let E be a normed vector space and E^* its topological dual. Let $F, G : E \rightarrow \mathbb{R} \cup \{+\infty\}$ be proper convex functions and denote by $F^*, G^* : E^* \rightarrow \mathbb{R} \cup \{+\infty\}$ their Legendre–Fenchel transforms. Assume that there is $z_0 \in E$ such that G is continuous at z_0 and $F(z_0), G(z_0) < \infty$. Then we have:*

$$\sup_{z \in E} [-F(z) - G(z)] = \min_{z^* \in E^*} [F^*(z^*) + G^*(-z^*)]. \quad (3.4)$$

Proof of Theorem 3.4 Let us first note that, by the convexity of the constraint (3.1), any $\phi \in \text{HJ}_{\mathcal{X}}^1$ can be approximated uniformly by C^1 functions satisfying (3.1) by convolution (after scaling the function to a slightly larger interval $[-\delta, 1+\delta]$ via Remark 3.2). Therefore, the final part of the theorem follows.

To show the dual representation with C^1 functions, we will apply Theorem 3.5 in the following situation. Let E be the Banach space

$$E = C^1([0, 1], \mathbb{R}^{\mathcal{X}}) \times L^2((0, 1), \mathbb{R}^{\mathcal{X} \times \mathcal{X}}).$$

Since we can identify $C^1([0, 1], \mathbb{R}^{\mathcal{X}})$ with $\mathbb{R}^{\mathcal{X}} \times C^0([0, 1], \mathbb{R}^{\mathcal{X}})$ via the map $I : \phi \mapsto (\phi_0, \dot{\phi})$, the dual space E^* can be identified with

$$E^* = \mathbb{R}^{\mathcal{X}} \times M([0, 1], \mathbb{R}^{\mathcal{X}}) \times L^2((0, 1), \mathbb{R}^{\mathcal{X} \times \mathcal{X}}),$$

where the duality pairing between $(\phi, \Phi) \in E$ and $(b, \sigma, V) \in E^*$ is given by

$$\langle (\phi, \Phi), (b, \sigma, V) \rangle = \langle \phi_0, b \rangle + \int_0^1 \langle \dot{\phi}_t, d\sigma(t) \rangle + \int_0^1 \langle \Phi_t, V_t \rangle dt,$$

keeping in mind that σ is a vector-valued measure.

Define the functionals $F, G : E \rightarrow \mathbb{R} \cup \{+\infty\}$ by

$$F(\phi, \Phi) = \begin{cases} -\langle \phi_1, \mu_1 \rangle + \langle \phi_0, \mu_0 \rangle, & \Phi = \nabla \phi, \\ +\infty, & \text{otherwise,} \end{cases}$$

$$G(\phi, \Phi) = \begin{cases} 0, & (\phi, \Phi) \in \mathcal{D}, \\ +\infty, & \text{otherwise.} \end{cases}$$

Here we say that a pair $(\phi, \Phi) \in E$ belongs to $\mathcal{D}$ if for all continuous curves $t \mapsto \eta_t \in \mathbb{R}_+^{\mathcal{X}}$ we have $\int_0^1 \langle \dot{\phi}_t, \eta_t \rangle + \frac{1}{2} \|\Phi_t\|_{\eta_t}^2 dt \leq 0$. It is readily checked that F and G are convex. Moreover, setting $\bar{\phi}(t) = t(-1, \dots, -1)$ and $\bar{\Phi} \equiv 0$, both F and G are finite at $(\bar{\phi}, \bar{\Phi})$ and G is continuous at $(\bar{\phi}, \bar{\Phi})$. Note that for $\phi \in C^1([0, 1], \mathbb{R}^{\mathcal{X}})$ we have $(\phi, \nabla \phi) \in \mathcal{D}$ if and only if $\phi \in \text{HJ}_{\mathcal{X}}^1$, which follows from a simple localisation argument in t . Hence, the supremum in the left-hand side of (3.4) coincides with the supremum in the right-hand side of (3.2).

We will calculate the Legendre–Fenchel transforms of F and G . For F we obtain

$$\begin{aligned} F^*(b, \sigma, V) &= \sup_{(\phi, \Phi) \in E} \{ \langle (\phi, \Phi), (b, \sigma, V) \rangle - F(\phi, \Phi) \} \\ &= \sup_{\phi} \left\{ \langle \phi_0, b \rangle + \int_0^1 \langle \dot{\phi}_t, d\sigma(t) \rangle + \int_0^1 \langle \nabla \phi_t, V_t \rangle dt + \langle \phi_1, \mu_1 \rangle - \langle \phi_0, \mu_0 \rangle \right\}. \end{aligned}$$

Thus, by homogeneity of the last expression in ϕ , one has $F^*(b, \sigma, v) = +\infty$ unless (σ, V) satisfies the continuity equation $\partial_t \sigma + \nabla \cdot V = 0$ with boundary values $-(\mu_0 - b)$ and $-\mu_1$, in the sense that

$$\langle \phi_1, -\mu_1 \rangle - \langle \phi_0, -(\mu_0 - b) \rangle = \int_0^1 \langle \dot{\phi}_t, d\sigma(t) \rangle + \int_0^1 \langle \nabla \phi_t, V_t \rangle dt \quad (3.5)$$

for all $\phi \in C^1([0, 1], \mathbb{R}^{\mathcal{X}})$. In particular, the distributional derivative of σ belongs to $L^2([0, 1]; \mathbb{R}^{\mathcal{X}})$. Since the antiderivative of a distribution is unique up to a constant, the fundamental theorem of Lebesgue calculus implies that σ has the form $d\sigma(t) = \sigma_t dt$ for some curve $(\sigma_t)_t \in H^1([0, 1]; \mathbb{R}^{\mathcal{X}})$. Moreover, (3.5) implies $\sigma_0 = -(\mu_0 - b)$ and $\sigma_1 = -\mu_1$. Thus, we obtain

$$F^*(b, \sigma, V) = \begin{cases} 0, & (-\sigma, -V) \in \text{CE}'(\mu_0 - b, \mu_1), \\ +\infty, & \text{otherwise,} \end{cases} \quad (3.6)$$

where CE' is defined by dropping the positivity and normalisation condition (i) in the definition of CE , and we have identified the measure σ with the H^1 -map σ_t .

As it suffices to calculate the transform of G at points (b, σ, V) where $F^*(-b, -\sigma, -V)$ is finite, we can assume that $d\sigma(t) = \sigma_t dt$ with $(\sigma_t)_t \in H^1([0, 1]; \mathbb{R}^{\mathcal{X}})$. We claim that:

$$G^*(b, \sigma, V) = \begin{cases} \frac{1}{2} \int_0^1 \mathcal{A}(\sigma_t, V_t) dt, & b = 0, \\ +\infty, & \text{otherwise.} \end{cases} \quad (3.7)$$

Indeed, it follows that

$$\begin{aligned} G^*(b, \sigma, V) &= \sup_{(\phi, \Phi) \in E} \{ \langle (\phi, \Phi), (b, \sigma, V) \rangle - G(\phi, \Phi) \} \\ &= \sup_{(\phi, \Phi) \in \mathcal{D}} \left\{ \langle \phi_0, b \rangle + \int_0^1 \langle \dot{\phi}_t, \sigma_t \rangle + \langle \Phi_t, V_t \rangle dt \right\}. \end{aligned}$$

Since $(\phi, \Phi) \in \mathcal{D}$ implies $(\phi + c, \Phi) \in \mathcal{D}$ for all $c \in \mathbb{R}^{\mathcal{X}}$, we have $G^*(b, \sigma, V) = +\infty$ unless $b = 0$. Moreover, from the definition of $\mathcal{D}$ we infer that $G^*(b, \sigma, V) = +\infty$ unless $\sigma_t \in \mathbb{R}_+^{\mathcal{X}}$ for a.e. t .

Let us assume that $b = 0$ and $\int_0^1 \mathcal{A}(\sigma_t, V_t) dt < \infty$. Then we obtain

$$\begin{aligned} G^*(0, \sigma, V) &= \sup_{(\phi, \Phi) \in \mathcal{D}} \left\{ \int_0^1 \langle \dot{\phi}, \sigma_t \rangle + \langle \Phi_t, V_t \rangle dt \right\} \\ &\leq \sup_{(\phi, \Phi) \in \mathcal{D}} \left\{ \int_0^1 -\frac{1}{2} \|\Phi_t\|_{\sigma_t}^2 + \langle \Phi_t, V_t \rangle dt \right\} \leq \frac{1}{2} \int_0^1 \mathcal{A}(\sigma_t, V_t) dt, \end{aligned} \quad (3.8)$$

where the first inequality follows from the definition of $\mathcal{D}$ and the second from (2.4).

It remains to show that we have in fact equality. First we consider a convolution in time yielding smooth pairs $\sigma_t^\varepsilon, V_t^\varepsilon$ converging to σ_t, V_t as $\varepsilon \rightarrow 0$. Then we set for $\delta > 0$, $\sigma_t^{\delta, \varepsilon} = \sigma_t^\varepsilon + \delta\pi$. By convexity of the action and monotonicity of the mean Λ we have

$$\int_0^1 \mathcal{A}(\sigma_t^{\delta, \varepsilon}, V_t^\varepsilon) dt \leq \int_0^1 \mathcal{A}(\sigma_t^\varepsilon, V_t^\varepsilon) dt \leq \int_0^1 \mathcal{A}(\sigma_t, V_t) dt. \quad (3.9)$$

The convexity and lower semicontinuity of A further implies the lower semicontinuity of the action; see [1, Theorem 3.4.3] for a general result on lower semicontinuity of integral functionals and the proof of [5, Theorem 3.2] for the application to the action functional $\mathcal{A}$. Consequently,

$$\int_0^1 \mathcal{A}(\sigma_t, V_t) dt = \lim_{\delta \rightarrow 0} \lim_{\varepsilon \rightarrow 0} \int_0^1 \mathcal{A}(\sigma_t^{\delta, \varepsilon}, V_t^\varepsilon) dt.$$

Now, we can choose in particular $(\phi^{\delta, \varepsilon}, \Phi^{\delta, \varepsilon})$ such that

$$\Phi_t^{\delta, \varepsilon} = \frac{V_t^\varepsilon}{\dot{\sigma}_t^{\delta, \varepsilon}}, \quad \dot{\phi}^{\delta, \varepsilon}(x) = -\frac{1}{2} \sum_y \partial_1 \Lambda(\rho_t^{\delta, \varepsilon}(x), \rho_t^{\delta, \varepsilon}(y)) |\Phi_t^{\delta, \varepsilon}(x, y)|^2 Q(x, y),$$

where $\sigma^{\delta, \varepsilon} = \rho^{\delta, \varepsilon} \pi$.

We claim that $(\phi^{\delta, \varepsilon}, \Phi^{\delta, \varepsilon}) \in \mathcal{D}$. To see this, we use the inequality

$$\partial_1 \Lambda(s, t)u + \partial_2 \Lambda(s, t)v \geq \Lambda(u, v) \quad \forall s, t > 0, u, v \geq 0,$$

which is an identity for $s = v, t = u$, see [5, Lemma 2.2]. From this we infer that for any $\mu = \tilde{\rho}\pi \in \mathcal{P}(\mathcal{X})$ we have

$$\begin{aligned} \langle \dot{\phi}_t^{\delta, \varepsilon}, \mu \rangle &= -\frac{1}{2} \sum_{x, y} \partial_1 \Lambda(\rho_t^{\delta, \varepsilon}(x), \rho_t^{\delta, \varepsilon}(y)) \tilde{\rho}(x) |\Phi_t^{\delta, \varepsilon}(x, y)|^2 Q(x, y) \pi(x) \\ &= -\frac{1}{4} \sum_{x, y} \left[\partial_1 \Lambda(\rho^{\delta, \varepsilon}(x), \rho^{\delta, \varepsilon}(y)) \tilde{\rho}(x) + \partial_2 \Lambda(\rho^{\delta, \varepsilon}(x), \rho^{\delta, \varepsilon}(y)) \tilde{\rho}(y) \right] \\ &\quad \times |\Phi_t^{\delta, \varepsilon}(x, y)|^2 Q(x, y) \pi(x) \\ &\leq -\frac{1}{4} \sum_{x, y} \Lambda(\tilde{\rho}(x), \tilde{\rho}(y)) |\Phi_t^{\delta, \varepsilon}(x, y)|^2 Q(x, y) \pi(x) = -\frac{1}{2} \|\Phi_t^{\delta, \varepsilon}\|_\mu^2, \end{aligned} \quad (3.10)$$

which proves the claim. Note that for $\tilde{\rho} = \rho_t^{\delta, \varepsilon}$ we obtain equality.

Next we claim that

$$\lim_{\delta \rightarrow 0} \lim_{\varepsilon \rightarrow 0} \int_0^1 \langle \dot{\phi}^{\delta, \varepsilon}, \sigma_t \rangle dt = -\frac{1}{2} \int_0^1 \mathcal{A}(\sigma_t, V_t) dt.$$

To prove this, we compare the left-hand side and the second line in (3.10). The limit $\varepsilon \rightarrow 0$ is justified by dominated convergence, since (3.10) yields the majorant

$$\frac{1}{2} \|\Phi_t^{\delta, \varepsilon}\|_{\sigma_t}^2 \leq \frac{C}{\delta} \mathcal{A}(\sigma_t^\varepsilon, V_t^\varepsilon),$$

where C depends on Q and π . The right-hand side converges as $\varepsilon \rightarrow 0$ by (3.9). The limit $\delta \rightarrow 0$ is justified by monotone convergence. Similarly, we have

$$\lim_{\delta \rightarrow 0} \lim_{\varepsilon \rightarrow 0} \int_0^1 \langle V_t, \Phi_t^{\varepsilon, \delta} \rangle dt = \int_0^1 \mathcal{A}(\sigma_t, V_t) dt.$$

Here, we can use the estimate $|ab| \leq \frac{1}{2}a^2 + \frac{1}{2}b^2$ to obtain a majorant that converges by (3.9) as before. Thus the expression in the first braced bracket of (3.8) converges to the right-hand side of (3.8) with this choice of $(\phi^{\delta, \varepsilon}, \Psi^{\delta, \varepsilon})$ as $\delta, \varepsilon \rightarrow 0$. A similar argument yields $G^*(0, \sigma, V) = \infty$ if $\int_0^1 \mathcal{A}(\sigma_t, V_t) dt = \infty$. Combining (3.6), (3.7) and the fact that $\mathcal{A}(\sigma, V) = +\infty$ unless $\sigma \in \mathbb{R}_+^{\mathcal{X}}$, we obtain

$$F^*(-b, -\sigma, -V) + G^*(b, \sigma, V) = \begin{cases} \frac{1}{2} \int_0^1 \mathcal{A}(\sigma_t, V_t) dt, & (\sigma, V) \in \text{CE}(\mu_0, \mu_1), b = 0, \\ +\infty, & \text{otherwise.} \end{cases}$$

Thus the infimum in the right-hand side of (3.4) coincides with $\frac{1}{2} \mathcal{W}(\mu_0, \mu_1)^2$. An application of Theorem 3.5 concludes the proof. $\square$

4 Locality of optimal curves

In this section we investigate locality properties for discrete transport geodesics. More precisely, we study the following question: Given two probability measures supported in a subset $\mathcal{Y}$ of a state space $\mathcal{X}$, is there an optimal curve connecting them that is supported in $\mathcal{Y}$? The crucial tool to analyse this question is the dual characterisation of the transport problem given in the previous section. We prove two types of results.

Firstly, we show that the question can be answered affirmatively, under a simple condition (the *retraction property* of the subgraph $\mathcal{Y}$), which will be introduced below. This property ensures that any competitor in the dual problem on the subgraph can be extended to a competitor on the full graph. We present several examples where this property is satisfied. Later, in Sect. 6, we will show that locality may fail if the retraction property is not satisfied.

We start by introducing the retraction property and we give several examples. To increase readability, we often write subscripts instead of parentheses, e.g., $Q_{xy} = Q(x, y)$.

A subset $\mathcal{Y} \subseteq \mathcal{X}$ is said to be *connected* if any two distinct points $y, y' \in \mathcal{Y}$ can be connected by a path $\{y_i\}_{i=0}^n \subseteq \mathcal{Y}$ satisfying $y_0 = y$, $y_n = y'$, and $Q(y_{i-1}, y_i) > 0$ for $i = 1, \dots, n$.

Definition 4.1 (*Retraction property*) A connected subset $\mathcal{Y} \subseteq \mathcal{X}$ has the *retraction property* if there exists a map $T: \mathcal{X} \rightarrow \mathcal{Y}$ such that

- (R1) $T(y) = y$ for all $y \in \mathcal{Y}$;
- (R2) For all $y, y' \in \mathcal{Y}$ with $y \neq y'$, and all $x \in T^{-1}(y)$, we have

$$\sum_{x' \in T^{-1}(y')} Q(x, x') \leq Q(y, y').$$

Fig. 1 Retraction of a 9-cycle $\mathcal{X}$ onto a 4-point set $\mathcal{Y}$. The labels indicate the image of the corresponding vertex under the retraction T

The map T is called a *retraction* of $\mathcal{X}$ onto $\mathcal{Y}$.

Remark 4.2 If the Markov triple $(\mathcal{X}, Q, \pi)$ corresponds to a simple random walk (i.e., $Q(x, y) \in \{0, 1\}$ for all $x, y \in \mathcal{X}$), the retraction property can be rephrased in graph theoretical terms. Indeed, it is readily verified that the retraction property holds if and only if there exists a map $T : \mathcal{X} \rightarrow \mathcal{Y}$ with the following properties:

- (R1') $T(y) = y$ for all $y \in \mathcal{Y}$;
- (R2') If $x \sim x'$, then $T(x) = T(x')$ or $T(x) \sim T(x')$;
- (R3') If $x'_1 \sim x$, $x'_2 \sim x$, and $T(x'_1) = T(x'_2)$ for some $x'_1 \neq x'_2$, then $T(x) = T(x'_1)$.

Definition 4.3 (*Restriction*) The *restriction* of a Markov triple $(\mathcal{X}, Q, \pi)$ to a connected subset $\mathcal{Y} \subseteq \mathcal{X}$ is the Markov triple $(\mathcal{Y}, Q|_{\mathcal{Y}}, \pi|_{\mathcal{Y}})$, where $Q|_{\mathcal{Y}}$ is the restriction of Q to $\mathcal{Y} \times \mathcal{Y}$, and $\pi|_{\mathcal{Y}}$ is the normalised restriction of π to $\mathcal{Y}$.

Connectedness of $\mathcal{Y}$ implies that the Markov triple $(\mathcal{Y}, Q|_{\mathcal{Y}}, \pi|_{\mathcal{Y}})$ is irreducible, and the detailed balance relation is obviously inherited. The following result implies that if $\mathcal{Y}$ has the retraction property as a subset of $\mathcal{X}$, it also has this property as a subset of any set $\mathcal{X}'$ with $\mathcal{Y} \subseteq \mathcal{X}' \subseteq \mathcal{X}$.

Lemma 4.4 Let $(\mathcal{X}, Q, \pi)$ be a Markov triple and $\mathcal{Y} \subseteq \mathcal{X}' \subseteq \mathcal{X}$. If $T : \mathcal{X} \rightarrow \mathcal{Y}$ is a retraction, then its restriction $T|_{\mathcal{X}'} : \mathcal{X}' \rightarrow \mathcal{Y}$ is a retraction as well.

Proof This follows immediately from the definition. $\square$

We present some examples of sets with the retraction property.

Example 4.5 (*Cycle*) For $n \geq 2$, let $\mathcal{X} = \mathbb{Z}/n\mathbb{Z}$, and set $Q_{j,j+1} = Q_{j+1,j} = 1$ and $Q_{ij} = 0$ otherwise. All computations are to be understood modulo n . We claim that the subset $\{1, \dots, k\}$ of $\mathcal{X}$ has the retraction property if and only if $2k \leq n$. In this case, a retraction is given as follows (cf. Fig. 1):

$$T : \mathcal{X} \rightarrow \{1, \dots, k\}, \quad T(j) = \begin{cases} j & \text{if } 1 \leq j \leq k, \\ 2k - j + 1 & \text{if } k + 1 \leq j \leq 2k, \\ 1 & \text{if } 2k + 1 \leq j \leq n. \end{cases}$$

Indeed, to check sufficiency, note that $(R1')$ is trivial, $(R2')$ holds since $n \geq 2k$, and $(R3')$ is readily checked as well. Necessity follows from a simple argument.

Fig. 2 An x - y -cut associated with a retraction

Example 4.6 (Grid) Consider $\mathbb{Z}^d$ with the usual graph structure given by $Q_{xy} = 1$ if $|x - y| = 1$ and $Q_{xy} = 0$ otherwise. Let $\mathcal{Y} \subseteq \mathbb{Z}^d$ be a nonempty subset of the form $\mathcal{Y} = \mathcal{R} \cap \mathbb{Z}^d$, where $\mathcal{R} = \prod_{j=1}^d [a_j, b_j]$ is a hyperrectangle, and let $\mathcal{X}$ be a connected subgraph of $\mathbb{Z}^d$ containing $\mathcal{Y}$. We claim that $\mathcal{Y}$ has the retraction property. Indeed, it is readily checked that a retraction from $\mathcal{X}$ to $\mathcal{Y}$ can be obtained by mapping $x \in \mathcal{X}$ to the point in $\mathcal{Y}$ that is closest to x with respect to the Euclidean distance.

Example 4.7 (2-Point space) Assume that Q takes values in $\{0, 1\}$ and let $x, y \in \mathcal{X}$ with $Q_{xy} = 1$. A disjoint decomposition $\mathcal{X} = A_x \cup A_y$ with $x \in A_x$ and $y \in A_y$ is called an x - y cut. An edge $(u, v) \in \mathcal{E}$ is a cross if $u \in A_x$ and $v \in A_y$. The subset $\{x, y\}$ has the retraction property if and only if there exists an x - y cut such that no distinct crosses share a point. The correspondence between x - y cuts with this property and retractions is given by $T^{-1}(x) = A_x$, $T^{-1}(y) = A_y$ (Fig. 2).

Example 4.8 (Honeycomb lattice) Let $(\mathcal{X}, \mathcal{E})$ be a connected subgraph of the honeycomb lattice and define transition rates by setting $Q_{xy} = 1$ if $(x, y) \in \mathcal{E}$ and zero otherwise. Then each fundamental cell $\mathcal{Y} = \{y_1, \dots, y_6\}$ (see Fig. 3) has the retraction property. Indeed, to obtain a retraction of $\mathcal{X}$ onto $\mathcal{Y}$, we partition the plane into 6 sectors separated by rays that originate at the centre of $\mathcal{Y}$ and intersect the midpoints of the sides of $\mathcal{Y}$ orthogonally. A retraction is then obtained by mapping each $x \in \mathcal{X}$ to the unique $y \in \mathcal{Y}$ that belongs to the same sector (cf. Fig. 3).

Example 4.9 (Trees) Assume that the graph $(\mathcal{X}, \mathcal{E})$ is a tree, i.e., it does not contain a cycle. Every subtree $\mathcal{Y}$ of $\mathcal{X}$ has the retraction property, and a retraction can be constructed as follows: Fix a vertex $y \in \mathcal{Y}$. Since $\mathcal{X}$ is a tree, for every $x \in \mathcal{X}$ there is a unique path γ without self-intersections connecting x and y . The map assigning to x the first point where the path γ meets $\mathcal{Y}$ is a retraction of $\mathcal{X}$ onto $\mathcal{Y}$. Note that the retraction property depends only on the graph $(\mathcal{X}, \mathcal{E})$ and not on the choice of the transition rates Q (as long as they give rise to the same graph).

Theorem 4.10 (Extension of Hamilton–Jacobi subsolutions) *Let $(\mathcal{X}, Q, \pi)$ be a Markov triple, and let $\mathcal{Y}$ be a connected subset of $\mathcal{X}$. If $\mathcal{Y}$ has the retraction property, then every Hamilton–Jacobi subsolution on $\mathcal{Y}$ can be extended to a Hamilton–Jacobi subsolution on $\mathcal{X}$.*

Proof Let ϕ be a Hamilton–Jacobi subsolution on $\mathcal{Y}$, and let T be a retraction of $\mathcal{X}$ onto $\mathcal{Y}$. Define $\bar{\phi} : \mathcal{X} \rightarrow \mathbb{R}$ by $\bar{\phi} := \phi \circ T$, so that $\bar{\phi}|_{\mathcal{Y}} = \phi$ by (R1). We will show that for any $\bar{v} \in \mathcal{P}(\mathcal{X})$, there exists $v \in \mathcal{P}(\mathcal{Y})$ such that

$$\langle \dot{\bar{\phi}}_t, \bar{v} \rangle + \frac{1}{2} \|\nabla \bar{\phi}_t\|_{\bar{v}}^2 \leq \langle \dot{\phi}_t, v \rangle + \frac{1}{2} \|\nabla \phi_t\|_v^2 \quad (4.1)$$

Fig. 3 Part of the honeycomb lattice with a fundamental cell $\mathcal{Y}$. The labels indicate the image of the corresponding vertex under the retraction

for a.e. t . To improve readability, we omit the subscript t . As $\phi \in \text{HJ}_{\mathcal{Y}}$, the right-hand side of (4.1) is nonpositive, so this suffices to prove the theorem.

For $\bar{v} \in \mathcal{P}(\mathcal{X})$ define $v \in \mathcal{P}(\mathcal{Y})$ by $v := T_{\#}\bar{v}$. Clearly,

$$\langle \dot{\bar{\phi}}, \bar{v} \rangle = \langle \dot{\phi} \circ T, \bar{v} \rangle = \langle \dot{\phi}, T_{\#}\bar{v} \rangle = \langle \dot{\phi}, v \rangle.$$

It thus remains to show that $\|\nabla \bar{\phi}_t\|_{\bar{v}} \leq \|\nabla \phi_t\|_v$.

Splitting the sum we obtain

$$\begin{aligned} \|\nabla \bar{\phi}\|_{\bar{v}}^2 &= \frac{1}{2} \sum_{x, x' \in \mathcal{X}} \Lambda(\bar{v}_x Q_{xx'}, \bar{v}_{x'} Q_{x'x})(\bar{\phi}_x - \bar{\phi}_{x'})^2 \\ &= \frac{1}{2} \sum_{\substack{y, y' \in \mathcal{Y} \\ y \neq y'}} (\phi_y - \phi_{y'})^2 \sum_{\substack{x \in T^{-1}(y) \\ x' \in T^{-1}(y')}} \Lambda(\bar{v}_x Q_{xx'}, \bar{v}_{x'} Q_{x'x}). \end{aligned}$$

The concavity and homogeneity of Λ imply

$$\begin{aligned} &\sum_{\substack{x \in T^{-1}(y) \\ x' \in T^{-1}(y')}} \Lambda(\bar{v}_x Q_{xx'}, \bar{v}_{x'} Q_{x'x}) \\ &\leq \Lambda\left(\sum_{x \in T^{-1}(y)} \sum_{x' \in T^{-1}(y')} \bar{v}_x Q_{xx'}, \sum_{x' \in T^{-1}(y')} \sum_{x \in T^{-1}(y)} \bar{v}_{x'} Q_{x'x}\right). \end{aligned}$$

Given $x \in \mathcal{X}$ and $y, y' \in \mathcal{Y}$ with $y \neq y'$ and $T(x) = y$, the retraction property (R2) implies that $\sum_{x' \in T^{-1}(y')} Q_{xx'} \leq Q_{yy'}$ (and the same holds with primed and unprimed variables interchanged). Hence the monotonicity of Λ yields

$$\begin{aligned} &\Lambda\left(\sum_{x \in T^{-1}(y)} \sum_{x' \in T^{-1}(y')} \bar{v}_x Q_{xx'}, \sum_{x' \in T^{-1}(y')} \sum_{x \in T^{-1}(y)} \bar{v}_{x'} Q_{x'x}\right) \\ &\leq \Lambda\left(Q_{yy'} \sum_{x \in T^{-1}(y)} \bar{v}_x, Q_{yy'} \sum_{x' \in T^{-1}(y')} \bar{v}_{x'}\right) = \Lambda(v_y Q_{yy'}, v_{y'} Q_{y'y}). \end{aligned}$$

Combining these inequalities, we infer that

$$\|\nabla \bar{\phi}\|_{\bar{\nu}}^2 \leq \frac{1}{2} \sum_{y, y' \in \mathcal{Y}} (\phi_y - \phi_{y'})^2 \Lambda(v_y Q_{yy'}, v_{y'} Q_{y'y}) = \|\nabla \phi\|_{\bar{\nu}}^2,$$

which completes the proof. $\square$

The following result shows that any pair of measures supported in a set $\mathcal{Y}$ with the retraction property can be connected by a geodesic supported in $\mathcal{Y}$.

Theorem 4.11 (Weak locality under the retraction property) *Let $(\mathcal{X}, Q, \pi)$ be a Markov triple, and let $\mathcal{Y}$ be a subset of $\mathcal{X}$ with the retraction property. For all $\mu^0, \mu^1 \in \mathcal{P}(\mathcal{X})$ with support in $\mathcal{Y}$ there exists a minimising $\mathcal{W}$ -geodesic $(\mu_t)_{t \in [0,1]} \subseteq \mathcal{P}(\mathcal{X})$ connecting μ^0 and μ^1 such that μ_t has support in $\mathcal{Y}$ for all $t \in [0, 1]$.*

In fact, we will show that any $\mathcal{W}_{\mathcal{Y}}$ -geodesic $(\mu_t)_t \subseteq \mathcal{P}(\mathcal{Y})$ is also a $\mathcal{W}_{\mathcal{X}}$ -geodesic when regarded as a curve in $\mathcal{P}(\mathcal{X})$.

Proof Let $(\mu_t)_t$ be a minimising geodesic in $\mathcal{P}(\mathcal{Y})$ satisfying the continuity Eq. (2.1) with momentum vector field $(V_t)_t$. Consider the extension to $\mathcal{X}$ defined by $\bar{\mu}_t(x) = 0$ if $x \notin \mathcal{Y}$ and $\bar{V}_t(x, x') = 0$ if $x \notin \mathcal{Y}$ or $x' \notin \mathcal{Y}$. Clearly, $(\bar{\mu}_t, \bar{V}_t)_t$ has the same action as $(\mu_t, V_t)_t$.

Let $\varepsilon > 0$. Since $(\mu_t, V_t)_t$ is a geodesic in $\mathcal{P}(\mathcal{Y})$, Theorem 3.4 (applied in $\mathcal{P}(\mathcal{Y})$) implies that there exists a Hamilton–Jacobi subsolution $\phi \in \text{HJ}_{\mathcal{Y}}$ such that

$$\langle \phi_1, \mu^1 \rangle - \langle \phi_0, \mu^0 \rangle \geq \frac{1}{2} \int_0^1 \mathcal{A}(\mu_t, V_t) \, dt - \varepsilon.$$

By Theorem 4.10, ϕ can be extended to a Hamilton–Jacobi subsolution $\bar{\phi} \in \text{HJ}_{\mathcal{X}}$. In particular, using Theorem 3.4 once more (this time in $\mathcal{P}(\mathcal{X})$),

$$\langle \phi_1, \mu^1 \rangle - \langle \phi_0, \mu^0 \rangle = \langle \bar{\phi}_1, \mu^1 \rangle - \langle \bar{\phi}_0, \mu^0 \rangle \leq \frac{1}{2} \mathcal{W}_{\mathcal{X}}^2(\mu^0, \mu^1).$$

Since $\varepsilon > 0$ is arbitrary, it follows that $\int_0^1 \mathcal{A}(\mu_t, V_t) \, dt \leq \mathcal{W}_{\mathcal{X}}^2(\mu^0, \mu^1)$, which yields the result. $\square$

It follows from the previous result that Ricci curvature bounds in the sense of [5, 18] are inherited by subsets with the retraction property. We recall that a Markov triple $(\mathcal{X}, Q, \pi)$ is said to have *Ricci curvature bounded from below* by $\kappa \in \mathbb{R}$ if for any $\mu_0, \mu_1 \in \mathcal{P}(\mathcal{X})$, and for some (equivalently, for any) $\mathcal{W}$ -geodesic (μ_t) connecting μ_0 and μ_1 , the relative entropy $\mu \mapsto \text{Ent}_{\pi}(\mu) := \sum_{x \in \mathcal{X}} \mu(x) \log \left(\frac{\mu(x)}{\pi(x)} \right)$ satisfies the following κ -convexity inequality, for any $0 \leq t \leq 1$:

$$\text{Ent}_{\pi}(\mu_t) \leq (1-t) \text{Ent}_{\pi}(\mu_0) + t \text{Ent}_{\pi}(\mu_1) - \frac{\kappa}{2} t(1-t) \mathcal{W}(\mu_0, \mu_1)^2.$$

In this case we write $\text{Ric}(\mathcal{X}, Q, \pi) \geq \kappa$; cf. [5, 18] for further details.

Corollary 4.12 *Let $(\mathcal{X}, Q, \pi)$ be a Markov triple, and let $\mathcal{Y}$ be a subset of $\mathcal{X}$ with the retraction property. If $\text{Ric}(\mathcal{X}, Q, \pi) \geq \kappa$ for some $\kappa \in \mathbb{R}$, then $\text{Ric}(\mathcal{Y}, Q|_{\mathcal{Y}}, \pi|_{\mathcal{Y}}) \geq \kappa$ as well.*

Proof Take $\mu_0, \mu_1 \in \mathcal{P}(\mathcal{Y})$, and let $(\mu_t)_t$ be a $\mathcal{W}_{\mathcal{Y}}$ -geodesic connecting them. By Theorem 4.11, $(\mu_t)_t$ is also a geodesic in $\mathcal{P}(\mathcal{X})$. Since $\text{Ric}(\mathcal{X}, Q, \pi) \geq \kappa$, it follows that $t \mapsto \text{Ent}(\mu_t|\pi)$ is κ -convex. As $\text{Ent}_{\pi|_{\mathcal{Y}}}(\mu_t) = \text{Ent}_{\pi}(\mu_t) + \log(\pi(\mathcal{Y}))$ and $\mathcal{W}_{\mathcal{Y}}(\mu_0, \mu_1) = \mathcal{W}_{\mathcal{X}}(\mu_0, \mu_1)$, we infer that $t \mapsto \text{Ent}_{\pi|_{\mathcal{Y}}}(\mu_t)$ is κ -convex as well, which yields the result. $\square$

5 Optimal transport avoids dead ends

In this section we prove the intuitively natural statement that optimal curves do not transport mass into “dead ends”. We formalise this concept by considering the gluing of two Markov triples along a vertex.

Definition 5.1 (*Gluing of Markov triples*) Let $(\mathcal{X}_1, Q_1, \pi_1)$ and $(\mathcal{X}_2, Q_2, \pi_2)$ be Markov triples, and fix $x_1 \in \mathcal{X}_1, x_2 \in \mathcal{X}_2$. The *gluing* of the two triples at x_1, x_2 is the Markov triple $(\mathcal{X}, Q, \pi)$ defined by setting

$$\mathcal{X} = (\mathcal{X}_1 \sqcup \mathcal{X}_2) / \{x_1, x_2\}$$

and $* = [x_1] = [x_2]$. For brevity, let us write $\mathcal{X}'_i := \mathcal{X}_i \setminus \{x_i\}$. We have canonical injections $\mathcal{X}'_1 \rightarrow \mathcal{X}, \mathcal{X}'_2 \rightarrow \mathcal{X}$, and we identify elements of $(\mathcal{X}_1 \sqcup \mathcal{X}_2) \setminus \{x_1, x_2\}$ with their respective images. We define transition rates $Q : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ by

$$Q(x, y) = \begin{cases} Q_i(x, y) & \text{if } x, y \in \mathcal{X}'_i, \\ Q_i(x, x_i) & \text{if } x \in \mathcal{X}'_i \text{ and } y = *, \\ Q_i(x_i, y) & \text{if } x = * \text{ and } y \in \mathcal{X}'_i, \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to see that Q is irreducible and reversible, and the unique invariant probability measure is given by

$$\pi(x) = \frac{1}{1 - \pi_1(\mathcal{X}'_1)\pi_2(\mathcal{X}'_2)} \times \begin{cases} \pi_1(x)\pi_2(x_2) & \text{if } x \in \mathcal{X}'_1, \\ \pi_1(x_1)\pi_2(x) & \text{if } x \in \mathcal{X}'_2, \\ \pi_1(x_1)\pi_2(x_2) & \text{if } x = *. \end{cases}$$

Definition 5.2 (*Dead end*) Let $(\mathcal{X}, Q, \pi)$ be a Markov triple, and let $\mathcal{X}_1, \mathcal{X}_2 \subseteq \mathcal{X}$. We say that $\mathcal{X}_2$ is a *dead end* for $\mathcal{X}_1$ (and vice versa) if the intersection of $\mathcal{X}_1$ and $\mathcal{X}_2$ contains exactly one point (denoted “*”), and moreover, $Q(x, y) = Q(y, x) = 0$ whenever $x \in \mathcal{X}'_1$ and $y \in \mathcal{X}'_2$. Here, we write $\mathcal{X}'_i = \mathcal{X}_i \setminus \{*\}$.

Remark 5.3 The notions of dead end and gluing of Markov triples are compatible in the following sense: Let $(\mathcal{X}, Q, \pi)$ be a Markov triple, and suppose that $\mathcal{X}_2 \subseteq \mathcal{X}$ is a dead end for $\mathcal{X}_1 \subseteq \mathcal{X}$ with intersection point *. Then one recovers $(\mathcal{X}, Q, \pi)$ by gluing together the restrictions of $\mathcal{X}$ to $\mathcal{X}_1$ and $\mathcal{X}_2$ at *.

Proposition 5.4 Let $(\mathcal{X}_1, Q_1, \pi_1)$ and $(\mathcal{X}_2, Q_2, \pi_2)$ be Markov triples, and let $(\mathcal{X}, Q, \pi)$ be the Markov triple obtained by gluing the triples at $x_1 \in \mathcal{X}_1$ and $x_2 \in \mathcal{X}_2$. Then $\mathcal{X}_1$ and $\mathcal{X}_2$ have the retraction property as subsets of $\mathcal{X}$.

Proof Define $T : \mathcal{X} \rightarrow \mathcal{X}_1$ by $T(x) = x$ for $x \in \mathcal{X}_1$ and $T(x) = *$ for $x \in \mathcal{X}'_2$. One verifies that T indeed defines a retraction by distinguishing cases. $\square$

In view of Theorem 4.11, the previous result implies that any two measures μ_0, μ_1 supported in (the image of) $\mathcal{X}_1$ can be connected by a geodesic that is supported in $\mathcal{X}_1$ for all times; i.e., weak locality holds. We will now show that in fact strong locality holds: any geodesic connecting μ_0 and μ_1 has to be supported in $\mathcal{X}_1$.

Theorem 5.5 Let $(\mathcal{X}_1, Q_1, \pi_1)$ and $(\mathcal{X}_2, Q_2, \pi_2)$ be Markov triples, and let $(\mathcal{X}, Q, \pi)$ be the Markov triple obtained by gluing the triples at $x_1 \in \mathcal{X}_1$ and $x_2 \in \mathcal{X}_2$. If $(\mu_t)_{t \in [0,1]}$ is a geodesic in $(\mathcal{P}(\mathcal{X}), \mathcal{W})$ with $\text{supp } \mu_0, \text{supp } \mu_1 \subseteq \mathcal{X}_1$, then $\text{supp } \mu_t \subseteq \mathcal{X}_1$ for all $t \in [0, 1]$.

Proof Let $t \mapsto V_t \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ be an anti-symmetric momentum vector field such that (μ, V) is a solution to the continuity equation with $\int_0^1 \mathcal{A}(\mu_t, V_t) dt = \mathcal{W}^2(\mu_0, \mu_1)$. We define a new curve $t \mapsto \bar{\mu}_t \in \mathcal{P}(\mathcal{X})$ by

$$\bar{\mu}_t(x) = \begin{cases} \mu_t(x) & \text{if } x \in \mathcal{X}'_1, \\ \mu_t(*) + \sum_{y \in \mathcal{X}'_2} \mu_t(y) & \text{if } x = *, \\ 0 & \text{otherwise,} \end{cases}$$

and a new anti-symmetric momentum vector field $t \mapsto \bar{V}_t \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ by

$$\bar{V}_t(x, y) = \begin{cases} V_t(x, y) & \text{if } x, y \in \mathcal{X}'_1 \cup \{*\}, \\ 0 & \text{otherwise.} \end{cases}$$

We claim that $(\bar{\mu}, \bar{V})$ solves the continuity Eq. (2.1) as well.

Indeed, this statement trivially holds for any $x \in \mathcal{X} \setminus \{*\}$. To prove the claim at $*$, we note that for any $y \in \mathcal{X}'_2$,

$$\frac{d}{dt} \mu_t(y) = \sum_{x \in \mathcal{X}'_2 \cup \{*\}} V_t(x, y).$$

Therefore, using the anti-symmetry of V_t ,

$$\sum_{y \in \mathcal{X}'_2} \frac{d}{dt} \mu_t(y) = \sum_{y \in \mathcal{X}'_2} \sum_{x \in \mathcal{X}'_2 \cup \{*\}} V_t(x, y) = \sum_{y \in \mathcal{X}'_2} V_t(*, y).$$

Furthermore,

$$\frac{d}{dt} \mu_t(*) = \sum_{y \in \mathcal{X}'_1} V_t(y, *) + \sum_{y \in \mathcal{X}'_2} V_t(y, *),$$

hence by another application of the anti-symmetry,

$$\frac{d}{dt} \bar{\mu}_t(*) = \frac{d}{dt} \mu_t(*) + \sum_{y \in \mathcal{X}'_2} \frac{d}{dt} \mu_t(y) = \sum_{y \in \mathcal{X}'_1} V_t(y, *) = \sum_{y \in \mathcal{X}} \bar{V}_t(y, *),$$

which proves the claim.

For all $t \in (0, 1)$ and $x, y \in \mathcal{X}$, we clearly have

$$A(\mu_t(x) Q(x, y), \mu_t(y) Q(y, x), V_t(x, y)) \geq A(\bar{\mu}_t(x) Q(x, y), \bar{\mu}_t(y) Q(y, x), \bar{V}_t(x, y)).$$

Moreover, if $\mu_t(\mathcal{X}'_2) > 0$ for some $t \in (0, 1)$, then there exists $z \in \mathcal{X}'_2$ such that $V_t(*, z) > 0$ and $\Lambda(\mu_t(*) Q(*, z), \mu_t(z) Q(z, *)) > 0$ for all t on a set of positive measure in $(0, 1)$. Therefore,

$$A(\mu_t(*) Q(*, z), \mu_t(z) Q(z, *), V_t(*, z)) > 0 = A(\bar{\mu}_t(*) Q(*, z), \bar{\mu}_t(z) Q(z, *), \bar{V}_t(*, z)).$$

This strict inequality contradicts the fact that $(\mu_t)_{t \in [0, 1]}$ is a geodesic. $\square$

6 Nonlocality of optimal transport on the triangle

Consider a Markov triple $(\mathcal{X}, Q, \pi)$ and a connected subset $\mathcal{Y} \subseteq \mathcal{X}$. In this section we show that locality of geodesics in $\mathcal{P}(\mathcal{Y})$ may fail if $\mathcal{Y}$ does not have the retraction property. We consider the simplest possible setting, where $(\mathcal{X}, Q, \pi)$ corresponds to simple random walk on a triangle, and $\mathcal{Y} \subseteq \mathcal{X}$ is a two-point set. We show that the canonical lift of a geodesic between Dirac measures on the two-point space is *not* an optimal curve in $\mathcal{P}(\mathcal{X})$, by constructing a competitor that transports mass along all edges.

Throughout this section we make the following additional assumption on the mean Λ .

Assumption 6.1 For any $s > 0$ we have

$$\Lambda(s, t) \rightarrow \infty \quad \text{as } t \rightarrow \infty. \quad (6.1)$$

If $\Lambda(0, t) > 0$ for $t > 0$, then (6.1) also holds for $s = 0$.

Clearly, this assumption is satisfied for the arithmetic, geometric, and logarithmic means, but not for the harmonic mean.

The main result of this section relies on the following lemma concerning the variation of the action functional on cycles of arbitrary length.

Lemma 6.2 For $n \geq 3$, let $\mathcal{X} = \mathbb{Z}/n\mathbb{Z}$ be equipped with transition rates Q_{ij} such that $Q_{i,i+1}, Q_{i+1,i} > 0$ for all $i \in \mathcal{X}$ and $Q_{ij} = 0$ otherwise. Let $\mu, \nu \in \mathcal{P}(\mathcal{X})$, and let $V, U \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ be anti-symmetric, and such that both $A(\mu, V)$ and $A(\nu, U)$ are finite. Assume that $\mu_1, \mu_2 > 0$ and $\mu_i = 0$ for all $i \neq 1, 2$, $V_{12} \neq 0$, and $V_{ij} = 0$ for all $\{i, j\} \neq \{1, 2\}$ and that $U_{12} = 0$. For $\alpha \in [0, 1]$ we define $\mu^\alpha = (1 - \alpha)\mu + \alpha\nu$ and $V^\alpha = (1 - \alpha)V + \alpha U$. Then we have:

$$\begin{aligned} \lim_{\alpha \rightarrow 0} \frac{1}{\alpha} (A(\mu^\alpha, V^\alpha) - A(\mu, V)) &= \frac{-V_{12}^2}{\Lambda(\mu_1 Q_{12}, \mu_2 Q_{21})} \left[1 + \frac{\partial_1 \Lambda(\mu_1 Q_{12}, \mu_2 Q_{21}) \nu_1 Q_{12} + \partial_2 \Lambda(\mu_1 Q_{12}, \mu_2 Q_{21}) \nu_2 Q_{21}}{\Lambda(\mu_1 Q_{12}, \mu_2 Q_{21})} \right] \\ &\quad + \sum_{i=3}^{n-1} \frac{U_{i,i+1}^2}{\Lambda(\nu_i Q_{i,i+1}, \nu_{i+1} Q_{i+1,i})}. \end{aligned} \quad (6.2)$$

Proof First note that

$$A(\mu, V) = \frac{V_{12}^2}{\Lambda(\mu_1 Q_{12}, \mu_2 Q_{21})}.$$

Using the 1-homogeneity of Λ we observe that

$$\begin{aligned} A(\mu^\alpha, V^\alpha) &= \frac{(1 - \alpha) V_{12}^2}{\Lambda((\mu_1 + \frac{\alpha}{1-\alpha} \nu_1) Q_{12}, (\mu_2 + \frac{\alpha}{1-\alpha} \nu_2) Q_{21})} + \alpha \sum_{i=3}^{n-1} \frac{U_{i,i+1}^2}{\Lambda(\nu_i Q_{i,i+1}, \nu_{i+1} Q_{i+1,i})} \\ &\quad + \frac{\alpha U_{23}^2}{\Lambda((\frac{1-\alpha}{\alpha} \mu_2 + \nu_2) Q_{23}, \nu_3 Q_{32})} + \frac{\alpha U_{n1}^2}{\Lambda(\nu_n Q_{n1}, (\frac{1-\alpha}{\alpha} \mu_1 + \nu_1) Q_{1n})} \\ &=: T_1 + T_2 + T_3 + T_4. \end{aligned}$$

Since $\mu_1, \mu_2 > 0$, the first term in (6.2) is well-defined and easily seen to be the limit of $(T_1 - A(\mu, V))/\alpha$. Obviously, T_2/α converges to the second term in (6.2). Finally, T_3

vanishes unless $U_{23} \neq 0$. But in this case, since $(1 - \alpha)/\alpha \rightarrow \infty$ as $\alpha \rightarrow 0$, we see that T_3/α converges to zero as $\alpha \rightarrow 0$ by Assumption 6.1. A similar argument applies to T_4 . $\square$

Now we can prove the nonlocality result.

Theorem 6.3 *Let $(\mathcal{X}, Q, \pi)$ be a Markov triple with $\mathcal{X} = \{1, 2, 3\}$ and such that $Q(x, y) > 0$ for all $x \neq y$. Let $(\mu_t)_{t \in [0,1]}$ be a W -geodesic connecting $\mu_0 = \delta_1$ to $\mu_1 = \delta_2$. Then, $\mu_t(3) > 0$ for some $0 < t < 1$.*

As $\mu_t(3) > 0$ for some $0 < t < 1$, the result implies that mass is transported along the edges $(1, 3)$ and $(3, 2)$.

Proof Suppose that the geodesic $(\mu_t, V_t)_{t \in [0,1]}$ transports only along the edge $(1, 2)$, i.e., $V_t(2, 3) = V_t(3, 1) = 0$ for a.e. $t \in (0, 1)$. Then (μ_t, V_t) must be given by the corresponding geodesic on the two point space $\{1, 2\}$. Obviously, we have $\mu_t(1), \mu_t(2) > 0$ for all $t \in (0, 1)$. Let $(v, U) \in \text{CE}_1(\delta_1, \delta_2)$ be a curve of finite action such that $U_t(1, 2) = 0$ for a.e. t and $v_t(1), v_t(2), v_t(3) > 0$ for all $0 < t < 1$. Define $(\mu^\alpha, V^\alpha) \in \text{CE}_1(\delta_1, \delta_2)$ by $\mu^\alpha = (1 - \alpha)\mu + \alpha v$ and $V^\alpha = (1 - \alpha)V + \alpha U$ for $\alpha \in [0, 1]$. Then Lemma 6.2 yields for a.e. t :

$$\lim_{\alpha \rightarrow 0} \frac{1}{\alpha} [\mathcal{A}(\mu_t^\alpha, V_t^\alpha) - \mathcal{A}(\mu_t, V_t)] < 0.$$

Consequently, there exists $\alpha > 0$ such that

$$\int_0^1 \mathcal{A}(\mu_t^\alpha, V_t^\alpha) dt < \int_0^1 \mathcal{A}(\mu_t, V_t) dt,$$

contradicting the optimality of (μ, V) . $\square$

Acknowledgements Open access funding provided by Institute of Science and Technology (IST Austria). Matthias Erbar gratefully acknowledges support by the German Research Foundation through the Hausdorff Centre for Mathematics and the Collaborative Research Centre 1060, *The Mathematics of Emergent Effects*. Jan Maas gratefully acknowledges support by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 716117), and by the Austrian Science Fund (FWF), Project SFB F65. Melchior Wirth gratefully acknowledges financial support by the German Academic Scholarship Foundation (Studienstiftung des deutschen Volkes) and by the German Research Foundation (DFG) via RTG 1523. We thank an anonymous referee for suggesting to include the statement of Corollary 4.12 in the paper.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Buttazzo, G.: Semicontinuity, Relaxation and Integral Representation in the Calculus of Variations. Pitman Research Notes in Mathematics Series, vol. 207. Longman Scientific & Technical, Harlow (1989)
2. Chow, S.-N., Huang, W., Li, Y., Zhou, H.: Fokker–Planck equations for a free energy functional or Markov process on a graph. Arch. Ration. Mech. Anal. **203**(3), 969–1008 (2012)
3. Crandall, M.G., Lions, P.-L.: Viscosity solutions of Hamilton–Jacobi equations. Trans. Am. Math. Soc. **277**(1), 1–42 (1983)
4. Erbar, M., Henderson, C., Menz, G., Tetali, P.: Ricci curvature bounds for weakly interacting Markov chains. Electron. J. Probab. **22**, Paper No. 40, 23 (2017)

5. Erbar, M., Maas, J.: Ricci curvature of finite Markov chains via convexity of the entropy. *Arch. Ration. Mech. Anal.* **206**(3), 997–1038 (2012)
6. Erbar, M., Maas, J.: Gradient flow structures for discrete porous medium equations. *Discrete Contin. Dyn. Syst.* **34**, 1355–1374 (2014)
7. Erbar, M., Maas, J., Tetali, P.: Discrete Ricci curvature bounds for Bernoulli–Laplace and random transposition models. *Ann. Fac. Sci. Toulouse Math.* (6) **24**(4), 781–800 (2015)
8. Fathi, M., Maas, J.: Entropic Ricci curvature bounds for discrete interacting systems. *Ann. Appl. Probab.* **26**(3), 1774–1806 (2016)
9. Gangbo, W., Li, W., Mou, C.: Geodesic of minimal length in the set of probability measures on graphs. *ArXiv e-prints*, December (2017)
10. Gozlan, N., Roberto, C., Samson, P.-M.: Hamilton Jacobi equations on metric spaces and transport entropy inequalities. *Rev. Mat. Iberoam.* **30**(1), 133–163 (2014)
11. Jordan, R., Kinderlehrer, D., Otto, F.: The variational formulation of the Fokker–Planck equation. *SIAM J. Math. Anal.* **29**(1), 1–17 (1998)
12. Lisini, S.: Characterization of absolutely continuous curves in Wasserstein spaces. *Calc. Var. Partial Differ. Equ.* **28**(1), 85–120 (2007)
13. Lott, J., Villani, C.: Hamilton–Jacobi semigroup on length spaces and applications. *J. Math. Pures Appl.* (9) **88**(3), 219–229 (2007)
14. Lott, J., Villani, C.: Ricci curvature for metric-measure spaces via optimal transport. *Ann. Math.* (2) **169**(3), 903–991 (2009)
15. Maas, J.: Gradient flows of the entropy for finite Markov chains. *J. Funct. Anal.* **261**(8), 2250–2292 (2011)
16. McCann, R.J.: A convexity principle for interacting gases. *Adv. Math.* **128**(1), 153–179 (1997)
17. Mielke, A.: A gradient structure for reaction-diffusion systems and for energy-drift-diffusion systems. *Nonlinearity* **24**(4), 1329–1346 (2011)
18. Mielke, A.: Geodesic convexity of the relative entropy in reversible Markov chains. *Calc. Var. Partial Differ. Equ.* **48**(1–2), 1–31 (2013)
19. Maas, J., Matthes, D.: Long-time behavior of a finite volume discretization for a fourth order diffusion equation. *Nonlinearity* **29**(7), 1992–2023 (2016)
20. Otto, F.: The geometry of dissipative evolution equations: the porous medium equation. *Commun. Partial Differ. Equ.* **26**(1–2), 101–174 (2001)
21. Shu, Y.: Hamilton–Jacobi equations on graph and applications. *Potential Anal.* **48**(2), 125–157 (2018)
22. Sturm, K.-Th.: On the geometry of metric measure spaces. I and II. *Acta Math.* **196**(1), 65–177 (2006)
23. Villani, C.: *Topics in Optimal Transportation*. Graduate Studies in Mathematics, vol. 58. American Mathematical Society, Providence (2003)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.