

GRADIENT FLOWS OF THE ENTROPY FOR FINITE MARKOV CHAINS

JAN MAAS

ABSTRACT. Let K be an irreducible and reversible Markov kernel on a finite set \mathcal{X} . We construct a metric \mathcal{W} on the set of probability measures on \mathcal{X} and show that with respect to this metric, the law of the continuous time Markov chain evolves as the gradient flow of the entropy. This result is a discrete counterpart of the Wasserstein gradient flow interpretation of the heat flow in \mathbb{R}^n by Jordan, Kinderlehrer, and Otto (1998). The metric \mathcal{W} is similar to, but different from, the L^2 -Wasserstein metric, and is defined via a discrete variant of the Benamou-Brenier formula.

1. INTRODUCTION

Since the seminal work of Jordan, Kinderlehrer and Otto [14], it is known that the heat flow on \mathbb{R}^n is the gradient flow of the Boltzmann-Shannon entropy with respect to the L^2 -Wasserstein metric on the space of probability measures on \mathbb{R}^n . This discovery has been the starting point for many developments in evolution equations, probability theory and geometry. We refer to the monographs [1, 27, 28] for an overview. By now a similar interpretation of the heat flow has been established in a wide variety of settings, including Riemannian manifolds [10], Hilbert spaces [2], Wiener spaces [11], Finsler spaces [19], Alexandrov spaces [13] and metric measure spaces [12, 25].

Let $(K(x, y))_{x, y \in \mathcal{X}}$ be an irreducible and reversible Markov transition kernel on a finite set \mathcal{X} , and consider the continuous time semigroup $(H(t))_{t \geq 0}$ associated with K . This semigroup is defined by $H(t) = e^{t(K-I)}$, and can be interpreted as the ‘heat semigroup’ on \mathcal{X} with respect to the geometry determined by the Markov kernel K . Therefore it seems natural to ask whether the heat flow can also be identified as the gradient flow of an entropy functional with respect to some metric on the space of probability densities on \mathcal{X} . Unfortunately, it is easily seen that the L^2 -Wasserstein metric over a discrete space is not appropriate for this purpose. In fact, since the metric derivative of the heat flow in the Wasserstein metric is typically infinite in a discrete setting, the heat flow can not be interpreted as the gradient flow of *any* functional in the L^2 -Wasserstein metric. (We refer to Section 2 for a more detailed discussion.)

2000 *Mathematics Subject Classification.* Primary 60J27; Secondary: 28A33, 49Q20, 60B10.

Key words and phrases. Markov chains, entropy, gradient flows, Wasserstein metric, optimal transportation.

The author is supported by Rubicon subsidy 680-50-0901 of the Netherlands Organisation for Scientific Research (NWO).

The main contribution of this paper is the construction of a metric \mathcal{W} on the space of probability densities on \mathcal{X} , which allows to extend the interpretation of the heat flow as the gradient flow of the entropy to the setting of finite Markov chains.

Notation. As before, let $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ be a Markov kernel on a finite space \mathcal{X} , i.e.,

$$K(x, y) \geq 0 \quad \forall x, y \in \mathcal{X}, \quad \sum_{y \in \mathcal{X}} K(x, y) = 1 \quad \forall x \in \mathcal{X}.$$

We assume that K is *irreducible*, which implies the existence of a unique steady state π . Thus π is a probability measure on \mathcal{X} , represented by a row vector that is invariant under right-multiplication by K :

$$\pi(y) = \sum_{x \in \mathcal{X}} \pi(x) K(x, y).$$

It follows from elementary Markov chain theory that π is strictly positive. We shall assume that K is *reversible*, i.e., $\pi(x)K(x, y) = \pi(y)K(y, x)$ for any $x, y \in \mathcal{X}$. Consider the set

$$\mathcal{P}(\mathcal{X}) := \left\{ \rho : \mathcal{X} \rightarrow \mathbb{R} \mid \rho(x) \geq 0 \quad \forall x \in \mathcal{X} ; \sum_{x \in \mathcal{X}} \pi(x) \rho(x) = 1 \right\}$$

consisting of all *probability densities* on \mathcal{X} . The subset consisting of those probability densities that are strictly positive is denoted by $\mathcal{P}_*(\mathcal{X})$. The *relative entropy* of a probability density $\rho \in \mathcal{P}(\mathcal{X})$ with respect to π is defined by

$$\mathcal{H}(\rho) = \sum_{x \in \mathcal{X}} \pi(x) \rho(x) \log \rho(x). \quad (1.1)$$

with the usual convention that $\rho(x) \log \rho(x) = 0$ if $\rho(x) = 0$.

Wasserstein-like metrics in a discrete setting. To motivate the definition of the metric \mathcal{W} , recall that for probability densities ρ_0, ρ_1 on \mathbb{R}^n , the Benamou-Brenier formula [3] asserts that the squared Wasserstein distance W_2 satisfies the identity

$$W_2(\rho_0, \rho_1)^2 = \inf_{\rho, \psi} \left\{ \int_0^1 \int_{\mathbb{R}^n} |\nabla \psi_t(x)|^2 \rho_t(x) \, dx \, dt \right\}, \quad (1.2)$$

where the infimum runs over sufficiently regular curves $\rho : [0, 1] \rightarrow \mathcal{P}(\mathbb{R}^n)$ and $\psi : [0, 1] \times \mathbb{R}^n \rightarrow \mathbb{R}$ satisfying the continuity equation

$$\begin{cases} \partial_t \rho + \nabla \cdot (\rho \nabla \psi) = 0, \\ \rho(0) = \rho_0, \quad \rho(1) = \rho_1. \end{cases} \quad (1.3)$$

Here, by a slight abuse of notation, $\mathcal{P}(\mathbb{R}^n)$ denotes the set of probability densities on \mathbb{R}^n . At least formally, the Benamou-Brenier formula has been interpreted by Otto [23] as a Riemannian metric on the space of probability densities on \mathbb{R}^n .

In the discrete setting, we shall define a class of pseudo-metrics \mathcal{W} (i.e., metrics which possibly attain the value $+\infty$) by mimicking the formulas

(1.2) and (1.3). In order to obtain a metric with the desired properties, it turns out to be necessary to define, for $\rho \in \mathcal{P}(\mathcal{X})$ and $x, y \in \mathcal{X}$,

$$\rho(x, y) := \theta(\rho(x), \rho(y)) ,$$

where $\theta : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a function satisfying (A1) – (A7) below. At this stage we remark that typical examples of admissible functions are the logarithmic mean $\theta(s, t) = \int_0^1 s^{1-pt} t^p \, dp$, the geometric mean $\theta(s, t) = \sqrt{st}$ and, more generally, the functions $\theta(s, t) = s^\alpha t^\alpha$ for $\alpha > 0$.

Now we are ready to state the definition of \mathcal{W} :

Definition. For $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$ we set

$$\mathcal{W}(\rho_0, \rho_1)^2 := \inf_{\rho, \psi} \left\{ \frac{1}{2} \int_0^1 \sum_{x, y \in \mathcal{X}} (\psi_t(x) - \psi_t(y))^2 K(x, y) \rho_t(x, y) \pi(x) \, dt \right\} ,$$

where the infimum runs over all piecewise C^1 curves $\rho : [0, 1] \rightarrow \mathcal{P}(\mathcal{X})$ and all measurable functions $\psi : [0, 1] \rightarrow \mathbb{R}^{\mathcal{X}}$ satisfying, for a.e. $t \in [0, 1]$,

$$\begin{cases} \frac{d}{dt} \rho_t(x) + \sum_{y \in \mathcal{X}} (\psi_t(y) - \psi_t(x)) K(x, y) \rho_t(x, y) = 0 & \forall x \in \mathcal{X} , \\ \rho(0) = \rho_0 , \quad \rho(1) = \rho_1 . \end{cases} \quad (1.4)$$

Remark. Similar to the Wasserstein metric, $\mathcal{W}(\rho_0, \rho_1)^2$ can be interpreted as the cost of transporting mass from its initial configuration ρ_0 to the final configuration ρ_1 . However, unlike the Wasserstein metric, the cost of transporting a unit mass from x to y depends on the amount of mass already present at x and y . In a continuous setting, metrics with these properties have been studied in the recent papers [6, 9]. The essential new feature of the metric considered in this paper is the fact that the dependence is *non-local*.

In order to state the first main result of the paper, we introduce some notation. Fix a probability density $\rho \in \mathcal{P}(\mathcal{X})$. We shall write $x \sim_\rho y$ if $x, y \in \mathcal{X}$ belong to the same connected component of the support of ρ . More formally, we say that $x \sim_\rho y$ if $x = y$, or if there exist $k \geq 1$ and $x_1, \dots, x_k \in \mathcal{X}$ such that

$$\rho(x, x_1) K(x, x_1), \rho(x_1, x_2) K(x_1, x_2), \dots, \rho(x_k, y) K(x_k, y) > 0 .$$

Furthermore, we set

$$C_\theta := \int_0^1 \frac{1}{\sqrt{\theta(1-r, 1+r)}} \, dr \in [0, \infty] .$$

It turns out that C_θ is the \mathcal{W} -distance between a Dirac mass and the uniform density on a two-point space $\{a, b\}$ endowed with the Markov kernel defined by $K(a, b) = K(b, a) = \frac{1}{2}$. Note that C_θ is finite if θ is the logarithmic or geometric mean. If $\theta(s, t) = s^\alpha t^\alpha$, then C_θ is finite for $0 < \alpha < 2$ and infinite for $\alpha \geq 2$.

For $\sigma \in \mathcal{P}(\mathcal{X})$ we shall write

$$\mathcal{P}_\sigma(\mathcal{X}) := \{ \rho \in \mathcal{P}(\mathcal{X}) : \mathcal{W}(\rho, \sigma) < \infty \} .$$

The first main result of this paper reads as follows:

Theorem 1.1. *The following assertions hold:*

- (1) \mathcal{W} defines a pseudo-metric on $\mathcal{P}(\mathcal{X})$.
- (2)
 - If $C_\theta < \infty$, then $\mathcal{W}(\rho_0, \rho_1) < \infty$ for all $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$.
 - If $C_\theta = \infty$, the following are equivalent for $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$:
 - (a) $\mathcal{W}(\rho_0, \rho_1) < \infty$;
 - (b) For all $x \in \mathcal{X}$ we have

$$\sum_{y \sim_{\rho_0} x} \rho_0(y) \pi(y) = \sum_{y \sim_{\rho_1} x} \rho_1(y) \pi(y).$$

- (3) For all $\sigma \in \mathcal{P}(\mathcal{X})$, \mathcal{W} metrizes the topology of weak convergence on $\mathcal{P}_\sigma(\mathcal{X})$.
- (4)
 - If $C_\theta < \infty$ and θ is concave, the metric space $(\mathcal{P}_*(\mathcal{X}), \mathcal{W})$ is a Riemannian manifold.
 - If $C_\theta = \infty$, the metric space $(\mathcal{P}_\sigma(\mathcal{X}), \mathcal{W})$ is a complete Riemannian manifold for all $\sigma \in \mathcal{P}(\mathcal{X})$.

Remark (Finiteness). Part (2) of the theorem above provides a complete characterisation of finiteness of \mathcal{W} for general Markov kernels, in terms of the behaviour of \mathcal{W} for kernels on a two-point space. If $C_\theta = \infty$, the statement can be rephrased informally by saying that the distance $\mathcal{W}(\rho_0, \rho_1)$ is finite if and only if the following conditions hold: ρ_0 and ρ_1 have equal support, and both measures assign the same mass to each connected component of their support. In particular, it is important to note that the distance between two *strictly* positive densities is finite.

Remark (Weak convergence). Although (3) asserts that \mathcal{W} metrizes the topology of weak convergence on $\mathcal{P}_\sigma(\mathcal{X})$ for every $\sigma \in \mathcal{P}(\mathcal{X})$, it follows from (2) that \mathcal{W} does *not* metrize this topology on the full space $\mathcal{P}(\mathcal{X})$ if $C_\theta = \infty$. In fact, a weakly convergent sequence in $\mathcal{P}_\sigma(\mathcal{X})$ converges in \mathcal{W} -metric if and only if the weak limit belongs to $\mathcal{P}_\sigma(\mathcal{X})$.

Remark (Non-compactness). If $C_\theta = \infty$, we hasten to point out that the Riemannian manifold $(\mathcal{W}, \mathcal{P}_\sigma(\mathcal{X}))$ can be a singleton. According to (2), this happens if and only if $K(x, y)\sigma(x, y) = 0$ for every $x \in \text{supp } \sigma$ and every $y \in \mathcal{X}$, which is for instance the case if σ is the density of a Dirac measure. If $\mathcal{P}_\sigma(\mathcal{X})$ consists of more than one element, it turns out that $(\mathcal{P}_\sigma(\mathcal{X}), \mathcal{W})$ is non-compact. By contrast, the L^2 -Wasserstein space over a compact metric space is compact.

Remark (Riemannian metric). The Riemannian metric on $(\mathcal{P}_*(\mathcal{X}), \mathcal{W})$ is a natural discrete analogue of the formal Riemannian metric on the Wasserstein space over \mathbb{R}^n . In fact, consider a smooth curve $(\rho_t)_{t \in [0,1]}$ in $\mathcal{P}_*(\mathcal{X})$ and take $t \in [0, 1]$. In Section 3 we shall prove that there exists a unique discrete gradient $\nabla \psi_t = (\psi_t(x) - \psi_t(y))_{x, y \in \mathbb{R}^n}$ such that the continuity equation (1.4) holds. In view of this observation, we shall identify the tangent space at $\rho \in \mathcal{P}_*(\mathcal{X})$ with the collection of discrete gradients

$$T_\rho := \{ \nabla \psi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}} : \psi \in \mathbb{R}^{\mathcal{X}} \}.$$

We shall regard the discrete gradient $\nabla \psi_t$ as being the tangent vector along the curve $t \mapsto \rho_t$. The distance \mathcal{W} is the Riemannian distance induced by

the inner product $\langle \cdot, \cdot \rangle_\rho$ on T_ρ given by

$$\langle \nabla \varphi, \nabla \psi \rangle_\rho = \frac{1}{2} \sum_{x, y \in \mathcal{X}} (\varphi(x) - \varphi(y))(\psi(x) - \psi(y))K(x, y)\rho(x, y)\pi(x).$$

This formula is analogous to the corresponding expression in the continuous case [23]. In Section 3 we obtain a similar description of the Riemannian metric on each of the components of $\mathcal{P}(\mathcal{X})$. If ρ is not strictly positive, the tangent space shall be identified with the collection of discrete gradients of an appropriate subset of functions on \mathcal{X} .

Remark (Two-point space). If K is a reversible Markov kernel on a space \mathcal{X} consisting of only two points, it is possible to obtain an explicit formula for the metric \mathcal{W} . We refer to Section 2 for an extensive discussion.

Example. If $C_\theta = \infty$, it follows from Theorem 1.1 that the incidence graph associated with the Markov kernel K determines the topology of $(\mathcal{P}(\mathcal{X}), \mathcal{W})$. Let us illustrate this fact by two simple examples on a three-point space $\mathcal{X} = \{x_1, x_2, x_3\}$.

If $K(x_i, x_j) > 0$ for all $i \neq j$, then the space $\mathcal{P}(\mathcal{X})$ consists of 7 distinct Riemannian manifolds:

- one 2-dimensional manifold: $\mathcal{P}_*(\mathcal{X})$;
- three 1-dimensional manifolds: for $i = 1, 2, 3$,

$$C_i := \{\rho \in \mathcal{P}(\mathcal{X}) : \rho(x_j) = 0 \text{ iff } j = i\}.$$

- three singletons: for $i = 1, 2, 3$,

$$D_i := \{\rho \in \mathcal{P}(\mathcal{X}) : \rho(x_j) = 0 \text{ iff } j \neq i\}.$$

If $K(x_1, x_2), K(x_2, x_3) > 0$ and $K(x_1, x_3) = 0$, then the space $\mathcal{P}(\mathcal{X})$ consists of infinitely many distinct Riemannian manifolds:

- one 2-dimensional manifold: $\mathcal{P}_*(\mathcal{X})$;
- two 1-dimensional manifolds: C_1 and C_3 ;
- infinitely many singletons: the three singletons D_i for $i = 1, 2, 3$, and the infinite collection

$$\{\{\rho\} : \rho(x_1) > 0, \rho(x_3) > 0, \rho(x_2) = 0\}.$$

The gradient flow of the entropy. Since the entropy functional \mathcal{H} restricts to a smooth functional on the Riemannian manifold $(\mathcal{P}_*(\mathcal{X}), \mathcal{W})$, it makes sense to consider the associated gradient flow. Let $D_t \rho$ denote the tangent vector field along a smooth curve $\rho : (0, \infty) \rightarrow \mathcal{P}_*(\mathcal{X})$ and let $\text{grad } \varphi$ denote the gradient of a smooth functional $\varphi : \mathcal{P}_*(\mathcal{X}) \rightarrow \mathbb{R}$.

Consider the continuous time Markov semigroup $H(t) = e^{t(K-I)}$, $t \geq 0$, associated with K . It follows from the theory of Markov chains that $H(t)$ maps $\mathcal{P}(\mathcal{X})$ into $\mathcal{P}_*(\mathcal{X})$. The second main result of this paper asserts that the ‘heat flow’ determined by $H(t)$ is the gradient flow of the entropy \mathcal{H} with respect to \mathcal{W} , if θ is the logarithmic mean.

Theorem 1.2 (Heat flow is gradient flow of entropy). *Let θ be the logarithmic mean. For $\rho \in \mathcal{P}(\mathcal{X})$ and $t \geq 0$, set $\rho_t := e^{t(K-I)}\rho$. Then the gradient flow equation*

$$D_t \rho = -\text{grad } \mathcal{H}(\rho_t)$$

holds for all $t > 0$.

Remark. The choice of the logarithmic mean is essential in Theorem 1.2 if one wishes to identify the heat flow as the gradient flow of the entropy associated with the function $f(\rho) = \rho \log \rho$. In section 4 we prove that analogous results can be proved for certain different functions f , if one replaces the logarithmic mean by $\theta(s, t) = \frac{s-t}{f'(s)-f'(t)}$. The appearance of the logarithmic mean in discrete heat flow problems is not surprising. In fact, the ‘‘Log Mean Temperature Difference’’, usually called LMTD, plays an important rôle in the engineering literature on heat and mass transfer problems (see, e.g., [18]), in particular in heat flow through long cylinders (see also [4, Section 4.5] for a discussion).

Remark. For Markov chains on a two-point space $\{-1, 1\}$ we shall show in Section 2 that (under mild additional assumptions) the metric \mathcal{W} is the *unique* metric for which the gradient flow of the entropy coincides with the heat flow. We refer to Proposition 2.13 below for a precise statement.

Ricci curvature in a discrete setting. A synthetic theory of Ricci curvature in metric measure spaces has been developed recently by Lott-Sturm-Villani [17, 26]. These authors defined lower bounds on the Ricci curvature of a *geodesic* metric measure space in terms of convexity properties of the entropy functional along geodesics in the L^2 -Wasserstein metric. For long there has been interest to define and study a notion of Ricci curvature on discrete spaces, but unfortunately the Lott-Sturm-Villani definition cannot be applied directly. The reason is that geodesics in the L^2 -Wasserstein space do typically not exist if the underlying metric space is discrete, even in the simplest possible example of the two-point space (see Section 2 below for more details).

The metric \mathcal{W} constructed in this paper does not have this defect. By a lower-semicontinuity argument it can be shown that every pair of probability densities in $\mathcal{P}(\mathcal{X})$ can be joined by a constant speed geodesic. Since \mathcal{W} takes over the rôle of the L^2 -Wasserstein metric if θ is the logarithmic mean, the following modification of the Lott-Sturm-Villani definition of Ricci curvature seems natural:

Definition 1.3 (Ricci curvature lower bound). *Let $K = (K(x, y))_{x, y \in \mathcal{X}}$ be an irreducible and reversible Markov kernel on a finite space \mathcal{X} . Then K is said to have Ricci curvature bounded from below by $\kappa \in \mathbb{R}$, if for every $\bar{\rho}_0, \bar{\rho}_1 \in \mathcal{P}(\mathcal{X})$ there exists a constant speed geodesic $(\rho_t)_{t \in [0, 1]}$ in $(\mathcal{P}(\mathcal{X}), \mathcal{W})$ satisfying $\rho_0 = \bar{\rho}_0$, $\rho_1 = \bar{\rho}_1$, and*

$$\mathcal{H}(\rho_t) \leq (1-t)\mathcal{H}(\rho_0) + t\mathcal{H}(\rho_1) - \frac{\kappa}{2}t(1-t)\mathcal{W}(\rho_0, \rho_1)^2$$

for all $t \in [0, 1]$. We set

$$\text{Ric}(K) := \sup\{\kappa \in \mathbb{R} : K \text{ has Ricci curvature bounded from below by } \kappa.\}$$

Calculating or estimating $\text{Ric}(K)$ in concrete situations does not appear to be an easy task. We shall address this topic in a forthcoming publication.

Several other approaches to Ricci curvature in a discrete setting have been considered recently.

Bonciocat and Sturm [5] adapted the definition based on displacement convexity of the entropy from [17, 26] to the discrete setting. The non-existence of geodesics in the L^2 -Wasserstein space is circumvented by considering approximate midpoints between measures in the L^2 -Wasserstein metric. Using this approach it is shown that certain planar graphs have non-negative Ricci curvature.

Ollivier [20, 21] defined a notion of Ricci curvature by comparing transportation distances between small balls and their centers. This notion coincides with the usual notion of Ricci curvature lower boundedness on Riemannian manifolds and is very well adapted to study Ricci curvature on discrete spaces. In particular, it is easy to show that the Ricci curvature of the n -dimensional discrete hypercube is proportional to $\frac{1}{n}$. However, as has been discussed in [22], the relation with displacement convexity remains to be clarified.

Very recently Y. Lin and S.-T. Yau [16] studied Ricci curvature on graphs by taking a characterisation in terms of the heat semigroup due to Bakry and Emery as a definition. With this definition it is shown that the Ricci curvature on locally finite graphs is bounded from below by -1 .

Structure of the paper. Section 2 contains a detailed analysis of the metric \mathcal{W} associated with Markov kernels on a two-point space. In section 3 we study the metric \mathcal{W} in a general setting and prove Theorem 1.1. In Section 4 we study gradient flows and present the proof of Theorem 1.2.

Note added. After completion of this paper, the author has been informed about the recent preprint [7] where a related class of metrics has been studied independently. The results obtained in both papers are largely complementary.

Acknowledgement. The author is grateful to Matthias Erbar, Nicola Gigli, Nicolas Juillet, Giuseppe Savaré, and Karl-Theodor Sturm for stimulating discussions on this paper and related topics.

2. ANALYSIS ON THE TWO-POINT SPACE

In this section we shall carry out a detailed analysis of the metric \mathcal{W} in the simplest case of interest, where the underlying space is a two-point space, say $\mathcal{X} = \mathcal{Q}^1 = \{a, b\}$. The reason for discussing the two-point space separately is twofold. Firstly, it is possible to perform explicit calculations, which lead to simple proofs and more precise results than in the general case. Secondly, some of the results obtained in this section shall be used in Section 3, where results for more general Markov chains are obtained by comparison arguments involving Markov chains on a two-point space.

Markov chains on the two-point space. Consider a Markov kernel K with transition probabilities

$$K(a, b) = p, \quad K(b, a) = q \tag{2.1}$$

for some $p, q \in (0, 1]$. Then the associated continuous time semigroup $H(t) = e^{t(K-I)}$ is given by

$$H(t) = \frac{1}{p+q} \left(\begin{bmatrix} q & p \\ q & p \end{bmatrix} + e^{-(p+q)t} \begin{bmatrix} p & -p \\ -q & q \end{bmatrix} \right).$$

and the stationary distribution π satisfies

$$\pi(a) = \frac{q}{p+q}, \quad \pi(b) = \frac{p}{p+q}.$$

Since $K(a, b)\pi(a) = K(b, a)\pi(b)$, we observe that K is reversible. Every probability measure on \mathcal{Q}^1 is of the form $\frac{1}{2}((1-\beta)\delta_a + (1+\beta)\delta_b)$ for some $\beta \in [-1, 1]$. The corresponding density ρ^β with respect to π is then given by

$$\rho^\beta(a) := \frac{p+q}{q} \frac{1-\beta}{2}, \quad \rho^\beta(b) := \frac{p+q}{p} \frac{1+\beta}{2}.$$

It follows that $H(t)\rho^\beta = \rho^{\beta_t}$ where

$$\beta_t := \frac{p-q}{p+q} (1 - e^{-(p+q)t}) + \beta e^{-(p+q)t}, \quad (2.2)$$

thus β solves the differential equation

$$\dot{\beta}_t = p(1 - \beta_t) - q(1 + \beta_t). \quad (2.3)$$

Remark 2.1 (Limitations of the L^2 -Wasserstein distance). Before introducing a new class of (pseudo)-metrics on $\mathcal{P}(\mathcal{Q}^1)$, we shall argue why the L^2 -Wasserstein metric W_2 is not appropriate for the purposes of this paper. First we shall show that – as we already mentioned in the introduction – the metric derivative of the heat flow is infinite with respect to the L^2 -Wasserstein metric. To see this, take $\beta \in [-1, 1] \setminus \{\frac{p-q}{p+q}\}$, and let $u(t) := H(t)\rho^\beta = \rho^{\beta_t}$ be the heat flow starting at ρ^β . Since $W_2(\rho^\alpha, \rho^\beta) = \sqrt{2|\beta - \alpha|}$ for $\alpha, \beta \in [-1, 1]$, we have

$$\begin{aligned} |u|(t) &:= \limsup_{s \rightarrow t} \frac{W_2(u(t), u(s))}{|t-s|} = \sqrt{2} \limsup_{s \rightarrow t} \frac{\sqrt{|\beta_t - \beta_s|}}{|t-s|} \\ &= \sqrt{2 \left| \beta - \frac{p-q}{p+q} \right|} \limsup_{s \rightarrow t} \frac{\sqrt{|e^{-(p+q)t} - e^{-(p+q)s}|}}{|t-s|} = +\infty. \end{aligned}$$

In particular, the heat flow is *not* a curve of maximal slope (see, e.g., [1] for this concept of gradient flow) for any functional on $\mathcal{P}(\mathcal{Q}^1)$.

Furthermore, the Lott-Sturm-Villani definition of Ricci curvature [17, 26] cannot be applied in the discrete setting, since W_2 -geodesics between distinct elements of $\mathcal{P}(\mathcal{Q}_1)$ do not exist. To see this, let $\{\rho^{\beta(t)}\}_{0 \leq t \leq 1}$ be a constant speed geodesic in $\mathcal{P}(\mathcal{Q}^1)$. For $s, t \in [0, 1]$ we then have

$$\begin{aligned} \sqrt{2|\beta(t) - \beta(s)|} &= W_2(\rho^{\beta(t)}, \rho^{\beta(s)}) \\ &= |t-s| W_2(\rho^{\beta(0)}, \rho^{\beta(1)}) = |t-s| \sqrt{2|\beta(0) - \beta(1)|}, \end{aligned}$$

which implies that $t \mapsto \beta(t)$ is 2-Hölder, hence constant on $[0, 1]$. It thus follows that all constant speed W_2 -geodesics are constant.

A new metric. Given a fixed Markov chain K on $\{a, b\}$ we shall define a (pseudo-)metric \mathcal{W} on $\mathcal{P}(\{a, b\})$ that depends on the choice of a function $\theta : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$. The following assumptions will be in force throughout this section:

Assumption 2.2. *The function $\theta : [0, \infty) \times [0, \infty) \rightarrow [0, \infty)$ has the following properties:*

- (A1) θ is continuous on $[0, \infty) \times [0, \infty)$;
- (A2) θ is C^∞ on $(0, \infty) \times (0, \infty)$;
- (A3) $\theta(s, t) = \theta(t, s)$ for $s, t \geq 0$;
- (A4) $\theta(s, t) > 0$ for $s, t > 0$.

The most interesting choice for the purposes of this paper is the case where θ is the logarithmic mean defined by $\theta(s, t) := \int_0^1 s^{1-pt} t^p dp$.

To simplify notation we define, for $\beta \in [-1, 1]$,

$$\widehat{\rho}(\beta) = \theta(\rho^\beta(a), \rho^\beta(b)) .$$

On the two-point space the variational definition of \mathcal{W} given in the introduction can be simplified as follows:

Lemma 2.3. *For $\alpha, \beta \in [-1, 1]$ we have*

$$\mathcal{W}(\rho^\alpha, \rho^\beta)^2 = \inf_{\gamma} \left\{ \frac{p+q}{4pq} \int_0^1 \frac{\dot{\gamma}_t^2}{\widehat{\rho}(\gamma_t)} \mathbf{1}_{\{\widehat{\rho}(\gamma_t) > 0\}} dt \right\}, \quad (2.4)$$

where the infimum runs over all piecewise C^1 -functions $\gamma : [0, 1] \rightarrow [-1, 1]$.

Proof. Substituting $\chi(t) = \psi_t(b) - \psi_t(a)$ in the definition of \mathcal{W} , one obtains

$$\mathcal{W}(\rho^\alpha, \rho^\beta)^2 = \inf_{\gamma, \chi} \left\{ \frac{pq}{p+q} \int_0^1 \widehat{\rho}(\gamma_t) \chi_t^2 dt \right\},$$

where the infimum runs over all piecewise C^1 -functions $\gamma : [0, 1] \rightarrow [-1, 1]$ and all measurable functions $\chi : [0, 1] \rightarrow \mathbb{R}$ satisfying $\gamma_0 = \alpha$, $\gamma_1 = \beta$ and

$$\dot{\gamma}_t = \frac{2pq}{p+q} \widehat{\rho}(\gamma_t) \chi_t .$$

The result follows by inserting the latter constraint in the expression for $\mathcal{W}(\rho^\alpha, \rho^\beta)$. \square

Lemma 2.3 provides a representation of $\mathcal{W}(\rho^\alpha, \rho^\beta)$ in terms of a one-dimensional variational problem. Note that some care needs to be taken when solving this problem, since for some choices of θ (including the logarithmic mean) the denominator in (2.4) tends to 0 as β_t tends to ± 1 . The following result provides an explicit formula for \mathcal{W} :

Theorem 2.4. *For $-1 \leq \alpha \leq \beta \leq 1$ we have*

$$\mathcal{W}(\rho^\alpha, \rho^\beta) = \frac{1}{2} \sqrt{\frac{1}{p} + \frac{1}{q}} \int_{\alpha}^{\beta} \frac{1}{\sqrt{\widehat{\rho}(r)}} dr \in [0, \infty] .$$

Proof. Suppose first that α and β belong to $(-1, 1)$. (If $\widehat{\rho}$ is bounded away from 0, this distinction is not necessary.) It is easily checked that the infimum in (2.4) may be restricted to monotone functions β . Since $g : r \mapsto \frac{1}{\widehat{\rho}(r)}$ is bounded on compact intervals in $(-1, 1)$, (2.4) reduces to an elementary

one-dimensional variational problem, which admits a minimizer, say ξ , that solves the Euler-Lagrange equation

$$2\ddot{\xi}g(\xi_t) + \dot{\xi}_t^2 g'(\xi_t) = 0 .$$

This equation implies that $t \mapsto \dot{\xi}_t \sqrt{g(\xi_t)}$ is constant, say equal to C . Since $\alpha \leq \beta$, it follows that $C > 0$. We infer that

$$\mathcal{W}(\rho^\alpha, \rho^\beta)^2 = \frac{p+q}{4pq} \int_0^1 \frac{\dot{\xi}_t^2}{\widehat{\rho}(\xi_t)} dt = \frac{p+q}{4pq} C^2 .$$

Moreover, ξ is monotone, hence invertible. It follows from the inverse function theorem that its inverse $\gamma : [\alpha, \beta] \rightarrow [0, 1]$ satisfies $\gamma'(r) = C^{-1} \sqrt{g(r)}$. We thus obtain

$$1 = \gamma(\beta) - \gamma(\alpha) = \int_\alpha^\beta \gamma'(r) dr = C^{-1} \int_\alpha^\beta \sqrt{g(r)} dr ,$$

hence

$$\mathcal{W}(\rho^\alpha, \rho^\beta) = \frac{C}{2} \sqrt{\frac{1}{p} + \frac{1}{q}} = \frac{1}{2} \sqrt{\frac{1}{p} + \frac{1}{q}} \int_\alpha^\beta \sqrt{g(r)} dr ,$$

which implies the desired identity.

The general case $-1 \leq \alpha \leq \beta \leq 1$ follows from a straightforward continuity argument. \square

For $\beta \in [-1, 1]$ it will be useful to define

$$\varphi(\beta) := \frac{1}{2} \sqrt{\frac{1}{p} + \frac{1}{q}} \int_0^\beta \frac{1}{\sqrt{\widehat{\rho}(r)}} dr \in [-\infty, \infty] , \quad (2.5)$$

so that Theorem 2.4 implies that

$$\mathcal{W}(\rho^\alpha, \rho^\beta) = |\varphi(\alpha) - \varphi(\beta)|$$

for $\alpha, \beta \in [-1, 1]$. It follows from the assumptions on θ that φ is real-valued, continuous and strictly increasing on $(-1, 1)$. Moreover, $\varphi(\pm 1) = \lim_{\beta \rightarrow \pm 1} \varphi(\beta)$ is possibly $\pm\infty$, depending on the behaviour of θ near 0.

In order to avoid having to distinguish between several cases in the results below, we set

$$(-1, 1)_* = \{\beta \in [-1, 1] : |\varphi(\beta)| < \infty\} , \quad I = \{\varphi(\beta) : \beta \in (-1, 1)_*\} ,$$

and

$$\mathcal{P}_1(\mathcal{Q}^1) := \{\rho^\beta \in \mathcal{P}(\mathcal{Q}^1) : \beta \in (-1, 1)_*\} .$$

It follows from the remarks above that $(-1, 1) \subseteq (-1, 1)_* \subseteq [-1, 1]$ and that I is a (possibly infinite) closed interval in \mathbb{R} . The following result, which summarises this discussion, is now obvious:

Proposition 2.5. *The function \mathcal{W} defines a pseudo-metric on $\mathcal{P}(\mathcal{Q}^1)$ that restricts to a metric on $\mathcal{P}_1(\mathcal{Q}^1)$. The mapping*

$$J : \rho^\beta \mapsto \varphi(\beta)$$

defines an isometry from $(\mathcal{P}_1(\mathcal{Q}^1), \mathcal{W})$ onto I endowed with the euclidean metric. In particular, $(\mathcal{P}_1(\mathcal{Q}^1), \mathcal{W})$ is complete.

The most interesting case for the purposes of this paper is the following:

Example 2.6 (Logarithmic mean). If θ is the logarithmic mean, i.e., $\theta(s, t) = \int_0^1 s^{1-r} t^r dr$, then $\widehat{\rho}(-1) = \widehat{\rho}(1) = 0$ and for $\beta \in (-1, 1)$ we have

$$\widehat{\rho}(\beta) = \frac{p+q}{2pq} \frac{q(1+\beta) - p(1-\beta)}{\log q(1+\beta) - \log p(1-\beta)}.$$

In this case we have $(-1, 1)_* = [-1, 1]$ and $I = [\varphi(-1), \varphi(1)]$ is a compact interval. Furthermore, for $-1 \leq \alpha \leq \beta \leq 1$,

$$\mathcal{W}(\rho^\alpha, \rho^\beta) = \frac{1}{\sqrt{2}} \int_\alpha^\beta \sqrt{\frac{\log q(1+r) - \log p(1-r)}{q(1+r) - p(1-r)}} dr.$$

If moreover $p = q$, we have

$$\widehat{\rho}(\beta) = \frac{\beta}{\operatorname{arctanh} \beta}.$$

and

$$\mathcal{W}(\rho^\alpha, \rho^\beta) = \frac{1}{\sqrt{2p}} \int_\alpha^\beta \sqrt{\frac{\operatorname{arctanh} r}{r}} dr.$$

Recall that a *constant speed geodesic* in a metric space (M, d) is a curve $u : [0, 1] \rightarrow M$ satisfying

$$d(u(s), u(t)) = |t - s|d(u(0), u(1))$$

for all $s, t \in [0, 1]$.

The next result gives a characterisation of \mathcal{W} -geodesics in $\mathcal{P}_1(\mathcal{Q}^1)$.

Proposition 2.7 (Characterisation of geodesics). *Let $\rho, \sigma \in \mathcal{P}_1(\mathcal{Q}^1)$. There exists a unique constant speed geodesic $\{\rho^{\gamma(t)}\}_{0 \leq t \leq 1}$ in $\mathcal{P}_1(\mathcal{Q}^1)$ with $\rho^{\gamma(0)} = \rho$ and $\rho^{\gamma(1)} = \sigma$. Moreover, the function γ belongs to $C^1([0, 1]; \mathbb{R})$ and satisfies the differential equation*

$$\gamma'(t) = 2w \sqrt{\frac{pq}{p+q} \widehat{\rho}(\gamma(t))} \quad (2.6)$$

for $t \in [0, 1]$, where $w := \operatorname{sgn}(\beta - \alpha) \mathcal{W}(\rho^\alpha, \rho^\beta)$.

Proof. Since the mapping J is an isometry from $\mathcal{P}_1(\mathcal{Q}^1)$ onto I , existence and uniqueness of geodesics follow directly from the corresponding facts in I .

Take now $\alpha, \beta \in (-1, 1)_*$ and let $\gamma \in C^1([0, 1]; \mathbb{R})$ be the solution to (2.6) with initial condition $\gamma(0) = \alpha$. For $0 \leq s < t \leq 1$ we then obtain by (2.5),

$$\varphi(\gamma(t)) - \varphi(\gamma(s)) = \int_s^t \varphi'(\gamma(r)) \gamma'(r) dr = w(t - s),$$

which implies that $\mathcal{W}(\rho^{\gamma(t)}, \rho^{\gamma(s)}) = |w|(t - s)$ and $\gamma(1) = \beta$, hence $t \mapsto \rho^{\gamma(t)}$ is a constant speed geodesic between ρ^α and ρ^β . \square

Gradient flows. In order to identify the heat flow as a gradient flow in $\mathcal{P}(\mathcal{Q}^1)$, we make the following assumption:

Assumption 2.8. *In addition to (A1) - (A4) we assume that there exists a function $f \in C([0, \infty); \mathbb{R}) \cap C^\infty((0, \infty); \mathbb{R})$ satisfying $f''(t) > 0$ for $t > 0$, and*

$$\theta(s, t) = \frac{s - t}{f'(s) - f'(t)}, \quad (2.7)$$

for all $s, t > 0$ with $s \neq t$.

Example 2.9. Note that this assumption is satisfied in Example 2.6 with $f(t) = t \log t$.

Consider the functional $\mathcal{F} : \mathcal{P}(\mathcal{Q}^1) \rightarrow \mathbb{R}$ defined by

$$\mathcal{F}(\rho) := \sum_{x \in \mathcal{Q}^1} f(\rho(x)) \pi(x)$$

where $f : \mathbb{R}_+ \rightarrow \mathbb{R}$ has been defined above. It thus follows that

$$\mathcal{F}(\rho^\beta) := \frac{q}{p+q} f(\rho^\beta(a)) + \frac{p}{p+q} f(\rho^\beta(b)). \quad (2.8)$$

Proposition 2.5 implies that $(\mathcal{P}_*(\mathcal{Q}^1), \mathcal{W})$ is a 1-dimensional Riemannian manifold. In particular, it makes sense to study gradient flows in $(\mathcal{P}_*(\mathcal{Q}^1), \mathcal{W})$.

Proposition 2.10 (Heat flow is the gradient flow of the entropy). *For $\beta \in [-1, 1]$ let $u : t \mapsto \rho^{\beta t} = H(t)\rho^\beta$ be the heat flow trajectory starting from ρ^β . Then u is a gradient flow trajectory of the functional \mathcal{F} in the Riemannian manifold $(\mathcal{P}_*(\mathcal{Q}^1), \mathcal{W})$.*

Proof. Recall that the function $J : \rho^\beta \mapsto \varphi(\beta)$ maps $\mathcal{P}_1(\mathcal{Q}^1)$ isometrically onto a closed interval $I \subseteq \mathbb{R}$. Therefore it suffices to show that the gradient flow equation

$$\frac{d}{dt} \varphi(\beta_t) = -\tilde{\mathcal{F}}'(\varphi(\beta_t)) \quad (2.9)$$

holds for $t > 0$, where $\tilde{\mathcal{F}} := \mathcal{F} \circ J^{-1}$.

To prove this, we set

$$c_{pq} := \frac{1}{2} \sqrt{\frac{1}{p} + \frac{1}{q}}, \quad \ell(\beta) := \rho^\beta(a), \quad r(\beta) := \rho^\beta(b),$$

for brevity. Using (2.5) and (2.7) we obtain

$$\varphi'(\beta) = \frac{c_{pq}}{\sqrt{\tilde{\rho}(\beta)}} = c_{pq} \sqrt{\frac{f'(r(\beta)) - f'(\ell(\beta))}{r(\beta) - \ell(\beta)}}. \quad (2.10)$$

Since

$$\tilde{\mathcal{F}}(\varphi(\beta)) = \tilde{\mathcal{F}}(J(\rho^\beta)) = \mathcal{F}(\rho^\beta) = \frac{q}{p+q} f(\ell(\beta)) + \frac{p}{p+q} f(r(\beta)),$$

it follows that $\tilde{\mathcal{F}}$ is continuously differentiable on I and

$$\begin{aligned}\tilde{\mathcal{F}}'(\varphi(\beta)) &= \frac{f'(r(\beta)) - f'(\ell(\beta))}{2\varphi'(\beta)} \\ &= \frac{1}{2c_{pq}} \sqrt{(r(\beta) - \ell(\beta))(f'(r(\beta)) - f'(\ell(\beta)))}.\end{aligned}$$

On the other hand, (2.3) and (2.10) imply that

$$\begin{aligned}\frac{d}{dt}\varphi(\beta_t) &= (p(1 - \beta_t) - q(1 + \beta_t))\varphi'(\beta_t) \\ &= -\frac{1}{2c_{pq}^2}(r(\beta_t) - \ell(\beta_t))\varphi'(\beta_t) \\ &= -\frac{1}{2c_{pq}} \sqrt{(r(\beta_t) - \ell(\beta_t))(f'(r(\beta_t)) - f'(\ell(\beta_t)))}.\end{aligned}$$

Combining the latter two identities we obtain (2.9), which completes the proof. \square

In order to investigate the convexity of \mathcal{F} along \mathcal{W} -geodesics, we consider the function $K : (-1, 1) \rightarrow \mathbb{R}$ defined by

$$K(\beta) := \frac{p+q}{2} + \frac{1}{2}\hat{\rho}(\beta)(qf''(\rho^\beta(b)) + pf''(\rho^\beta(a)))$$

and

$$\kappa := \inf\{K(\beta) : \beta \in (-1, 1)\}. \quad (2.11)$$

Since $f'' > 0$, it follows that $\kappa \geq \frac{p+q}{2}$.

Remark 2.11. If $f(\rho) = \rho \log \rho$, straightforward calculus shows that

$$K(\beta) = \frac{p+q}{2} + \frac{1}{1-\beta^2} \frac{q(1+\beta) - p(1-\beta)}{\log q(1+\beta) - \log p(1-\beta)}$$

If moreover $p = q$, one has

$$K(\beta) = p \left(1 + \frac{1}{1-\beta^2} \frac{\beta}{\operatorname{arctanh} \beta} \right) \quad \text{and} \quad \kappa = 2p.$$

It turns out that κ determines the convexity of the functional \mathcal{F} :

Proposition 2.12 (Convexity of \mathcal{F} along \mathcal{W} -geodesics). *Let κ be defined by (2.11). The functional \mathcal{F} is κ -convex along geodesics. More explicitly, let $\bar{\rho}_0, \bar{\rho}_1 \in \mathcal{P}_1(\mathcal{Q}^1)$ and let $\{\rho_t\}_{0 \leq t \leq 1}$ be the unique constant speed geodesic satisfying $\rho_0 = \bar{\rho}_0$ and $\rho_1 = \bar{\rho}_1$. Then the inequality*

$$\mathcal{F}(\rho_t) \leq (1-t)\mathcal{F}(\rho_0) + t\mathcal{F}(\rho_1) - \frac{\kappa}{2}t(1-t)\mathcal{W}^2(\rho_0, \rho_1)$$

holds for all $t \in [0, 1]$.

Proof. Let $\alpha, \beta \in (-1, 1)_*$ be such that $\bar{\rho}_0 = \rho^\alpha$ and $\bar{\rho}_1 = \rho^\beta$ and set $w := \mathcal{W}(\rho^\alpha, \rho^\beta)$. Without loss of generality we assume that $\alpha \leq \beta$. Proposition 2.7 implies that $\rho_t = \rho^{\gamma(t)}$, where γ satisfies (2.6).

Set $\zeta(t) := \mathcal{F}(\rho_t)$. It suffices to show that $\zeta''(t) \geq w^2\kappa$ for $t \in [0, 1]$. By (2.8) we have

$$\zeta'(t) = \frac{1}{2}\gamma'(t)(f'(\rho^{\gamma(t)}(b)) - f'(\rho^{\gamma(t)}(a))),$$

and therefore (2.6) implies that

$$\zeta'(t) = w \sqrt{\frac{pq}{p+q}} \sqrt{(\rho^{\gamma(t)}(b) - \rho^{\gamma(t)}(a)) (f'(\rho^{\gamma(t)}(b)) - f'(\rho^{\gamma(t)}(a)))}.$$

Differentiating this identity and using (2.6) once more, we obtain

$$\zeta''(t) = w^2 K(\gamma(t)) \geq w^2 \kappa,$$

which completes the proof. \square

The question arises whether the metric \mathcal{W} constructed above is the *unique* geodesic metric on $\mathcal{P}(\mathcal{Q}^1)$ for which the heat flow is the gradient flow of the entropy. The answer is affirmative, provided that one requires that the left part $\{\rho^\beta : \beta < \bar{\beta}\}$ and the right part $\{\rho^\beta : \beta > \bar{\beta}\}$ of $\mathcal{P}_1(\mathcal{Q}^1)$ are patched together in a ‘reasonable’ way. Here $\bar{\beta} := \frac{p-q}{p+q}$, so that $\rho^{\bar{\beta}}$ corresponds to equilibrium. Such a condition is necessary, since the heat flow starting at ρ^β with $\beta > \bar{\beta}$ does not ‘see’ the measures ρ^α with $\alpha < \bar{\beta}$, and vice versa.

A precise uniqueness statement is given below. Since we shall not use this result elsewhere in the paper, we postpone its technical proof to Appendix B, where the notions of 2-absolute continuity and $\text{EVI}_0(\mathcal{F})$ are defined as well.

Proposition 2.13 (Uniqueness of the metric). *Let \mathcal{M} be a geodesic metric on $\mathcal{P}_1(\mathcal{Q}^1)$ with the following properties:*

- (1) *For $\beta \in (-1, 1)_*$, the heat flow $t \mapsto \rho^{\beta t}$ given by (2.2), is a 2-absolutely continuous curve satisfying $\text{EVI}_0(\mathcal{F})$.*
- (2) *For $\alpha, \beta \in (-1, 1)_*$ with $\alpha \leq \bar{\beta} \leq \beta$, we have*

$$\mathcal{M}(\rho^\alpha, \rho^\beta) = \mathcal{M}(\rho^\alpha, \rho^{\bar{\beta}}) + \mathcal{M}(\rho^{\bar{\beta}}, \rho^\beta).$$

Then $\mathcal{M} = \mathcal{W}$.

Note that (1) and (2) of Proposition 2.13 are satisfied if $\mathcal{M} = \mathcal{W}$. Indeed, since \mathcal{F} is convex by Proposition 2.12, (1) follows from [28, Proposition 23.1]. Furthermore (2) follows from the explicit expression for \mathcal{W} obtained in Theorem 2.4.

3. A WASSERSTEIN-LIKE METRIC FOR MARKOV CHAINS

In this section we consider a Markov kernel $K = (K(x, y))_{x, y \in \mathcal{X}}$ on a finite state space \mathcal{X} . We assume that K is irreducible, and denote its unique steady state by π . For all $x \in \mathcal{X}$ we then have $\pi(x) > 0$. We also assume that K is reversible, or equivalently, that the *detailed balance equations*

$$K(x, y)\pi(x) = K(y, x)\pi(y) \tag{3.1}$$

hold for all $x, y \in \mathcal{X}$.

Definition of the (pseudo-)metric. We start with the definition of a class of Wasserstein-like pseudo-metrics on $\mathcal{P}(\mathcal{X})$. As in Section 2, the metric depends on the choice of a function $\theta : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$, which we fix from now on. To simplify notation, we set

$$\rho(x, y) := \theta(\rho(x), \rho(y))$$

for $\rho \in \mathcal{P}(\mathcal{X})$ and $x, y \in \mathcal{X}$.

Assumption 3.1. *Throughout this section we shall assume that θ satisfies Assumption 2.2. In addition we impose the following assumptions:*

- (A5) *(Zero at the boundary): $\theta(0, t) = 0$ for all $t \geq 0$.*
- (A6) *(Monotonicity): $\theta(r, t) \leq \theta(s, t)$ for all $0 \leq r \leq s$ and $t \geq 0$.*
- (A7) *(Doubling property): for any $T > 0$ there exists a constant $C_d > 0$ such that*

$$\theta(2s, 2t) \leq 2C_d\theta(s, t)$$

whenever $0 \leq s, t \leq T$.

Remark 3.2. Actually, the additional assumptions (A5) – (A7) shall not be used until Theorem 3.12.

At some places, in particular in Lemmas 3.14 and 3.16 below, it is possible to obtain sharper results by imposing one or both of the following assumptions as well. Note that (A7') implies (A7).

(A7') *(Positive homogeneity): $\theta(\lambda s, \lambda t) = \lambda\theta(s, t)$ for $\lambda > 0$ and $s, t \geq 0$.*

(A8) *(Concavity): the function $\theta : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is concave.*

Observe that (A7') and (A8) hold if θ is the logarithmic mean.

Definition 3.3 (of the pseudo-metric \mathcal{W}). *For $\bar{\rho}_0, \bar{\rho}_1 \in \mathcal{P}(\mathcal{X})$ we define*

$$\mathcal{W}(\bar{\rho}_0, \bar{\rho}_1)^2 := \inf \left\{ \frac{1}{2} \int_0^1 \sum_{x, y \in \mathcal{X}} (\psi_t(x) - \psi_t(y))^2 K(x, y) \rho_t(x, y) \pi(x) dt : \right. \\ \left. (\rho, \psi) \in \mathcal{CE}_1(\bar{\rho}_0, \bar{\rho}_1) \right\},$$

where, for $T > 0$, $\mathcal{CE}_T(\bar{\rho}_0, \bar{\rho}_1)$ denotes the collection of pairs (ρ, ψ) satisfying the following conditions:

$$\left\{ \begin{array}{l} (i) \quad \rho : [0, T] \rightarrow \mathbb{R}^{\mathcal{X}} \text{ is piecewise } C^1 ; \\ (ii) \quad \rho_0 = \bar{\rho}_0, \quad \rho_T = \bar{\rho}_1 ; \\ (iii) \quad \rho_t \in \mathcal{P}(\mathcal{X}) \text{ for all } t \in [0, T] ; \\ (iv) \quad \psi : [0, T] \rightarrow \mathbb{R}^{\mathcal{X}} \text{ is measurable ;} \\ (v) \quad \text{For all } x \in \mathcal{X} \text{ and a.e. } t \in (0, T) \text{ we have} \\ \quad \dot{\rho}_t(x) + \sum_{y \in \mathcal{X}} (\psi_t(y) - \psi_t(x)) K(x, y) \rho_t(x, y) = 0 . \end{array} \right. \quad (3.2)$$

The latter equation may be thought of as a ‘continuity equation’. For simplicity we shall often write

$$\mathcal{CE}(\rho_0, \rho_1) := \mathcal{CE}_1(\rho_0, \rho_1) .$$

Remark 3.4 (Matrix reformulation). It will be very useful to reformulate Definition 3.3 in terms of matrices. For $\rho \in \mathcal{P}(\mathcal{X})$ consider the matrices $A(\rho)$ and $B(\rho)$ in $\mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ defined by

$$A_{x, y}(\rho) := \begin{cases} \sum_{z \neq x} K(x, z) \rho(x, z) \pi(x), & x = y, \\ -K(x, y) \rho(x, y) \pi(x), & x \neq y, \end{cases}$$

and

$$B_{x, y}(\rho) := \begin{cases} \sum_{z \neq x} K(x, z) \rho(x, z), & x = y, \\ -K(x, y) \rho(x, y), & x \neq y. \end{cases}$$

Definition 3.3 can then be rewritten as

$$\mathcal{W}(\bar{\rho}_0, \bar{\rho}_1)^2 = \inf \left\{ \int_0^1 [A(\rho_t)\psi_t, \psi_t] dt : (\rho, \psi) \in \mathcal{CE}(\bar{\rho}_0, \bar{\rho}_1) \right\}, \quad (3.3)$$

and the ‘continuity equation’ in (3.2) reads as

$$\dot{\rho}_t = B(\rho_t)\psi_t. \quad (3.4)$$

Here and in the sequel we use square brackets $[\cdot, \cdot]$ to denote the standard inner product in $\mathbb{R}^{\mathcal{X}}$. It follows from the detailed balance equations (3.1) that $A(\rho)$ is symmetric, but $B(\rho)$ is not necessarily symmetric. Since $\sum_{y \neq x} |A_{x,y}(\rho)| = A_{x,x}(\rho) \geq 0$ for all $x \in \mathcal{X}$, the matrix $A(\rho)$ is diagonally dominant, which implies that

$$[A(\rho)\psi, \psi] \geq 0 \quad (3.5)$$

for all $\psi \in \mathbb{R}^{\mathcal{X}}$. Note that

$$A(\rho) = \Pi B(\rho),$$

where the diagonal matrix $\Pi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ is defined by

$$\Pi := \text{diag}(\pi(x))_{x \in \mathcal{X}}$$

Geometric interpretation. Before continuing we present another, more geometric reformulation of Definition 3.3 which makes the connection to the Benamou-Brenier formula 1.2 (even) more apparent. We introduce some notation that will be used throughout the remainder of the paper.

For $\psi \in \mathbb{R}^{\mathcal{X}}$ we consider the discrete gradient $\nabla\psi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ defined by

$$\nabla\psi(x, y) := \psi(x) - \psi(y),$$

and for $\Psi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ we consider the divergence $\nabla \cdot \Psi \in \mathbb{R}^{\mathcal{X}}$ defined by

$$(\nabla \cdot \Psi)(x) := \frac{1}{2} \sum_{y \in \mathcal{X}} K(x, y)(\Psi(y, x) - \Psi(x, y)) \in \mathbb{R}.$$

It is easily checked that the ‘integration by parts formula’ holds:

$$\langle \nabla\psi, \Psi \rangle_{\pi} = -\langle \psi, \nabla \cdot \Psi \rangle_{\pi},$$

where, for $\varphi, \psi \in \mathbb{R}^{\mathcal{X}}$ and $\Phi, \Psi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$,

$$\begin{aligned} \langle \varphi, \psi \rangle_{\pi} &= \sum_{x \in \mathcal{X}} \varphi(x)\psi(x)\pi(x), \\ \langle \Phi, \Psi \rangle_{\pi} &= \frac{1}{2} \sum_{x, y \in \mathcal{X}} \Phi(x, y)\Psi(x, y)K(x, y)\pi(x). \end{aligned}$$

Furthermore, for $\rho \in \mathcal{P}(\mathcal{X})$ we write

$$\begin{aligned} \langle \Phi, \Psi \rangle_{\rho} &:= \frac{1}{2} \sum_{x, y \in \mathcal{X}} \Phi(x, y)\Psi(x, y)K(x, y)\rho(x, y)\pi(x), \\ \|\Phi\|_{\rho} &:= \sqrt{\langle \Phi, \Phi \rangle_{\rho}}, \end{aligned} \quad (3.6)$$

and note that $\langle \cdot, \cdot \rangle_{\pi} = \langle \cdot, \cdot \rangle_{\rho}$ if $\rho(x) = 1$ for all $x \in \mathcal{X}$.

For a probability density $\rho \in \mathcal{P}(\mathcal{X})$ and $x \in \mathcal{X}$ we consider the matrix $\widehat{\rho} \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$ defined by

$$\widehat{\rho}(x, y) := \rho(x, y) .$$

Given two matrices $M, N \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$, let $M \bullet N$ denote their entrywise product defined by

$$(M \bullet N)(x, y) := M(x, y)N(x, y)$$

The definition of \mathcal{W} can now be reformulated as follows:

Lemma 3.5 (Geometric reformulation). *For $\bar{\rho}_0, \bar{\rho}_1 \in \mathcal{P}(\mathcal{X})$ we have*

$$\mathcal{W}(\bar{\rho}_0, \bar{\rho}_1)^2 = \inf_{\rho, \psi} \left\{ \int_0^1 \|\nabla \psi_t\|_{\rho_t}^2 dt : (\rho, \psi) \in \mathcal{CE}(\bar{\rho}_0, \bar{\rho}_1) \right\} ,$$

and the differential equation in (3.2) can be rewritten as

$$\dot{\rho}_t + \nabla \cdot (\widehat{\rho}_t \bullet \nabla \psi_t) = 0 . \quad (3.7)$$

Proof. This follows directly from the definitions. \square

For the L^2 -Wasserstein metric on Euclidean space, it is well known that one can take the infimum in the Benamou-Brenier formula (1.2) over all vector fields $\Psi : \mathbb{R}^n \rightarrow \mathbb{R}^n$, rather than only considering gradients $\Psi = \nabla \psi$. In order to formulate a similar result in the discrete setting, we replace (iv) and (v) in (3.2) by

$$\begin{aligned} (iv') \quad & \Psi : [0, T] \rightarrow \mathbb{R}^{\mathcal{X} \times \mathcal{X}} \text{ is measurable ;} \\ (v') \quad & \text{For all } x \in \mathcal{X} \text{ and a.e. } t \in (0, T) \text{ we have} \\ & \dot{\rho}_t(x) + \frac{1}{2} \sum_{y \in \mathcal{X}} (\Psi_t(x, y) - \Psi_t(y, x)) K(x, y) \rho_t(x, y) = 0 ; \end{aligned} \quad (3.8)$$

and define

$$\mathcal{CE}'_T(\bar{\rho}_0, \bar{\rho}_1) := \{(\rho, \Psi) : (i), (ii), (iii), (iv'), (v') \text{ hold}\} .$$

With this notation the following result holds.

Lemma 3.6. *For $\bar{\rho}_0, \bar{\rho}_1 \in \mathcal{P}(\mathcal{X})$ we have*

$$\mathcal{W}(\bar{\rho}_0, \bar{\rho}_1)^2 = \inf \left\{ \frac{1}{2} \int_0^1 \sum_{x, y \in \mathcal{X}} \Psi_t(x, y)^2 K(x, y) \rho_t(x, y) \pi(x) dt : \right. \\ \left. (\rho, \Psi) \in \mathcal{CE}'_1(\bar{\rho}_0, \bar{\rho}_1) \right\} .$$

Proof. As the inequality “ \geq ” is trivial, it suffices to prove the inequality “ \leq ”. For this purpose, fix $\rho \in \mathcal{P}(\mathcal{X})$ and let \mathcal{H}_ρ denote the set of all equivalence classes of functions $\Psi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$, where we identify functions that agree on $\{(x, y) \in \mathcal{X} \times \mathcal{X} : \rho(x, y) K(x, y) > 0\}$. Endowed with the inner product $\langle \cdot, \cdot \rangle_\rho$ defined in (3.6), \mathcal{H}_ρ is a finite-dimensional Hilbert space. The discrete gradient $\nabla \varphi(x, y) := \varphi(x) - \varphi(y)$ defines a linear operator $\nabla : L^2(\mathcal{X}, \pi) \rightarrow \mathcal{H}_\rho$, whose adjoint is given by

$$\nabla_\rho^* \Psi(x) := \frac{1}{2} \sum_{y \in \mathcal{X}} (\Psi(x, y) - \Psi(y, x)) K(x, y) \rho(x, y) . \quad (3.9)$$

Let P_ρ denote the orthogonal projection in \mathcal{H}_ρ onto the range of ∇ .

Now suppose that $((\rho_t), (\Psi_t)) \in \mathcal{CE}'(\bar{\rho}_0, \bar{\rho}_1)$ and let $\psi : [0, 1] \rightarrow \mathbb{R}^{\mathcal{X}}$ be such that $P_{\rho_t} \Psi_t = \nabla \psi_t$ for $t \in [0, 1]$. In view of the orthogonal decomposition

$$\mathcal{H}_{\rho_t} = \text{Ran}(\nabla) \oplus^\perp \text{Ker}(\nabla_{\rho_t}^*), \quad (3.10)$$

it follows that $(I - P_{\rho_t})\Psi_t \in \text{Ker}(\nabla_{\rho_t}^*)$. This implies that $\nabla_{\rho_t}^* \Psi_t = \nabla_{\rho_t}^* (\nabla \psi_t)$, hence $(\rho, \psi) \in \mathcal{CE}(\bar{\rho}_0, \bar{\rho}_1)$. Using the decomposition (3.10) once more, we infer that $\langle \nabla \psi_t, \nabla \psi_t \rangle_{\rho_t} \leq \langle \Psi_t, \Psi_t \rangle_{\rho_t}$, from which the result follows. \square

Remark 3.7 (Distance between positive measures). It is of course possible, and occasionally useful, to extend the definition of $\mathcal{W}(\rho_0, \rho_1)$ to densities $\rho_0, \rho_1 : \mathcal{X} \rightarrow \mathbb{R}_+$ having equal mass $m = \sum_{x \in \mathcal{X}} \rho_i(x) \pi(x) \in (0, \infty) \setminus \{1\}$. A straightforward argument based on Lemma 3.6 and the doubling property (A7) shows that

$$c\mathcal{W}(\rho_0, \rho_1) \leq \mathcal{W}\left(\frac{1}{m}\rho_0, \frac{1}{m}\rho_1\right) \leq C\mathcal{W}(\rho_0, \rho_1),$$

where the constants $c, C > 0$ do not depend on ρ_0 and ρ_1 . If (A7') holds, it follows that $\mathcal{W}(\rho_0, \rho_1) = \sqrt{m}\mathcal{W}(\frac{1}{m}\rho_0, \frac{1}{m}\rho_1)$.

Basic properties of the metric. The main result of this subsection reads as follows:

Theorem 3.8. *The mapping $\mathcal{W} : \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{X}) \rightarrow \mathbb{R}$ defines a pseudo-metric on $\mathcal{P}(\mathcal{X})$.*

To prove this result we need some lemmas.

Lemma 3.9. *For $\bar{\rho}_0, \bar{\rho}_1 \in \mathcal{P}(\mathcal{X})$ and $T > 0$ we have*

$$\mathcal{W}(\bar{\rho}_0, \bar{\rho}_1) = \inf \left\{ \int_0^T [A(\rho_t)\psi_t, \psi_t]^{\frac{1}{2}} dt : (\rho, \psi) \in \mathcal{CE}_T(\bar{\rho}_0, \bar{\rho}_1) \right\}.$$

Proof. This follows from a standard argument based on parametrisation by arc-length. We refer to [1, Lemma 1.1.4] or [9, Theorem 5.4] for the details in a very similar situation. \square

The next lemma provides a lower bound for \mathcal{W} in terms of the total variation distance, defined for $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$ by

$$d_{TV}(\rho_0, \rho_1) = \sum_{x \in \mathcal{X}} \pi(x) |\rho_0(x) - \rho_1(x)|.$$

Lemma 3.10 (Lower bound by total variation distance). *For $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$ we have*

$$d_{TV}(\rho_0, \rho_1) \leq \sqrt{2\|\theta\|_\infty} \mathcal{W}(\rho_0, \rho_1),$$

where

$$\|\theta\|_\infty = \sup \left\{ \theta(s, t) : 0 \leq s, t \leq \left(\min_{x \in \mathcal{X}} \pi(x) \right)^{-1} \right\}.$$

Proof. We assume that $\mathcal{W}(\rho_0, \rho_1) < \infty$, since otherwise there is nothing to prove. Let $\varepsilon > 0$, let $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$ and take $(\rho, \psi) \in \mathcal{CE}(\rho_0, \rho_1)$ satisfying

$$\int_0^1 [A(\rho_t)\psi_t, \psi_t] dt < \mathcal{W}^2(\rho_0, \rho_1) + \varepsilon. \quad (3.11)$$

Using the continuity equation (3.4) we obtain for any $\varphi : \mathcal{X} \rightarrow \mathbb{R}$,

$$\begin{aligned} & \left| \sum_{x \in \mathcal{X}} \varphi(x) (\rho_0(x) - \rho_1(x)) \pi(x) \right| = \left| \int_0^1 [\Pi \varphi, \dot{\rho}_t] dt \right| \\ & = \left| \int_0^1 [\Pi \varphi, B(\rho_t) \psi_t] dt \right| = \left| \int_0^1 [A(\rho_t) \varphi, \psi_t] dt \right| \\ & \leq \left(\int_0^1 [A(\rho_t) \psi_t, \psi_t] dt \right)^{1/2} \left(\int_0^1 [A(\rho_t) \varphi, \varphi] dt \right)^{1/2}, \end{aligned}$$

where the appeal to the Cauchy-Schwarz inequality is justified by (3.5). The latter integrand can be estimated brutally by

$$\begin{aligned} [A(\rho_t) \varphi, \varphi] &= \frac{1}{2} \sum_{x, y \in \mathcal{X}} (\varphi(x) - \varphi(y))^2 K(x, y) \rho_t(x, y) \pi(x) \\ &\leq 2 \|\theta\|_\infty \|\varphi\|_\infty^2 \sum_{x, y \in \mathcal{X}} K(x, y) \pi(x) = 2 \|\theta\|_\infty \|\varphi\|_\infty^2, \end{aligned}$$

where we used the stationarity of π to obtain the latter identity. Taking (3.11) into account, and noting that $\varepsilon > 0$ is arbitrary, we thus obtain

$$\left| \sum_{x \in \mathcal{X}} \varphi(x) (\rho_0(x) - \rho_1(x)) \pi(x) \right| \leq \sqrt{2 \|\theta\|_\infty} \|\varphi\|_\infty \mathcal{W}(\rho_0, \rho_1).$$

Using the duality between $\ell^1(\mathcal{X})$ and $\ell^\infty(\mathcal{X})$, the result follows. \square

Proof of Theorem 3.8. The symmetry of \mathcal{W} is obvious, and Lemma 3.10 implies that $\mathcal{W}(\rho_0, \rho_1) > 0$ whenever $\rho_0 \neq \rho_1$. Finally, the triangle inequality easily follows using Lemma 3.9. \square

Characterisation of finiteness. In the study of finiteness of the metric \mathcal{W} , a crucial role will be played by the quantity

$$C_\theta := \int_0^1 \frac{1}{\sqrt{\theta(1-r, 1+r)}} dr \in [0, \infty].$$

Note that $C_\theta = \sqrt{2} \varphi(1)$, where φ denotes the function defined in (2.5) with $p = q = 1$. Therefore C_θ is finite if and only if Dirac measures on the two-point space lie at finite \mathcal{W} -distance from the uniform measure. Observe that $C_\theta < \infty$ if (A7') holds, since in that case

$$\theta(1-r, 1+r) \geq \theta(1-r, 1-r) = (1-r)\theta(1, 1),$$

for $r \in [0, 1)$.

The next result provides a characterisation of finiteness of the metric in terms of the support of the densities. For $\rho \in \mathcal{P}(\mathcal{X})$ we shall write

$$\text{supp } \rho := \{x \in \mathcal{X} : \rho(x) > 0\}.$$

Before stating the result we recall the following definition:

Definition 3.11. Let $\rho \in \mathcal{P}(\mathcal{X})$. For $x, y \in \mathcal{X}$ we write ' $x \sim_\rho y$ ' if

- (i) $x = y$; or,
- (ii) there exist $k \geq 1$ and $x_1, \dots, x_k \in \mathcal{X}$ such that

$$\rho(x, x_1)K(x, x_1), \rho(x_1, x_2)K(x_1, x_2), \dots, \rho(x_k, y)K(x_k, y) > 0.$$

It is easy to see that for each $\rho \in \mathcal{P}(\mathcal{X})$, \sim_ρ defines an equivalence relation on \mathcal{X} , which depends only on the support of ρ . Furthermore, if ρ is strictly positive, then $x \sim_\rho y$ for any $x, y \in \mathcal{X}$, since K is irreducible by assumption.

Now we are ready to state the main result of this subsection.

Theorem 3.12 (Characterisation of finiteness).

- (1) If $C_\theta < \infty$, then $\mathcal{W}(\rho_0, \rho_1) < \infty$ for all $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$.
- (2) If $C_\theta = \infty$, the following assertions are equivalent for $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$:
 - (a) $\mathcal{W}(\rho_0, \rho_1) < \infty$;
 - (b) For any $x \in \mathcal{X}$ we have

$$\sum_{y \sim_{\rho_0} x} \rho_0(y) \pi(y) = \sum_{y \sim_{\rho_1} x} \rho_1(y) \pi(y) . \quad (3.12)$$

Before turning to the proof of this result we record some immediate consequences:

Corollary 3.13. *Suppose that $C_\theta = \infty$. For $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$ the following assertions hold:*

- (1) If $\mathcal{W}(\rho_0, \rho_1) < \infty$, then $\text{supp } \rho_0 = \text{supp } \rho_1$.
- (2) If $\text{supp } \rho_0 = \text{supp } \rho_1 = \mathcal{X}$, then $\mathcal{W}(\rho_0, \rho_1) < \infty$.

Proof. (1) Suppose that $\rho_0(x) = 0$ for a certain $x \in \mathcal{X}$. In view of (A5) it then follows that $x \not\sim_{\rho_0} y$ for any $y \neq x$, hence by Theorem 3.12,

$$\rho_1(x) \pi(x) \leq \sum_{y \sim_{\rho_1} x} \rho_1(y) \pi(y) = \sum_{y \sim_{\rho_0} x} \rho_0(y) \pi(y) = \rho_0(x) \pi(x) = 0 .$$

It follows that $\rho_1(x) = 0$, which shows that $\text{supp } \rho_0 \supseteq \text{supp } \rho_1$. The reverse inclusion follows by reversing the roles of ρ_0 and ρ_1 .

(2) If $\text{supp } \rho_0 = \text{supp } \rho_1 = \mathcal{X}$, then $x \sim_{\rho_i} y$ for every $y \neq x$ and $i = 0, 1$ by irreducibility. It follows that

$$\sum_{y \sim_{\rho_0} x} \rho_0(y) \pi(y) = 1 = \sum_{y \sim_{\rho_1} x} \rho_1(y) \pi(y) ,$$

hence $\mathcal{W}(\rho_0, \rho_1) < \infty$ by Theorem 3.12. \square

The proof of Theorem 3.12 relies on a sequence of lemmas of independent interest.

First we prove two comparison results, which relate the pseudo-metric \mathcal{W} on $\mathcal{P}(\mathcal{X})$ to the pseudo-metric $\mathcal{W}_{p,q}$ on $\mathcal{P}(\mathcal{Y})$, where $\mathcal{Y} = \{a, b\}$ is a two-point space endowed with the Markov kernel (2.1) with parameters p and q .

Lemma 3.14 (Comparison to the two-point space I). *Let $a, b \in \mathcal{X}$ be distinct points with $K(a, b) > 0$, and set $p := K(a, b) \pi(a)$. Suppose that $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$ satisfy $\rho_0(x) = \rho_1(x)$ for all $x \in \mathcal{X} \setminus \{a, b\}$. Consider the two-point space $\mathcal{Q}^1 = \{\alpha, \beta\}$ endowed with the Markov kernel defined by $K(\alpha, \beta) := K(\beta, \alpha) := p$. For $i = 0, 1$, let $\bar{\rho}_i : \mathcal{Q}^1 \rightarrow \mathbb{R}_+$ be defined by*

$$\bar{\rho}_i(\alpha) := 2\rho_i(a) \pi(a) , \quad \bar{\rho}_i(\beta) := 2\rho_i(b) \pi(b) .$$

Then we have

$$\mathcal{W}(\rho_0, \rho_1) \leq \sqrt{C_d} \mathcal{W}_{p,p}(\bar{\rho}_0, \bar{\rho}_1) .$$

where C_d is the constant from (A7). In particular, if (A7') holds, then

$$\mathcal{W}(\rho_0, \rho_1) \leq \mathcal{W}_{p,p}(\bar{\rho}_0, \bar{\rho}_1).$$

Remark 3.15. Note that $\bar{\rho}_0$ and $\bar{\rho}_1$ are not necessarily probability densities on $\{\alpha, \beta\}$, but they do have equal mass, since

$$\bar{\rho}_i(\alpha)\pi(\alpha) + \bar{\rho}_i(\beta)\pi(\beta) = \rho_j(a)\pi(a) + \rho_j(b)\pi(b)$$

for $i, j \in \{0, 1\}$. Therefore $\mathcal{W}_{p,p}(\bar{\rho}_0, \bar{\rho}_1)$ can be interpreted in the sense of Remark 3.7.

Proof of Lemma 3.14. Let $\varepsilon > 0$ and take $(\bar{\rho}, \bar{\psi}) \in \mathcal{CE}(\bar{\rho}_0, \bar{\rho}_1)$. It then follows that

$$\begin{aligned} \dot{\bar{\rho}}_t(\alpha) + (\bar{\psi}_t(\beta) - \bar{\psi}_t(\alpha))K(\alpha, \beta)\bar{\rho}_t(\alpha, \beta) &= 0, \\ \dot{\bar{\rho}}_t(\beta) + (\bar{\psi}_t(\alpha) - \bar{\psi}_t(\beta))K(\beta, \alpha)\bar{\rho}_t(\alpha, \beta) &= 0. \end{aligned} \quad (3.13)$$

For $t \in (0, 1)$ define $\rho_t \in \mathcal{P}(\mathcal{X})$ by

$$\rho_t(a) := \frac{\bar{\rho}_t(\alpha)}{2\pi(a)}, \quad \rho_t(b) := \frac{\bar{\rho}_t(\beta)}{2\pi(b)}, \quad \rho_t(x) := \rho_0(x),$$

for $x \in \mathcal{X} \setminus \{a, b\}$. Furthermore, we define $\Psi_t : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ by

$$\begin{aligned} \Psi_t(a, b) &:= -\Psi_t(b, a) := \frac{\bar{\rho}_t(\alpha, \beta)}{2\rho_t(a, b)} (\bar{\psi}_t(\beta) - \bar{\psi}_t(\alpha)) \mathbf{1}_{\{\rho_t(a, b) > 0\}}, \\ \Psi_t(x, y) &:= 0, \end{aligned}$$

for all other values of $x, y \in \mathcal{X}$. Using (3.13) it then follows that $(\rho, \Psi) \in \mathcal{CE}'(\rho_0, \rho_1)$. Using Lemma 3.6 we thus obtain

$$\begin{aligned} \mathcal{W}(\rho_0, \rho_1)^2 &\leq \int_0^1 \Psi_t(a, b)^2 \rho_t(a, b) K(a, b) \pi(a) dt \\ &= \frac{1}{2} \int_0^1 (\bar{\psi}_t(\alpha) - \bar{\psi}_t(\beta))^2 \frac{\bar{\rho}_t(\alpha, \beta)^2}{\rho_t(a, b)} \mathbf{1}_{\{\rho_t(a, b) > 0\}} K(\alpha, \beta) \pi(\alpha) dt. \end{aligned}$$

From (A6) and (A7) we infer that

$$\bar{\rho}_t(\alpha, \beta) = \theta(2\pi(a)\rho_t(a), 2\pi(b)\rho_t(b)) \leq 2C_d\theta(\rho_t(a), \rho_t(b)) = 2C_d\rho_t(a, b),$$

which yields

$$\mathcal{W}(\rho_0, \rho_1)^2 \leq C_d \int_0^1 (\bar{\psi}_t(\alpha) - \bar{\psi}_t(\beta))^2 \bar{\rho}_t(\alpha, \beta) K(\alpha, \beta) \pi(\alpha) dt.$$

Minimising the right-hand side over all $(\bar{\rho}, \bar{\psi}) \in \mathcal{CE}(\bar{\rho}_0, \bar{\rho}_1)$, the result follows. \square

Lemma 3.16 (Comparison to the two-point space II). *Let $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$ and set $\beta_i(x) = 1 - 2\rho_i(x)\pi(x)$ for $i = 0, 1$ and $x \in \mathcal{X}$. Then the bound*

$$\mathcal{W}(\rho_0, \rho_1) \geq c \sup_{x \in \mathcal{X}} \mathcal{W}_{1,1}(\rho^{\beta_0(x)}, \rho^{\beta_1(x)})$$

holds, for some $c > 0$ depending only on K, π and θ . If (A7') and (A8) hold, then

$$\mathcal{W}(\rho_0, \rho_1) \geq \sup_{x \in \mathcal{X}} \mathcal{W}_{1,1}(\rho^{\beta_0(x)}, \rho^{\beta_1(x)}).$$

Proof. First we shall prove the result under the assumption that (A7') and (A8) hold. Fix $o \in \mathcal{X}$ and let $\mathcal{Y} = \{a, b\}$ be a two-point space endowed with the Markov kernel (2.1) with $p = q = 1$. For $\rho \in \mathcal{P}(\mathcal{X})$ and $\psi \in \mathbb{R}^{\mathcal{X}}$ we define, by a slight abuse of notation, $\rho \in \mathcal{P}(\mathcal{Y})$ and $\psi \in \mathbb{R}^{\mathcal{Y}}$ by

$$\begin{aligned} \rho(a) &:= 2\rho(o)\pi(o), & \rho(b) &:= 2 \sum_{x \neq o} \rho(x)\pi(x), \\ \psi(a) &:= \psi(o), & \psi(b) &:= \frac{\sum_{x \neq o} \psi(x)K(o, x)\rho(o, x)}{\sum_{x \neq o} K(o, x)\rho(o, x)}. \end{aligned}$$

In the definition of $\psi(b)$ we use the convention that $0/0 = 0$. Observe that ρ indeed belongs to $\mathcal{P}(\mathcal{Y})$ since $\pi(a) = \pi(b) = \frac{1}{2}$ and $\rho(a) + \rho(b) = 2$. We set $\tilde{\rho}(a, b) := 2\pi(o) \sum_{x \neq o} K(o, x)\rho(o, x)$ and claim that

$$\tilde{\rho}(a, b) \leq \rho(a, b), \quad (3.14)$$

$$[A(\rho)\psi, \psi] \geq \frac{1}{2}(\psi(a) - \psi(b))^2 \tilde{\rho}(a, b). \quad (3.15)$$

In the proof of both claims we shall assume that $\tilde{\rho}(a, b) > 0$, since otherwise there is nothing to prove. To prove (3.14), note first that for any $x \in \mathcal{X}$ with $K(o, x) > 0$,

$$\frac{\pi(x)}{\pi(o)} = \frac{K(o, x)}{K(x, o)} \geq K(o, x). \quad (3.16)$$

Using this inequality together with (A6), (A7') and (A8),

$$\begin{aligned} \rho(a, b) &= \theta \left(2\rho(o)\pi(o), 2 \sum_{x \neq o} \rho(x)\pi(x) \right) \\ &= 2\pi(o)\theta \left(\rho(o), \sum_{x \neq o} \rho(x) \frac{\pi(x)}{\pi(o)} \right) \\ &\geq 2\pi(o)\theta \left(\rho(o), \sum_{x \neq o} K(o, x)\rho(x) \right) \\ &\geq 2\pi(o) \sum_{x \neq o} K(o, x)\theta(\rho(o), \rho(x)) = \tilde{\rho}(a, b), \end{aligned}$$

which proves (3.14).

To prove (3.15), write $k(x) := K(o, x)\rho(o, x)$ for brevity and note that

$$\sum_{x \neq o} \psi(x)^2 k(x) \geq \frac{(\sum_{x \neq o} \psi(x)k(x))^2}{\sum_{x \neq o} k(x)} = \frac{\psi(b)^2 \tilde{\rho}(a, b)}{2\pi(o)}.$$

Using the detailed balance equations (3.1) in the first inequality, we obtain

$$\begin{aligned}
[A(\rho)\psi, \psi] &= \frac{1}{2} \sum_{x,y \in \mathcal{X}} (\psi(x) - \psi(y))^2 K(x,y) \rho(x,y) \pi(x) \\
&\geq \sum_{x \neq o} (\psi(o) - \psi(x))^2 K(o,x) \rho(o,x) \pi(o) \\
&= \left(\psi(o)^2 \sum_{x \neq o} k(x) - 2\psi(o) \sum_{x \neq o} \psi(x) k(x) + \sum_{x \neq o} \psi(x)^2 k(x) \right) \pi(o) \\
&\geq \frac{1}{2} \psi(a)^2 \tilde{\rho}(a,b) - \psi(a)\psi(b) \tilde{\rho}(a,b) + \frac{1}{2} \psi(b)^2 \tilde{\rho}(a,b) \\
&= \frac{1}{2} (\psi(a) - \psi(b))^2 \tilde{\rho}(a,b),
\end{aligned}$$

which proves (3.15).

Take $(\rho, \psi) \in \mathcal{CE}(\rho_0, \rho_1)$. Since

$$\dot{\rho}_t(o) + \sum_{x \neq o} (\psi_t(x) - \psi_t(o)) K(o,x) \rho_t(o,x) = 0,$$

it follows that

$$\dot{\rho}_t(a) + (\psi_t(b) - \psi_t(a)) \tilde{\rho}_t(a,b) = 0. \quad (3.17)$$

Set $\beta_t := 1 - 2\rho_t(o)\pi(o)$ for $t \in [0, 1]$ and note that $\dot{\beta}_t = 0$ if $\tilde{\rho}_t(a,b) = 0$. Using (3.15), (3.17), (3.14) and Lemma 2.3 we obtain

$$\begin{aligned}
\int_0^1 [A(\rho_t)\psi_t, \psi_t] dt &\geq \frac{1}{2} \int_0^1 (\psi_t(a) - \psi_t(b))^2 \tilde{\rho}_t(a,b) dt \\
&= \frac{1}{2} \int_0^1 \frac{\dot{\beta}_t^2 \mathbf{1}_{\{\tilde{\rho}_t(a,b) > 0\}}}{\tilde{\rho}_t(a,b)} dt \geq \frac{1}{2} \int_0^1 \frac{\dot{\beta}_t^2 \mathbf{1}_{\{\rho_t(a,b) > 0\}}}{\rho_t(a,b)} dt \\
&\geq \mathcal{W}_{1,1}^2(\rho^{\beta_0}, \rho^{\beta_1}).
\end{aligned}$$

Taking the infimum over all pairs $(\rho, \psi) \in \mathcal{CE}(\rho_0, \rho_1)$, we infer that

$$\mathcal{W}^2(\rho_0, \rho_1) \geq \mathcal{W}_{1,1}^2(\rho^{\beta_0}, \rho^{\beta_1}).$$

The result follows by taking the supremum over $o \in \mathcal{X}$.

Finally, without assuming (A7') and (A8), the same argument applies, if one replaces (3.14) by the following estimate, which uses the doubling property (A7), (3.16) and (A5):

$$\begin{aligned}
\tilde{\rho}(a,b) &= 2\pi(o) \sum_{x \neq o} K(o,x) \theta(\rho(o), \rho(x)) \\
&\leq C \sum_{x \neq o} \theta(2\rho(o)K(o,x)\pi(o), 2\rho(x)K(o,x)\pi(o)) \\
&\leq C \sum_{x \neq o} \theta(2\rho(o)\pi(o), 2\rho(x)\pi(x)) \\
&\leq C |\mathcal{X}| \theta\left(2\rho(o)\pi(o), 2 \sum_{x \neq o} \rho(x)\pi(x)\right) \\
&= C |\mathcal{X}| \rho(a,b).
\end{aligned}$$

□

The next lemma provides a useful characterisation of the kernel and the range of the matrices $A(\rho)$ and $B(\rho)$.

Lemma 3.17. *For $\rho \in \mathcal{P}(\mathcal{X})$ we have*

$$\begin{aligned} \text{Ker } A(\rho) &= \text{Ker } B(\rho) = \{ \psi \in \mathbb{R}^{\mathcal{X}} \mid \psi(x) = \psi(y) \text{ whenever } x \sim_{\rho} y \}, \\ \text{Ran } A(\rho) &= \left\{ \psi \in \mathbb{R}^{\mathcal{X}} \mid \forall x \in \mathcal{X} : \sum_{y \sim_{\rho} x} \psi(y) = 0 \right\}, \\ \text{Ran } B(\rho) &= \left\{ \psi \in \mathbb{R}^{\mathcal{X}} \mid \forall x \in \mathcal{X} : \sum_{y \sim_{\rho} x} \psi(y) \pi(y) = 0 \right\}. \end{aligned}$$

Proof. Recall that (A3) and (A5) imply that $\rho(x, y) = 0$ whenever $\rho(x) = 0$ or $\rho(y) = 0$. Therefore the assertions concerning $A(\rho)$ follow directly from Lemma A.1. Since $B(\rho) = \Pi^{-1}A(\rho)$, one has

$$\text{Ker } B(\rho) = \text{Ker } A(\rho), \quad \text{Ran } B(\rho) = \Pi^{-1} \text{Ran } A(\rho),$$

hence the remaining assertions follow as well. □

For $\sigma \in \mathcal{P}(\mathcal{X})$ and $a \geq 0$ we shall use the notation

$$\begin{aligned} \mathcal{P}_{\sigma}^a(\mathcal{X}) &:= \{ \rho \in \mathcal{P}(\mathcal{X}) \mid \forall x \in \mathcal{X} : (3.12) \text{ holds with } \rho_0 = \rho \text{ and } \rho_1 = \sigma; \\ &\quad \forall z \in \text{supp}(\sigma) : \rho(z) \geq a \}. \end{aligned}$$

Lemma 3.18. *For $\rho \in \mathcal{P}(\mathcal{X})$, $B(\rho)$ restricts to an isomorphism from $\text{Ran } A(\rho)$ onto $\text{Ran } B(\rho)$. Moreover, for $\sigma \in \mathcal{P}(\mathcal{X})$ and $a > 0$ there exist constants $0 < c < C < \infty$ such that the bound*

$$c\|\psi\| \leq \|B(\rho)\psi\| \leq C\|\psi\| \tag{3.18}$$

holds for all $\rho \in \mathcal{P}_{\sigma}^a(\mathcal{X})$ and all $\psi \in \text{Ran}(\sigma)$.

Proof. Since $A(\rho)$ is self-adjoint, $A(\rho)$ restricts to an isomorphism on its range. Since Π is an isomorphism from $\text{Ran } A(\rho)$ onto $\text{Ran } B(\rho)$ and $B(\rho) = \Pi^{-1}A(\rho)$, the first assertion follows.

Lemma 3.17 implies that $\text{Ran } A(\rho) = \text{Ran } A(\sigma)$ and $\text{Ran } B(\rho) = \text{Ran } B(\sigma)$ for all $\rho \in \mathcal{P}_{\sigma}^a(\mathcal{X})$. Thus $B(\rho)$ restricts to an isomorphism, denoted by B_{ρ} , from $\text{Ran } A(\sigma)$ onto $\text{Ran } B(\sigma)$. Since the mapping $\mathcal{P}_{\sigma}^a(\mathcal{X}) \ni \rho \mapsto \|B_{\rho}^{-1}\|$ is continuous w.r.t. the euclidean metric and strictly positive, the lower bound in (3.18) follows by compactness. The upper bound is clear, since the entries of $B(\rho)$ are bounded uniformly in ρ . □

The next result provides a partial converse to Lemma 3.10.

Lemma 3.19. *Fix $\sigma \in \mathcal{P}(\mathcal{X})$ and $a > 0$. There exist constants $0 < c < C < \infty$ such that for all $\rho_0, \rho_1 \in \mathcal{P}_{\sigma}^a(\mathcal{X})$ we have*

$$cd_{TV}(\rho_0, \rho_1) \leq \mathcal{W}(\rho_0, \rho_1) \leq Cd_{TV}(\rho_0, \rho_1).$$

Proof. Since the lower bound for \mathcal{W} has been proved in Lemma 3.10, it remains to prove the upper bound.

For $t \in [0, 1]$ set $\rho_t := (1-t)\rho_0 + t\rho_1$ and note that $\rho_t \in \mathcal{P}_{\sigma}^a(\mathcal{X})$. Since

$$\dot{\rho}_t = \rho_1 - \rho_0 \in \text{Ran } B(\rho_t) = \text{Ran } B(\sigma)$$

by Lemma 3.17, Lemma 3.18 implies that, for each $t \in [0, 1]$, there exists a unique element $\psi_t \in \text{Ran } A(\rho_t)$ satisfying

$$\dot{\rho}_t = B(\rho_t)\psi_t .$$

Moreover, Lemma 3.18 implies that

$$\|\psi_t\| \leq C\|\rho_1 - \rho_0\|$$

for some constant $C > 0$ that does not depend on ρ_0, ρ_1 and t . It thus follows that

$$\mathcal{W}(\rho_0, \rho_1)^2 \leq \int_0^1 [A(\rho_t)\psi_t, \psi_t] dt \leq C^2 C' \|\rho_1 - \rho_0\|^2 \leq C^2 C' C'' d_{TV}^2(\rho_0, \rho_1) ,$$

where $C' := \sup_{\rho \in \mathcal{P}(\mathcal{X})} \|A(\rho)\| < \infty$ and $C'' > 0$ depends only on π . \square

Now we are ready to prove the main result of this subsection.

Proof of Theorem 3.12. Since K is irreducible, (1) follows from Lemma 3.14, Remark 3.7 and the triangle inequality for \mathcal{W} .

The implication (b) \Rightarrow (a) of (2) follows from Lemma 3.19.

In order to prove the converse implication, we take $\rho_0, \rho_1 \in \mathcal{P}(\mathcal{X})$ with $\mathcal{W}(\rho_0, \rho_1) < \infty$ and claim that $\text{supp } \rho_0 = \text{supp } \rho_1$. Indeed, if the claim were false, then there would exist $x \in \mathcal{X}$ with $\rho_0(x) = 0$ and $\rho_1(x) > 0$ (or vice versa). Set $\beta = 1 - 2\pi(x)\rho_1(x)$ and note that $\beta \in [-1, 1]$. Lemma 3.16 implies that $\mathcal{W}(\rho_0, \rho_1) \geq c\mathcal{W}_{1,1}(\rho^1, \rho^\beta)$ for some $c > 0$. Since $C_\theta = \infty$, the right-hand side is infinite, which contradicts our assumption and thus proves the claim.

Let $(\rho, \psi) \in \mathcal{CE}(\rho_0, \rho_1)$ with $\int_0^1 [A(\rho_t), \psi_t, \psi_t] dt < \infty$. The claim implies that $\text{supp } \rho_0 = \text{supp } \rho_t$ for all $t \in [0, 1]$ and therefore $x \sim_{\rho_t} y$ if and only if $x \sim_{\rho_0} y$. Fix $z \in \text{supp } \rho_0$ and take $x \in \mathcal{X}$ with $x \sim_{\rho_0} z$. Since $K(x, y)\rho_t(x, y) = 0$ whenever $y \not\sim_{\rho_0} z$, we have

$$\dot{\rho}_t(x) + \sum_{y \sim_{\rho_t} z} (\psi_t(y) - \psi_t(x))K(x, y)\rho_t(x, y) = 0 .$$

Multiplying this identity by $\pi(x)$ and summing over $x \in \mathcal{X}$ with $x \sim_{\rho_t} z$, it follows using the detailed balance equations (3.1) that

$$\sum_{x \sim_{\rho_0} z} \dot{\rho}_t(x)\pi(x) = 0 ,$$

which implies (3.12). \square

Remark 3.20. Alternatively, the implication (b) \Rightarrow (a) in the proof of Theorem 3.12 can be proved as an application of Lemma 3.14.

We continue to prove the remaining parts of Theorem 1.1.

Theorem 3.21 (Topology). *Let $\sigma \in \mathcal{P}(\mathcal{X})$. For $\rho, \rho_\alpha \in \mathcal{P}_\sigma(\mathcal{X})$, the following assertions are equivalent:*

$$(1) \quad \lim_{\alpha} d_{TV}(\rho_\alpha, \rho) = 0 ; \quad (2) \quad \lim_{\alpha} \mathcal{W}(\rho_\alpha, \rho) = 0 .$$

Proof. It follows from Lemma 3.10 that (2) implies (1).

Conversely, suppose that (1) holds. If $C_\theta < \infty$, then (2) follows easily using Lemma 3.14. If $C_\theta = \infty$, there exists an index $\bar{\alpha}$ and a constant $b > 0$ such that ρ and ρ_α belong to $\mathcal{P}_\sigma^b(\mathcal{X})$ for every $\alpha \geq \bar{\alpha}$. Lemma 3.19 implies then that there exists a constant $C > 0$ such that

$$\mathcal{W}(\rho_\alpha, \rho) \leq C d_{TV}(\rho_\alpha, \rho)$$

for all $\alpha \geq \bar{\alpha}$, which yields the result. \square

Theorem 3.22 (Completeness). *For every $\sigma \in \mathcal{P}(\mathcal{X})$ the metric space $(\mathcal{P}_\sigma(\mathcal{X}), \mathcal{W})$ is complete.*

Proof. If $C_\theta < \infty$, this follows directly from Lemma 3.10 and Theorem 3.21. If $C_\theta = \infty$, take a sequence $(\rho_n)_n$ in $\mathcal{P}_\sigma(\mathcal{X})$ which is Cauchy with respect to \mathcal{W} . In particular, $(\rho_n)_n$ is bounded in the \mathcal{W} -metric, hence by Lemma 3.16 there exists a constant $a > 0$ such that ρ_n belongs to $\mathcal{P}_\sigma^a(\mathcal{X})$ for every n . By Lemma 3.10 $(\rho_n)_n$ is Cauchy in the total variation metric, hence ρ_n converges to some $\bar{\rho} \in \mathcal{P}(\mathcal{X})$ in total variation. Since $\mathcal{P}_\sigma^a(\mathcal{X})$ is a d_{TV} -closed subset of $\mathcal{P}(\mathcal{X})$, it follows that $\bar{\rho}$ belongs to $\mathcal{P}_\sigma^a(\mathcal{X})$. From Theorem 3.21 we then infer that ρ_n converges to $\bar{\rho}$ in \mathcal{W} -metric, which yields the desired result. \square

Riemannian structure. Fix a probability density $\sigma \in \mathcal{P}(\mathcal{X})$ and consider the space

$$\mathcal{P}'_\sigma(\mathcal{X}) := \left\{ \rho \in \mathcal{P}(\mathcal{X}) \mid \forall x \in \mathcal{X} : \sum_{y \sim_\rho x} \rho(y) \pi(y) = \sum_{y \sim_\sigma x} \sigma(y) \pi(y) \right\}.$$

Note that $\mathcal{P}'_1(\mathcal{X}) = \mathcal{P}_*(\mathcal{X})$ where $\mathbf{1}$ denotes the uniform density with respect to π . Moreover, if $C_\theta = \infty$, Theorem 3.12 implies that $\mathcal{P}'_\sigma(\mathcal{X}) = \mathcal{P}_\sigma(\mathcal{X})$ for all $\sigma \in \mathcal{P}(\mathcal{X})$.

Our next aim is to show that the metric space $(\mathcal{P}'_\sigma(\mathcal{X}), \mathcal{W})$ is a Riemannian manifold. First, we have the following result:

Proposition 3.23. *The metric space $(\mathcal{P}'_\sigma(\mathcal{X}), \mathcal{W})$ is a smooth manifold of dimension*

$$d(\sigma) := |\text{supp } \sigma| - n(\sigma),$$

where $|\text{supp } \sigma|$ is the cardinality of $\text{supp } \sigma$, and $n(\sigma)$ is the number of equivalence classes in the support of σ for the equivalence relation \sim_σ .

Proof. It follows from Theorem 3.12 and Lemma 3.17 that $\mathcal{P}'_\sigma(\mathcal{X})$ is a relatively open subset of the affine subspace

$$S_\sigma := \sigma + \text{Ran } B(\sigma) \subseteq \mathbb{R}^\mathcal{X}.$$

Theorem 3.21 implies that the topology induced by \mathcal{W} coincides with the euclidean topology on $\mathcal{P}'_\sigma(\mathcal{X})$, hence $(\mathcal{P}'_\sigma(\mathcal{X}), \mathcal{W})$ endowed with the euclidean structure is a smooth manifold.

The assertion concerning the dimension follows immediately, since $d(\sigma)$ is the dimension of $\text{Ran } B(\sigma)$. \square

Fix $\sigma \in \mathcal{P}(\mathcal{X})$ and $\rho \in \mathcal{P}'_\sigma(\mathcal{X})$. Since $\mathcal{P}'_\sigma(\mathcal{X})$ is an open subset of the affine space $\sigma + \text{Ran } B(\sigma)$, the tangent space of $\mathcal{P}'_\sigma(\mathcal{X})$ at ρ can be naturally identified with $\text{Ran } B(\sigma) = \text{Ran } B(\rho)$. Our next aim is to show that the tangent space can be identified with a space of gradients, in the spirit of the Otto calculus developed in [23]. In fact, we shall construct an isomorphism \mathcal{I}_ρ from $\text{Ran } B(\sigma)$ onto

$$T_\rho := \{\nabla\psi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}} : \psi \in \text{Ran } A(\rho)\}.$$

Remark 3.24. Note that if ρ belongs to $\mathcal{P}_*(\mathcal{X})$, we have

$$T_\rho = \{\nabla\psi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}} : \psi \in \mathbb{R}^{\mathcal{X}}\}.$$

However, it is easy to see that this is no longer true if $\rho \notin \mathcal{P}_*(\mathcal{X})$.

Proposition 3.25. *Let $\rho \in \mathcal{P}'_\sigma(\mathcal{X})$. The mapping*

$$\mathcal{I}_\rho : \text{Ran } B(\sigma) \rightarrow T_\rho, \quad B(\rho)\psi \mapsto \nabla\psi$$

defined for $\psi \in \text{Ran } A(\rho)$, is a linear isomorphism.

Proof. To show that \mathcal{I}_ρ is well-defined, consider the following mappings:

$$\begin{aligned} F_\rho : \text{Ran } A(\rho) &\rightarrow \text{Ran } B(\rho), & \psi &\mapsto B(\rho)\psi, \\ G : \text{Ran } A(\rho) &\rightarrow T_\rho, & \psi &\mapsto \nabla\psi. \end{aligned}$$

We claim that F_ρ and G are linear isomorphisms. Once this has been established, the proposition follows at once. The claim for F_ρ has been proved in Lemma 3.18. To prove the claim for G , suppose that $\nabla\psi = 0$ for some $\psi \in \text{Ran}(A)$. It then follows that

$$[A(\rho)\psi, \psi] = \langle \nabla\psi, \nabla\psi \rangle_\rho = 0,$$

Since $A(\rho)$ is symmetric and $\psi \in \text{Ran } A(\rho)$, it follows that $\psi = 0$, which completes the proof. \square

The following statement clarifies the connection with the Otto calculus in the continuous setting:

Proposition 3.26. *Let $\rho : [0, 1] \rightarrow \mathcal{P}'_\sigma(\mathcal{X})$ be differentiable at $t \in [0, 1]$. Then $\mathcal{I}_{\rho_t}\dot{\rho}_t$ is the unique element $\nabla\psi_t \in T_{\rho_t}$ satisfying the identity*

$$\dot{\rho}_t + \nabla \cdot (\hat{\rho}_t \bullet \nabla\psi_t) = 0.$$

Proof. Since $B(\rho)\psi = -\nabla \cdot (\hat{\rho} \bullet \nabla\psi)$ for $\rho \in \mathcal{P}(\mathcal{X})$ and $\psi \in \mathbb{R}^{\mathcal{X}}$, this is an immediate consequence of Proposition 3.25. \square

Henceforth we shall identify the tangent space of $\mathcal{P}'_\sigma(\mathcal{X})$ at ρ with T_ρ by means of the isomorphism \mathcal{I}_ρ .

Definition 3.27. *Let $\rho \in \mathcal{P}'_\sigma(\mathcal{X})$. We endow T_ρ with the inner product*

$$\langle \nabla\varphi, \nabla\psi \rangle_\rho = \frac{1}{2} \sum_{x, y \in \mathcal{X}} (\varphi(x) - \varphi(y))(\psi(x) - \psi(y))K(x, y)\rho(x, y)\pi(x),$$

defined for $\varphi, \psi \in \text{Ran } A(\rho)$.

Note that, for $\rho \in \mathcal{P}'_\sigma(\mathcal{X})$ and $\varphi, \psi \in \text{Ran } A(\rho)$

$$\langle \nabla\varphi, \nabla\psi \rangle_\rho = [A(\rho)\varphi, \psi]. \quad (3.19)$$

Remark 3.28. It is clear from the definition that $\langle \nabla\varphi, \nabla\psi \rangle_{\mathcal{P}}$ is well-defined. Moreover, (3.19) implies that if $\langle \nabla\psi, \nabla\psi \rangle_{\mathcal{P}} = 0$ for some $\psi \in \text{Ran } A(\rho)$, then $\psi = 0$, thus the expression indeed defines an inner product on T_{ρ} .

Theorem 3.29. *The following statements hold:*

- If $C_{\theta} < \infty$ and (A8) holds, then $(\mathcal{P}_*(\mathcal{X}), \mathcal{W})$ is a Riemannian manifold.
- If $C_{\theta} = \infty$, then $(\mathcal{P}'_{\sigma}(\mathcal{X}), \mathcal{W})$ is a complete Riemannian manifold for every $\sigma \in \mathcal{P}(\mathcal{X})$.

The Riemannian metric is given by Definition 3.27.

Proof. Suppose first that $C_{\theta} = \infty$. Then Proposition 3.23 asserts that $(\mathcal{P}_{\sigma}(\mathcal{X}), \mathcal{W})$ is a smooth manifold and the completeness has been proved in Theorem 3.22. The result would follow immediately from Lemma 3.5 and Definition 3.27, if we were allowed to add the following requirements to the definition of $\mathcal{CE}(\rho_0, \rho_1)$ without changing the value of $\mathcal{W}(\rho_0, \rho_1)$:

- (i) $\rho_t \in \mathcal{P}_{\sigma}(\mathcal{X})$ for all $t \in [0, 1]$;
- (ii) $\psi_t \in \text{Ran } A(\rho_t)$ for all $t \in [0, 1]$.

But (i) may be added by Theorem 3.12 and (ii) may be added in view of the orthogonal decomposition $\mathcal{X} = \text{Ran } A(\rho) \oplus \text{Ker } A(\rho)$.

If $C_{\theta} < \infty$ the same argument applies, with Lemma 3.30 below providing the analogue of (i). \square

The next result asserts that in the definition of \mathcal{W} , only curves consisting of strictly positive densities need to be considered if the endpoints are strictly positive as well.

Lemma 3.30. *Suppose that (A8) holds. For $\rho_0, \rho_1 \in \mathcal{P}_*(\mathcal{X})$, we may replace (iii) in Definition 3.3 by “(iii') : $\rho_t \in \mathcal{P}_*(\mathcal{X})$ for all $t \in [0, T]$ ”.*

Proof. For notational reasons, let us write

$$\mathcal{A}(\rho, \Psi) := \|\Psi\|_{\rho}^2 = \frac{1}{2} \sum_{x, y \in \mathcal{X}} \Psi(x, y)^2 K(x, y) \rho(x, y) \pi(x)$$

for $\rho \in \mathcal{P}(\mathcal{X})$ and $\Psi \in \mathbb{R}^{\mathcal{X} \times \mathcal{X}}$. Let $0 < \varepsilon < 1$ and let $(\rho, \Psi) \in \mathcal{CE}'(\rho_0, \rho_1)$ be such that

$$\int_0^1 \mathcal{A}(\rho_t, \Psi_t) dt < \mathcal{W}^2(\rho_0, \rho_1) + \varepsilon.$$

We set $\rho_i^{\varepsilon} = (1 - \varepsilon)\rho_i + \varepsilon$ for $i = 0, 1$.

Firstly, we define $(\rho^{\varepsilon}, \Psi^{\varepsilon}) \in \mathcal{CE}'(\rho_0^{\varepsilon}, \rho_1^{\varepsilon})$ by

$$\begin{aligned} \rho_t^{\varepsilon}(x) &:= (1 - \varepsilon)\rho_t(x) + \varepsilon, \\ \Psi_t^{\varepsilon}(x, y) &:= (1 - \varepsilon) \frac{\rho_t(x, y)}{\rho_t^{\varepsilon}(x, y)} \Psi_t(x, y). \end{aligned}$$

The concavity assumption (A8) implies the convexity of the function

$$\mathbb{R} \times \mathbb{R}_+ \times \mathbb{R}_+ \ni (x, s, t) \mapsto \frac{x^2}{\theta(s, t)},$$

which yields

$$\int_0^1 \mathcal{A}(\rho_t^\varepsilon, \Psi_t^\varepsilon) dt \leq (1 - \varepsilon) \int_0^1 \mathcal{A}(\rho_t, \Psi_t) dt < (1 - \varepsilon) \mathcal{W}^2(\rho_0, \rho_1) + \varepsilon .$$

Secondly, for $i = 0, 1$, we define $(\rho^{i,\varepsilon}, \Psi^{i,\varepsilon}) \in \mathcal{CE}'(\rho_i, \rho_i^\varepsilon)$ by linear interpolation, i.e.,

$$\rho_t^{i,\varepsilon} := (1 - t)\rho_i + t\rho_i^\varepsilon .$$

As in the proof of Lemma 3.19, for $t \in (0, 1)$, let $\psi_t^{i,\varepsilon}$ be the unique element in $\text{Ran } A(\rho_t^{i,\varepsilon})$ satisfying $\dot{\rho}_t^{i,\varepsilon} = B(\rho_t^{i,\varepsilon})\psi_t^{i,\varepsilon}$. Setting $\Psi^{i,\varepsilon} := \nabla\psi^{i,\varepsilon}$, it then follows that $(\rho^{i,\varepsilon}, \Psi^{i,\varepsilon}) \in \mathcal{CE}'(\rho_i, \rho_i^\varepsilon)$. Lemma 3.19 and its proof imply that there exists a constant $C > 0$, independent of $\varepsilon > 0$, such that

$$\int_0^1 \mathcal{A}(\rho_t^{i,\varepsilon}, \Psi_t^{i,\varepsilon}) dt \leq C d_{TV}^2(\rho_i, \rho_i^\varepsilon) \leq 4C\varepsilon^2 .$$

Finally, it remains to rescale the three curves in time and glue them together. We thus define

$$(\bar{\rho}_t^\varepsilon, \bar{\Psi}_t^\varepsilon) := \begin{cases} (\rho_{t/\varepsilon}^{0,\varepsilon}, \varepsilon^{-1}\Psi_{t/\varepsilon}^{0,\varepsilon}) , & t \in [0, \varepsilon] , \\ (\rho_{(t-\varepsilon)/(1-2\varepsilon)}^\varepsilon, (1-2\varepsilon)^{-1}\Psi_{(t-\varepsilon)/(1-2\varepsilon)}^\varepsilon) , & t \in (\varepsilon, 1-\varepsilon) , \\ (\rho_{(1-t)/\varepsilon}^{1,\varepsilon}, \varepsilon^{-1}\Psi_{(1-t)/\varepsilon}^{1,\varepsilon}) , & t \in [1-\varepsilon, 1] , \end{cases}$$

so that $(\bar{\rho}^\varepsilon, \bar{\Psi}^\varepsilon) \in \mathcal{CE}(\rho_0, \rho_1)$. We infer that

$$\begin{aligned} \int_0^1 \mathcal{A}(\bar{\rho}_t^\varepsilon, \bar{\Psi}_t^\varepsilon) dt &\leq \int_0^1 \frac{\mathcal{A}(\rho_t^{0,\varepsilon}, \Psi_t^{0,\varepsilon})}{\varepsilon} + \frac{\mathcal{A}(\rho_t^\varepsilon, \Psi_t^\varepsilon)}{1-2\varepsilon} + \frac{\mathcal{A}(\rho_t^{1,\varepsilon}, \Psi_t^{1,\varepsilon})}{\varepsilon} dt \\ &\leq 4C\varepsilon + \frac{(1-\varepsilon)\mathcal{W}^2(\rho_0, \rho_1) + \varepsilon}{1-2\varepsilon} + 4C\varepsilon . \end{aligned}$$

Since the right-hand side tends to $\mathcal{W}^2(\rho_0, \rho_1)$ as $\varepsilon \rightarrow 0$, the result follows from the observation that $\bar{\Psi}_t^\varepsilon$ may be replaced by $P_{\bar{\rho}_t^\varepsilon}\bar{\Psi}_t^\varepsilon$, as in the proof of Lemma 3.6. \square

In the next result we will slightly abuse notation and write

$$\partial_1\rho(x, y) := \partial_1\theta(\rho(x), \rho(y)) .$$

Theorem 3.31 (Geodesics). *Suppose that $C_\theta = \infty$ and let $\sigma \in \mathcal{P}(\mathcal{X})$. The following assertions hold:*

- (1) *For each $\bar{\rho}_0, \bar{\rho}_1 \in \mathcal{P}_\sigma(\mathcal{X})$ there exists a constant speed geodesic $\rho : [0, 1] \rightarrow \mathcal{P}(\mathcal{X})$ with $\rho_0 = \bar{\rho}_0$ and $\rho_1 = \bar{\rho}_1$.*
- (2) *Let $\rho : [0, 1] \rightarrow \mathcal{P}_\sigma(\mathcal{X})$ be a constant speed geodesic and let $\psi_t = \mathcal{I}_{\rho_t}\dot{\rho}_t$. Then the following equations hold for $t \in [0, 1]$ and $x \in \mathcal{X}$:*

$$\begin{cases} \partial_t\rho_t(x) = \sum_{y \in \mathcal{X}} (\psi_t(x) - \psi_t(y))K(x, y)\rho_t(x, y) , \\ \partial_t\psi_t(x) = \frac{1}{2} \sum_{y \in \mathcal{X}} (\psi_t(x) - \psi_t(y))^2 K(x, y)\partial_1\rho_t(x, y) . \end{cases} \quad (3.20)$$

Proof. Since $(\mathcal{P}_\sigma(\mathcal{X}), \mathcal{W})$ is a complete Riemannian manifold, (1) follows from the Hopf-Rinow theorem. The equations in (2) are the equations for the cogeodesic flow (see, e.g., [15, Theorem 1.9.3]) and follow directly from the representation of \mathcal{W} as a Riemannian metric given in this section. \square

Remark 3.32. The equations (3.20) should be compared to the geodesic equations for the L^2 -Wasserstein metric over \mathbb{R}^n (see [3], [23], [24]), which are given under appropriate assumptions by

$$\begin{cases} \partial_t \rho + \nabla \cdot (\rho \nabla \psi) = 0, \\ \partial_t \psi + \frac{1}{2} |\nabla \psi|^2 = 0. \end{cases} \quad (3.21)$$

The equations (3.20) are a natural discrete analogue of (3.21). Note however that the equations for ψ in the discrete case depend on ρ .

4. GRADIENT FLOWS OF ENTROPY FUNCTIONALS

We continue in the setting of Section 3, where K is an irreducible and reversible Markov kernel on a finite set \mathcal{X} . We fix a function $\theta : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ satisfying Assumption 3.1 and consider the associated (pseudo-)metric defined in Section 3. If $C_\theta < \infty$, we shall also assume that (A8) holds.

Since $\mathcal{P}_*(\mathcal{X})$ is a Riemannian manifold, as has been shown in Theorem 3.29, we are in a position to study gradient flows of smooth functionals defined on $\mathcal{P}_*(\mathcal{X})$. Let

$$\Delta := K - I$$

denote the generator of the continuous time Markov semigroup $(e^{t\Delta})_{t \geq 0}$ associated with K . The main result in this section is Theorem 4.7, which asserts that solutions to the ‘‘heat equation’’ $\dot{\rho}_t = \Delta \rho_t$ are gradient flow trajectories of the entropy \mathcal{H} with respect to the metric \mathcal{W} .

Notation. In view of Proposition 3.25, we shall always regard T_ρ as being the tangent space of $\mathcal{P}_*(\mathcal{X})$ at $\rho \in \mathcal{P}_*(\mathcal{X})$. The tangent vector field along a smooth curve $t \mapsto \rho_t \in \mathcal{P}_*(\mathcal{X})$ will be denoted by

$$t \mapsto D_t \rho \in T_{\rho_t}.$$

The gradient of a smooth functional $\mathcal{G} : \mathcal{P}_*(\mathcal{X}) \rightarrow \mathbb{R}$ at $\rho \in \mathcal{P}_*(\mathcal{X})$ is denoted by

$$\text{grad } \mathcal{G}(\rho) \in T_\rho.$$

Functionals. We shall consider the following types of functionals:

- For a function $V : \mathcal{X} \rightarrow \mathbb{R}$ we consider the *potential energy functional* $\mathcal{V} : \mathcal{P}_*(\mathcal{X}) \rightarrow \mathbb{R}$ defined by

$$\mathcal{V}(\rho) := \sum_{x \in \mathcal{X}} V(x) \rho(x) \pi(x).$$

- For a differentiable function $f : (0, \infty) \rightarrow \mathbb{R}$, we consider the *generalised entropy* $\mathcal{F} : \mathcal{P}_*(\mathcal{X}) \rightarrow \mathbb{R}$ defined by

$$\mathcal{F}(\rho) := \sum_{x \in \mathcal{X}} f(\rho(x)) \pi(x).$$

Proposition 4.1 (Gradient of potential energy functionals). *The functional $\mathcal{V} : \mathcal{P}_*(\mathcal{X}) \rightarrow \mathbb{R}$ is differentiable, and for $\rho \in \mathcal{P}_*(\mathcal{X})$ we have*

$$\text{grad } \mathcal{V}(\rho) = \nabla V .$$

Proof. Clearly, \mathcal{V} is differentiable. Let $t \mapsto \rho_t \in \mathcal{P}_*(\mathcal{X})$ be a differentiable curve and let $\psi_t \in \text{Ran } A(\rho_t)$ be such that $\nabla \psi_t := D_t \rho$. Then

$$\begin{aligned} \frac{d}{dt} \mathcal{V}(u_t) &= \sum_{x \in \mathcal{X}} V(x) \dot{\rho}_t(x) \pi(x) = \sum_{x \in \mathcal{X}} V(x) (B(\rho_t) \psi_t)(x) \pi(x) \\ &= -\langle V, \nabla \cdot (\hat{\rho}_t \bullet \nabla \psi_t) \rangle_\pi = \langle \nabla V, \hat{\rho}_t \bullet \nabla \psi_t \rangle_\pi = \langle \nabla V, \nabla \psi_t \rangle_{\rho_t} , \end{aligned}$$

which yields the result. \square

Proposition 4.2 (Gradient of generalised entropy functionals). *The functional $\mathcal{F} : \mathcal{P}_*(\mathcal{X}) \rightarrow \mathbb{R}$ is differentiable, and for $\rho \in \mathcal{P}_*(\mathcal{X})$ we have*

$$\text{grad } \mathcal{F}(\rho) = \nabla (f' \circ \rho) .$$

Proof. The differentiability of \mathcal{F} is clear from its definition. Let $t \mapsto \rho_t \in \mathcal{P}_*(\mathcal{X})$ be a differentiable curve and let $\psi_t \in \text{Ran } A(\rho_t)$ be such that $\nabla \psi_t := D_t \rho$. Since f is differentiable, we obtain

$$\begin{aligned} \frac{d}{dt} \mathcal{F}(u_t) &= \sum_{x \in \mathcal{X}} f'(\rho_t(x)) \dot{\rho}_t(x) \pi(x) = \sum_{x \in \mathcal{X}} f'(\rho_t(x)) (B(\rho_t) \psi_t)(x) \pi(x) \\ &= -\langle f'(\rho_t), \nabla \cdot (\hat{\rho}_t \bullet \nabla \psi_t) \rangle_\pi = \langle \nabla f'(\rho_t), \hat{\rho}_t \bullet \nabla \psi_t \rangle_\pi \\ &= \langle \nabla f'(\rho_t), \nabla \psi_t \rangle_{\rho_t} , \end{aligned}$$

which yields the result. \square

In the special case where $\mathcal{F} = \mathcal{H}$ is the entropy functional from (1.1) we obtain:

Corollary 4.3. *The functional $\mathcal{H} : \mathcal{P}_*(\mathcal{X}) \rightarrow \mathbb{R}$ is differentiable, and for $\rho \in \mathcal{P}_*(\mathcal{X})$ we have*

$$\text{grad } \mathcal{H}(\rho) = \nabla \log \rho .$$

Proof. This follows directly from Proposition 4.2. \square

Gradient flows. In order to study gradient flows, we impose the following assumption which will be in force throughout the remainder of this section.

Assumption 4.4. *In addition to Assumption 3.1 we assume:*

(A9) *There exists a function $k \in C^\infty((0, \infty); \mathbb{R})$ such that*

$$\theta(s, t) = \frac{s - t}{k(s) - k(t)}$$

for all $s, t > 0$ with $s \neq t$.

Recall that this assumption is satisfied if θ is the logarithmic mean, in which case $k(t) = \log(t)$.

Proposition 4.5 (Tangent vector field along the heat flow). *Let $\rho \in \mathcal{P}(\mathcal{X})$ and let $\rho_t = e^{t\Delta} \rho$, $t \geq 0$ denote the heat flow. Then $t \mapsto \rho_t$ is C^∞ on $(0, \infty)$ and for $t > 0$ we have*

$$D_t \rho = -\nabla (k \circ \rho_t) .$$

Proof. The differentiability assertion follows from general Markov chain theory. For any $\rho \in \mathcal{P}_*(\mathcal{X})$, we have

$$\rho(x, y) = \frac{\rho(x) - \rho(y)}{k(\rho(x)) - k(\rho(y))},$$

and therefore

$$\Delta\rho = \nabla \cdot (\nabla\rho) = \nabla \cdot (\widehat{\rho} \bullet \nabla(k \circ \rho)).$$

Since $t \mapsto \rho_t$ solves the heat equation $\dot{\rho}_t = \Delta\rho_t$, it follows that

$$\dot{\rho}_t - \nabla \cdot (\widehat{\rho}_t \bullet \nabla(k \circ \rho_t)) = 0,$$

hence $D_t\rho = -\nabla(k \circ \rho_t)$ by Proposition 3.26. \square

We slightly modify the usual definition of a gradient flow trajectory, as we wish to allow for initial values that do not belong to $\mathcal{P}_*(\mathcal{X})$:

Definition 4.6 (Gradient flow). *Let $\mathcal{F} : \mathcal{P}_*(\mathcal{X}) \rightarrow \mathbb{R}$ be differentiable. A curve $\rho : [0, \infty) \rightarrow \mathcal{P}(\mathcal{X})$ is said to be a gradient flow trajectory for \mathcal{F} starting from $\bar{\rho} \in \mathcal{P}(\mathcal{X})$ if the following assertions hold:*

- (1) $t \mapsto \rho_t$ is differentiable on $(0, \infty)$, for every $t > 0$ we have $\rho_t \in \mathcal{P}_*(\mathcal{X})$ and

$$D_t\rho = -\text{grad } \mathcal{F}(\rho_t).$$

- (2) $t \mapsto \rho_t$ is continuous in total variation at $t = 0$ and $\rho_0 = \bar{\rho}$.

Theorem 4.7. *Let $f \in C^2((0, \infty); \mathbb{R})$ be such that $f' = k$ and let $\rho \in \mathcal{P}(\mathcal{X})$. Then the heat flow $t \mapsto e^{t\Delta}\rho$ is a gradient flow trajectory for the functional \mathcal{F} with respect to \mathcal{W} .*

Proof. The first condition in Definition 4.6 is a consequence of Propositions 4.2 and 4.5. The second one follows from general Markov chain theory. \square

Corollary 4.8 (Heat flow is gradient flow of the entropy). *Let θ be the logarithmic mean defined by $\theta(s, t) = \int_0^1 s^{1-pt} p \, dp$ and let $\rho \in \mathcal{P}(\mathcal{X})$. Then the heat flow $t \mapsto e^{t\Delta}\rho$ is a gradient flow trajectory for the entropy \mathcal{H} with respect to \mathcal{W} .*

Proof. This is a special case of Theorem 4.7 with $k(t) = 1 + \log t$ and $f(t) = t \log t$. \square

APPENDIX A. A RESULT FROM THE THEORY OF DIAGONALLY DOMINANT MATRICES

The following result from the theory of diagonally dominant matrices is a special case of [8]. For the convenience of the reader we present a simple proof.

Lemma A.1. *Let $A = (a_{ij})_{i,j=1,\dots,n}$ be a real matrix satisfying*

- (1) $\forall i : a_{ii} \geq 0$, (2) $\forall i \neq j : a_{ij} = a_{ji} \leq 0$, (3) $\forall i : \sum_j a_{ij} = 0$.

Consider the equivalence relation \sim on $I = \{1, \dots, n\}$ defined by

$$i \sim j \quad :\Leftrightarrow \quad \begin{cases} i = j, & \text{or} \\ \exists k \geq 1 \exists i_1, \dots, i_k \in I : a_{i,i_1}, a_{i_1,i_2}, \dots, a_{i_k,j} < 0, \end{cases}$$

and let $(I_\alpha)_\alpha \subseteq I$ denote the corresponding equivalence classes. Then the following identities hold:

$$\text{Ker } A = \{(x_i) \in \mathbb{R}^n \mid x_i = x_j \text{ whenever } i \sim j\}, \quad (\text{A.1})$$

$$\text{Ran } A = \left\{ (x_i) \in \mathbb{R}^n \mid \forall \alpha : \sum_{i \in I_\alpha} x_i = 0 \right\}. \quad (\text{A.2})$$

Proof. First we remark that the assumptions (1) – (3) imply that $a_{ij} = 0$ if $i \in I_\alpha$ and $j \in I_\beta$ for some $\alpha \neq \beta$. Furthermore, it suffices to show (A.1), since (A.2) then follows by duality.

To show “ \supseteq ”, suppose that $x = (x_i)$ satisfies $x_i = x_j$ whenever $i \sim j$. Fix $k \in I$ and take β such that $k \in I_\beta$. Using the remark and (3), it follows that

$$\sum_{j \in I} a_{kj} x_j = \sum_{j \in I_\beta} a_{kj} x_j = x_k \sum_{j \in I_\beta} a_{kj} = x_k \sum_{j \in I} a_{kj} = 0,$$

which yields the desired inclusion.

Conversely, to show “ \subseteq ”, we use the identity

$$2x_i x_j = x_i^2 + x_j^2 - (x_i - x_j)^2$$

to write, for $x = (x_i)$,

$$\begin{aligned} 2\langle Ax, x \rangle &= 2 \sum_{i,j \in I} a_{ij} x_i x_j \\ &= \sum_{i \in I} x_i^2 \sum_{j \in I} a_{ij} + \sum_{j \in I} x_j^2 \sum_{i \in I} a_{ij} - \sum_{i,j \in I} a_{ij} (x_i - x_j)^2. \end{aligned}$$

Using (3) and the symmetry of A we infer that

$$\langle Ax, x \rangle = -\frac{1}{2} \sum_{i,j \in I} a_{ij} (x_i - x_j)^2.$$

Consequently, if $Ax = 0$, it follows that $\langle Ax, x \rangle = 0$, hence $x_i = x_j$ whenever $i \sim j$, which completes the proof. \square

APPENDIX B. UNIQUENESS OF THE METRIC ON THE TWO-POINT SPACE

In this appendix we shall prove Proposition 2.13. First we need two definitions. Let (M, d) be a metric space.

Definition B.1. Let $I \subseteq \mathbb{R}$ be an interval and let $1 \leq p < \infty$. A curve $\gamma : I \rightarrow M$ is said to be p -absolutely continuous if there exists a function $m \in L^p(I; \mathbb{R})$ such that

$$d(\gamma(s), \gamma(t)) \leq \int_s^t m(r) \, dr$$

for all $s, t \in I$ with $s \leq t$. The curve γ is locally p -absolutely continuous if it is p -absolutely continuous on each compact subinterval of I .

We shall use the notation $\gamma \in AC^p(I; M)$ and $\gamma \in AC_{\text{loc}}^p(I; M)$ respectively.

The following notion of gradient flow in a metric space (M, d) has been studied in great detail in [1].

Definition B.2. Let $\mathcal{F} : M \rightarrow \mathbb{R} \cup \{+\infty\}$ be lower-semicontinuous and not identically $+\infty$. A curve $\gamma \in C([0, \infty); M) \cap AC_{\text{loc}}^2((0, \infty); M)$ is said to satisfy the evolution variational inequality (EVI $_{\lambda}(\mathcal{F})$) if, for any $y \in \text{D}(\mathcal{F})$, the inequality

$$\frac{1}{2} \frac{d}{dt} d^2(\gamma(t), y) + \frac{\lambda}{2} d^2(\gamma(t), y) \leq \mathcal{F}(y) - \mathcal{F}(\gamma(t)) \quad (\text{B.1})$$

holds a.e. on $(0, \infty)$.

Proof of Proposition 2.13. Recall that $\bar{\beta} = \frac{p-q}{p+q}$. Let $\beta \in (\bar{\beta}, 1)$ and suppose that there exists $\alpha \in (-1, 1)$ such that

$$\mathcal{M}(\rho^{\bar{\beta}}, \rho^{\beta}) = \mathcal{M}(\rho^{\bar{\beta}}, \rho^{\alpha}) + \mathcal{M}(\rho^{\alpha}, \rho^{\beta}). \quad (\text{B.2})$$

We claim that $\alpha \in [\bar{\beta}, \beta]$. To prove this, suppose first – to obtain a contradiction – that $\alpha > \beta$. Then there exists $T > 0$ such that $e^{T(K-I)}\rho^{\alpha} = \rho^{\beta}$, hence (B.1) implies that

$$\mathcal{M}(\rho^{\beta}, \rho^{\bar{\beta}})^2 - \mathcal{M}(\rho^{\alpha}, \rho^{\bar{\beta}})^2 \leq 2T(\mathcal{H}(\rho^{\bar{\beta}}) - \mathcal{H}(\rho^{\beta})) \leq 0.$$

In view of (B.2), it follows that $\mathcal{M}(\rho^{\alpha}, \rho^{\beta}) = 0$, thus $\alpha = \beta$, which contradicts the assumption. Suppose now that $\alpha < \bar{\beta}$. Adding (B.2) and the inequality in (2) we infer that $\rho^{\alpha} = \rho^{\bar{\beta}}$, hence $\alpha = \bar{\beta}$, which proves the claim.

Now, fix $\beta \in (\bar{\beta}, 1)$ and let $t \mapsto \rho^{\psi(t)}$ be a speed-1 geodesic with $\psi(0) = \bar{\beta}$ and $\psi(T) = \beta$ where $T = \mathcal{M}(\rho^{\bar{\beta}}, \rho^{\beta})$. For $0 \leq s < t \leq T$ we then have $\mathcal{M}(\rho^{\bar{\beta}}, \rho^{\psi(t)}) = \mathcal{M}(\rho^{\bar{\beta}}, \rho^{\psi(s)}) + \mathcal{M}(\rho^{\psi(s)}, \rho^{\psi(t)})$, thus the claim implies that $\psi(s) \leq \psi(t)$. Since ψ is a geodesic, we have $\psi(s) \neq \psi(t)$, thus ψ is strictly increasing on $[0, T]$.

Now we claim that ψ is continuous on $[0, T]$. To show this, take $t \in (0, T)$. Since ψ is increasing, the limits $\psi(t-)$ and $\psi(t+)$ exist and for any $\varepsilon > 0$ we have $\mathcal{M}(\rho^{\psi(t-)}, \rho^{\psi(t+)}) \leq \mathcal{M}(\rho^{\psi(t-\varepsilon)}, \rho^{\psi(t-\varepsilon)}) = 2\varepsilon$, thus $\psi(t-) = \psi(t+)$. A similar argument shows that ψ is continuous at 0 and T , thus ψ is continuous on $[0, T]$. Since ψ is continuous and strictly increasing we infer that the mapping $\psi : [0, T] \rightarrow [\bar{\beta}, \beta]$ is surjective. As a consequence, the inverse mapping $\varphi : [\bar{\beta}, \beta] \rightarrow [0, T]$ is well-defined, and continuous and strictly increasing as well.

Note that the mapping

$$I : t \mapsto \rho^{\psi(t)} \quad (\text{B.3})$$

defines an isometry from $[0, T]$ endowed with the euclidean metric onto $\{\rho^{\alpha} : \alpha \in [\bar{\beta}, \beta]\} \subseteq \mathcal{P}_*(\mathcal{X})$ endowed with the metric \mathcal{M} . The inverse mapping is given by

$$J : \rho^{\alpha} \rightarrow \varphi(\alpha).$$

Since $u : t \mapsto \rho^{\beta t}$ is a 2-absolutely continuous curve satisfying $\text{EVI}_0(\mathcal{H})$ for the metric \mathcal{M} , (B.3) implies that the mapping

$$t \mapsto \tilde{u}(t) := J(u(t)) = \varphi(\beta t)$$

is a 2-absolutely continuous curve satisfying $\text{EVI}_0(\tilde{\mathcal{H}})$ where $\tilde{\mathcal{H}} := \mathcal{H} \circ I$, for the euclidean metric. It follows that the mapping $\varphi : [\bar{\beta}, \beta] \rightarrow [0, T]$ itself is

absolutely continuous, hence almost everywhere differentiable, and the same holds for its inverse ψ . Moreover, the identity

$$\psi'(\varphi(\alpha))\varphi'(\alpha) = 1 \quad (\text{B.4})$$

holds for a.e. $\alpha \in [\bar{\beta}, \beta]$.

For any $\alpha \in [\bar{\beta}, \beta]$ we have

$$\begin{aligned} \tilde{\mathcal{H}}(\varphi(\alpha)) &= \mathcal{H}(I(\varphi^\alpha)) = \mathcal{H}(\rho^\alpha) \\ &= \frac{q}{p+q} f\left(\frac{p+q}{q} \frac{1-\alpha}{2}\right) + \frac{p}{p+q} f\left(\frac{p+q}{p} \frac{1+\alpha}{2}\right), \end{aligned}$$

thus, for $r \in (0, T)$,

$$\tilde{\mathcal{H}}(r) = \frac{q}{p+q} f\left(\frac{p+q}{q} \frac{1-\psi(r)}{2}\right) + \frac{p}{p+q} f\left(\frac{p+q}{p} \frac{1+\psi(r)}{2}\right).$$

It follows that $\tilde{\mathcal{H}}$ is a.e. differentiable and the identity

$$\tilde{\mathcal{H}}'(r) = \frac{\psi'(r)}{2} \left[f'\left(\frac{p+q}{p} \frac{1+\psi(r)}{2}\right) - f'\left(\frac{p+q}{q} \frac{1-\psi(r)}{2}\right) \right] \quad (\text{B.5})$$

holds a.e.

Since $t \mapsto \tilde{u}(t)$ is a 2-absolutely continuous curve satisfying $\text{EVI}_0(\tilde{\mathcal{H}})$ and since the functional $\tilde{\mathcal{H}}$ is differentiable a.e., it follows from [1, Proposition 1.4.1] that the gradient flow equation

$$\tilde{u}'(t) = -\tilde{\mathcal{H}}'(\tilde{u}(t))$$

holds almost everywhere.

Since φ is differentiable a.e., the left-hand side equals a.e.

$$\tilde{u}'(t) = \frac{d}{dt} \varphi(\beta_t) = (p(1-\beta_t) - q(1+\beta_t)) \varphi'(\beta_t).$$

Taking (B.4) into account, it follows from (B.5) that the right-hand side equals a.e.

$$\tilde{\mathcal{H}}'(\tilde{u}(t)) = \frac{1}{2\varphi'(\beta_t)} \left[f'(\rho^{\beta_t}(b)) - f'(\rho^{\beta_t}(a)) \right]$$

Combining the latter two inequalities we infer that for a.e. $\alpha \in [\bar{\beta}, \beta]$,

$$(q(1+\alpha) - p(1-\alpha)) \varphi'(\alpha) = \frac{1}{2\varphi'(\alpha)} \left[f'(\rho^\alpha(b)) - f'(\rho^\alpha(a)) \right]$$

Since φ is absolutely continuous,

$$\varphi(\beta) = \int_{\bar{\beta}}^{\beta} \varphi'(\alpha) d\alpha = \int_{\bar{\beta}}^{\beta} \sqrt{\frac{f'(\rho^\alpha(b)) - f'(\rho^\alpha(a))}{2(q(1+\alpha) - p(1-\alpha))}} d\alpha.$$

hence, since $t \mapsto \psi(t)$ is a geodesic, we obtain for $\bar{\beta} < \alpha < \beta$,

$$\mathcal{M}(\rho^\alpha, \rho^\beta) = \mathcal{M}(\rho^{\psi(\varphi(\alpha))}, \rho^{\psi(\varphi(\beta))}) = C(\varphi(\beta) - \varphi(\alpha)).$$

Thus the distance between ρ^α and ρ^β is uniquely determined for all $\alpha, \beta \geq \bar{\beta}$. The same argument shows that the distance is uniquely determined for $\alpha, \beta \leq \bar{\beta}$. The case $\alpha < \bar{\beta} < \beta$ follows from the assumption (2). \square

REFERENCES

- [1] L. AMBROSIO, N. GIGLI, AND G. SAVARÉ, *Gradient flows in metric spaces and in the space of probability measures*, second ed., Lectures in Mathematics ETH Zürich, Birkhäuser Verlag, Basel, 2008.
- [2] L. AMBROSIO, G. SAVARÉ, AND L. ZAMBOTTI, *Existence and stability for Fokker-Planck equations with log-concave reference measure*, Probab. Theory Related Fields **145** (2009), no. 3-4, 517–564.
- [3] J.-D. BENAMOU AND Y. BRENIER, *A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem*, Numer. Math. **84** (2000), no. 3, 375–393.
- [4] R. BHATIA, *Positive definite matrices*, Princeton Series in Applied Mathematics, Princeton University Press, Princeton, NJ, 2007.
- [5] A.-I. BONCIOCAT AND K.-TH. STURM, *Mass transportation and rough curvature bounds for discrete spaces*, J. Funct. Anal. **256** (2009), no. 9, 2944–2966.
- [6] J. A. CARRILLO, S. LISINI, G. SAVARÉ, AND D. SLEPČEV, *Nonlinear mobility continuity equations and generalized displacement convexity*, J. Funct. Anal. **258** (2010), no. 4, 1273–1309.
- [7] S.-N. CHOW, W. HUANG, Y. LI, AND H. ZHOU, *Fokker-Planck equations for a free energy functional or Markov process on a graph*, preprint.
- [8] G. DAHL, *A note on diagonally dominant matrices*, Linear Algebra Appl. **317** (2000), no. 1-3, 217–224.
- [9] J. DOLBEAULT, B. NAZARET, AND G. SAVARÉ, *A new class of transport distances between measures*, Calc. Var. Partial Differential Equations **34** (2009), no. 2, 193–231.
- [10] M. ERBAR, *The heat equation on manifolds as a gradient flow in the Wasserstein space*, Ann. Inst. Henri Poincaré Probab. Stat. **46** (2010), no. 1, 1–23.
- [11] S. FANG, J. SHAO, AND K.-TH. STURM, *Wasserstein space over the Wiener space*, Probab. Theory Related Fields **146** (2010), no. 3-4, 535–565.
- [12] N. GIGLI, *On the heat flow on metric measure spaces: existence, uniqueness and stability*, Calc. Var. Partial Differential Equations **39** (2010), no. 1-2, 101–120.
- [13] N. GIGLI, K. KUWADA, AND S.-I. OHTA, *Heat flow on Alexandrov spaces*, preprint at arXiv:1008.1319 (2010).
- [14] R. JORDAN, D. KINDERLEHRER, AND F. OTTO, *The variational formulation of the Fokker-Planck equation*, SIAM J. Math. Anal. **29** (1998), no. 1, 1–17.
- [15] J. JOST, *Riemannian geometry and geometric analysis*, fifth ed., Universitext, Springer-Verlag, Berlin, 2008.
- [16] Y. LIN AND S.-T. YAU, *Ricci curvature and eigenvalue estimate on locally finite graphs*, Math. Res. Lett. **17** (2010), no. 2, 343–356.
- [17] J. LOTT AND C. VILLANI, *Ricci curvature for metric-measure spaces via optimal transport*, Ann. of Math. (2) **169** (2009), no. 3, 903–991.
- [18] W.H. MCADAMS, *Heat transmission*, vol. 532, McGraw-Hill New York, 1954.
- [19] S.-I. OHTA AND K.-TH. STURM, *Heat flow on Finsler manifolds*, Comm. Pure Appl. Math. **62** (2009), no. 10, 1386–1433.
- [20] Y. OLLIVIER, *Ricci curvature of metric spaces*, C. R. Math. Acad. Sci. Paris **345** (2007), no. 11, 643–646.
- [21] Y. OLLIVIER, *Ricci curvature of Markov chains on metric spaces*, J. Funct. Anal. **256** (2009), no. 3, 810–864.
- [22] Y. OLLIVIER AND C. VILLANI, *A curved Brunn-Minkowski inequality on the discrete hypercube*, preprint at arXiv:1011.4779 (2010).
- [23] F. OTTO, *The geometry of dissipative evolution equations: the porous medium equation*, Comm. Partial Differential Equations **26** (2001), no. 1-2, 101–174.
- [24] F. OTTO AND C. VILLANI, *Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality*, J. Funct. Anal. **173** (2000), no. 2, 361–400.
- [25] G. SAVARÉ, *Gradient flows and diffusion semigroups in metric spaces under lower curvature bounds*, C. R. Math. Acad. Sci. Paris **345** (2007), no. 3, 151–154.
- [26] K.-TH. STURM, *On the geometry of metric measure spaces. I and II*, Acta Math. **196** (2006), no. 1, 65–177.

- [27] C. VILLANI, *Topics in optimal transportation*, Graduate Studies in Mathematics, vol. 58, American Mathematical Society, Providence, RI, 2003.
- [28] C. VILLANI, *Optimal transport, old and new*, Grundlehren der Mathematischen Wissenschaften, vol. 338, Springer-Verlag, Berlin, 2009.

UNIVERSITY OF BONN, INSTITUTE FOR APPLIED MATHEMATICS, ENDENICHER ALLEE
60, 53115 BONN, GERMANY

E-mail address: `maas@iam.uni-bonn.de`

`http://www.janmaas.org`