

Monotonicity in Quadratically Regularized Linear Programs*

Alberto González-Sanz[†] Marcel Nutz[‡] Andrés Riveros Valdevenito[§]

April 16, 2025

Abstract

In optimal transport, quadratic regularization is a sparse alternative to entropic regularization: the solution measure tends to have small support. Computational experience suggests that the support decreases monotonically to the unregularized counterpart as the regularization parameter is relaxed. We find it useful to investigate this monotonicity more abstractly for linear programs over polytopes, regularized with the squared norm. Here, monotonicity can be stated as an invariance property of the curve mapping the regularization parameter to the solution: once the curve enters a face of the polytope, does it remain in that face forever? We show that this invariance is equivalent to a geometric property of the polytope, namely that each face contains the minimum norm point of its affine hull. Returning to the optimal transport problem and its associated Birkhoff polytope, we verify this property for low dimensions, but show that it fails for marginals with five or more point masses. As a consequence, the conjectured monotonicity of the support *fails* in general, even if experiments suggest that monotonicity holds for many cost matrices. Separately, we apply our geometric point of view to a problem of Erdős, namely to characterize the doubly stochastic matrices whose maximal trace equals their squared norm.

Keywords Linear Program, Quadratic Regularization, Optimal Transport, Sparsity

AMS 2020 Subject Classification 49N10; 49N05; 90C25

1 Introduction

Let $\mathcal{P} \subset \mathbb{R}^d$ be a polytope and $c \in \mathbb{R}^d$. The regularized linear program

$$x^\delta = \operatorname{argmin}_{x \in \mathcal{P}: \|x\| \leq \delta} \langle c, x \rangle \quad (1)$$

has a unique solution as long as $\delta \geq 0$ belongs to a certain interval $[\delta_{\min}, \delta_{\max}]$, cf. Lemma 2.1, hence we can consider the curve $\delta \mapsto x^\delta \in \mathcal{P}$ which travels across various faces of $\mathcal{P}$ as δ increases (i.e., as the regularizing constraint relaxes). We are interested in the following invariance property: once $\delta \mapsto x^\delta \in \mathcal{P}$ enters a given face, does it ever leave that face? (Cf. Figure 1.) In the optimal transport applications that motivate our study (see below), this property corresponds to the *monotonicity of the optimal support* $\operatorname{spt} x^\delta$ wrt. the *regularization strength*. (The support is defined as the set of locations with nonzero mass,

*The authors thank three anonymous referees for their careful reading and constructive comments.

[†]Department of Statistics, Columbia University, ag4855@columbia.edu

[‡]Departments of Mathematics and Statistics, Columbia University, mnutz@columbia.edu. Research supported by NSF Grants DMS-1812661, DMS-2106056, DMS-2407074.

[§]Department of Statistics, Columbia University, ar4151@columbia.edu.




Figure 1: A non-monotone (left) and a monotone (right) polytope with cost c (red). The curve $[\delta_{\min}, \delta_{\max}] \ni \delta \mapsto x^\delta$ follows the blue arrows and has the invariance property only in the right example.

$\text{spt } x := \{i : x_i > 0\}$ for $x = (x_1, \dots, x_d) \in \mathbb{R}^d$.) Our abstract result, Theorem 3.2, geometrically characterizes all polytopes such that the invariance property holds (for any cost c). We show that this property holds for the set of probability measures (unit simplex), but fails for the set of transport plans (Birkhoff polytope) when the dimension $d \geq 25$. As a consequence—which may be surprising given numerical experience—the optimal support in quadratically regularized optimal transport problems is *not* always monotone.

1.1 Motivation

This study is motivated by optimal transport and related minimization problems over probability measures. In its simplest form, the transport problem between probability measures μ and ν is

$$\inf_{\gamma \in \Gamma(\mu, \nu)} \int \hat{c}(x, y) \gamma(dx, dy), \quad (\text{OT})$$

where $\hat{c}(x, y)$ is a given “cost” function and $\Gamma(\mu, \nu)$ denotes the set of couplings; i.e., joint probability measures γ with marginals (μ, ν) . See [35] for a detailed discussion. In many applications such as machine learning or statistics (see [24, 32] for surveys), the marginals encode observed data points $X_1, \dots, X_N$ and $Y_1, \dots, Y_N$ which are represented by their empirical measures $\mu = \frac{1}{N} \sum_i \delta_{X_i}$ and $\nu = \frac{1}{N} \sum_i \delta_{Y_i}$. Denoting by $c_{ij} = \hat{c}(X_i, Y_j)$ the resulting $N \times N$ cost matrix, (OT) then becomes a linear program

$$\inf_{x \in \mathcal{P}} \langle c, x \rangle \quad (\text{LP})$$

where the polytope $\mathcal{P}$ is (up to a constant factor) the Birkhoff polytope of doubly stochastic matrices of size $N \times N$. For different choices of polytope, (LP) includes other problems of recent interest, such as multi-marginal optimal transport and Wasserstein barycenters [2] or adapted Wasserstein distances [3]. The optimal transport problem is computationally costly when N is large. The impactful paper [14] proposed to regularize (OT) by penalizing with Kullback–Leibler divergence (entropy). Then, solutions can be computed using the Sinkhorn–Knopp algorithm, which has lead to an explosion of high-dimensional applications

(see [33]). More generally, [15] introduced regularized optimal transport with regularization by a divergence. Different divergences give rise to different properties of the solution. Entropic regularization always leads to couplings whose support contains all data pairs (X_i, Y_j) , even though (OT) typically has a sparse solution. In some applications that is undesirable; for instance, it may correspond to blurrier images in an image processing task [6]. The second-most prominent regularization is χ^2 -divergence or equivalently the squared norm, as proposed in [6],¹ which gives rise to sparse solutions. In the Eulerian formulation of [15], this is exactly our problem (1) with $\mathcal{P}$ being the (scaled) Birkhoff polytope. Alternately, the problem can be stated in Lagrangian form as in [6], making the squared-norm penalty explicit:

$$\inf_{\gamma \in \Gamma(\mu, \nu)} \int \hat{c}(x, y) \gamma(dx, dy) + \frac{1}{2\eta} \left\| \frac{d\gamma}{d(\mu \otimes \nu)} \right\|_{L^2(\mu \otimes \nu)}^2 \quad (\text{QOT})$$

where $d\gamma/d(\mu \otimes \nu)$ denotes the density of γ wrt. the product measure $\mu \otimes \nu$ and the regularization strength is now parameterized by $\eta \in [0, \infty]$. In the general setting, this corresponds to

$$\inf_{x \in \mathcal{P}} \langle c, x \rangle + \frac{1}{2\eta} \|x\|^2 \quad (2)$$

which has a unique solution x_η for any $\eta \in (0, \infty)$. The curves (x^η) and (x_η) of solutions to (1) and (2), respectively, coincide up to a simple reparametrization.

In optimal transport, the regularized problem is often solved to approximate the linear problem (OT). The latter has a generically unique (see [13]) solution x , which is recovered as $\delta \rightarrow \delta_{\max}$ (or $\eta \rightarrow \infty$): $x = x^{\delta_{\max}} = x_\infty$.² In particular, the support of x^δ converges to the support of x , which is generically sparse (N out of N^2 possible points, again by [13]). In numerous experiments, it has been observed not only that $\text{spt } x^\delta$ is sparse when δ is large, but also that $\text{spt } x^\delta$ monotonically decreases to $\text{spt } x$ (e.g., [6]). If this monotonicity holds, then in particular $\text{spt } x \subset \text{spt } x^\delta$, meaning that $\text{spt } x^\delta$ can be used as a multivariate confidence band for the (unknown) solution x .

1.2 Summary of Results

When $\mathcal{P}$ is the unit simplex or the Birkhoff polytope, its elements can be interpreted as probability measures and thus carry a natural notion of support, $\text{spt } x := \{i : x_i > 0\}$ for $x = (x_1, \dots, x_d) \in \mathbb{R}^d$. We can then ask whether $\text{spt } x^\delta$ is monotone wrt. the strength δ of regularization. When $\mathcal{P}$ is a general polytope, the set of vertices of the minimal face containing x yields a similar notion, replacing the concept of support. The question of monotonicity then becomes to the aforementioned invariance property: if $x^\delta \in F$ for some face F , does it follow that $x^{\delta'} \in F$ for all $\delta' \geq \delta$? (This indeed specializes to the monotonicity of $\text{spt } x^\delta$ when $\mathcal{P}$ is the simplex or the Birkhoff polytope; cf. Lemma 4.2.) The answer may of course depend on the cost c ; we call $\mathcal{P}$ monotone if the invariance holds for any $c \in \mathbb{R}^d$.

We show that monotonicity can be characterized by the geometry of $\mathcal{P}$. Namely, monotonicity of $\mathcal{P}$ is equivalent to two properties: for any proper face F , the minimum-norm point

¹See also [18] for a similar formulation of minimum-cost flow problems, and the predecessors referenced therein. Our model with a general polytope includes such problems, and many others.

²By “generic” we mean that the set of c where uniqueness fails is a Lebesgue nullset. For problems with non-unique solution, $x^{\delta_{\max}}$ recovers the minimum-norm solution.

of the affine hull of F must lie in F , and moreover, the minimum-norm point of $\mathcal{P}$ must lie in the relative interior $\text{ri } \mathcal{P}$. See Theorem 3.2. Once the right point of view is taken, the proof is quite elementary. An example satisfying both properties, and hence monotonicity, is the unit simplex $\Delta \subset \mathbb{R}^d$, for any $d \geq 1$. For that choice of polytope, x^δ is a sparse soft-min of $c = (c_1, \dots, c_d)$, converging to $\frac{1}{\#\arg\min c} 1_{\{\arg\min c\}}$ in the unregularized limit. On the other hand, we show that the Birkhoff polytope violates the first condition whenever the marginals have $N \geq 5$ data points (meaning that $d \geq 25$). As a result, the optimal support in quadratically regularized optimal transport problems is *not* always monotone, even if it appears so in numerous experiments. In fact, we exhibit a particularly egregious failure of monotonicity where the support of the limiting (unregularized) optimal coupling x is not contained in $\text{spt } x^\delta$ for some reasonably large δ . Our counterexample is constructed by hand, based on theoretical considerations. Our numerical experiments using random cost matrices have failed to locate counterexamples, suggesting that there are, in some sense, “few” faces violating our condition on the minimum-norm point. (See also the proof of Lemma 4.10.)

This is an interesting problem for further study in combinatorics, where the Birkhoff polytope has remained an active topic even after decades of research (e.g., [31] and the references therein). In that spirit, we also apply our point of view to a problem of Erdős, namely to characterize the doubly stochastic $N \times N$ matrices A satisfying $\|A\|^2 = \max \text{tr } A := \max_\sigma \sum_{i=1}^N a_{i,\sigma(i)}$ where σ ranges over all permutations of $\{1, \dots, N\}$. We show in Theorem 5.4 that this is precisely the set of minimum-norm points of certain faces F of the Birkhoff polytope. We use this geometric characterization to prove a conjecture of [7], namely that such matrices must have rational entries.

The remainder of this note is organized as follows. Section 2 introduces the regularized linear program and its solution curve in detail. Section 3 contains the main abstract result, whereas Section 4 reports the applications to optimal transport and soft-min. We conclude in Section 5 with the application to Erdős’ problem about doubly stochastic matrices.

2 Preliminaries

This section collects notation and well-known or straightforward results for ease of reference. Let $\langle \cdot, \cdot \rangle$ be an inner product on $\mathbb{R}^d$ and $\|\cdot\|$ the induced norm. Let $\emptyset \neq \mathcal{P} \subseteq \mathbb{R}^d$ be a polytope; i.e., the convex hull of finitely many points. The minimal set of such points, called vertices or extreme points, is denoted $\text{ext } \mathcal{P}$. A face of $\mathcal{P}$ is a subset $F \subset \mathcal{P}$ such that any open line segment $L = (x, y) \subset \mathcal{P}$ with $L \cap F \neq \emptyset$ satisfies $\overline{L} \subset F$, where $\overline{L} = [x, y]$ denotes closure. Alternately, a nonempty face F of $\mathcal{P}$ is the intersection $F = \mathcal{P} \cap H$ of $\mathcal{P}$ with a tangent hyperplane H . Here a hyperplane $H = \{x \in \mathbb{R}^d : \langle a, x \rangle = b\}$ is called tangent if $H \cap \mathcal{P} \neq \emptyset$ and $\mathcal{P} \subset \{x \in \mathbb{R}^d : \langle a, x \rangle \leq b\}$. We denote by $\text{ri } K$ and $\text{rbd } K = K \setminus \text{ri } K$ the relative interior and relative boundary of a set K , and by $\text{aff } K$ its affine span. See [9] for background on polytopes and related notions.

Next, we introduce the regularized linear program and the interval of parameters δ where the set of feasible points is nonempty and the constraint is binding.

Lemma 2.1 (Eulerian Formulation). *Denote $\mathcal{P}(\delta) := \{x \in \mathcal{P} : \|x\| \leq \delta\}$ for $\delta \geq 0$ and*

$$\delta_{\min} := \min\{\delta \geq 0 : \mathcal{P}(\delta) \neq \emptyset\}, \quad \delta_{\max} := \min\left\{\delta \geq 0 : \min_{x \in \mathcal{P}(\delta)} \langle c, x \rangle = \min_{x \in \mathcal{P}} \langle c, x \rangle\right\}.$$

Then

$$D := \left\{ \delta > \delta_{\min} : \operatorname{argmin}_{x \in \mathcal{P}(\delta)} \langle c, x \rangle \cap \operatorname{argmin}_{x \in \mathcal{P}} \langle c, x \rangle = \emptyset \right\} = (\delta_{\min}, \delta_{\max}).$$

For each $\delta \in \overline{D} = [\delta_{\min}, \delta_{\max}]$, the problem

$$\inf_{x \in \mathcal{P}(\delta)} \langle c, x \rangle \quad \text{has a unique minimizer } x^\delta. \quad (3)$$

Moreover, $[\delta_{\min}, \delta_{\max}] \ni \delta \mapsto x^\delta$ is continuous and $\|x^\delta\| = \delta$. In particular, $x^{\delta_{\min}}$ is the singleton $\mathcal{P}(\delta_{\min})$, and $x^{\delta_{\max}}$ is the minimum-norm solution of $\min_{x \in \mathcal{P}} \langle c, x \rangle$.

Proof. Compactness of $\mathcal{P}$ implies that the minima defining $\delta_{\min}$ and $\delta_{\max}$ are indeed attained. Moreover, the identity $D = (\delta_{\min}, \delta_{\max})$ follows directly from the definition of $\delta_{\max}$.

Fix $\delta \in D$. Let $x^\delta \in \operatorname{argmin}_{x \in \mathcal{P}(\delta)} \langle c, x \rangle$ and $x^* \in \operatorname{argmin}_{x \in \mathcal{P}} \langle c, x \rangle$. If $\|x^\delta\| < \delta$, then $x := \lambda x^\delta + (1 - \lambda)x^* \in \mathcal{P}(\delta)$ for sufficiently large $\lambda \in (0, 1)$. By optimality it follows that $x^\delta \in \operatorname{argmin}_{x \in \mathcal{P}} \langle c, x \rangle$, meaning that $\delta \notin D$. Thus, for $\delta \in D$, we have $\|x^\delta\| = \delta$. In particular, the set $\operatorname{argmin}_{x \in \mathcal{P}(\delta)} \langle c, x \rangle$ is contained in the sphere $\{x : \|x\| = \delta\}$. As the set is also convex, it must be a singleton by the strict convexity of $\|\cdot\|$. Continuity of $\delta \mapsto x^\delta$ is straightforward.

It remains to deal with the boundary cases. For $\delta = \delta_{\min}$, we note that $\|x^\delta\| < \delta$ is trivially ruled out, hence we can conclude as above. Clearly $\{x^{\delta_{\min}}\} = \mathcal{P}(\delta_{\min})$. Define x^* as the minimum-norm solution of $\min_{x \in \mathcal{P}} \langle c, x \rangle$, which is unique by strict convexity. Clearly $\|x^*\| = \delta_{\max}$. Thus when $\delta = \delta_{\max}$, we must have $x^\delta = x^*$ for any $x^\delta \in \operatorname{argmin}_{x \in \mathcal{P}(\delta)} \langle c, x \rangle$. Continuity and $\|x^\delta\| = \delta$ are again straightforward. $\square$

The next lemma recalls the standard projection theorem (e.g., [8, Theorem 5.2]).

Lemma 2.2 (Projection). *Let $\emptyset \neq K \subseteq \mathbb{R}^d$ be closed and convex. Given $x \in \mathbb{R}^d$, there exists a unique $x_K \in K$, called the projection of x onto K and denoted $x_K = \operatorname{proj}_K(x)$, such that*

$$\|x - x_K\| = \inf_{x' \in K} \|x - x'\|.$$

Moreover, x_K is characterized within K by

$$\langle x - x_K, x' - x_K \rangle \leq 0 \quad \text{for all } x' \in K. \quad (4)$$

If $x_K \in \operatorname{ri} K$, then $x_K = \operatorname{proj}_{\operatorname{aff} K}(x)$ and (4) can be sharpened to

$$\langle x - x_K, x' - x_K \rangle = 0 \quad \text{for all } x' \in K. \quad (5)$$

In particular, (5) holds when K is an affine subspace. In that case, $x \mapsto \operatorname{proj}_K(x)$ is affine.

Below, we often use $\operatorname{proj}_K(0)$ as a convenient notation for $\operatorname{argmin}_{x \in K} \|x\|$, the minimum-norm point of K . In the Lagrangian formulation of our regularized linear program, the solution can be expressed as a projection onto $\mathcal{P}$ as follows.

Lemma 2.3 (Lagrangian Formulation). *Define*

$$x_\eta := \operatorname{proj}_{\mathcal{P}}(-\eta c), \quad \eta \in [0, \infty), \quad x_\infty := \lim_{\eta \rightarrow \infty} x_\eta. \quad (6)$$

Then

$$x_\eta = \operatorname{argmin}_{x \in \mathcal{P}} \langle c, x \rangle + \frac{1}{2\eta} \|x\|^2, \quad \eta \in (0, \infty), \quad (7)$$

$$x_0 = \operatorname{argmin}_{x \in \mathcal{P}} \|x\|, \quad x_\infty = \operatorname{argmin}_{x' \in \operatorname{argmin}_{x \in \mathcal{P}} \langle c, x \rangle} \|x'\|. \quad (8)$$

The limit $x_\eta \rightarrow x_\infty$ is stationary; i.e., there exists $\bar{\eta} \in \mathbb{R}_+$ such that $x_\eta = x_\infty$ for all $\eta \geq \bar{\eta}$.

Proof. Let $\eta \in (0, \infty)$. Then (7) follows from

$$\frac{1}{2\eta} \|- \eta c - x\|^2 = \frac{\eta}{2} \|c\|^2 + \langle c, x \rangle + \frac{1}{2\eta} \|x\|^2.$$

The first claim in (8) is trivial. For the second claim and the stationary convergence $x_\eta \rightarrow x_\infty$ (which will not be used directly), see [28, Theorem 2.1], or [21] for the exact threshold $\bar{\eta}$. $\square$

The algorithm of [23] solves the problem of projecting a point onto a polyhedron, hence can be used to find $x_\eta = \operatorname{proj}_{\mathcal{P}}(-\eta c)$ numerically. The solutions x^δ of the Eulerian formulation (3) and x_η of the Lagrangian formulation (6) are related as follows.

Lemma 2.4 (Euler $\leftrightarrow$ Lagrange). *Given $\eta \in [0, \infty]$, there exists a unique $\delta = \delta(\eta) \in \overline{D}$ such that $x_\eta = x^\delta$, namely $\delta = \|x_\eta\|$. The function $\eta \mapsto \delta(\eta)$ is nondecreasing.*

Conversely, let $\delta \in \overline{D}$. Then there exists $\eta \in [0, \infty)$ (possibly non-unique) such that $x_\eta = x^\delta$. Define $\eta(\delta)$ as the minimal η with $x_\eta = x^\delta$.³ Then $\delta \mapsto \eta(\delta)$ is strictly increasing on $\overline{D}$.

Proof. We first prove the second part. Let $\delta \in \overline{D}$ and recall that $\|x^\delta\| = \delta$. Note also that $\eta \mapsto x_\eta$ is continuous and that $\eta \mapsto \|x_\eta\|$ is nondecreasing, with range $[\|x_0\|, \|x_\infty\|] = \overline{D}$. Hence there exists $\eta = \eta(\delta)$ such that $\|x_\eta\| = \delta$. We see from (7) that x_η minimizes $\langle c, x \rangle$ among all $x \in \mathcal{P}(\delta)$, which is to say that $x_\eta = x^\delta$.

To see the first part, fix $\eta \in [0, \infty]$ and define $\delta(\eta) := \|x_\eta\|$. The above shows that $x_\eta = x^{\delta(\eta)}$. Moreover, $\delta(\eta)$ must be unique with that property, due to $\|x^\delta\| = \delta$. $\square$

3 Abstract Result

Recall that x^δ and x_η denote the solutions of the Eulerian (3) and Lagrangian (6) formulation, respectively. Lemma 2.4 shows that the curves $(x^\delta)_{\delta_{\min} \leq \delta \leq \delta_{\max}}$ and $(x_\eta)_{\eta \geq 0}$ are reparametrizations of one another. In particular, they trace out the same trajectory, and only the trajectory matters for the subsequent definition.

Definition 3.1. Let $\mathcal{P} \subset \mathbb{R}^d$ be a polytope and $c \in \mathbb{R}^d$. We say that $\mathcal{P}$ is *c-monotone* if for any face F of $\mathcal{P}$,

$$x^\delta \in F \implies x^{\delta'} \in F \quad \text{for all } \delta' \geq \delta \text{ in } [\delta_{\min}, \delta_{\max}], \quad \text{or equivalently if} \quad (9)$$

$$x_\eta \in F \implies x_{\eta'} \in F \quad \text{for all } \eta' \geq \eta \text{ in } [0, \infty]. \quad (10)$$

We say that $\mathcal{P}$ is *monotone* if it is *c-monotone* for all $c \in \mathbb{R}^d$.

³This is merely for concreteness; any choice of selector will do.

This definition means that the “support” of x^δ is monotonically decreasing in δ , in the following sense. For any $x \in \mathcal{P}$, there is a unique face $F = F(x)$ such that $x \in \text{ri } F$; moreover, $x \in \text{ri } F$ if and only if x is a convex combination of the vertices $\text{ext } F$ with strictly positive weights [9, Exercise 3.1, Theorem 5.6]. Thus $\text{ext } F(x)$ is a notion of support for x , and then the property (9) indeed means that the support of x^δ is monotonically decreasing for inclusion. When $\mathcal{P}$ is the unit simplex $\{x : x_i \geq 0, \sum_i x_i = 1\}$, then $\text{ext } F(x)$ boils down to the usual measure-theoretic notion of support, namely $\{i : x_i > 0\}$ for $x = (x_1, \dots, x_d) \in \mathcal{P}$. This identification breaks down when $\mathcal{P}$ is the Birkhoff polytope, but monotonicity is nevertheless equivalent for the two notions of support (see Lemma 4.2).⁴ Next, we characterize monotonicity in geometric terms.

Theorem 3.2. *A polytope $\mathcal{P}$ is monotone (Definition 3.1) if and only if*

(H1) $\text{proj}_{\mathcal{P}}(0) \in \text{ri } \mathcal{P}$ and

(H2) $\text{proj}_{\text{aff } F}(0) \in F$ for each face $\emptyset \neq F \subset \mathcal{P}$.

The two conditions are similar, with the requirement for the proper faces F being less stringent than for the improper face $\mathcal{P}$: while $\text{proj}_F(0) \in \text{ri } F$ is a sufficient condition for (H2), the condition (H2) includes the boundary case where $\text{proj}_F(0) = \text{proj}_{\text{aff } F}(0) \in \text{rbd } F$. We shall see that (H1) and (H2) play different roles in the proof. Simple examples show that changing $\text{ri } \mathcal{P}$ to $\mathcal{P}$ in (H1) or F to $\text{ri } F$ in (H2) invalidates the equivalence in Theorem 3.2.

3.1 Proof of Theorem 3.2

We first prove the “if” implication, using the Lagrangian formulation (Lemma 2.3). This parametrization is convenient because $\eta \mapsto x_\eta$ is piecewise affine. While the latter is well-known (even for some more general norms, see [19] and the references therein); we detail the statement for completeness.

Lemma 3.3. *Let $\mathcal{P} \subseteq \mathbb{R}^d$ be a polytope, $c \in \mathbb{R}^d$ and let x_η be the solution of (2) with $\eta > 0$. The curve $[0, \infty) \ni \eta \mapsto x_\eta$ is piecewise affine. The affine pieces correspond to faces of $\mathcal{P}$ as follows. Fix $\eta_0 \in [0, \infty)$ and let F be the unique face of $\mathcal{P}$ such that $x_{\eta_0} \in \text{ri } F$. Let $I = \{\eta : x_\eta \in \text{ri } F\}$. Then I is an interval containing η_0 and $\bar{I} \ni \eta \mapsto x_\eta \in F$ is affine. In particular, the curve $\eta \mapsto x_\eta$ does not return to $\text{ri } F$ after leaving it.*

Proof. Consider $\eta_1 < \eta_2$ in I and let $H = \text{aff } F$. As $x_{\eta_i} \in \text{ri } F$, we have $x_{\eta_i} = \text{proj}_F(-\eta_i c) = \text{proj}_H(-\eta_i c)$. Since H is an affine subspace, the curve $\eta \mapsto \text{proj}_H(-\eta c)$ is affine. In particular, convexity of $\text{ri } F$ implies that $\text{proj}_H(-\eta c) \in \text{ri } F$ for all $\eta \in [\eta_1, \eta_2]$. Thus I is an interval and $I \ni \eta \mapsto x_\eta \in \text{ri } F$ is affine, and this extends to the closure by continuity. $\square$

Recall that $\text{rbd } K := K \setminus \text{ri } K$ denotes the relative boundary of $K \subset \mathbb{R}^d$. Consider the first time the curve $\eta \mapsto x_\eta$ touches a given face F of $\mathcal{P}$. That point may lie in $\text{ri } F$ or in $\text{rbd } F$. As seen in Lemma 3.5 below, the boundary situation indicates that $\eta \mapsto x_\eta$ left another face F_* when it entered $F \not\subset F_*$, meaning that $\mathcal{P}$ is not monotone. Hence, we analyze that situation in detail in the next lemma.

⁴Clearly this equivalence does not hold in general. E.g., when $\mathcal{P} \subset (0, 1)^d$, all points have the same measure-theoretic support.

Lemma 3.4. Let $\mathcal{P} \subseteq \mathbb{R}^d$ be a polytope, $c \in \mathbb{R}^d$ and let x_η be the solution of (2) with $\eta > 0$. Let F be a face of $\mathcal{P}$ and $I = \{\eta : x_\eta \in \text{ri } F\} \neq \emptyset$. Let $\eta_0 := \inf I$. If $x_{\eta_0} \in \text{rbd } F$, then exactly one of the following holds true:

(i) $\text{proj}_{\text{aff } F}(0) \notin F$;

(ii) $\eta_0 = 0$, and then $x_{\eta_0} = x_0 = \text{proj}_{\mathcal{P}}(0) = \text{proj}_{\text{aff } F}(0) \in \text{rbd } F$. In particular, $x_0 \notin \text{ri } \mathcal{P}$.

Proof. Recall from Lemma 3.3 that I is an interval and $\bar{I} \ni \eta \mapsto x_\eta \in F$ is affine. Let $H = \text{aff } F$ and consider $z_\eta := \text{proj}_H(-\eta c)$. Recalling Lemma 2.2, we have for $\eta \in I$ that $x_\eta = \text{proj}_{\mathcal{P}}(-\eta c) = \text{proj}_F(-\eta c) = \text{proj}_H(-\eta c) = z_\eta$ and in particular $z_\eta \in \text{ri } F$.

As $x_\eta = z_\eta$ for $\eta \in I$ and both curves are continuous, it follows that $x_{\eta_0} = z_{\eta_0}$. Therefore, our assumption that $x_{\eta_0} \in \text{rbd } F$ implies that $z_{\eta_0} \in \text{rbd } F$.

Since H is an affine space, $\mathbb{R}_+ \ni \eta \mapsto z_\eta$ is affine. We have $z_{\eta_0} \in \text{rbd } F$ and $z_\eta \in \text{ri } F$ for all $\eta \in (\eta_0, \eta_1)$, where $\eta_1 := \sup I > \eta_0$. (Note that I cannot be a singleton when $z_{\eta_0} \in \text{rbd } F$.) As F is convex, it follows that $z_\eta \notin F$ for $0 \leq \eta < \eta_0$. If $z_0 \notin F$, we are in the first case. Whereas if $z_0 \in F$, it follows that $\eta_0 = 0$. Thus $x_{\eta_0} = x_0 = \text{proj}_{\mathcal{P}}(0) = \text{proj}_F(0)$. Moreover, $z_{\eta_0} = x_{\eta_0} \in F$ implies $x_{\eta_0} = \text{proj}_{\text{aff } F}(0)$. $\square$

Lemma 3.5. Let $\mathcal{P} \subseteq \mathbb{R}^d$ be a polytope, $c \in \mathbb{R}^d$ and let x_η be the solution of (2) with $\eta > 0$. Let $\eta_* \geq 0$ and let F_* be the unique face of $\mathcal{P}$ with $x_{\eta_*} \in \text{ri } F_*$. Suppose there exists $\eta > \eta_*$ such that $x_\eta \notin F_*$ and let $\eta_0 = \max\{\eta \geq \eta_* : x_\eta \in F_*\}$. For sufficiently small $\eta_1 > \eta_0$, there is a unique face F such that $x_\eta \in \text{ri } F$ for all $\eta \in (\eta_0, \eta_1)$.⁵ If $x_0 := \text{proj}_{\mathcal{P}}(0) \in \text{ri } \mathcal{P}$, then F satisfies $\text{proj}_{\text{aff } F}(0) \notin F$.

Proof. For $x \in \mathcal{P}$, let $F(x)$ denote the unique face such that $x \in \text{ri } F(x)$. By Lemma 3.3, the map $(\eta_0, \infty) \ni \eta \mapsto F(x_\eta)$ has finitely many values and the preimage of each value is an interval. Hence the map must be constant on (η_0, η_1) for sufficiently small $\eta_1 > \eta_0$. Let $F = F(x_\eta)$, $\eta \in (\eta_0, \eta_1)$ be the corresponding face. As F_* and F are faces with $F_* \not\subset F$, we have $F \cap F_* \subset \text{rbd } F$. Thus $x_{\eta_0} \in F \cap F_*$ implies $x_{\eta_0} \in \text{rbd } F$, and now Lemma 3.4 applies. $\square$

Proof of Theorem 3.2. Step 1: (H1), (H2) $\Rightarrow$ monotone. If $\mathcal{P}$ is not monotone, then Lemma 3.5 shows that either (H1) or (H2) must be violated.

Step 2: Not (H1) $\Rightarrow$ not monotone. Suppose that $x_0 := \text{proj}_{\mathcal{P}}(0) \notin \text{ri } \mathcal{P}$. As $x_0 \in \mathcal{P} \setminus (\text{ri } \mathcal{P})$, there exists a hyperplane $H = \{x \in \mathbb{R}^d : \langle a, x \rangle = b\}$ with

$$x_0 \in H, \quad \mathcal{P} \subseteq \{x \in \mathbb{R}^d : \langle a, x \rangle \leq b\}, \quad \text{ri } \mathcal{P} \subseteq \{x \in \mathbb{R}^d : \langle a, x \rangle < b\}.$$

Define $c = a$ and $F = H \cap \mathcal{P}$ and $\delta_0 = \|x_0\|$, and denote $\bar{B}_\delta = \{x \in \mathbb{R}^d : \|x\| \leq \delta\}$. Then $\mathcal{P} \cap \bar{B}_{\delta_0} = \{x_0\}$ and in particular $x^{\delta_0} = x_0$. Moreover,

$$\langle c, x \rangle = b \quad \text{for all } x \in F, \quad \text{while} \quad \langle c, x \rangle < b \quad \text{for all } x \in \text{ri } \mathcal{P}. \quad (11)$$

For any $\delta > \delta_0$, we have $\bar{B}_\delta \cap \text{ri } \mathcal{P} \neq \emptyset$, and then (11) implies $x^\delta \notin F$. In particular, $\mathcal{P}$ is not monotone for the cost c .

⁵When traveling along the curve $(x_\eta)_{\eta \geq \eta_*}$, F is the first face outside F_* whose interior is visited.

Step 3: Not (H2) $\Rightarrow$ not monotone. Given the assumption, there exists a face F of $\mathcal{P}$ such that $\text{proj}_{\text{aff } F}(0)$ does not belong to F . Set $x_{\min} := \text{proj}_F(0)$. As F is a face of $\mathcal{P}$, there is a hyperplane $H = \{x \in \mathbb{R}^d : \langle a, x \rangle = b\}$ such that

$$F = H \cap \mathcal{P} \quad \text{and} \quad \mathcal{P} \subset \{x \in \mathbb{R}^d : \langle a, x \rangle \geq b\}.$$

Define $c = a - x_{\min}$ and $\delta_1 = \|x_{\min}\|$. Then for every $x \in \mathcal{P} \cap \overline{B}_{\delta_1}$,

$$\langle c, x \rangle = \langle a, x \rangle - \langle x_{\min}, x \rangle \geq b - \|x\| \|x_{\min}\| \geq b - \|x_{\min}\|^2 = \langle c, x_{\min} \rangle,$$

showing that $x^{\delta_1} = x_{\min}$. Recall that $x_{\min} := \text{proj}_F(0)$, so that $\langle x_{\min}, x - x_{\min} \rangle \geq 0$ for all $x \in F$ (Lemma 2.2). Thus $x_{\min} \neq \text{proj}_{\text{aff } F}(0)$ implies that

$$F' := \{x \in \mathbb{R}^d : \langle x_{\min}, x - x_{\min} \rangle = 0\} \cap F$$

is a proper face of F —and *a fortiori* a face of $\mathcal{P}$ —such that

$$\emptyset \neq F \setminus F' \subset \{x \in \mathbb{R}^d : \langle x_{\min}, x - x_{\min} \rangle > 0\}.$$

We claim that $x^\delta \notin F'$ for all $\delta > \delta_1$. Indeed, for any $x' \in F'$ it holds that

$$\langle c, x' \rangle = \langle a, x' \rangle - \langle x_{\min}, x' \rangle = b - \langle x_{\min}, x' \rangle = \langle c, x_{\min} \rangle,$$

whereas for any $x \in F \setminus F'$, it holds that

$$\langle c, x \rangle = b - \langle x_{\min}, x \rangle = \langle c, x_{\min} \rangle - \langle x_{\min}, x - x_{\min} \rangle < \langle c, x_{\min} \rangle.$$

In summary, $\langle c, x \rangle < \langle c, x' \rangle$ for all $x \in F \setminus F'$ and $x' \in F'$. In view of $x^{\delta_1} = x_{\min} \in F'$, it follows that $x^\delta \notin F'$ for all $\delta > \delta_1$ and in particular that $\mathcal{P}$ is not monotone for the cost c . $\square$

4 Applications

4.1 Soft-Min

To begin with a straightforward example, let $\mathcal{P} = \Delta$ be the unit simplex in Euclidean space $\mathbb{R}^d$ and $c = (c_1, \dots, c_d) \in \mathbb{R}^d$. Then

$$\min_{x \in \Delta} \langle c, x \rangle = \min_{1 \leq i \leq d} c_i$$

corresponds to finding the minimum value of c . More specifically, $x = (x_1, \dots, x_d)$ is a minimizer if and only if x is supported in $\text{argmin } c$; i.e., $\text{spt } x = \{i : x_i \neq 0\} \subset \text{argmin } c$. When this linear program is regularized with the entropy of x , we obtain the usual soft-min (counterpart of log-sum-exp) which gives large weights to the small values of c but non-zero weights to all values. The quadratic regularization,

$$x^\delta = \underset{x \in \Delta : \|x\| \leq \delta}{\text{argmin}} \langle c, x \rangle \quad \text{or} \quad x_\eta = \underset{x \in \Delta}{\text{argmin}} \langle c, x \rangle + \frac{1}{2\eta} \|x\|^2, \quad (12)$$

yields a sparse soft-min: as $\delta \rightarrow \delta_{\max}$ (or as $\eta \rightarrow \infty$), the support of the solution tends to $\text{argmin } c$. In this example, $\delta_{\min} = \frac{1}{\sqrt{d}}$ and $\delta_{\max} = \frac{1}{\sqrt{\#\text{argmin } c}}$.

Corollary 4.1. *The unit simplex $\Delta \subset \mathbb{R}^d$ is monotone for all $d \geq 1$; i.e., the support of the soft-min x^δ (or x_η) defined in (12) decreases monotonically to $\operatorname{argmin} c$ as $\delta \in [\delta_{\min}, \delta_{\max}]$ (or $\eta \geq 0$) increases.*

Proof. We have $\operatorname{ext} \Delta = \{e_1, \dots, e_d\}$, the canonical basis of $\mathbb{R}^d$. Let $F \subset \Delta$ be a face, then there exist $\{i_1, \dots, i_k\} \subseteq \{1, \dots, d\}$ such that $F = \operatorname{conv}(e_{i_1}, \dots, e_{i_k})$. Note that $\operatorname{proj}_F(0) = \sum_{j=1}^k \lambda_j e_{i_j}$, where $\lambda = (\lambda_1, \dots, \lambda_k)$ is the solution of

$$\min_{\lambda \in \mathbb{R}^d} \left\| \sum_{j=1}^k \lambda_j e_{i_j} \right\|^2 \quad \text{s.t.} \quad \sum_{j=1}^k \lambda_j = 1,$$

namely $\lambda = (k^{-1}, \dots, k^{-1})$. In particular, $\operatorname{proj}_F(0) = k^{-1} \sum_{j=1}^k e_{i_j} \in \operatorname{ri} F$. Now Theorem 3.2 applies. $\square$

4.2 Optimal Transport

We recall the quadratically regularized optimal transport problem

$$\inf_{\gamma \in \Gamma(\mu, \nu)} \int \hat{c}(x, y) d\gamma(x, y) + \frac{1}{2\eta} \left\| \frac{d\gamma}{d(\mu \otimes \nu)} \right\|_{L^2(\mu \otimes \nu)}^2 \quad (\text{QOT})$$

where $\Gamma(\mu, \nu)$ denotes the set of couplings of the probability measures (μ, ν) . We will be concerned with the discrete case as introduced in [6, 15, 18]. See also [26, 27, 30] for more general theory, [4, 17, ?] for asymptotic aspects, [16, 20, 22, 25, 34] for computational approaches and [25, 36] for some recent applications.

Fix $N \in \mathbb{N}$ and two sets of distinct points, $\{X_1, \dots, X_N\} \subset \mathbb{R}^D$ and $\{Y_1, \dots, Y_N\} \subset \mathbb{R}^D$. Let $\mu = \frac{1}{N} \sum_{i=1}^N \delta_{X_i}$ and $\nu = \frac{1}{N} \sum_{i=1}^N \delta_{Y_i}$ denote the associated empirical measures, and let $\hat{c} : \mathbb{R}^D \times \mathbb{R}^D \rightarrow \mathbb{R}$ be a (cost) function. We note that $\mu \otimes \nu = \frac{1}{N^2} \sum_{i,j=1}^N \delta_{(X_i, Y_j)}$ while $d\gamma/d(\mu \otimes \nu)$ is simply the ratio of the probability mass functions. Any coupling $\gamma \in \Gamma(\mu, \nu)$ can be identified with its matrix of probability weights

$$(\gamma_{i,j})_{i,j=1}^N \in \Gamma_N = \left\{ \gamma \in \mathbb{R}^{N \times N} : \sum_{i=1}^N \gamma_{i,j} = \frac{1}{N}, \sum_{j=1}^N \gamma_{i,j} = \frac{1}{N}, \gamma_{i,j} \geq 0 \right\}$$

via $\gamma = \sum_{i,j=1}^N \gamma_{i,j} \delta_{(X_i, Y_j)}$. Writing similarly $c_{ij} = \hat{c}(X_i, Y_j)$, (QOT) can be identified with the problem

$$\inf_{\gamma \in \Gamma_N} \langle c, \gamma \rangle + \frac{\varepsilon N}{2} \|\gamma\|^2.$$

Clearly $\Gamma_N = \frac{1}{N} \Pi_N$ where

$$\Pi_N = \left\{ \pi \in \mathbb{R}^{N \times N} : \sum_{i=1}^N \pi_{i,j} = 1, \sum_{j=1}^N \pi_{i,j} = 1, \pi_{i,j} \geq 0 \right\}$$

denotes the *Birkhoff polytope* of doubly stochastic matrices; see, e.g., [10]. The monotonicity (in the sense of Definition 3.1) of Π_N is equivalent to that of Γ_N .

The following result connects the measure-theoretic support $\operatorname{spt} \pi = \{(i, j) : \pi_{ij} > 0\}$ with the geometric notion discussed below Definition 3.1.

Lemma 4.2. *Let $\pi, \pi' \in \Pi_N$. Denote by $F(\pi)$ the unique face $F \subset \Pi_N$ such that $\pi \in \text{ri } F$. For $\pi, \pi' \in \Pi_N$, the following are equivalent:*

- (i) $\text{spt } \pi' \subset \text{spt } \pi$,
- (ii) $F(\pi') \subset F(\pi)$,
- (iii) if $\pi \in F$ for some face $F \subset \Pi_N$, then $\pi' \in F$.

Proof. The equivalence of (ii) and (iii) follows from the fact that $F(\pi)$ is the smallest face containing π [9, Theorem 5.6]. The equivalence of (i) and (iii) is deferred to Lemma 4.8 below. (The implication (ii) $\Rightarrow$ (i) is straightforward, and holds for any polytope $\mathcal{P}$ of measures.) $\square$

As a consequence of Lemma 4.2, we have the following equivalence.

Corollary 4.3. *Let $\gamma_{\eta, \hat{c}}$ denote the solution of (QOT) for $\hat{c} : \mathbb{R}^N \times \mathbb{R}^N \rightarrow \mathbb{R}$. The Birkhoff polytope is monotone if and only if the optimal support $\text{spt } \gamma_{\eta, \hat{c}}$ is monotone decreasing in η for all $\hat{c}$.*

Theorem 4.4. *The Birkhoff polytope Π_N is monotone for $1 \leq N \leq 4$ but not monotone for $N \geq 5$. In particular, when $N \geq 5$, the optimal support $\eta \mapsto \text{spt}(\gamma_{\eta, \hat{c}})$ is not monotone for some costs $\hat{c}$.*

Example 4.5. We exhibit an example for $N = 5$ where the support is not monotone. (This is not easily achieved by brute-force numerical experiment). Our starting point is the matrix A in (14) below, which is used to describe a face where (H2) fails. Step 3 in the proof of Theorem 3.2 then suggests to perturb $-A$ in a suitable direction to find a cost c exhibiting non-monotonicity. With some geometric considerations, this leads us to propose the cost matrix

$$c = \begin{bmatrix} -1.1 & -1 & -1 & -1 & -1 \\ -1 & -1.1 & 0 & 0 & 0 \\ -1 & 0 & -1.1 & 0 & 0 \\ -1 & 0 & 0 & -1.1 & 0 \\ -1 & 0 & 0 & 0 & -1.1 \end{bmatrix}.$$

(If one prefers a nonnegative cost matrix, one can obtain that by adding 1.1 to all entries. This translation does not change the solution γ_η to (QOT).) For $\eta = 2.5$, or equivalently regularization strength $1/(2\eta) = 0.2$, the corresponding problem (QOT) has an exact solution coupling γ_η with probability weights given by

$$(\gamma_{\eta=2.5}) = \begin{bmatrix} 0 & 0.05 & 0.05 & 0.05 & 0.05 \\ 0.05 & 0.15 & 0 & 0 & 0 \\ 0.05 & 0 & 0.15 & 0 & 0 \\ 0.05 & 0 & 0 & 0.15 & 0 \\ 0.05 & 0 & 0 & 0 & 0.15 \end{bmatrix}.$$

(This can be verified by noticing that the stated coupling is induced by the dual potentials $f = g = (-0.575, -0.175, -0.175, -0.175, -0.175)$; see for instance [30, Theorem 2.2 and Remark 2.3].) We observe in particular that the location (X_1, Y_1) is not in the support. On the other hand, because the diagonal of c features the smallest costs, the solution γ_∞ of the

unregularized transport problem is to put all mass on the diagonal; i.e., to transport all mass from X_i to Y_i for each i . Because of the stationary convergence mentioned in Lemma 2.3, that is also the solution of the regularized problem for large enough value of η (e.g., $\eta = 100$ will do):

$$(\gamma_{\eta=100}) = (\gamma_{\infty}) = \begin{bmatrix} 0.2 & 0 & 0 & 0 & 0 \\ 0 & 0.2 & 0 & 0 & 0 \\ 0 & 0 & 0.2 & 0 & 0 \\ 0 & 0 & 0 & 0.2 & 0 \\ 0 & 0 & 0 & 0 & 0.2 \end{bmatrix}.$$

In particular, (X_1, Y_1) is part of the support for large η , completing the example. Figure 2 shows in more detail the weight at (X_1, Y_1) as a function of (the inverse of) η .




Figure 2: Probability mass $\gamma_{\eta}(X_1, Y_1)$ plotted against $1/(2\eta)$, showing that (X_1, Y_1) is in the support for small and large values of η but outside for an intermediate interval.

Remark 4.6. The early work of [18] considered a minimum-cost flow problem with quadratic regularization that predates the optimal transport literature for this regularization. The authors point out that the solution can have non-monotone support already in the minimal setting of 2×2 points [18, Figure 1]. Similarly, it is not hard to obtain non-monotonicity if instead of $\mu \otimes \nu$ we use a different measure to define the L^2 -norm in (QOT), similarly as in [30, Example A.1]. In those examples, the mechanism causing non-monotonicity is straightforward and quite different than in the present work, where non-monotonicity arises only in higher dimensions.

4.3 Proof of Lemma 4.2 and Theorem 4.4

The celebrated Birkhoff's theorem [5] (or, e.g., [?, p. 30]) states that the vertices $\text{ext } \Pi_N$ are the permutation matrices of $\{1, \dots, N\}$; that is, the elements of Π_N with binary entries. Following [12], the faces of Π_N can be described using the so-called permanent function. If A is a binary $N \times N$ matrix, its *permanent* $\text{per}(A) \in \mathbb{N}$ is defined as the number of permutation matrices P with $P \leq A$ (meaning that $P_{i,j} \leq A_{i,j}$ for all i, j). Denoting by $\text{conv}(\cdot)$ the convex hull, the following characterization is contained in [12, Theorem 2.1].

Lemma 4.7. *Let $t \in \mathbb{N}$ and $P^{(1)}, \dots, P^{(t)} \in \text{ext}(\Pi_N)$. Let $A = (A_{i,j})_{i,j=1}^N$ be the matrix such that $A_{i,j} = 1$ if there exists $s \in \{1, \dots, t\}$ with $P_{i,j}^{(s)} = 1$ and $A_{i,j} = 0$ otherwise. Then $\text{conv}(\{P^{(1)}, \dots, P^{(t)}\})$ is a face of Π_N if and only if $\text{per}(A) = t$.*

We use Lemma 4.7 to relate faces with the distribution of zeros.

Lemma 4.8. *Let $\pi, \pi' \in \Pi_N$. The following are equivalent:*

- (i) *there exists (i, j) such that $\pi_{i,j} = 0$ and $\pi'_{i,j} > 0$;*
- (ii) *there exists a face F of Π_N such that $\pi \in F$ and $\pi' \notin F$.*

Proof. (i) $\Rightarrow$ (ii): Let (i, j) be such that $\pi_{i,j} = 0$ and $\pi'_{i,j} > 0$. Assume w.l.o.g. that $(i, j) = (1, 1)$. Then π belongs to the set

$$F = \{\pi \in \Pi_N : \langle \pi, A \rangle = N\},$$

where A is the matrix

$$A = \begin{pmatrix} 0 & 1 & 1 & \cdots & 1 \\ 1 & 1 & 1 & \cdots & 1 \\ 1 & 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \cdots & 1 \end{pmatrix}.$$

As $\Pi_N \subset \{\pi \in \Pi_N : \langle \pi, A \rangle \leq N\}$, we see that F is a face. Clearly π' does not belong to F , proving the claim.

(ii) $\Rightarrow$ (i): Let $P^{(1)}, \dots, P^{(t)} \in \text{ext} \Pi_N$ be such that $F = \text{conv}(\{P^{(1)}, \dots, P^{(t)}\})$ and let $A = (A_{i,j})_{i,j=1}^N$ be the matrix such that $A_{i,j} = 1$ if there exists $s \in \{1, \dots, t\}$ such that $P_{i,j}^{(s)} = 1$ and $A_{i,j} = 0$ otherwise. We have $\text{per}(A) = t$ by Lemma 4.7. As an element of Π_N , π' can be written as a convex combination of permutation matrices, i.e.,

$$\pi' = \sum_{P \in \text{ext}(\Pi_N)} \lambda_P P, \quad \text{with } \lambda_P \in [0, 1] \text{ and } \sum_{P \in \text{ext}(\Pi_N)} \lambda_P = 1.$$

As $\pi' \notin F$, there exists $P \in \text{ext}(\Pi_N) \setminus \{P^{(1)}, \dots, P^{(t)}\}$ with $\lambda_P > 0$. Using the fact that $\text{per}(A) = t$, we derive the existence of (i, j) such that $P_{i,j} = 1$ but $A_{i,j} = 0$. In particular, $\pi'_{i,j} \geq \lambda_P > 0$ but $\pi_{i,j} = 0$, proving the claim. $\square$

Lemma 4.9. *Let $N \geq 1$. Define the permutation matrices $P^k \in \mathbb{R}^{N \times N}$, $1 \leq k \leq N$ by*

$$P_{i,j}^k = \begin{cases} 1 & \text{if } i = 1 \text{ and } j = k, \\ 1 & \text{if } i = k \text{ and } j = 1, \\ 1 & \text{if } 1 \neq i = j \neq k, \\ 0 & \text{otherwise.} \end{cases}$$

That is, P^k permutes the first with the k -th element. Then $F := \text{conv}(P^1, \dots, P^N)$ is a face of Π_N and

$$\text{proj}_{\text{aff } F}(0) = \sum_{i=1}^N \lambda_i P^i \quad \text{for } \lambda_1 = \frac{4-N}{N+2}, \quad \lambda_2 = \dots = \lambda_N = \frac{2}{N+2}. \quad (13)$$

As a consequence,

$$\text{proj}_{\text{aff } F}(0) \in \begin{cases} \text{ri } F, & \text{if } N = 1, 2, 3, \\ \text{rbd } F, & \text{if } N = 4, \\ \mathbb{R}^{N \times N} \setminus F & \text{if } N \geq 5, \end{cases}$$

and in particular (H2) is violated for $N \geq 5$.

Proof. Define $A \in \mathbb{R}^{N \times N}$ by

$$A_{i,j} = \begin{cases} 1 & \text{if } i = j, \\ 1 & \text{if } i = 1 \text{ or } j = 1, \\ 0 & \text{otherwise,} \end{cases} \quad (14)$$

so that A is the entry-wise maximum $A = \max(P^1, \dots, P^N)$. One readily verifies that $\text{per}(A) = N$. Hence by Lemma 4.7, $F := \text{conv}(P^1, \dots, P^N)$ is a face of Π_N . To determine $\text{proj}_{\text{aff } F}(0)$, we consider the minimization problem

$$\min_{\lambda=(\lambda_1, \dots, \lambda_N) \in \mathbb{R}^N} \left\| \sum_{i=1}^N \lambda_i P^i \right\|^2 \quad \text{s.t.} \quad \sum_{i=1}^N \lambda_i = 1,$$

where $\|\cdot\|$ denotes the Frobenius norm. The Lagrangian for this problem is

$$L(\lambda, \rho) = \lambda_1^2 + 2 \sum_{j=2}^N \lambda_j^2 + \sum_{j=2}^N \left(\sum_{k \neq j} \lambda_k \right)^2 + \rho \left(1 - \sum_{j=1}^N \lambda_j \right).$$

Here the first three terms arise from the objective: the first is the value $(\sum_{i=1}^N \lambda_i P^i)_{1,1}$, the second is the sum of all remaining terms in the first row and column, and the third is the sum of the remaining terms in the diagonal. Finally, the fourth term arises from the constraint. The resulting optimality equations are

$$\begin{aligned} \rho &= 2N\lambda_1 + 2(N-2) \sum_{j=2}^N \lambda_j, & \sum_{j=1}^N \lambda_j &= 1, \\ \rho &= 2N\lambda_i + 2(N-2)\lambda_1 + 2(N-3) \sum_{\substack{j=2 \\ j \neq i}}^N \lambda_j & \text{for } 2 \leq i \leq N. \end{aligned}$$

By symmetry, the unique optimal λ satisfies $\lambda_i = \lambda_j =: \lambda_0$ for all $i, j \geq 2$, so that

$$\begin{aligned} 2N\lambda_0 + 2(N-2)\lambda_1 + 2(N-2)(N-3)\lambda_0 &= 2N\lambda_1 + 2(N-1)(N-2)\lambda_0, \\ \lambda_1 + (N-1)\lambda_0 &= 1 \end{aligned}$$

and finally

$$\lambda_1 = \lambda_0 \frac{(4-N)}{2}, \quad \lambda_1 + (N-1)\lambda_0 = 1.$$

Solving for λ_1 and λ_0 yields (13). Note that $\lambda_0 > 0$ for any N whereas $\lambda_1 > 0$ for $N \in \{1, 2, 3\}$, $\lambda_1 = 0$ for $N = 4$ and $\lambda_1 < 0$ for $N \geq 5$, implying the last claim. $\square$

The matrix A of (14) is inspired by the 4×4 matrix in [11, Example 1.5] where the authors are interested in a different problem (and, in turn, credit [1]).⁶ For the present question of monotonicity, this matrix yields a counterexample only for $N \geq 5$. As an aside, the subsequent proof that for $N = 4$, the face F considered in Lemma 4.9 is (up to symmetries) the only face where $\text{proj}_{\text{aff } F}(0) \in \text{rbd } F$, whereas all other faces F' satisfy $\text{proj}_{\text{aff } F'}(0) \in \text{ri } F'$.

Lemma 4.10. *Let $1 \leq N \leq 4$. Then Π_N is monotone.*

Proof. We verify (H1) and (H2). Note that $\text{proj}_{\Pi_N}(0)$ is the matrix with all entries equal to $1/N$ (corresponding to the product measure). This matrix is clearly in the relative interior of Π_N , showing (H1). The property (H2) is trivial for $N = 1, 2$. For $N = 3$ and $N = 4$, we give a computer-assisted proof in the interest of brevity.⁷⁸ Specifically, we generate all $N \times N$ permutation matrices and determine all families $\{P_1, \dots, P_m\}$ of permutation matrices that form the vertices of a nonempty face F by using the permanent function as in Lemma 4.7. There are 49 nonempty faces for $N = 3$ and 7443 for $N = 4$. For each face F , we can numerically compute $\text{proj}_{\text{aff } F}(0)$ or more specifically scalar coefficients $(\lambda_k)_{1 \leq k \leq m}$ such that $\text{proj}_{\text{aff } F}(0) = \sum_k \lambda_k P_k$ and $\sum_k \lambda_k \leq 1$. The coefficients λ_k are non-unique for some faces; in that case, we choose positive weights if possible. Note that as P_k are binary matrices and $N \leq 4$, the computation can be done with accuracy close to machine precision. It turns out that for $N = 3$, all coefficients satisfy $\lambda_k > 0.01$ (much larger than machine precision), establishing that $\text{proj}_{\text{aff } F}(0) \in \text{ri } F$ and in particular (H2). For $N = 4$, most faces have coefficients $\lambda_k > 0.01$, whereas for 96 faces, one coefficient is numerically close to zero, hence requiring an analytic argument. We verify that all those 96 faces are equivalent up to permutations of rows and columns, corresponding to a relabeling of the points X_i and Y_j . Specifically, they are all equivalent to the particular face analyzed in Lemma 4.9, where we have seen that $\text{proj}_{\text{aff } F}(0) \in F$ for $N \leq 4$ (and in fact $\text{proj}_{\text{aff } F}(0) \in \text{rbd } F$ for $N = 4$). We conclude that (H2) holds for all faces F when $N = 4$, completing the proof. $\square$

Proof of Theorem 4.4. In view of Theorem 3.2, the last statement of Lemma 4.9 implies that Π_N is not monotone for $N \geq 5$, whereas we have seen in Lemma 4.10 that Π_N is monotone for $1 \leq N \leq 4$. $\square$

5 On a Problem of Erdős

Let $A = (a_{ij})$ be an $N \times N$ matrix. Then its maximal trace is defined as

$$\max \text{tr } A := \max_{\sigma} \sum_{i=1}^N a_{i, \sigma(i)}$$

where σ ranges over all permutations of $\{1, \dots, N\}$. When $A \in \Pi_N$ (i.e., A is doubly stochastic) a result of the early paper [29] states that

$$\|A\|^2 \leq \max \text{tr } A, \quad (15)$$

⁶In [11], the conclusion is that “not every zero pattern of a fully indecomposable $(0, 1)$ -matrix is realizable as the zero pattern of a doubly stochastic matrix whose diagonal sums avoiding the 0’s are constant.”

⁷An analytic proof is also available, but requires us to go through 52 different cases for $N = 4$.

⁸The cases $N \leq 3$ can also be obtained as a corollary of the case $N = 4$. One can check directly that if the Birkhoff polytope Π_N is (not) monotone for some $N \in \mathbb{N}$, then $\Pi_{N'}$ is also (not) monotone for all $N' \leq (\geq) N$.

where $\|A\|^2 = \sum_{i,j} a_{ij}^2$ is the squared Frobenius norm. For a simple proof of (15), note that the maximal trace can be expressed as

$$\max \operatorname{tr} A = \max_{P \in \operatorname{ext} \Pi_N} \langle P, A \rangle \quad (16)$$

where P ranges over the set $\operatorname{ext} \Pi_N$ of all $N \times N$ permutation matrices (Birkhoff's theorem) and the inner product is the Frobenius one, $\langle B, A \rangle = \sum_{i,j} b_{ij} a_{ij}$ for $A = (a_{ij})$ and $B = (b_{ij})$. Consider the linear optimization problem $\max_{B \in \Pi_N} \langle B, A \rangle$. Its maximum must be attained at an extreme point. On the other hand, $A \in \Pi_N$, hence

$$\|A\|^2 = \langle A, A \rangle \leq \max_{B \in \Pi_N} \langle B, A \rangle = \max_{P \in \operatorname{ext} \Pi_N} \langle P, A \rangle = \max \operatorname{tr} A,$$

proving (15). Erdős posed the following problem; see also [7] for further background.

Question 5.1 (Erdős). For which doubly stochastic matrices A is inequality (15) saturated?

In formulas, the problem is to describe the set $\{A \in \Pi_N : \|A\|^2 = \max \operatorname{tr} A\}$. To simplify the discussion, let us introduce the following terminology.

Definition 5.2. We call $A \in \mathbb{R}^{N \times N}$ an *Erdős matrix* if it is doubly stochastic (i.e., $A \in \Pi_N$) and $\|A\|^2 = \max \operatorname{tr} A$.

Definition 5.3. Let F be a face of Π_N . We say that F is *centered* if $\operatorname{proj}_{\operatorname{aff} F}(0) \in F$.

In this terminology, Theorem 3.2 states that Π_N is monotone iff every face of Π_N is centered. The following theorem shows that Erdős matrices can be parametrized by the centered faces of Π_N ; in particular, we provide a geometric answer to Erdős' question. As a reminder, for $N \leq 4$, all faces of Π_N are centered, whereas for $N \geq 5$, not all faces are centered (see Theorem 4.4).

Theorem 5.4. For $A \in \mathbb{R}^{N \times N}$, the following are equivalent:

- (i) A is an Erdős matrix,
- (ii) $A = \operatorname{proj}_{\operatorname{aff} F}(0)$ for some centered face F of Π_N , and $A \neq \operatorname{proj}_{F'}(0)$ for any non-centered face F' containing F .

Proof. (i) $\Rightarrow$ (ii): Let A be an Erdős matrix and $\{P_1, \dots, P_k\} = \operatorname{argmax}_{P \in \operatorname{ext} \Pi_N} \langle P, A \rangle$. Then we have $\langle A, A \rangle = \langle P_i, A \rangle$ for $i = 1, \dots, k$ and $\langle A, A \rangle > \langle P, A \rangle$ for all $P \in \operatorname{ext} \Pi_N \setminus \{P_1, \dots, P_k\}$. Hence,

$$F := \operatorname{conv}(P_1, \dots, P_k) = \operatorname{argmax}_{P \in \Pi_N} \langle P, A \rangle$$

is a face containing A . As $\langle A, A \rangle = \langle P_i, A \rangle$ for $i = 1, \dots, k$, we deduce $A = \operatorname{proj}_{\operatorname{aff} F}(0)$; cf. Lemma 2.2. Recalling that A is doubly stochastic, this implies $A \in (\operatorname{aff} F) \cap \Pi_N = F$, and hence that F is centered.

To see the second part of (ii), suppose that $A = \operatorname{proj}_{F'}(0)$ for some face F' . When F' is not centered, it follows that $A \neq \operatorname{proj}_{\operatorname{aff} F'}(0)$. By Lemma 2.2, the combination of both facts implies that

$$0 < \langle A, P - A \rangle = \langle A, P \rangle - \|A\|^2$$

for some $P \in \text{ext } F'$, contradicting that A is Erdős. (We did not use the condition that $F' \supset F$. Thus, the theorem remains true if in (ii) we require $A \neq \text{proj}_{F'}(0)$ for all non-centered faces F' of Π_N , and not just the ones containing F .)

(ii) $\Rightarrow$ (i): Assume that $A = \text{proj}_{\text{aff } F}(0)$ for some centered face F of Π_N , hence also $A \in F$, and assume that A is not Erdős. We shall construct a face $F' \supset F$ with $A = \text{proj}_{F'}(0)$ that is not centered. Indeed, $A = \text{proj}_{\text{aff } F}(0)$ implies in particular that

$$F \subset \{\pi \in \Pi_N : \langle A, \pi - A \rangle = 0\}. \quad (17)$$

We claim that there exists a face F' of Π_N such that $F \subset F'$,

$$F' \subset \{\pi \in \Pi_N : \langle A, \pi - A \rangle \geq 0\} \quad \text{and} \quad \{\pi \in \Pi_N : \langle A, \pi - A \rangle > 0\} \cap F' \neq \emptyset. \quad (18)$$

This claim will complete the proof: Indeed, as $A \in F \subset F'$, the first part of (18) implies that $A = \text{proj}_{F'}(0)$. If F' were centered, this would yield $A = \text{proj}_{\text{aff } F'}(0)$, contradicting the second part of (18).

It remains to prove the claim. Let $\{\pi : \langle B, \pi - A \rangle = 0\}$ be a supporting hyperplane such that

$$\Pi_N \subset \{\pi : \langle B, \pi - A \rangle \leq 0\} \quad \text{and} \quad \Pi_N \cap \{\pi : \langle B, \pi - A \rangle = 0\} = F. \quad (19)$$

We define

$$H_\alpha = \{\pi : \langle \alpha B + (1 - \alpha)A, \pi - A \rangle \leq 0\}, \quad \alpha \in [0, 1]$$

and the function

$$[0, 1] \ni \alpha \mapsto \Phi(\alpha) := \begin{cases} 1 & \text{if } \Pi_N \subset H_\alpha, \\ 0 & \text{otherwise.} \end{cases}$$

Let $\pi \in \Pi_N$. As $\langle B, \pi - A \rangle \leq 0$, there are two cases: if $\langle A, \pi - A \rangle \leq 0$, then $\pi \in H_\alpha$ for all $\alpha \in [0, 1]$, whereas if $\langle A, \pi - A \rangle > 0$, then $\alpha \mapsto \langle \alpha B + (1 - \alpha)A, \pi - A \rangle$ is nonincreasing. It follows that $\alpha \mapsto H_\alpha \cap \Pi_N$ is nondecreasing (for inclusion) and hence that Φ is nondecreasing. Notice also that Φ is right-continuous (by continuity of the inner product) with $\Phi(1) = 1$, and that

$$\{\pi \in \Pi_N : \langle A, \pi - A \rangle \leq 0\} \subset H_\alpha \quad \text{for all } \alpha \in [0, 1]. \quad (20)$$

In view of (17) and (19), the latter implies that

$$F \subset (\text{rbd } H_\alpha) \cap \Pi_N \subset \{\pi \in \Pi_N : \langle A, \pi - A \rangle \geq 0\} \quad \text{for all } \alpha \in [0, 1]. \quad (21)$$

Define

$$\alpha^* := \min\{\alpha \in [0, 1] : \Phi(\alpha) = 1\} \quad \text{and} \quad F' := (\text{rbd } H_{\alpha^*}) \cap \Pi_N,$$

where the minimum is attained thanks to right-continuity and $\Phi(1) = 1$. The fact that $\Phi(\alpha^*) = 1$ yields that F' is a face of Π_N . Moreover, (21) for $\alpha = \alpha^*$ shows $F \subset F'$ and the first part of (18). It remains to show the second part of (18). Suppose for contradiction that it fails, then by (20) we have $F' \subset H_\alpha$ for all $\alpha \in [0, 1]$. Thus

$$F' := (\text{rbd } H_{\alpha^*}) \cap \Pi_N \subset H_\alpha \cap \Pi_N \quad \text{for all } \alpha \in [0, 1].$$

On the other hand, the definition of α^* implies⁹

$$(\text{rbd } H_{\alpha^*}) \cap \Pi_N \not\subset H_\alpha \cap \Pi_N \quad \text{for all } \alpha \in (\alpha^* - \epsilon, \alpha^*) \cap [0, 1]$$

for some $\epsilon > 0$. The two displays form a contradiction unless $\alpha^* = 0$. Hence $\Phi(0) = 1$; that is, $\Pi_N \subset H_0$. The latter means that $\langle A, \pi - A \rangle \leq 0$ for all $\pi \in \Pi_N$, which in turn implies that A is Erdős. This contradiction completes the proof. $\square$

As a consequence of Lemma 4.10, we deduce the following simpler statement for $N \leq 4$.

Corollary 5.5. *For $A \in \mathbb{R}^{N \times N}$ and $N \leq 4$, the following are equivalent:*

- (i) *A is an Erdős matrix,*
- (ii) *$A = \text{proj}_{\text{aff } F}(0)$ for some face F of Π_N .*

Arguably Theorem 5.4 is only a partial answer to Erdős' question, as it is not straightforward to list all centered faces of Π_N when N is large. To illustrate how Theorem 5.4 can be useful, we use this geometric point of view to answer a question recently posed in [7, Question 7.2]: Do all Erdős matrices have rational entries? (The main contribution of [7] is to list explicitly all 3×3 Erdős matrices, and it is observed that they have rational entries. Of course this list can also be recovered using our description.) Theorem 5.4 allows us to answer this affirmatively in a general and elegant way.

Corollary 5.6. *Erdős matrices have rational entries.*

Proof. Let A be an Erdős matrix. By Theorem 5.4, $A = \text{proj}_{\text{aff } F}(0)$ for some face F of Π_N . Let $P_1, \dots, P_k$ be the vertices of F and note that

$$A = \text{proj}_{\text{aff } F}(0) = \text{proj}_{\text{aff}\{0, P_2 - P_1, \dots, P_k - P_1\}}(-P_1) = \text{proj}_{\text{lin}\{P_2 - P_1, \dots, P_k - P_1\}}(-P_1),$$

the projection of $-P_1$ onto the linear span of $P_2 - P_1, \dots, P_k - P_1$. Eliminating some of the latter matrices if necessary, we may assume that they are linearly independent. If we see $N \times N$ matrices as vectors of length N^2 by stacking their columns, the projection admits the following formula, well known as the solution of the least squares linear regression problem:

$$A = (X^\top X)^{-1} X^\top (-P_1), \tag{22}$$

where X is the $N^2 \times (k-1)$ matrix with columns $P_2 - P_1, \dots, P_k - P_1$. Because P_i are permutation matrices, X has integer entries and it follows that $(X^\top X)^{-1}$ has rational entries (as can be seen from Cramer's rule for computing the inverse). It is now apparent from (22) that A also has rational entries. $\square$

⁹Specifically, let $S' \subset \text{ext } \Pi_N$ be the set of extreme points of Π_N that are contained in F' , or equivalently are contained in $\text{rbd } H_{\alpha^*}$. Note that all of $\text{ext } \Pi_N$ is contained in H_{α^*} as $\Phi(\alpha^*) = 1$. As $\text{ext } \Pi_N$ is a finite set and $\alpha \mapsto H_\alpha$ is continuous, there exists $\epsilon > 0$ such that $(\text{ext } \Pi_N \setminus S') \subset H_\alpha \cap \Pi_N$ for all $\alpha \in (\alpha^* - \epsilon, \alpha^*) \cap [0, 1]$. We choose such an ϵ . If $\alpha \in (\alpha^* - \epsilon, \alpha^*) \cap [0, 1]$, then as $\Phi(\alpha) = 0$, we know that $\text{ext } \Pi_N \not\subset H_\alpha$, and it follows that $S' \not\subset H_\alpha$; i.e., $F' \not\subset H_\alpha$.

References

- [1] E. Achilles. Doubly stochastic matrices with some equal diagonal sums. *Linear Algebra Appl.*, 22:293–296, 1978.
- [2] M. Agueh and G. Carlier. Barycenters in the Wasserstein space. *SIAM J. Math. Anal.*, 43(2):904–924, 2011.
- [3] J. Backhoff-Veraguas, D. Bartl, M. Beiglböck, and M. Eder. All adapted topologies are equal. *Probab. Theory Related Fields*, 178(3-4):1125–1172, 2020.
- [4] E. Bayraktar, S. Eckstein, and X. Zhang. Stability and sample complexity of divergence regularized optimal transport. *Bernoulli*, 31(1):213–239, 2025.
- [5] G. Birkhoff. Tres observaciones sobre el algebra lineal. *Univ. Nac. Tucumán. Revista A.*, 5:147–151, 1946.
- [6] M. Blondel, V. Seguy, and A. Rolet. Smooth and sparse optimal transport. volume 84 of *Proceedings of Machine Learning Research*, pages 880–889, 2018.
- [7] L. Bouthat, J. Mashreghi, and F. Morneau-Guérin. On a question of Erdős on doubly stochastic matrices. *Linear Multilinear Algebra*, 72(17):2823–2844, 2024.
- [8] H. Brezis. *Functional analysis, Sobolev spaces and partial differential equations*. Universitext. Springer, New York, 2011.
- [9] A. Brøndsted. *An introduction to convex polytopes*, volume 90 of *Graduate Texts in Mathematics*. Springer-Verlag, New York-Berlin, 1983.
- [10] R. A. Brualdi. *Combinatorial matrix classes*, volume 108 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge, 2006.
- [11] R. A. Brualdi and G. Dahl. Diagonal sums of doubly stochastic matrices. *Linear Multilinear Algebra*, 70(20):4946–4972, 2022.
- [12] R. A. Brualdi and P. M. Gibson. Convex polyhedra of doubly stochastic matrices. I. Applications of the permanent function. *J. Combinatorial Theory Ser. A*, 22(2):194–230, 1977.
- [13] J. A. Cuesta-Albertos and A. Tuero-Díaz. A characterization for the solution of the Monge-Kantorovich mass transference problem. *Statist. Probab. Lett.*, 16(2):147–152, 1993.
- [14] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems 26*, pages 2292–2300. 2013.
- [15] A. Dessein, N. Papadakis, and J.-L. Rouas. Regularized optimal transport and the rot mover’s distance. *J. Mach. Learn. Res.*, 19(15):1–53, 2018.
- [16] S. Eckstein and M. Kupper. Computation of optimal transport and related hedging problems via penalization and neural networks. *Appl. Math. Optim.*, 83(2):639–667, 2021.
- [17] S. Eckstein and M. Nutz. Convergence rates for regularized optimal transport via quantization. *Math. Oper. Res.*, 49(2):1223–1240, 2024.
- [18] M. Essid and J. Solomon. Quadratically regularized optimal transport on graphs. *SIAM J. Sci. Comput.*, 40(4):A1961–A1986, 2018.
- [19] M. Finzel and W. Li. Piecewise affine selections for piecewise polyhedral multifunctions and metric projections. *J. Convex Anal.*, 7(1):73–94, 2000.
- [20] A. Genevay, M. Cuturi, G. Peyré, and F. Bach. Stochastic optimization for large-scale optimal transport. In *Advances in Neural Information Processing Systems 29*, pages 3440–3448, 2016.
- [21] A. González-Sanz and M. Nutz. Quantitative convergence of quadratically regularized linear

programs. *Preprint arXiv:2408.04088v1*, 2024.

- [22] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville. Improved training of Wasserstein GANs. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 5769–5779, 2017.
- [23] W. W. Hager and H. Zhang. Projection onto a polyhedron that exploits sparsity. *SIAM J. Optim.*, 26(3):1773–1798, 2016.
- [24] S. Kolouri, S. R. Park, M. Thorpe, D. Slepcev, and G. K. Rohde. Optimal mass transport: Signal processing and machine-learning applications. *IEEE Signal Processing Magazine*, 34(4):43–59, 2017.
- [25] L. Li, A. Genevay, M. Yurochkin, and J. Solomon. Continuous regularized Wasserstein barycenters. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 17755–17765. Curran Associates, Inc., 2020.
- [26] D. Lorenz and H. Mahler. Orlicz space regularization of continuous optimal transport problems. *Appl. Math. Optim.*, 85(2):Paper No. 14, 33, 2022.
- [27] D. Lorenz, P. Manns, and C. Meyer. Quadratically regularized optimal transport. *Appl. Math. Optim.*, 83(3):1919–1949, 2021.
- [28] O. L. Mangasarian. Normal solutions of linear programs. *Math. Programming Stud.*, 22:206–216, 1984. Mathematical programming at Oberwolfach, II (Oberwolfach, 1983).
- [29] M. Markus and H. Minc. Diagonal products in doubly stochastic matrices. *Quart. J. Math. Oxford Ser. (2)*, 16:32–34, 1965.
- [30] M. Nutz. Quadratically regularized optimal transport: Existence and multiplicity of potentials. *SIAM J. Math. Anal.*, to appear, 2024. Preprint arXiv:2404.06847v1.
- [31] A. Paffenholz. Faces of Birkhoff polytopes. *Electron. J. Combin.*, 22(1):Paper 1.67, 36, 2015.
- [32] V. M. Panaretos and Y. Zemel. Statistical aspects of Wasserstein distances. *Annu. Rev. Stat. Appl.*, 6:405–431, 2019.
- [33] G. Peyré and M. Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019.
- [34] V. Seguy, B. B. Damodaran, R. Flamary, N. Courty, A. Rolet, and M. Blondel. Large scale optimal transport and mapping estimation. In *International Conference on Learning Representations*, 2018.
- [35] C. Villani. *Optimal transport, old and new*, volume 338 of *Grundlehren der Mathematischen Wissenschaften*. Springer-Verlag, Berlin, 2009.
- [36] S. Zhang, G. Mordant, T. Matsumoto, and G. Schiebinger. Manifold learning with sparse regularised optimal transport. *Preprint arXiv:2307.09816v1*, 2023.