

SOLUTIONS TO MULTI-MARGINAL OPTIMAL TRANSPORT PROBLEMS CONCENTRATED ON SEVERAL GRAPHS *

ABBAS MOAMENI¹ AND BRENDAN PASS²

Abstract. We study solutions to the multi-marginal Monge–Kantorovich problem which are concentrated on several graphs over the first marginal. We first present two general conditions on the cost function which ensure, respectively, that any solution must concentrate on either finitely many or countably many graphs. We show that local differential conditions on the cost, known to imply local d -rectifiability of the solution, are sufficient to imply a local version of the first of our conditions. We exhibit two examples of cost functions satisfying our conditions, including the Coulomb cost from density functional theory in one dimension. We also prove a number of results relating to the uniqueness and extremality of optimal measures. These include a sufficient condition on a collection of graphs for any competitor in the Monge–Kantorovich problem concentrated on them to be extremal, and a general negative result, which shows that when the problem is symmetric with respect to permutations of the variables, uniqueness cannot occur except under very special circumstances.

Mathematics Subject Classification. 49K20, 49K30.

Received July 21, 2015. Revised December 25, 2015. Accepted December 28, 2015.

1. INTRODUCTION

In this paper, we consider the multi-marginal Monge–Kantorovich transport problem. Given Borel probability measures μ_1, \dots, μ_n (the marginals) on smooth manifolds X_1, \dots, X_n , respectively, let $\Pi(\mu_1, \dots, \mu_n)$ be the set of Borel probability measures on $X_1 \times X_2 \times \dots \times X_n$ which have X_i -marginal μ_i for each $i \in \{1, 2, \dots, n\}$. Letting

$$c : X_1 \times X_2 \times \dots \times X_n \rightarrow \mathbb{R},$$

Keywords and phrases. Multi-marginal optimal transport, Moge–Kantorovich problem, extremal points of convex sets, m -twist, c -splitting set.

* *B.P. is pleased to acknowledge the support of a University of Alberta start-up grant and National Sciences and Engineering Research Council of Canada Discovery Grant number 412779-2012. A. M. greatly acknowledges the support of a grant from the Natural Sciences and Engineering Research Council of Canada (417885-2013).*

¹ School of Mathematics and Statistics, Carleton University, Ottawa, Ontario, Canada K1S 5B6. momeni@math.carleton.ca.

² Department of Mathematical and Statistical Sciences, 632 CAB, University of Alberta, Edmonton, Alberta, Canada, T6G 2G1. pass@ualberta.ca.

be a measurable cost function, the transport cost associated to a transport plan $\pi \in \Pi(\mu_1, \dots, \mu_n)$ is defined by

$$I_c(\pi) = \int_{X_1 \times X_2 \times \dots \times X_n} c(x_1, \dots, x_n) d\pi.$$

The multi-marginal Monge–Kantorovich transport problem is then the following minimization problem:

$$\inf \left\{ I_c(\pi); \pi \in \Pi(\mu_1, \dots, \mu_n) \right\}. \quad (MK)$$

Under quite mild conditions on c and the μ_i , it is well known that (MK) admits a minimizer. Our purpose here is to study the geometric structure of optimal measures $\pi \in \Pi(\mu_1, \dots, \mu_n)$.

Note that when $n = 2$, (MK) is precisely the classical optimal transport problem of Monge and Kantorovich. In this case, assuming that the first marginal μ_1 is absolutely continuous with respect to local coordinates, the celebrated *twist* condition implies that the solution is concentrated on a graph over x_1 , and is unique [5, 13, 14, 20, 30]. While this condition holds for many important cost functions, there are also a variety where it fails, including *any* smooth cost on a compact manifold. Recently, there have been some developments on two marginal problems where the twist condition fails [2, 9, 25]. In particular, one of the present authors developed two conditions, known as the *m-twist* and *generalized twist* conditions, which ensure respectively that the solution concentrates on several or countably many (rather than one) graphs [25]. A related line of recent research concerns sufficient conditions for uniqueness of the optimal measure, one interesting goal being to exhibit smooth costs on arbitrary manifolds for which optimal plans are unique, despite the fact that such plans are not generally concentrated on graphs. The interested reader is referred to a recent manuscript [22] and also to [1, 2, 9, 17] for more examples. We recall that a measure γ in the convex and weakly compact set $\Pi(\mu_1, \mu_2)$ is called *extremal* if $\gamma = t\gamma_1 + (1-t)\gamma_2$ with $\gamma_1, \gamma_2 \in \Pi(\mu_1, \mu_2)$ and $0 < t < 1$ imply that $\gamma_1 = \gamma_2$. It is well known that if (MK) admits a unique solution, then this solution must be extremal in $\Pi(\mu_1, \mu_2)$; an important aspect of the research on uniqueness is the characterization of extremal points of $\Pi(\mu_1, \mu_2)$. It is known that the extremal points of $\Pi(\mu_1, \mu_2)$ are not necessarily concentrated on a single graph. Indeed, any extremal point of $\Pi(\mu_1, \mu_2)$ is known to concentrate on the union of the graph of a function from X_1 to X_2 and the antigraph of a function from X_2 to X_1 . We refer the interested reader to [24] where a sufficient and almost necessary condition for the structure of the support of extremal measures in $\Pi(\mu_1, \mu_2)$ is established.

Interest in the $n \geq 3$ case (also known as the multi-marginal case) has picked up recently. Here, an analogue of the twist condition, known as the *twist on splitting sets* condition has recently been developed; this condition ensures that the solution to the multi-marginal problem concentrates on a graph over the first variable x_1 (and is unique) [19]. In contrast to the two marginal setting, where the twist condition is reasonably mild (at least for spaces with trivial topology), the twist on splitting sets condition is very strong and there are many important examples for which it fails. In particular, examples arising in density functional theory in physics [4, 11] and roommate type problems in economics [8] involve equal marginals and cost functions which are symmetric under all permutations of the arguments (see also [17]). For problems with this symmetric structure, it is known that there are always solutions that do not concentrate on a single graph (except in the very simple special case when the diagonal $\{(x, x, \dots, x)\}$ supports a unique optimizer) [29]. In this paper, we focus on what can be said in the multi-marginal case when the solution does not concentrate on a single graph; in particular, we investigate the possibility that the solution may be concentrated on several graphs.

Specifically, we consider the *m-twist* condition on splitting sets, and the related generalized twist on splitting sets condition, (see Def. 2.2) which are analogues of the conditions in [25] in the two marginal case. These conditions were recently introduced by one of the authors in [23], and assuming the first marginal is absolutely continuous with respect to local coordinates, they imply, respectively, that any optimal plan π concentrates on at most m graphs, or countably many graphs, over the first variable.

The *m-twist* condition on splitting sets is somewhat more flexible than the twist on splitting sets condition and, in particular, it seems relevant to symmetric costs which arise in the applications outlined above. We show, for instance, that the Coulomb cost (see (3.3) below) from density functional theory is *m-twisted* on splitting

sets in one dimension. Our work therefore offers a different perspective on a recent result of Colombo *et al.* [10], who constructed explicit solutions to (MK) for this cost, when the marginals are all equal; their solutions concentrated on $(n - 1)!$ graphs (where n is the number of marginals). In fact, our result also applies to the case when the marginals are not equal; on the other hand, we are able to prove only that solutions concentrate on at most $n! - (n - 1)!$ graphs. It is unclear to us whether there are in fact marginals for which our upper bound on the number of graphs is sharp; if so, the marginals cannot all be equal.

We also demonstrate that a local differential condition on c , which is known to imply that the optimizer concentrates on a low dimensional subset of the product space, also implies a local version of the twist condition on splitting sets. This local twist on splitting sets in turn implies the generalized twist on splitting sets, so that each optimal plan concentrates on the union of the graphs of at most a countable number of maps.

Following the recent developments on uniqueness issues for non graphical solutions in the two marginal case, we establish here some results which can be useful to derive uniqueness of solutions concentrated on several graphs in the multi-marginal setting, and exhibit some examples. We also prove a general, negative result, implying that the permutation symmetric problems described above can have unique solutions only in very special circumstances. For solutions concentrated on several graphs, however, these circumstances are somewhat more flexible than in the single graph case (where only the trivial, diagonal solution $(Id, Id, \dots Id)_{\#}\mu$ has the possibility of being a unique optimizer); we provide an example of a symmetric three marginal problem where a unique solution concentrates on two graphs and at least one of the components of each is not the identity mapping.

The next section is devoted to preliminary definitions. In the third section, we state our results for costs which satisfy the m and generalized twist on splitting sets conditions. We then prove that the local differential conditions alluded to above imply local twist on splitting sets, and outline several examples. In the fourth section, we prove several results related to uniqueness of solutions concentrated on multiple graphs. We show that under the generalized twist on splitting sets condition, a collection of several graphs containing the support of every optimizer is uniquely determined. We then establish conditions under which any solution concentrated on several graphs must be extremal, and prove that solutions to symmetric problems cannot be unique when $n \geq 3$, except under very particular conditions.

2. DEFINITIONS AND PRELIMINARIES

Here we recall several definitions which will be needed throughout the paper. We begin with the definition of c -splitting sets, modified slightly from [19].

Definition 2.1. A set $S \subset X_1 \times X_2 \times \dots \times X_n$ is a c -splitting set (or simply a splitting set) if there exist Borel functions $u_i : X_i \rightarrow [-\infty, \infty)$ such that for all (x_1, x_2, \dots, x_n) ,

$$\sum_{i=1}^n u_i(x_i) \leq c(x_1, x_2, \dots, x_n) \tag{2.1}$$

with equality whenever $(x_1, x_2, \dots, x_n) \in S$. The n -tuple (u_1, \dots, u_n) is called the c -splitting tuple for S .

The relevance of splitting sets to (MK) is that, as can be easily seen from the dual formulation found in [12], any optimal measure is concentrated on a c -splitting set.

Next, we recall the following relaxations, introduced in [23], of the twist on splitting sets conditions from [19]. In what follows, $D_{x_1}c(x_1, x_2, \dots, x_n)$ denotes the differential of c with respect to the first variable x_1 .

Definition 2.2. Let c be a Borel measurable function and denote by $D_1(c)$ the set of points at which c is differentiable with respect to the first variable.

1. **m -twist condition:** We say that c is m -twisted on splitting sets if for any c -splitting set $S \subset X_1 \times X_2 \times \dots \times X_n$, any $x_1 \in X_1$ and any p in the cotangent space $T_{x_1}^*X_1$ to X_1 at x_1 , the cardinality of the set

$$\left\{ (x_1, x_2, \dots, x_n) \in S \cap D_1(c); D_{x_1}c(x_1, x_2, \dots, x_n) = p \right\},$$

is at most m . We say that c locally satisfies the m -twist condition on splitting sets if for any $(x_1, x_2, \dots, x_n) \in S \cap D_1(c)$ there exists a neighborhood $U \subseteq X_2 \times \dots \times X_n$ of (x_2, \dots, x_n) such that the cardinality of the set

$$\left\{ y \in U, (x_1, y) \in S \cap D_1(c), D_{x_1}c(x_1, y) = D_{x_1}c(x_1, x_2, \dots, x_n) \right\},$$

is at most m .

2. **Generalized-twist condition:** We say that c satisfies the generalized twist condition on splitting sets if for any c -splitting set $S \subset X_1 \times X_2 \times \dots \times X_n$, any $x_1 \in X_1$ and any p in the cotangent space $T_{x_1}^*X_1$ to X_1 at x_1 , the set

$$\left\{ (x_1, x_2, \dots, x_n) \in S \cap D_1(c); D_{x_1}c(x_1, x_2, \dots, x_n) = p \right\},$$

is a finite subset of S .

Remark 2.3. Given a c -splitting set S and splitting functions u_1, \dots, u_n , it is well known that one can construct new splitting functions, $\bar{u}_1, \dots, \bar{u}_n$ for S such that, for each $i = 1, 2, \dots, n$, we have

$$\bar{u}_i(x_i) = \inf_{x_j, j \neq i} c(x_1, \dots, x_n) - \sum_{j \neq i} \bar{u}_j(x_j). \tag{2.2}$$

Assuming c is globally Lipschitz, it is well known that each u_i is as well [21] (unless u_i is identically $-\infty$, in which case the splitting set is empty). Therefore, the set $\{c(x_1, \dots, x_n) = \sum_{i=1}^n \bar{u}_i(x_i)\}$, which contains S , is closed (and hence compact if each X_i is).

The following result provides a connection between the generalized twist condition on splitting sets and local m -twistedness on splitting sets.

Proposition 2.4. *Assume that c is continuously differentiable with respect to the first variable and each X_i is compact. If, for some $m \in \mathbb{N}$, the function c is locally m -twisted on splitting sets, then c satisfies the generalized-twist condition on splitting sets.*

Proof. Assume that $S \subset X_1 \times \dots \times X_n$ is a c -splitting set.

Fix $(\bar{x}_1, \dots, \bar{x}_n) \in S$. We need to show that the set

$$L = \left\{ (\bar{x}_1, x_2, \dots, x_n) \in S; D_1c(\bar{x}_1, \bar{x}_2, \dots, \bar{x}_n) = D_1c(\bar{x}_1, x_2, \dots, x_n) \right\},$$

is finite. By Remark 2.3, we can find a compact splitting set, \tilde{S} , such that $S \subseteq \tilde{S}$; we will in fact prove the stronger statement, that

$$\tilde{L} = \left\{ (\bar{x}_1, x_2, \dots, x_n) \in \tilde{S}; D_1c(\bar{x}_1, \bar{x}_2, \dots, \bar{x}_n) = D_1c(\bar{x}_1, x_2, \dots, x_n) \right\},$$

is finite (note that $L \subseteq \tilde{L}$). If \tilde{L} is not finite there exists a countably infinite subset $\{(\bar{x}_1, x_2^k, \dots, x_n^k)\}_{k \in \mathbb{N}} \subset \tilde{L}$. Since \tilde{S} is compact then the sequence $\{(\bar{x}_1, x_2^k, \dots, x_n^k)\}_{k \in \mathbb{N}}$ has an accumulation point $(\bar{x}_1, x_2^0, \dots, x_n^0) \in S$ and there exists a subsequence still denoted by $\{(\bar{x}_1, x_2^k, \dots, x_n^k)\}_{k \in \mathbb{N}}$ such that $x_i^k \rightarrow x_i^0$ as $k \rightarrow \infty$ for $i = 2, \dots, n$. Since D_1c is continuous it follows that $(\bar{x}_1, x_2^0, \dots, x_n^0) \in \tilde{L}$. Since c is locally m -twisted on \tilde{S} , this leads to a contradiction as $(\bar{x}_1, x_2^0, \dots, x_n^0)$ is an accumulation point of the sequence $\{(\bar{x}_1, x_2^k, \dots, x_n^k)\}_{k \in \mathbb{N}}$ and

$$D_1c(\bar{x}_1, x_2^0, \dots, x_n^0) = D_1c(\bar{x}_1, x_2^k, \dots, x_n^k), \quad \forall k \in \mathbb{N}.$$

This completes the proof. □

We recall next an important property of splitting sets, which will be useful at various points throughout the paper.

Definition 2.5. A set $S \subset X_1 \times X_2 \times \dots \times X_n$ is c -monotone if the following holds. Whenever $x, \bar{x} \in S$, and P_+, P_- are two non empty disjoint sets of indices such that $P_+ \cup P_- = \{1, 2, \dots, n\}$, we have

$$c(x) + c(\bar{x}) \leq c(x_+, \bar{x}_-) + c(\bar{x}_+, x_-).$$

Here, we have decomposed $x = (x_+, x_-)$ and $\bar{x} = (\bar{x}_+, \bar{x}_-)$ in the obvious way; that is, $x_+ = (x_i)_{i \in P_+}$, with analogous definitions for x_-, \bar{x}_+ and \bar{x}_- .

It is well known that any c -splitting set is c -monotone (see [19, 28]).

For probability spaces (X, μ) and (Y, ν) , we say that a map $T : X \rightarrow Y$ transports or pushes forward μ to ν , and write $T\#\mu = \nu$, if

$$\forall f \in \mathcal{L}^1(\nu), \quad f \circ T \in \mathcal{L}^1(\mu) \ \& \ \int_X f(Tx) \, d\mu = \int_Y f(y) \, d\nu,$$

where $\mathcal{L}^1(\nu)$ is the set of ν -integrable functions on Y . Given a map T , the measure $(Id, T)\#\mu$ on $X \times Y$ is often described as being concentrated on the graph of T , and belongs to $\Pi(\mu, \nu)$ whenever $T\#\mu = \nu$.

We close this section by precisely defining some analogous notation describing measures concentrated on several graphs.

Definition 2.6. Let X and Y be Polish spaces with Borel probability measures μ on X and ν on Y .

We say that a measure $\gamma \in \Pi(\mu, \nu)$ is concentrated on the graphs of measurable maps $\{T_i\}_{i=1}^k$ from X to Y , if there exists a sequence of measurable non-negative functions $\{\alpha_i\}_{i=1}^k$ from X to \mathbb{R} with $\sum_{i=1}^k \alpha_i(x) = 1$ (μ -almost surely) such that for each measurable set $S \subset X \times Y$,

$$\gamma(S) = \sum_{i=1}^k \int_X \alpha_i(x) \chi_S(x, T_i x) \, d\mu,$$

where χ_S is the indicator function of the set S . In this case we write $\gamma = \sum_{i=1}^k \alpha_i (Id \times T_i)\#\mu$.

3. MULTI-GRAPH SOLUTIONS

The following characterization and its proof can be found in [23]; it asserts that, under the m -twist on splitting sets condition, any solution to the Kantorovich problem concentrates on (at most) m graphs over the first variable.

Theorem 3.1. Assume that the cost function c satisfies the m -twist condition on splitting sets, μ_1 is non-atomic and any function $u_1(x_1)$ that is of form (2.2) and not identically infinite is differentiable μ_1 -almost surely on its domain. Then for each optimal plan γ of (MK) with $\text{Supp}(\gamma) \subset D_1(c)$, there exist $k \leq m$, a sequence of non-negative measurable real functions $\{\alpha_i\}_{i=1}^k$ on X_1 and, Borel measurable maps $G_1, \dots, G_k : X_1 \rightarrow X_2 \times \dots \times X_n$ such that

$$\gamma = \sum_{i=1}^k \alpha_i (Id \times G_i)\#\mu, \tag{3.1}$$

where $\sum_{i=1}^k \alpha_i(x) = 1$ for μ_1 -a.e. $x \in X_1$.

Moreover, if one replaces the m -twist condition by the generalized-twist condition then each optimal plan γ of (MK) is of the form (3.1) for some $k \in \mathbb{N} \cup \{\infty\}$.

Although this result asserts quite strong conclusions about the structure of optimizers, the m -twist on splitting sets (respectively, generalized twist on splitting sets) hypothesis seems quite difficult to verify. Below, we establish that certain local differential conditions on the cost actually suffice for the generalized twist on splitting sets condition.

3.1. A differential condition for local twist on splitting sets

Here we establish a differential condition on c which is sufficient for the local 1-twist on splitting sets condition. The condition first appeared in [28], where it was shown to imply that the solutions concentrate on low dimensional subsets of the product space. Assume that each X_i is a d -dimensional smooth manifold. As in [28], we let g be the off diagonal part of the Hessian of c ; that is, in block form,

$$g = \begin{bmatrix} 0 & D^2_{x_1x_2}c & \cdots & \cdots & D^2_{x_1x_n}c \\ D^2_{x_2x_1}c & 0 & D^2_{x_2x_3}c & \cdots & D^2_{x_2x_n}c \\ \cdots & \cdots & \cdots & \cdots & \cdots \\ D^2_{x_nx_1}c & D^2_{x_nx_2}c & \cdots & \cdots & 0 \end{bmatrix}, \tag{3.2}$$

where each $D^2_{x_ix_j}c = (\frac{\partial^2 c}{\partial x_i^\alpha \partial x_j^\beta})_{\alpha\beta}$, for $i \neq j$, is the $d \times d$ matrix of mixed, second order partial derivatives. We recall that the signature of a symmetric matrix is the ordered triplet, $(\lambda_+, \lambda_-, \lambda_0)$, where λ_+ , λ_- and λ_0 denote respectively the number of positive and negative eigenvalues, and the multiplicity of the zero eigenvalue.

We then have the following.

Proposition 3.2. *Let c be twice differentiable. The following assertions hold:*

- (1) *If the symmetric, $dn \times dn$ matrix g has signature $(nd - d, d, 0)$ then c is locally 1-twisted on splitting sets.*
- (2) *If in addition each X_i is compact then c satisfies the generalized twist condition on splitting sets.*

In fact, the only property of splitting sets used in the proof of this result is their c -monotonicity. Therefore, analogous results hold when the splitting set condition on S is replaced with the weaker c -monotone condition.

Proof. We first prove assertion 1. The proof is by contradiction; to this end, assume that c is not locally 1-twisted on splitting sets. It follows that there exists a splitting set S and a point $(\bar{x}_1, \bar{x}_2, \dots, \bar{x}_n) \in S$ such that the set

$$\left\{ (x_2, \dots, x_n) : (\bar{x}_1, x_2, \dots, x_n) \in S, D_{x_1}c(\bar{x}_1, \bar{x}_2, \dots, \bar{x}_n) = D_{x_1}c(\bar{x}_1, x_2, \dots, x_n) \right\}$$

intersects any open neighbourhood of $(\bar{x}_2, \dots, \bar{x}_n)$ (at a point other than $(\bar{x}_2, \dots, \bar{x}_n)$). We can therefore take a sequence (x_2^k, \dots, x_n^k) in this set converging to $(\bar{x}_2, \dots, \bar{x}_n)$, such that $(x_2^k, \dots, x_n^k) \neq (\bar{x}_2, \dots, \bar{x}_n)$ for all k .

Set

$$v^k = \frac{(x_2^k, \dots, x_n^k) - (\bar{x}_2, \dots, \bar{x}_n)}{|(x_2^k, \dots, x_n^k) - (\bar{x}_2, \dots, \bar{x}_n)|}.$$

As each v^k has unit norm, we may pass to a convergent subsequence, so that $v^k \rightarrow v = (v_2, \dots, v_n) \neq 0$. Note that the vector $(0, v) = (0, v_2, \dots, v_n)$ is then tangent to S at $(\bar{x}_1, \bar{x}_2, \dots, \bar{x}_n)$.

Furthermore, the splitting set S is necessarily c -monotone, and so, by a result in [28] we have

$$(0, v)^T \cdot g \cdot (0, v) \leq 0$$

This is equivalent to

$$\sum_{i,j=2, i \neq j}^n v_i \cdot D^2_{x_ix_j}c \cdot v_j \leq 0.$$

Now, since for each k we have

$$\frac{D_{x_1}c(\bar{x}_1, \bar{x}_2, \dots, \bar{x}_n) - D_{x_1}c(\bar{x}_1^k, x_2^k, \dots, x_n^k)}{|(x_2^k, \dots, x_n^k) - (\bar{x}_2, \dots, \bar{x}_n)|} = 0,$$

taking the limit as $k \rightarrow \infty$ yields

$$0 = \sum_{j=2}^n D^2_{x_1x_j}c \cdot v_j.$$

But now let $w \in T_{\bar{x}_1} X_1$, $s \in \mathbb{R}$, and consider the vector $u = (w, 0, 0, \dots, 0) + s(0, v)$. We then have

$$u^t \cdot g \cdot u = \sum_{i,j=1, i \neq j}^k u_i \cdot D_{x_i x_j}^2 c \cdot u_j = 2sw \cdot \sum_{i=2}^n D_{x_1 x_j}^2 c \cdot v_j + \sum_{i,j=2, i \neq j}^n v_i \cdot D_{x_i x_j}^2 c \cdot v_j = s^2 v^t \cdot g \cdot v \leq 0.$$

We have therefore found a $d+1$ dimensional vector space, $T_{\bar{x}_1} X_1 \oplus \text{span}(v) \subseteq T_{\bar{x}_1 \times \bar{x}_2, \dots, \bar{x}_n} X_1 \times X_2 \times \dots \times X_n$, on which g is negative definite, contradicting the assumption that g has only d timelike directions, and completing the proof of assertion 1. Proposition 2.4 then immediately implies the second assertion. \square

The last proposition together with Theorem 3.1 imply that each optimal plan is concentrated on the union of the graphs of at most a countable number of maps.

3.2. Application: The one dimensional Coulomb cost with n marginals

Let $c : \mathbb{R}^n \rightarrow \mathbb{R}$ be the one dimensional Coulomb cost; that is, take each $X_i = \mathbb{R}$ and set

$$c(x_1, \dots, x_n) = \sum_{1 \leq i < j \leq n} \frac{1}{|x_i - x_j|}, \quad \forall (x_1, \dots, x_n) \in \mathbb{R}^n. \tag{3.3}$$

This cost has important applications in density functional theory in physics [4,11]. A recent paper by Colombo, De Pascale *et al.* [10] solves the optimal transport problem with this cost explicitly when the marginals μ_i are all the same, which is the most important case in view of physical applications. The solutions they exhibit all concentrate on several graphs. Here, we show that this cost satisfies the m twist on splitting sets condition, which, by virtue of Theorem 3.1, uncovers a new perspective on their result, as well as showing that the multi-graph solution structure persists to the setting where the marginals differ.

Proposition 3.3. *The one dimensional Coulomb cost (3.3) is $n! - (n - 1)!$ twisted on splitting sets.*

As with Proposition 3.2, the proof of this result uses only the c -monotonicity property of splitting sets, and so we also obtain that the Coulomb cost is $n! - (n - 1)!$ twisted on c -monotone sets.

Proof. Let S be a splitting set and u_1, \dots, u_n the corresponding splitting functions. Fix $x_1 \in \mathbb{R}$ and $p \in \mathbb{R}$. As the splitting functions never take on the value $+\infty$, the inequality (2.1) implies that there is no point $x_1, \dots, x_n \in S$ with $x_i = x_j$ for some $i \neq j$. We first prove that, for each permutation σ on n letters, there is at most one $(x_2, \dots, x_n) \in X_2 \times \dots \times X_n$, such that $D_{x_1} c(x_1, \dots, x_n) = p$ and

$$(x_1, x_2, \dots, x_n) \in S \cap A_\sigma$$

where

$$A_\sigma = \{(x_1, x_2, \dots, x_n) : x_{\sigma(i)} < x_{\sigma(j)}, \forall i < j\}.$$

Now, suppose we have two such points (x_1, x_2, \dots, x_n) and $(x_1, \bar{x}_2, \dots, \bar{x}_n)$ in $S \cap A_\sigma$ with

$$D_{x_1} c(x_1, x_2, \dots, x_n) = p = D_{x_1} c(x_1, \bar{x}_2, \dots, \bar{x}_n).$$

We claim that either $x_i \geq \bar{x}_i$ for all $i = 2, \dots, n$ or *vice versa*. As it is easy to check that $x_i \mapsto D_{x_1} c(x_1, x_2, \dots, x_n)$ is strictly monotone decreasing on A_σ for each i , this will establish the result (as if, for example, $x_i \geq \bar{x}_i$ for all $i = 2, \dots, n$, then the monotonicity implies $D_{x_1} c(x_1, x_2, \dots, x_n) \leq D_{x_1} c(x_1, \bar{x}_2, \dots, \bar{x}_n)$, with strict inequality as long as $x_i > \bar{x}_i$ for at least one i).

To see the claim, we use the fact that the splitting set S is c -monotone. Assume that the claim is false; then there exists non empty disjoint sets $P_+, P_- \subseteq \{1, 2, \dots, n\}$ such that $P_+ \cup P_- = \{1, 2, \dots, n\}$, and $x_i \geq \bar{x}_i$ for i in P_+ , $x_i \leq \bar{x}_i$ for i in P_- , and at least one of the inequalities is strict in each set.

For ease of notation, we decompose $(x_1, x_2, \dots, x_n) = (x_+, x_-)$ and $(\bar{x}_1, \bar{x}_2, \dots, \bar{x}_n) = (\bar{x}_+, \bar{x}_-)$, (with the same meaning for x_+, x_-, \bar{x}_+ and \bar{x}_- as in Definition 2.5).

By the c -monotonicity property, we have

$$c(x_+, x_-) + c(\bar{x}_+, \bar{x}_-) \leq c(\bar{x}_+, x_-) + c(x_+, \bar{x}_-),$$

which is equivalent to

$$\sum_{i \in P_+, j \in P_-} \frac{1}{|x_i - x_j|} + \frac{1}{|\bar{x}_i - \bar{x}_j|} \leq \sum_{i \in P_+, j \in P_-} \frac{1}{|x_i - \bar{x}_j|} + \frac{1}{|\bar{x}_i - x_j|}. \tag{3.4}$$

Now, for a fixed $i \in P_+, j \in P_-$, consider the paths $x_i(t) = x_i + t(\bar{x}_i - x_i)$ and $x_j(s) = x_j + s(\bar{x}_j - x_j)$. As both x and \bar{x} are in the same A_σ we can assume without loss of generality that $x_i < x_j$ and $\bar{x}_i < \bar{x}_j$. Together with the fact that $(\bar{x}_i - x_j)(\bar{x}_j - x_j) \leq 0$, (which follows from i being in P_+ and j in P_-) this implies that for each point $(t, s) \in [0, 1]^2$ we have $x_i(t) < x_j(s)$. We therefore have

$$\begin{aligned} \frac{\partial^2}{\partial s \partial t} \left[\frac{1}{|x_i(t) - x_j(s)|} \right] &= \frac{\partial^2}{\partial s \partial t} \left[\frac{1}{x_j(s) - x_i(t)} \right] \\ &= -2 \frac{1}{[x_j(s) - x_i(t)]^3} (\bar{x}_i - x_i)(\bar{x}_j - x_j) \\ &\geq 0. \end{aligned}$$

Now note that this implies

$$\frac{1}{|x_i - x_j|} + \frac{1}{|\bar{x}_i - \bar{x}_j|} - \frac{1}{|x_i - \bar{x}_j|} - \frac{1}{|\bar{x}_i - x_j|} = \int_0^1 \int_0^1 \frac{\partial^2}{\partial s \partial t} \left[\frac{1}{|x_i(t) - x_j(s)|} \right] ds dt \geq 0.$$

Furthermore, the inequality is strict for at least one $i \in P_+$ and one $j \in P_-$; summing over $i \in P_+$ and $j \in P_-$, this violates (3.4); this contradiction establishes the claim.

So far, we have proven that we can have at most one point (x_2, \dots, x_n) such that $(x_1, x_2, \dots, x_n) \in S \cap A_\sigma$ and $D_{x_1}c(x_1, \dots, x_n) = p$ for each permutation σ .

Now, note that if $p \geq 0$, we cannot have any solutions to $p = D_{x_1}c(x_1, x_2, \dots, x_n)$ in the region where $x_1 < x_i$ for all $i = 2, 3, \dots, n$ (as clearly $D_{x_1}c(x_1, x_2, \dots, x_n) < 0$ there). A similar argument applies to the region $x_1 < x_i$ for all $i = 2, 3, \dots, n$ when $p \leq 0$. Therefore, we do not have any solutions in A_σ when $\sigma(1) = 1$ in the first case, or when $\sigma(1) = n$ in the second case. In either case, there are $(n - 1)!$ permutations with the appropriate property.

This means that there are at most $n! - (n - 1)!$ permutations σ for which we potentially have one solution (x_2, \dots, x_n) to $D_{x_1}c(x_1, \dots, x_n) = p$ on $S \cap A_\sigma$. We therefore have at most $n! - (n - 1)!$ solutions to the equation $p = D_{x_1}c(x_1, x_2, \dots, x_n)$ on the splitting set S . □

Corollary 3.4. *Assume that c is finite $\mu_1 \otimes \dots \otimes \mu_n$ almost everywhere and there exists a transport plan $\pi \in \Pi(\mu_1, \mu_2, \dots, \mu_n)$ with finite total cost, $\int_{\mathbb{R}^n} c d\pi < \infty$. If μ_1 is absolutely continuous with respect to the 1-dimensional Lebesgue measure then for any optimal π_0 in (C) , there exist $k \leq n! - (n - 1)!$, a sequence of non-negative measurable real functions $\{\alpha_i\}_{i=1}^k$ on \mathbb{R} and, Borel measurable maps $G_1, \dots, G_k : \mathbb{R} \rightarrow \mathbb{R}^{n-1}$ such that*

$$\pi_0 = \sum_{i=1}^k \alpha_i (Id \times G_i) \# \mu_1, \tag{3.5}$$

where $\sum_{i=1}^k \alpha_i(x) = 1$ for μ_1 -a.e. $x \in \mathbb{R}$.

Proof. By a similar argument as in the proof of theorem (4) in [4], there exist potentials u_1, \dots, u_n with $u_i \in \mathcal{L}^1(\mu_i)$, taking values in $[-\infty, \infty)$ and $c(x_1, \dots, x_n) \geq \sum_{i=1}^n u(x_i)$ such that

$$\int_{\mathbb{R}^n} c(x_1, \dots, x_n) \, d\pi_0 = \sum_{i=1}^n \int_{\mathbb{R}} u(x_i) \, d\mu_i,$$

and u_1 is differentiable μ_1 almost everywhere. It also follows from Proposition 3.3 that c is $n! - (n - 1)!$ twisted in splitting sets. Thus, the result follows from Theorem 3.1. \square

As shown in [10], when the marginals are all equal (the physically relevant case) the solution is concentrated on $(n - 1)!$ graphs; our result guarantees it is concentrated on *at most* $n! - (n - 1)!$, even when the marginals differ. It is unclear to us whether this is sharp; that is, whether there really are solutions which concentrate on $n! - (n - 1)!$ graphs. If so, this can only happen when the marginals are not all equal.

3.3. A cost satisfying local 1-twistedness

We exhibit now an example of a multi-dimensional cost function which is locally (but not globally) 1-twisted on splitting sets, together with a solution to the optimal transport problem which concentrates on several (rather than one) graphs.

Proposition 3.5. *Take, for $x = (x^1, x^2), y = (y^1, y^2), z = (z^1, z^2) \in \mathbb{R}^2$,*

$$c(x, y, z) = -e^{x^1+y^1} \cos(x^2 - y^2) - e^{x^1+z^1} \cos(x^2 - z^2) - e^{y^1+z^1} \cos(z^2 - y^2). \tag{3.6}$$

Then c is locally twisted on splitting sets.

Proof. We first recall that, for three marginal costs, the condition

$$D_{xy}^2 c [D_{zy}^2 c]^{-1} D_{zx}^2 c < 0$$

is equivalent to the matrix g in (3.2) having signature $(4, 2, 0)$ (see [26]); hence, by Proposition 3.2, this condition implies the local twist on splitting sets property). We will show that this condition holds for the cost (3.6). Note that

$$D_{xy}^2 c = - \begin{bmatrix} e^{x^1+y^1} \cos(x^2 - y^2) & e^{x^1+y^1} \sin(x^2 - y^2) \\ -e^{x^1+y^1} \sin(x^2 - y^2) & e^{x^1+y^1} \cos(x^2 - y^2) \end{bmatrix}.$$

Up to multiplication by $-e^{x^1+y^1}$, this is a rotation matrix through the angle $x^2 - y^2$. Similarly, $D_{yz}^2 c$ is $-e^{y^1+z^1}$ multiplied by a rotation through $y^2 - z^2$, so its inverse transpose, $[D_{zy}^2 c]^{-1}$, is $-\frac{1}{e^{y^1+z^1}}$ multiplied by the rotation through $y^2 - z^2$ (recall that, for a rotation matrix A , $A^T = A^{-1}$). Finally, $D_{zx}^2 c$ is $-e^{x^1+z^1}$ multiplied by a rotation through $z^2 - x^2$. Therefore, the product $D_{xy}^2 c [D_{zy}^2 c]^{-1} D_{zx}^2 c$ is, up to a negative multiplicative constant, rotation through $x^2 - y^2 + y^2 - z^2 + z^2 - x^2 = 0$; that is, the product is a multiple of the identity:

$$D_{xy}^2 c [D_{zy}^2 c]^{-1} D_{zx}^2 c = -e^{x^1+y^1} \frac{1}{e^{y^1+z^1}} e^{x^1+z^1} I < 0.$$

Therefore, by Proposition 3.2, this cost is locally twisted on splitting sets. \square

In fact, for this cost, we can exhibit explicitly an optimizer which concentrates on several graphs. To this end, note that

$$-e^{x^1+y^1} \cos(x^2 - y^2) \geq -e^{x^1+y^1} \geq -e^{2x^1}/2 - e^{2y^1}/2$$

and we have equality only when $x^1 = y^1$ and $x^2 - y^2$ is an integer multiple of 2π . Applying similar reasoning to the other pieces of c , we get

$$c(x, y, z) \geq -e^{2x^1} - e^{2y^1} - e^{2z^1}$$

with equality only when

$$(x, y, z) \in S = \{(x, y, z) : x^1 = y^1 = z^1 \text{ and } x^2 - y^2 = 2k\pi, x^2 - z^2 = 2l\pi \text{ for some integers } k, l\}$$

This easily implies that any measure γ concentrated on the set S is optimal in (MK) for its marginals, as $-e^{2x^1}$, $-e^{2y^1}$ and $-e^{2z^1}$ serve as Kantorovich potentials.

Locally, the set S looks like a graph, but globally, it is the union of countably many graphs (or finitely many, if we restrict to compact domains).

4. UNIQUENESS ISSUES

In this section, we consider the uniqueness of solutions concentrated on several graphs. We begin with a criterion for extremality of measures of type (3.1) in $\Pi(\mu_1, \dots, \mu_n)$. Note that as the minimization of a linear functional over a convex, weakly compact set $\Pi(\mu_1, \dots, \mu_n)$, the Kantorovich problem necessarily has at least one solution which is extremal in that set. Extremality of a solution $\gamma \in \Pi(\mu_1, \dots, \mu_n)$ to (MK) is therefore a *necessary* condition for γ to be the unique solution.

Theorem 4.1. *Let X_1, X_2, X_3 be Polish spaces equipped with Borel probability measures μ_i on X_i , and let $\{G_i = (H_i, K_i)\}_{i=1}^k$ be a finite sequence of measurable maps from X_1 to $X_2 \times X_3$. Assume that the following assertions hold:*

- (i) *For each $i \geq 1$ the map H_i is injective and onto.*
- (ii) *For each $j \geq 2$ the map K_j is injective and $\text{Ran}(K_i) \cap \text{Ran}(K_j) = \emptyset$ for all $i \geq 2$ with $i \neq j$.*
- (iii) *There exists a bounded measurable function $\theta : X_3 \rightarrow \mathbb{R}$ with the property that for each $i \geq 2$,*

$$\theta(K_1 \circ H_1^{-1}(x_2)) > \theta(K_i \circ H_i^{-1}(x_2)) \quad \forall x_2 \in \text{Dom}(K_1 \circ H_1^{-1}) \cap \text{Dom}(K_i \circ H_i^{-1}).$$

Then each $\gamma \in \Pi(\mu_1, \mu_2, \mu_3)$ that is concentrated on $\cup_{i=1}^k \text{Graph}(G_i)$ is an extremal point of $\Pi(\mu_1, \mu_2, \mu_3)$.

The proof of the preceding theorem is fairly long and so is deferred to the end of this section (see Sect. 4.2). We note that we have stated the result for $n = 3$ only to keep the presentation as simple as possible; it has a straight forward generalization to larger n .

The theorem may be used in certain circumstances to deduce uniqueness of the optimizer; in particular, if one can show that *all* solutions must be concentrated on a single collection of several graphs, satisfying the hypotheses in the Theorem, then uniqueness follows immediately. The following example illustrates this concept.

To the best of our knowledge, this is the first example of a solution to a multi-marginal optimal transport problem which is both 1) unique and 2) not concentrated on a single graph. Although the example is admittedly somewhat artificial, we still believe it is of interest as this type of behavior has not been observed before in multi-marginal problems.

Example 4.2 (A unique solution concentrated on two graphs). Let $c(x, y, z) = (x-y)^2 + (x-z)^2(x-z+1/2)^2$ on $X_1 \times X_2 \times X_3 = [0, 1] \times [0, 1] \times [0, 3/2]$. Consider the maps $(H_1, K_1) : X_1 \rightarrow X_2 \times X_3$ given by $(H_1, K_1)(x) = (x, x)$ and $(H_2, K_2) : X_1 \rightarrow X_2 \times X_3$ given by $(H_2, K_2)(x) = (x, x + \frac{1}{2})$. We note that

$$c(x, y, z) \geq 0$$

with equality if and only if $(y, z) = (H_i, K_i)(x)$, for $i = 1$ or 2 . Therefore, letting γ be uniform measure (renormalized to have total mass 1) on the union $S = \text{graph}(H_1, K_1) \cup \text{graph}(H_2, K_2)$ of these two graphs, we have that γ is optimal for its marginals μ_1, μ_2, μ_3 . To see that γ is in fact the unique optimizer, assume

that $\bar{\gamma}$ is any other optimizer; it must necessarily concentrate on the same set S . The interpolant $\frac{1}{2}(\gamma + \bar{\gamma})$ then also concentrates on S and must be optimal too. However, it is easy to see that (H_i, K_i) satisfy the conditions in Theorem 4.1 (taking, for instance, $\theta(z) = -z$), and so $\frac{1}{2}(\gamma + \bar{\gamma})$ must be extremal in $\Pi(\mu_1, \mu_2, \mu_3)$, a contradiction as it is the average of two distinct measures $\gamma, \bar{\gamma} \in \Pi(\mu_1, \mu_2, \mu_3)$. Therefore, γ is the unique optimal measure, as desired.

We now record a result that shows that, under the m -twisted on splitting sets criterion, uniqueness can only fail by reorganizing the mass within a single collection of m graphs.

Theorem 4.3. *Suppose that the continuously differentiable function c satisfies the m -twist condition on splitting sets, μ_1 is non-atomic and any function $u_1(x_1)$ of the form (2.2) which is not identically infinite is differentiable μ_1 -almost surely on its domain. Assume that $\{G_i\}_{i=1}^m$ is a sequence of measurable functions from X_1 to $X_2 \times \dots \times X_n$ such that for each $i \neq j$ the set $\{x; G_i(x) = G_j(x)\}$ is μ_1 -negligible. Let $\bar{\gamma}$ be an optimal plan for (MK) such that*

$$\bar{\gamma} = \sum_{k=1}^m \alpha_k (Id \times G_k) \# \mu, \quad (\alpha_i(x_1) \geq 0 \text{ and } \alpha_1(x_1)\alpha_2(x_1)\dots\alpha_m(x_1) \neq 0 \text{ for } \mu_1 - a.e. \ x_1 \in X_1). \quad (4.1)$$

Then for any other optimal plan γ we have

$$\text{Supp}(\gamma) \subseteq \text{Supp}(\bar{\gamma}).$$

Proof. By Kantorovich duality (see [7, 12] for a proof in the multi-marginal case), there exist functions $\varphi_i \in \mathcal{L}^1(\mu_i)$, $i = 1, \dots, n$ such that

$$\varphi_i(x_i) = \inf_{x_j \in X_j, j \neq i} \left\{ c(x_1, \dots, x_n) - \sum_{j \neq i} \varphi_j(x_j) \right\},$$

for all $x_i \in X_i$ and

$$\int c \, d\bar{\gamma} = \sum_{i=1}^n \int_{X_i} \varphi_i(x_i) \, d\mu_i.$$

Let S be the c -splitting set generated by the n -tuple $(\varphi_1, \dots, \varphi_n)$, that is,

$$S = \left\{ (x_1, \dots, x_n); c(x_1, \dots, x_n) = \sum_{i=1}^n \varphi_i(x_i) \right\}.$$

It follows that

$$\int c(x_1, \dots, x_n) \, d\bar{\gamma} = \int \sum_{i=1}^n \varphi_i(x_i) \, d\bar{\gamma},$$

from which we obtain

$$\sum_{k=1}^m \int_{X_1} \alpha_i(x_1) c(x_1, G_i x_1) \, d\mu_1 = \sum_{k=1}^m \int_{X_1} \alpha_i(x_1) \left[\varphi_1(x_1) + \psi(G_i x_1) \right] \, d\mu_1,$$

where $\psi(x_2, \dots, x_n) = \sum_{i=2}^n \varphi_i(x_i)$. It then follows that

$$\sum_{k=1}^m \int_{X_1} \alpha_i(x_1) \left[c(x_1, G_i x_1) - \varphi_1(x_1) - \psi(G_i x_1) \right] \, d\mu_1 = 0.$$

Since each integrand in the latter expression is non-negative and μ_1 almost surely $\alpha_1(x_1)\alpha_2(x_1)\dots\alpha_m(x_1) \neq 0$, we have that

$$c(x_1, G_i x_1) = \varphi_1(x_1) + \psi(G_i x_1) \quad \mu_1 - a.e. \quad \forall i \in \{1, \dots, m\}.$$

Consequently we obtain,

$$D_{x_1}c(x_1, G_i x_1) = D\varphi_1(x_1) \quad \mu_1 - a.e. \quad \forall i \in \{1, \dots, m\}. \tag{4.2}$$

Note also that for $i \neq j$ the set $\{x_1 \in X_1; G_i(x_1) = G_j(x_1)\}$ is a null set with respect to the measure μ_1 . This together with (4.2) and the m -twist condition on S imply that the cardinality of the set $\{G_1 x_1, \dots, G_m x_1\}$ is m for μ_1 -a.e. $x_1 \in X_1$.

Now assume that γ is also an optimal plan of (MK) . It follows from Theorem 3.1 that there exist a sequence of non-negative functions $\{\beta_i\}_{i=1}^m$ and, Borel measurable maps $T_1, \dots, T_m : X_1 \rightarrow X_2 \times \dots \times X_n$ such that

$$\gamma = \sum_{i=1}^m \beta_i (\text{Id} \times T_i) \# \mu_1.$$

By a similar argument as above one obtains

$$\beta_i(x) [D_{x_1}c(x_1, T_i x) - D\varphi_1(x_1)] = 0 \quad \mu_1 - a.e. \quad \forall i \in \{1, \dots, m\}.$$

For each i define $\Omega_i = \{x_1 \in X_1; \beta_i(x_1) \neq 0\}$. Since the cardinality of the set $\{G_1 x_1, \dots, G_m x_1\}$ is m for μ_1 -a.e. $x_1 \in X_1$ and since c satisfies the m -twist condition on S we have that for each i , $\{T_i x_1\} \subseteq \{G_1 x_1, \dots, G_m x_1\}$ for μ_1 -a.e. $x_1 \in \Omega_i$. This completes the proof. \square

4.1. Permutation symmetric problems

Finally, we turn our attention to permutation symmetric problems; that is, problems for which the cost is symmetric under permutation of its arguments, and the marginal are all the same. As discussed in the introduction, these problems have important applications in physics and economics.

We first demonstrate with an example that it is in fact possible to have unique solutions to these problems concentrated on several graphs. In [29], it was noted that a problem of this type could not admit a unique solution concentrated on a single graph, $\{(x, F_2(x), \dots, F_n(x))\}$, unless each component of that graph was the identity almost everywhere, $F_i(x) = x$ for a.e. x , for each $i = 2, \dots, m$. The example below illustrates that one can have somewhat less trivial unique solutions if we relax the structural requirements to being concentrated on two graphs instead of one. However, as Proposition 4.5 below implies, unique optimizers are still fairly special; this proposition asserts that one can not have unique solutions, except under very particular conditions.

Example 4.4 (Another unique solution concentrated on two graphs). Take $c(x, y, z) = xyz$ on $[-1, 1]^3$, with each marginal equal to

$$\mu = \frac{1}{3}[\mathcal{L}_{[-1,0]} + 2\mathcal{L}_{[0,1]}].$$

Define $G_1 : [-1, 1] \rightarrow [-1, 1]^2$ by $G_1(x) = (-x, |x|)$ and $G_2 : [-1, 1] \rightarrow [-1, 1]^2$ by $G_2(x) = (x, -|x|)$. Note that by the geometric-arithmetic mean inequality, we have

$$|xyz| \leq \frac{|x|^3}{3} + \frac{|y|^3}{3} + \frac{|z|^3}{3},$$

with equality only when $|x| = |y| = |z|$ and so

$$c(x, y, z) = xyz \geq -|xyz| \geq -\frac{|x|^3}{3} - \frac{|y|^3}{3} - \frac{|z|^3}{3}$$

with equality only when either one or three of x, y and z are non-positive, and $|x| = |y| = |z|$; that is, we have equality precisely on the graphs of G_1 and G_2 . Therefore, if we can find a $\gamma \in \Pi(\mu, \mu, \mu)$ concentrated on these two graphs, it is optimal; if there is a unique such γ , then the optimal measure is unique (as any probability measure concentrated outside these sets has larger total cost).

Now, finding a measure $\gamma \in \Pi(\mu, \mu, \mu)$ of the form

$$\gamma = \alpha_1(x)(Id \times G_1)_{\#}\mu + \alpha_2(x)(Id \times G_2)_{\#}\mu \tag{4.3}$$

amounts to solving, for almost every $x \geq 0$, the system of equations:

$$\begin{aligned} \alpha_1(x) + \alpha_2(x) &= 1 \\ \alpha_1(-x) + \alpha_2(-x) &= 1 \\ \frac{1}{3}\alpha_1(-x) + \frac{2}{3}\alpha_2(x) &= \frac{2}{3} \\ \frac{2}{3}\alpha_1(x) + \frac{1}{3}\alpha_2(-x) &= \frac{1}{3} \\ \frac{2}{3}\alpha_1(x) + \frac{1}{3}\alpha_1(-x) &= \frac{2}{3} \\ \frac{2}{3}\alpha_2(x) + \frac{1}{3}\alpha_2(-x) &= \frac{1}{3}. \end{aligned}$$

A straightforward calculation shows that the unique solution to these equations is given by

$$\alpha_1(x) = \begin{cases} \frac{1}{2} & x > 0 \\ 1 & x < 0, \end{cases}$$

and

$$\alpha_2(x) = \begin{cases} \frac{1}{2} & x > 0 \\ 0 & x < 0. \end{cases}$$

The measure given by equation (4.3), with these coefficients, is therefore the unique solution.

The preceding example demonstrates that it is possible to have a unique solution to a symmetric problem which concentrates on two graphs. The components of these graphs need not be the identity, but note that for each $G_i = (H_i, K_i)$ and each x , one of the following holds:

- (1) $H_i(x) = x$;
- (2) $K_i(x) = x$, or;
- (3) $H_i(x) = K_i(x)$.

In particular, $\gamma(S_1 \times S_2 \times S_3) = 0$ whenever the sets $S_1, S_2, S_3 \subseteq [-1, 1]$ are pairwise disjoint. This is essentially the only way that multiple maps can support a unique solution to a symmetric problem, as the following result shows.

Proposition 4.5. *Let $n = 3$. Suppose that the spaces $X_i := X$ and marginals $\mu_i := \mu$ are the same for each i , and that $c(x, y, z)$ is symmetric with respect to any permutation of the arguments. Assume that there exist mutually disjoint sets $S_1, S_2, S_3 \subseteq X$, and an optimal measure γ that charges $S := S_1 \times S_2 \times S_3$; that is, $\gamma(S) > 0$. Then the solution is nonunique.*

As should be clear from the proof, a similar result holds for $n \geq 4$.

Proof. The proof is by contradiction; assume that γ is the unique solution. Letting σ be any permutation on the variables (x, y, z) , this uniqueness implies that $\sigma_{\#}\gamma = \gamma$.

Now, note that

$$\sigma_{\#}(\gamma_{S_1 \times S_2 \times S_3}) = \gamma_{\sigma(S_1 \times S_2 \times S_3)},$$

where $\gamma_{S_1 \times S_2 \times S_3}$ denotes the restriction of γ to the set $S_1 \times S_2 \times S_3$. Therefore, by the symmetry of c , we have

$$\int_{X^3} c(x, y, z) d\gamma_{S_1 \times S_2 \times S_3} = \int_{X^3} c(x, y, z) d(\sigma_{\#}(\gamma_{S_1 \times S_2 \times S_3})) = \int_{X^3} c(x, y, z) d\gamma_{\sigma(S_1 \times S_2 \times S_3)}. \tag{4.4}$$

It is clear that each $\gamma_{\sigma(S_1 \times S_2 \times S_3)} \leq \gamma$, and that the measures

$$\frac{1}{6} \sum_{\sigma \in S_3} \gamma_{\sigma(S_1 \times S_2 \times S_3)}$$

and

$$\frac{1}{3} [\gamma_{S_1 \times S_2 \times S_3} + \gamma_{S_3 \times S_1 \times S_2} + \gamma_{S_2 \times S_3 \times S_1}]$$

share the same marginals, where S_3 is the set of permutations on 3 letters. It then follows that

$$\bar{\gamma} = \gamma + \frac{1}{6} \sum_{\sigma \in S_3} \gamma_{\sigma(S_1 \times S_2 \times S_3)} - \frac{1}{3} [\gamma_{S_1 \times S_2 \times S_3} + \gamma_{S_3 \times S_1 \times S_2} + \gamma_{S_2 \times S_3 \times S_1}]$$

is a nonnegative measure with the same marginals as γ . On the other hand, $\bar{\gamma} \neq \gamma$, as

$$\frac{1}{6} \sum_{\sigma \in S_3} \gamma_{\sigma(S_1 \times S_2 \times S_3)} \neq \frac{1}{3} [\gamma_{S_1 \times S_2 \times S_3} + \gamma_{S_3 \times S_1 \times S_2} + \gamma_{S_2 \times S_3 \times S_1}].$$

Finally, as by (4.4),

$$\int_{X^3} c(x, y, z) d\gamma = \int_{X^3} c(x, y, z) d\bar{\gamma},$$

$\bar{\gamma}$ is also optimal, which contradicts uniqueness. □

4.2. Extremal measures with fixed marginals

We shall now prove Theorem 4.1; the proof requires a few preliminary results. We begin with the following result from [24].

Theorem 4.6. *Let X and Y be Polish spaces equipped with Borel probability measures μ on X and ν on Y , and let $\{T_i\}_{i=1}^k$ be a (possibly infinite) sequence of measurable maps from X to Y . Assume that the following assertions hold:*

(i) *For each $i \geq 2$ the map T_i is injective on the set*

$$D_i := \{x \in \text{Dom}(T_1) \cap \text{Dom}(T_i); T_1x \neq T_ix\},$$

and $\text{Ran}(T_i) \cap \text{Ran}(T_j) = \emptyset$ for all $i, j \geq 2$ with $i \neq j$.

(ii) *There exists a bounded measurable function $\theta : Y \rightarrow \mathbb{R}$ with the property that $\theta(T_1x) > \theta(T_ix)$ on D_i .*

Then there exists at most one $\gamma \in \Pi(\mu, \nu)$ that is concentrated on $\cup_{i=1}^k \text{Graph}(T_i)$. Moreover, any $\gamma \in \Pi(\mu, \nu)$ that is concentrated on $\cup_{i=1}^k \text{Graph}(T_i)$ is an extremal point of $\Pi(\mu, \nu)$.

Our proof exploits this result, together with a connection between extremality in the two and three marginal cases, established in the following lemma.

Lemma 4.7. *Let γ be a measure in $\Pi(\mu_1, \mu_2, \mu_3)$ and let ν be the projection of γ onto $X_2 \times X_3$. If γ is an extremal point of $\Pi(\mu_1, \nu)$ and ν is an extremal point of $\Pi(\mu_2, \mu_3)$ then γ is an extremal point of $\Pi(\mu_1, \mu_2, \mu_3)$.*

Proof. Let $\gamma_1, \gamma_2 \in \Pi(\mu_1, \mu_2, \mu_3)$ and $0 < t < 1$ be such that $\gamma = t\gamma_1 + (1-t)\gamma_2$. For $i = 1, 2$ denote by ν_i the projection of γ_i on $X_2 \times X_3$. It is obvious that $\nu_1, \nu_2 \in \Pi(\mu_2, \mu_3)$, and that $\nu = t\nu_1 + (1-t)\nu_2$. Since ν is an extremal point of $\Pi(\mu_2, \mu_3)$ and $0 < t < 1$ we must have $\nu = \nu_1 = \nu_2$, which means that $\gamma_1, \gamma_2 \in \Pi(\mu_1, \nu)$. On the other hand γ is an extremal point of $\Pi(\mu_1, \nu)$; together with $\gamma = t\gamma_1 + (1-t)\gamma_2$, this implies $\gamma = \gamma_1 = \gamma_2$, from which we obtain the desired conclusion. □

Proof of Theorem 4.1. Let $\gamma \in \Pi(\mu_1, \mu_2, \mu_3)$ be a triply stochastic measure that is concentrated on $\cup_{i=1}^k \text{Graph}(G_i)$. By definition there exist non-negative Borel measurable real functions $\alpha_i : X_1 \rightarrow \mathbb{R}$ with $\sum_{i=1}^k \alpha_i(x_1) = 1$ such that $\gamma = \sum_{i=1}^k \alpha_i(\text{Id} \times G_i) \# \mu_1$, i.e.,

$$\int f \, d\gamma = \sum_{i=1}^k \int_{X_1} \alpha_i(x_1) f(x_1, H_i x_1, K_i x_1) \, d\mu_1.$$

Denote by ν the projection of γ on $X_2 \times X_3$. We proceed with the proof in several steps.

Step 1. In this step we show that the projection $\nu \in \Pi(\mu_2, \mu_3)$ of γ onto $X_2 \times X_3$ is concentrated on $\cup_{i=1}^k \text{Graph}(K_i \circ H_i^{-1})$. For $1 \leq i \leq k$, define the measures $\mu_{i,1}$ and $\mu_{2,i}$ by $d\mu_{i,1} = \alpha_i d\mu_1$ and $\mu_{2,i} = H_i \# \mu_{i,1}$. For every measurable set $B \in \mathcal{B}(X_2)$ we have

$$\begin{aligned} \mu_2(B) &= \gamma(X_1 \times B \times X_3) = \sum_{i=1}^k \int_{X_1} \alpha_i(x_1) \chi_B(H_i x_1) \, d\mu_1 = \sum_{i=1}^k \int_{X_1} \chi_B(H_i x_1) \, d\mu_{i,1} \\ &= \sum_{i=1}^k \int_{X_2} \chi_B(x_2) \, d\mu_{2,i} \\ &= \sum_{i=1}^k \int_B d\mu_{2,i} = \sum_{i=1}^k \mu_{2,i}(B). \end{aligned}$$

Thus, for all $B \in \mathcal{B}(X_2)$ we have $\mu_2(B) = \sum_{i=1}^k \mu_{2,i}(B)$. It implies that each $\mu_{2,i}$ is absolutely continuous with respect to μ_2 . Therefore, for each i there exists a nonnegative real function $\beta_i : X_2 \rightarrow \mathbb{R}$ such that $d\mu_{2,i} = \beta_i d\mu_2$. This together with $\mu_2(B) = \sum_{i=1}^k \mu_{2,i}(B)$ yield that

$$\mu_2(B) = \sum_{i=1}^k \int_B \beta_i(x_2) \, d\mu_2, \quad \forall B \in \mathcal{B}(X_2),$$

from which we obtain $\sum_{i=1}^k \beta_i(x_2) = 1$ for μ_2 almost every $x_2 \in X_2$. Take a measurable bounded function $f : X_2 \times X_3 \rightarrow \mathbb{R}$. It follows that

$$\begin{aligned} \int_{X_2 \times X_3} f(x_2, x_3) \, d\nu &= \int_{X_1 \times X_2 \times X_3} f(x_2, x_3) \, d\gamma = \sum_{i=1}^k \int_{X_1} \alpha_i(x_1) f(H_i x_1, K_i x_1) \, d\mu_1 \\ &= \sum_{i=1}^k \int_{X_1} f(H_i x_1, K_i x_1) \, d\mu_{i,1} \\ &= \sum_{i=1}^k \int_{X_2} f(x_2, K_i \circ H_i^{-1} x_2) \, d\mu_{2,i} \\ &= \sum_{i=1}^k \int_{X_2} \beta_i(x_2) f(x_2, K_i \circ H_i^{-1} x_2) \, d\mu_2. \end{aligned}$$

This implies that $\nu = \sum_{i=1}^k \beta_i(\text{Id} \times K_i \circ H_i^{-1}) \# \mu_2$.

Step 2. In this step we show that ν is an extremal point of $\Pi(\mu_2, \mu_3)$. We shall make use of Theorem 4.6 to prove this. By assumption we have that K_i is injective and $\text{Ran}(K_i) \cap \text{Ran}(K_j) = \emptyset$ for all $i, j \geq 2$ with $i \neq j$.

This implies that $K_i \circ H_i^{-1}$ is injective for all $i \geq 2$ and $\text{Ran}(K_i \circ H_i^{-1}) \cap \text{Ran}(K_j \circ H_j^{-1}) = \emptyset$ from which assumption (i) of Theorem 4.6 follows. The second assumption in Theorem 4.6 also follows due to hypothesis (iii) in Theorem 4.1 and therefore the extremality of ν in $\Pi(\mu_2, \mu_3)$ follows.

Step 3. We show that γ is an extremal point of $\Pi(\mu_1, \nu)$. We shall again use Theorem 4.6 to prove this part. Note that $G_i = (H_i, K_i)$ is injective for $i \geq 2$ as both H_i and K_i are injective. For $i, j \geq 2$ with $i \neq j$ we have $\text{Ran}(G_i) \cap \text{Ran}(G_j) = \emptyset$ as by assumption $\text{Ran}(K_i) \cap \text{Ran}(K_j) = \emptyset$. Define $\varphi : X_2 \times X_3 \rightarrow \mathbb{R}$ by

$$\varphi(x_2, x_3) = \begin{cases} (k-1)\theta(x_3) - \sum_{i=2}^k \theta(K_i \circ H_i^{-1}x_2), & x_3 \in X_3 \setminus \bigcup_{i=2}^k \text{Ran}(K_i), \\ \theta(x_3) - \theta(K_j \circ H_j^{-1}x_2), & x_3 \in \text{Ran}(K_j) \text{ for some } j \geq 2. \end{cases}$$

Since K_i is injective for $i \geq 2$ it follows that $\text{Ran}(K_i)$ is Borel measurable ([3], Thm. 6.8.6). Thus, φ is a bounded Borel measurable function. We show that $\varphi(G_1(x_1)) > \varphi(G_j(x_1))$ for all $j \geq 2$ and $x_1 \in \text{Dom}(G_1) \cap \text{Dom}(G_j)$. Fix $x_1 \in \text{Dom}(G_1) \cap \text{Dom}(G_j)$. If $K_1x_1 \in X_3 \setminus \bigcup_{i=2}^k \text{Ran}(K_i)$ it follows that

$$\begin{aligned} \varphi(G_1(x_1)) &= \varphi(H_1x_1, K_1x_1) = (k-1)\theta(K_1x_1) - \sum_{i=2}^k \theta(K_i \circ H_i^{-1} \circ H_1x_1) \\ &> (k-1)\theta(K_1x_1) - \sum_{i=2}^k \theta(K_1 \circ H_1^{-1} \circ H_1x_1) \\ &= (k-1)\theta(K_1x_1) - (k-1)\theta(K_1x_1) = 0, \end{aligned}$$

where in the second line we have used the fact that $\theta \circ K_i \circ H_i^{-1} < \theta \circ K_1 \circ H_1^{-1}$. If now $K_1x_1 \in \text{Ran}(K_j)$ for some $j \geq 2$ we obtain

$$\begin{aligned} \varphi(G_1(x_1)) &= \varphi(H_1x_1, K_1x_1) = \theta(K_1x_1) - \theta(K_j \circ H_j^{-1} \circ H_1x_1) \\ &> \theta(K_1x_1) - \theta(K_1 \circ H_1^{-1} \circ H_1x_1) \\ &= \theta(K_1x_1) - \theta(K_1x_1) = 0. \end{aligned}$$

It then follows that in both cases we have $\varphi(G_1(x_1)) > 0$. We now show that $\varphi(G_j(x_1)) = 0$. In fact,

$$\varphi(G_j(x_1)) = \varphi(H_jx_1, K_jx_1) = \theta(K_jx_1) - \theta(K_j \circ H_j^{-1} \circ H_jx_1) = 0.$$

Therefore for each $x_1 \in \text{Dom}(G_1) \cap \text{Dom}(G_j)$ we have

$$\varphi(G_1(x_1)) > 0 = \varphi(G_j(x_1)).$$

It now follows from Theorem 4.6 that γ is an extremal point of $\Pi(\mu_1, \nu)$. Therefore, it follows from Lemma 4.7 together with Steps 2 and 3 that γ is an extremal point of $\Pi(\mu_1, \mu_2, \mu_3)$. \square

REFERENCES

- [1] N. Ahmad, H.K. Kim and R.J. McCann, Optimal transportation, topology and uniqueness. *Bull. Math. Sci.* **1** (2011) 13–32.
- [2] S. Bianchini and L. Caravenna. On the extremality, uniqueness, and optimality of transference plans. *Bull. Inst. Math. Acad. Sin. (N.S.)* **4** (2009) 353–454.
- [3] V.I. Bogachev, *Measure Theory*. Vol. I, II. Springer-Verlag, Berlin (2007).
- [4] G. Buttazzo, L. De Pascale and P. Gori-Giorgi. Optimal transport formulation of electronic density-functional theory. *Phys. Rev. A* **85** (2012) 062502.
- [5] L. Caffarelli, Allocation maps with general cost functions. In *Partial Differential Equations and Applications*. Vol. 177 of *Lect. Notes Pure Appl. Math.* Dekker, New York (1996) 2935.
- [6] G. Carlier, On a class of multidimensional optimal transportation problems. *J. Convex Anal.* **10** (2003) 517–529.
- [7] G. Carlier and B. Nazaret, Optimal transportation for the determinant. *ESAIM: COCV* **14** (2008) 678–698.

- [8] P.A. Chiappori, A. Galichon and B. Salanie, *The roommate problem is more stable than you think*. Tech. Report (2014).
- [9] P.-A. Chiappori, R.J. McCann and L.P. Nesheim, Hedonic price equilibria, stable matching, and optimal transport: equivalence, topology, and uniqueness. *Econom. Theory* **42** (2010) 317–354.
- [10] M. Colombo, L. De Pascale and S. Di Marino, Multimarginal Optimal Transport Maps for 1-dimensional Repulsive Costs. *Canad. J. Math.* **67** (2015) 350–368.
- [11] C. Cotar, G. Friesecke and C. Kluppelberg, Density functional theory and optimal transportation with Coulomb cost. *Commun. Pure Appl. Math.* **66** (2013) 548599.
- [12] H.G. Kellerer, Duality theorems for marginal problems. *Z. Wahrsch. Verw. Gebiete* **67** (1984) 399–432.
- [13] W. Gangbo, *Habilitation Thesis*. Université de Metz (1995).
- [14] W. Gangbo and R.J. McCann, The geometry of optimal transportation. *Acta Math.* **177** (1996) 113–161.
- [17] W. Gangbo and R.J. McCann, Shape recognition via Wasserstein distance. *Quart. Appl. Math.* **58** (2000) 705–737.
- [16] W. Gangbo and A. Swiech, Optimal maps for the multidimensional mongekantorovich problem. *Commun. Pure Appl. Math.* **51** (1998) 23–45.
- [17] N. Ghoussoub and A. Moameni, Symmetric Monge–Kantorovich problems and polar decompositions of vector fields. *Geom. Funct. Anal.* **24** (2014) 1129–1166.
- [18] H. Heinich, Probleme de Monge pour n probabilités. *C. R. Math. Acad. Sci. Paris* **334** (2002) 793–795.
- [19] Y.-H. Kim and B. Pass, A general condition for Monge solutions in the multi-marginal optimal transport problem. *SIAM J. Math. Anal.* **46** (2014) 1538–1550.
- [20] V. Levin, Abstract cyclical monotonicity and Monge solutions for the general Monge–Kantorovich problem. *Set-Valued Analysis* **7** (1999) 7–32.
- [21] R. McCann, Polar factorization of maps on Riemannian manifolds. *Geom. Funct. Anal.* **11** (2001) 589–608.
- [22] R. McCann and L. Rifford, The intrinsic dynamics of optimal transport. *J. Ecole Polytechnique – Math.* **3** (2016) 67–98.
- [23] A. Moameni, Multi-marginal Monge–Kantorovich transport problems: A characterization of solutions. *C. R. Math. Acad. Sci. Paris* **352** (2014) 993–998.
- [24] A. Moameni, *Support of extremal doubly stochastic measures*. Preprint (2014).
- [25] A. Moameni, A characterization for solutions of the Monge–Kantorovich mass transport problem. *Math. Ann.* **365** (2016) 1279–1304.
- [26] B. Pass, *Structural results on optimal transportation plans*. Ph.D. thesis, University of Toronto (2011). Available at <http://www.ualberta.ca/pass/thesis.pdf>
- [27] B. Pass, Uniqueness and monge solutions in the multimarginal optimal transportation problem. *SIAM J. Math. Anal.* **43** (2011) 2758–2775.
- [28] B. Pass, On the local structure of optimal measures in the multimarginal optimal transportation problem. *Calc. Var. Partial Differ. Eq.* **43** (2012) 529–536.
- [29] B. Pass, Multi-marginal optimal transport: theory and applications. *ESAIM: M2AN* **49** (2015) 1771–1790.
- [30] C. Villani, Optimal transport, Old and new. *Grund. Math. Wiss.* Springer-Verlag, Berlin (2009).