Optimal and Better Transport Plans

Mathias Beiglböck [b,1], Martin Goldstern [a], Gabriel Maresch [a,2], Walter Schachermayer [b,3]

[a] *Institut für Diskrete Mathematik und Geometrie, Technische Universität Wien Wiedner Hauptstraße 8-10/104,1040 Wien, Austria*

[b] *Fakultät für Mathematik, Universität Wien Nordbergstraße 15, 1090 Wien, Austria*

Abstract

We consider the Monge-Kantorovich transport problem in a purely measure theoretic setting, i.e. without imposing continuity assumptions on the cost function. It is known that transport plans which are concentrated on c-monotone sets are optimal, provided the cost function c is either lower semi-continuous and finite, or continuous and may possibly attain the value ∞. We show that this is true in a more general setting, in particular for merely Borel measurable cost functions provided that $\{c = \infty\}$ is the union of a closed set and a negligible set. In a previous paper Schachermayer and Teichmann considered strongly c-monotone transport plans and proved that every strongly c-monotone transport plan is optimal. We establish that transport plans are strongly c-monotone if and only if they satisfy a "better" notion of optimality called robust optimality.

Key words: Monge-Kantorovich problem, c-cyclically monotone, strongly c-monotone, measurable cost function

1 Introduction

We consider the *Monge-Kantorovich transport problem* (μ, ν, c) for Borel probability measures μ, ν on Polish spaces X, Y and a Borel measurable cost function-

[1] Supported by the Austrian Science Fund (FWF) under grant S9612.
[2] Supported by the Austrian Science Fund (FWF) under grant Y328 and P18308.
[3] Supported by the Austrian Science Fund (FWF) under grant P19456, from the Vienna Science and Technology Fund (WWTF) under grant MA13 and from the Christian Doppler Research Association (CDG)

10 June 2009

tion $c : X \times Y \to [0, \infty]$. As standard references on the theory of mass transport we mention [1,9,14,15]. By $\Pi(\mu, \nu)$ we denote the set of all probability measures on $X \times Y$ with X-marginal μ and Y-marginal ν. For a Borel measurable *cost function* $c : X \times Y \to [0, \infty]$ the transport costs of a given transport plan $\pi \in \Pi(\mu, \nu)$ are defined by

$$I_c[\pi] := \int_{X \times Y} c(x, y) d\pi. \tag{1}$$

π is called a *finite* transport plan if $I_c[\pi] < \infty$.

A nice interpretation of the Monge-Kantorovich transport problem is given by Cédric Villani in Chapter 3 of the impressive monograph [15]:

> "Consider a large number of bakeries, producing breads, that should be transported each morning to cafés where consumers will eat them. The amount of bread that can be produced at each bakery, and the amount that will be consumed at each café are known in advance, and can be modeled as probability measures (there is a "density of production" and a "density of consumption") on a certain space, which in our case would be Paris (equipped with the natural metric such that the distance between two points is the length of the shortest path joining them). The problem is to find in practice where each unit of bread should go, in such a way as to minimize the total transport cost."

We are interested in *optimal* transport plans, i.e. minimizers of the functional $I_c[\cdot]$ and their characterization via the notion of c-monotonicity.

Definition 1.1 *A Borel set $\Gamma \subseteq X \times Y$ is called* c-monotone *if*

$$\sum_{i=1}^{n} c(x_i, y_i) \le \sum_{i=1}^{n} c(x_i, y_{i+1}) \tag{2}$$

for all pairs $(x_1, y_1), \dots, (x_n, y_n) \in \Gamma$ using the convention $y_{n+1} := y_1$. A transport plan π is called c-monotone if there exists a c-monotone Γ with $\pi(\Gamma) = 1$.

In the literature (e.g. [1,3,7,8,13]) the following characterization was established under various continuity assumptions on the cost function. Our main result states that those assumptions are not required.

Theorem 1 *Let X, Y be Polish spaces equipped with Borel probability measures μ, ν and let $c : X \times Y \to [0, \infty]$ a Borel measurable cost function.*

a. *Every finite optimal transport plan is c-monotone.*
b. *Every finite c-monotone transport plan is optimal if there exist a closed set F and a $\mu \otimes \nu$-null set N such that $\{(x, y) : c(x, y) = \infty\} = F \cup N$.*

Thus in the case of a cost function which does not attain the value ∞ the equivalence of optimality and c-monotonicity is valid without any restrictions beyond the obvious measurability conditions inherent in the formulation of the problem.

The subsequent construction due to Ambrosio and Pratelli in [1, Example 3.5] shows that if c is allowed to attain ∞ the implication "c-monotone $\Rightarrow$ optimal" does not hold without some additional assumption as in Theorem 1.b.

Example 1.2 (Ambrosio and Pratelli) *Let $X = Y = [0,1]$, equipped with Lebesgue measure $\lambda = \mu = \nu$. Pick $\alpha \in [0,1)$ irrational. Set*

$$\Gamma_0 = \{(x,x) : x \in X\}, \quad \Gamma_1 = \{(x, x \oplus \alpha) : x \in X\},$$

where $\oplus$ is addition modulo 1. Let $c : X \times Y \to [0,\infty]$ be such that $c = a \in [0,\infty)$ on Γ_0, $c = b \in [0,\infty)$ on Γ_1 and $c = \infty$ otherwise. It is then easy to check that Γ_0 and Γ_1 are c-monotone sets. Using the maps $f_0, f_1 : X \to X \times Y$, $f_0(x) = (x,x), f_1(x) = (x, x \oplus \alpha)$ one defines the transport plans $\pi_0 = f_{0\#}\lambda$, $\pi_1 = f_{1\#}\lambda$ supported by Γ_0 respectively Γ_1. Then π_0 and π_1 are finite c-monotone transport plans, but as $I_c[\pi_0] = a, I_c[\pi_1] = b$ it depends on the choice of a and b which transport plan is optimal. Note that in contrast to the assumption in Theorem 1.b the set $\{(x,y) \in X \times Y : c = \infty\}$ is open.

We want to remark that rather trivial (folkloristic) examples show that no optimal transport has to exist if the cost function doesn't satisfy proper continuity assumptions.

Example 1.3 *Consider the task to transport points on the real line (equipped with the Lebesgue measure) from the interval $[0,1)$ to $[1,2)$ where the cost of moving one point to another is the squared distance between these points ($X = [0,1), Y = [1,2)$, $c(x,y) = (x-y)^2$, $\mu = \nu = \lambda$). The simplest way to achieve this transport is to shift every point by 1. This results in transport costs of 1 and one easily checks that all other transport plans are more expensive.*

If we now alter the cost function to be 2 whenever two points have distance 1, i.e. if we set

$$\tilde{c}(x,y) = \begin{cases} 2 & \text{if } y = x+1 \\ c(x,y) & \text{otherwise} \end{cases},$$

it becomes impossible to find a transport plan $\pi \in \Pi(\mu,\nu)$ with total transport costs $I_{\tilde{c}}[\pi] = 1$, but it is still possible to achieve transport costs arbitrarily close to 1. (For instance, shift $[0, 1-\varepsilon)$ to $[1+\varepsilon, 2)$ and $[1-\varepsilon, 1)$ to $[1, 1+\varepsilon)$ for small $\varepsilon > 0$.)

1.1 History of the problem

The notion of c-monotonicity originates in convex analysis. The well known Rockafellar Theorem (see for instance [11, Theorem 3] or [14, Theorem 2.27]) and its generalization, Rüschendorf's Theorem (see [12, Lemma 2.1]), characterize c-monotonicity in $\mathbb{R}^n$ in terms of integrability. The definitions of c-concave functions and super-differentials can be found for instance in [14, Section 2.4].

Theorem (Rockafellar) *A non-empty set $\Gamma \subseteq \mathbb{R}^n \times \mathbb{R}^n$ is cyclically monotone (that is, c-monotone with respect to the squared euclidean distance) if and only if there exists a l.s.c. concave function $\varphi : \mathbb{R}^n \to \mathbb{R}$ such that Γ is contained in the super-differential $\partial(\varphi)$.*

Theorem (Rüschendorf) *Let X, Y be abstract spaces and $c : X \times Y \to [0, \infty]$ arbitrary. Let $\Gamma \subseteq X \times Y$ be c-monotone. Then there exists a c-concave function $\varphi : X \to Y$ such that Γ is contained in the c-super-differential $\partial^c(\varphi)$.*

Important results of Gangbo and McCann [3] and Brenier [14, Theorem 2.12] use these potentials to establish uniqueness of the solutions of the Monge-Kantorovich transport problem in $\mathbb{R}^n$ for different types of cost functions subject to certain regularity conditions.

Optimality implies c-monotonicity: This is evident in the discrete case if X and Y are finite sets. For suppose that π is a transport plan for which c-monotonicity is violated on pairs $(x_1, y_1), \ldots, (x_n, y_n)$ where all points $x_1, \ldots, x_n$ and $y_1, \ldots, y_n$ carry positive mass. Then we can reduce costs by sending the mass $\alpha > 0$, for α sufficiently small, from x_i to y_{i+1} instead of y_i, that is, we replace the original transport plan π with

$$\pi^\beta = \pi + \alpha \sum_{i=1}^{n} \delta_{(x_i, y_{i+1})} - \alpha \sum_{i=1}^{n} \delta_{(x_i, y_i)}. \tag{3}$$

(Here we are using the convention $y_{n+1} = y_1$.)

Gangbo and McCann ([3, Theorem 2.3]) show how continuity assumptions on the cost function can be exploited to extend this to an abstract setting. Hence one achieves:

Let X and Y be Polish spaces equipped with Borel probability measures μ, ν. Let $c : X \times Y \to [0, \infty]$ be a l.s.c. cost function. Then every finite optimal transport plan is c-monotone.

Using measure theoretic tools, as developed in the beautiful paper by Kellerer [6], we are able to extend this to Borel measurable cost functions (Theorem

1.a.) without any additional regularity assumption.

c-monotonicity implies optimality: In the case of finite spaces X, Y this again is nothing more than an easy exercise ([14, Exercise 2.21]). The problem gets harder in the infinite setting. It was first proved in [3] that for X, Y compact subsets of $\mathbb{R}^n$ and c a continuous cost function, c-monotonicity implies optimality. In a more general setting this was shown in [1, Theorem 3.2] for l.s.c. cost functions which additionally satisfy the moment conditions

$$\mu\left(\left\{x : \int_Y c(x,y)d\nu < \infty\right\}\right) > 0,$$
$$\nu\left(\left\{y : \int_X c(x,y)d\mu < \infty\right\}\right) > 0.$$

Further research into this direction was initiated by the following problem posed by Villani in [14, Problem 2.25]:

For $X = Y = \mathbb{R}^n$ and $c(x,y) = \|x-y\|^2$, the squared euclidean distance, does c-monotonicity of a transport plan imply its optimality?

A positive answer to this question was given independently by Pratelli in [8] and by Schachermayer and Teichmann in [13]. Pratelli proves the result for countable spaces and shows that it extends to the Polish case by means of approximation if the cost function $c : X \times Y \to [0, \infty]$ is continuous. The paper [13] pursues a different approach: The notion of *strong c-monotonicity* is introduced. From this property optimality follows fairly easily and the main part of the paper is concerned with the fact that strong c-monotonicity follows from the usual notion of c-monotonicity in the Polish setting if c is assumed to be l.s.c. and finitely valued.

Part (b) of Theorem 1 unifies these statements: Pratelli's result follows from the fact that for continuous $c : X \times Y \to [0, \infty]$ the set $\{c = \infty\} = c^{-1}[\{\infty\}]$ is closed; the Schachermayer-Teichmann result follows since for finite c the set $\{c = \infty\}$ is empty.

Similar to [13] our proofs are based on the concept of *strong c-monotonicity.* In Section 1.2 we present *robust optimality* which is a variant of optimality that we shall show to be equivalent to strong c-monotonicity. As not every optimal transport plan is also robustly optimal, this accounts for the somewhat provocative concept of "better than optimal" transport plans alluded to in the title of this paper.

Correspondingly the notion of strong c-monotonicity is in fact stronger than ordinary c-monotonicity (at least if c is allowed to assume the value ∞).

It turns out that optimality of a transport plan is intimately connected with the notion of *strong c-monotonicity* introduced in [13].

Definition 1.4 *A Borel set* $\Gamma \subseteq X \times Y$ *is* strongly c-monotone *if there exist Borel measurable functions* $\varphi : X \to [-\infty, \infty)$ *and* $\psi : Y \to [-\infty, \infty)$ *such that* $\varphi(x) + \psi(y) \leq c(x, y)$ *for all* $(x, y) \in X \times Y$ *and* $\varphi(x) + \psi(y) = c(x, y)$ *for all* $(x, y) \in \Gamma$. *A transport plan* $\pi \in \Pi(\mu, \nu)$ *is* strongly c-monotone *if* π *is concentrated on a strongly c-monotone Borel set* Γ.

Strong c-monotonicity implies c-monotonicity since

$$\sum_{i=1}^{n} c(x_{i+1}, y_i) \geq \sum_{i=1}^{n} \varphi(x_{i+1}) + \psi(y_i) = \sum_{i=1}^{n} \varphi(x_i) + \psi(y_i) = \sum_{i=1}^{n} c(x_i, y_i) \tag{4}$$

whenever $(x_1, y_1), \ldots, (x_n, y_n) \in \Gamma$.

If there are *integrable* functions φ and ψ witnessing that π is strongly c-monotone, then for every $\tilde{\pi} \in \Pi(\mu, \nu)$ we can estimate:

$$\begin{aligned} I_c[\pi] &= \int_\Gamma c(x, y) d\pi = \int_\Gamma [\varphi(x) + \psi(y)] d\pi = \\ &= \int_\Gamma \varphi(x) d\mu + \int_\Gamma \psi(y) d\nu = \int_\Gamma [\varphi(x) + \psi(y)] d\tilde{\pi} \leq I_c[\tilde{\pi}]. \end{aligned}$$

Thus in this case strong c-monotonicity implies optimality. However there is no reason why the Borel measurable functions φ, ψ appearing in Definition 1.4 should be integrable. In [13, Proposition 2.1] it is shown that for l.s.c. cost functions, there is a way of truncating which allows to also handle non-integrable functions φ and ψ. The proof extends to merely Borel measurable functions; hence we have:

Proposition 1.5 *Let* X, Y *be Polish spaces equipped with Borel probability measures* μ, ν *and let* $c : X \times Y \to [0, \infty]$ *be Borel measurable. Then every finite transport plan which is strongly c-monotone is optimal.*

No new ideas are required to extend [13, Proposition 2.1] to the present setting but since Proposition 1.5 is a crucial ingredient of several proofs in this paper we provide an outline of the argument in Section 3.

As it will turn out, strongly c-monotone transport plans even satisfy a "better" notion of optimality, called *robust optimality.*

Definition 1.6 *Let* X, Y *be Polish spaces equipped with Borel probability measures* μ, ν *and let* $c : X \times Y \to [0, \infty]$ *be a Borel measurable cost function. A transport plan* $\pi \in \Pi(\mu, \nu)$ *is* robustly optimal *if, for any Polish space* Z *and* *any finite Borel measure $\lambda \geq 0$ on Z, there exists a Borel measurable extension $\tilde{c} : (X \cup Z) \times (Y \cup Z) \to [0, \infty]$ satisfying*

$$\tilde{c}(a,b) = \begin{cases} c(a,b) & \textit{for} \quad a \in X, b \in Y \\ 0 & \textit{for} \quad a, b \in Z \\ < \infty & \textit{otherwise} \end{cases}$$

such that the measure $\tilde{\pi} := \pi + (id_Z \times id_Z)_{\#}\lambda$ is optimal on $(X \cup Z) \times (Y \cup Z)$. Note that $\tilde{\pi}$ is not a probability measure, but has total mass $1 + \lambda(Z) \in [1, \infty)$.

Note that since we allow the possibility $\lambda(Z) = 0$ every robustly optimal transport plan is in particular optimal in the usual sense.

Robust optimality has a colorful "economic" interpretation: a tycoon wants to enter the Parisian croissant consortium. She builds a storage of size $\lambda(Z)$ where she buys up croissants and sends them to the cafés. Her hope is that by offering low transport costs, the previously optimal transport plan π will not be optimal anymore, so that the traditional relations between bakeries and cafés will collapse. Of course, the authorities of Paris will try to defend their structure by imposing (possibly very high, but still finite) tolls for all transports to and from the tycoon's storage, thus resulting in finite costs $\tilde{c}(a, b)$ for $(a, b) \in (X \times Z) \cup (Z \times Y)$. In the case of robustly optimal π they can successfully defend themselves against the intruder.

Every robustly optimal transport π plan is optimal in the usual sense and hence also c-monotone. The crucial feature is that robust optimality implies strong c-monotonicity. In fact, the two properties are equivalent.

Theorem 2 *Let X, Y be Polish spaces equipped with Borel probability measures μ, ν and $c : X \times Y \to [0, \infty]$ a Borel measurable cost function. For a finite transport plan π the following assertions are equivalent:*

a. π is strongly c-monotone.
b. π is robustly optimal.

Example 5.1 below shows that robust optimality resp. strong c-monotonicity is in fact a stronger property than usual optimality.

1.3 Putting things together

Finally we want to point out that in the situation where c is finite all previously mentioned notions of monotonicity and optimality coincide. We can even pass to a slightly more general setting than finite cost functions and obtain the following result.

Theorem 3 *Let X, Y be Polish spaces equipped with Borel probability measures μ, ν and let $c : X \times Y \to [0, \infty]$ be Borel measurable and $\mu \otimes \nu$-a.e. finite. For a finite transport plan π the following assertions are equivalent:*

(1) π is optimal.
(2) π is c-monotone.
(3) π is robustly optimal.
(4) π is strongly c-monotone.

The equivalence of (1), (2) and (4) was established in [13] under the additional assumption that c is l.s.c. and finitely valued.

We sum up the situation under fully general assumptions. The upper line (1 and 2) relates to the optimality of a transport plan π. The lower line (3 and 4) contains the two equivalent strong concepts and implies the upper line but - without additional assumptions - not vice versa.

$$\begin{tikzcd}
1 \arrow[r, shift left] & 2 \arrow[l, dotted, shift left, "\text{Thm. 1}"'] \arrow[d, dotted, "\text{Thm. 3}"] \\
3 \arrow[u] \arrow[r, leftrightarrow, "\text{Thm. 2}"'] & 4
\end{tikzcd}$$

(1)	optimal
(2)	c-monotone
(3)	robustly optimal
(4)	strongly c-monotone

Fig. 1. Implications between properties of transport plans

Note that the implications symbolized by dotted lines in Figure 1 are not true without additional assumptions ((2) $\not\Rightarrow$ (1): Example 1.2, (1) $\not\Rightarrow$ (3) resp. (4): Example 5.1).

The paper is organized as follows: In Section 2 we prove that every optimal transport plan π is c-monotone (Theorem 1.a). In Section 3 we introduce an auxilliary property (*connectedness*) of the support of a transport plan and show that it allows to pass from c-monotonicity to strong c-monotonicity. Moreover we establish that strong c-monotonicity implies optimality (Proposition 1.5). Section 4 is concerned with the proof of Theorem 1.b. Finally we complete the proofs of Theorems 2 and 3 in Section 5.

We observe that in all the above discussion we only referred to the Borel structure of the Polish spaces X, Y, and never referred to the topological structure. Hence the above results (with the exception of Theorem 1.b.) hold true for standard Borel measure spaces.

In fact it seems likely that our results can be transferred to the setting of perfect measure spaces. (See [10] for a general overview resp. [9] for a treatment of problems of mass transport in this framework.) However we do not pursue this direction.

Acknowledgement. The authors are indebted to the extremely careful referee who noticed many inaccuracies resp. mistakes and whose insightful suggestions led to a more accessible presentation of several results in this paper.

2 Improving Transports

Assume that some transport plan $\pi \in \Pi(\mu, \nu)$ is given. From a purely heuristic point of view there are either few tupels $((x_1, y_1), \dots, (x_n, y_n))$ along which c-monotonicity is violated, or there are many such tuples, in which case π can be enhanced by rerouting the transport along these tuples. As the notion of c-monotonicity refers to n-tuples it turns out that it is necessary to consider finitely many measure spaces to properly formulate what is meant by "few" resp. "many".

Let $X_1, \dots, X_n$ be Polish spaces equipped with finite Borel measures $\mu_1, \dots, \mu_n$. By $\Pi(\mu_1, \dots, \mu_n) \subseteq \mathcal{M}(X_1 \times \dots \times X_n)$ we denote the set of all Borel measures on $X_1 \times \dots \times X_n$ such that the i-th marginal measure coincides with the Borel measure μ_i for $i = 1, \dots, n$. By $p_{X_i} : X_1 \times \dots \times X_n \to X_i$ we denote the projection onto the i-th component. $B \subseteq X_1 \times \dots \times X_n$ is called an *L-shaped null set* if there exist null sets $N_1 \subseteq X_1, \dots, N_n \subseteq X_n$ such that $B \subseteq \bigcup_{i=1}^n p_{X_i}^{-1}[N_i]$.

The Borel sets of $X_1 \times \dots \times X_n$ satisfy a nice dichotomy. They are either L-shaped null sets or they carry a positive measure whose marginals are absolutely continuous with respect to $\mu_1, \dots, \mu_n$:

Proposition 2.1 *Let $X_1, \dots, X_n, n \geq 2$ be Polish spaces equipped with Borel probability measures $\mu_1, \dots, \mu_n$. Then for any Borel set $B \subseteq X_1 \times \dots \times X_n$ let*

$$P(B) := \sup \{\pi(B) : \pi \in \Pi(\mu_1, \dots, \mu_n)\} \tag{5}$$

$$L(B) := \inf \left\{ \sum_{i=1}^n \mu_i(B_i) : B_i \subseteq X_i \text{ and } B \subseteq \bigcup_{i=1}^n p_{X_i}^{-1}[B_i] \right\}. \tag{6}$$

Then $P(B) \geq 1/n \, L(B)$. In particular B satisfies one of the following alternatives:

a. B is an L-shaped null set.
b. There exists $\pi \in \Pi(\mu_1, \ldots, \mu_n)$ such that $\pi(B) > 0$.

The main ingredient in the proof Proposition 2.1 is the following duality theorem due to Kellerer (see [6, Lemma 1.8(a), Corollary 2.18]).

Theorem (Kellerer) *Let $X_1, \ldots, X_n, n \geq 2$ be Polish spaces equipped with Borel probability measures $\mu_1, \ldots, \mu_n$ and assume that $c : X = X_1 \times \cdots \times X_n \to \mathbb{R}$ is Borel measurable and that $\overline{c} := \sup_X c, \underline{c} := \inf_X c$ are finite. Set*

$$I(c) = \inf\left\{\int_X c \; d\pi : \pi \in \Pi(\mu_1, \ldots, \mu_n)\right\},$$

$$S(c) = \sup\left\{\sum_{i=1}^{n} \int_{X_i} \varphi_i \; d\mu_i : c(x_1, \ldots, x_n) \geq \sum_{i=1}^{n} \varphi_i(x_i), \frac{\overline{c}}{n} - (\overline{c} - \underline{c}) \leq \varphi_i \leq \frac{\overline{c}}{n}\right\}.$$

Then $I(c) = S(c)$.

PROOF of Proposition 2.1. Observe that $-I(-\mathbb{1}_B) = P(B)$ and that

$$-S(-\mathbb{1}_B) = \inf\left\{\sum_{i=1}^{n} \int_{X_i} \chi_i d\mu_i : \mathbb{1}_B(x_1, \ldots, x_n) \leq \sum_{i=1}^{n} \chi_i(x_i), 0 \leq \chi_i \leq 1\right\}. \quad (7)$$

By Kellerer's Theorem $-S(-\mathbb{1}_B) = -I(-\mathbb{1}_B)$. Thus it remains to show that $-S(-\mathbb{1}_B) \geq \frac{1}{n} L(B)$. Fix functions $\chi_1, \ldots, \chi_n$ as in (7). Then for $(x_1, \ldots, x_n) \in B$ one has $1 = \mathbb{1}_B(x_1, \ldots, x_n) \leq \sum_{i=1}^{n} \chi_i(x_i)$ and hence there exists some i such that $\chi_i(x_i) \geq \frac{1}{n}$. Thus $B \subseteq \bigcup_{i=1}^{n} p_{X_i}^{-1}[\{\chi_i \geq \frac{1}{n}\}]$. It follows that

$$\begin{aligned} -S(-\mathbb{1}_B) &\geq \inf\left\{\sum_{i=1}^{n} \int_{X_i} \chi_i d\mu_i : B \subseteq \bigcup_{i=1}^{n} p_{X_i}^{-1}[\{\chi_i \geq \tfrac{1}{n}\}], 0 \leq \chi_i \leq 1\right\} \\ &\geq \inf\left\{\sum_{i=1}^{n} \tfrac{1}{n}\mu_i(\{\chi_i \geq \tfrac{1}{n}\}) : B \subseteq \bigcup_{i=1}^{n} p_{X_i}^{-1}[\{\chi_i \geq \tfrac{1}{n}\}]\right\} \geq \tfrac{1}{n} L(B) \end{aligned}$$

From this we deduce that either $L(B) = 0$ or there exists $\pi \in \Pi(\mu_1, \ldots, \mu_n)$ such that $\pi(B) > 0$. The last assertion of Proposition 2.1 now follows from the following Lemma due to Richárd Balka and Márton Elekes (private communication). $\square$

Lemma 2.2 *Suppose that $L(B) = 0$ for a Borel set $B \subseteq X_1 \times \cdots \times X_n$. Then B is an L-shaped null set.*

PROOF. Fix $\varepsilon > 0$ and Borel sets $B_1^{(k)}, \ldots, B_n^{(k)}$ with $\mu_i(B_i^{(k)}) \leq \varepsilon\, 2^{-k}$ such that for each k

$$B \subseteq p_{X_1}^{-1}[B_1^{(k)}] \cup \ldots \cup p_{X_n}^{-1}[B_n^{(k)}].$$

Let $B_i := \bigcup_{k=1}^{\infty} B_i^{(k)}$ for $i = 2, \ldots, n$ such that

$$B \subseteq p_{X_1}^{-1}[B_1^{(k)}] \cup p_{X_2}^{-1}[B_2] \cup \ldots \cup p_{X_n}^{-1}[B_n]$$

for each $k \in \mathbb{N}$. Thus with $B_1 := \bigcap_{k=1}^{\infty} B_1^{(k)}$,

$$B \subseteq p_{X_1}^{-1}[B_1] \cup p_{X_2}^{-1}[B_2] \cup \ldots \cup p_{X_n}^{-1}[B_n].$$

Hence we can assume from now on that $\mu_1(B_1) = 0$ and that $\mu_i(B_i)$ is arbitrarily small for $i = 2, \ldots, n$. Iterating this argument in the obvious way we get the statement. $\square$

Remark 2.3 *In the case $n = 2$ it was shown in [6, Proposition 3.3] that $L(B) = P(B)$ for every Borel set $B \subseteq X_1 \times X_2$. However, for $n > 2$, equality does not hold true, cf. [6, Example 3.4].*

Definition 2.4 *Let X, Y be Polish spaces. For a Borel measurable cost function $c : X \times Y \to [0, \infty], n \in \mathbb{N}$ and $\varepsilon > 0$ we set*

$$B_{n,\varepsilon} := \left\{ (x_i, y_i)_{i=1}^n \in (X \times Y)^n : \sum_{i=1}^n c(x_i, y_i) \geq \sum_{i=1}^n c(x_i, y_{i+1}) + \varepsilon \right\}. \tag{8}$$

The definition of the sets $B_{n,\varepsilon}$ is implicitly given in [3, Theorem 2.3]. The idea behind it is, that $(x_i, y_i)_{i=1}^n \in B_{n,\varepsilon}$ tells us that transport costs can be reduced if "x_i is transported to y_{i+1} instead of y_i" (recall the conventions $x_{n+1} = x_1$ resp. $y_{n+1} = y_1$). In what follows we make this statement precise and give a coordinate free formulation.

Denote by $\sigma, \tau : (X \times Y)^n \to (X \times Y)^n$ the shifts defined via

$$\sigma : (x_i, y_i)_{i=1}^n \mapsto (x_{i+1}, y_{i+1})_{i=1}^n \tag{9}$$
$$\tau : (x_i, y_i)_{i=1}^n \mapsto \quad (x_i, y_{i+1})_{i=1}^n. \tag{10}$$

Observe that $\sigma^n = \tau^n = \mathrm{Id}_{(X \times Y)^n}$ and that σ and τ commute. Also note that the set $B_{n,\varepsilon}$ from (8) is σ-invariant (i.e. $\sigma(B_{n,\varepsilon}) = B_{n,\varepsilon}$), but in general not τ-invariant. Denote by $p_i : (X \times Y)^n \to X \times Y$ the projection on the i-th component of the product. The projections $p_X : X \times Y \to X, (x, y) \mapsto x$ and $p_Y : X \times Y \to Y, (x, y) \mapsto y$ are defined as usual and there will be no danger of confusion.

Lemma 2.5 *Let X, Y be Polish spaces equipped with Borel probability measures μ, ν. Let π be a transport plan. Then one of the following alternatives holds:*

a. π is c-monotone,

b. there exist $n \in \mathbb{N}$, $\varepsilon > 0$ and a measure $\kappa \in \Pi(\pi, \dots, \pi)$ such that $\kappa(B_{n,\varepsilon}) > 0$. Moreover κ can be taken to be both σ and τ invariant.

PROOF. Suppose that $B_{n,\varepsilon}$ is an L-shaped null set for all $n \in \mathbb{N}$ and every $\varepsilon > 0$. Then there are Borel sets $S^1_{n,\varepsilon}, \dots, S^n_{n,\varepsilon} \subseteq X \times Y$ of full π-measure such that

$$\left(S^1_{n,\varepsilon} \times \ldots \times S^n_{n,\varepsilon}\right) \cap B_{n,\varepsilon} = \emptyset$$

and π is concentrated on the c-monotone set

$$S = \bigcap_{k=1}^{\infty} \bigcap_{n=1}^{\infty} \bigcap_{i=1}^{n} S^i_{n,1/k}.$$

If there exist $n \in \mathbb{N}$ and $\varepsilon > 0$ such that $B_{n,\varepsilon}$ is not an L-shaped null set, we apply Proposition 2.1 to conclude the existence of a measure $\kappa \in \Pi(\pi, \dots, \pi)$ with $\kappa(B_{n,\varepsilon}) > 0$. To achieve the desired invariance, simply replace κ by

$$\frac{1}{n^2} \sum_{i,j=1}^{n} (\sigma^i \circ \tau^j)_{\#}\kappa \qquad \square \tag{11}$$

We are now in the position to prove **Theorem 1.a**, i.e.

Let X, Y be Polish spaces equipped with Borel probability measures μ, ν and let $c : X \times Y \to [0, \infty]$ be a Borel measurable cost function. If π is a finite optimal transport plan, then π is c-monotone.

PROOF. Suppose by contradiction that π is optimal, $I_c[\pi] < \infty$ but π is not c-monotone. Then by Lemma 2.5 there exist $n \in \mathbb{N}$, $\varepsilon > 0$ and an invariant measure $\kappa \in \Pi(\pi, \dots, \pi)$ which gives mass $\alpha > 0$ to the Borel set $B_{n,\varepsilon} \subseteq (X \times Y)^n$. Consider now the restriction of κ to $B_{n,\varepsilon}$ defined via $\hat{\kappa}(A) := \kappa(A \cap B_{n,\varepsilon})$ for Borel sets $A \subseteq (X \times Y)^n$. $\hat{\kappa}$ is σ-invariant since both the measure κ and the Borel set $B_{n,\varepsilon}$ are σ-invariant. Denote the marginal of $\hat{\kappa}$ in the first coordinate $(X \times Y)$ of $(X \times Y)^n$ by $\hat{\pi}$. Due to σ-invariance we have

$$p_{i\#}\hat{\kappa} = p_{i\#}(\sigma_{\#}\hat{\kappa}) = (p_i \circ \sigma)_{\#}\hat{\kappa} = p_{i+1\#}\hat{\kappa},$$

i.e. all marginals coincide and we have $\hat{\kappa} \in \Pi(\hat{\pi}, \dots, \hat{\pi})$. Furthermore, since $\hat{\kappa} \leq \kappa$, the same is true for the marginals, i.e. $\hat{\pi} \leq \pi$. Denote the marginal of $\tau_{\#}\hat{\kappa}$ in the first coordinate $(X \times Y)$ of $(X \times Y)^n$ by $\hat{\pi}_\beta$. As σ and τ commute, $\tau_{\#}\hat{\kappa}$ is σ-invariant, so the marginals in the other coordinates coincide with $\hat{\pi}_\beta$. An easy calculation shows that $\hat{\pi}$ and $\hat{\pi}_\beta$ have the same marginals in X resp. Y:

$$\begin{aligned} p_{X\#}\hat{\pi}_\beta &= p_{X\#}(p_{i\#}(\tau_{\#}\hat{\kappa})) = (p_X \circ p_i \circ \tau)_{\#}\hat{\kappa} = (p_X \circ p_i)_{\#}\hat{\kappa} = p_{X\#}\hat{\pi}, \\ p_{Y\#}\hat{\pi}_\beta &= p_{Y\#}(p_{i\#}(\tau_{\#}\hat{\kappa})) = (p_Y \circ p_i \circ \tau)_{\#}\hat{\kappa} = (p_Y \circ p_{i+1})_{\#}\hat{\kappa} = p_{Y\#}\hat{\pi}. \end{aligned}$$

The equality of the total masses is proved similarly:

$$\alpha = \hat{\pi}_\beta(X \times Y) = (p_i \circ \tau)_\# \hat{\kappa}(X \times Y) = p_{i\#}\hat{\kappa}(X \times Y) = \hat{\pi}(X \times Y).$$

Next we compute the transport costs associated to $\hat{\pi}_\beta$:

$$\begin{aligned}
\int_{X\times Y} c\, d\hat{\pi}_\beta = & \int_{(X\times Y)^n} c \circ p_1 \, d(\tau_\# \hat{\kappa}) && \text{(marginal property)} \\
= & \frac{1}{n}\sum_{i=1}^{n} \int_{(X\times Y)^n} c \circ p_i \, d(\tau_\# \hat{\kappa}) && \text{(σ-invariance)} \\
= & \frac{1}{n}\sum_{i=1}^{n} \int_{(X\times Y)^n} (c \circ p_i \circ \tau)\, d\hat{\kappa} && \text{(push-forward)} \\
= & \frac{1}{n}\sum_{i=1}^{n} \int_{B_{n,\varepsilon}} (c \circ p_i \circ \tau)\, d\kappa && \text{(definition of $\hat{\kappa}$)} \\
\leq & \frac{1}{n} \int_{B_{n,\varepsilon}} \left[\sum_{i=1}^{n} (c \circ p_i) \; - \varepsilon \right] d\kappa && \text{(definition of $B_{n,\varepsilon}$)} \\
= & \int_{X\times Y} c d\hat{\pi} - \varepsilon \frac{\alpha}{n}. && \text{(definition of $\hat{\pi}$)}
\end{aligned}$$

To improve the transport plan π we define

$$\pi_\beta := (\pi - \hat{\pi}) + \hat{\pi}_\beta. \tag{12}$$

Recall that $\pi - \hat{\pi}$ is a positive measure, so π_β is a positive measure. As $\hat{\pi}$ and $\hat{\pi}_\beta$ have the same total mass, π_β is a probability measure. Furthermore $\hat{\pi}$ and $\hat{\pi}_\beta$ have the same marginals, so π_β is indeed a transport plan. We have

$$I_c[\pi_\beta] = I_c[\pi] + \int_{X\times Y} c\, d(\hat{\pi}_\beta - \hat{\pi}) \leq I_c[\pi] - \varepsilon \frac{\alpha}{n} < I_c[\pi]. \qquad \square \tag{13}$$

3 Connecting c-monotonicity and strong c-monotonicity

The Ambrosio-Pratelli example (Example 1.2) shows that c-monotonicity need not imply strong c-monotonicity in general. Subsequently we shall present a condition which ensures that this implication is valid.

A c-monotone transport plan resists the attempt of enhancement by means of cyclically rerouting. This, however, may be due to the fact that cyclical rerouting is a priori impossible due to infinite transport costs on certain routes.

Continuing Villani's interpretation, a situation where rerouting in this consortium of bakeries and cafés is possible in a satisfactory way is as follows:

Suppose that bakery $x = x_0$ is able to produce one more croissant than it already does and that café $\tilde{y}$ is short of one croissant. It might not be possible to transport the additional croissant itself to the café in need, as the costs $c(x, \tilde{y})$ may be infinite. Nevertheless it might be possible to find another bakery x_1 (which usually supplies café y_1) such that bakery x can transport (with finite costs!) the extra croissant to y_1; this leaves us with a now unused item from bakery x_1, which can be transported to $\tilde{y}$ with finite costs. Of course we allow not only one, but finitely many intermediate pairs $(x_1, y_1), \ldots, (x_n, y_n)$ of bakeries/cafés to achieve this relocation of the additional croissant.

In the Ambrosio-Pratelli example we can reroute from a point $(x, x \oplus \alpha) \in \Gamma_1$ to a point $(\tilde{x}, \tilde{x} \oplus \alpha) \in \Gamma_1$ only if there exists $n \in \mathbb{N}$ such that $x \oplus (n\alpha) = \tilde{x}$. In particular, irrationality of α implies that if we can redirect with finite costs from $(x, x \oplus \alpha)$ to $(\tilde{x}, \tilde{x} \oplus \alpha)$ we never can redirect back from $(\tilde{x}, \tilde{x} \oplus \alpha)$ to $(x, x \oplus \alpha)$.

Definition 3.1 *Let X, Y be Polish spaces equipped with Borel probability measures μ, ν, let $c : X \times Y \to [0, \infty]$ be a Borel measurable cost function and $\Gamma \subseteq X \times Y$ a Borel measurable set on which c is finite. We define*

a. $(x, y) \lesssim (\tilde{x}, \tilde{y})$ *if there exist pairs* $(x_0, y_0), \ldots, (x_n, y_n) \in \Gamma$ *such that* $(x, y) = (x_0, y_0)$ *and* $(\tilde{x}, \tilde{y}) = (x_n, y_n)$ *and* $c(x_1, y_0), \ldots, c(x_n, y_{n-1}) < \infty$.
b. $(x, y) \approx (\tilde{x}, \tilde{y})$ *if* $(x, y) \lesssim (\tilde{x}, \tilde{y})$ *and* $(x, y) \gtrsim (\tilde{x}, \tilde{y})$.

We call (Γ, c) connecting *if c is finite on Γ and $(x, y) \approx (\tilde{x}, \tilde{y})$ for all (x, y), $(\tilde{x}, \tilde{y}) \in \Gamma$.*

These relations were introduced in [15, Chapter 5, p.75] and appear in a construction due to Stefano Bianchini.

When there is any danger of confusion we will write $\lesssim_{c,\Gamma}$ and $\approx_{c,\Gamma}$, indicating the dependence on Γ and c. Note that $\lesssim$ is a pre-order, i.e. a transitive and reflexive relation, and that $\approx$ is an equivalence relation. We will also need the projections $\lesssim_X, \approx_X$ resp. $\lesssim_Y, \approx_Y$ of these relations onto the set $p_X[\Gamma] \subseteq X$ resp. $p_Y[\Gamma] \subseteq Y$. The projection is defined in the obvious way: $x \lesssim_X \tilde{x}$ if there exist $y, \tilde{y}$ such that $(x, y), (\tilde{x}, \tilde{y}) \in \Gamma$ and $(x, y) \lesssim (\tilde{x}, \tilde{y})$ holds.

The other relations are defined analogously. The projections of $\lesssim$ are again pre-orders and the projections of $\approx$ are again equivalence relations, provided c is finite on Γ. The equivalence classes of $\approx$ and its projections are compatible in the sense that $[(x, y)]_\approx = ([x]_{\approx_X} \times [y]_{\approx_Y}) \cap \Gamma$. The elementary proofs of these facts are left to the reader.

The main objective of this section is to prove Proposition 3.2, based on several lemmas which will be introduced throughout the section.

Proposition 3.2 *Let X, Y be Polish spaces equipped with Borel probability measures μ, ν and let $c : X \times Y \to [0, \infty]$ be a Borel measurable cost function. Let π be a finite transport plan. Assume that there exists a c-monotone set $\Gamma \subseteq X \times Y$ with $\pi(\Gamma) = 1$ on which c is finite, such that (Γ, c) is connecting. Then π is strongly c-monotone.*

In the proof of Proposition 3.2 we will establish the existence of the functions φ, ψ using the construction given in [12], see also [14, Chapter 2] and [1, Theorem 3.2]. As we do not impose any continuity assumptions on the cost function c, we can not prove the Borel measurability of φ and ψ by using limiting procedures similar to the methods used in [1,12–14]. Instead we will use the following projection theorem, a proof of which can be found in [2, Theorem III.23] by analyists or in [5, Section 29.B] by readers who have some interest in set theory.

Proposition 3.3 [4] *Let X and Y be Polish spaces, $A \subseteq X$ a Borel measurable set and $f : X \to Y$ a Borel measurable map. Then $B := f(A)$ is* universally measurable, *i.e. B is measurable with respect to the completion of every σ-finite Borel measure on Y.*

The system of universally measurable sets is a σ-algebra. If X is a Polish space, we call a function $f : X \to [-\infty, \infty]$ universally measurable if the pre-image of every Borel set is universally measurable.

Lemma 3.4 *Let X be a Polish space and μ a finite Borel measure on X. If $\varphi : X \to [-\infty, \infty)$ is universally measurable, then there exists a Borel measurable function $\tilde{\varphi} : X \to [-\infty, \infty)$ such that $\tilde{\varphi} \le \varphi$ everywhere and $\varphi = \tilde{\varphi}$ almost everywhere.*

PROOF. Let $(I_n)_{n=1}^{\infty}$ be an enumeration of the intervals $[a, b)$ with endpoints in $\mathbb{Q}$ and denote the completion of μ by $\tilde{\mu}$. Then for each $n \in \mathbb{N}$, $\varphi^{-1}[I_n]$ is $\tilde{\mu}$-measurable and hence the union of a Borel set B_n and a $\tilde{\mu}$-null set N_n. Let N be a Borel null set which covers $\bigcup_{n=1}^{\infty} N_n$. Let $\tilde{\varphi}(x) = \varphi(x) - \infty \cdot \mathbf{1}_N(x)$. Clearly $\tilde{\varphi}(x) \le \varphi(x)$ for all $x \in X$ and $\varphi(x) = \tilde{\varphi}(x)$ for $\tilde{\mu}$-almost all $x \in X$. Furthermore, $\tilde{\varphi}$ is Borel measurable since $(I_n)_{n=1}^{\infty}$ is a generator of the Borel σ-algebra on $[-\infty, \infty)$ and for each $n \in \mathbb{N}$ we have that $\tilde{\varphi}^{-1}[I_n] = B_n \setminus N$ is a Borel set. □

The following definition of the functions $\varphi_n, n \in \mathbb{N}$ resp. φ is reminiscent of the construction in [12].

[4] Sets which are images of Borel sets under measurable functions are called *analytic* in descripitive set theory. Lusin first noticed that analytic sets are universally measurable. Details can be found for instance in [5].

Lemma 3.5 *Let X, Y be Polish spaces, $c : X \times Y \to [0, \infty]$ a Borel measurable cost function and $\Gamma \subseteq X \times Y$ a Borel set. Fix $(x_0, y_0) \in \Gamma$ and assume that c is finite on Γ. For $n \in \mathbb{N}$, define $\varphi_n : X \times \Gamma^n \to (-\infty, \infty]$ by*

$$\varphi_n(x; x_1, y_1, \ldots, x_n, y_n) = [c(x, y_n) - c(x_n, y_n)] + \sum_{i=0}^{n-1} [c(x_{i+1}, y_i) - c(x_i, y_i)] \quad (14)$$

Then the map $\varphi : X \to [-\infty, \infty]$ defined by

$$\varphi(x) = \inf\{\varphi_n(x; x_1, y_1, \ldots, x_n, y_n) : n \geq 1, (x_i, y_i)_{i=1}^n \in \Gamma^n\} \quad (15)$$

is universally measurable.

PROOF. First note that the Borel σ-algebra on $[-\infty, \infty]$ is generated by intervals of the form $[-\infty, \alpha)$, thus it is sufficient to determine the pre-images of those sets under φ. We have

$$\varphi(x) < \alpha \leftrightarrow \exists n \in \mathbb{N}\ \exists (x_1, y_1), \ldots, (x_n, y_n) \in \Gamma : \varphi_n(x; x_1, y_1, \ldots, x_n, y_n) < \alpha.$$

The set $\varphi_n^{-1}[[-\infty, \alpha)]$ is Borel measurable. Hence

$$\varphi^{-1}[[-\infty, \alpha)] = \bigcup_{n \in \mathbb{N}} p_X[\varphi_n^{-1}[[-\infty, \alpha)]]$$

is the countable union of projections of Borel sets. Since projections of Borel sets are universally measurable by Proposition 3.3, $\varphi^{-1}[[-\infty, \alpha)]$ belongs also to the σ-algebra of universally measurable sets. $\square$

Lemma 3.6 *Let X, Y be Polish spaces and $c : X \times Y \to [0, \infty]$ a Borel measurable cost function. Suppose Γ is c-monotone, c is finite on Γ and (Γ, c) is connecting. Fix $(x_0, y_0) \in \Gamma$. Then the map φ from (15) is finite on $p_X[\Gamma]$. Furthermore*

$$\varphi(x) \leq \varphi(x') + c(x, y) - c(x', y) \quad \forall x \in X,\ (x', y) \in \Gamma. \quad (16)$$

PROOF. Fix $x \in p_X[\Gamma]$. Since $x_0 \lesssim x$ (recall Definition 3.1), we can find $x_1, y_1, \ldots, x_n, y_n$ such that $\varphi_n(x; x_1, y_1, \ldots, x_n, y_n) < \infty$. Hence $\varphi(x) < \infty$. Proving $\varphi(x) > -\infty$ involves some wrestling with notation but, not very surprisingly, it comes down to applying the fact that $x \lesssim x_0$. Let $a_1 = x$ and choose $b_1, a_2, b_2, \ldots, a_m, b_m$ such that $(a_1, b_1), \ldots, (a_m, b_m) \in \Gamma$ and $c(a_2, b_1), \ldots, c(a_m, b_{m-1}), c(x, b_m) < \infty$. Assume now that $x_1, y_1, \ldots, x_n, y_n$ are given such that $\varphi_n(x; x_1, y_1, \ldots, x_n, y_n) < \infty$. Put $x_{n+i} = a_i$ and $y_{n+i} = b_i$ for $i \in \{1, \ldots, m\}$. Due to c-monotonicity of Γ and the finiteness of all involved terms we have:

$$0 \leq [c(x_0, y_{n+m}) - c(x_{n+m}, y_{n+m})] + \sum_{i=0}^{n+m-1} [c(x_{i+1}, y_i) - c(x_i, y_i)],$$

which, after regrouping yields

$$\begin{aligned}\alpha := [c(x_0, b_m) - c(a_m, b_m)] + \sum_{i=1}^{m-1} [c(a_{i+1}, b_i) - c(a_i, b_i)] \\ \leq [c(x, y_n) - c(x_n, y_n)] + \sum_{i=0}^{n-1} [c(x_{i+1}, y_i) - c(x_i, y_i)]. \quad (17)\end{aligned}$$

Note that the right hand side of (17) is just $\varphi_n(x; x_1, y_1, \ldots, x_n, y_n)$. Thus passing to the infimum we see that $\varphi(x) \geq \alpha > -\infty$. To prove the remaining inequality, observe that the right hand side of (16) can be written as

$$\inf\{\varphi_n(x; x_1, y_1, \ldots, x_n, y_n) : n \geq 1, (x_i, y_i)_{i=1}^n \in \Gamma^n \text{ and } (x_n, y_n) = (x', y)\}$$

whereas the left hand side of (16) is the same, without the restriction $(x_n, y_n) = (x', y)$. $\square$

Lemma 3.7 *Let X, Y be Polish spaces and $c : X \times Y \to [0, \infty]$ a Borel measurable cost function. Let $X_0 \subseteq X$ be a non-empty Borel set and let $\varphi : X_0 \to \mathbb{R}$ be a Borel measurable function. Then the c-transform $\psi : Y \to [-\infty, \infty)$, defined as*

$$\psi(y) := \inf_{x \in X_0} [c(x, y) - \varphi(x)] \tag{18}$$

is universally measurable.

PROOF. As in the proof of Lemma 3.4 we consider the set $\psi^{-1}[[-\infty, \alpha)]$:

$$\psi(y) < \alpha \leftrightarrow \exists x \in X_0 : c(x, y) - \varphi(x) < \alpha.$$

Note that the set $\{(x, y) \in X_0 \times Y : c(x, y) - \varphi(x) < \alpha\}$ is Borel. Thus

$$\psi^{-1}[[-\infty, \alpha)] = p_X[\{(x, y) \in X_0 \times Y : c(x, y) - \varphi(x) < \alpha\}]$$

is the projection of a Borel set, hence universally measurable. $\square$

We are now able to prove the main result of this section.

PROOF of Proposition 3.2. Let $\Gamma \subseteq X \times Y$ be a c-monotone Borel set such that $\pi(\Gamma) = 1$ and the pair (Γ, c) is connecting. Let φ be the map from

Lemma 3.5. Using Lemma 3.4 and Lemma 3.6, and eventually passing to a subset of full π-measure, we may assume that φ is Borel measurable, that $X_0 := p_X[\Gamma]$ is a Borel set and that

$$c(x', y) - \varphi(x') \leq c(x, y) - \varphi(x) \quad \forall x \in X_0,\, (x', y) \in \Gamma. \tag{19}$$

Note that (19) follows from (16) in Lemma 3.6. Here we consider $x \in X_0$ in order to ensure that $\varphi(x)$ is finite on X_0. Now consider the c-transform

$$\psi(y) := \inf_{x \in X_0} [c(x, y) - \varphi(x)], \tag{20}$$

which by Lemma 3.7 is universally measurable. Fix $y \in p_Y[\Gamma]$. Using (19) we see that the infimum in (20) is attained at a point $x_0 \in X_0$ satisfying $(x_0, y) \in \Gamma$. This implies that $\varphi(x) + \psi(y) = c(x, y)$ on Γ and $\varphi(x) + \psi(y) \leq c(x, y)$ on $p_X[\Gamma] \times p_Y[\Gamma]$. To guarantee this inequality on the whole product $X \times Y$, one has to redefine φ and ψ to be $-\infty$ on the complement of $p_X[\Gamma]$ resp. $p_Y[\Gamma]$. Applying Lemma 3.4 once more, we find that there exists a Borel set $N \subseteq Y$ of zero ν-measure, such that $\tilde{\psi}(y) = \psi(y) - \infty \cdot \mathbb{1}_N(y)$ is Borel measurable. Finally, replace Γ by $\Gamma \cap (X \times (Y \setminus N))$ and ψ by $\tilde{\psi}$. $\square$

We conclude this section by proving that every strongly c-monotone transport plan is optimal (**Proposition 1.5**).

Let X, Y be Polish spaces equipped with Borel probability measures μ, ν and let $c : X \times Y \to [0, \infty]$ be Borel measurable. Then every finite transport plan which is strongly c-monotone is optimal.

PROOF. Let π_0 be a strongly c-monotone transport plan. Then, according to the definition, there exist Borel functions $\varphi(x)$ and $\psi(y)$ taking values in $[-\infty, \infty)$ such that

$$\varphi(x) + \psi(y) \leq c(x, y) \tag{21}$$

everywhere on $X \times Y$ and equality holds π_0-a.e. We define the truncations $\varphi_n = (n \wedge (\varphi \vee -n)), \psi_n = (n \wedge (\psi \vee -n))$ and let $\xi_n(x, y) := \varphi_n(x) + \psi_n(y)$ resp. $\xi(x, y) := \varphi(x) + \psi(y)$. Note that $\varphi_n, \psi_n, \xi_n, \xi$ are Borel measurable. By elementary considerations which are left the reader, we get pointwise monotone convergence $\xi_n \uparrow \xi$ on the set $\{\xi \geq 0\}$ resp. $\xi_n \downarrow \xi$ on the set $\{\xi \leq 0\}$ Let π_1 be an arbitrary finite transport plan; to compare $I_c[\pi_0]$ and $I_c[\pi_1]$ we make the following observations:

a. By monotone convergence

$$\int_{\{\xi \geq 0\}} \xi_n \, d\pi_i \uparrow \int_{\{\xi \geq 0\}} \xi \, d\pi_i \leq I_c[\pi_i] < \infty \text{ and} \tag{22}$$

$$\int_{\{\xi < 0\}} \xi_n \, d\pi_i \downarrow \int_{\{\xi < 0\}} \xi \, d\pi_i \tag{23}$$

for $i \in \{0, 1\}$, hence $\lim_{n\to\infty} \int \xi_n \, d\pi_i = \int \xi \, d\pi_i$.

b. By the assumption on equal marginals of π_0 and π_1 we obtain for $n \geq 0$

$$\int \xi_n \, d\pi_0 = \int \varphi_n \, d\pi_0 + \int \psi_n \, d\pi_0 \tag{24}$$

$$= \int \varphi_n \, d\pi_1 + \int \psi_n \, d\pi_1 = \int \xi_n \, d\pi_1. \tag{25}$$

Thus $I_c[\pi_0] = \int \xi \, d\pi_0 = \lim_{n\to\infty} \int \xi_n \, d\pi_0 = \lim_{n\to\infty} \int \xi_n \, d\pi_1 = \int \xi \, d\pi_1 \leq I_c[\pi_1]$; since π_1 was arbitrary, this implies optimality of π_0. $\square$

4 From c-monotonicity to optimality

This section is devoted to the proof of Theorem 1.b. Our argument starts with a finite c-monotone transport plan π and we aim for showing that π is at least as good as any other finite transport plan. The idea behind the proof is to partition X and Y into cells $C_i, i \in I$ resp. $D_i, i \in I$ in such a way that π is strongly c-monotone on "diagonal" sets of the form $C_i \times D_i$ while regions $C_i \times D_j, i \neq j$ can be ignored, because no finite transport plan will give positive measure to the set $C_i \times D_j$.

Thus it will be necessary to apply previously established results to some restricted transport problems on a space $C_i \times D_i$ equipped with some relativized transport plan $\pi \restriction C_i \times D_i$. As in general the cells C_i, D_i are plainly Borel sets they may fail to be Polish spaces with respect to the topologies inherited from X resp. Y. However, for us it is only important that there exist *some* Polish topologies that generate the same Borel sets on C_i resp. D_i (see e.g. [5, Theorem 13.1]). At this point it is crucial that our results only need measurability of the cost function and do not ask for any form of continuity (cf. the remarks at the end of the introduction). Before we give the proof of Theorem 1.b we will need some preliminary lemmas.

Lemma 4.1 *Let X, Y be Polish spaces equipped with Borel probability measures μ, ν and let $c : X \times Y \to [0, \infty]$ be a Borel measurable cost function. Let π, π_0 be finite transport plans and $\Gamma \subseteq X \times Y$ a Borel set with $\pi(\Gamma) = 1$ on which c is finite. Let $I = \{0, \ldots, n\}$ or $I = \mathbb{N}$ and assume that $C_i, i \in I$ are mutually disjoint Borel sets in X, $D_i, i \in I$ are mutually disjoint Borel sets in Y such that the equivalence classes of $\approx_{c,\Gamma}$ are of the form $\Gamma \cap (C_i \times D_i)$. Then also $\pi_0(\bigcup_{i\in I} C_i \times D_i) = 1$.*

In the proof we will need the following simple lemma. (For a proof see for instance [4, Proposition 8.13].)

Lemma 4.2 *Let $I = \{0, \ldots, n\}$ or $I = \mathbb{N}$ and let $P = (p_{ij})_{i,j\in I}$ be a matrix with non-negative entries such that $\sum_{j\in I} p_{i_0 j} = 1$ for each $i_0 \in I$. Assume that* *there exists a vector $(p_i)_{i\in I}$ with strictly positive entries such $p\cdot P = p$.*[5] *Then whenever $p_{i_0 i_1} > 0$ for $i_0, i_1 \in I$, there exists a finite sequence $i_0, i_1, \ldots, i_n = i_0$ such that for all $0 \le k < n$ one has $p_{i_k i_{k+1}} > 0$.*

PROOF of Lemma 4.1. As $\approx_{\Gamma,c}$ is an equivalence relation and π is concentrated on Γ, the sets C_i, $i \in I$ are a partition of X modulo μ-null sets. Likewise the sets D_i, $i \in I$ form a partition of Y modulo ν-null sets. In particular the quantities

$$p_i := \mu(C_i) = \nu(D_i) = \pi(C_i \times D_i), \quad i \in I \tag{26}$$

add up to 1. Without loss of generality we may assume that $p_i > 0$ for all $i \in I$. We define

$$p_{ij} := \frac{\pi_0(C_i \times D_j)}{\mu(C_i)}, \quad i, j \in I. \tag{27}$$

Then $\sum_{j\in I} p_{i_0 j} = \frac{\pi_0(C_i \times Y)}{\mu(C_i)} = 1$ for each $i_0 \in I$. By the condition on the marginals of π_0 we have for the i-th component of $p\cdot P$

$$(p\cdot P)_i = \sum_{j\in I} \mu(C_j) \frac{\pi_0(C_j \times D_i)}{\mu(C_j)} = \pi_0(X \times D_i) = \nu(D_i) = p_i$$

i.e. $p\cdot P = p$. Hence P satisfies the assumptions of Lemma 4.2. We claim that $p_{ii} = 1$ for all $i \in I$. Suppose not. Pick $i_0 \in I$ such that $p_{i_0 i_0} < 1$. Then there exists some index $i_1 \neq i_0$ such that $p_{i_0 i_1} > 0$. Pick a finite sequence $i_0, i_1, \ldots, i_n = i_0$ according to Lemma 4.2. Fix $k \in \{1, \ldots, n-1\}$. Then

$$\pi_0(C_{i_k} \times D_{i_{k+1}}) = p_{i_k i_{k+1}} > 0.$$

Since π_0 is a finite transport plan, there exist $x_k \in C_{i_k} \cap p_X[\Gamma]$ and $y'_{k+1} \in D_{i_{k+1}} \cap p_Y[\Gamma]$ such that $c(x_k, y'_{k+1}) < \infty$. Choose $y_k \in D_{i_k}$ and $x'_{k+1} \in C_{i_{k+1}}$ such that $(x_k, y_k), (x'_{k+1}, y'_{k+1}) \in \Gamma$. Then

$$(x_0, y_0) \lesssim (x'_1, y'_1) \approx (x_1, y_1) \lesssim (x'_2, y'_2) \approx (x_2, y_2) \lesssim \ldots \lesssim (x'_n, y'_n) \approx (x_0, y_0).$$

But this implies that $(x_0, y_0) \approx (x_1, y_1)$, contradicting the assumption that $(C_{i_0} \times D_{i_0}) \cap \Gamma, (C_{i_1} \times D_{i_1}) \cap \Gamma$ are different equivalence classes of $\approx_{\Gamma,c}$. Hence we have indeed $p_{ii} = 1$ for all $i \in I$, thus $\pi_0(C_i \times D_i) = \mu(C_i)$ which implies $\pi_0(\bigcup_{i\in I} C_i \times D_i) = 1$. $\square$

Lemma 4.3 *Let X, Y be Polish spaces equipped with Borel probability measures μ, ν and let $c : X \times Y \to [0, \infty]$ be a Borel measurable cost function which is $\mu \otimes \nu$-a.e. finite. For every finite transport plan π and every Borel set $\Gamma \subseteq X \times Y$ with $\pi(\Gamma) = 1$ on which c is finite, there exist Borel sets $O \subseteq X, U \subseteq Y$ such that $\Gamma' = \Gamma \cap (O \times U)$ has full π-measure and (Γ', c) is connecting.*

[5] Such a matrix P is often called a *stochastic matrix* while p is a *stochastic vector*.

PROOF. By Fubini's Theorem for μ-almost all $x \in X$ the set $\{y : c(x,y) < \infty\}$ has full ν-measure and for ν-almost all $y \in Y$ the set $\{x : c(x,y) < \infty\}$ has full μ-measure. In particular the set of points (x_0, y_0) such that both $\mu(\{x : c(x,y_0) < \infty\}) = 1$ and $\nu(\{y : c(x_0,y) < \infty\}) = 1$ has full π-measure. Fix such a pair $(x_0, y_0) \in \Gamma$ and let $O = \{x \in X : c(x,y_0) < \infty\}, U = \{y \in Y : c(x_0,y) < \infty\}$. Then $\Gamma' = \Gamma \cap (O \times U)$ has full π-measure and for every $(x,y) \in \Gamma'$ both quantities $c(x,y_0)$ and $c(x_0,y)$ are finite. Hence $x \approx_X x_0$, for every $x \in p_X[\Gamma']$. Similarly we obtain $y \approx_Y y_0$, for every $y \in p_Y[\Gamma']$. Hence (Γ', c) is connecting. $\square$

Finally we prove the statement of **Theorem 1.b**:

Let X, Y be Polish spaces equipped with Borel probability measures μ, ν and $c : X \times Y \to [0, \infty]$ a Borel measurable cost function. Every finite c-monotone transport plan is optimal if there exist a closed set F and a $\mu \otimes \nu$-null set N such that $\{(x,y) : c(x,y) = \infty\} = F \cup N$.

PROOF. Let π be a finite c-monotone transport plan and pick a c-monotone Borel set $\Gamma \subseteq X \times Y$ with $\pi(\Gamma) = 1$ on which c is finite.

Let $O_n, U_n, n \in \mathbb{N}$ be open sets such that $\bigcup_{n\in\mathbb{N}}(O_n \times U_n) = (X \times Y) \setminus F$. Fix $n \in \mathbb{N}$ and interpret $\pi \restriction O_n \times U_n$ as a transport plan on the spaces (O_n, μ_n) and (U_n, ν_n) where μ_n and ν_n are the marginals corresponding to $\pi \restriction O_n \times U_n$. Apply Lemma 4.3 to $\Gamma \cap (O_n \times U_n)$ and the cost function $c \restriction O_n \times U_n$ to find $O'_n \subseteq O_n, U'_n \subseteq U_n$ and $\Gamma_n = \Gamma \cap (O'_n \times U'_n)$ with $\pi(\Gamma_n) = \pi(\Gamma \cap (O_n \times U_n))$ such that (Γ_n, c) is connecting. Then $\tilde{\Gamma} = \bigcup_{n\in\mathbb{N}} \Gamma_n$ is a subset of Γ of full measure and every equivalence class of $\approx_{\tilde{\Gamma},c}$ can be written in the form $((\bigcup_{n\in N} O'_n) \times (\bigcup_{n\in N} U'_n)) \cap \Gamma$ for some non-empty index set $N \subseteq \mathbb{N}$. Thus there are at most countably many equivalence classes which we can write in the form $(C_i \times D_i) \cap \Gamma, i \in I$ where $I = \{1, \ldots, n\}$ or $I = \mathbb{N}$. Note that by shrinking the sets $C_i, D_i, i \in I$ we can assume that $C_i \cap C_j = D_i \cap D_j = \emptyset$ for $i \neq j$.

Assume now that we are given another finite transport plan π_0. Apply Lemma 4.1 to π, π_0 and $\tilde{\Gamma}$ to achieve that π_0 is concentrated on $\bigcup_{i\in I} C_i \times D_i$. For $i \in I$ we consider the restricted problem of transporting $\mu \restriction C_i$ to $\nu \restriction D_i$. We know that $\pi \restriction C_i \times D_i$ is optimal for this task by Propositions 1.5 and 3.2, hence $I_c[\pi] \leq I_c[\pi_0]$. $\square$

Remark 4.4 *In fact the following somewhat more general (but also more complicated to state) result holds true:* Assume that $\{(x,y) : c(x,y) = \infty\} \subseteq F \cup N$ where F is closed and N is a $\mu \otimes \nu$-null set. Then every c-monotone transport plan π with $\pi(F \cup N) = 0$ is optimal.

5 Completing the picture

First we give the proof of **Theorem 2**.

Let X, Y be Polish spaces equipped with Borel probability measures μ, ν and $c : X \times Y \to [0, \infty]$ a Borel measurable cost function. For a finite transport plan π the following assertions are equivalent:

a. π is robustly optimal.
b. π is strongly c-monotone.

PROOF. *a.* $\Rightarrow$ *b.*: Let Z and $\lambda \neq 0$ be according to the definition of robust optimality. As $\tilde{\pi} = (\mathrm{Id}_Z \times \mathrm{Id}_Z)_{\#}\lambda + \pi$ is optimal, Theorem 1.a ensures the existence of a $\tilde{c}$-monotone Borel set $\tilde{\Gamma} \subseteq (X \cup Z) \times (Y \cup Z)$ such that $\tilde{c}$ is finite on $\tilde{\Gamma}$ and $\tilde{\pi}$ is concentrated on $\tilde{\Gamma}$. Note that $(z, z) \in \tilde{\Gamma}$ for λ-a.e. $z \in Z$. We claim that for λ-a.e. $z \in Z$ and all $(x, y) \in \Gamma = \tilde{\Gamma} \cap (X \times Y)$ the relation

$$(x, y) \approx_{\tilde{\Gamma},\tilde{c}} (z, z) \tag{28}$$

holds true. Indeed, since $\tilde{c}$ is finite on $Z \times Y$ we have $c(z, y) < \infty$ hence $(x, y) \lesssim_{\tilde{\Gamma},\tilde{c}} (z, z)$. Analogously finiteness of $\tilde{c}$ on $X \times Z$ implies $c(x, z) < \infty$ such that also $(z, z) \lesssim_{\tilde{\Gamma},\tilde{c}} (x, y)$. By transitivity of $\approx_{\tilde{\Gamma},\tilde{c}}$, $(\tilde{\Gamma}, \tilde{c})$ is connecting. Applying Proposition 3.2 to the spaces $X \cup Z$ and $Y \cup Z$ we get that $\tilde{\pi}$ is strongly $\tilde{c}$-monotone, i.e. there exist $\tilde{\varphi}$ and $\tilde{\psi}$ such that $\tilde{\varphi}(a) + \tilde{\psi}(b) \leq \tilde{c}(a, b)$ for $(a, b) \in (X \cup Z) \times (Y \cup Z)$ and equality holds $\tilde{\pi}$-almost everywhere. By restricting $\tilde{\varphi}$ and $\tilde{\psi}$ to X resp. Y we see that π is strongly c-monotone.

b. $\Rightarrow$ *a.*: Let Z be a Polish space and let λ be a finite Borel measure on Z. We extend c to $\tilde{c} : (X \cup Z) \times (Y \cup Z) \to [0, \infty]$ via

$$\tilde{c}(a, b) = \begin{cases} c(a, b) & \text{for } (a, b) \in X \times Y \\ \max(\varphi(a), 0) & \text{for } (a, b) \in X \times Z \\ \max(\psi(b), 0) & \text{for } (a, b) \in Z \times Y \\ 0 & \text{otherwise.} \end{cases}$$

Define $\tilde{\varphi}(a) := \begin{cases} \varphi(a) & \text{for } a \in X \\ 0 & \text{for } a \in Z \end{cases}$ and $\tilde{\psi}(b) := \begin{cases} \psi(b) & \text{for } b \in Y \\ 0 & \text{for } b \in Z. \end{cases}$

Then $\tilde{\varphi}$ resp. $\tilde{\psi}$ are extensions of φ resp. ψ to $X \cup Z$ resp. $Y \cup Z$ which satisfy $\tilde{\varphi}(a) + \tilde{\psi}(b) \leq \tilde{c}(a, b)$ and equality holds on $\tilde{\Gamma} = \Gamma \cup \{(z, z) : z \in Z\}$. Hence $\tilde{\Gamma}$ is strongly $\tilde{c}$-monotone. Since $\tilde{\pi}$ is concentrated on $\tilde{\Gamma}$, $\tilde{\pi}$ is optimal by Proposition 1.5. □

Next consider **Theorem 3**.

Let X, Y be Polish spaces equipped with Borel probability measures μ, ν and let $c : X \times Y \to [0, \infty]$ be Borel measurable and $\mu \otimes \nu$-a.e. finite. For a finite transport plan π the following assertions are equivalent:

(1) π is optimal.
(2) π is c-monotone.
(3) π is robustly optimal.
(4) π is strongly c-monotone.

PROOF. By Theorem 2, (3) and (4) are equivalent and they trivially imply (1) and (2) which are equivalent by Theorem 1. It remains to see that (2) $\Rightarrow$ (4). Let π be a finite c-monotone transport plan. Pick a c-monotone Borel set $\Gamma \subseteq X \times Y$ such that c is finite on Γ and $\pi(\Gamma) = 1$. By Lemma 4.3 there exists a Borel set $\Gamma' \subseteq \Gamma$ such that $\pi(\Gamma') = 1$ and (Γ', c) is connecting, hence Proposition 3.2 applies. □

Finally the example below shows that the ($\mu \otimes \nu$-a.e.) finiteness of the cost function is essential to be able to pass from the "weak properties" (optimality, c-monotonicity) to the "strong properties" (robust optimality, strong c-monotonicity).

Example 5.1 *(Optimality does not imply strong c-monotonicity.) Let $X = Y = [0, 1]$ and equip both spaces with Lebesgue measure $\lambda = \mu = \nu$. Define c to be ∞ above the diagonal and $1 - \sqrt{x - y}$ for $y \leq x$. The optimal (in this case the only finite) transport plan is the Lebesgue measure π on the diagonal Δ. We claim that π is not strongly c-monotone. Striving for a contradiction we assume that there exist φ and ψ witnessing the strong c-monotonicity. Let Δ_1 be the full-measure subset of Δ on which $\varphi + \psi = c$, and write $p_X[\Delta_1]$ for the projection of Δ_1. We claim that*

$$\forall x, x' \in p_X[\Delta_1] : \text{ If } x < x', \text{ then } \varphi(x) - \varphi(x') \geq \sqrt{x' - x}, \tag{29}$$

which will yield a contradiction when combined with the fact that $p_X[\Delta_1]$ is dense.

Our claim (29) *follows directly from*

$$\varphi(x') + \psi(x) \leq c(x', x) = 1 - \sqrt{x' - x} \quad \text{and} \quad \varphi(x) + \psi(x) = c(x, x) = 1. \tag{30}$$

Now let $x < x+a$ be elements of $p_X[\Delta_1]$, let $b := \varphi(x)-\varphi(x')$, and let $n \in \mathbb{N}$ be a sufficiently large number, say satisfying $n > 2\frac{b^2}{a^2}$. Using the fact that $p_X[\Delta_1]$ is dense, we can find real numbers $x = x_0 < x_1 < \cdots < x_n = x+a$ in Δ_1 satisfying $x_k - x_{k-1} < 2/n$ for $k = 1, \ldots, n$.

Let $\varepsilon_k := x_k - x_{k-1}$ for $k = 1, \ldots, n$. Then we have $\varepsilon_k < \frac{2}{n} < \frac{a^2}{b^2}$ for all k, hence $\sqrt{\varepsilon_k} > \frac{b}{a}\varepsilon_k$. So we get

$$b = \varphi(x) - \varphi(x') = \sum_{k=1}^{n} \varphi(x_{k-1}) - \varphi(x_k) \geq \sum_{k=1}^{n} \sqrt{\varepsilon_k} > \sum_{k=1}^{n} \frac{b}{a}\varepsilon_k = \frac{b}{a}\sum_{k=1}^{n} \varepsilon_k = b,$$

a contradiction. (By letting $c = 0$ below the diagonal the argument could be simplified, but then we would lose lower semi-continuity of c.)

References

[1] L. Ambrosio, A. Pratelli, Existence and stability results in the L^1 theory of optimal transportation, in: Optimal transportation and applications (Martina Franca, 2001), vol. 1813 of Lecture Notes in Math., Springer, Berlin, 2003, pp. 123–160.

[2] C. Castaing, M. Valadier, Convex analysis and measurable multifunctions, Lecture Notes in Mathematics, Vol. 580, Springer-Verlag, Berlin, 1977.

[3] W. Gangbo, R. J. McCann, The geometry of optimal transportation, Acta Math. 177 (2) (1996) 113–161.

[4] O. Kallenberg, Foundations of modern probability, Probability and its Applications (New York), Springer-Verlag, New York, 1997.

[5] A. S. Kechris, Classical descriptive set theory, vol. 156 of Graduate Texts in Mathematics, Springer-Verlag, New York, 1995.

[6] H. G. Kellerer, Duality theorems for marginal problems, Z. Wahrsch. Verw. Gebiete 67 (4) (1984) 399–432.

[7] M. Knott, C. Smith, On Hoeffding-Fréchet bounds and cyclic monotone relations, J. Multivariate Anal. 40 (2) (1992) 328–334.

[8] A. Pratelli, On the sufficiency of c-cyclical monotonicity for optimality of transport plans, Math. Z. In press.

[9] S. T. Rachev, L. Rüschendorf, Mass transportation problems. Vol. I, Probability and its Applications (New York), Springer-Verlag, New York, 1998.

[10] D. Ramachandran, Perfect measures and related topics, in: Handbook of measure theory, Vol. I, II, North-Holland, Amsterdam, 2002, pp. 765–786.

[11] R. T. Rockafellar, Characterization of the subdifferentials of convex functions, Pacific J. Math. 17 (1966) 497–510.

[12] L. Rüschendorf, On c-optimal random variables, Statist. Probab. Lett. 27 (3) (1996) 267–270.

[13] W. Schachermayer, J. Teichmann, Characterization of optimal transport plans for the Monge-Kantorovich problem, Proceedings of the A.M.S. In press.

[14] C. Villani, Topics in optimal transportation, vol. 58 of Graduate Studies in Mathematics, American Mathematical Society, Providence, RI, 2003.

[15] C. Villani, Optimal Transport Old and New, Series: Grundlehren der mathematischen Wissenschaften, Vol 338, Springer Verlag, Berlin, 2009.