

ON A PROBLEM OF OPTIMAL TRANSPORT UNDER MARGINAL MARTINGALE CONSTRAINTS

MATHIAS BEIGLBÖCK AND NICOLAS JUILLET

ABSTRACT. The basic problem of optimal transportation consists in minimizing the expected costs $\mathbb{E}[c(X_1, X_2)]$ by varying the joint distribution (X_1, X_2) where the marginal distributions of the random variables X_1 and X_2 are fixed.

Inspired by recent applications in mathematical finance and connections with the peacock problem we study this problem under the additional condition that $(X_i)_{i=1,2}$ is a martingale, i.e. $\mathbb{E}[X_2|X_1] = X_1$.

We establish a *variational principle* for this problem which enables us to determine optimal martingale transport plans for specific cost functions. In particular we identify a martingale coupling that resembles the classic monotone quantile coupling in several respects. In analogy with the celebrated Theorem of Brenier the following behavior can be observed: if the initial distribution is continuous, then this “monotone martingale” is supported by the graphs of *two* functions $T_1, T_2 : \mathbb{R} \rightarrow \mathbb{R}$.

1. INTRODUCTION

1.1. Presentation of the martingale transport problem. We will denote by \mathcal{P} the set of probability measures on \mathbb{R} having finite first moments. We are given measures $\mu, \nu \in \mathcal{P}$, and a (measurable) *cost function* $c : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ which will be continuous in most of our applications. We assume moreover that $c(x, y) \geq a(x) + b(y)$ where a (resp. b) is integrable with respect to μ (resp. ν). Hence if (X, Y) is a joint law with marginal distributions law $X = \mu$ and law $Y = \nu$, the expectation of $c(X, Y) \geq a(X) + b(Y)$ is well defined, taking its value in $[\mathbb{E}[a(X)] + \mathbb{E}[b(Y)], +\infty]$. We will refer to this technical hypothesis as *the sufficient integrability condition*. The basic problem of optimal transport consists in the minimization problem

$$(1) \quad \text{Minimize } \mathbb{E}[c(X, Y)] \quad \text{subject to } \text{law}(X) = \mu, \text{law}(Y) = \nu.$$

where the infimum is taken over all joint distributions. We denote the infimum in (1) by $C(\mu, \nu)$. The joint laws on $\mathbb{R} \times \mathbb{R}$ are usually called *transport plans* after the classical concrete problem of Monge [22]: how can one transport a heap of soil distributed according to μ to a target distribution ν ? A transport plan π prescribes that for $(x, y) \in \mathbb{R}^2$ a quantity of mass $\pi(dx dy)$ is transported from x to y . Minimizers of the problem (1) are called optimal transport plans. Note that we will also use the more probabilistic term *coupling* for transport plans. Following [28] we denote the set of all transport plans by $\Pi(\mu, \nu)$ so that one has the alternative definition

$$C(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \iint c(x, y) d\pi(x, y).$$

2010 *Mathematics Subject Classification.* 60G42, 49N05.

Key words and phrases. optimal transport, convex order, martingales, model-independence, peacocks.

The first author gratefully acknowledges financial support from FWF under grant P21209. The second author is partially supported by the Agence Nationale de la Recherche, ANR-09-BLANC-0364-01 and ANR JCJC - SIMI 1 - ANR 2011 JS01 011 01.

Our main interest lies in a *martingale version* of the transport problem. That is, our aim is to minimize $\mathbb{E}[c(X, Y)]$ over the set of all *martingale transport plans*

$$\Pi_M(\mu, \nu) = \{\pi \in \Pi(\mu, \nu) : \pi = \text{law}(X, Y) \text{ and } \mathbb{E}[Y|X] = X\}.$$

A transport plan π is equivalently described through its disintegration $(\pi_x)_{x \in \mathbb{R}}$ with respect to the initial distribution μ . The probabilistic interpretation is that $(x, A) \mapsto \pi_x(A)$ is the transition kernel of the two-step process $(X_i)_{i=1,2}$ where $X_1 = X$ and $X_2 = Y$, i.e. $\pi_x(A) = \mathbb{P}(Y \in A | X = x)$. In these terms, π is an element of $\Pi_M(\mu, \nu)$, if and only if $\int y d\pi_x(y) = x$ holds μ -a.s. Hence in this paper we study the minimization problem

$$(2) \quad \text{Minimize } \mathbb{E}_\pi[c] = \iint c(x, y) d\pi(x, y) \quad \text{subject to } \pi \in \Pi_M(\mu, \nu)$$

for various costs. Let $C_M(\mu, \nu)$ denote the infimum $\inf\{\mathbb{E}_\pi[c] : \pi \in \Pi_M(\mu, \nu)\}$.

Our optimal transport approach permits to distinguish some special couplings of $\Pi_M(\mu, \nu)$ that are comparable to the monotone (or Hoeffding-Fréchet) coupling $\pi_{\text{HF}} \in \Pi(\mu, \nu)$. Indeed we have developed our martingale transport theory parallel to the classical theory and the optimizer of (2) will enjoy canonical properties. Nevertheless notable differences occur between the theories. An obvious one is the fact that $\Pi_M(\mu, \nu)$ can be empty while $\Pi(\mu, \nu)$ always contains the element $\mu \otimes \nu$. The existence of a martingale transport plan is actually quite an old topic that is present (but under different names) at least since the study of Muirhead's inequality by Hardy, Littlewood, and Pólya [11]. Several articles in different fields (analysis, combinatorics, potential theory, and probability) deal with this question in different settings, often for marginal distributions in spaces much more general than the real line (see e.g. [3, 26, 21, 5, 27, 8, 19, 9]). The interest in finding an explicit coupling has appeared recently in the peacock problem (see [12] and the references therein): a peacock is a stochastic process $(X_t)_{t \in I}$ such that there exists at least one martingale $(M_t)_{t \in I}$ satisfying $\text{law}(X_t) = \text{law}(M_t)$ for every t . The problem consists in building as explicitly as possible such a martingale (M_t) from (X_t) . The martingale transport problem is maybe even closer linked to the theory of model-independent pricing in mathematical finance.¹ Indeed the problem (2) has been first studied in this context by Hobson and Neuberger [16] for the specific cost function $c(x, y) = -|y - x|$. The link between optimal transport and model independent pricing has been made explicit in [2] in a discrete time framework and by Galichon, Henry-Labordere, and Touzi [10] in a continuous time setup.

We note that several of the basic features of the problem (2) are similar to the usual optimal transport problem. This appeals for instance to the weak compactness of $\Pi(\mu, \nu)$ and $\Pi_M(\mu, \nu)$. If c is lower semi-continuous, this carries over to the mapping $\pi \mapsto \mathbb{E}_\pi[c]$ for either space of transport plans. In particular the infimum is attained. Note also that as in the standard setup the problem has a natural dual formulation [2]. However as we already mentioned in the previous paragraph, while there is always a transport plan which moves μ to ν , the marginal distributions need to satisfy additional assumptions to guarantee that a *martingale* transport plan exists: the set $\Pi_M(\mu, \nu)$ is non-empty if and only if μ is smaller than ν in the *convex order* (see Definition 2.1). More details are provided in Section 2 along with a construction of a martingale transport plan between two given marginals.

¹We refer to the recent survey by Hobson [14] for a very readable introduction to this area. Arguably, the most important tool in model-independent finance is the Skorokhod-embedding approach; an extensive overview is given by Oblój in [23].

1.2. Summary on the classical transport problem on \mathbb{R} . A cornerstone in the modern theory of optimal transportation is Brenier's Theorem (or Brenier-Rachev-Rüschendorf Theorem), see [4, 24]. It treats the optimal transport problem in the particular case $c(x, y) = |y - x|^2$, where $|\cdot|$ denotes the Euclidean norm on \mathbb{R}^n . This is simply Problem (1) when μ and ν are interpreted as measures on \mathbb{R}^n . Under appropriate regularity conditions on μ , the optimal transport $\pi \in \Pi(\mu, \nu)$ is unique and supported by the graph of a function $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that is the gradient of some convex function. In particular the optimal transport is realized by a mapping. Note that in dimension one the gradient of a convex function is simply a monotonically increasing function so that the optimal coupling is the usual monotone coupling. This fact can be directly proved without too many difficulties (see for instance [17]) but nevertheless it is interesting as one of the rare cases where an optimal transport plan can be so easily understood. Moreover even without any assumption on μ , the monotone coupling is the unique optimal transport plan. In this paper we will see that similar results are valid in the martingale case; e.g., the uniqueness of the minimizer or the fact that the optimal coupling is concentrated on a special set comparable to the graph of a monotone mapping.

We present the classical (non martingale) optimal transport problem on the real line that will serve as a guideline to our paper. The results are given for an arbitrary strictly convex cost. Any cost of this type activates the same theory, which again is characteristic of dimension one.

Theorem 1.1. *Let μ, ν be probability measures and c a cost function defined by $c(x, y) = h(y - x)$, where $h : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly convex function. We assume that c satisfies the sufficient integrability condition with respect to μ and ν and that $C(\mu, \nu) < \infty$. The following statements are equivalent*

- (1) *The measure π is optimal.*
- (2) *The transport preserves the order, i.e. there is a set Γ with $\pi(\Gamma) = 1$ such that whenever $(x, y), (x', y') \in \Gamma$, if $x < x'$ one has also $y \leq y'$.*

We have the two following corollaries.

Corollary 1.2. *For given measures μ and ν , if $C(\mu, \nu)$ is finite then there exists a unique minimizer to the transport problem (1) and it is the monotone (Hoeffding-Fréchet-) coupling π_{HF} .*

One has in fact $\pi_{\text{HF}} = (G_\mu \otimes G_\nu)_{\#} \lambda_{[0,1]}$ where λ is the Lebesgue measure and G_μ and G_ν are the *quantile functions* of μ and ν , i.e. the non-decreasing and left-continuous functions obtained from the cumulative distribution functions F_μ and F_ν as a generalized inverse by the formula $G(s) = \inf\{t \in \mathbb{R} : s \leq F(t)\}$.² This observation is the reason why the coupling π_{HF} is also known under the alternative name *quantile coupling*.

For the following corollary we recall that a measure μ is said to be continuous if $\mu(\{x\}) = 0$ for every $x \in \mathbb{R}$.

Corollary 1.3. *Under the assumptions of Corollary 1.2, if μ is continuous then the optimal transport plan π_{HF} is concentrated on the graph of an increasing mapping $T : \mathbb{R} \rightarrow \mathbb{R}$. Moreover $T_{\#}\mu = \nu$.*

It is straightforward to see that $T = G_\nu \circ F_\mu$. This formula determines T , μ -a.s.

Quadratic costs in the martingale setting. While $c(x, y) = (y - x)^2$ is arguably the most important cost function in the theory of optimal transport we stress that it plays a rather different role in the martingale setup. Assume that $\text{law}(X) = \mu$ and

²Note that the function G may take infinite values at the boundary of its domain $[0, 1]$.

$\text{law}(Y) = \nu$ are linked by a martingale coupling π and possess second moments. Then

$$\mathbb{E}[XY] = \mathbb{E}[\mathbb{E}[XY|X]] = \mathbb{E}[X^2],$$

hence we have the Pythagorean relation

$$\int (y-x)^2 d\pi(x,y) = \mathbb{E}[(Y-X)^2] = \mathbb{E}[Y^2] - \mathbb{E}[X^2].$$

Thus the cost associated to π depends only on the marginal distributions, i.e. not on the particular choice of $\pi \in \Pi_M(\mu, \nu)$.

We record the following consequence: Let c be a cost function and assume that

$$\tilde{c}(x,y) = c(x,y) + p \cdot (y-x)^2 + q \cdot (y-x)$$

for some real constants p and q . Then in Problem (2) the minimizers are the same for the costs c and \tilde{c} . In particular, if $c(x,y) = h(y-x)$, we do not expect that monotonicity or convexity properties of the function h are relevant for the structure of the optimizer.

1.3. A new coupling: the monotone martingale coupling, main results.

In this section we will discuss a particular coupling which may be viewed as a martingale analogue to the monotone (Hoeffding-Fréchet) coupling. Notable similarities are that it is canonical with respect to the convex order as well as that it is optimal for a range of different cost functions.

Definition 1.4. *A martingale transport plan π on $\mathbb{R} \times \mathbb{R}$ is left-monotone or simply monotone if there exists a Borel set $\Gamma \subseteq \mathbb{R} \times \mathbb{R}$ with $\pi(\Gamma) = 1$ such that whenever $(x, y^-), (x, y^+), (x', y') \in \Gamma$ we cannot have (see Figure 1 where this situation is represented)*

$$(3) \quad x < x' \quad \text{and} \quad y^- < y' < y^+.$$

Respectively π is said to be right-monotone if there exists Γ such that if $(x, y^-), (x, y^+)$ and (x', y') are elements of Γ then we do not have

$$x > x' \quad \text{and} \quad y^- < y' < y^+.$$

We will refer to the set Γ as the monotonicity set of π .

In this paper we will only state the results for (left-)monotone couplings. The corresponding results for right-monotone couplings can be deduced easily. We illustrate the forbidden situation (3) in Figure 1. Note that the top line represents the measure μ while ν is distributed on the bottom line; this convention will also be used in the subsequent pictures.




FIGURE 1. The forbidden mapping.

The next theorem is proved in Section 5.

Theorem 1.5. *Let μ, ν be probability measures in convex order. Then there exists a unique (left-)monotone transport plan in $\Pi_M(\mu, \nu)$. We denote this coupling by π_{lc} and call it left-curtain³ coupling.*

Of course one does not expect that a martingale is concentrated on the graph of a deterministic mapping T ; this holds only in the trivial case when $\mu = \nu$ and $T(x) \equiv x$. Rather we have the following result.

Corollary 1.6. *Let μ, ν be probability measures in convex order and assume that μ is continuous. Then there exist a Borel set $S \subseteq \mathbb{R}$ and two measurable functions $T_1, T_2 : S \rightarrow \mathbb{R}$ such that*

- (1) π_{lc} is concentrated on the graphs of T_1 and T_2 .
- (2) For all $x \in \mathbb{R}$, $T_1(x) \leq x \leq T_2(x)$.
- (3) For all $x < x' \in \mathbb{R}$, $T_2(x) < T_2(x')$ and $T_1(x') \notin]T_1(x), T_2(x)[$.

The following picture (Figure 2) illustrates the coupling π_{lc} in a specific case. The measures μ and ν are Gaussian distributions having the same mean, the variance of ν being greater than the variance of μ . There exist two points at which the density of μ (w.r.t. Lebesgue measure) equals the density of ν . Denote the smaller of these points by x_0 . Then we have $T_1(x) = T_2(x) = x$ for $x < x_0$. For $x > x_0$, the map T_1 is strictly decreasing and T_2 is strictly increasing.




FIGURE 2. Scheme of the left-curtain π_{lc} coupling between two Gaussian measures.

The subsequent result states that the transport plan π_{lc} is optimal for a variety of different cost functions. (see Theorem 6.1 below.)

Theorem 1.7 (π_{lc} is optimal). *Let μ, ν be probability measures in convex order. Assume that $c(x, y) = h(y - x)$ for some differentiable function h whose derivative is strictly convex and that c satisfies the sufficient integrability condition. If $C_M(\mu, \nu) < \infty$ then π_{lc} is the unique optimizer.*

³This name is explained in some detail before Theorem 4.18.

Natural examples of cost functions to which the result applies are given by $c(x, y) = (y - x)^3$ and $c(x, y) = \exp(y - x)$.

We discuss a further characteristic property of the transport plan π_{lc} . For a real number t and $\pi \in \Pi(\mu, \nu)$ consider the measure

$$\nu_t^\pi := \text{proj}_{\#}^y \pi|_{]-\infty, t] \times \mathbb{R}}$$

where $\text{proj}^y : (a, b) \in \mathbb{R}^2 \mapsto b \in \mathbb{R}$. Loosely speaking, the mass $\mu|_{]-\infty, t]}$ is moved to ν_t^π by the transport plan π . It is intuitively clear (and not hard to verify) that a transport plan $\pi \in \Pi(\mu, \nu)$ is uniquely determined by the family $(\nu_t^\pi)_{t \in \mathbb{R}}$.

Using this notation, the classic monotone transport plan π_{HF} is characterized by the fact that for each t , the measure $\nu_t^{\pi_{\text{HF}}}$ is as *left* as possible. More precisely, for every t the measure $\nu_t^{\pi_{\text{HF}}}$ is minimal with respect to the first order stochastic dominance in the family

$$\{\nu_t^\pi : \pi \in \Pi(\mu, \nu)\}.$$

We have the following, analogous characterization for the monotone martingale coupling π_{lc} . This is in fact the way we will formally define π_{lc} in Theorem 4.18.

Theorem 1.8 (π_{lc} is canonical with respect to the convex order). *For every real number t the measure $\nu_t^{\pi_{\text{lc}}}$ is minimal with respect to the convex order (i.e. second order stochastic dominance) in the family*

$$\{\nu_t^\pi : \pi \in \Pi_M(\mu, \nu)\}.$$

The next theorem summarizes the properties of π_{lc} .

Theorem 1.9. *Let μ, ν be probability measures in convex order. Let $h : \mathbb{R} \rightarrow \mathbb{R}$ be a differentiable function such that h' is strictly convex and assume that the cost function $c : (x, y) \mapsto h(y - x)$ satisfies the sufficient integrability condition.*

We assume moreover $C_M(\mu, \nu) < +\infty$. Let π be a martingale coupling in $\Pi_M(\mu, \nu)$. The following statements are equivalent:

- *The coupling π is monotone.*
- *The coupling π is optimal.*
- *The coupling π is the left-curtain coupling π_{lc} : for every $(\pi', t) \in \Pi_M(\mu, \nu) \times \mathbb{R}$, the measure ν_t^π is smaller than $\nu_t^{\pi'}$ in the convex order.*

Note that Theorem 1.9 is a consequence of the other results stated above.

1.4. A “variational principle” for the martingale transport problem. An important basic tool in optimal transport is the notion of *c-cyclical monotonicity* (see [29, Chapter 4]) which links the optimality of transport plans to properties of the support of the transport plan. A parallel statement holds true in the present setup and plays a fundamental role in our considerations. Heuristically we expect that if $\pi \in \Pi_M(\mu, \nu)$ is optimal, then it will prescribe optimal movements for single particles. To make this precise we use the following notion.

Definition 1.10. *Let α be a measure on $\mathbb{R} \times \mathbb{R}$ with finite first moment in the second variable. We say that α' , a measure on the same space, is a competitor of α if α' has the same marginals as α and for $(\text{proj}_{\#}^x \alpha)$ -a.e. $x \in \mathbb{R}$*

$$\int y \, d\alpha_x(y) = \int y \, d\alpha'_x(y),$$

where $(\alpha_x)_{x \in \mathbb{R}}$ and $(\alpha'_x)_{x \in \mathbb{R}}$ are disintegrations of the measures with respect to $\text{proj}_{\#}^x \alpha$.

We can now formulate a “variational principle” for the martingale transport problem.

Lemma 1.11 (variational lemma). *Assume that μ, ν are probability measures in convex order and that $c : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a Borel measurable cost function satisfying the sufficient integrability condition. Assume that $\pi \in \Pi_M(\mu, \nu)$ is an optimal martingale transport plan which leads to finite costs. Then there exists a Borel set Γ with $\pi(\Gamma) = 1$ such that the following holds:*

If α is a measure on $\mathbb{R} \times \mathbb{R}$ with $|\text{spt}(\alpha)| < \infty$ and $\text{spt}(\alpha) \subseteq \Gamma$ then we have $\int c d\alpha \leq \int c d\alpha'$ for every competitor α' of α .

Indeed, under the additional assumption that the cost function c is continuous and bounded we can prove that the condition given in the variational lemma is not only necessary but also sufficient to guarantee that a measure is optimal, see Lemma 8.2 in Appendix A.

The variational lemma 1.11 is one of the key ingredients in our investigation of the monotone martingale transport plan π_{1c} introduced above. Moreover it turns out to be very useful if one seeks to derive results on the optimizers for various specific cost functions. Assuming for simplicity that μ is continuous, Lemma 1.11 allows us to derive the following results.

- (1) If $c(x, y) = (y - x)^4$, then $\text{card}(\text{spt } \pi_x) \leq 3$, $\mu(x)$ -a.s.
- (2) Assume that $c(x, y) = h(y - x)$ for some continuously differentiable function h and that the derivative h' intersects every affine function at most in $k \in \mathbb{N}$ points. Then $\text{card}(\text{spt } \pi_x) \leq k$, $\mu(x)$ -a.s. for the optimizing π . (See Theorem 7.1, and also Theorem 7.2 for a similar result which appeals to the classical transport problem.)
- (3) If $c(x, y) = -|y - x|$, then there is a unique optimizer $\pi \in \Pi_M(\mu, \nu)$. Moreover $\text{card}(\text{spt } \pi_x) \leq 2$, $\mu(x)$ -a.s. (This was first shown in [16], see Theorem 7.3.)
- (4) If $c(x, y) = |y - x|$, then there is a unique optimizer $\pi \in \Pi_M(\mu, \nu)$. Moreover $\text{card}(\text{spt } \pi_x) \leq 3$ and $\text{card}(\text{spt } \pi_x \setminus \{x\}) \leq 2$, $\mu(x)$ -a.s. (see Theorem 7.4)

Having financial applications in mind, the cost functions $c(x, y) = |y - x|$ and $c(x, y) = -|y - x|$ are particularly relevant, we refer to the work of Hobson and Neuberger [16].

1.5. Organisation of the paper. We will start with a warm up section (Section 2) in which we derive some basic properties and explain a procedure that allows to find a martingale coupling for two given measures in convex order. Then, in Section 3, we establish the variational lemma 1.11 which will play a crucial role throughout the paper. In Section 4 we introduce and study the shadow projection, which permits us to introduce the left-curtain transport plan π_{1c} . We define it in Theorem 4.18 through its canonical property with respect to the convex order, we explain the name “left-curtain” and prove that it is monotone in Theorem 4.21. The particular properties of the transport plan π_{1c} are established in Sections 5 and 6. In section 7 we present results related to other costs and other couplings. Finally, in the Appendix, we present a converse to the variational lemma 1.11. We also provide an alternative derivation of Lemma 1.11 which is longer than argument presented in Section 3 but has the advantage to be constructive and self-contained.

Acknowledgments. The authors wish to thank Michel Émery, Martin Goldstern, Claus Griessler, Martin Keller-Ressel, Vincent Vigon, and the participants of the Winter school 2012 in Regen for enlightening discussions on the topic of this paper. We are also indebted to a particularly careful referee for numerous valuable suggestions and for pointing out a mistake in the initial version of this manuscript.

2. CONSTRUCTION OF A MARTINGALE TRANSPORT PLAN FOR MEASURES.

In this section we extend the martingale optimal transport problem to general finite measures with finite first moment and we define the convex order on this space. We prove that there exists a martingale transport plan between two measures in convex order and give a very short description of the duality theory linked to our optimization problem.

2.1. Basic notions. Denote by \mathcal{M} the set of finite measures on \mathbb{R} having finite first moment. We consider it with the usual topology, i.e. we say that a sequence $(\nu_n)_n$ converges weakly in \mathcal{M} to an element $\nu \in \mathcal{M}$ if

- (1) $(\nu_n)_n$ converges weakly in the usual sense, i.e. using continuous bounded functions as test functions;
- (2) the sequence $\int |x| d\nu_n$ converges to $\int |x| d\nu$.

Note that this is the same as adding all functions that grow at most linearly in $\pm\infty$ to the set \mathcal{C}_b of continuous and bounded test functions.

The reason we are interested in the space \mathcal{M} is that we will need to consider also transport plans between measures $\mu, \nu \in \mathcal{M}$ which have (the same) mass k , where k is possibly different from 1. In direct generalization of the earlier definition the set of transport plans $\Pi(\mu, \nu)$ then consists of all Borel measures π on $\mathbb{R} \times \mathbb{R}$ satisfying $\text{proj}_{\#}^x \pi = \mu$, $\text{proj}_{\#}^y \pi = \nu$. As a consequence of Prohorov's Theorem the set $\Pi(\mu, \nu)$ is compact; see e.g. [29, Lemma 4.4] for details. If c is a continuous (or lower semi-continuous) cost function satisfying the sufficient integrability condition with respect to μ and ν then the cost functional

$$\pi \in \Pi(\mu, \nu) \mapsto \int c d\pi \in]-\infty, +\infty]$$

is lower semi-continuous w.r.t. the weak topology ([29, Lemma 4.3]). It follows that the infimum in the classic transport problem is attained.

We proceed analogously in the martingale setup. If μ and ν are not necessarily probabilities, we define $\Pi_M(\mu, \nu)$ to consist of all transport plans π such that the disintegration in probability measures $(\pi_x)_{x \in \mathbb{R}}$ w.r.t. μ satisfies

$$\int y d\pi_x(y) = x$$

for μ -almost every x . Then $\pi \in \Pi(\mu, \nu)$ is a martingale measure if and only if

$$(4) \quad \int \rho(x)(y - x) d\pi(x, y) = 0.$$

for all bounded measurable functions $\rho : \mathbb{R} \rightarrow \mathbb{R}$. To see whether π is a martingale measure it is of course enough to test (4) for a sufficiently rich class of functions: e.g. for all functions of the form $\rho = \mathbb{1}_{]-\infty, x]}$, $x \in \mathbb{R}$ or for all continuous bounded functions (see [2, Lemma 2.3]).

Hence the set $\Pi_M(\mu, \nu)$ is compact in the weak topology (see [2, Proposition 2.4]). Precisely as in the usual setup it follows that the value of the minimization problem (2) is attained provided that the set $\Pi_M(\mu, \nu)$ is non-empty.

Of course it is a fundamental question on which conditions martingale transport plans exist. In the usual optimal transport setup the problem is simple enough: the properly renormalized product measure $\frac{1}{\mu(\mathbb{R})} \mu \otimes \nu$ witnesses that $\Pi(\mu, \nu)$ is non-empty. As mentioned in the introduction, the proper notion which guarantees existence of a martingale transport plan is the convex order. As it plays a crucial role throughout the paper we will discuss it in some detail.

2.2. The convex order of measures. Let us start with the definition.

Definition 2.1. *Two measures μ and ν are said to be in convex order⁴ if*

- (1) *they have finite mass and finite first moments, i.e. lie in \mathcal{M} ,*
- (2) *for convex functions φ defined on \mathbb{R} , $\int \varphi d\mu \leq \int \varphi d\nu$.*

In that case we will write $\mu \preceq_C \nu$.

Note that if $\mu \preceq_C \nu$, then one can apply (2) to all affine functions. Using the particular choices $\varphi(x) \equiv 1$ and $\varphi(x) \equiv -1$ one obtains that μ and ν have the same total mass and considering the functions $\varphi(x) \equiv x$ and $\varphi(x) \equiv -x$ one finds that μ and ν have the same barycenter.⁵

It is useful to know that it is sufficient to test hypothesis (2) against suitable subclasses of the convex functions. For instance, measures μ, ν having the same finite mass and the same first moments are in convex order if and only if

$$\int (x - k)_+ d\mu(x) \leq \int (x - k)_+ d\nu(x)$$

for all real k . This follows from simple approximation arguments (see [13] and also Paragraph 4.1) using monotone convergence. In particular it is sufficient to check (2) for positive convex functions with finite asymptotic slope in $-\infty$ and $+\infty$.

We give some examples of measures in convex order.

Example 2.2. *If δ is an atom of mass $\alpha > 0$ at the point x , then $\delta \preceq_C \nu$ simply means that ν has mass α and barycenter x .*

Example 2.3. *If $\mu_i \preceq_C \nu_i$ for $i = 1, \dots, n$ then $\sum_{i=1}^n \mu_i \preceq_C \sum_{i=1}^n \nu_i$.*

Example 2.4. *If two measures μ and μ' have the same barycenter and the same mass, μ is concentrated on $[a, b]$ and μ' is concentrated on $\mathbb{R} \setminus]a, b[$ then $\mu \preceq_C \mu'$. Indeed it can be proved for convex functions φ defined on \mathbb{R} that*

$$\int \varphi d\mu \leq \int \psi d\mu = \int \psi d\mu' \leq \int \varphi d\mu'$$

where ψ is the linear function satisfying $\psi = \varphi$ in a and b .

Example 2.5. *If two measures μ and μ' have the same barycenter and the same mass, $\mu - (\mu \wedge \mu')$ is concentrated on $[a, b]$ and $\mu' - (\mu \wedge \mu')$ is concentrated on $\mathbb{R} \setminus]a, b[$ then we have $\mu \preceq_C \mu'$. To see this, apply Example 2.4 to the two reduced measures and note that adding $\mu \wedge \mu'$ preserves the order.*

The following result formally states the connection between the convex order and the existence of martingale transport plans.

Theorem 2.6. *Let $\mu, \nu \in \mathcal{M}$. The condition $\mu \preceq_C \nu$ is necessary and sufficient for the existence of a martingale transport plan in $\Pi_M(\mu, \nu)$.*

It is a simple consequence of Jensen's inequality that the condition $\mu \preceq_C \nu$ is necessary to have $\Pi_M(\mu, \nu) \neq \emptyset$: if π is a martingale transport plan and φ is convex then

$$\begin{aligned} \int \varphi(y) d\nu(y) &= \int \varphi(y) d\pi(x, y) = \\ &= \iint \varphi(y) d\pi_x(y) d\mu(x) \geq \int \varphi(x) d\mu(x). \end{aligned}$$

The fact that the condition is also sufficient is well known and goes back at least to a paper by Strassen [27]. Nevertheless we think that it is worthwhile to describe a procedure which allows to obtain a martingale transport plan. This is what we do in the next subsection.

⁴The convex order is also called Choquet order or second order stochastic dominance.

⁵The barycenter or mean of a measure μ is $\frac{1}{\mu(\mathbb{R})} \int x d\mu(x)$.

2.3. Construction of a martingale transport. We fix finite measures μ, ν having finite first moments and satisfying $\mu \preceq_C \nu$; our aim is to show that $\Pi_M(\mu, \nu)$ is non-empty. The desired result will first be given in the case where μ is concentrated on finitely many points. The construction in Proposition 2.7 will rely on the elementary fact (related to Example 2.3) that $\pi_1 \in \Pi_M(\mu_1, \nu_1)$, $\pi_2 \in \Pi_M(\mu_2, \nu_2)$ implies that $\pi_1 + \pi_2 \in \Pi_M(\mu_1 + \mu_2, \nu_1 + \nu_2)$.

Proposition 2.7. *Assume that $\mu = \sum_{i=1}^n \delta_i$, where each δ_i is an atomic measure. If ν satisfies $\mu \preceq_C \nu$, then $\Pi_M(\mu, \nu)$ is non-empty.*

First note that by Example 2.2 this proposition is clear if $n = 1$. The general case will be established by induction. To perform the inductive step we need to understand how to couple a single atom, say $\delta := \delta_1$, with a properly chosen portion ν' of ν so that the other atoms ($\sum_{i=2}^n \delta_i$) are smaller than $\nu - \nu'$ in convex order. Assume that δ has mass α and is concentrated on x . Recalling Example 2.2, we should pick ν' so that it has mass α and barycenter x . Clearly, it also needs to satisfy $\nu' \leq \nu$, where \leq refers to the usual pointwise order of measures.

As δ is a part of μ and $\mu \preceq_C \nu$, we can introduce the measure $\tilde{\mu} = \mu - \delta$ which has mass $t = \nu(\mathbb{R}) - \alpha$. Obviously we then have $\delta + \tilde{\mu} \preceq_C \nu$. We are looking for the measure ν' among the measures $\{\nu_s : s \in [0, t]\}$ obtained as the restriction of ν between two quantiles s and $s' = s + \alpha$. More precisely we consider $\nu_s = G_{\#} \lambda_{[s, s+\alpha]}$ where $G : [0, t + \alpha] \rightarrow \mathbb{R}$ is the quantile function of ν , and $\lambda_{[s, s']}$ is the Lebesgue measure restricted to $[s, s']$. In Paragraph 1.2 we have discussed quantile functions only for probability measures but of course the notion carries over to measures in \mathcal{M} . For completeness note that $\nu = G_{\#} \lambda_{[0, t+\alpha]}$.

The barycenter $B(s, \nu)$ of ν_s depends continuously on the parameter $s \in [0, t]$ and we claim that

$$(5) \quad B(0, \nu) \leq x, \quad B(t, \nu) \geq x.$$

This is a consequence of the convex order relation $(\delta + \tilde{\mu}) \preceq_C \nu$ applied to the convex and non-negative functions $u \mapsto (u - G(\alpha))_-$ and $u \mapsto (u - G(t))_+$. For instance

$$\int u - G(t) d\delta(u) \leq \int (u - G(t))_+ d\delta(u) \leq \int (u - G(t))_+ d\nu(u) = \int u - G(t) d\nu_t(u).$$

By the intermediate value theorem, the continuity of $s \mapsto B(s, \nu)$ implies that there exists some $s \in [0, t]$ such that ν_s has barycenter x . Moreover if $B(s, \nu) = B(s', \nu)$, the measures ν_s and $\nu_{s'}$ are equal so that there exists a unique measure with barycenter x . We denote it by ν' .

This discussion leads us to the following lemma.

Lemma 2.8. *Let μ be of the form $\mu = \tilde{\mu} + \delta$, where δ is an atom and assume that $\mu \preceq_C \nu$. Then there exists a unique splitting of the measure ν into two positive measures ν' and $\tilde{\nu} = \nu - \nu'$ in such a way that*

- (1) $\delta \preceq_C \nu'$,
- (2) $\tilde{\nu}(I) = 0$ where $I = \text{conv}(\text{spt}(\nu'))$ is the interior of the smallest interval containing the support of ν' .

Moreover the measures $\tilde{\mu}$ and $\tilde{\nu}$ satisfy $\tilde{\mu} \preceq_C \tilde{\nu}$.

Proof. Having already constructed ν' (and I , i.e. $]G(s), G(s + \alpha)[$) in the paragraph above Lemma 2.8 it remains to show (2): $\tilde{\mu}$ is smaller than $\tilde{\nu}$ in the convex order. Let φ be a non-negative convex function which satisfies

$$\limsup_{|x| \rightarrow +\infty} |\varphi(x)/x| < +\infty.$$

We will prove that $\int \varphi d\tilde{\mu} \leq \int \varphi d\tilde{\nu}$. To this end we introduce a new function ψ which equals φ on $\mathbb{R} \setminus I$ and is linear on I . The function ψ can be chosen to be convex and satisfy $\psi \geq \varphi$. (Note that this is possible also in the case where I is unbounded.) The functions φ and ψ coincide on the border of I . We have

$$\int \varphi d\tilde{\mu} \leq \int \psi d\tilde{\mu} = \int \psi d\mu - \int \psi d\delta.$$

But as ψ is linear on I , one has $\int \psi d\delta = \int \psi d\nu'$ and because $\mu \preceq_C \nu$ one has $\int \psi d\mu \leq \int \psi d\nu$. It follows that

$$\int \varphi d\tilde{\mu} \leq \int \psi d\nu - \int \psi d\nu' = \int \psi d\tilde{\nu} = \int \varphi d\tilde{\nu}.$$

The last equality is due to the fact that $\tilde{\nu}$ is concentrated on $\mathbb{R} \setminus I$. We have thus established our claim that $\tilde{\mu} \preceq_C \tilde{\nu}$. \square

Proof of Proposition 2.7. In the first step we apply Lemma 2.8 to the measures $\delta = \delta_1$ and $\tilde{\mu} = \sum_{i=2}^n \delta_i$ to obtain a splitting $\nu = \hat{\nu}_1 + \tilde{\nu}$ that satisfies $\delta_1 \preceq_C \hat{\nu}_1$ and $\tilde{\mu} \preceq_C \tilde{\nu}$. Trivially $\Pi_M(\delta_1, \hat{\nu}_1)$ consists of a single element π_1 .

In the next step we repeat the procedure with $\tilde{\mu}$ and $\tilde{\nu}$ in the place of μ, ν and continue until the n -th step where δ_n can be martingale-transported to the remaining part of ν because the convex order relation $\delta_n \preceq_C (\nu - \sum_{i=1}^{n-1} \hat{\nu}_i)$ is satisfied in Example 2.2. Hence we have obtained recursively a sequence $(\hat{\nu}_i)_{i=1}^n$ such that $\delta_i \preceq_C \hat{\nu}_i$ and $\hat{\nu}_1 + \dots + \hat{\nu}_n = \nu$. We have constructed n martingale transport plans π_1, \dots, π_n where π_i is the unique element of $\Pi_M(\delta_i, \hat{\nu}_i)$. Thus $\pi_1 + \dots + \pi_n$ is an element of $\Pi_M(\mu, \nu)$. \square

To extend Proposition 2.7 to the case of general $\mu \in \mathcal{M}$ we need the following simple and straight-forward fact that will also be useful in Section 4.

Lemma 2.9 (Approximation of a measure in the convex order). *Assume $\gamma \in \mathcal{M}$. There exists a sequence $(\gamma^{(n)})_n$ of finitely supported measures such that $\gamma^{(n+1)} \succeq_C \gamma^{(n)}$, the sequence $(\gamma^{(n)})_n$ converges weakly to γ in \mathcal{M} and $\gamma^{(n)} \preceq_C \gamma$ holds for every n .*

Proof. To any partition \mathcal{J} of \mathbb{R} into finitely many intervals we can associate some $\gamma_{\mathcal{J}}$ smaller than γ in the convex order. We simply replace $\gamma = \sum_{I \in \mathcal{J}} \gamma|_I$ by $\gamma_{\mathcal{J}} = \sum_{I \in \mathcal{J}} \delta_I$ where δ_I is an atom with the same mass and same barycenter as $\gamma|_I$. Note that if \mathcal{J}' is finer than \mathcal{J} (the intervals of \mathcal{J} are broken in sub-intervals) then $\gamma_{\mathcal{J}} \preceq_C \gamma_{\mathcal{J}'}$. For $k, N \in \mathbb{N}$ we consider the partition

$$\mathcal{J}_{k,N} = \left(\bigcup_{i=-2^k N}^{(2^k-1)N} \left] \frac{i}{2^k}, \frac{i+1}{2^k} \right] \right) \cup]N, +\infty[\cup]-\infty, -N],$$

and set $\gamma_{k,N} = \gamma_{\mathcal{J}_{k,N}}$. We have $\gamma_{k,N} \preceq_C \gamma_{k+1,N}$ and $\gamma_{k,N} \preceq_C \gamma_{k,N+1}$. Write $\gamma^{(n)}$ for $\gamma_{n,n}$. Let f be a continuous function that grows less than linearly in $\pm\infty$. There exist $a, b > 0$ such that $|f(x)| \leq a|x| + b$. Let $\varepsilon > 0$ and N be such that $\int_{|x| \geq N} a|x| + b d\gamma(x) \leq \varepsilon/3$. The function f is uniformly continuous on $[-N, N]$. Thus there exists ω such that if $x, y \in [-N, N]$ and $|x - y| \leq \omega$ we have $|f(x) - f(y)| \leq \varepsilon/3$. Let k be such that $1/2^k \leq \omega$. For $n \geq \max\{k, N\}$ we have

$$\begin{aligned} |\gamma(f) - \gamma^{(n)}(f)| &\leq \left| \int_{-N}^N f d\gamma - \int_{-N}^N f d\gamma^{(n)} \right| \\ &\quad + \left| \int_{|x| \geq N} f d\gamma \right| + \left| \int_{|x| \geq N} f d\gamma^{(n)} \right| \leq \frac{\varepsilon}{3} + \frac{\varepsilon}{3} + \frac{\varepsilon}{3}. \end{aligned}$$

The first two estimates are a consequence of our preparations: to see this, note that

$$\left| \int_{|x| \geq N} f d\gamma^{(n)} \right| \leq \int_{|x| \geq N} a|x| + b d\gamma^{(n)} \leq \int_{|x| \geq N} a|x| + b d\gamma$$

where the convexity of $x \mapsto a|x| + b$ and $\gamma^{(n)}|_{\{|x| \geq N\}} \preceq_C \gamma|_{\{|x| \geq N\}}$ are used. \square

We are now finally in the position to conclude the proof of Theorem 2.6.

Proof of sufficiency in Theorem 2.6. Pick a sequence of finitely supported measures $(\mu_n)_{n \geq 1}$ satisfying $\mu_n \preceq_C \nu$ such that μ_n converges to μ weakly. (By Lemma 2.9, the sequence could be chosen to be increasing in the convex order, but we do not need this here.) We have already solved the problem of transporting a discrete distribution. Pick martingale measures $(\pi_n)_{n \geq 1}$ which transport μ_n to ν for each n . To be able to pass to a limit, we note that the set

$$\Omega := \Pi_M(\mu, \nu) \cup \bigcup_{n=1}^{\infty} \Pi_M(\mu_n, \nu)$$

is compact. Hence the sequence $(\pi_n)_{n \geq 1}$ has an accumulation point π in Ω and of course π is as desired: its marginals are μ and ν and it is a martingale transport plan. \square

We have thus seen a self-contained proof to Theorem 2.6. Of course the reader may object that the martingale established in the course of the proof was in no sense canonical and that the derivation was not constructive since we have invoked a compactness argument to prove the existence in the case of a general measure μ . In Section 4 we will be concerned with a modification of the above ideas which does not suffer from these shortfalls.

2.4. A dual problem. We mention that the martingale transport problem (2) admits a dual formulation. In analogy to the dual part of the optimal transport problem one may consider

$$\text{Maximize } \int \varphi d\mu + \int \psi d\nu$$

where one maximizes over all functions $\varphi \in L^1(\mu), \psi \in L^1(\nu)$ such that there exists $\Delta \in \mathcal{C}_b(\mathbb{R})$ satisfying

$$(6) \quad c(x, y) \geq \varphi(x) + \psi(y) + \Delta(x)(y - x)$$

for all $x, y \in \mathbb{R}$. Denote the corresponding supremal value by D . The inequality $D \leq C_M(\mu, \nu)$ then follows by integrating (6) against $\pi \in \Pi_M(\mu, \nu)$. In the case of lower semi-continuous costs c the duality relation $D = C_M(\mu, \nu)$ is established in [2, Theorem 1.1]. We also note that the dual part of the problem appears naturally in mathematical finance where it has a canonical interpretation in terms of *replication*. We refer to [2] for more details on this topic.

Duality results for a continuous time martingale transport problem are obtained by Galichon, Henry-Labordere, Touzi [10] and Dolinsky, Soner [7].

3. A SHORT PROOF OF THE VARIATIONAL LEMMA

The aim of this section is to establish the *variational lemma*, Lemma 1.11. That is, for a given optimal martingale transport plan π we want to construct a Borel set Γ , $\pi(\Gamma) = 1$ such that the following holds: if α is a measure on $\mathbb{R} \times \mathbb{R}$ with $|\text{spt}(\alpha)| < \infty$ and $\text{spt}(\alpha) \subseteq \Gamma$ then we have $\int c d\alpha \leq \int c d\alpha'$ for every competitor α' of α .

As mentioned above, this result can be viewed as a substitute for the characterization of optimality through the notion of c -cyclical monotonicity in the classical setup. Under mild regularity assumptions it is not too hard to show that a transport plan π which is optimal for the (usual) transport problem is c -cyclically monotone, we refer to [29, Theorem 5.10]. However this approach does not translate effortlessly to the martingale case. Roughly speaking the main problem in the present setup is

that the martingale condition makes manipulation of transport plans a relatively delicate issue.

Instead we give here a proof of Lemma 1.11 that is based on certain measure theoretic tools: it requires a general duality theorem of Kellerer [20, Lemma 1.8(a), Corollary 2.18] which in turn requires Choquet's capability theorem [6].⁶ See the Appendix for an alternative and constructive proof of the variational lemma.

The crucial ingredient is the following result:

Theorem 3.1. *Let (Z, ζ) be a Polish probability space and $M \subseteq Z^n$. Then either of the following holds true:*

- (1) *There exist subsets $(M_i)_i$ of Z^n such that $\zeta(\text{proj}^i M_i) = 0$ for $i = 1, \dots, n$ and*

$$M \subseteq \bigcup_{i=1}^n M_i.$$

- (2) *There exists a measure γ on Z^n such that $\gamma(M) > 0$ and $\text{proj}_{\#}^i \gamma \leq \zeta$ for $i = 1, \dots, n$.*

We refer to [1, Proposition 2.1] for a detailed proof of Theorem 3.1 from Kellerer's result.

Proof of Lemma 1.11. Fix a number $n \in \mathbb{N}$. We want to construct a set Γ_n for which the optimality property holds for all α satisfying $|\text{spt } \alpha| \leq n$. This set Γ_n will satisfy $\pi(\Gamma_n) = 1$. Clearly $\Gamma = \bigcap_{n \in \mathbb{N}} \Gamma_n$ is then as required to establish the lemma.

For a fixed $n \in \mathbb{N}$, define a Borel set M by

$$M := \left\{ (x_i, y_i)_{i=1}^n : \exists \alpha \text{ s.t. } \begin{array}{l} (1) \alpha \text{ is a measure on } \mathbb{R} \times \mathbb{R}, \\ (2) \text{spt } \alpha \subseteq \{(x_i, y_i) : i = 1, \dots, n\}, \text{ and} \\ (3) \exists \text{ competitor } \alpha' \text{ satisfying } \int c \, d\alpha' < \int c \, d\alpha. \end{array} \right\}.$$

We then apply Theorem 3.1 to the space $(Z, \zeta) = (\mathbb{R}^2, \pi)$ and the set M .

If we are in case (1), let N be $\bigcup_{i=1}^n \text{proj}^i(M_i)$ so that $\pi(N) = 0$ and $M \subseteq (N \times Z^{n-1}) \cup \dots \cup (Z^{n-1} \times N) = Z^n \setminus (Z \setminus N)^n$. We can then simply define $\Gamma_n := Z \setminus N = \mathbb{R}^2 \setminus N$ to obtain a set which does not support any non-optimal α with $|\text{spt } \alpha| \leq n$. Moreover $\pi(\Gamma_n) = 1$ as we want, hence the proof is complete.

It remains to show that case (2) cannot occur. Striving for a contradiction we assume that there is a measure γ such that $\gamma(M) > 0$ and $\text{proj}_{\#}^i \gamma \leq \pi$ for $i = 1, \dots, n$. Restricting γ to M , we may of course assume that $\gamma(\mathbb{R} \times \mathbb{R} \setminus M) = 0$. Rescaling γ if necessary, we may also assume that $\text{proj}_{\#}^i \gamma \leq \frac{1}{n} \pi$.

Consider the measure $\omega = \sum_{i=1}^n \text{proj}_{\#}^i \gamma$ on \mathbb{R}^2 . It is smaller than π and has positive mass. In particular $\mu_\omega = \text{proj}_{\#}^x \omega \leq \mu$. We will find a competitor ω' (recall Definition 1.10) such that ω' leads to smaller costs than ω , i.e.

$$\int c(x, y) \, d\omega' < \int c(x, y) \, d\omega.$$

If such a measure ω' exists then the measure $\pi - \omega + \omega'$ is a martingale transport plan which leads to smaller costs than π , contradicting the optimality of π . It remains to explain how ω' is obtained. For each $p = ((x_1, y_1), \dots, (x_n, y_n)) \in (\mathbb{R} \times \mathbb{R})^n$ let α_p be the measure which is uniformly distributed on the set $\{(x_1, y_1), \dots, (x_n, y_n)\}$. Then

$$\omega = \int_{p \in (\mathbb{R} \times \mathbb{R})^n} \alpha_p \, d\gamma(p).$$

⁶This approach is inspired by [1] where c -cyclical monotonicity is linked to optimality with the help of Kellerer's result.

For each $p \in (\mathbb{R} \times \mathbb{R})^n$ let α'_p be an optimizer of the problem

$$\text{Minimize } \int_{(x,y) \in \mathbb{R} \times \mathbb{R}} c(x,y) d\beta(x,y), \quad \beta \text{ competitor of } \alpha_p.$$

We emphasize that α'_p exists and can be taken to depend measurably on p . This follows for instance by calculating α'_p using the simplex algorithm.⁷

As γ is concentrated on M , for γ -almost all points p the measure α'_p satisfies

$$\int_{(x,y) \in \mathbb{R} \times \mathbb{R}} c(x,y) d\alpha'_p(x,y) < \int_{(x,y) \in \mathbb{R} \times \mathbb{R}} c(x,y) d\alpha_p(x,y).$$

(Note that α'_p is in general not concentrated on the same set as α_p .) Then ω' defined by

$$\omega' = \int_{p \in (\mathbb{R} \times \mathbb{R})^n} \alpha'_p d\gamma(p)$$

satisfies the above conditions as required. For instance we have

$$\begin{aligned} \int_{\mathbb{R} \times \mathbb{R}} c d\omega' &= \int_{p \in (\mathbb{R} \times \mathbb{R})^n} \int_{(x,y) \in \mathbb{R} \times \mathbb{R}} c(x,y) d\alpha'_p(x,y) d\gamma(p) \\ &< \int_{p \in (\mathbb{R} \times \mathbb{R})^n} \int_{(x,y) \in \mathbb{R} \times \mathbb{R}} c(x,y) d\alpha_p(x,y) d\gamma(p) = \int_{\mathbb{R} \times \mathbb{R}} c d\omega. \end{aligned}$$

The other properties are checked analogously. \square

We note that the just given proof of Lemma 1.11 is likely to extend to more general setups. In particular we expect that the result remains valid if martingale transport plans between higher dimensional spaces and with a finite number of time steps (i.e. $(X_i)_{i=1}^n$ rather than just $X_1 = X$ and $X_2 = Y$) are considered.

Subsequently Lemma 1.11 will several times be applied in conjunction with the following technical assertion. Given $\Gamma \subseteq \mathbb{R}^2$ we will use the notation Γ_x for $\{y \in \mathbb{R} : (x,y) \in \Gamma\}$.

Lemma 3.2. *Let k be a positive integer and $\Gamma \subseteq \mathbb{R}^2$. Assume also that there are uncountably many $a \in \mathbb{R}$ satisfying $|\Gamma_a| \geq k$.*

There exist a and $b_1 < \dots < b_k \in \Gamma_a$ such that for every $\varepsilon > 0$ one may find $a' > a$ and $b'_1 < \dots < b'_k \in \Gamma_{a'}$ with

$$\max(|a - a'|, |b_1 - b'_1|, \dots, |b_k - b'_k|) < \varepsilon.$$

Moreover one may also find $a'' < a$ and $b''_1 < \dots < b''_k \in \Gamma_{a''}$ with

$$\max(|a - a''|, |b_1 - b''_1|, \dots, |b_k - b''_k|) < \varepsilon.$$

Proof. Write A for the set of all a such that $|\Gamma_a| \geq k$ and pick for each $a \in A$ distinct elements $b_1^a, \dots, b_k^a \in \Gamma_a$. Set $\Gamma_A = \{(a, b_1^a, \dots, b_k^a) : a \in A\}$. We call $(a, b_1^a, \dots, b_k^a) \in \Gamma_A$ a *right-accumulation* point if for every $\varepsilon > 0$ there exists $a' \in]a, a + \varepsilon[$ such that $|b_i^a - b_i^{a'}| < \varepsilon$ for every i . We call it *right-isolated* otherwise. If p belongs to the set of right-isolated points $I_r \subseteq \Gamma_A$ then there exists some $\varepsilon_p > 0$ such that

$$[\{p\} + (]0, \varepsilon_p[\times] - \varepsilon_p, \varepsilon_p[{}^k)] \cap \Gamma_A = \emptyset,$$

where $+$ refers to the Minkowski sum of sets.

Assume for contradiction that the set I_r is uncountable. Then there exists some $\zeta > 0$ such that $K = \{p \in I_r : \varepsilon_p > \zeta\}$ is uncountable. Given $p_1, p_2 \in K$, we have

⁷It is well known that the optimal transport problem for finite spaces falls into the realm of linear programming, see for instance [28, p 23]. The same holds true in the martingale case.

$p_2 \notin p_1 + ((0, \zeta) \times (-\zeta, \zeta)^k)$. Since p_1 and p_2 have different first coordinates, this implies

$$[\{p_1\} + (]0, \zeta/2[\times] - \zeta/2, \zeta/2[^k) \cap [\{p_2\} + (]0, \zeta/2[\times] - \zeta/2, \zeta/2[^k) = \emptyset.$$

This is a contradiction since there cannot be uncountably many disjoint open sets in \mathbb{R}^{k+1} .

It follows that all but countably many elements of A are right-accumulation points. Arguing the same way with left replacing right we obtain the desired conclusion. \square

4. EXISTENCE OF A MONOTONE MARTINGALE TRANSPORT PLAN: THE LEFT-CURTAIN TRANSPORT PLAN

A short way to prove that there exists some monotone martingale transport plan would be to take a minimizer of Problem (2) for $c(x, y) = h(y - x)$ where h is chosen appropriately. Then one may apply Lemma 1.11 to prove that this minimizer is monotone. This kind of argument will be encountered in Sections 6 and 7 below. Here however we find it useful to give a construction which yields more insight in the structure of the martingale transport plan. In particular it will also allow us to prove the uniqueness of a monotone martingale transport plan in Section 5 and it will not require any assumptions on μ and ν .

For our argument, we reconsider the construction used in Proposition 2.7 and decide to transport the atoms δ_i of $\mu = \sum_i \delta_i$ to ν in a particular order, starting with the left-most atom and continuing to the right. It turns out that one can characterize the martingale coupling that we obtain in terms of an *extended convex order* and *shadow* introduced below (see Definition 4.3 and Lemma 4.6). These notions enable us to adapt the construction directly to the continuous case, thus making the approximation procedure used in Paragraph 2.3 obsolete.

4.1. Potential functions. An important tool in this section will be the so called *potential functions*. For each $\mu \in \mathcal{M}$ we define the potential function $u_\mu : \mathbb{R} \rightarrow \mathbb{R}$ by

$$u_\mu(x) = \int_{-\infty}^{\infty} |y - x| d\mu(y)$$

for $x \in \mathbb{R}$. Set $k = \mu(\mathbb{R})$ and $m = \frac{1}{k} \int x d\mu$.

Proposition 4.1. *If μ is in \mathcal{M} and $k = \mu(\mathbb{R})$, $m = \frac{1}{k} \int x d\mu$, then u_μ has the following properties:*

- (i) u_μ is convex,
- (ii) $\lim_{x \rightarrow -\infty} u_\mu(x) - k|x - m| = 0$ and $\lim_{x \rightarrow +\infty} u_\mu(x) - k|x - m| = 0$.

Conversely, if f is a function satisfying these properties for some numbers $m \in \mathbb{R}$ and $k \in [0, +\infty[$, then there exists a unique measure $\mu \in \mathcal{M}$ such that $f = u_\mu$. The measure μ is one half the second derivative f'' in the sense of distributions.

Proof. See for instance the proof of Proposition 2.1 in [13]. \square

Let us list some relevant properties of potential functions.

Proposition 4.2. *Let μ and ν be in \mathcal{M} .*

- *If μ and ν have the same mass, $\mu \preceq_C \nu$ is equivalent to $u_\mu \leq u_\nu$.*
- *We have $\mu \leq \nu$ if and only if u_μ has smaller curvature than u_ν . More precisely $\mu \leq \nu$ if and only if $u_\nu - u_\mu$ is convex.*
- *A sequence of measures $(\mu_n)_n$ in \mathcal{M} with mass k and mean m converges weakly in \mathcal{M} to some μ if and only if $(u_{\mu_n})_n$ converges pointwise to the potential function of some $\mu' \in \mathcal{M}$. In that case $\mu = \mu'$.*

Proof. For the first property see [12, Exercise 1.7], for the third [13, Proposition 2.3]. The second property is a consequence Proposition 4.1. Namely 2μ and 2ν are the second derivatives of u_μ and u_ν . \square

We will need the following generalization of the convex order.

Definition 4.3 (Extended convex order on \mathcal{M}). *Let μ and ν be measures in \mathcal{M} . We write $\mu \preceq_E \nu$ and say that ν is greater than μ in the extended convex order if for any non-negative convex function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ we have*

$$\int \varphi d\mu \leq \int \varphi d\nu.$$

The partial order \preceq_C on \mathcal{M} is extended by the order \preceq_E in the sense that \preceq_E keeps the old relations and gives rise to new ones. By definition, if $\mu \preceq_C \nu$ then we have $\mu \preceq_E \nu$ (since non-negative convex functions are convex). But if $\mu \preceq_E \nu$, we will also have $\mu \preceq_C \nu$ (as non-negative convex functions are non-negative). Note that in this second case the two measures may have neither the same mass nor the same barycenter.

As $x \mapsto 1$ is a convex function, a trivial consequence of $\mu \preceq_E \nu$ is $\mu(\mathbb{R}) \leq \nu(\mathbb{R})$. More precisely, let us prove that if the two measures have the same mass, $\mu \preceq_E \nu$ is equivalent to $\mu \preceq_C \nu$. Indeed if $\mu \preceq_E \nu$, for a convex function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ and any (negative) constant y , the convex function $\varphi_y : x \mapsto \varphi(x) \vee y$ satisfies $\int \varphi_y d\mu \leq \int \varphi_y d\nu$ because $\int \varphi_y - y d\mu \leq \int \varphi_y - y d\nu$. Letting y go to $-\infty$ we obtain $\int \varphi d\mu \leq \int \varphi d\nu$. Hence $\mu \preceq_C \nu$.

In terms of \preceq_C , the extended convex order can be characterized as follows:

Proposition 4.4. *Assume that $\mu \preceq_E \nu$. Then there exists a measure $\theta \leq \nu$ such that $\mu \preceq_C \theta$.*

Of course the converse statement is true as well: if there exists θ such that $\mu \preceq_C \theta$ and $\theta \leq \nu$ then we have also $\mu \preceq_E \nu$.

Proof. Let μ and ν satisfy $\mu \preceq_E \nu$. We can assume that ν is a probability measure and denote by k and m the mass resp. the mean of μ . We define a measure $\theta \leq \nu$ as follows. Consider the quantile function G_ν of ν . Recall that λ is the Lebesgue measure on \mathbb{R} . For a parameter $\zeta \in [0, k]$, we denote by λ^ζ the restriction of λ to $[0, 1] \setminus [\zeta, \zeta + (1 - k)]$. This measure has mass k as well as does $\theta = (G_\nu)_\# \lambda^\zeta$. We now pick ζ such that θ has mean m . To see that this can be done we will apply the intermediate value theorem in the same fashion as in the discussion preceding Lemma 2.8: To see that m is indeed an intermediate value between the means of θ obtained for $\zeta = 0$ and $\zeta = k$, we consider the non-negative and convex functions $x \mapsto (x - G_\nu(1 - k))_+$ and $x \mapsto (G_\nu(k) - x)_+$ and integrate them against μ and ν in the same way as we did above to obtain the inequalities in (5). Clearly the mean of θ depends continuously on ζ and hence the intermediate value theorem yields the existence of the desired ζ .

We are now given two measures μ and θ of the same mass and the same mean. Consider a convex function φ . We want to prove that its integral with respect to μ is smaller than the one with respect to θ . For that we can assume without loss of generality $\varphi(G_\nu(\zeta)) = \varphi(G_\nu(\zeta + (1 - k))) = 0$. Then

$$\begin{aligned} \int \varphi d\mu(x) &\leq \int \varphi_+(x) d\mu(x) \\ &\leq \int \varphi_+(x) d\nu(x) = \int \varphi_+(x) d\theta(x) = \int \varphi(x) d\theta(x). \end{aligned}$$

This concludes the proof. \square

4.2. Maximal and minimal elements. For $\mu \preceq_E \nu$, let F_μ^ν be the set of measures η such that $\mu \preceq_C \eta$ and $\eta \leq \nu$. Note that the measures in F_μ^ν have the same mass and the same barycenter as μ . In the next lemmas we consider the partially ordered set (F_μ^ν, \preceq_C) and show that it has both a maximal and a minimal element.

Lemma 4.5. *For $\mu \preceq_E \nu$, the set F_μ^ν has an element which is maximal w.r.t. the convex order, i.e. there exists $T^\nu(\mu)$ such that*

(i) $T^\nu(\mu) \leq \nu$.

(ii) $\mu \preceq_C T^\nu(\mu)$.

(iii) *If η is another measure satisfying (i) and (ii) then we have $\eta \preceq_C T^\nu(\mu)$.*

Proof. Consider the measure θ defined as in the proof of Proposition 4.4 and let η be another measure in F_μ^ν . We know that θ is concentrated outside an open interval I and that it coincides with ν on $\mathbb{R} \setminus \bar{I}$ so that $\theta|_{\mathbb{R} \setminus \bar{I}} \geq \eta|_{\mathbb{R} \setminus \bar{I}}$. Thus $\eta - (\eta \wedge \theta)$ is concentrated on \bar{I} whereas $\theta - (\eta \wedge \theta)$ is concentrated on $\overline{\mathbb{R} \setminus \bar{I}}$. It follows from Example 2.5 that $\eta \preceq_C \theta$. \square

The existence of a minimal element is more involved and will play an important role subsequently.

Lemma 4.6 (Shadow embedding). *Let $\mu, \nu \in \mathcal{M}$ and assume $\mu \preceq_E \nu$. Then there exists a measure $S^\nu(\mu)$, called the shadow of μ in ν , such that*

(i) $S^\nu(\mu) \leq \nu$.

(ii) $\mu \preceq_C S^\nu(\mu)$.

(iii) *If η is another measure satisfying (i) and (ii) then we have $S^\nu(\mu) \preceq_C \eta$.*

As a consequence of (iii), the measure $S^\nu(\mu)$ is uniquely determined. Moreover it satisfies the following property:

(iii') *If η is a measure such that $\eta \leq \nu$ and $\mu \preceq_E \eta$ then we have $S^\nu(\mu) \preceq_E \eta$.*

Note that if $\mu \preceq_C \nu$, i.e. if μ and ν have the same mass, then the shadow $S^\nu(\mu)$ is just ν itself because this is the only measure η with mass $\mu(\mathbb{R}) = \nu(\mathbb{R})$ that satisfies $\eta \leq \nu$.

Proof of Lemma 4.6. First observe that (iii') follows from Lemma 4.4 applied to μ and η .

We write k (resp. m) for the mass (resp. the mean) of μ . The principal strategy of our proof is to rewrite the problem in terms of potential functions. Set $f = u_\mu$ and $g = u_\nu$.

The task is to find a convex function h (corresponding to $u_{S^\nu(\mu)}$) such that

- (1) $h - g$ is concave, i.e. $h'' \leq g''$ in a weak sense.
- (2) $f \leq h$ and $\lim_{|x| \rightarrow \infty} h(x) - k|x - m| = 0$.
- (3) We have $h \leq h_2$ for all functions h_2 in the set

$$U_F = \{h \text{ is convex and satisfies (1) and (2)}\} = \{h = u_\eta : \eta \in F\}$$

We note that by Lemma 4.4 there exist functions satisfying Conditions (1) and (2). Hence the sets $F = \{\eta : \mu \preceq_C \eta, \eta \leq \nu\}$ and U_F are not empty. Looking for a function which also satisfies the third property we define

$$(7) \quad \bar{h} = \inf_{h \in U_F} h.$$

If this function is convex, which we shall show below, it will satisfy the three required conditions. Conditions (2) and (3) are clear, let us briefly prove (1): every function $h \in U_F$ is “less convex” than g , i.e. the function $h - g$ is concave. Hence $\bar{h} - g = (\inf_{h \in U_F} h) - g = \inf_{h \in U_F} (h - g)$ is also concave.

The convexity of \bar{h} will be proved if we can establish that its epigraph $\mathcal{E}(\bar{h})$ is convex, i.e. that every segment of \mathbb{R}^2 with both ends in $\mathcal{E}(\bar{h})$ is included in this set.

This will be the case if U_F is stable under the following operation: take h_1, h_2 in U_F and let h_{\min} be the convex hull of $x \mapsto \min(h_1(x), h_2(x))$. More precisely

$$h_{\min}(x) = \inf_{ab \geq 0, (a,b) \neq (0,0)} \frac{bh_1(x-a) + ah_2(x+b)}{a+b}.$$

Since $\lim_{|c| \rightarrow \infty} (h_1 - h_2)(x+c) = 0$ this infimum is in fact a minimum. Condition (2) holds for h_{\min} . It remains to prove that $h_{\min} - g$ is concave.

We use a non-usual but clear characterization of concavity: a real function is concave if and only if has locally an upper tangent in every point. More precisely f is concave if for every $x \in \mathbb{R}$, there exists an affine function l with $l(x) = f(x)$ and $l \geq f$ in a neighborhood of x . With respect to the definition of h_{\min} , there are two kinds of real x . A point x such that $h_{\min}(x)$ equals $h_i(x)$ for some $i \in \{1, 2\}$ is of the first kind. In this case the property is true because $h_i \geq h_{\min}$ so that $h_i - g \geq h_{\min} - g$ where the first function is concave. These relations even hold globally. In the other case there exist a, b with $ab > 0$ such that $h_{\min}(x) = \frac{bh_1(x-a) + ah_2(x+b)}{a+b}$. Without loss of generality we may assume $a > 0$ and $b > 0$. As h_{\min} both is convex and its graph is below the cord $[(x-a, h_1(x-a)), (x+b, h_2(x+b))]$ we can conclude that it is affine on $[x-a, x+b]$. Hence $h_{\min} - g$ is concave in a neighborhood of x . Summing up, the property holds for the two kinds of real x . Finally $h_{\min} - g$ is concave and $h_{\min} \in U_F$. Hence \bar{h} is convex and satisfies Conditions (1)–(3). \square

Note that in Lemma 2.8 we have implicitly encountered the shadow in the case where the starting distribution consists of an atom.

Example 4.7 (Shadow of an atom). *Let δ be an atom of mass α at a point x . Assume that $\delta \preceq_E \nu$. Then $S^\nu(\delta)$ is the restriction of ν between two quantiles, i.e. it is $\nu' = (G_\nu)_{\#} \lambda_{[s, s']}$ where $s' - s = \alpha$ and the barycenter of ν' is x . In particular for another measure $\eta \in \mathcal{M}$ with $\delta \preceq_C \eta$ and $\eta \leq \nu$, applying the observation from Example 2.5 to ν' and η we obtain $\nu' \preceq_C \eta$.*

4.3. Associativity of shadows. In this section we will establish the following associativity property of the shadow.

Theorem 4.8 (Shadow of a sum). *Let γ_1, γ_2 and ν be elements of \mathcal{M} and assume that $\mu = \gamma_1 + \gamma_2 \preceq_E \nu$. Then we have $\gamma_2 \preceq_E \nu - S^\nu(\gamma_1)$ and*

$$S^\nu(\gamma_1 + \gamma_2) = S^\nu(\gamma_1) + S^{\nu - S^\nu(\gamma_1)}(\gamma_2).$$


In Figure 3 we can see the shadow of $\mu = \gamma_1 + \gamma_2$ in ν for two different ways of labeling the γ_i 's. In both cases $\nu_1 := S^\nu(\gamma_1)$ is simply γ_1 . On the left part of the figure $S^{\nu - \nu_1}(\gamma_2)$ is quite intuitive while on the right part it is deduced from the associativity of the shadow projection. Of course it has to be $S^\nu(\mu) - \nu_1$.

Our proof of Theorem 4.8 will rely on approximations of μ by atomic measures and we need several auxiliary results. In our argument we will require a certain continuity property of the mapping $\nu \mapsto S^\nu(\delta)$ stated in Lemma 4.10. We will derive it now with the help of the Kantorovich metric.

Proposition 4.9 (Metric on \mathcal{M}). *The function W , defined on \mathcal{M} by*

$$(8) \quad W(\nu, \hat{\nu}) = \begin{cases} +\infty & \text{if } \nu(\mathbb{R}) \neq \hat{\nu}(\mathbb{R}), \\ \sup_f (\int f d\nu - \int f d\hat{\nu}) & \text{otherwise,} \end{cases}$$

where the supremum is taken over all 1-Lipschitz functions $f : \mathbb{R} \rightarrow \mathbb{R}$ is a metric with values in $[0, +\infty]$. For $k > 0$, the associated topology on the subspaces of measure of mass k coincides with the weak topology introduced in Paragraph 2.1.

FIGURE 3. Shadow of $\mu = \gamma_1 + \gamma_2$ in ν .

In the case where $\nu, \hat{\nu}$ are probability measures, $W(\nu, \hat{\nu})$ is the classical Kantorovich metric (also called 1-Wasserstein distance, or transport distance). We state here two useful relations that well known (and straight forward) in the case of probability measures and extended to finite measures through normalization. If $\nu(\mathbb{R}) = \hat{\nu}(\mathbb{R})$, we have

$$W(\nu, \hat{\nu}) = \|F_\nu - F_{\hat{\nu}}\|_1 = \|G_\nu - G_{\hat{\nu}}\|_1$$

where $F_\nu, F_{\hat{\nu}}$ and $G_\nu, G_{\hat{\nu}}$ are the cumulative distribution functions and the quantile functions of ν and $\hat{\nu}$ respectively. The norm $\|\cdot\|_1$ refers to the L^1 -norm for the Lebesgue measure on \mathbb{R} , resp. $[0, \nu(\mathbb{R})]$. Recall that $\nu = (G_\nu)_\# \lambda$ and $\hat{\nu} = (G_{\hat{\nu}})_\# \lambda$.

Let us now fix some notations in preparation to Lemma 4.10. First, let ν and $\hat{\nu}$ be of mass 1. We also fix a quantity $\alpha \leq 1$ and set $t = 1 - \alpha$. As in the discussion preceding Lemma 2.8 we consider for $s \in [0, t]$ the restriction $\nu_s = (G_\nu)_\# \lambda_{[s, s+\alpha]}$ of ν between the quantiles s and $s + \alpha$. We adopt the same convention for $\hat{\nu}$. Note that the barycenter of ν_s can be written

$$(9) \quad B(s, \nu) = \frac{1}{\alpha} \int_{\mathbb{R}} x d\nu_s(t) \quad \text{or} \quad B(s, \nu) = \frac{1}{\alpha} \int_0^\alpha G_\nu(s+t) d\lambda(t).$$

Indeed, the function $t \in [0, \alpha] \mapsto G_\nu(s+t)$ is simply G_{ν_s} and $\nu_s = (G_{\nu_s})_\# \lambda_{[0, \alpha]}$.

Together with (8) applied to the functions $f : x \mapsto \pm x$, the first formula for the barycenter implies

$$|B(s, \nu) - B(s, \hat{\nu})| \leq \frac{1}{\alpha} W(\nu_s, \hat{\nu}_s).$$

Moreover we can prove that

$$W(\nu_r, \nu_s) = \alpha |B(r, \nu) - B(s, \nu)|$$

without difficulty by using $W(\nu_r, \nu_s) = \|G_{\nu_r} - G_{\nu_s}\|_1$ and the fact that G_{ν_s} and G_{ν_r} are equal to the non-decreasing function G_ν up to translation. Another simple property is

$$W(\nu_s, \hat{\nu}_s) \leq W(\nu, \hat{\nu}).$$

Again this can be seen as a consequence of the representation of W by quantile functions: we have $W(\nu_s, \hat{\nu}_s) = \|G_{\nu_s} - G_{\hat{\nu}_s}\|_1 \leq \|G_\nu - G_{\hat{\nu}}\|_1$.

Let x be an element of \mathbb{R} and consider the subset of measures $\nu \in \mathcal{P}$ such that $B(0, \nu) \leq x \leq B(t, \nu)$. These are exactly the measures such that there exists $s \in \mathbb{R}$ satisfying $B(s, \nu) = x$; for such ν the shadow $S^\nu(\delta) = \nu_s$ is well defined.

Lemma 4.10. *Let $\delta = \alpha \delta_x$ be an atom of mass $\alpha < 1$. The map $\nu \mapsto S^\nu(\delta)$ is continuous on its domain of definition inside the probability measures.*

Proof. Let $\nu, \hat{\nu}$ be probability measures in \mathcal{M} and assume that $S^\nu(\delta), S^{\hat{\nu}}(\delta)$ exist. Let r, s be such that $\nu_r = S^\nu(\delta)$ and $\hat{\nu}_s = S^{\hat{\nu}}(\delta)$. Of course both measures have the same barycenter. Then

$$\begin{aligned} W(S^\nu(\delta), S^{\hat{\nu}}(\delta)) &= W(\nu_r, \hat{\nu}_s) \\ &\leq W(\nu_r, \nu_s) + W(\nu_s, \hat{\nu}_s) \\ &= \alpha |B(r, \nu) - B(s, \nu)| + W(\nu_s, \hat{\nu}_s) \\ &= \alpha |B(s, \hat{\nu}) - B(s, \nu)| + W(\nu_s, \hat{\nu}_s) \\ &\leq W(\nu_s, \hat{\nu}_s) + W(\nu_s, \hat{\nu}_s) \leq 2W(\nu, \hat{\nu}). \quad \square \end{aligned}$$

Lemma 4.11. *Let δ be an atom and assume $\delta \preceq_E \eta$, where $\eta \leq \nu$. Then we have*

$$\eta - S^\eta(\delta) \leq \nu - S^\nu(\delta).$$

Proof. First note that $S^\eta(\delta) \leq \eta \leq \nu$. Hence $\delta \preceq_E \nu$ and $S^\nu(\delta)$ is well defined. As explained in Example 4.7, there exists an interval $Q \subseteq [0, \nu(\mathbb{R})]$ such that $S^\nu(\mu)$ equals $G_{\nu\#}\lambda_Q$. The same is true for δ, η, G_η and some interval of $[0, \eta(\mathbb{R})]$ but we will represent the “quantile coordinates” of $S^\eta(\delta)$ under η in a slightly different way. Indeed $S^\eta(\delta)$ is the restriction of η to a real interval plus possibly some atomic parts of η at the ends of this interval. In any case, it is smaller than η and ν . Thus we can parametrize it with a subinterval Q' of $[0, \nu(\mathbb{R})]$ such that $S^\eta(\delta) = (G_{\nu\#}\lambda_{Q'}) \wedge \eta$. Note that the length of Q' is greater than the length of Q which equals the mass of δ . The measures $S^\nu(\delta)$ and $S^\eta(\delta)$ have the same mass and the same barycenter and both are smaller than ν .

We prove by contradiction that $Q \subseteq Q'$. By symmetry it is enough to prove $b' \geq b$ where we denote Q and Q' by $[a, b]$ and $[a', b']$ respectively. If it were not the case, $S^\eta(\delta)$ would be stochastically strictly smaller than $S^\nu(\delta)$, which is the right-most measure that stays smaller than quantile b , has the same mass as δ and is smaller than ν . In particular the barycenters would be strictly ordered (see the discussion before Lemma 2.8 for a similar and more detailed argument). This is a contradiction since the barycenters coincide by the definition of the shadow. Finally

$$\eta - S^\eta(\delta) = \eta - [(G_{\nu\#}\lambda_{Q'}) \wedge \eta] \leq \nu - G_{\nu\#}\lambda_{Q'} \leq \nu - G_{\nu\#}\lambda_Q.$$

Here we used the fact that for three measures α, β, γ satisfying the relations $\alpha \leq \gamma$ and $\beta \leq \gamma$, the measure $\gamma - \alpha$ is greater than the positive part of $\beta - \alpha$, which is $\beta - (\alpha \wedge \beta)$. \square

Lemma 4.12 (Shadow of one atom and one measure). *Consider now $\delta + \gamma$ where δ is an atom. Assume $(\delta + \gamma) \preceq_E \nu$. Then we have $\gamma \preceq_E \nu - S^\nu(\delta)$ and*

$$(10) \quad S^\nu(\delta + \gamma) = S^\nu(\delta) + S^{\nu - S^\nu(\delta)}(\gamma).$$

Proof. We first prove that γ is smaller than $\nu' := \nu - S^\nu(\delta)$ in the extended order. Note that there exists an interval I such that $S^\nu(\delta)$ is concentrated on \bar{I} and $\nu'(I) = 0$. Let φ be a non-negative convex function which satisfies $\limsup_{|x| \rightarrow +\infty} |\varphi(x)/x| < +\infty$. We will prove $\int \varphi d\gamma \leq \int \varphi d\nu'$. For that we introduce ψ which equals φ on $\mathbb{R} \setminus I$ and is linear on I . We can assume that ψ is convex and $\psi \geq \varphi$ (even if I is unbounded). Note that φ and ψ coincide on the border of I . We have

$$\int \varphi d\gamma \leq \int \psi d\gamma \leq \int \psi d\nu - \int \psi d\delta.$$

But $\int \psi d\delta = \int \psi dS^\nu(\delta)$ because ψ is linear on I and this quantity is greater than $\int \varphi dS^\nu(\delta)$. Moreover $\int \psi d\nu' = \int \varphi d\nu'$ because ν' is concentrated on $\mathbb{R} \setminus I$. It follows that

$$\int \varphi d\gamma \leq \int \psi d\nu - \int \psi d\delta \leq \int \varphi d\nu'.$$

As in the case of the usual convex order, it is of course sufficient to test against convex functions of linear growth, hence $\gamma \preceq_E \nu'$.

It remains to establish (10). It is clear (see for instance Example 2.3) that both sides of the equation are greater than $\delta + \gamma$ in the convex order and $\leq \nu$. Hence by the definition of the shadow it follows $S^\nu(\delta + \gamma) \preceq_C S^\nu(\delta) + S^{\nu - S^\nu(\delta)}(\gamma)$. The other inequality is shown as follows: we will prove that for $\eta \succeq_C \delta + \gamma$ and satisfying $\eta \leq \nu$ we have $S^\nu(\delta) + S^{\nu - S^\nu(\delta)}(\gamma) \preceq_C \eta$. In fact if $\eta \succeq_C \delta + \gamma$ then $S^\eta(\delta) \leq \eta$ and $S^{\eta - S^\eta(\delta)}(\gamma) \leq \eta - S^\eta(\delta)$ so that, since measures in the convex order have the same mass,

$$\eta = S^\eta(\delta) + S^{\eta - S^\eta(\delta)}(\gamma).$$

(Note that we have already proved that all terms exist in this decomposition since \succeq_E extends \succeq_C .) But it follows from $\eta \leq \nu$ and $\eta - S^\eta(\delta) \leq \nu - S^\nu(\delta)$ (proved in Lemma 4.11) that $F_\gamma^\eta \subseteq F_\gamma^\nu$ and $F_\gamma^{\eta - S^\eta(\delta)} \subseteq F_\gamma^{\nu - S^\nu(\delta)}$ so that $S^\eta(\delta) \succeq_C S^\nu(\delta)$ and $S^{\eta - S^\eta(\delta)}(\gamma) \succeq_C S^{\nu - S^\nu(\delta)}(\gamma)$. As in Example 2.3 the compatibility of sum and convex order concludes the proof. \square

Lemma 4.13 (Shadow of finitely many atoms). *Let $(\delta_i)_{i \in \mathbb{N}}$ be a family of atoms at point x_i and of mass $\alpha_i \in [0, +\infty[$ (where we allow the weight α_i to be 0). For every $n \geq 1$, let $\mu_n = \delta_1 + \dots + \delta_n$ and assume that $\mu_n \preceq_E \nu$. The sequence $(\nu_n)_{n \in \mathbb{N}}$ defined by $\nu_n = S^\nu(\mu_n)$ satisfies the following recurrence relation:*

- $\nu_0 = 0$
- $\nu_n = \nu_{n-1} + S^{\nu - \nu_{n-1}}(\delta_n)$ for every $n \geq 1$.

Proof. The lemma is proved by induction. The basis holds with $\nu_1 = S^\nu(\delta_1)$. Fix $n \geq 1$ and assume that the recurrence relation holds until n . Let $(\mu_i)_i$, ν and $(\nu_i)_i$ be as in the statement of the lemma. Denote $\sum_{i=2}^{n+1} \delta_i$ by μ'_n and more generally $\sum_{i=2}^{i+1} \delta_i$ by μ'_i . As $\mu_{n+1} \preceq_E \nu$ we can apply Lemma 4.12 to the decomposition $\mu_{n+1} = \delta_1 + \mu'_n$. So $\mu'_n \preceq_E \nu - \nu_1$ and

$$(11) \quad S^\nu(\mu_{n+1}) = S^\nu(\delta_1) + S^{\nu'}(\mu'_n)$$

where we denoted $\nu - \nu_1$ by ν' . But because of the inductive hypothesis applied to μ'_n and ν' , the shadow $S^{\nu'}(\mu'_n)$ is $\nu'_n = \nu'_{n-1} + S^{\nu' - \nu'_{n-1}}(\delta_{n+1})$ where the measures ν'_i denote the shadows of μ'_i in ν' . Note also that $\nu_n = \nu_1 + \nu'_{n-1}$ by Lemma 4.12. Starting from (11) we now have

$$\nu_{n+1} = \nu_1 + \nu'_n = \nu_1 + \nu'_{n-1} + S^{\nu' - \nu'_{n-1}}(\delta_{n+1}) = \nu_n + S^{\nu' - \nu'_{n-1}}(\delta_{n+1}).$$

But $\nu' - \nu'_{n-1} = (\nu_1 + \nu') - (\nu_1 + \nu'_{n-1}) = \nu - \nu_n$. This concludes the proof. \square

Remark 4.14. *An important consequence of the lemma above is that $\nu_n - \nu_k$ is the shadow of $\mu_n - \mu_k$ in $\nu - S^\nu(\mu_k)$. Even though the above construction is of inductive nature, when permuting the n first atoms, the measure $\nu_n = \sum_{i=1}^n \nu_i - \nu_{i-1}$ is always the same: it simply equals $S^\nu(\mu_n)$. The same assertions apply to Proposition 4.17 below.*

Proposition 4.15. *Assume that $(\mu_n)_n$ is increasing in the convex order and $\mu_n \preceq_C \mu \preceq_E \nu$ for every $n \in \mathbb{N}$. Then both $(\mu_n)_n$ and $(S^\nu(\mu_n))_n$ converge in \mathcal{M} . If we call μ_∞ , respectively S_∞ the limits, then the measure S_∞ is the shadow of μ_∞ in ν .*

Proof. First note that the assumptions imply $u_{\mu_0} \leq u_{\mu_1} \leq \dots \leq u_{\mu_n}$ and $u_{\mu_n} \leq u_\mu$. The limit $u_\infty := \lim_{n \in \mathbb{N}} u_{\mu_n}$ exists because for every $x \in \mathbb{R}$, $(u_{\mu_n}(x))_n$ is increasing and bounded from above. Of course the limit u_∞ is a convex function and since u_μ is an upper bound it has the correct asymptotic behavior. Therefore u_∞ is

a potential function and by Proposition 4.1 it is the potential function of some $\mu_\infty \in \mathcal{M}$ with the same mass and mean as μ and the μ_n 's.

On the other hand, for $n \in \mathbb{N}$ we consider the set $F_{\mu_n}^\nu$ of measures η_n satisfying $\mu_n \preceq_C \eta_n$ and $\eta_n \leq \nu$. (We are using the notations of the proof of Lemma 4.6.) The measure $S^\nu(\mu_n)$ is the smallest element of $F_{\mu_n}^\nu$ with respect to the convex order. The family $F_{\mu_n}^\nu$ is decreasing in n and it is not difficult to see that $F_\mu^\nu \subseteq \bigcap F_{\mu_n}^\nu$ so that it is not empty. Hence $S^\nu(\mu_n)$ is increasing in the convex order and it is bounded from above by $S^\nu(\mu)$. Exactly for the same reasons as for the sequence $(\mu_n)_n$, it converges to some S_∞ in \mathcal{M} . We now have to conclude that $S^\nu(\mu_\infty) = S_\infty$. We will in fact prove that $S_\infty \preceq_C S^\nu(\mu_\infty)$ and $S^\nu(\mu_\infty) \preceq_C S_\infty$.

For every n we have $\mu_n \preceq_C \mu_\infty \preceq_C S^\nu(\mu_\infty)$ and $S^\nu(\mu_\infty) \leq \nu$. Thus $S^\nu(\mu_n) \preceq_C S^\nu(\mu_\infty)$. By Proposition 4.2 we have $S_\infty \preceq_C S^\nu(\mu_\infty)$. Conversely, using again Proposition 4.2, the relation $\mu_n \preceq_C S^\nu(\mu_n)$ yields $\mu_\infty \preceq_C S_\infty$ as n goes to $+\infty$. But $S_\infty \leq \nu$ (the limit of a converging sequence $(u_\nu - u_{S^\nu(\mu_n)})_n$ is convex). Hence $S^\nu(\mu_\infty) \preceq_C S_\infty$. \square

Lemma 4.16 (Shadow of one measure and one atom). *Consider now $\gamma + \delta$ where δ is an atom. Assume $(\gamma + \delta) \preceq_E \nu$. Then we have $\delta \preceq_E S^\nu(\gamma + \delta) - S^\nu(\gamma)$ and*

$$(12) \quad S^\nu(\gamma + \delta) = S^\nu(\gamma) + S^{\nu - S^\nu(\gamma)}(\delta).$$

Proof. If γ is the sum of finitely many atoms, the result follows from Lemma 4.13. Let us consider an approximating sequence $(\gamma^{(n)})_n$ of γ as in Lemma 2.9. We can write the decomposition of the shadow of $\gamma^{(n)} + \delta$ in ν as in the statement of the lemma and apply Proposition 4.15 to the sequence $(S^\nu(\gamma^{(n)}))_n$. It follows that the limit exists and equals $S^\nu(\gamma)$. Write $\nu^{(n)}$ for $S^\nu(\gamma^{(n)})$ and $\nu^{(\infty)}$ for $S^\nu(\gamma)$. For the same reasons as above the shadows of $\gamma^{(n)} + \delta$ converge to $S^\nu(\gamma + \delta)$.

We still have to show that $S^{\nu - \nu^{(n)}}(\delta)$ converges to $S^{\nu - \nu^{(\infty)}}(\delta)$. We know that $\nu^{(n)}$ converges to $\nu^{(\infty)}$ in \mathcal{M} so $\nu - \nu^{(n)}$ tends to $\nu - \nu^{(\infty)}$ and all these measures are bounded by ν (in particular they have a density smaller than 1 with respect to ν). We also know from Example 4.7 that $S^{\nu - \nu^{(n)}}(\delta)$ is the restriction of $\nu - \nu^{(n)}$ to the (uniquely determined) ‘‘quantile interval’’ with the correct mass and barycenter. Rescaling masses if necessary, the continuity Lemma 4.10 implies that $S^{\nu - \nu^{(n)}}(\delta)$ converges to $S^{\nu - \nu^{(\infty)}}(\delta)$. \square

We are now finally in the position to prove the desired associativity property of the shadow mapping.

Proof of Theorem 4.8. If γ_2 is the sum of finitely many atoms, the property holds since by Lemma 4.16 it is possible to construct recursively $S^\nu(\gamma_1 + \gamma_2)$ using a decomposition with one atom from γ_2 and the rest of $\gamma_1 + \gamma_2$ as the second measure. Let us consider a sequence $(\gamma_2^{(n)})_n$ of measures consisting of finitely many atoms that weakly converge to γ_2 and satisfy $\gamma_2^{(n)} \preceq_C \gamma_2$. Moreover we may assume that $(\gamma_2^{(n)})_n$ is increasing in the convex order as in Lemma 2.9.

We can write the decomposition of the shadow of $\gamma_1 + \gamma_2^{(n)}$ in ν as in the statement of the theorem and apply Proposition 4.15 to the sequence $(S^{\nu - S^\nu(\gamma_1)}(\gamma_2^{(n)}))_n$. We obtain that the limit exists and equals $S^{\nu - S^\nu(\gamma_1)}(\gamma_2)$. For the same reasons the shadow of $\gamma_1 + \gamma_2^{(n)}$ converges to $S^\nu(\gamma_1 + \gamma_2)$. This concludes the proof. \square

Before we define the left-curtain transport plan it seems worthwhile to record the following result.

Proposition 4.17 (Shadow of the sum of finitely many measures). *Let $(\gamma_i)_i$ be a family of measures (that possibly vanish identically). Let $\mu_n = \gamma_1 + \dots + \gamma_n$. Assume*

also that $\mu_n \preceq_E \nu$ for every $n \geq 1$. The sequence $(\nu_n)_{n \in \mathbb{N}}$ defined by $\nu_n = S^\nu(\mu_n)$ satisfies the following recurrence relation:

- $\nu_0 = 0$
- $\nu_n - \nu_{n-1} = S^{\nu - \nu_{n-1}}(\gamma_n)$.

Proof. The statement is the same as Lemma 4.13 except that we do not require the measures γ_i to be atoms. Lemma 4.13 relies on Lemma 4.12 which characterizes the shadow of $\gamma_1 + \gamma_2$ under the assumption that γ_1 is an atom. Substituting it with Theorem 4.8 the present claim follows verbatim. \square

Let us now formally define the left-curtain coupling π_{1c} that has been discussed in the introduction and whose properties will be derived in the sequel. We baptize it the “left-curtain transport plan” because it projects shadow measures as a curtain that one closes starting from the left side.

Note that given measures $\mu \leq \mu' \preceq_E \nu$, Theorem 4.8 implies that $S^\nu(\mu) \leq S^\nu(\mu')$. This property is essential for the definition of π_{1c} .

Theorem 4.18 (Definition of π_{1c}). *Assume that $\mu \preceq_C \nu$. There is a unique probability measure π_{1c} on $\mathbb{R} \times \mathbb{R}$ which transports $\mu|_{]-\infty, x]}$ to $S^\nu(\mu|_{]-\infty, x]}$, i.e. satisfies $\text{proj}_{\#}^x(\pi_{1c}|_{]-\infty, x]} \times \mathbb{R}) = \mu|_{]-\infty, x]}$ and $\text{proj}_{\#}^y(\pi_{1c}|_{]-\infty, x]} \times \mathbb{R}) = S^\nu(\mu|_{]-\infty, x]}$ for all $x \in \mathbb{R}$. Moreover π_{1c} is a martingale transport plan which takes μ to ν , i.e. $\pi_{1c} \in \Pi_M(\mu, \nu)$.*

Proof. Plainly, the condition given in the statement prescribes the value of

$$\pi_{1c}(]-\infty, x] \times A) = S^\nu(\mu|_{]-\infty, x]})(A)$$

for $x \in \mathbb{R}$ and every Borel set $A \subseteq \mathbb{R}$, thus giving rise to a unique measure on the product space. Here we use that, by Theorem 4.8, $S^\nu(\mu|_{]-\infty, x]}) \leq S^\nu(\mu|_{]-\infty, x']})$ whenever $x \leq x'$.

Clearly, the first marginal of π_{1c} equals μ . By construction, the second marginal satisfies $\text{proj}_{\#}^y \pi_{1c} \leq \nu$. Since μ and ν have the same mass, this implies $\text{proj}_{\#}^y \pi_{1c} = \nu$ as required.

To establish the martingale property we show that property (4) holds for any function $\rho = \mathbb{1}_{]-\infty, x']}$, $x' \in \mathbb{R}$. Indeed we have

$$\int (y - x)\rho(x) d\pi_{1c}(x, y) = \int y dS^\nu(\mu|_{]-\infty, x']})(y) - \int x d\mu|_{]-\infty, x']}(x) = 0.$$

\square

Remark 4.19. *The family of intervals $(]-\infty, x])_{x \in \mathbb{R}}$ is totally ordered with respect to \subseteq and it spans the σ -field of Borel measurable sets. In the proof of Theorem 4.18 we used these properties to show that there is a unique martingale transport plan which transports $\mu|_{]-\infty, x]}$ to $S^\nu(\mu|_{]-\infty, x]}$. This construction can be applied to more general families of sets: let I be some index set and $(C_\iota)_{\iota \in I}$ a family of Borel sets that both is totally ordered with respect to \subseteq and spans the σ -field of Borel sets. Then a measure $\pi \in \Pi_M(\mu, \nu)$ is defined uniquely by the relations $\pi(C_\iota \times A) = S^\nu(\mu|_{C_\iota})(A)$ for all indices $\iota \in I$ and Borel sets $A \subseteq \mathbb{R}$.*

Example 4.20. *In the case of a finitely supported measure $\mu = \sum_{i=1}^n \delta_i$, it follows that if the ordering is done so that the support of δ_i is $\{x_i\}$ with $x_1 \leq \dots \leq x_n$, then the π_{1c} -coupling is $\pi_{1c} = \sum_{i=1}^n \tilde{\delta}_i \otimes S^{\nu - \nu_{i-1}}(\delta_i)$ where $\tilde{\delta}_i = \delta_i / \delta_i(x_i)$ are the properly renormalized versions of δ_i and the measures ν_i are $S^\nu(\mu_i)$ with $\mu_i = \delta_1 + \dots + \delta_i$ as in Lemma 4.13.*

Theorem 4.21. *The martingale π_{1c} is left-monotone in the sense of Definition 1.4.*

Proof. Note that π_{1c} is simultaneously a minimizer for all cost functions of the form $c_{s,t}(x,y) = \mathbb{1}_{]-\infty,s]}(x)|y-t|$, where s,t are real numbers. Indeed if π is an arbitrary martingale transport plan then

$$\iint c_{s,t}(x,y) d\pi(x,y) = \iint_{]-\infty,s] \times \mathbb{R}} |y-t| d\pi(x,y) = \int |y-t| d(\text{proj}_{\#}^y \pi|_{]-\infty,s] \times \mathbb{R}})(y).$$

Setting $\nu_s^\pi = \text{proj}_{\#}^y \pi|_{]-\infty,s] \times \mathbb{R}}$ we have $\nu_s^\pi \leq \nu$ and $\mu|_{]-\infty,s]} \preceq_C \nu_s^\pi$ which implies $S^\nu(\mu|_{]-\infty,s]}) \preceq_C \nu_s^\pi$. Therefore

$$\int |y-t| dS^\nu(\mu|_{]-\infty,s]})(y) \leq \int |y-t| d\nu_s^\pi(y),$$

where equality holds for all $s,t \in \mathbb{R}$ if (and only if) $\pi = \pi_{1c}$.

Applying Lemma 1.11 to the costs $c_{s,t}$ for $s,t \in \mathbb{Q}$, we obtain a Borel set $\Gamma_{s,t}$ of π_{1c} -measure 1. Set $\Gamma = \bigcap_{s,t \in \mathbb{Q}} \Gamma_{s,t}$. We claim that a configuration as in (3) cannot appear in Γ . Indeed if (x,y^-) , (x,y^+) and (x',y') are in Γ and satisfy $x < x'$ and $y^- < y' < y^+$, they are also in $\Gamma_{s,t}$ where (s,t) satisfies $s \in]x,x'[$ and $t \in]y',y^+[$. Let $\lambda \in]0,1[$ be such that $y' = \lambda y^+ + (1-\lambda)y^-$. The measure $\alpha = \lambda \delta_{(x,y^+)} + (1-\lambda)\delta_{(x,y^-)} + \delta_{(x',y')}$ is concentrated on Γ but the competitor $\alpha' = \lambda \delta_{(x',y^+)} + (1-\lambda)\delta_{(x',y^-)} + \delta_{(x,y')}$ leads to a lower global cost. This yields the desired contradiction. \square

5. UNIQUENESS OF THE MONOTONE MARTINGALE TRANSPORT

In this section, we establish that the left-curtain coupling π_{1c} is the unique monotone martingale coupling. Our proof of this result is specific to the present setup. We will also explain a more classical argument that is often invoked in the optimal transport theory to establish some uniqueness property. This so called *half sum* argument will be used several times subsequently but requires the initial distribution μ to be continuous.

We start with two preliminary lemmas which are required to derive the main result of this part, Theorem 5.3.

Lemma 5.1. *If $\mu \preceq_C \nu$ then one of the following statements holds true:*

- we have $\mu(]a, +\infty[) > 0$ and $\nu(]a, +\infty[) > 0$ for every a ;
- the number $a = \sup(\text{spt}(\mu))$ is finite and $\nu(]a, +\infty[) > 0$;
- the number $a = \sup(\text{spt}(\mu))$ is finite and $\nu(]a, +\infty[) = 0$. Moreover $\nu(\{a\}) \geq \mu(\{a\})$.

The corresponding result for intervals of the form $] -\infty, b[$ is true as well.

Proof. Integrating the convex function $x \mapsto (x - a')^+$ for different values of a' we obtain $\sup(\text{spt}(\mu)) \leq \sup(\text{spt}(\nu))$. Therefore the first case corresponds to $\sup(\text{spt}(\mu)) = \sup(\text{spt}(\nu)) = +\infty$, the second to $\sup(\text{spt}(\mu)) < \sup(\text{spt}(\nu))$ and the third to $\sup(\text{spt}(\mu)) = \sup(\text{spt}(\nu)) < +\infty$.

Let us prove that in the third case we also have $\mu(\{a\}) \leq \nu(\{a\})$. If $\mu(\{a\}) = 0$ we are done. If $\mu(\{a\}) > 0$, the conditional transport measure π_a must be the static transport because it is a martingale transport plan and $\sup(\text{spt}(\nu)) = a$. This concludes the proof. \square

For $u, v \in \mathbb{R}, u < v$ let $g_{u,v}$ be defined by

$$(13) \quad g_{u,v}(x) = \begin{cases} v - x & \text{if } x \in [u, v], \\ 0 & \text{otherwise.} \end{cases}$$

Lemma 5.2. *Let σ be a non trivial signed measure of mass 0 and denote its Hahn decomposition by $\sigma = \sigma^+ - \sigma^-$. There exist $a \in \text{spt}(\sigma^+)$ and $b > a$ such that $\int g_{a,b}(x) d\sigma(x) > 0$.*

Proof. First notice that $u \mapsto \int g_{u,u+1}(x) d\sigma(x)$ does not vanish identically. Since, by Fubini's theorem,

$$\iint g_{u,u+1}(x) d\sigma(x) du = 0$$

there exists $u \in \mathbb{R}$ such that $\int g_{u,u+1}(x) d\sigma(x) > 0$. The set $\text{spt}(\sigma^+ \cap [u, u+1])$ cannot be empty, so let $a = \min(\text{spt}(\sigma^+ \cap [u, u+1]))$. It follows that

$$0 < \int g_{u,u+1} d\sigma \leq \int g_{a,u+1} d\sigma. \quad \square$$

Theorem 5.3 (Uniqueness of the monotone martingale coupling.). *Let π be a monotone martingale transport plan and $\mu = \text{proj}_{\#}^x \pi$ and $\nu = \text{proj}_{\#}^y \pi$. Then π is the left-curtain coupling π_{1c} from μ to ν .*

Proof. Let π be left-monotone with monotonicity set Γ as in Definition 1.4 and let π_{1c} be the left-curtain transport plan between μ and ν . We consider the target measures ν_x^π and $\nu_x^{\pi_{1c}}$ obtained when transporting the μ -mass of $] - \infty, x]$ into ν , i.e.

$$\nu_x^\pi = \text{proj}_{\#}^y \pi|_{]-\infty, x] \times \mathbb{R}}$$

and

$$\nu_x^{\pi_{1c}} = S^\nu(\mu|_{]-\infty, x]) = \text{proj}_{\#}^y \pi_{1c}|_{]-\infty, x] \times \mathbb{R}}.$$

If $\nu_x^\pi = \nu_x^{\pi_{1c}}$ for every x then $\pi = \pi_{1c}$ by the definition of the curtain-coupling in Theorem 4.18.

Assume for contradiction that there exists some x with $\nu_x^\pi \neq \nu_x^{\pi_{1c}}$. This means in particular that $\sigma_x = (\nu_x^{\pi_{1c}} - \nu_x^\pi) \neq 0$. The shadow property implies that $\nu_x^{\pi_{1c}} \preceq_C \nu_x^\pi$. By Lemma 5.2 we can pick $u \in \text{spt}(\sigma_x^+)$ and $v > u$ such that

$$\int g_{u,v} d\sigma_x > 0.$$

As $u \in \text{spt} \sigma_x^+$, $\sigma_x^+ \leq \nu - \nu_x^\pi = \text{proj}_{\#}^y \pi|_{]x, +\infty[\times \mathbb{R}}$, and $\pi(\Gamma) = 1$, there is a sequence $(x'_n, u_n)_n$ such that

- $x'_n > x$
- $(x'_n, u_n) \in \Gamma$
- $u_n \rightarrow u$.

By the monotonicity property of Γ , for every $t \leq x$ and $n \in \mathbb{N}$, the set Γ_t defined by $\{y \in \mathbb{R} : (t, y) \in \Gamma\}$ cannot intersect $] - \infty, u_n[$ and $]u_n, +\infty[$. Hence for $t \leq x$,

$$(14) \quad \Gamma_t \cap] - \infty, u_n[= \emptyset \quad \text{or} \quad \Gamma_t \cap]u_n, +\infty[= \emptyset.$$

This remark will be important in the sequel of the proof.

We distinguish two cases depending on the respective positions of u and x .

(1) First case: $u < x$. Note that we have

$$\nu_x^\pi - \nu_u^\pi = \text{proj}_{\#}^y \pi|_{]u, x] \times \mathbb{R}}$$

and

$$\nu_x^{\pi_{1c}} - \nu_u^{\pi_{1c}} = \text{proj}_{\#}^y \pi_{1c}|_{]u, x] \times \mathbb{R}} = S^{\nu - \nu_u^{\pi_{1c}}}(\mu|_{]u, x]).$$

As a consequence of (14) and of the fact that π is a martingale transport plan, π transports the mass of $] - \infty, u]$ to $] - \infty, u]$ and the mass of $]u, x]$ to $]u, +\infty[$. We show below that the same applies to π_{1c} , more precisely that $\nu_u^{\pi_{1c}} \preceq_C \nu_u^\pi$ and $(\nu_x^{\pi_{1c}} - \nu_u^{\pi_{1c}}) \preceq_C (\nu_x^\pi - \nu_u^\pi)$.

- The measure $\nu_u^{\pi_{1c}}$ is the shadow of $\mu|_{]-\infty, u]}$ in ν . We have also $\mu|_{]-\infty, u]} \preceq_C \nu_u^\pi$ and $\nu_u^\pi \leq \nu$ so that $\nu_u^{\pi_{1c}} \preceq_C \nu_u^\pi$. We apply now Lemma 5.1 and obtain that $\nu_u^{\pi_{1c}}$ is concentrated on $] - \infty, u]$ and $\nu_u^{\pi_{1c}}(\{u\}) \leq \nu_u^\pi(\{u\})$.

- We have $\pi|_{]u,x] \times \mathbb{R}} \in \Pi_M(\mu|_{]u,x]}, \eta)$ where $\eta := \text{proj}_{\#}^y \pi|_{]u,x] \times \mathbb{R}} = \nu_x^\pi - \nu_u^\pi$ is concentrated on $[u, +\infty[$. More precisely we have

$$\eta \leq (\nu - \nu_u^\pi)|_{[u, +\infty[} \leq (\nu - \nu_u^{\pi_{1c}})|_{[u, +\infty[} \leq \nu - \nu_u^{\pi_{1c}}$$

because $\nu_u^{\pi_{1c}}$ and ν_u^π are concentrated on $] - \infty, x]$ and $\nu_u^{\pi_{1c}}(\{u\}) \leq \nu_u^\pi(\{u\})$ as we have seen above. Moreover we have $\mu|_{]u,x]} \preceq_C \eta$. Hence

$$\nu_x^{\pi_{1c}} - \nu_u^{\pi_{1c}} = S^{\nu - \nu_u^{\pi_{1c}}}(\mu|_{]u,x]}) \preceq_C \eta = \nu_x^\pi - \nu_u^\pi.$$

Note that $g_{u,v}$ is convex on $[u, +\infty[$ so that $\int g_{u,v} d(\nu_x^{\pi_{1c}} - \nu_u^{\pi_{1c}}) \leq \int g_{u,v} d(\nu_x^\pi - \nu_u^\pi)$. Moreover we have $\int g_{u,v} d\nu_u^{\pi_{1c}} \leq \int g_{u,v} d\nu_u^\pi$ because $\nu_u^{\pi_{1c}}(\{u\}) \leq \nu_u^\pi(\{u\})$. Summing these inequalities we obtain $\int g_{u,v} d\nu_x^{\pi_{1c}} \leq \int g_{u,v} d\nu_x^\pi$, which is a contradiction to $\int g_{u,v} d\sigma_x > 0$.

- (2) Second case: $x \leq u$. The measure π cannot transport mass from $] - \infty, x]$ to $]u, +\infty[$. Indeed because of the martingale property it then would also transport mass to the set $] - \infty, u[$, contradicting (14). Thus ν_x^π is concentrated on $] - \infty, u]$. But we have $\nu_x^{\pi_{1c}} \preceq_C \nu_x^\pi$ so that considering Lemma 5.1, $\int g_{u,v} d\nu_x^{\pi_{1c}} \leq \int g_{u,v} d\nu_x^\pi$ holds (even in the third case of this lemma where $a = u$). This contradicts $\int g_{u,v} d\nu_x^\pi > 0$.

□

Remark 5.4. *The two cases in the proof are actually not very different. In both of them, $\pi|_{]-\infty, x] \times \mathbb{R}}$ and $\pi_{1c}|_{]-\infty, x] \times \mathbb{R}}$ (roughly speaking the transport plans restricted to $\mu|_{]-\infty, x]}$) are concentrated on*

$$([-\infty, u] \times]-\infty, u]) \cup (]u, +\infty[\times [u, +\infty[)$$

and this lies at the core of the argument.

5.1. Structure of the monotone martingale coupling. It remains to establish Corollary 1.6 which states that if μ is continuous, then π_{1c} is concentrated on the graph of two functions. We need the following lemma.

Lemma 5.5. *Assume that $\Gamma \subseteq \mathbb{R}^2$ is a Borel set such that for each $x \in \mathbb{R}$ we have $|\Gamma_x| \leq 2$. Then $S = \text{proj}_x(\Gamma)$ is a Borel set and there exist Borel functions $T_1, T_2 : S \rightarrow \mathbb{R}$ with $T_1 \leq T_2$ such that*

$$\Gamma = \text{graph}(T_1) \cup \text{graph}(T_2).$$

Proof. This is a consequence of [18, Theorem 18.11]. □

We can now complete the proof.

Proof of Corollary 1.6. Consider the left-curtain coupling π_{1c} between measures $\mu \preceq_C \nu$, where μ is continuous. As π_{1c} is left-monotone there exists a Borel monotonicity set Γ as in Definition 1.4. Note that if $\mu(A) = 0$, the set $\Gamma \setminus (A \times \mathbb{R})$ is still a monotonicity set. This applies in particular to all countable sets since μ is continuous.

With the notations of Lemma 3.2 let us show that $A = \{x \in \mathbb{R} : |\Gamma_x| \geq 3\}$ is countable. If not, we can apply this lemma and obtain $x \in \mathbb{R}$ with three points $y^- < y < y^+$ in the set Γ_x that can be approximated from the right side. In particular there exists $(x', y') \in \Gamma$ with $x' > x$ and $y' \in]y^-, y^+[$, which is the forbidden configuration (3). Therefore A is countable so that we can assume that $|\Gamma_x| \leq 2$ for every x . Applying Lemma 5.5 we obtain the desired assertion. □

The following lemma permits to obtain uniqueness of the optimal martingale transport plan, provided that we know that every optimal martingale transport is concentrated on the graphs of two mappings (see Section 7). We can apply it to

the martingale transport plans when μ is continuous and recover the uniqueness of the monotone transport plan in this particular case.

Lemma 5.6. *Let μ and ν be in convex order and \mathcal{E} a non-empty convex set of martingale transport plans. Assume that every $\pi \in \mathcal{E}$ is concentrated on some $\Gamma^\pi \subseteq \mathbb{R}^2$ with $|\Gamma_x^\pi| \leq 2$ for every $x \in \mathbb{R}$. Then the set \mathcal{E} consists of a single point.*

Proof. Let π and π' be elements of \mathcal{E} . We consider $\bar{\pi} = \frac{\pi + \pi'}{2} \in \mathcal{E}$ and $\Gamma^{\bar{\pi}}$, which can be seen as the graph of two functions according to Lemma 5.5. The measures π and π' are also concentrated on $\Gamma^{\bar{\pi}}$. For two disintegrations $(\pi_x)_{x \in \mathbb{R}}$ and $(\pi'_x)_{x \in \mathbb{R}}$ with respect to μ , we know that μ -a.s. π_x and π'_x are probability measures concentrated on $\Gamma_x^{\bar{\pi}}$ and with the same barycenter, namely x . It follows that $\pi'_x = \pi_x$, μ -a.s. so that $\pi' = \pi$. \square

6. OPTIMALITY PROPERTIES OF THE MONOTONE MARTINGALE TRANSPORT

In this section we prove that π_{1c} is the unique optimal coupling for the martingale optimal transport problem (2) associated to two different kinds of cost functions. The special case $c(x, y) = \exp(y - x)$ is in the intersection of these two families of cost functions.

Theorem 6.1. *Assume that $c(x, y) = h(y - x)$ for some differentiable function h whose derivative is strictly convex and that c satisfies the sufficient integrability condition. If there exists a finite martingale transport plan, then π_{1c} is the unique optimizer.*

Proof. We have to show that every finite optimizer π is monotone. Pick a set Γ such that $\pi(\Gamma) = 1$ and Γ resists improvements by barycenter preserving reroutings as in Lemma 1.11. Pick $(x, y^-), (x, y^+), (x', y') \in \Gamma$. Striving for a contradiction we assume that they satisfy (3). Let us define a transport α on these edges and a competitor α' of it. We pick $\lambda \in]0, 1[$ such that $\lambda y^+ + (1 - \lambda)y^- = y'$. The measure α puts mass λ on (x, y^+) , mass $1 - \lambda$ on (x, y^-) and mass 1 on (x', y') . Our candidate for α' will assert mass $1 - \lambda$ on (x', y^-) , mass λ on (x', y^+) and mass 1 on (x, y') . Clearly α' is a competitor of α . It leads to smaller costs if and only if

$$\lambda c(x, y^+) + (1 - \lambda)c(x, y^-) + c(x', y') > \lambda c(x', y^+) + (1 - \lambda)c(x', y^-) + c(x, y').$$

A sufficient condition for this is that

$$(15) \quad d(t) := \lambda c(t, y^+) + (1 - \lambda)c(t, y^-) - c(t, y')$$

is strictly decreasing in x . In terms of h the function d can be written as

$$d(t) = \lambda h(y^+ - t) + (1 - \lambda)h(y^- - t) - h(y' - t).$$

To have it decreasing it is sufficient that

$$\begin{aligned} 0 > d'(t) &= -\lambda h'(y^+ - t) - (1 - \lambda)h'(y^- - t) + h'(y' - t) \\ &= h'(\lambda(y^+ - t) + (1 - \lambda)(y^- - t)) - [\lambda h'(y^+ - t) + (1 - \lambda)h'(y^- - t)]. \end{aligned}$$

Finally it is sufficient to know that h' is strictly convex which holds by assumption. \square

Remark 6.2. *The left-curtain transport plan is also a solution to the problem of minimizing the essential supremum of $y - x$ among all martingale transport plans with the same marginals. To see this, note that the function $h_n : x \mapsto \exp(nx)$ has a strictly convex derivative for every $n > 0$ and that $\frac{1}{n} \ln(\int \exp(n(y - x)) d\pi(x, y))$ tends to $\text{essup}_\pi(y - x)$ as $n \rightarrow +\infty$ for every martingale transport plan π .⁸*

⁸We thank Fillipo Santambrogio for pointing this out to us.

We mention another class of cost functions for which the monotone martingale transport plan π_{1c} is optimal.

Theorem 6.3. *Let ψ be a non-negative strictly convex function and φ a non-negative decreasing function. Consider the cost function $c(x, y) = \varphi(x)\psi(y) \geq 0$. For two finite measures μ and ν in convex order the left-curtain coupling π_{1c} is the unique optimal transport.*

One could show that optimal martingale couplings are monotone in a very similar way as in the proof of Theorem 6.1. We prefer to give an alternative proof relying on the order properties of the left-curtain coupling.

Proof. Let π be optimal for the problem and assume that $\int c d\pi < +\infty$. We want to prove $\int c d\pi_{1c} \leq \int c d\pi$ with equality if and only if $\pi = \pi_{1c}$. First of all note that for positive measurable functions f

$$\int f(x) \varphi(x) d\mu(x) = \int_0^{+\infty} \left(\int \mathbb{1}_{] -\infty, \varphi^{-1}(t)]} f(x) d\mu(x) \right) dt$$

where $\varphi^{-1}(t)$ means $\sup\{x \in \mathbb{R} : t \leq \varphi(x)\}$. Taking $f(x) = \int \psi(y) d\pi_x(y)$ we obtain

$$(16) \quad \int c(x, y) d\pi(x, y) = \int_0^{+\infty} \left(\int \psi(y) d\nu^\pi|_{\varphi^{-1}(t)}(y) \right) dt$$

where ν_u^π denotes $\text{proj}_{\#}^y \pi|_{]-\infty, u]}$ as in the introduction or in Section 5. In particular $\nu_u^{\pi_{1c}}$ equals $S^\nu(\mu|_{]-\infty, u]})$. Of course the representation (16) remains true if we replace all occurrences of π by π_{1c} .

The measures $\nu_u^{\pi_{1c}}$ and ν_u^π are in convex order and ψ is strictly convex. Thus $\int \psi d\nu_u^{\pi_{1c}} \leq \int \psi d\nu_u^\pi$ and equality holds if and only if the two measures coincide. This follows from Strassen's theorem (Theorem 2.6) and the equality case in Jensen's inequality. Finally it follows from (16) that π is the left-curtain coupling. \square

7. OTHER COST FUNCTIONS – OTHER OPTIMAL MARTINGALE COUPLINGS

In this section we use Lemma 1.11 to derive results that appeal to general cost functions.

7.1. Cost functions of the form $c(x, y) = h(y - x)$.

Theorem 7.1. *Assume that the cost function $c(x, y)$ is given by $h(y - x)$ for some function h which is twice continuously differentiable. If affine functions $x \mapsto ax + b$ meet $h'(x)$ in at most k points and π is an optimal transport plan, then there exists a disintegration $(\pi_x)_{x \in \mathbb{R}}$ such that for any $x \in \mathbb{R}$ at least one of the two following statements holds*

$$\mu(\{x\}) > 0 \quad \text{or} \quad \text{card}(\text{spt}(\pi_x)) \leq k.$$

In particular if μ is continuous then $\text{card}(\text{spt}(\pi_x)) \leq k$ is satisfied μ -a.s. for any disintegration of π .

Proof. Let π be optimal and Γ according to Lemma 1.11. If there are only countably many continuity points of μ such that $\text{card}(\Gamma_x) \geq k + 1$, then we can remove them. Assume for contradiction that there are uncountably many. Consider the set

$$\tilde{\Gamma} = \{(x, y) \in \Gamma : \mu(\{x\}) = 0\}$$

to obtain $a \in \mathbb{R}$ and $b_0 < \dots < b_k \in \Gamma_a$ verifying the assertions of Lemma 3.2.

Let $a' \in \mathbb{R}$, $\lambda \in]0, 1[$ and set $b_\lambda = (1 - \lambda)b_0 + \lambda b_k$. We will compare

$$(17) \quad h(b_\lambda - a) + \lambda h(b_k - a') + (1 - \lambda)h(b_0 - a')$$

and

$$(18) \quad h(b_\lambda - a') + \lambda h(b_k - a) + (1 - \lambda)h(b_0 - a).$$

As a' tends to a , $b_i - a'$ tends to $b_i - a$. Considering a Taylor expansion of h at $b_i - a$ we find some $\varepsilon > 0$ such that $|a - a'| < \varepsilon$ implies

$$|[h(b_i - a') - h(b_i - a)] - h'(b_i - a) \cdot (a - a')| \leq |h''(b_i - a)|(a - a')^2$$

for $i \in \{0, \lambda, k\}$. Hence if we subtract (17) from (18) we obtain

$$(19) \quad (h'(b_\lambda - a) - [(1 - \lambda)h'(b_0 - a) + \lambda h'(b_k - a)])(a' - a)$$

up to an error of

$$[(1 - \lambda)|h''(b_0 - a)| + \lambda|h''(b_k - a)| + |h''(b_\lambda - a)|] \cdot (a - a')^2.$$

But h' is not linear so that (19) is not identically zero. Moreover according to the assumption on h' and the affine functions there is an index $i \in \{1, \dots, k - 1\}$ such that if $b_\lambda = b_i$ and $a' \neq a$ then (19) is not zero. More precisely as h'' is continuous there exists some $\varepsilon_1 < \varepsilon$ such that if $|b_i - b_\lambda| < \varepsilon_1$ and $0 < |a - a'| < \varepsilon_1$ then the difference of (17) and (18) is not zero and its sign is determined by the one of $a - a'$.

Since a, b_0, \dots, b_k were chosen according to Lemma 3.2, we may pick a' and $b_\lambda \in \Gamma_{a'}$ such that (a', b_λ) is sufficiently close to (a, b_i) and a' is on the correct side of a , making (17) smaller than (18).

Setting

$$\begin{aligned} \alpha &= \lambda \delta_{(a, b_k)} + (1 - \lambda) \delta_{(a, b_0)} + \delta_{(a', b_\lambda)}, \\ \alpha' &= \lambda \delta_{(a', b_k)} + (1 - \lambda) \delta_{(a', b_0)} + \delta_{(a, b_\lambda)}, \end{aligned}$$

we have thus found a competitor α' which has lower costs than α , contradicting the choice of Γ . \square

7.2. The cost function $h(y - x)$ in the usual setup. It seems worthwhile to mention that Theorem 7.1 is the martingale variant of a result that belongs to the theory of the classical problem (1). We mention it below in Theorem 7.2 because we are not aware that it has been recorded in the literature in this form. In fact for a family of special costs we can bound the number of parts the mass can split in if it is transported optimally. Note that this number is not attained for every pair (μ, ν) (see [25]). The similarity with Theorem 7.1 lies in the fact that we want to count the number of intersection points of $\text{graph}(h')$ with affine lines in the martingale case, and with horizontal lines in the classical setup.

Theorem 7.2. *Let k be a positive integer and let $h : \mathbb{R} \rightarrow \mathbb{R}$ be a twice continuously differentiable function such that the cost function $c : (x, y) \mapsto h(y - x)$ satisfies the sufficient integrability condition with respect to probability measures μ and ν . Assume also that $C(\mu, \nu) < +\infty$.*

If the equation $h'(x) = b$ has at most k different solutions for $b \in \mathbb{R}$, then there exists a disintegration $(\pi_x)_{x \in \mathbb{R}}$ such that for any $x \in \mathbb{R}$ at least one of the two statements

$$\mu(\{x\}) > 0 \quad \text{or} \quad \text{card}(\text{spt}(\pi_x)) \leq k$$

holds. In particular if μ is continuous then $\text{card}(\text{spt}(\pi_x)) \leq k$ is satisfied μ -a.s. for any disintegration.

7.3. (Counter)examples based on the cost function $c(x, y) = (y - x)^4$. In this section we give two counterexamples that distinguish the general behavior from the one of the curtain transport plan: the optimizer is in general not unique and it may very well split into more than two parts even if the starting distribution is continuous (see Corollary 1.6 resp. Theorem 7.1). Throughout this subsection we consider the cost function $c(x, y) = (y - x)^4$.

7.3.1. Example of non uniqueness of the transport. Let μ be uniformly distributed on $\{-1; 1\}$ and ν uniformly distributed on $\{-2; 0; 2\}$. We denote -1 and 1 by $(x_i)_{i=1,2}$ and $-2, 0$ and 2 by $(y_j)_{j=1,2,3}$. To any matrix $A = (a_{i,j})$ of two rows and three columns satisfying $\sum_j a_{i,j} = 1/2$ and $\sum_i a_{i,j} = 1/3$ we associate the transport plan defined by $\pi(\{(x_i, y_j)\}) = a_{i,j}$. For such a transport plan the accumulated costs equal

$$\begin{aligned} \sum_{i,j} a_{i,j} \cdot |x_i - y_j|^4 &= (a_{1,1} + a_{1,2} + a_{2,2} + a_{2,3}) + 3^4 \cdot (a_{1,3} + a_{2,1}) \\ &= 1 + 80(a_{1,3} + a_{2,1}). \end{aligned}$$

The matrices associated to a martingale transport plan are

$$A_\lambda = \begin{pmatrix} 1/4 & 1/4 & 0 \\ 1/12 & 1/12 & 1/3 \end{pmatrix} + \lambda \begin{pmatrix} 1/12 & -1/6 & 1/12 \\ -1/12 & 1/6 & -1/12 \end{pmatrix}$$

where $\lambda \in [0, 1]$. Therefore the martingale transport plan associated to the parameter λ gives rise to total costs of $1 + 80(\lambda/12 + 1/12 - \lambda/12) = 23/3$, independently of λ . We conclude that every martingale transport plan is optimal.

7.3.2. Example of splitting in exactly three points in the continuous case. Roughly speaking we have proved in Theorem 7.1 that if μ is continuous, $d\mu(x)$ -mass elements split in at most three points. Indeed $t \mapsto t^4$ has derivative $t \mapsto 4t^3$ which is of degree 3. In this paragraph we give a numerical example showing that this upper bound is sharp. The construction is inspired by the dual theory of the martingale transport problem mentioned in Paragraph 2.4. Briefly, Figure 4 depicts a family of curves indexed by x . These curves touch three envelope curves at three moving points y_1, y_2 and y_3 close to $-1, 0$ and 1 . The optimal martingale transport plan that we construct is supported by the union of the graphs $\Gamma_i = \{(x, y_i(x)) \in \mathbb{R}^2 : x \in]0, 1/5[\}$ for $i = 1, 2, 3$. Let $\psi : \mathbb{R} \rightarrow \mathbb{R}$ be defined by




FIGURE 4. Graphs and envelope of the functions $y \mapsto F(x, y)$ for $x \in [0, 1/5]$.

$$(20) \quad \psi(y) = y^4 - \max_{x \in [0, 1/2]} \left\{ 4x \left(y + \frac{x}{2} \right) (y + 1 - x) (y - 1 - x) \right\}.$$

Hence for any $(x, y) \in [0, 1/2] \times \mathbb{R}$

$$y^4 - \psi(y) \geq 4xy^3 - 6x^2y^2 + a_1(x)y + b_1(x),$$

where $a_1(x) = 4x - 4x^2 - 4x^3$ and $b_1(x) = 2x^2 - 2x^4$. But $y^4 = (y - x)^4 + 4xy^3 - 6x^2y^2 + a_2(x)y + b_2(x)$ so that

$$(21) \quad (y - x)^4 \geq a_3(x) + b_3(x)y + \psi(y)$$

for $a_3 = a_1 - a_2$ and $b_3 = b_1 - b_2$. Here (21) is an equality at the point (x_0, y_0) if and only $\psi(y_0)$ is realized in (20) by $x = x_0$. Integrating (21) against a transport plan π one obtains

$$\iint (y - x)^4 d\pi(x, y) \geq \int a_3(x) d\mu(x) + \iint b_3(x)y d\pi(x, y) - \int \psi(y) d\nu(y)$$

and the equality holds if and only if π is concentrated on

$$\{(x, y) \in [0, 1/2] \times \mathbb{R} : (y - x)^4 = a_3(x) + b_3(x)y + \psi(y)\}.$$

Moreover as we are considering a martingale transport plan we have

$$\iint (y - x)^4 d\pi(x, y) \geq \int a_3(x) d\mu(x) + \int b_3(x)x d\mu(x) + \int \psi(y) d\nu(y).$$

Here the lower bound on the right-hand side is the same for every martingale transport plan π . It follows that martingale transport plans concentrated on $\{(x, y) \in [0, 1/2] \times \mathbb{R} : (y - x)^4 = a_3(x) + b_3(x)y + \psi(y)\}$ are optimal with respect to their marginals. We set $F(x, y) = 4x(y + \frac{x}{2})(y + 1 - x)(y - 1 - x)$ so that (20) is $\psi(y) = y^4 - \sup_{x \in [0, 1/2]} F(x, y)$. In Figure 4, one can see the graphs of $F(x, \cdot)$ for values of x between 0 and 1/5.

We will prove that for $y \in]-1, 0[\cup]1, 2[$, $F(\cdot, y) : [0, 1/2] \rightarrow \mathbb{R}$ has a unique global maximum in $]0, 1/2[$. Actually $F(\cdot, y)$ has main term $2x^4$. Therefore it is sufficient to prove that $\partial_x F(\cdot, y)$ is positive for $x = 0$ and negative for $x = 1/2$. Indeed this means that we are analyzing the variation of the polynomial function $F(\cdot, y)$ of degree 4 on an interval where its variations are different from the asymptotic ones. In particular $F(\cdot, y)$ will have a unique maximum on $]0, 1/2[$. This turns out to be true. Indeed

$$(22) \quad \partial_x F(x, y) = 4((x + y)[(x - y)^2 - 1] + x(x + 2y)(x - y)),$$

so that for any parameter y in $] - 1, 0[\cup]1, 2[$, the function $\partial_x F(\cdot, y)$ is positive in $x = 0$ since it equals $y \mapsto 4(y(y^2 - 1))$. For $x = 1/2$, straightforward considerations show that $\partial_x F(1/2, y)$ is negative for all $y \in] - \infty, 2[$.

We will now show that for a given parameter $x \in]0, 1/5[$, x is the maximum of $F(\cdot, y)$ on $[0, 1/2]$ for exactly three elements y of $] - 1, 0[\cup]1, 2[$. For this purpose we consider $y \mapsto \partial_x F(x, y)$. We prove that it vanishes exactly three times on $] - 1, 0[\cup]1, 2[$. For fixed $x \in]0, 1/5[$ this function is indeed negative in 0 and -1 while it is positive in $-1/2$. The sign is also different for $y = 1$ and $y = 2$ so that we have found the three zeros of $y \mapsto \partial_x F(x, y)$. But as explained in the previous step, for $y \in] - 1, 0[\cup]1, 2[$ being a maximum of $F(\cdot, y)$ is exactly the same as having zero derivate.

Therefore any $x \in]0, 1/5[$ gives rise to the maximum of $F(\cdot, y)$ for three different $y \in] - 1, 0[\cup]1, 2[$. Hence there are y_1, y_2, y_3 such that $\psi(y_i) = y_i^4 - F(x, y_i)$ for $i = 1, 2, 3$. Notice that x is in the convex hull of these points because y_1 is close to -1 , y_2 is close to 0 and y_3 close to 1. Hence there exists a martingale transport plan π concentrated on $[0, 1/5] \times ([-1, 0] \cup [1, 2])$ such that π_x is supported on $\{y_1, y_2, y_3\}(x)$ with positive μ -probability. Moreover it follows from the explanations above that

this martingale transport plan is optimal. Namely (20) holds π -a.s. Hence we have proved that the bound $k = 3$ of Theorem 7.1 is sharp in the case $c(x, y) = (y - x)^4$.

7.4. The Hobson-Neuberger cost function and its converse. As mentioned in the introduction, Hobson and Neuberger [16] study the case $c(x, y) = -|y - x|$, motivated by applications in mathematical finance. They identify the minimizer π_{HN} based on a construction of the maximizers for the dual problem. Here some conditions on the underlying measures are necessary; an example in [2, Proposition 5.2] shows that the dual maximizers need not always exist. Based on Lemma 1.11 we partly recover their result. Throughout this part we will only deal with the case of a continuous starting distribution μ (see Remark 7.6 on this hypothesis).

Theorem 7.3. *Assume that μ and ν are in convex order and that μ is continuous. There exists a unique optimal martingale transport plan π_{HN} for the cost function $c(x, y) = -|y - x|$.*

Moreover, there exist two non-decreasing functions $T_1, T_2 : \mathbb{R} \rightarrow \mathbb{R}$ such that $T_1(x) \leq x \leq T_2(x)$ and π_{HN} is concentrated on the graphs of these functions.

A similar behavior holds for the cost function $c(x, y) = |y - x|$ built on the absolute value $h : x \mapsto |x|$. We have learned about the structure of the optimizer for this cost function from D. Hobson and M. Klimmek [15]. Recall that $\Gamma_x = \{y : (x, y) \in \Gamma\}$ for $\Gamma \subseteq \mathbb{R}^2$.

Theorem 7.4. *Assume that μ and ν are in convex order and that μ is continuous. There exists a unique optimal martingale transport plan π_{abs} for the cost function $c(x, y) = |y - x|$.*

Moreover there is a set Γ such that π_{abs} is concentrated on Γ and $|\Gamma_x| \leq 3$ for every $x \in \mathbb{R}$. More precisely, π_{abs} can be decomposed into $\pi_{\text{stay}} + \pi_{\text{go}}$ where $\pi_{\text{stay}} = (\text{Id} \otimes \text{Id})_{\#}(\mu \wedge \nu)$ (this measure is concentrated on the diagonal of \mathbb{R}^2) and π_{go} is concentrated on $\text{graph}(T_1) \cup \text{graph}(T_2)$ where T_1, T_2 are real functions.

The ‘‘combinatorial core’’ of the proofs to Theorem 7.3 and Theorem 7.4 is contained in the following lengthy but simple lemma.

Lemma 7.5. *Let $x, y^-, y^+, y' \in \mathbb{R}$ such that $y^- < x, y' < y^+$. Pick λ such that $\lambda y^+ + (1 - \lambda)y^- = y'$. For $x' \in \mathbb{R}$ we want to compare the quantities*

$$A := \lambda|x - y^+| + (1 - \lambda)|x - y^-| + |x' - y'|, \quad B := \lambda|x' - y^+| + (1 - \lambda)|x' - y^-| + |x - y'|.$$

- (1) *Assume that $y' < x$. Then there exists $x_0 \in]y^-, y'[$ such that $(A - B)$ seen as a function of x' exactly vanishes at x_0 and x , is strictly positive outside $[x_0, x]$ and strictly negative in $]x_0, x[$.*

$$\frac{x'}{(A - B)(x')} \left| \begin{array}{cccccc} -\infty & y^- & x_0 & y' & x & +\infty \\ + & 0 & - & 0 & + & \end{array} \right.$$

- (2) *Assume that $y' > x$. Then there exists $x_1 \in]y', y^+[$ such that $(A - B)$ vanishes if $x' \in \{x_1, x\}$, is strictly positive outside $[x, x_1]$ and strictly negative in $]x, x_1[$.*

$$\frac{x'}{(A - B)(x')} \left| \begin{array}{cccccc} -\infty & x & y' & x_1 & y^+ & +\infty \\ + & 0 & - & 0 & + & \end{array} \right.$$

- (3) *Assume that $y' = x$. Then $(A - B)$ is non-negative and vanishes exactly in x .*

$$\frac{x'}{(A - B)(x')} \left| \begin{array}{cccccc} -\infty & y^- & x = y' & y^+ & +\infty \\ + & 0 & + & + & \end{array} \right.$$

Proof. Consider the function

$$f(t) = \lambda|t - y^+| + (1 - \lambda)|t - y^-| - |t - y'|.$$

Then $A > B$ is equivalent to $f(x) > f(x')$ and $A = B$ is equivalent to $f(x) = f(x')$.

The behavior of the function f is easy enough to understand. On the intervals $]-\infty, y^-]$, $[y^+, \infty[$, the function is zero. On the interval $[y^-, y']$ it increases linearly from 0 to $2\lambda(1 - \lambda)(y^+ - y^-)$. On the interval $[y', y^+]$ it decreases linearly from $2\lambda(1 - \lambda)(y^+ - y^-)$ to 0.

The above assertions are simple consequences of this behavior. Moreover it is easy to calculate x_0, x_1 explicitly. For instance in the case $y' < x$ pick $t \in]0, 1[$ such that $x = y' + t(y^+ - y')$. Then $x_0 = y' + t(y^- - y')$. \square

Proof of Theorem 7.3. Pick Γ according to Lemma 1.11 and $(x, y^-), (x, y^+), (x', y') \in \Gamma$, with $y^- < y' < y^+$. Then it cannot happen that

$$(23) \quad y' \leq x' < x \quad \text{or} \quad x < x' \leq y'.$$

Indeed choosing $\lambda \in]0, 1[$ and α resp. α' as in the proof of Theorem 6.1 we find that an improvement is possible if

$$-\lambda|x - y^+| - (1 - \lambda)|x - y^-| - |x' - y'| > -\lambda|x' - y^+| - (1 - \lambda)|x' - y^-| - |x - y'|.$$

But this inequality holds in the just mentioned cases by Lemma 7.5.

Consider the set A of points a such that Γ_a contains more than two points and assume by contradiction that this set is uncountable. According to Lemma 3.2 there is an accumulation effect at some $a \in A$ together with $b^-, b, b^+ \in \Gamma_a$ in the order $b^- < b < b^+$. (Without loss of generality one may assume $b \leq a$.) In particular Lemma 3.2 provides $(a_0, b_0^-), (a_0, b_0^+) \in \Gamma$ such that $a < a_0 < b_0^+$ and $b_0^- < b$. We have settled the first forbidden situation of (23) for $(x, y^-) = (a_0, b_0^-)$, $(x, y^+) = (a_0, b_0^+)$ and $(x', y') = (a, b)$, which provides the desired contradiction. Hence A is countable and $\mu(A) = 0$. It follows that one can assume $|\Gamma_a| \leq 2$ for every $a \in \mathbb{R}$.

We may thus assume that there exist T_1 and T_2 from $\text{proj}^x(\Gamma)$ to \mathbb{R} such that $\Gamma_x = \{T_1(x), T_2(x)\}$ where $T_1(x) \leq x \leq T_2(x)$ for μ -almost every $x \in \text{proj}^x(\Gamma)$. It remains to show that T_1 and T_2 are monotone. Let $x, x' \in \mathbb{R}$ with $x < x'$. We necessarily have $T_2(x) \leq T_2(x')$ since the opposite inequality leads to the second forbidden inequality in (23) taking $y^- = T_1(x), y' = T_2(x')$ and $y^+ = T_2(x)$. The monotonicity of T_1 is established in the same way.

It remains to show that the optimizer is unique. Due to the linear structure of the optimization problem the set of solutions is convex. Hence Lemma 5.6 applies. \square

Remark 7.6. *If μ is not continuous, there may be more than one minimizer. This is the case for example if μ and ν are chosen as in Paragraph 7.3.1. In fact if h is an even function then for the cost function $c(x, y) = h(y - x)$ (e.g. $x \mapsto -|y - x|$) every martingale transport plan is optimal. Hence it seems that it is not directly possible to define the Hobson–Neuberger transport plan for a general starting distribution μ in an unambiguous way.*

Proof of Theorem 7.4. Let π be an optimal martingale transport plan. Pick Γ according to Lemma 1.11 and $(x, y^-), (x, y^+), (x', y') \in \Gamma$, with $y^- < y' < y^+$. Then it cannot happen that

$$(24) \quad x' < x \leq y' \quad \text{or} \quad y' \leq x < x' \quad \text{or} \quad x' \notin [y^-, y^+].$$

Indeed choosing $\lambda \in]0, 1[$, α and α' as in the proof of Theorem 6.1 above we find that an improvement of α by α' is possible if

$$\lambda|x - y^+| + (1 - \lambda)|x - y^-| + |x' - y'| > \lambda|x' - y^+| + (1 - \lambda)|x' - y^-| + |x - y'|.$$

Indeed this inequality holds in the just mentioned cases by Lemma 7.5. Note in particular that one of the forbidden cases of (24) occurs if $x \neq x'$ and $x = y'$. This will be crucial in the following argument which establishes that as much mass as possible is transported by the identity mapping. (Roughly speaking the following is forbidden: some mass goes from x to y^- and y^+ while some mass goes from x' to $y' = x$.)

Set $\pi_0 = \pi|_{\Delta}$, where Δ is the diagonal $\{(x, y) \in \mathbb{R}^2 : x = y\}$ and $\bar{\pi} = \pi - \pi_0$, let ρ be the projection of π_0 onto the first (or the second) coordinate. As $\rho \leq \mu$ and $\rho \leq \nu$, we have $\rho \leq \mu \wedge \nu$. We want to prove that $\rho = \mu \wedge \nu$, i.e. π_0 is $(\text{Id} \otimes \text{Id})_{\#}(\mu \wedge \nu)$. Let us define the reduced measures $\bar{\mu} = \mu - \rho$, $\bar{\nu} = \nu - \rho$ and $\kappa = \mu \wedge \nu - \rho$. Note that $\bar{\pi} \in \Pi_M(\bar{\mu}, \bar{\nu})$ and that $\bar{\pi}$ is concentrated on $\bar{\Gamma} = \Gamma \setminus \Delta$. Hence we have the following:

- For $\bar{\mu}$ -almost every a there exist b^- and b^+ such that $a \in]b^-, b^+[$ and $(a, b^-), (a, b^+) \in \bar{\Gamma}$.
- For κ -almost every b there exists some $a \neq b$ such that $(a, b) \in \bar{\Gamma}$.

As $\kappa \leq \bar{\mu}$ we conclude that κ -almost every real number satisfies both of these conditions. Thus for κ -almost every x there exist y^-, y^+ and x' such that the points (x, y^-) , (x, y^+) and (x', x) are included in $\bar{\Gamma}$ and one has $x' \neq x$ and $x \in]y^-, y^+[$. This coincides with one of the forbidden situations of (24). Hence κ has mass 0 and $\pi_0 = (\text{Id} \otimes \text{Id})_{\#}(\mu \wedge \nu)$ as claimed above.

Our next goal is to establish that, removing countably many points if necessary, we have $|\bar{\Gamma}_x| \leq 2$ for every $x \in \mathbb{R}$. Indeed if this is not true, then there exist a, b', b^- and b^+ with $b^- < b < b^+ \in \bar{\Gamma}_a$ to which the assertion of Lemma 3.2 applies. We know that $b < a$ or $a < b$; assume without loss of generality that $a < b$. But then there exist a' with $b^- < a' < a$ and b' with $a < b' < b$ such that $(a', b') \in \Gamma$. This contradicts (24) (with $x = a, y^- = b^-, y^+ = b^+, x' = a', y' = b'$).

It remains to establish that there exists at most one optimizer. For optimal transports π the static part $\pi_0 = \pi|_{\Delta}$ equals $(\text{Id} \otimes \text{Id})_{\#}(\mu \wedge \nu)$. Hence the reduced measure $\bar{\pi} = \pi - \pi_0$ is a minimizer of the martingale transport problem between $\bar{\mu} = \mu - \mu \wedge \nu$ and $\bar{\nu} = \nu - \mu \wedge \nu$. Note that $\bar{\mu} \wedge \bar{\nu} = 0$ so that the optimal martingale couplings are concentrated on two Borel graphs. We conclude by Lemma 5.6. \square

Remark 7.7. *Exactly as in Remark 7.6, the hypothesis that μ is continuous is needed to prove uniqueness of the optimizer; π_{abs} is not well-defined otherwise.*

8. APPENDIX A: A CONVERSE TO THE VARIATIONAL LEMMA

In this section we prove that the optimality criterion given in the Variational Lemma 1.11 is not only necessary but also *sufficient* provided that the cost function is assumed to be bounded and continuous. We conjecture that these regularity assumptions can be relaxed. Before we state the Variational Lemma let us give a definition.

Definition 8.1. *Let c be cost function with values in \mathbb{R} . We say that a Borel set Γ is finitely optimal for c if for every measure α on $\mathbb{R} \times \mathbb{R}$ with $|\text{spt}(\alpha)| < \infty$ and $\text{spt}(\alpha) \subseteq \Gamma$ and every competitor α' of α we have $\int c d\alpha \leq \int c d\alpha'$.*

As c only takes finite values the integrals exist.

Lemma 8.2 (Variational Lemma, Part II). *Assume that $\mu, \nu \in \mathcal{P}$ are in convex order and that $c : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a continuous bounded cost function. Let $\pi \in \Pi_M(\mu, \nu)$. It there exists a finitely optimal set Γ such that $\pi(\Gamma) = 1$, then π is an optimal martingale transport plan.*

The strategy of our proof will be to establish *dual maximizers* (see Paragraph 2.4). Such dual maximizers do not exist in general as follows from [2, Proposition

4.1]. However, the following simple lemma allows us to reduce the martingale transport problem to “irreducible components”. It turns out that on each of these components it is possible to construct the desired dual maximizers.⁹

8.1. Irreducible decompositions. Let us now introduce some of the necessary vocabulary.

Definition 8.3. Let μ, ν be elements of \mathcal{M} such that $\mu \preceq_C \nu$. We say that (μ, ν) is irreducible if there exists an open interval I (bounded or not) such that $\mu(I)$ and $\nu(\bar{I})$ have the total mass and $u_\mu < u_\nu$ on I .

Note that on $\mathbb{R} \setminus I$ we have $u_\mu = u_\nu$ so that I is exactly $\{u_\mu < u_\nu\}$.

Theorem 8.4 (Decomposition of (μ, ν) into irreducible components). Let μ, ν be elements of \mathcal{M} such that $\mu \preceq_C \nu$. Let $(I_k)_k$ be the (in essence unique) sequence of disjoint open intervals such that $\bigcup_k I_k = \{u_\mu < u_\nu\}$ and write F for the closed set $\mathbb{R} \setminus \bigcup_k I_k$. Set $\mu_k = \mu|_{I_k}$ and define $\eta = \mu|_F$ such that $\mu = (\sum_k \mu_k) + \eta$.

There exists a unique decomposition $\nu = (\sum_k \nu_k) + \nu$ such that $\mu_k \preceq_C \nu_k$ for each k and $\eta \preceq_C \nu$.

For this decomposition $\eta = \nu$ and (μ_k, ν_k) is irreducible with $\{u_{\mu_k} < u_{\nu_k}\} = I_k$. Moreover any martingale transport plan $\pi \in \Pi_M(\mu, \nu)$ can be decomposed in the form

$$(25) \quad \pi = \left(\sum_k \pi_k \right) + \pi_F$$

where π_k is a martingale transport from μ_k to ν_k . This decomposition is unique and $\pi_F = (\text{Id} \otimes \text{Id})_{\#} \eta$.

Note that the measure $\eta \wedge \nu_k$ does not necessarily vanish.

Proof. To establish the uniqueness part we need two auxiliary results:

Lemma 8.5. Assume that μ, ν are elements of \mathcal{P} and let $\pi \in \Pi_M(\mu, \nu), s \in \mathbb{R}$. The following are equivalent.

- (i) $\pi(\cdot - \infty, s[\cdot \times] - \infty, s] \cup \{(s, s)\} \cup]s, \infty[\times[s, \infty[) = 1$.
- (ii) $u_\mu(s) = u_\nu(s)$.

Consequently, as (ii) does not depend on π , if (i) holds for one measure in $\Pi_M(\mu, \nu)$, then it applies to all elements of $\Pi_M(\mu, \nu)$.

Proof. This is essentially [2, Lemma 4.2]; the only difference is that the formulation in [2] refers to the function $u_\mu^+(x) := \int (y - x)_+ d\mu(y)$ rather than to u_μ . However, the proof goes through in the same way if $(\cdot)_+$ is replaced by $|\cdot|$. \square

We record the following consequence.

Lemma 8.6. Let I be an open interval such that $u_\mu = u_\nu$ on the boundary of I . Let μ_I be $\mu|_I$ and π be a transport plan of $\Pi_M(\mu, \nu)$. Set also $\nu_I := \text{proj}_{\#}^y(\pi|_{I \times \mathbb{R}})$.

The measure ν_I is concentrated on \bar{I} and does not actually depend on the particular choice of π . Moreover we have $u_{\nu_I} - u_{\mu_I} = 0$ on $\mathbb{R} \setminus I$ and $u_{\nu_I} - u_{\mu_I} = u_\nu - u_\mu$ on I .

Proof. Pick $\pi \in \Pi_M(\mu, \nu)$ and apply Lemma 8.5 to every $s \in \partial I$. Then

$$(26) \quad \pi((I \times \bar{I}) \cup (\mathbb{R} \setminus I)^2) = 1.$$

⁹Roughly speaking the construction given in [2, Proposition 4.1] uses an infinite number of such irreducible components. While it is possible to construct optimizers on each component, it turns out to be impossible to glue them together.

Set $\pi_I := \pi|_{I \times \mathbb{R}}$. Relation (26) asserts that no mass of μ is moved from $\mathbb{R} \setminus I$ to I and that the mass of I is transported into \bar{I} . Thus $\mu_I \preceq_C \nu_I = \text{proj}_{\#}^y \pi_I$ (so that the two measures have the same integral against linear functions) and ν_I is concentrated on \bar{I} . It follows directly from the definition of the potential functions that $u_{\nu_I} = u_{\mu_I}$ on $\mathbb{R} \setminus I$. Applying similar arguments to $\mu|_J$ and $\nu_J = \text{proj}_{\#}^y \pi|_{J \times \mathbb{R}}$ for every (closed) connected component J of $\mathbb{R} \setminus I$ and recalling that $\alpha \mapsto u_\alpha$ is linear we obtain $u_{\mu - \mu_I} = u_{\nu - \nu_I}$ on I . Hence $u_{\nu_I} - u_{\mu_I} = u_\nu - u_\mu$ holds on this interval. \square

We first prove the existence of some decomposition of ν . We fix some $\pi \in \Pi_M(\mu, \nu)$ and for every k , we define μ_k and ν_k as the marginals of $\pi_k := \pi|_{I_k \times \mathbb{R}}$. Denote by η, v the marginals of $\pi_F := \pi|_{F \times \mathbb{R}}$. The transport plans π_k and π_F are martingale transport plans so that $\mu_k \preceq_C \nu_k$ and $\eta \preceq_C v$.

For the uniqueness part, we take for $i = 1, 2$ a decomposition $(\nu_k^i)_k, v^i$ of ν such that $\mu_k \preceq_C \nu_k^i$ and $\eta \preceq_C v^i$. According to Example 2.3 there exists a martingale transport plan π^i that transports every μ_k on ν_k^i and η on v^i . But the μ_k 's are concentrated on disjoint intervals so that $\nu_k^i = \text{proj}_{\#}^y \pi^i|_{I_k \times \mathbb{R}}$ and $v^i = \text{proj}_{\#}^y \pi^i|_{F \times \mathbb{R}}$. It follows from Lemma 8.6 that $\text{proj}_{\#}^y \pi|_{I_k \times \mathbb{R}}$ does not depend on the particular choice of $\pi \in \Pi_M(\mu, \nu)$. Hence $\nu_k^1 = \nu_k^2$ for every k and $v^1 = \nu - \sum_k \nu_k^1 = v^2$.

Let's now prove the properties listed in the second part of Theorem 8.4. We continue to use the notations of the existence part ($\pi, \pi_k, \pi_F, \mu_k, \nu_k, \eta$ and v). As a consequence of Lemma 8.6 (applied to μ, ν and I_k) we have the following:

- i) ν_k is concentrated on \bar{I}_k ;
- ii) $u_{\nu_k} - u_{\mu_k}$ is 0 on $\mathbb{R} \setminus I_k$ and $u_\nu - u_\mu$ on I_k .

As the I_k 's are disjoint we have

$$u_{\sum \nu_k} - u_{\sum \mu_k} = \sum_k (u_{\nu_k} - u_{\mu_k}) = \begin{cases} u_\nu - u_\mu & \text{on } \bigcup_k I_k, \\ 0 = u_\nu - u_\mu & \text{on } F = \bigcap_k \bar{I}_k. \end{cases}$$

Hence

$$u_v = u_\nu - u_{\sum \nu_k} = u_\mu - u_{\sum \mu_k} = u_\eta$$

on the whole real line. Thus we have $v = \eta$. The fact that (μ_k, ν_k) is irreducible and $\{u_{\mu_k} < u_{\nu_k}\} = I_k$ follows directly from Definition 8.3 and what has been proved so far. Finally concerning π , note that $\pi = (\sum_k \pi_k) + \pi_F$ where π_k has marginals μ_k and ν_k . As π_F is a martingale transport plan from η to $v = \eta$ it is the identical transport plan $(\text{Id} \otimes \text{Id})_{\#} \eta$. The uniqueness of the decomposition (25) follows from the fact that the μ_k 's are concentrated on disjoint intervals. \square

As a consequence of Theorem 8.4 we have the straightforward corollary:

Corollary 8.7 (Reducing the transport problem). *Let μ, ν be elements of \mathcal{M} and $\mu \preceq_C \nu, \pi \in \Pi_M(\mu, \nu)$ with decompositions $(\mu_k)_k, (\nu_k)_k, \eta, \pi = (\sum_k \pi_k) + (\text{Id} \otimes \text{Id})_{\#} \eta$ as in Theorem 8.4. Let c be a cost function such that the martingale transport problem satisfies the sufficient integrability condition and leads to finite costs. Then the transport π is optimal if and only if every π_k is optimal for the transport problem between μ_k and ν_k .*

Recall that in Lemma 8.2, the main result of this section, one is assuming that some particular finitely optimal set exists for the cost c . We will need several times to assume that this set satisfies some additional properties that we introduce in the next definition. Recall for the sequel that for a set $G \subseteq \mathbb{R}^2$ we write $G_x = \{y : (x, y) \in G\}$ and denote the projections of G by X_G and Y_G , respectively.

Definition 8.8. *Let I be an open interval. A set G satisfies the regularity property (Reg.) on I if $G \subseteq I \times \bar{I}$ and for every $x \in I$ we have $G_x = \emptyset$ or $G_x = \{x\}$ or $x \in] \inf G_x, \sup G_x[$.*

A set G satisfies the irreducibility property (Irred.) on I if $G \subseteq I \times \bar{I}$ and for every $y \in I$ there exist $x \in I$ and $y^-, y^+ \in G_x$ so that $y^- < y < y^+$.

Note that if G satisfies (Irred.) on I , we can apply this property to points $y \in I$ close to the boundary of I . Therefore we have $I = \text{conv}(Y_G)$.

Lemma 8.9. *Let μ, ν be elements of \mathcal{P} such that (μ, ν) is irreducible with $I = \{u_\mu < u_\nu\}$. Let c be a cost function. Let moreover G be a finitely optimal set and π a martingale transport plan with $\pi(G) = 1$. Then there exists a Borel set $G' \subseteq G \cap (I \times \bar{I})$ that satisfies (Reg.) and (Irred.) on I and such that $\pi(G') = 1$. Moreover G' is finitely optimal.*

Proof. Let G and π be as in the statement. Since π is a martingale transport plan we find that for μ -almost all $x \in I$

$$x \in \text{conv}(G_x) \quad \text{or} \quad \{x\} = G_x.$$

Erasing a negligible set if necessary, we can assume that (Reg.) is satisfied on I . Let G' be the resulting set. Assume by contradiction that G' does not satisfy (Irred.) on I . Hence there exists $y \in I$ such that for every $x \in I$, the set G_x is included in $] -\infty, y]$ or in $[y, +\infty[$. Since (Reg.) holds, $G_x \subseteq] -\infty, y]$ if $x \leq y$ and $G_x \subseteq [y, +\infty[$ otherwise. Hence $\pi(] -\infty, y]^2 \cup [y, +\infty[^2) = 1$ so that $u_\mu(y) = u_\nu(y)$, according to Lemma 8.5. But $y \in I = \{u_\mu < u_\nu\}$, which yields a contradiction. Therefore the set G' satisfies (Reg.) and (Irred.) on I . Each subset of G is finitely optimal, hence so is G' . \square

8.2. Existence of dual maximizers φ, ψ, Δ on an irreducible component.

In this paragraph we aim to prove Proposition 8.10. The cost function c , the sets $\Gamma \subseteq \mathbb{R}^2$ and I are fixed accordingly throughout Subsections 8.2 and 8.3.

Proposition 8.10. *Assume that $c : \mathbb{R} \rightarrow \mathbb{R}$ is continuous and let Γ be a finitely optimal set satisfying the conditions (Reg.) and (Irred.) on some open interval I .*

Then there exist upper semi-continuous functions $\varphi : I \rightarrow] -\infty, \infty[$, $\psi : J = \text{conv}(Y_\Gamma) \rightarrow] -\infty, \infty[$ and a measurable function $\Delta : I \rightarrow \mathbb{R}$ such that

$$\varphi(x) + \psi(y) + \Delta(x)(y - x) \leq c(x, y)$$

for all $x \in I, y \in J$, with equality holding whenever $(x, y) \in \Gamma$.

We emphasize that the functions appearing in Proposition 8.10 can be interpreted as a sort of maximizer for the dual problem described in Section 2.4.

Throughout Section (8.2) we will work under the assumptions of Proposition 8.10; some preparations will be necessary to establish the result.

Definition 8.11. *Let ψ be a function from a subset of \mathbb{R} into \mathbb{R} and let G be a subset of $\mathbb{R} \times \mathbb{R}$ such that ψ is defined on $Y_G = \text{proj}^y(G)$. The function ψ is called G -good if the following holds true:*

For every $x \in X_G = \text{proj}^x(G)$ there exists an affine function $y \mapsto a_x(y)$ such that

$$(27) \quad a_x(y) \leq -\psi(y) + c(x, y)$$

for all $y \in Y_G$ with equality holding true if $y \in G_x = \{y \in \mathbb{R} : (x, y) \in G\}$.

Note that the function a_x is uniquely determined if $|G_x| \geq 2$. Clearly, a function ψ is G -good if and only if there exist functions φ, Δ (defined on some set containing X_G) such that

$$\varphi(x) + \psi(y) + \Delta(x)(y - x) \leq c(x, y)$$

for all $x \in X_G$ and $y \in Y_G$ with equality being satisfied whenever $(x, y) \in G$.

Subsequently we will show that in Proposition 8.10 there exists a Γ -good function ψ . We want to explain already at this stage that for a given Γ -good function ψ , suitable functions φ and Δ can be defined rather explicitly in terms of the function ψ : Fix $x \in X_\Gamma$. By the property (Reg.), there exists y^-, y^+ with $y^- < x < y^+$, $(x, y^-), (x, y^+) \in \Gamma$ and a unique affine function a_x such that $a_x(y^-) = -\psi(y^-) + c(x, y^-)$ and $a_x(y^+) = -\psi(y^+) + c(x, y^+)$, moreover a_x lies below the function $y \mapsto -\psi(y) + c(x, y)$. Writing $g(\cdot)^{**}$ for the convex hull of a function $y \mapsto g(y)$, we find further that $a_x(y)$ is also smaller or equal than $(-\psi(\cdot) + c(x, \cdot))^{**}(y)$, with equality holding true for all $y \in [y^-, y^+]$. This implies that $a_x(y) = \varphi(x) + \Delta(x)(y - x)$, where

$$(28) \quad \varphi(x) := (-\psi(\cdot) + c(x, \cdot))^{**}(x),$$

and $\Delta(x)$ denotes the derivative of $y \mapsto (-\psi(\cdot) + c(x, \cdot))^{**}(y)$ at the point $y = x$.

The first step towards the existence of a Γ -good function in Proposition 8.10 is the following auxiliary result.

Lemma 8.12. *Let $G \subseteq \Gamma$ be a finite set. Then there exists a G -good function.*

Proof. As Γ satisfies (Reg.), there exists a finite set \tilde{G} , $G \subseteq \tilde{G} \subseteq \Gamma$ such that \tilde{G} satisfies (Reg.). As a consequence of (Reg.), there exists a probability measure α which has support \tilde{G} and is a martingale transport plan between its marginals, i.e. satisfies $\alpha \in \Pi_M(\mu_0, \nu_0)$ for $\mu_0 := \text{proj}_\#^x \alpha$, $\nu_0 := \text{proj}_\#^y \alpha$. As Γ is finitely optimal, every competitor of α leads at least to the same amount of costs as α , i.e. α is an optimal martingale measure. By the duality theorem of linear programming, there exist functions $\varphi, \Delta : X_{\tilde{G}} \rightarrow \mathbb{R}, \psi : Y_{\tilde{G}} \rightarrow \mathbb{R}$ such that

$$\varphi(x) + \psi(y) + \Delta(x)(y - x) \leq c(x, y)$$

for all $(x, y) \in X_{\tilde{G}} \times Y_{\tilde{G}}$ with equality holding for all elements of the set \tilde{G} . In particular ψ is a G -good function. \square

The following technical lemma will give us some control over the variety of different G -good functions which can exist for a specified set G .

Lemma 8.13. *Let $G = \{(x_i, y_i^-), (x_i, y_i^+) : i = 1, 2\}$, where $y_i^- < x_i < y_i^+$. Assume that $]y_1^-, y_1^+[\cap]y_2^-, y_2^+[\neq \emptyset$. Given bounded intervals K_1^\pm there exist bounded intervals K_2^\pm such that the following holds: if ψ is G -good and $\psi(y_1^\pm) \in K_1^\pm$, then $\psi(y_2^\pm) \in K_2^\pm$.*

Let $G = \{(x_1, y_1^-), (x_1, y_1^+), (x_2, y_2)\}$, where $y_1^- < x_1 < y_1^+$. Assume that $y_2 \in]y_1^-, y_1^+]$. Given bounded intervals K_1^\pm there exists a bounded interval K_2 such that the following holds: if ψ is G -good and $\psi(y_1^\pm) \in K_1^\pm$, then $\psi(y_2) \in K_2$.

Proof. We will only prove the first part of the lemma, the second is similar. Moreover we will assume that $y_1^- < y_2^- < y_2^+ < y_1^+$. If these numbers are ordered in a different way, the argument can be adapted easily. Since ψ is G -good, there is an affine function a_{x_1} such that

$$(29) \quad a_{x_1}(y_1^-) = -\psi(y_1^-) + c(x_1, y_1^-) \in -K_1^- + c(x_1, y_1^-),$$

$$(30) \quad a_{x_1}(y_1^+) = -\psi(y_1^+) + c(x_1, y_1^+) \in -K_1^+ + c(x_1, y_1^+),$$

$$(31) \quad a_{x_1}(y_2^-) \leq -\psi(y_2^-) + c(x_1, y_2^-),$$

$$(32) \quad a_{x_1}(y_2^+) \leq -\psi(y_2^+) + c(x_1, y_2^+).$$

From (29) and (30) we have a good control over the possible positions of the affine function a_{x_1} . By (31) and (32) this translates to a lower bounded for the value of $-\psi(y_2^-)$ (, resp. $-\psi(y_2^+)$). More precisely we obtain that there exists a real

number q which depends on $K_1^\pm, x_1, y_1^\pm, y_2^\pm$ and c (but not on the particular values of $\psi(y_2^\pm)$) such that $q \leq -\psi(y_2^\pm)$.

On the other hand there exists an affine function a_{x_2} such that

$$\begin{aligned} a_{x_2}(y_1^-) &\leq -\psi(y_1^-) + c(x_1, y_1^-) \in -K_1^- + c(x_2, y_1^-), \\ a_{x_2}(y_1^+) &\leq -\psi(y_1^+) + c(x_1, y_1^+) \in -K_1^+ + c(x_2, y_1^+), \\ a_{x_2}(y_2^-) &= -\psi(y_2^-) + c(x_2, y_2^-), \\ a_{x_2}(y_2^+) &= -\psi(y_2^+) + c(x_2, y_2^+). \end{aligned}$$

This implies the existence of a constant p such that $p \geq -\psi(y_2^\pm)$. Summing up, we may choose $K_2^+ = K_2^- = [-p, -q]$. \square

Lemma 8.14. *There exists a Γ -good function ψ .*

Proof. In Lemma 8.12 we have already seen that for every finite set $G \subseteq \Gamma$ there exists a G -good function. The idea of the proof is thus to pass to some sort of limit of these functions. To do so we aim to confine (properly chosen) G -good functions to a compact subset of the space \mathbb{R}^{Y_G} . The existence of this compact set will be a consequence of Lemma 8.13 and Tychonoff's Theorem.

We claim that there exist compact intervals $(K_y)_{y \in Y_\Gamma}$ such that for any finite set $G \subseteq \Gamma$ there is a G -good function ψ such that $\psi(y) \in K_y$ for $y \in Y_G$.

We give the proof under the assumption that $Y_\Gamma \subseteq I$ is such that $\text{conv}(Y_\Gamma)$ is open (such that $\text{conv}(Y_\Gamma) = I$), the other cases are similar. Properties (Irred.) and (Reg.) imply that for every $y \in I$ there exist $(x, y^-), (x, y^+) \in \Gamma$ such that $y^- < y < y^+$ and $y^- < x < y^+$. That is, I is the union of intervals of the form $]y^-, y^+[$, where $(x, y^-), (x, y^+) \in \Gamma$ and $y^- < x < y^+$. Using that the set I can be written as a countable union of compact sets, it is straightforward that there exist sequences $(x_k)_{k \in \mathbb{N}}, (y_k^-)_{k \in \mathbb{N}}, (y_k^+)_{k \in \mathbb{N}}$ such that the points (x_k, y_k^-) and (x_k, y_k^+) are in Γ , we have $y_k^- < x_k < y_k^+$,

$$\bigcup_{i=0}^k]y_i^-, y_i^+[\cap]y_{k+1}^-, y_{k+1}^+[\neq \emptyset, k \in \mathbb{N}, \quad \text{and} \quad \bigcup_{k \in \mathbb{N}}]y_k^-, y_k^+[= I.$$

Given an arbitrary set G , a G -good function ψ and an affine function a , the function $\psi' = \psi - a$ is again a G -good function. Thus for all finite G satisfying $(x_0, y_0^-), (x_0, y_0^+) \in G$, there is a G -good function ψ such that $\psi(y_0^-) = \psi(y_0^+) = 0$. Iterating (the first part of) Lemma 8.13 for $k \in \mathbb{N}$ we find the desired intervals K_y for $y \in \{y_k^-, y_k^+ : k \in \mathbb{N}\}$.

For every $y \in Y_\Gamma$ there exist $x \in \mathbb{R}$ and $k \in \mathbb{N}$ such that $(x, y) \in \Gamma$ and $y \in (y_k^-, y_k^+)$. Hence (the second part of) Lemma 8.13 yields the existence of the desired interval K_y for $y \in Y_\Gamma \setminus \{y_k^-, y_k^+ : k \in \mathbb{N}\}$.

We can view the set $\mathcal{K} := \prod_{y \in Y_\Gamma} K_y$ as a subset of the space of all functions from Y_G to \mathbb{R} . In the topology of pointwise convergence the set \mathcal{K} is compact by Tychonoff's theorem.

For every finite $G \subseteq \Gamma$, the set

$$\Psi_G := \left\{ \psi \in \mathcal{K} : \psi \text{ is } G\text{-good} \right\}$$

is a non-empty closed subset of the set \mathcal{K} . Moreover the family $(\Psi_G)_G$ has the finite intersection property. For instance, given finite sets $G_1, G_2 \subseteq \Gamma$ the intersection of Ψ_{G_1} and Ψ_{G_2} contains $\Psi_{G_1 \cup G_2}$ and is therefore non-empty. By compactness of \mathcal{K} , the intersection

$$\bigcap_{G \subseteq \Gamma, |G| < \infty} \Psi_G =: \Psi_\Gamma$$

of all these sets is nonempty as well. Obviously, any element $\psi \in \Psi_\Gamma$ is Γ -good. \square

Proof of Proposition 8.10. By Lemma 8.14, there exists a Γ -good function ψ . We have to show that ψ can be replaced by an upper semi-continuous function and that there exist appropriate functions φ and Δ . We start with the latter task.

Recall that we write $J = \text{conv}(Y_\Gamma)$ and note that $I \subseteq J \subseteq \bar{I}$.

For fixed $x \in X_\Gamma$, consider the function $y \mapsto g_x(y) = -\psi(y) + c(x, y)$, $y \in Y_\Gamma$. For any $x \in I$ let $g_x^{**} : \mathbb{R} \rightarrow [-\infty, +\infty]$ be the largest convex function which is smaller than g_x on the set Y_Γ for $x \in X_\Gamma$ and $g_x^{**} = +\infty$ if $x \in I \setminus X_\Gamma$. For $x \in X_\Gamma$, there exists an affine function which is smaller than g_x . Hence g_x^{**} does not take the value $-\infty$ in this case.

Since $I = \text{conv}(Y_\Gamma)$ the function g_x^{**} is continuous and finitely valued on the set J for $x \in X_\Gamma$. As a function on the set \mathbb{R} , g_x^{**} may possibly assume the value $+\infty$. Moreover, if $x \in I \setminus X_\Gamma$ then g_x^{**} can take the value $-\infty$.

We now define the function $H : I \times \mathbb{R} \rightarrow [-\infty, \infty]$ by

$$H(x, y) := (-\psi(\cdot) + c(x, \cdot))^{**}(y)$$

and emphasize that H takes finite values on $X_\Gamma \times J$. Thus the function $\varphi : I \rightarrow [-\infty, \infty]$, defined by

$$(33) \quad \varphi(x) := (-\psi(\cdot) + c(x, \cdot))^{**}(x) = H(x, x)$$

takes finite values on the set X_Γ .

To prove that φ is upper semi-continuous, consider for $n \in \mathbb{N}$ the function

$$H_n(x, y) := ((-\psi(\cdot) \vee (-n)) + c(x, \cdot))^{**}(y).$$

It is straightforward to prove that H_n is *continuous* on the set $I \times J$. Thus $H = \inf_{n \in \mathbb{N}} H_n$ is upper-semicontinuous and hence φ is upper semi-continuous as well.

For each $x \in I$ denote by $\Delta(x)$ the right-derivative of the convex function $y \mapsto H(x, y)$ in the point x if $H(x, x) > -\infty$ and set $\Delta(x) = 0$ otherwise.

By construction we then have

$$\varphi(x) + \psi(y) + \Delta(x)(y - x) \leq c(x, y),$$

for all $(x, y) \in X_\Gamma \times J$. Moreover, as ψ was assumed to be Γ -good, equality holds for all $(x, y) \in \Gamma$. (See the discussion preceding (28).)

Next we define a function $\tilde{\psi}$ by

$$\tilde{\psi}(y) = \inf_x c(x, y) - [\varphi(x) + \Delta(x)(y - x)].$$

For every x , the function $y \mapsto c(x, y) - [\varphi(x) + \Delta(x)(y - x)]$ is continuous, hence $\tilde{\psi}$ is upper semi-continuous. As above $\varphi(x) + \tilde{\psi}(y) + \Delta(x)(y - x) \leq c(x, y)$ holds by construction and since $\tilde{\psi}(y)$ is greater or equal to $\psi(y)$ for all $y \in I$ we conclude that the inequality is indeed an equality on the set Γ . \square

8.3. Integrating the duality relation between φ , ψ , Δ and c on the irreducible components. Paragraph 8.2 was a first step in the direction of the proof of Lemma 8.2. Unfortunately, the functions φ, ψ constructed in Proposition 8.10 are measurable but not necessarily integrable. The following lemma will provide a remedy for this.

Lemma 8.15. *Let χ be a convex or concave function on some (possibly unbounded) interval I and assume that μ, ν are in convex order and concentrated on I . Then*

$$(34) \quad \int \left[\int \chi(y) d\pi_x(y) - \chi(x) \right] d\mu(x) = \int \left[\int \chi(y) d\tilde{\pi}_x(y) - \chi(x) \right] d\mu(x)$$

for all measures $\pi, \tilde{\pi} \in \Pi_M(\mu, \nu)$.

Proof. We will give the proof in the case where $I = \mathbb{R}$ and χ convex, the other cases being similar. Note that, leaving integrability issues aside, the left as well as the right-hand side of (34) equal $\int \chi d\nu - \int \chi d\mu$ and in particular we expect them to be equal. To give a formal proof we approximate χ by functions which grow at most linearly so that all involved integrals do exist.

Denote by χ_n the smallest convex function which agrees with χ on the interval $[-n, n]$. (So χ_n is affine on the complement of $[-n, n]$.) We have to show that for each $\pi \in \Pi_M(\mu, \nu)$.

$$\int \underbrace{\left[\int \chi(y) d\pi_x(y) - \chi(x) \right]}_{=:f(x)} d\mu(x) = \lim_n \int \underbrace{\left[\int \chi_n(y) d\pi_x(y) - \chi_n(x) \right]}_{=:f_n(x)} d\mu(x)$$

Applying Jensen's inequality to the functions χ, χ_n we see that $f, f_n \geq 0$ and applying Jensen's inequality to the convex function $\chi_n - \chi_m$, we see that $f_n \leq f_m$ for $n \leq m$. Hence the desired equality follows from the monotone convergence theorem. \square

As a consequence of this lemma, the following definition is well.

Definition 8.16. *Assume that φ, ψ are measurable functions and that μ, ν are in convex order. Let χ be a convex¹⁰ function such that $\varphi_0 = \varphi + \chi$, $\psi_0 = \psi - \chi$ are uniformly bounded. Then we set*

$$\int \varphi d\mu + \int \psi d\nu := \int \varphi_0 d\mu + \int \psi_0 d\nu + \int \left[\int \chi(y) d\pi_x(y) - \chi(x) \right] d\mu(x),$$

where π is some martingale transport plan.

Corollary 8.17. *Assume that we are given measurable functions φ, ψ, Δ and a convex function χ such that*

$$(35) \quad \varphi(x) + \psi(y) + \Delta(x)(y - x) \leq c(x, y)$$

for all $x, y \in I$ and such that φ and $-\psi$ differ from χ only by some bounded functions. Then we have

$$\int \varphi d\mu + \int \psi d\nu \leq \int c d\pi$$

for any martingale transport plan π . Furthermore, if equality holds π -a.s. in (35), then $\int \varphi d\mu + \int \psi d\nu = \int c d\pi$.

We are now finally in the position to establish the main result of this section.

8.4. Proof of Lemma 8.2. We will first give the proof assuming that (μ, ν) is irreducible on the open interval I (bounded or not). According to Lemma 8.9, we may assume that the finitely optimal set Γ is included in $I \times \bar{I}$ and satisfies (Reg.) and (Irred.) on I . It follows from Proposition 8.10 that there exist upper semi-continuous functions $\varphi, \psi : I \rightarrow]-\infty, \infty]$ and a measurable function $\Delta : I \rightarrow \mathbb{R}$ such that

$$\varphi(x) + \psi(y) + \Delta(x)(y - x) \leq c(x, y)$$

for all $x, y \in I$, with equality holding for (x, y) in Γ . Recall that the function ψ constructed in Proposition 8.10 is of the form

$$\inf_x c(x, y) - [\varphi(x) + \Delta(x)(y - x)].$$

This leads us to define the convex function $\chi : I \rightarrow \mathbb{R}$ by

$$\chi(y) = \sup_x \varphi(x) + \Delta(x)(y - x).$$

¹⁰Of course the assertion is also true in the case where χ is concave, but we do not need this.

Since c is assumed to be bounded, it follows that ψ differs from $-\chi$ only by a bounded function (i.e. $\psi + \chi$ is bounded). Replacing φ by

$$(-\psi(\cdot) + c(x, \cdot))^{**}(x),$$

it follows also that φ differs from χ only by a bounded function (i.e. $\varphi - \chi$ is bounded). Thus Corollary 8.17 implies that π is an optimal transport plan.

Consider now the general case and the decomposition $\pi = (\sum_k \pi_k) + \eta$ of Theorem 8.4, (25), where $(\text{proj}^x \pi_k, \text{proj}^y \pi_k) =: (\mu_k, \nu_k)$ is irreducible. But Γ has full measure for π_k (if not $\pi(\Gamma)$ would be smaller than 1) and it is finitely optimal for the cost c . According to the first part of the proof π_k is an optimal martingale transport plan from μ_k to ν_k . By Theorem 8.7, π is optimal and this concludes the proof of Lemma 8.2.

9. APPENDIX B: A SELF-CONTAINED APPROACH TO THE VARIATIONAL LEMMA

In this appendix we provide a self-contained proof of the variational lemma (Lemma 1.11, established in Section 3). Indeed we obtain a somewhat stronger conclusion in Theorem 9.4 below. The benefit of this second version is that Theorem 9.4 does not rely on the Choquet's capacability theorem and that the new approach provides an explicit set Γ . A drawback is that we have to assume that the cost function is continuous. Compared to the approach given in Section 3 another disadvantage is that the argument does not seem to be adaptable from $\mathbb{R} \times \mathbb{R}$ to more general product spaces.

9.1. Preliminaries based on Lebesgue's density theorem. Our aim is to establish Corollary 9.3 which may be viewed as an avatar of Lemma 3.2, the uncountable set of points a being replaced by a set A of positive measure. We start with the well-known Lebesgue density theorem. It asserts that for an integrable function f on $[0, 1]$ we have

$$(36) \quad \lim_{\varepsilon \rightarrow 0} \frac{1}{2\varepsilon} \int_{s-\varepsilon}^{s+\varepsilon} |f(s) - f(t)| dt = 0$$

for almost every $s \in]0, 1[$. In sloppy language, almost every point is a "good" point. Those points will be called *regular points* of f . In those regular points s we also have

$$(37) \quad \lim_{n \rightarrow +\infty} \frac{1}{\lambda(M_n)} \int_{M_n} |f(s) - f(t)| dt = 0$$

for every sequence (M_n) of measurable sets satisfying $M_n \subseteq [s - \varepsilon_n, s + \varepsilon_n]$ with $\frac{\lambda(M_n)}{\varepsilon_n}$ bounded from below and $\varepsilon_n \rightarrow 0$. Particular admissible choices are $M_n = [s, b_n]$ or $]s, b_n]$ and $M_n = [a_n, s]$ or $[a_n, s[$. As a consequence of (37) we have that

$$(38) \quad \lim_{n \rightarrow +\infty} \frac{1}{\lambda(M_n)} \int_{M_n} f(t) dt = f(s).$$

Intervals $B =]q, q']$ or $] -\infty, q']$ with $q, q' \in \mathbb{Q} \cup \{-\infty, +\infty\}$ will be called *rational semi-open intervals*. By Fubini's theorem, (37) implies the following result.

Lemma 9.1. *Let π be a probability measure on $\mathbb{R} \times \mathbb{R}$ with first marginal $\lambda_{[0,1]}$. Fix a disintegration $(\pi_x)_{x \in [0,1]}$. There exists a set $R \subseteq [0, 1]$, $\lambda(R) = 1$ such that for $s \in R$, any rational semi-open interval B and any two sequences $(a_n)_n, (b_n)_n$ satisfying $a_n, b_n \rightarrow s$ as well as $a_n \leq s < b_n$ or $a_n < s \leq b_n$, we have*

$$\lim_{n \rightarrow +\infty} \frac{1}{b_n - a_n} \int_{a_n}^{b_n} |\pi_t(B) - \pi_s(B)| d\lambda(t) = 0.$$

We now extend this lemma to the case where the first marginal of π is a general measure μ , not necessarily equal to $\lambda|_{[0,1]}$. Recall that G_μ denotes the quantile function of μ and F_μ the cumulative distribution function. See Figure 5 for the graphs of F_μ and G_μ in an example: here μ satisfies $\mu(\{1\}) = 1/3$ and is uniform of mass $2/3$ on $[0, 1] \cup [2, 3]$ (the axis are not scaled in the same way). Recall that the measure μ can be written as $(G_\mu)_\# \lambda$.




FIGURE 5. The quantile and cumulative distribution functions.

The map G_μ is increasing on $[0, 1]$ and hence continuous on the complement of a countable set D , the set of $s \in [0, 1]$ such that $F_\mu^{-1}(s)$ is a non trivial interval. For such a $s \in D$, the μ -measure of $F_\mu^{-1}(s)$ is zero so that $\mu(G_\mu(D)) \leq \mu(F_\mu^{-1}(D)) = 0$.

Consider a random variable $(U, G_\mu(U), Y)$ on $[0, 1] \times \mathbb{R} \times \mathbb{R}$ such that the law of U is λ and the law of $(G_\mu(U), Y)$ is π . Let $\tilde{\pi}$ be the law of (U, Y) and $(\tilde{\pi}_s)_{s \in [0,1]}$ a disintegration with respect to λ , i.e. $\tilde{\pi}_s$ is the conditional law of Y given the event $\{U = s\}$. Apply Lemma 9.1 to this disintegration of $\tilde{\pi}$ to obtain a set R . Let $S \subseteq \mathbb{R}$ be the set $G_\mu(R \setminus D)$ and let us call S the *set of regular points*.

Note that S has full measure and that it may depend on the disintegration of $\tilde{\pi}$.

Lemma 9.2. *Let π be a probability measure on \mathbb{R}^2 with first marginal μ and $(\pi_x)_{x \in \mathbb{R}}$ a disintegration of π . There exists a set $S \subseteq \mathbb{R}$ of measure $\mu(S) = 1$ satisfying the following: for any $x \in S$ and any rational semi-open interval B the limit*

$$\lim_{n \rightarrow +\infty} \frac{1}{\mu(N_n)} \int_{N_n} |\pi_t(B) - \pi_x(B)| d\mu(t)$$

is zero for any sequence $N_n = [x - \varepsilon_n, x + \varepsilon_n]$ with $\varepsilon_n \downarrow 0$.

Proof. We note that if the statement of the lemma holds for one particular disintegration of π , then it automatically carries over to any other disintegration.

Therefore we will consider a disintegration of π which is convenient for the proof. Let S and $\tilde{\pi}$ be as in the discussion preceding Lemma 9.2 and set for $x \in \mathbb{R}$

$$(39) \quad \pi_x = \begin{cases} \tilde{\pi}_{F_\mu(x)} & \text{if } \mu(\{x\}) = 0, \\ \frac{1}{\mu(\{x\})} \int_{G_\mu^{-1}(\{x\})} \tilde{\pi}_s ds & \text{if } \mu(\{x\}) > 0. \end{cases}$$

Let x be a point in S and $N_n = [x - \varepsilon_n, x + \varepsilon_n]$. We distinguish two cases depending on whether or not x is an atom of μ . The first case is quite straightforward. In the second case we will apply Lemma 9.1.

- Assume $\mu(\{x\}) > 0$. As $\bigcap_{n \in \mathbb{N}} N_n = \{x\}$ we have $\mu(N_n) \downarrow \mu(\{x\})$ as $\varepsilon_n \rightarrow 0$. Hence

$$\begin{aligned} \frac{1}{\mu(N_n)} \int_{N_n} |\pi_t(B) - \pi_x(B)| d\mu(t) &= \frac{1}{\mu(N_n)} \int_{\{x\}} |\pi_t(B) - \pi_x(B)| d\mu(t) \\ &\quad + \frac{1}{\mu(N_n)} \int_{N_n \setminus \{x\}} |\pi_t(B) - \pi_x(B)| d\mu(t). \end{aligned}$$

The first part of the sum equals 0 and the second part tends to 0 since $|\pi_t(B) - \pi_x(B)| \leq 2$ and $[\mu(N_n) - \mu(x)]/\mu(N_n) \downarrow 0$ as $\varepsilon_n \rightarrow 0$.

- Assume $\mu(\{x\}) = 0$. Set $s_0 \in G_\mu^{-1}(x)$ and $M_n = G_\mu^{-1}(N_n)$. As $x \in S$, we can assume that G_μ is continuous in s_0 . As G_μ is continuous in s_0 , the interval M_n cannot be trivial. Hence $\lambda(M_n) = \mu(N_n)$ is positive. Note that (s_0, x) is not on a vertical or horizontal part of the graphs of F_μ or G_μ . Hence $F_\mu(x) = s_0$. As x is the unique pre-image of s_0 and $x \in S$, we have $s_0 \in R$ using the notation of Lemma 9.1, i.e. s_0 is regular for the disintegration $(\tilde{\pi}_s)_s$.

We can separate the push-forward measure $\mu = (G_\mu)_\# \lambda$ into its atomic and its continuous part and integrate accordingly and thus obtain

$$\begin{aligned} \frac{1}{\mu(N_n)} \int_{N_n} |\pi_t(B) - \pi_x(B)| d\mu(t) &= \frac{1}{\lambda(M_n)} \int_{M_n} |\pi_{G_\mu(s)}(B) - \pi_x(B)| d\lambda(s) \\ (40) \qquad \qquad \qquad &\leq \frac{1}{\lambda(M_n)} \int_{M_n} |\tilde{\pi}_s(B) - \tilde{\pi}_{s_0}(B)| d\lambda(s). \end{aligned}$$

Here we used the following properties: (i) if $\mu(\{t\}) > 0$: Jensen's inequality for the integration on $\{s : G_\mu(s) = t\}$, (ii) if $\mu(\{t\}) = 0$: $G_\mu(s) = t$ implies that $F_\mu(t) = s$ or s is a discontinuity point of G_μ , such that $F_\mu(t) = s$ almost surely.

The right and left limit of the cumulative distribution function F_μ are equal in x so that F_μ is continuous in this point. Thus the interval $M_n = [F_\mu((x - \varepsilon_n)^-), F_\mu(x + \varepsilon_n)] \subseteq [F_\mu(x - 2\varepsilon_n), F_\mu(x + \varepsilon_n)]$ has length going to 0 as $\varepsilon_n \rightarrow 0$. Hence we can apply Lemma 9.1 in equation (40) at point s_0 for the disintegration $(\tilde{\pi}_s)_s$ and the sequence M_n . □

We remark that for $\pi \in \mathcal{P}(\mathbb{R}^2)$, if $y \in \text{spt}(\pi_x)$, it is not always true that $(x, y) \in \text{spt}(\pi)$. We have introduced S in order to obtain this conclusion for $x \in S$. More precisely, we have:

Corollary 9.3. *Let S be a set of regular points associated to $(\pi_x)_x$ as in Lemma 9.2 and let $x \in S$. Let B_1, \dots, B_k be a family of pairwise disjoint rational semi-open intervals such that $\pi_x(B_j) > 0$ for $j = 1, \dots, k$.*

For every $\varepsilon > 0$ there exists $A \subseteq S \cap [x - \varepsilon, x + \varepsilon]$ such that $\mu(A) > 0$ and $\pi_t(B_j) > 0$ for $(j, t) \in \{1, \dots, k\} \times A$.

Proof. Let π, x, ε and the sets B_j be given. Let $(\varepsilon_n)_n$ be a decreasing sequence of positive numbers tending to 0. For every j we have

$$\lim_{n \rightarrow +\infty} \frac{1}{\mu(N_n)} \int_{N_n} |\pi_x(B_j) - \pi_t(B_j)| d\mu(t) = 0,$$

where N_n is $[x - \varepsilon_n, x + \varepsilon_n]$ or, in the case $\mu(\{x\}) = 0$, one of the intervals $]x, x + \varepsilon_n]$ resp. $[x - \varepsilon_n, x[$. This implies

$$\mu(\{t \in N_n : |\pi_x(B_k) - \pi_t(B_k)| > \pi_x(B_k)/2\}) = o(\mu(N_n)).$$

Therefore

$$\mu(\{t \in N_n : \exists j \in \{1, \dots, k\}, |\pi_x(B_k) - \pi_t(B_k)| > \pi_x(B_k)/2\}) = o(\mu(N_n))$$

and

$$\mu(\{t \in N_n : \exists j \in \{1, \dots, k\}, \pi_t(B_k) = 0\}) = o(\mu(N_n)).$$

Hence for n sufficiently large the set

$$A = \{t \in N_n : \forall j \in \{1, \dots, k\}, \pi_t(B_k) > 0\}$$

has positive measure. For n large enough we also have $\varepsilon_n < \varepsilon$, which concludes the proof. \square

9.2. Construction of a better competitor when Γ supports a finite non optimal coupling. Let \mathcal{V} be the set of signed measures σ on \mathbb{R}^2 with Hahn decomposition $\sigma = \sigma^+ - \sigma^-$ such that the following conditions are satisfied:

- The total mass of σ is 0.
- The marginals $\text{proj}_{\#}^x \sigma$ and $\text{proj}_{\#}^y \sigma$ vanish identically.
- The measure $\text{proj}_{\#}^y(|\sigma|) = \text{proj}_{\#}^y \sigma^+ + \text{proj}_{\#}^y \sigma^-$ has finite first moment.
- σ has a disintegration $(\sigma_x)_x$ such that $\text{proj}_{\#}^x |\sigma|$ -a.s., the positive and the negative parts of σ_x have the same mean.

If only the three first conditions are satisfied σ will be an element of \mathcal{V}' .

Here the letter \mathcal{V} is reminiscent to the term *variation*. Indeed observe that if α is a positive measure on \mathbb{R}^2 such that $\text{proj}_{\#}^y \alpha$ has finite first moment and $\beta = \alpha - \sigma$ is a positive measure, then β is a competitor of α in the sense of Definition 1.10. Conversely, for a pair of competitors (α, β) , the measures $\alpha - \beta$ and $\beta - \alpha$ are elements of \mathcal{V} . A notable element of \mathcal{V} is $(\delta_x - \delta_{x'}) \otimes (\lambda \delta_{y^+} + (1 - \lambda) \delta_{y^-} - \delta_{\lambda y^+ + (1-\lambda)y^-})$, the kind of measure that we have used repeatedly in Sections 6 and 7. An element of \mathcal{V} will be called a *variation*. A variation σ is *positive* (resp. *negative*) if $\int c(x, y) d\sigma(x, y) > 0$ (resp. < 0).

For a cost function satisfying the sufficient integrability condition, it is not difficult to prove that the following statements are equivalent:

- (1) The martingale transport plan α is optimal for the cost c .
- (2) For any $\sigma \in \mathcal{V}$ such that $\sigma^+ \leq \alpha$, one has $\int c(x, y) d\sigma(x, y) \leq 0$.

We can now state the main result of this appendix.

Theorem 9.4. *Assume that μ, ν are probability measures in convex order and that $c : \mathbb{R}^2 \rightarrow \mathbb{R}$ is a continuous cost function satisfying the sufficient integrability condition. Assume that $\pi \in \Pi_M(\mu, \nu)$ is an optimal martingale transport plan which leads to finite costs. Let $(\pi_x)_x$ be a disintegration of π and $S \subseteq \mathbb{R}$ a set of regular points associated to $(\pi_x)_x$ in the sense of Lemma 9.2. We set*

$$\Gamma = \{(x, y) \in \mathbb{R}^2 : x \in S \text{ and } y \in \text{spt}(\pi_x)\}.$$

If α is a martingale transport plan such that

- *the support $\text{spt}(\alpha)$ of α is finite and*
- *the support $\text{spt}(\alpha)$ is included in Γ ,*

then the martingale transport plan α is optimal for c between $\text{proj}_{\#}^x \alpha$ and $\text{proj}_{\#}^y \alpha$.

Furthermore if σ is a measure of finite support in \mathcal{V} with $\text{spt}(\sigma^+) \subseteq \Gamma$, it is a non-positive variation.

Proof. Let α be as in the theorem and assume for contradiction that there exists a competitor β that leads to smaller costs. We will prove that π cannot be optimal, thus establishing the desired contradiction. In other words assume that there is a variation $\sigma \in \mathcal{V}$ with $\text{spt} \sigma^+ \subseteq \text{spt} \alpha$ and $\int c(x, y) d\sigma(x, y) > 0$. We will construct $\tilde{\sigma} \in \mathcal{V}$ by applying modifications to σ so that $\tilde{\sigma}^+ \leq \pi$ and $\int c(x, y) d\tilde{\sigma}(x, y) > 0$.

This yields a contradiction since the competitor $\pi - \tilde{\sigma}$ is cheaper than π with respect to the cost function c .

The argument is based on two lemmas and Proposition 9.6, whose proof is postponed to the next subsection. Let us introduce some notations. Assume first that $\text{spt} |\sigma|$ is included in $\{x_1, \dots, x_n\} \times \{y_1, \dots, y_m\}$ and define for $\varepsilon > 0$ the rectangle $R_{ij}(\varepsilon) = [x_i - \varepsilon, x_i + \varepsilon] \times [y_j - \varepsilon, y_j + \varepsilon]$.

Lemma 9.5. *There exists $\varepsilon > 0$ such that the sets $R_{ij}(\varepsilon)$ are disjoint and any measure $\sigma' \in \mathcal{V}$ satisfying*

- $|\sigma'|$ is concentrated on $\bigcup_{i,j} R_{ij}(\varepsilon)$ and
- for $(i, j) \in \{1, \dots, n\} \times \{1, \dots, m\}$

$$|\sigma - \sigma'| (R_{ij}) \leq \varepsilon,$$

is a positive variation.

Proof. The argument relies on the continuity of c and is straightforward. \square

Let us call $\mathcal{V}(\sigma, \varepsilon)$ the subset of the measures $\sigma' \in \mathcal{V}$ such that σ' satisfies the conditions of the above lemma. Elements of $\mathcal{V}(\sigma, \varepsilon)$ are positive variations and so are the elements of the cone $\mathcal{CV}(\sigma, \varepsilon) = \{w\sigma' \in \mathcal{V} : w > 0 \text{ and } \sigma' \in \mathcal{V}(\sigma, \varepsilon)\}$. We want to find a measure $\sigma' \in \mathcal{V}(\sigma, \varepsilon)$ and v such that $w\sigma'^+ \leq \pi$. For this purpose we will use the fact that σ^+ is concentrated on Γ .

Using the notations of Corollary 9.3, let A_i be the set A associated to x_i and consider an arbitrary family of rational semi-open intervals B_k with $y_j \in B_j \subseteq [y_j - \varepsilon, y_j + \varepsilon]$ and $\pi_{x_i}(B_j) > 0$ for each j . Moreover we take $A_i \subseteq S \cap [x_i - \varepsilon, x_i + \varepsilon]$ for every i .

Proposition 9.6. *Let $\varepsilon > 0$. There are sets A_1, \dots, A_n with $\mu(A_i) > 0$ and $A_i \subseteq [x_i - \varepsilon, x_i + \varepsilon]$ such that for $(t_1, \dots, t_n) \in A_1 \times \dots \times A_n$ there is a measure $\sigma_{t_1, \dots, t_n} \in E$ satisfying the following:*

- We have $\sigma_{t_1, \dots, t_n} \in \mathcal{CV}(\sigma, \varepsilon)$.
- The first marginal of $|\sigma_{t_1, \dots, t_n}|$ has support $\{t_1, \dots, t_n\}$.
- $\sigma_{t_1, \dots, t_n}^+ \leq \sum_{i=1}^n \mu(A_i) \times (\delta_{t_i} \otimes \pi_{t_i})$.

We postpone the proof of Proposition 9.6 to the next subsection.

Note that σ_{t_1, \dots, t_n} is not the measure $\tilde{\sigma}$ we are looking for. Nevertheless it satisfies almost all the conditions. It is in \mathcal{V} and even in $\mathcal{CV}(\sigma, \varepsilon)$ so that according to Lemma 9.5 it is a positive variation. The only missing condition is that $\sigma_{t_1, \dots, t_n}^+$ is not smaller than π . We provide a remedy in the following lemma:

Lemma 9.7 (A variation $\tilde{\sigma}$ leading to the contradiction). *The measure*

$$\tilde{\sigma} = \frac{1}{\mu(A_1) \times \dots \times \mu(A_n)} \iiint_{A_1 \times \dots \times A_n} \sigma_{t_1, \dots, t_n} d\mu(t_1) \otimes \dots \otimes d\mu(t_n)$$

is in $\mathcal{V}(\sigma, \varepsilon)$ and satisfies both $\iint c(x, y) d\tilde{\sigma}(x, y) > 0$ and $\tilde{\sigma}^+ \leq \pi$. Hence $\pi - \tilde{\sigma}$ gives rise to smaller costs than π .

Proof. As all σ_{t_1, \dots, t_n} are in $\mathcal{CV}(\sigma, \varepsilon)$, they are positive variations. Hence $\tilde{\sigma}$ which is an average of these measures in \mathcal{V} is also a positive variation. Let us prove that $\tilde{\sigma}^+ \leq \pi$. Observe that $\tilde{\sigma}^+$ is again the average of the positive parts $\sigma_{s_1, \dots, s_n}^+$. By

Proposition 9.6 this is smaller than

$$\begin{aligned}
& \frac{1}{\mu(A_1) \times \cdots \times \mu(A_n)} \iiint_{A_1 \times \cdots \times A_n} \sum_{i=1}^n \mu(A_i) (\delta_{t_i} \otimes \pi_{t_i}) d\mu(t_1) \otimes \cdots \otimes d\mu(t_n) \\
&= \sum_{i=1}^n \int_{A_i} \left(\frac{\iiint (\delta_{t_i} \otimes \pi_{t_i}) d\mu(t_1) \otimes \cdots \otimes \widehat{d\mu(t_i)} \otimes \cdots \otimes d\mu(t_n)}{\mu(A_1) \times \cdots \times \widehat{\mu(A_i)} \times \cdots \times \mu(A_n)} \right) d\mu(t_i) \\
&= \sum_{i=1}^n \int_{A_i} (\delta_{t_i} \otimes \pi_{t_i}) d\mu(t_i) = \pi|_{\bigcup_{i=1}^n A_i \times \mathbb{R}}. \quad \square
\end{aligned}$$

Up to Proposition 9.6 we have thus proved Theorem 9.4. \square

9.3. Proof of Proposition 9.6. Recall the definitions and notations of Theorem 9.4 and Proposition 9.6. In particular σ has finite support included in Γ . It is also included in some product set $\{x_1, \dots, x_n\} \times \{y_1, \dots, y_m\}$ where we choose m and n as small as possible. For $\tau \in \mathcal{V}$ we denote the support of $\text{proj}_{\#}^x(|\tau|)$ by $X(\tau)$ and the support of $\text{proj}_{\#}^y(|\tau|)$ by $Y(\tau)$ so that $\{x_1, \dots, x_n\} = X(\sigma)$ and $\{y_1, \dots, y_m\} = Y(\sigma)$. Let $d \leq n \cdot m$ be the cardinality of $\text{spt}(\sigma^+)$ and denote its elements by p_1, \dots, p_d .

For measures of finite support the conditions for being in \mathcal{V} can be simplified. A measure τ is in \mathcal{V} if

- (1) For every $y \in Y(\tau)$, $L_y(\tau)$ defined as $\sum_{x \in X} \tau(x, y)$ is zero,
- (2) for every $x \in X(\tau)$, $C_x(\tau)$ defined as $\sum_{y \in Y} \tau(x, y)$ is zero,
- (3) for every $x \in X(\tau)$, $M_x(\tau)$ defined as $\sum_{y \in Y} \tau(x, y) \times y$ is zero.

Moreover the measure τ is an element of \mathcal{V}' if the Conditions (1) and (2) are satisfied.

We introduce some further notations. For every $\tau \in \mathcal{V}'$ of finite support we introduce a relation between the points of $X(\tau)$. We write $x \rightarrow x'$ if there are $y, y', y > y'$ such that $\tau(x, y)$ and $\tau(x', y')$ are not zero. If $x \rightarrow x'$ and $x' \rightarrow x$ we write $x \leftrightarrow x'$ and will say that x *double-touches* x' . If $\tau \in \mathcal{V}$, for any point $x \in X(\tau)$ an important consequence of Condition (3) is that there exist three distinct points y, y', y'' such that $\tau(x, y)$, $\tau(x, y')$ and $\tau(x, y'')$ are not zero. Hence $x \leftrightarrow x$ if $x \in X(\tau)$. However the relation \leftrightarrow is not transitive. If $x \in X$ double-touches both x' and x'' we say that x *is a bridge over x' and x''* . In particular if $x \leftrightarrow x'$ the point x is a bridge over x' and x itself.

Roughly speaking for $\tau \in \mathcal{V}'$, the relation $x \rightarrow x'$ means that it is possible to replace τ (in a continuous manner) by a signed measure $\tau' \in \mathcal{V}'$ such that τ^+ and τ'^+ have the same support. Applying this modification $\tau \mapsto M_x(\tau)$ increases while $\tau \mapsto M_{x'}(\tau)$ decreases (and their sum remains constant). More precisely, consider $y, y', y > y'$ such that $\tau(x, y)$ and $\tau(x', y')$ are both non zero. Let m be the measure $(\delta_x - \delta_{x'}) \otimes (\delta_y - \delta_{y'})$. Notice that m is an element of $\mathcal{V}' \setminus \mathcal{V}$. Considering $\tau^h = \tau + h \cdot m$ and $h > 0$ we have

$$M_x(\tau^h) - M_x(\tau) = h \cdot M_x(m) = h \cdot (y - y') > 0.$$

We only consider positive h in order to keep the same support for $(\tau^h)^+$ and τ^+ . In particular this prohibits that $\tau(x, y') > 0$ and $\tau(x', y) > 0$. For the same reason we choose $h \in [0, h_0[$ where $h_0 = \max(|\tau(x, y)|, |\tau(x', y')|)$. Indeed if $\tau(x, y) < 0$ then the same applies to $\tau^h(x, y)$.

If we want to make M_x and $M_{x'}$ vary in the opposite direction we may consider the relation $x' \rightarrow x$ in place of $x \rightarrow x'$. Thus $x \leftrightarrow x'$ allows to make small variations of M_x and $M_{x'}$ in the one or the other direction. If there is a bridge $x'' \in X(\tau)$ over x and x' we have exactly the same freedom as if $x \leftrightarrow x'$. The next lemma is a tool for finding bridges between points when $\tau \in \mathcal{V}$.

Lemma 9.8. *Let τ be a finitely supported element of \mathcal{V} and $(x, y) \in X(\tau) \times Y(\tau)$ such that $\tau(x, y) > 0$. Let $G \subseteq X(\tau)$ be the subset of points x' such that*

- *there exists a bridge over x and x' ,*
- *$\tau(x', y) < 0$.*

Then

$$\tau(x, y) + \sum_{x' \in G} \tau(x', y) \leq 0.$$

Proof. Condition (1) implies that if every $x' \in X(\tau)$ satisfying $\tau(x', y) < 0$ is connected with x by a bridge, we are done. Conversely assume that there exists $x' \in X(\tau)$ such that $\tau(x', y) < 0$ and there is no bridge between x and x' . Then for any $x_0 \in X(\tau)$ the measure $|\tau|$ restricted to $\{x_0\} \times \mathbb{R}$ is concentrated on $\{x_0\} \times [y, +\infty[$ or $\{x_0\} \times]-\infty, y]$ (if not it would be a bridge between x and x'). Let $X^1 \sqcup X^2$ be the partition of $X(\tau)$ induced by this remark and τ^i the restriction of τ to $X^i \times \mathbb{R}$ for $i = 1, 2$. Without loss of generality we can assume $x \in X^1$. Let us prove that τ^1 and τ^2 are in \mathcal{V} . Actually they coincide with τ on vertical lines so that they satisfy Conditions (2) and (3). The total mass of τ on the horizontal lines that are not equal to $\mathbb{R} \times \{y\}$ is zero as well. Thus, as $\tau^i(\mathbb{R}^2) = 0$, we obtain $\tau^i(X^i \times \{y\}) = 0$ for $i = 1, 2$. This yields Condition (1) for τ^1 and τ^2 . Hence these measures are in \mathcal{V} .

As $\tau^1 \in \mathcal{V}$, applying Condition (1) we obtain that any $x'_1 \in X^1$ such that $\tau(x'_1, y) < 0$ is connected with x by a bridge. Indeed with Condition (2) and the definition of X^1 , we know that there are y' and y'' in $]y, +\infty[$ such that $\tau(x, y') \neq 0$ and $\tau(x'_1, y'') \neq 0$. Hence we have $x \leftrightarrow x'_1$. So we can apply the first remark to τ^1 in place of τ . Indeed, G is the set of points of $x_1 \in X(\tau^1)$ such that $\tau(x_1, y) = \tau^1(x_1, y) < 0$. \square

Lemma 9.9. *Let τ be a finitely supported positive variation and consider $\text{spt}(\tau^+) = \{p_1, \dots, p_d\} \subseteq \mathbb{R} \times \mathbb{R}$. There exists $\varepsilon > 0$ such that if $q_k \in \mathbb{R}^2$ has the same first coordinate as p_k and $|p_k - q_k| < \varepsilon$ for every $k \in \{1, \dots, d\}$, then there exists a sequence of positive variations $(\tau_k)_{k=1}^d$ such that $|\tau_k|$ has finite support and τ_k^+ has support $\{q_1, \dots, q_k, p_{k+1}, \dots, p_d\}$.*

Proof. Let ε be a positive real number. Let us denote by X the support of $\text{proj}_{\#}^x(|\tau_k|)$ for some $k \in \{1, \dots, d\}$ (which does not depend on k). We explain how to build τ_k from τ_{k-1} . Roughly speaking we are moving $p_k = (a, b)$ to a position $q_k = (a, b')$, where $|b' - b| < \varepsilon$. Doing this, we have to take care to stay in \mathcal{V} . The conditional measure $\tau_k|_x$ can easily be forced to preserve mass zero (Condition (2)) during this operation but there are two difficulties: for each y the conditional measures $\tau_k|_y$ must have mass zero (Condition (1)). The second problem is that for each $x \in X$ the positive and the negative part of $\tau_k|_x$ must have the same mean (Condition (3)).

Let us go into details. We define τ_k from τ_{k-1} in two steps: the first step is a vertical translation. Applying Lemma 9.8 to $p_k = (a, b)$ we obtain a measure m concentrated on $X(\tau) \times \{b\}$ that satisfies the following conditions:

- $m(\mathbb{R}^2) = 0$,
- m^+ is concentrated on the point $p_k = (a, b)$ and $m(a, b) = \tau_{k-1}(a, b)$,
- m^- is concentrated on a set $G \times \{b\}$ such that any $x \in G$ is connected with a by a bridge and $m^- \leq \tau_{k-1}^-$.

Let us denote m by $\zeta \otimes \delta_b$. We replace τ_{k-1} by $\tau'_{k-1} = \tau_{k-1} + \zeta \otimes (\delta_{b'} - \delta_b)$. Doing this we preserve Conditions (1) and (2), i.e. the measure is still in \mathcal{V}' , but Condition

(3) is possibly violated. Recall that ζ has mass zero. It follows that

$$M_a(\tau'_{k-1}) + \sum_{x \in G} M_x(\tau'_{k-1}) = 0.$$

Using the bridges between a and the elements of G (these bridges are available for τ'_{k-1} as they were for τ_{k-1} assuming that ε is sufficiently small) we can modify the measure and make M_a and M_x for $x \in G$ equal to 0. Call τ_k the result of this procedure. Observe that if the variations are sufficiently small then the points of positive mass are exactly $q_1, \dots, q_k, p_{k+1}, \dots, p_d$ as we want. As in Lemma 9.5, we also obtain that the variations $(\sigma_k)_{k=1}^d$ are positive provided that $\varepsilon > 0$ is sufficiently small. \square

We can now prove Proposition 9.6. Let $\sigma \in \mathcal{V}$ of finite support as in the proof of Theorem 9.4. Observe that σ can be written as a sum

$$\sum_{k=1}^d \zeta_k \otimes \delta_{y_k}$$

where for $k \in \{1, \dots, d\}$ the signed measure ζ_k has its positive part concentrated in one point. Given k , let ω_k be a probability measure on \mathbb{R} with expectation y_k (the same as δ_{y_k}). We consider

$$\sum_{k=1}^d \zeta_k \otimes \omega_k$$

and easily convince ourselves that this measure is an element of \mathcal{V} . We will apply this transformation not directly to σ but to a measure $\sigma_d \in \mathcal{V}(\sigma, \varepsilon)$, that we build in the following paragraph.

The proof of the proposition proceeds as follows. Consider the family of points (r_1, \dots, r_d) of the support of σ^+ and pick ε as in Lemma 9.5. For each point $r_k = (a, b)$ we consider a rational semi-open interval $B_k \ni b$ of diameter smaller than ε . Using Corollary 9.3 we obtain a family $(A_i)_{1 \leq i \leq n}$ and we can assume that these sets are included in $[x_i - \varepsilon, x_i + \varepsilon]$. We fix a point (t_1, \dots, t_n) of $A_1 \times \dots \times A_n$. For each $k \in \{1, \dots, d\}$ we can write r_k in the form (x_i, b) . We have $\pi_{t_i}(B_k) > 0$. Let now $p_k = (t_i, b)$ and $q_k = (t_i, \tilde{y})$ where $\tilde{y} = \frac{1}{\pi_{t_i}(B_k)} \int_{B_k} y d\pi_{t_i}(y)$. Apply Lemma 9.9 to the measure $\sigma_0 \in \mathcal{V}$ obtained from σ by translating horizontally the mass concentrated on the line $\{x_i\} \times \mathbb{R}$: the measure $\sigma|_{x_i}$ equals precisely $\sigma_0|_{t_i}$. The other parameters (p_1, \dots, p_d) and (q_1, \dots, q_d) have just been constructed. Applying Lemma 9.9 we obtain a measure $\sigma_d \in \mathcal{V}(\sigma, \varepsilon)$ concentrated on $\{t_1, \dots, t_n\} \times \mathbb{R}$ and $\text{spt } \sigma_d^+ = \{q_1, \dots, q_d\}$. Next we perform the transformation explained above where each ω_k has the form $\frac{1}{\pi_{t_i}(B_k)} \pi_{t_i}|_{B_k}$ for some (i, k) . The measure $\bar{\sigma}_d$ we obtain is in $\mathcal{V}(\sigma, \varepsilon)$ but it may not satisfy the condition $\bar{\sigma}_d^+ \leq \sum_{i=1}^n \mu(A_i) \delta_{t_i} \otimes \pi_{t_i}$. However this inequality does hold for $w \bar{\sigma}_d^+ \in \mathcal{CV}(\sigma, \varepsilon)$ if w is a sufficiently small positive constant.

REFERENCES

- [1] M. Beiglböck, M. Goldstern, G. Maresch, and W. Schachermayer. Optimal and better transport plans. *J. Funct. Anal.*, 256(6):1907–1927, 2009.
- [2] M. Beiglböck, P. Henry-Labordère, and F. Penkner. Model-independent bounds for option prices—a mass transport approach. *Finance Stoch.*, 17(3):477–501, 2013.
- [3] D. Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability, 1950*, pages 93–102, Berkeley and Los Angeles, 1951. University of California Press.
- [4] Y. Brenier. Décomposition polaire et réarrangement monotone des champs de vecteurs. *C. R. Acad. Sci. Paris Sér. I Math.*, 305(19):805–808, 1987.

- [5] P. Cartier, J. M. G. Fell, and P.-A. Meyer. Comparaison des mesures portées par un ensemble convexe compact. *Bull. Soc. Math. France*, 92:435–445, 1964.
- [6] G. Choquet. Forme abstraite du théorème de capacitabilité. *Ann. Inst. Fourier. Grenoble*, 9:83–89, 1959.
- [7] Y. Dolinsky and H. Mete Soner. Martingale Optimal Transport and Robust Hedging in Continuous Time. *Probab. Theory Relat. Fields*, page to appear, 2013.
- [8] J. L. Doob. Generalized sweeping-out and probability. *J. Funct. Anal.*, 2:207–225, 1968.
- [9] P. Fischer and J. A. R. Holbrook. Balayage defined by the nonnegative convex functions. *Proc. Amer. Math. Soc.*, 79(3):445–448, 1980.
- [10] A. Galichon, P. Henry-Labordere, and N. Touzi. A Stochastic Control Approach to No-Arbitrage Bounds Given Marginals, with an Application to Lookback Options. *Ann. Appl. Probab.*, page to appear, 2013.
- [11] G. H. Hardy, J. E. Littlewood, and G. Pólya. *Inequalities*. Cambridge, at the University Press, 1952. 2d ed.
- [12] F. Hirsch, C. Profeta, B. Roynette, and M. Yor. *Peacocks and associated martingales, with explicit constructions*, volume 3 of *Bocconi & Springer Series*. Springer, Milan, 2011.
- [13] F. Hirsch and B. Roynette. A new proof of Kellerer’s theorem. *ESAIM Probab. Stat.*, 16:48–60, 2012.
- [14] D. Hobson. The Skorokhod embedding problem and model-independent bounds for option prices. In *Paris-Princeton Lectures on Mathematical Finance 2010*, volume 2003 of *Lecture Notes in Math.*, pages 267–318. Springer, Berlin, 2011.
- [15] D. Hobson and M. Klimmek. private communication, 2012.
- [16] D. Hobson and A. Neuberger. Robust bounds for forward start options. *Math. Finance*, 22(1):31–56, 2012.
- [17] N. Juillet. On displacement interpolation of measures involved in Brenier’s theorem. *Proc. Amer. Math. Soc.*, 139(10):3623–3632, 2011.
- [18] A. Kechris. *Classical descriptive set theory*, volume 156 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.
- [19] H. Kellerer. Markov-Komposition und eine Anwendung auf Martingale. *Math. Ann.*, 198:99–122, 1972.
- [20] H. Kellerer. Duality theorems for marginal problems. *Z. Wahrsch. Verw. Gebiete*, 67(4):399–432, 1984.
- [21] L. Mirsky. Majorization of vectors and inequalities for convex functions. *Monatsh. Math.*, 65:159–169, 1961.
- [22] G. Monge. Mémoire sur la théorie des déblais et des remblais. *Histoire de l’académie Royale des Sciences de Paris*, 1781.
- [23] J. Oblój. The Skorokhod embedding problem and its offspring. *Probab. Surv.*, 1:321–390, 2004.
- [24] L. Rüschendorf and S. T. Rachev. A characterization of random variables with minimum L^2 -distance. *J. Multivariate Anal.*, 32(1):48–54, 1990.
- [25] L. Rüschendorf and L. Uckelmann. Numerical and analytical results for the transportation problem of Monge-Kantorovich. *Metrika*, 51(3):245–258 (electronic), 2000.
- [26] S. Sherman. On a theorem of Hardy, Littlewood, Polya, and Blackwell. *Proc. Nat. Acad. Sci. U. S. A.*, 37:826–831; errata: 38, 382 (1952), 1951.
- [27] V. Strassen. The existence of probability measures with given marginals. *Ann. Math. Statist.*, 36:423–439, 1965.
- [28] C. Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2003.
- [29] C. Villani. *Optimal Transport. Old and New*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2009.

MATHIAS BEIGLBÖCK
 FAKULTÄT FÜR MATHEMATIK, UNIVERSITÄT WIEN
 NORDBERGSTRASSE 15, 1090 WIEN, AUSTRIA
E-mail address: mathias.beiglboeck@univie.ac.at

NICOLAS JUILLET
 INSTITUT DE RECHERCHE MATHÉMATIQUE AVANCÉE, UNIVERSITÉ DE STRASBOURG ET CNRS
 7 RUE RENÉ DESCARTES, 67000 STRASBOURG, FRANCE
E-mail address: nicolas.juillet@math.unistra.fr