MARTINGALE BENAMOU–BRENIER: A PROBABILISTIC PERSPECTIVE

J. BACKHOFF, M. BEIGLBÖCK, M. HUESMANN, AND S. KÄLLBLAD

Abstract. In classical optimal transport, the contributions of Benamou–Brenier and McCann regarding the time-dependent version of the problem are cornerstones of the field and form the basis for a variety of applications in other mathematical areas.

We suggest a Benamou–Brenier type formulation of the martingale transport problem for given d-dimensional distributions μ, ν in convex order. The unique solution $M^* = (M_t^*)_{t\in[0,1]}$ of this problem turns out to be a Markov-martingale which has several notable properties: In a specific sense it mimics the movement of a Brownian particle as closely as possible subject to the conditions $M_0^* \sim \mu, M_1^* \sim \nu$. Similar to McCann's displacement-interpolation, M^* provides a time-consistent interpolation between μ and ν. For particular choices of the initial and terminal law, M^* recovers archetypical martingales such as Brownian motion, geometric Brownian motion, and the Bass martingale. Furthermore, it yields a natural approximation to the local vol model and a new approach to Kellerer's theorem.

This article is parallel to the work of Huesmann-Trevisan, who consider a related class of problems from a PDE-oriented perspective.

Keywords: Optimal Transport, Martingales, Brenier's Theorem, Benamou-Brenier, cyclical monotonicity, causal transport, Knothe Rosenblatt, Schrödinger problem.
Mathematics Subject Classification (2010): Primary 60G42, 60G44; Secondary 91G20.

1. Introduction

The roots of optimal transport as a mathematical field go back to Monge [46] and Kantorovich [36] who established its modern formulation. Important triggers for its steep development in the last decades were the seminal results of Benamou, Brenier, and McCann [18, 19, 15, 45]. Today the field is famous for its striking applications in areas ranging from mathematical physics and PDE-theory to geometric and functional inequalities. We refer to [54, 55, 2, 51] for comprehensive accounts of the theory.

Recently there has also been interest in optimal transport problems where the transport plan must satisfy additional martingale constraints. Such problems arise naturally in robust finance, but are also of independent mathematical interest, for example they have important consequences for the study of martingale inequalities (see e.g. [17, 30, 49]) and the Skorokhod embedding problem [7, 35]. Early papers to investigate such problems include [33, 11, 53, 23, 22, 20], and this topic is commonly referred to as martingale optimal transport.

In view of the central role taken by the seminal results of Benamou, Brenier, and McCann on optimal transport for squared euclidean distance, the related continuous time transport problem and McCann's displacement interpolation, it is intriguing to search for similar concepts also in the martingale context. While [13, 31] propose a martingale version of Brenier's monotone transport mapping, our starting point is the Benamou-Brenier continuous time transport problem which we restate here for comparison with the martingale analogues that we will consider subsequently.

1.1. Benamou-Brenier transport problem and McCann-interpolation in probabilistic terms.

In view of the probabilistic nature of the results we present subsequently, it is convenient to recall some classical concepts and results of optimal transport in probabilisitic language. Given probabilities μ, ν in the space $\mathcal{P}_2(\mathbb{R}^d)$ of d-dimensional distributions with

MB gratefully acknowledges support by FWF-grant Y00782, MH by the Deutsche Forschungsgemeinschaft through the CRC 1060 *The Mathematics of Emergent Effects* and the *Hausdorff Center for Mathematics*.

finite second moment consider

$$T_2(\mu,\nu) := \inf_{X_t = X_0 + \int_0^t v_s\,ds, X_0\sim\mu, X_1\sim\nu} \mathbb{E}\left[\int_0^1 |v_t|^2\,dt\right]. \tag{BB}$$

Then by [18] we have

Theorem 1.1. *Let $\mu,\nu \in \mathcal{P}_2(\mathbb{R}^d)$ and assume that μ is absolutely continuous with respect to Lebesgue measure. Then* (BB) *has a unique optimizer X^*.*

Remark 1.2. In Theorem 1.1 (and similarly below) the solution to (BB) is unique in the sense that there exists a unique probability measure on the pathspace $C([0,1])$ such that the canonical/identity process optimizes (BB).

In probabilistic terms, McCann's displacement interpolation can be defined by $[\mu,\nu]_t := law\,(X_t^*)$ where $t \in [0,1]$ and μ, ν, X^* as in Theorem 1.1.

Theorem 1.3. *Let $\mu,\nu \in \mathcal{P}_2(\mathbb{R}^d)$ and assume that μ is absolutely continuous with respect to Lebesgue measure. Let $s,t,\lambda \in [0,1]$, $s<t$. Then*

$$\big[[\mu,\nu]_s, [\mu,\nu]_t\big]_\lambda = [\mu,\nu]_{(1-\lambda)s+\lambda t}. \tag{1.1}$$

Moreover

$$(t-s)\,T_2^{1/2}(\mu,\nu) = T_2^{1/2}([\mu,\nu]_s, [\mu,\nu]_t). \tag{1.2}$$

Finally, the optimizer of (BB) is given through the gradient of a convex function. More precisely, by [15], we have

Theorem 1.4. *Assume that μ is absolutely continuous with respect to Lebesgue measure and $\mu,\nu \in \mathcal{P}_2(\mathbb{R}^d)$. A candidate process X, $X_0 \sim \mu, X_1 \sim \nu$ is an optimizer if and only if $X_1 = f(X_0)$, where f is the gradient of a convex function $\varphi : \mathbb{R}^n \to \mathbb{R}$ and all particles move with constant speed, i.e. $X_t = tX_1 + (1-t)X_0 = X_0 + t(X_1 - X_0)$.*

1.2. **Martingale counterparts.** Let $\mu,\nu \in \mathcal{P}_2(\mathbb{R}^d)$ be in convex order (denoted $\mu \preceq_c \nu$) and write B for Brownian motion on $\mathbb{R}^d$. We consider the optimization problem

$$MT(\mu,\nu) := \sup_{\substack{M_t = M_0 + \int_0^t \sigma_s\,dB_s \\ M_0\sim\mu, M_1\sim\nu}} \mathbb{E}\left[\int_0^1 \mathrm{tr}(\sigma_t)\,dt\right], \tag{MBB}$$

see the restatement of (MBB) below for more details. We have

Theorem 1.5. *Assume that $\mu,\nu \in \mathcal{P}_2(\mathbb{R}^d)$ satisfy $\mu \preceq_c \nu$. Then* (MBB) *has a unique optimizer M^*.*

At its face, the optimization problems (BB) and (MBB) look rather different. However it is not hard to see that both problems are equivalent to optimization problems that are much more obviously related. In Section 6 below we establish that

$$X^* = \mathrm{argmin}_{X_0\sim\mu, X_1\sim\nu} W^2(X, \text{constant speed particle}), \tag{1.3}$$

$$M^* = \mathrm{argmin}_{M_0\sim\mu, M_1\sim\nu} W_c^2(M, \text{constant volatility martingale}), \tag{1.4}$$

where W^2 denotes Wasserstein distance with respect to squared Cameron-Martin norm, while W_c^2 denotes its *causal* analogue[1] (in the terminology of Lasalle [41]).

The reformulation in (1.4) allows for the following interpretation: M^* is the process whose evolution follows the movement of a Brownian particle as closely as possible subject to the marginal conditions $M_0 \sim \mu, M_1 \sim \nu$. This motivates the name in following definition.

[1]Causal transport plans generalize adapted processes in the same way as classical Kantorovich transport plans extend Monge maps.

Definition 1.6. *Let μ, ν, M^* be as in Theorem 1.5. Then we call M^* the* stretched Brownian motion *(sBm) from μ to ν. We define the* martingale displacement interpolation *by*

$$[\mu, \nu]^M_t := law\, M^*_t\,, \tag{1.5}$$

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

In analogy to Theorem 1.3 we have

Theorem 1.7. *Assume that $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ satisfy $\mu \preceq_c \nu$. Let $s, t, \lambda \in [0, 1]$, $s < t$. Then*

$$\left[[\mu, \nu]^M_s, [\mu, \nu]^M_t\right]^M_\lambda = [\mu, \nu]^M_{(1-\lambda)s+\lambda t}\,. \tag{1.6}$$

Moreover

$$(t - s)\, MT^2\,(\mu, \nu) = MT^2([\mu, \nu]^M_s, [\mu, \nu]^M_t). \tag{1.7}$$

1.3. **Structure of stretched Brownian motion.** In the solution of the classical Benamou–Brenier transport problem, particles travel with constant speed along straight lines. In contrast, we will see that in the case of sBm the movement of individual particles mimics that of Brownian motion. Broadly speaking, the "direction" of these particles will be determined – similar to the classical case – by a mapping which is the gradient of a convex function.

For simplicity, we first consider the particular case where $\mu, \nu, \mu \preceq_c \nu$ are probabilities on the real line and μ is concentrated in a single point, i.e. $\mu = \delta_m$ where m is the center of ν. It turns out that in this case sBm M^* is precisely the "Bass martingale" [6] (or 'Brownian martingale') with terminal distribution ν. We briefly recall its construction: Pick $f : \mathbb{R} \to \mathbb{R}$ increasing such that $f(\gamma) = \nu$, where γ is the standard Gaussian distribution on $\mathbb{R}$. Then set for $t \in [0, 1]$

$$M_t := \mathbb{E}[f(B_1)|\mathcal{F}_t] = \mathbb{E}[f(B_1)|B_t] = f_t(B_t), \tag{1.8}$$

where $B = (B_t)_{t\in[0,1]}$ denotes Brownian motion started in $B_0 \sim \delta_0$, $(\mathcal{F}_t)_{t\in[0,1]}$ the Brownian filtration and $f_t(b) := \int f(b + y)\, d\gamma_{1-t}(y)$, $\gamma_s \sim N(0, s)$. Clearly M is a continuous Markov martingale such that $M_0 \sim \delta_m, M_1 \sim \nu$. As a particular consequence of the results below we will see that M is a stretched Brownian motion.

To state our results for the general case we need to consider an extension of the Bass construction. Let $F : \mathbb{R}^d \to \mathbb{R}$ be a convex function and set

$$f_t(b) = \int \nabla F(b + y)\gamma^d_{1-t}(dy), \tag{1.9}$$

where γ^d_s denotes the centered d-dimensional Gaussian with covariance matrix $s\,\mathrm{Id}$. If B denotes d-dimensional Brownian motion started in $B_0 \sim \alpha$, we have

$$\mathbb{E}[\nabla F(B_1)|\mathcal{F}_t] = f_t(B_t),\ t \in [0, 1]. \tag{1.10}$$

Definition 1.8. *A continuous $\mathbb{R}^d$-valued martingale M is a* standard stretched Brownian motion (s^2Bm) *from μ to ν if there exist a probability measure α on $\mathbb{R}^d$ and a convex function $F : \mathbb{R}^d \to \mathbb{R}$ with $\nabla F(\alpha * \gamma^d) = \nu$, such that*

$$M_t = E[\nabla F(B_1)|\mathcal{F}_t] \text{ and } M_0 \sim \mu,$$

where B is a Brownian motion with $B_0 \sim \alpha$.

Note, that for $\alpha, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ there exists a convex function F with $\nabla F(\alpha * \gamma^d) = \nu$ and F is $\alpha * \gamma^d$-unique up to an additive constant. (This is a consequence of Brenier's Theorem, see e.g. Theorem 1.1 or [54, Theorem 2.12].)

Remark 1.9. Both Brownian motion and geometric Brownian motion are examples of standard stretched Brownian motion.

We have the following results

Theorem 1.10. *Let $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ with $\mu \preceq_c \nu$. If M is a standard stretched Brownian motion from μ to ν, then M is an optimizer of* (MBB)*, i.e. M is the stretched Brownian motion from μ to ν.*

Theorem 1.11. *Let $\mu, \nu \in \mathcal{P}_2(\mathbb{R}^d)$ with $\mu \preceq_c \nu$. Let M^* be the stretched Brownian motion from μ to ν, i.e. the optimizer of* (MBB)*. Write $M^{*,x}$ for the martingale M conditioned on starting in $M_0 = x$. Then for μ-a.a. $x \in \mathbb{R}^d$ the martingale $M^{*,x}$ is a standard stretched Brownian motion.*

As a particular consequence of these results, the notions sBm and s^{2}Bm coincide if μ is concentrated in a single point. However the relation between sBm and s^{2}Bm is more complicated in general: A notable intricacy of the martingale transport problem is caused by the fact that, loosely speaking, certain regions of the space do not communicate with each other.

Consider for a moment the particular case where μ, ν are distributions on the real line. In this instance, a martingale transport problem can be decomposed into countably many "minimal" components and on each of these components the behaviour of the problem is fairly similar to the classical transport problem. We refer the reader to Section 3.1 for the precise definition and only provide an illustrative example at this stage.

Example 1.12. Let $\mu := 1/2(\lambda_{|[-3,-2]} + \lambda_{|[2,3]})$, $\nu := 1/6(\lambda_{|[-4,-1]} + \lambda_{|[1,4]})$. Then *any* martingale $M, M_0 \sim \mu, M_1 \sim \nu$ will satisfy the following: If $M_0 > 0$, then $M_1 > 0$ and if $M_0 \leq 0$ then $M_1 \leq 0$. I.e. the positive and the negative halfline do not "communicate," and a problem of martingale transport should be considered on either of these parts of space separately.

If the pair (μ, ν) decomposes into more than one minimal component, as in the previous example, there exists no s^{2}Bm from μ to ν. However for the one-dimensional case we will establish the following: A martingale is a sBm if and only if it behaves like a s^{2}Bm on each minimal component, see Theorem 3.1.

Notably, the challenges posed by non-communicating regions appear much more intricate for dimension $d \geq 2$, see the deep contributions of Ghoussoub–Kim–Lim [25], DeMarch–Touzi [21] and Obłój–Siorpaes [48]. In particular it is not yet fully understood how to break up a martingale transport problem into distinct pieces which mimic the behaviour of minimal components in the one dimensional case.

Below we will give special emphasis to the case $d = 2$ under the additional regularity assumption that ν is absolutely continuous. This instance seems of particular interest since it allows to recognize the geometric structure of the problem while avoiding the more intricate effects of non-minimality which are present in higher dimension. Based on the results of [21, 48] and a particular 'monotonicity principle' we will be able to largely recover the main one-dimensional result (Theorem 3.1) in the two-dimensional case, see Sections 3.2-3.3 below and specifically Theorem 3.12 therein. We conjecture that a similar structural characterization of sBm can be established in general dimensions, pending future developments in the direction of [21, 48].

1.4. **Further remarks.**

1.4.1. *Discrete time version and monotonicity principle.* The classical Benamou–Brenier transport formulation immediately reduces to the familiar discrete time transport problem for squared distance costs. Similarly, the martingale version (MBB) can be reformulated as discrete time problem, more precisely, a transport problem for general transport costs in the sense of [28].

The discrete time reformulation of (MBB) plays an important role in the derivation of our main results. To analyze the discrete problem we introduce a "monotonicity principle" for general transport costs. The origin of this approach is the characterization of optimal transport plans in terms of c-cyclical monotonicity. In optimal transport, the potential of this concept has been recognized by Gangbo–McCann [24]. More recently variants of this

idea have proved to be useful in a number of related situation, see [40, 10, 56, 29, 13, 7, 47, 14] among others. In view of this, it seems possible that the monotonicity principle for general transport costs could also be of interest in its own right.

1.4.2. *Schrödinger problem.* Our variational problem (MBB) is reminiscent of the celebrated Schrödinger problem, in which the idea is to minimize the relative entropy with respect to Wiener measure (or other Markov laws) over path-measures with fixed initial and final marginals. We refer to the survey [42] and the references therein. Among the similarities, let us mention that the solution to the Schrödinger problem is unique and is a Markov law, and furthermore this problem also has a transport-like discrete time reformulation which is fundamental to the dynamic path-space version. On the other hand, (MBB) and the Schrödinger problem are in particular sense at opposing ends of probabilistic variational problem, we optimize over "volatilities keeping the drift fixed" whereas the latter optimizes over "drifts keeping the volatility fixed."

1.4.3. *Bass-martingale and Skorokhod embedding.* The Bass-martingale (1.8) was used by Bass [6] to solve the Skorokhod embedding problem. Hobson asked whether there are natural optimality properties related to this construction and if one could give a version with a non trivial starting law. (MBB) yields such an optimality property of the Bass construction and stretched Brownian motion gives rise to a version of the Bass embedding with non trivial starting law. Notably a characterization of the Bass martingale in terms of an optimality property was first obtained in [8], the variational problem considered in that article refers to measure valued martingales and appears rather different from the one considered in (MBB).

1.4.4. *Geometric Brownian motion.* From the above results it is clear that Brownian motion is (up to an appropriate scaling of time) a s^2Bm between any of its marginals. We find it notable that same applies in the case of *geometric* Brownian motion.

1.4.5. *Kellerer's theorem and Lipschitz kernels.* Kellerer's theorem [39] states that if a family of distributions $(\mu_t)_{t\in[0,1]}$ on the real line satisfies $s \le t \Rightarrow \mu_s \preceq_c \mu_t$, there exists a Markovian martingale $(X_t)_{t\in\mathbb{R}_+}$ with $law\,(X_t) = \mu_t$ for every t. In contemporary terms (see [32]), $(\mu_t)_{t\in\mathbb{R}_+}$ is called a *peacock* and $(X_t)_{t\in\mathbb{R}_+}$ is a Markovian martingale associated to this peacock.

The technically most involved part in establishing Kellerer's theorem is to prove that for $\mu \preceq_c \nu$ there exists a martingale transition kernel P having the following *Lipschitz-property*: A kernel $P : x \mapsto \pi_x$, $\nu(dy) = \int \mu(dx)\pi_x(dy)$ is called Lipschitz (or more precisely 1-Lipschitz) if $\mathcal{W}_1(\pi_x, \pi_{x'}) \le |x - x'|$ for all x, x'. Kellerer's proof of the existence of Lipschitz-kernels is not constructive and employs Choquet's theorem. Other proofs are based on solutions to the Skorokhod problem for non-trivial starting law, see [43, 12].

Stretched Brownian motion yields a new construction of a Lipschitz-kernel: Given probabilities $\mu, \nu, \mu \preceq_c \nu$ on the real line and writing M^* for sBm from μ to ν, it is straightforward to see that $\pi := law\,(M_0^*, M_1^*)$ is a Lipschitz kernel.

The question whether Kellerer's theorem can be extended to the case of marginal measures on $\mathbb{R}^d, d \ge 2$ remains open. While all previously known constructions of kernels used for the proof of Kellerer's theorem were inherently limited to dimension $d = 1$, the approach sketched above seems more susceptible to generalization. We intend to pursue this question further in future work.

1.4.6. *Almost continuous diffusions / local volatility model.* Assume that $(\mu_t)_{t\in[0,1]}$ (where $\mu_t, t \in [0, 1]$ are probabilities on the real line) is a peacock such that $t \mapsto \mu_t$ is continuous in the weak topology. Lowther [43] establishes that an appropriate continuity condition makes the Markov martingale appearing in Kellerer's theorem unique. In his terms, there is a unique "almost continuous" martingale diffusion M^{ac} such that $M_t^{ac} \sim \mu_t, t \in [0, 1]$. Under further regularity conditions, M^{ac} is precisely Dupire's local volatility model.

Stretched Brownian motion yields a simple approximation scheme to M^{ac}. Write M^n for the Markov martingale satisfying that for each $k \in \{1, \dots, n\}$, $(M_t^n)_{t\in[(k-1)/n,k/n]}$ is (modulo

the obvious affine time-change) stretched Brownian motion between $\mu_{(k-1)/n}$ and $\mu_{k/n}$. M^n is then a continuous diffusion and based on Lowther's [43, 44] it is straightforward that

$$M^{ac} = \lim_{n\to\infty} M^n, \tag{1.11}$$

where the limit is in the sense of convergence of finite dimensional distribution (cf. [12]).

1.4.7. *Lévy processes.* Many arguments in this article rely only on the independence and stationarity of increments of Brownian motion. Therefore a problem similar to (MBB), but based on a reference Lévy process instead, should conceivably exhibit similar properties as we find in the Brownian case. In this direction it could be an interesting question to identify the outcome of the approximation procedure described in (1.11).

1.4.8. *Dual problem, related work.* Optimization problems similar to (MBB) were first studied from a general perspective by Tan and Touzi [53], in particular establishing a duality theory for these type of problems. The dual viewpoint is also emphasized in [34], which is parallel to the present work. Among other results, [34] derives a PDE that yields a sufficient condition for a flow of measures to optimize (MBB) or related cost criteria.

1.5. **Outline of the article:** In Section 2 we introduce the discrete-time variant of our optimization problem. We also prove some of the multidimensional results stated in the introduction, as well as introduce further properties of sBm (dynamic programming principle for (MBB), the Markov property of sBm). In Section 3 we state our main results regarding the structure of sBm in dimensions one and two. In Section 4 we present a monotonicity principle for generalized transport costs, which is crucial for our analysis in dimension two, but may also be of independent interest. In Section 5 we conclude the proofs of our main results. Finally in Section 6 we give further optimality properties of sBm and s^2Bm in terms of a (causal) optimal transport problem between martingale laws.

1.6. **Notation:** The set of probability measures on a set X will be denoted by $\mathcal{P}(\mathsf{X})$. For $\rho_1, \rho_2 \in \mathcal{P}(\mathsf{X})$ we write $\Pi(\rho_1, \rho_2)$ for the set of all couplings of ρ_1 and ρ_2, i.e. all measures on the product space with marginals ρ_1 and ρ_2 resp. $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$ are said to be in convex order, short $\mu \preceq_c \nu$ iff for all convex real valued functions φ it holds that $\int \varphi\, d\mu \le \int \varphi\, d\nu$. We fix $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$, assume that $\mu \preceq_c \nu$ and that both measures have finite second moment.

We denote by $\mathsf{M}(\mu, \nu)$ the set of all martingale couplings with marginals μ and ν (which is non-empty by Strassen's Theorem [52]), i.e.

$$\mathsf{M}(\mu,\nu) := \{\pi \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d) : \mathbb{E}^\pi[(y-x)h(x)] = 0 \text{ for all } h : \mathbb{R}^d \to \mathbb{R} \text{ Borel bounded}\}.$$

For a generic measure π on $\mathbb{R}^d \times \mathbb{R}^d$ we denote by $(\pi_x)_{x\in\mathbb{R}^d}$ the conditional transition kernel given the first coordinate or equivalently its disintegration w.r.t. the first marginal. For $\rho \in \mathcal{P}(\mathsf{X})$ and a measurable map $f : \mathsf{X} \to \mathsf{Y}$ we write $f(\rho) = \rho \circ f^{-1}$ for the pushforward of ρ under f.

For a set $A \subseteq \mathbb{R}^d$ we denote by $\mathrm{aff}(A)$ the smallest affine vector space containing it, $\dim(A)$ the dimension of $\mathrm{aff}(A)$, $\mathrm{ri}(A)$ the relative interior of A (i.e. interior of A with respect to the relative topology of $\mathrm{aff}(A)$ as inherited from the usual topology in $\mathbb{R}^d$), and $\partial A := \overline{A}\backslash\mathrm{ri}(A)$ the relative boundary. By $\mathrm{co}(A)$ and $\overline{\mathrm{co}}(A)$ we denote the convex hull and the closed convex hull of A respectively. The relative face of A at a is defined by $\mathrm{rf}_a(A) = \{y \in A : (a - \varepsilon(y-a), y + \varepsilon(y-a)) \subseteq A, \text{ some } \varepsilon > 0\}$. For a set $\Gamma \subseteq \mathbb{R}^d \times \mathbb{R}^d$ we denote $\Gamma_x := \{y : (x,y) \in \Gamma\}$ and $\mathrm{proj}_1(\Gamma)$ the projection of Γ onto the first coordinate. Given $\pi \in \mathsf{M}(\mu,\nu)$ we say that $\Gamma \subseteq \mathbb{R}^d \times \mathbb{R}^d$ is a martingale support for π if $\pi(\Gamma) = 1$ and $x \in \mathrm{ri}(\Gamma_x)$ for μ-a.e. x.

Finally, we denote by $\lambda^d, \gamma^d, \gamma^d_t$ resp. the Lebesgue, standard Gaussian, and the Gaussian measure with covariance matrix $t\,\mathrm{Id}$ in $\mathbb{R}^d$, and reserve the symbol $*$ for convolution.

2. Refined and auxiliary results in arbitrary dimensions

We start by restating our main optimization problem in more precise form.

$$MT := MT(\mu,\nu) := \sup_{M_t = M_0 + \int_0^t \sigma_s\, dB_s, M_0 \sim \mu, M_1 \sim \nu} \mathbb{E}\left[\int_0^1 \mathrm{tr}(\sigma_t)\, dt\right]. \tag{MBB}$$

Here the supremum is taken over the class of all filtered probability spaces $(\Omega, \mathcal{F}, \mathbb{P})$, with σ an $\mathbb{R}^{d\times d}$-valued $\mathcal{F}$-progressive process and B a d-dimensional $\mathcal{F}$-Brownian motion, such that M is a martingale. In fact, as a particular consequence of Theorem 2.2, the choice of the underlying probability space is not relevant, provided that $(\Omega, \mathcal{F}, \mathbb{P})$ is rich enough to support a $\mathcal{F}_0$-measurable random variable with continuous distribution.

By Doob's martingale representation theorem (see e.g. [37, Theorem 4.2]), the supremum above is the same if we optimized over all continuous d-dimensional local martingales from μ to ν with absolutely continuous cross variation matrix (one then replaces the cost by the trace of the root of the Radon Nikodym density of said matrix).

We will be also interested in a "static" version of the above problem, just as the Benamou-Brenier formula is associated to the static optimal transport problem with quadratic cost

$$GT := GT(\mu,\nu) := \sup_{\substack{\{\pi_x\}_x,\, \mathrm{mean}(\pi_x)=x \\ \int \mu(dx)\pi_x(dy)=\nu(dy)}} \int \mu(dx) \sup_{q\in\Pi(\pi_x,\gamma^d)} \int q(dm,db)\; m\cdot b\,. \tag{GOT}$$

The tag (GOT) reflects the fact that this is a general optimal transport problem (the cost function is non-linear in the optimization variable).

Remark 2.1. Completing the square in (GOT) yields

$$1 + \int |y|^2\, d\nu - 2\,GT = \inf_{\substack{\{\pi_x\}_x,\, \mathrm{mean}(\pi_x)=x \\ \int \mu(dx)\pi_x(dy)=\nu(dy)}} \int \mu(dx) \mathcal{W}_2(\pi_x,\gamma^d)^2, \tag{2.1}$$

where $\mathcal{W}_2$ is the usual L^2 Wasserstein distance on $\mathcal{P}(\mathbb{R}^d)$. The r.h.s. of (2.1) is clearly a general/weak transport problem in the setting of Gozlan et. al. [28, 27].

We start by establishing the link between the static and dynamic problems introduced so far, and moreover, establish the uniqueness of optimizers in either case. As a corollary, this proves Theorem 1.5 in the introduction.

Theorem 2.2. *The static and the dynamic problems* (GOT) *and* (MBB) *are equivalent. More precisely,*

(1) $GT = MT < \infty$,
(2) (GOT) *has a unique optimizer* π^**;*
(3) (MBB) *has a unique optimizer* M^**;*
(4) $\pi^* = law\,(M_0^*, M_1^*)$ *and* $M^* = G(\pi^*)$ *for some function G, i.e.* M^* *can be explicitly constructed from* π^*.

Proof. Let M be feasible for (MBB). By Itô's formula and the martingale property of M we have

$$\mathbb{E}\left[\int_0^1 \mathrm{tr}(\sigma_t)\, dt\right] = \mathbb{E}[M_1\cdot B_1 - M_0\cdot B_0] = \mathbb{E}[M_1\cdot(B_1-B_0)] = \mathbb{E}[\;\mathbb{E}[M_1\cdot(B_1-B_0)\,|\,M_0]\;].$$

Letting $q_x = law\,(M_1, B_1 - B_0 \,|\, M_0 = x)$ we find $q_x \in \Pi(\pi_x\,,\,\gamma^d)$ for $\pi_x = law\,(M_1 \,|\, M_0 = x)$ and

$$\mathbb{E}\left[\int_0^1 \mathrm{tr}(\sigma_t)\, dt\right] = \int \mu(dx) \int q_x(dm,db)\; m\cdot b.$$

From this we easily conclude $GT \le MT$.

Now let π be feasible for (GOT). For each x we can find $F^x(\cdot)$ convex such that $\nabla F^x(\gamma^d) = \pi_x$. We now define $M_t^x := E[\nabla F^x(B_1)|\mathcal{F}_t^B]$ for a given standard Brownian motion on $\mathbb{R}^d$ with Brownian filtration $\mathcal{F}^B$. Construct a filtration $\mathcal{F}$ by adding an independent random variable X to $\mathcal{F}^B$, with $X \sim \mu$. Since $M_0^x = \int y\pi_x(dy) = x$ and

$\int \mu(dx)\pi_x(dy) = \nu(dy)$ we conclude that $\{M_t^X\}_{t\in[0,1]}$ is a continuous martingale from μ to ν. By construction

$$\int \mu(dx) \sup_{q\in\Pi(\pi_x,\gamma^d)} \int q(dm,db)\ m\cdot b = \int \mu(dx) \int \gamma^d(db)\ b\cdot\nabla F^x(b) = \mathbb{E}\Big[\,\mathbb{E}\Big[B_1\cdot M_1^X|X\Big]\,\Big],$$

and the last term equals $\mathbb{E}\Big[\int_0^1 \mathrm{tr}(\sigma_t)dt\Big]$ as before (σ can easily be computed from ∇F^x). This proves $GT \geq MT$ and hence $GT = MT$. The finiteness $\infty > GT$ follows from $m\cdot b \leq |m|^2 + |b|^2$ and ν and γ having finite second moment; see (GOT).

To show that (GOT) is attained let us denote by $(\pi^n)_{n\in\mathbb{N}}$ with $\pi^n(dx,dy) = \pi_x^n(dy)\mu(dy)$ an optimizing sequence. The set $\Pi(\mu,\nu)$ is weakly compact in $\mathcal{P}(\mathbb{R}^d\times\mathbb{R}^d)$. Moreover, the convex subset $\mathsf{M}(\mu,\nu)$ is weakly closed (hence weakly compact), e.g. [54, Theorem 7,12 (iv)]. By [5, Theorem 3.7] we obtain the existence of a measurable kernel $x\mapsto\pi_x\in\mathcal{P}(\mathbb{R}^d)$ and a subsequence, still denoted by $(\pi^n)_n$, such that on a μ-full set

$$\tfrac{1}{N}\textstyle\sum_{n\leq N}\pi_x^n(dy) \to \pi_x(dy),$$

with respect to weak convergence in $\mathcal{P}(\mathbb{R}^d)$. In particular $\frac{1}{N}\sum_{n\leq N}\pi^n \to \pi$ in the weak topology in $\mathcal{P}(\mathbb{R}^d\times\mathbb{R}^d)$, where $\pi(dx,dy) := \mu(dx)\pi_x(dy)$. Since $\mathsf{M}(\mu,\nu)$ is closed, we have that $\pi\in\mathsf{M}(\mu,\nu)$. Finally,

$$\begin{aligned}
GT &= \lim_n \textstyle\int \mu(dx) \sup_{q\in\Pi(\pi_x^n,\gamma^d)} \int q(dm,db)\ m\cdot b\\
&= \lim_N \textstyle\int \mu(dx) \frac{1}{N}\sum_{n\leq N} \sup_{q\in\Pi(\pi_x^n,\gamma^d)} \int q(dm,db)\ m\cdot b\\
&\leq \lim_N \textstyle\int \mu(dx) \sup_{q\in\Pi(\frac{1}{N}\sum_{n\leq N}\pi_x^n,\gamma^d)} \int q(dm,db)\ m\cdot b\\
&\leq \textstyle\int \mu(dx) \limsup_N \sup_{q\in\Pi(\frac{1}{N}\sum_{n\leq N}\pi_x^n,\gamma^d)} \int q(dm,db)\ m\cdot b\\
&\leq \textstyle\int \mu(dx) \sup_{q\in\Pi(\pi_x,\gamma^d)} \int q(dm,db)\ m\cdot b \quad \leq GT\,.
\end{aligned}$$

The first inequality holds by concavity of $\eta\mapsto H(\eta) := \sup_{q\in\Pi(\eta,\gamma^d)}\int q(dm,db)\ m\cdot b$ w.r.t. convex combinations of measures. The second inequality is Fatou's lemma, noticing that the integrand is bounded in $L^1(\mu)$ (the bound equals the sum of the second moments of μ and γ). The third inequality follows by weak convergence of the averaged kernel on a μ-full set and upper semicontinuity of $H(\cdot)$. For uniqueness it suffices to notice that $H(\cdot)$ is actually strictly concave, which is an easy consequence of Brenier's Theorem. Hence, (GOT) is attained and we denote the unique optimizer by π^*.

Taking π^* we may build an optimizer M^* for (MBB) as in the first part of the proof (as the value of both problems agree).

We finally establish the uniqueness of optimizers for (MBB). Let $\tilde{M}$ be any such optimizer. From the previous considerations, we deduce that the law of $(\tilde{M}_0,\tilde{M}_1)$ is the unique optimizer π^* of (GOT). Thus conditioning on $\{\tilde{M}_0 = x\}$ we have that $\tilde{M}$ connects δ_x to π_x^*. It follows that $\mu(dx)$-a.s. $\tilde{M}$ conditioned on $\{\tilde{M}_0 = x\}$ is optimal between these marginals. Indeed,

$$\sup_{N_t = x+\int_0^t \sigma_s\, dB_s,\, N_1\sim\pi_x^*} \mathbb{E}\Big[\int_0^1 \mathrm{tr}(\sigma_t)\,dt\Big] = \sup_{q\in\Pi(\pi_x^*,\gamma^d)} \int q(dm,db)\ m\cdot b\ , \tag{2.2}$$

by the results obtained so far, since if $\tilde{M}$ conditioned on $\{\tilde{M}_0 = x\}$ was not optimal for the l.h.s. it could not deliver the equality $MT = GT$. So it suffices to show that the l.h.s. of (2.2) is uniquely attained. But any candidate martingale N with volatility σ satisfies $\mathbb{E}[\int_0^1 \mathrm{tr}(\sigma_t)dt] = \mathbb{E}[N_1B_1]$. Hence, Brenier's Theorem implies that $\tilde{M}_1 = \nabla F^x(B_1)$ on $\{\tilde{M}_0 = x\}$, for a convex function F^x. Since ∇F^x is unique, we finally get $\tilde{M} = M^*$. □

Remark 2.3. The proof of Theorem 2.2 shows how to build the optimizer for (MBB) via the following procedure, making the statement $M^* = G(\pi^*)$ there precise:

(1) Find the unique optimizer π^* of (GOT).
(2) Find convex functions F^x such that $\nabla F^x(\gamma^d) = \pi_x^*$.
(3) Define $M_t^x := \mathbb{E}[\nabla F^x(B_1)|B_t] = \int \nabla F^x(y+B_t)\gamma_{1-t}^d(dy)$.
(4) Take $X\sim\mu$ independent of B and let $M_t := M_t^X$.

In particular, this proves Theorem 1.11 in the introduction.

We now establish further properties of the optimizer M^* of (MBB), which hold likewise in any number of dimensions. The first two of them will be important for the proofs of the results yet to come, namely that (MBB) obeys a dynamic programming principle and that M^* is a strong Markov martingale. The final property, that M^* is an "optimal constant-speed" interpolation between its marginals, is crucial for the interpretation of our martingale as an analogue of displacement interpolation in classical transport and in particular proves Theorem 1.7 in the introduction.

Let us define

$$V(t,T,\mu,\nu) := \sup_{\substack{M_s = X + \int_t^s \sigma_r\, dB_r,\ t\le s\le T \\ X\sim\mu,\ M_T\sim\nu}} \mathbb{E}\left[\int_t^T \mathrm{tr}(\sigma_r)\, dr\right], \tag{2.3}$$

so that $MT = V(0,1,\mu,\nu)$.

Lemma 2.4 (Dynamic programming principle)**.** *For every $t \in [0,1]$*

$$\begin{aligned} &V(0,1,\mu,\nu) = \\ &\sup\left\{\mathbb{E}\left[\int_0^t tr(\sigma_r)dr\right] + V(t,1,law\,(M_t),\nu)\ :\ M_s = M_0 + \int_0^s \sigma_r dB_r,\ 0\le s\le t,\ M_0\sim\mu\right\}, \end{aligned} \tag{2.4}$$

with the convention that $\sup\emptyset = -\infty$*. In particular if M^* is optimal for $V(0,1,\mu,\nu)$, then:*

(1) $M^*|_{[t,1]}$ is optimal for $V(t,1,law\,(M_t^*),\nu)$,
(2) $M^*|_{[0,t]}$ is optimal for $V(0,t,\mu,law\,(M_t^*))$
(3) A.s. we have $law\,(M_1^*|M_0^*,M_t^*) = law\,(M_1^*|M_t^*)$.

Proof. Obviously the l.h.s. is smaller than the r.h.s. Take now M^1 feasible for the r.h.s. (so that M^1 is adapted to a filtration $\{\mathcal{F}_s^1\}_{s\le t}$, B is a Brownian motion on $[0,t]$ adapted to it, and $dM^1 = \sigma^{(1)}dB$). Let M^2 be optimal for $V(t,1,law\,(M_t^1),\nu)$. By Remark 2.3 we may build M^2 from the starting distribution M_t^1 and the filtration $\mathcal{F}^2$ of this random variable and a Brownian motion W independent of $\mathcal{F}^1$ (and so of M_t^1), so $dM^2 = \sigma^{(2)}dW$. We build then a continuous martingale M on $[0,1]$ by setting it to M^1 on $[0,t]$ and M^2 on $(t,1]$, obtaining easily that

$$\mathbb{E}\left[\int_0^t \mathrm{tr}(\sigma_r^{(1)})dr\right] + V(t,law\,(M_t^1),\nu) = \mathbb{E}\left[\int_0^t \mathrm{tr}(\sigma_r^{(1)})dr + \int_t^1 \mathrm{tr}(\sigma_r^{(2)})dr\right].$$

Observing that $\tilde{B}_s = 1_{[0,t]}(s)B_s + 1_{(t,1]}(s)[B_t + W_s - W_t]$ is a Brownian motion for the concatenation of filtrations $\mathcal{F}^1$ and $\mathcal{F}^2$, and $dM = (1_{[0,t]}(s)\sigma_s^{(1)} + 1_{(t,1]}(s)\sigma_s^{(2)})\, d\tilde{B}$, then the r.h.s. above is the cost of M as a martingale starting at μ and ending at ν, and so is smaller than $V(0,1,\mu,\nu)$.

Let M^* be optimal for $V(0,1,\mu,\nu)$. Using (2.4) it is trivial to show Points (1)-(2). But from this follows that $M^*|_{[0,t]}$ is optimal for the r.h.s. of (2.4). This, Point (1), and the arguments in the previous paragraph show how to stitch together $M^*|_{[0,t]}$ and $M^*|_{[t,1]}$ to produce an optimizer M for $V(0,1,\mu,\nu)$. But this must then coincide with M^*, by uniqueness. On the other hand M_1 is defined via M_t^* and a Brownian motion independent of $\{M_s^* : s\le t\}$, so $law\,(M_1|M_0^*,M_t^*) = law\,(M_1|M_t^*)$ and we conclude. □

Corollary 2.5. *The unique optimizer M^* of (MBB) has the strong Markov property.*

Proof. The arguments in Lemma 2.4 actually yield that $law\,(M_1^*|M_s^*, s\le t) = law\,(M_1^*|M_t^*)$. It is straightforward to extend this to the actual Markov property, and from this, to the strong Markov property. □

Proposition 2.6. *Let M^* be the optimizer of (MBB) and denote*

$$[\mu,\nu]_t^M := law\,(M_t^*).$$

Then law (M_0^*, M_t^*) *is optimal for* (GOT) *between the marginals* μ *and* $[\mu,\nu]_t^M$*. Similarly, the optimizer of* (MBB) *between the same marginals is the time-changed martingale* $s \in [0,1] \mapsto M_{st}^*$*. Finally, for* $0 \le r \le t \le 1$*, we have*

$$GT([\mu,\nu]_r^M, [\mu,\nu]_t^M) = MT([\mu,\nu]_r^M, [\mu,\nu]_t^M) = \sqrt{t-r}\, MT(\mu,\nu) = \sqrt{t-r}\, GT(\mu,\nu). \tag{2.5}$$

Proof. We use the notation in Remark 2.3, further writing $M_t^* = M_t^X = f_t^X(B_t)$ with

$$f_t^x(\cdot) := \int \nabla F^x(b+\cdot)\gamma_{1-t}^d(db).$$

Since $[\mu,\nu]_t = f_t^X(\sqrt{t}B_1)$, it is not difficult to see that

$$N_s^* := \mathbb{E}[f_t^X(\sqrt{t}B_1)|\mathcal{F}_s^B] = f_{st}^X(\sqrt{t}B_s),$$

is the optimizer of (MBB) from μ at $s=0$ to $[\mu,\nu]_t^M$ at $s=1$. Of course N^* coincides (in law) with the time-changed martingale $s \mapsto M_{st}^*$, and by Theorem 2.2 we get the optimality of *law* (M_0^*, M_t^*). We next remark that $J(f_s^X)(B_s)$ is a matrix-valued martingale, where J stands for Jacobian, as can be easily seen from the convolution structure or PDE arguments. Thus $\mathbb{E}[J(f_s^X)(B_s)] = \mathbb{E}[J(f_{st}^X)(\sqrt{t}B_s)]$. To recognize the "$\sigma$" of N^* and M^* we observe that

$$\begin{aligned} dN_s^* &= \sqrt{t}J(f_{st}^X)(\sqrt{t}B_s)dB_s, \\ dM_s^* &= J(f_s^X)(B_s)dB_s, \end{aligned}$$

by Itô formula. Putting all together we find

$$\begin{aligned} \mathbb{E}\Big[\int_0^1 \sqrt{t}J(f_{st}^X)(\sqrt{t}B_s)ds\Big] &= \sqrt{t}\int_0^1 \mathbb{E}\Big[J(f_{st}^X)(\sqrt{t}B_s)\Big]ds \\ &= \sqrt{t}\int_0^1 \mathbb{E}\Big[J(f_s^X)(B_s)\Big]ds \\ &= \sqrt{t}\mathbb{E}\Big[\int_0^1 J(f_s^X)(B_s)ds\Big], \end{aligned}$$

and again by Theorem 2.2 we get

$$MT([\mu,\nu]_0^M, [\mu,\nu]_t^M) = \sqrt{t}\, MT(\mu,\nu).$$

The general case of (2.5) follows similarly. □

Remark 2.7. The identities (2.5), at least for the continuous-time problems, have been obtained in [34, Remark 4.1] in a more general setting, via a scaling argument. The interpretation of (2.5) is clear: Our optimal martingale is a constant-speed geodesic when we think of the square of our optimization problems.

3. Main results in dimensions one and two

We want to further understand the structure of the unique optimizer found in the previous part. At the moment we can do this in dimension one (without assumptions) and dimension two (with further assumptions).

3.1. **The one-dimensional case.**

Let $\mu \preceq_c \nu$ be probability measures on the line with finite second moment. For a measure α on $\mathbb{R}$ on $x \in \mathbb{R}$ we write $u_\alpha(x) := \int |x-y|\, d\alpha(y)$. Using the this notation, the convex order relation $\mu \preceq_c \nu$ is equivalent to $u_\mu \le u_\nu$.

We recall from [13, Appendix A.1] that by the "irreducible components of (μ,ν)" are determined by the (unique) family of open disjoint intervals $\{I_k\}_{k\in\mathbb{N}}$ whose union equals open set

$$\{u_\mu < u_\nu\} := \Big\{x \in \mathbb{R} : \int |x-y|\, d(\mu-\nu)(y) \neq 0\Big\}.$$

One can then decompose

$$\mu = \eta + \textstyle\sum_k \mu_k \quad \text{and} \quad \nu = \eta + \sum_k \nu_k,$$

where $\mu_k = \mu|_{I_k}$, with $I_k = \{u_{\mu_k} < u_{\nu_k}\}$ and $\nu_k(\overline{I_k}) = \mu_k(I_k)$, whereas η is concentrated on $\mathbb{R}\setminus\bigcup_k I_k$. A useful straightforward result is that every martingale coupling from μ to ν (i.e. $\pi \in \mathsf{M}(\mu,\nu)$) is fully characterized by how it looks on the sets $I_k \times \overline{I_k}$. The restrictions $\pi_k :=$

$\pi_{|I_k\times\overline{I_k}} = \pi_{|I_k\times\overline{\mathbb{R}}}$ are still martingale couplings (in the sense that their respective disintegration satisfies $\int y\,(\pi_k)_x(y) = x$ for μ_k-a.a. x) but with total mass $\mu_k(I_k)$ and marginals μ_k, ν_k.

We can now state our main result for $d = 1$, where we better characterize the structure of stretched Brownian motion.

Theorem 3.1. *Let $\mu \preceq_c \nu$ be probability measures on the line with finite second moment. A candidate martingale M is an optimizer of (MBB) if and only if it is a* standard stretched Brownian motion *on each irreducible component (μ_k, ν_k) of (μ, ν). In particular,* stretched Brownian motion *(sBm) is a* standard stretched Brownian motion *(s^2Bm) in each irreducible component.*

Let us explain the terminology used here. Saying that M is s^2Bm on the irreducible components of (μ, ν) concretely means that, conditionally on $M_0 \in I_k$, M is a s^2Bm from $\frac{1}{\mu_k(I_k)}\mu_k$ to $\frac{1}{\mu_k(I_k)}\nu_k$. We stress that in the present $1-d$ case Theorem 3.1 is significantly stronger than Theorem 1.11.

We now proceed towards the subtler extension of Theorem 3.1 for $d = 2$. We omit the proof of Theorem 3.1 since it is easily derived from the two-dimensional considerations (with less effort and without the additional assumptions).

3.2. **Preliminaries.** We briefly discuss some of the aspects related to the decomposition of martingale couplings in arbitrary dimensions. Later this will be mostly used in dimension two. After this, we also provide an analysis lemma of much importance for the next sections.

Definition 3.2. *A convex paving C is a collection of disjoint relatively open convex sets from $\mathbb{R}^d$. Denoting $\bigcup C := \bigcup_{C\in C} C$, we will always assume $\mu(\bigcup C) = 1$ for such objects. For $x \in \bigcup C \subseteq \mathbb{R}^d$ we denote by $C(x)$ the unique element of C which contains x. We say that C is measurable (resp. μ-measurable, universally measurable) if the function $x \mapsto \overline{C(x)}$ is measurable as a map from $\mathbb{R}^d$ to the Polish space of all closed (convex) subsets of $\mathbb{R}^d$ equipped with the Wijsman topology (cf.* [21]*).*

Definition 3.3. *Let $\Gamma \subseteq \mathbb{R}^d \times \mathbb{R}^d$ and $\pi \in \mathsf{M}(\mu, \nu)$. We say that a convex paving C is*

- *π-invariant if $\pi_x(\overline{C(x)}) = 1$ for μ-a.e. x,*
- *Γ-invariant if $ri(\Gamma_x) \subseteq C(x)$ for all $x \in proj_1(\Gamma)$.*

Note that a natural order between convex pavings C, C' is given by

$$C \leq_\mu C' \iff C(x) \subseteq C'(x) \text{ for } \mu - a.e.\ x,$$

in which case we say that C is finer than C' (and the latter is coarser than the former). The following two theorems are shown in [26, 21, 48].

Theorem 3.4 (Ghoussoub-Kim-Lim [26])**.** *Given $\pi \in \mathsf{M}(\mu, \nu)$ and $\Gamma \subseteq \mathbb{R}^d \times \mathbb{R}^d$ a martingale support for π, there is a finest Γ-invariant convex paving. We denote it by $C_{\pi,\Gamma}$.*

Theorem 3.5 (De March-Touzi [21], Obłój-Siorpaes [48])**.** *There is a finest convex paving, denoted $C_{\mu,\nu}$, which is π-invariant for all $\pi \in \mathsf{M}(\mu, \nu)$ simultaneously. Writing $C_{\mu,\nu} = \{C_{\mu,\nu}(x)\}_{x\in\mathbb{R}^d}$, the function $x \mapsto C_{\mu,\nu}(x)$ is universally measurable.*

If we knew that these convex pavings coincide, this would streamline some of our proofs. For the case $d = 1$ this is indeed the case, but already for $d = 2$ this can fail. We will actually use another convex paving which incorporates ideas/properties from the above two.

Lemma 3.6. *Given $\pi \in \mathsf{M}(\mu, \nu)$ there is a finest measurable π-invariant convex paving, which we denote C_π.*

This can be established by a close reading of [21], and adapting the arguments therein (of course [21] achieves much more!). We give a self-contained, shorter argument under the following additional hypothesis, which we will also see appear in Section 3.3.

Assumption 3.7. *For all $\pi \in \mathsf{M}(\mu,\nu)$ and C convex paving we have*

$$\pi_x(\overline{C(x)}) = 1\,\mu - a.s. \;\Rightarrow \pi_x(C(x)) = 1\,\mu - a.s.$$

In particular, for such C and π, C is π-invariant iff $\pi_x(C(x)) = 1\,\mu$–a.s.

Proof of Lemmma 3.6 under Assumption 3.7. Inspired by [21], we introduce the optimization problem

$$\inf\{\textstyle\int \mu(dx)\, G(C(x)) : C \text{ is a } \pi\text{-invariant measurable convex paving}\},$$

where $G(C) := dim(C) + g_C(C)$ and g_C is the standard Gaussian measure on aff(C), i.e. as obtained from the dim(C)-dimensional Lebesgue measure on aff(C). Let C^n be an optimizing sequence of π-invariant convex pavings and let Ω be a set of μ-full measure on which we have $\pi_x(\overline{C^n(x)}) = 1$ for all n (here $C^n(x)$ denotes an element of C^n). Introduce for $x \in \Omega$ the relatively open convex sets $C_\pi(x) := \mathrm{rf}_x(\bigcap C^n(x))$. We have[2] $x \in C_\pi(x)$ since $x \in \bigcap C^n(x)$. Moreover we have that $C_\pi := \{C_\pi(x) : x \in \Omega\}$ forms a partition since already $\{\bigcap C^n(x) : x \in \Omega\}$ is a partition. Let us establish that $\pi_x(\overline{C_\pi(x)}) = 1$.

We start assuming

$$\forall\, K \text{ convex} : \; \mathrm{ri}\,\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x \subseteq K \Rightarrow \pi_x\left(\overline{\mathrm{rf}_x K}\right) = 1. \tag{3.1}$$

Let us take $K := \bigcap C^n(x)$. Since $\overline{C^n(x)}$ is closed, convex and satisfies $\pi_x(\overline{C^n(x)}) = 1$ we have $\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x \subseteq \overline{C^n(x)}$. On the other hand, $\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x$ cannot be contained in $\partial C^n(x)$ since by Assumption 3.7 we have $\pi_x(\partial C^n(x)) = 0$. By [50, Corollary 6.5.2] we must then have $\mathrm{ri}\,\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x \subseteq \mathrm{ri}\, C^n(x) = C^n(x)$ for all n, so $\mathrm{ri}\,\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x \subseteq \bigcap C^n(x) = K$. By (3.1) we get $\pi_x\left(\overline{\mathrm{rf}_x K}\right) = \pi_x\left(\overline{C_\pi(x)}\right) = 1$ as desired. All in all C_π is a π-invariant convex paving, and since $C_\pi(x) \subseteq C^n(x)$ we find $\int \mu(dx) G(C_\pi(x)) \leq \int \mu(dx) G(C^n(x))$ from which we get the optimality of C_π.

To finish the proof, let us establish (3.1). By the martingale property we easily see[3] that $x \in \mathrm{ri}\,\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x$. From this, $\mathrm{ri}\,\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x = \mathrm{rf}_x(\mathrm{ri}\,\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x) \subseteq \mathrm{rf}_x K$. Hence $\overline{\mathrm{ri}\,\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x} \subseteq \overline{\mathrm{rf}_x K}$, whose l.h.s. equals $\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x$ by [50, Theorem 6.3], so (3.1) follows. □

Remark 3.8. The same proof, modulo obvious changes, proves the existence of a finest measurable convex paving invariant for all $\pi \in \mathsf{M}(\mu,\nu)$ simultaneously. This however does not establish the existence of a maximally spreading martingale coupling as in [21].

Here is a sufficient criterion for Assumption 3.7 to hold.

Lemma 3.9. *Assumption 3.7 is satisfied if $d \in \{1,2\}$ and $\nu \ll \lambda^d$.*

Proof. This follows by similar arguments as in [26, Lemma C.1]. We omit the details. □

A direct consequence of Theorem 3.5 and Assumption 3.7 is the decomposition of a martingale into irreducible components. Notice the resemblance to the one-dimensional case explained in Section 3.1.

Proposition 3.10. *Let $C_{\mu,\nu} = \{C_{\mu,\nu}(x)\}_{x\in\mathbb{R}^d}$ be the convex paving of Theorem 3.5 and assume Assumption 3.7. Then*

(i) *we may decompose*

$$\mu = \textstyle\int \mu(\,\cdot\,|K) dC_{\mu,\nu}(\mu)(K), \text{ and } \nu = \int \nu(\,\cdot\,|K) dC_{\mu,\nu}(\mu)(K),$$

with $\mu(\,\cdot\,|K) \preceq_c \nu(\,\cdot\,|K)$ for $C_{\mu,\nu}(\mu)$-a.e. K;

[2]Recall that $A \subseteq A' \Rightarrow \mathrm{rf}_a A \subseteq \mathrm{rf}_a A'$, that $a \in A \iff a \in \mathrm{rf}_a(A)$ and that $\mathrm{rf}_a(A) = \mathrm{ri}\, A \iff a \in \mathrm{ri}\, A$.

[3]Let $m = \dim(\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x)$ and suppose $x \in \partial(\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x)$. We can then find an $(m-1)$-dimensional hyperplane supporting x and having $\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x$ contained in one associated half-space. By the martingale property one obtains that necessarily $\mathrm{supp}\,\pi_x$, and then $\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x$ too, must be actually contained in the hyperplane itself. Thus $\dim(\overline{\mathrm{co}}\,\mathrm{supp}\,\pi_x) \leq m-1$ yielding a contradiction.

(ii) *for any martingale coupling* $\pi \in \mathsf{M}(\mu,\nu)$ *we have that*

$$\pi(\,\cdot\,|K\times K) = \pi(\,\cdot\,|K\times\mathbb{R}^d)\ \textit{for}\ C_{\mu,\nu}(\mu) - a.e.\ K,$$

and this common measure has first and second marginals equal to $\mu(\,\cdot\,|K)$ *and* $\nu(\,\cdot\,|K)$ *respectively;*

(iii) *any martingale coupling* $\pi \in \mathsf{M}(\mu,\nu)$ *can be uniquely decomposed as*

$$\pi = \int \pi(\,\cdot\,|K\times K)dC_{\mu,\nu}(\mu)(K).$$

The proof is just as in [13, Appendix A.1], but simpler, thanks to the fact that under Assumption 3.7 we have that *martingales started on two neighbouring cells will not go on to reach the intersection of the boundaries of the cells*. We thus omit the proof.

We finally present a technical lemma which will be extremely useful in the proofs of the main results in dimension two.

Lemma 3.11. *Let* η *be a probability measure in* $\mathbb{R}^d$ *with finite second moment, and* $F : \mathbb{R}^d \to \mathbb{R}$ *convex such that* $\nabla F(\gamma^d) = \eta$. *Then* $\gamma^d - a.s.$ $\nabla F = \nabla\tilde{F}\circ P$, *where* P *is the orthogonal projection onto* $V := aff(supp(\eta))$ *and* $\tilde{F} : V \to \mathbb{R}$ *is convex.*

For all $s > 0$, *the function*

$$\mathbb{R}^d \ni b \mapsto f_s(b) := \int \nabla F(b+y)\gamma^d_s(dy) = \int \nabla\tilde{F}(Pb+z)P(\gamma^d_s)(dz) \in \mathbb{R}^d,$$

has the following properties:

(1) It is infinitely continuously differentiable.
(2) It is one-to-one on V,
(3) $\overline{f_s(\mathbb{R}^d)} = \overline{co}\,\nabla F(\mathbb{R}^d)$.
(4) $f_s(\gamma^d)$ *is equivalent to the* m*-dimensional Lebesgue measure on* V *restricted to* $co\,\nabla F(\mathbb{R}^d)$, *where* $m = dim(V)$.
(5) $supp(f_s(\gamma^d_t)) = \overline{co}\,\nabla F(\mathbb{R}^d)$ *is convex and does not depend on* $s > 0$ *nor* $t > 0$.

Proof. The $\gamma^d - a.s.$ equality $\nabla F = \nabla\tilde{F}\circ P$, follows from Brenier's Theorem, by taking $\nabla\tilde{F}$ mapping $P(\gamma^d)$ into η and observing that $\nabla(\tilde{F}\circ P) = P((\nabla\tilde{F})\circ P) = (\nabla\tilde{F})\circ P$. Point (1) follows by change of variables and differentiation under the integral sign. Alternatively, one can argue with the classical backwards heat equation. Points 2, 3 and 5 follow by the full-support property of γ^d in $\mathbb{R}^d$ and $P(\gamma^d)$ in V.

Point 4 is trivially true if η is a Dirac delta (then m=0). Otherwise, it suffices to study the claim for the smooth function $\tilde{f}_s(v) := \int \nabla\tilde{F}(v+z)\tilde{\gamma}(dz)$, with $\tilde{\gamma} = P(\gamma^d_s)$. Observe that the Jacobian of this function is $J\tilde{f}_s(v) = \int \mathrm{Hess}(\tilde{F})(v+z)\tilde{\gamma}(dz)$, where the (absolutely continnuous part of the) Hessian of $\tilde{F}$ exists $\tilde{\gamma}$-a.e. by Alexandrov's Theorem [54, Chapter 2.1.3(8)]. By Point 2, $\tilde{f}_s$ is one-to-one on V. Since η is not trivial, $\mathrm{Hess}(\tilde{F})$ must be (strictly) positive definite on a non-negligible set, so $J\tilde{f}_s$ is everywhere strictly positive definite and in particular invertible. By change of variables formula, it is easy to obtain

$$\frac{d\tilde{f}_s(\tilde{\gamma})}{d\lambda_V}(r) = \frac{d\tilde{\gamma}}{d\lambda_V}\left((\tilde{f}_s)^{-1}(r)\right)\frac{\mathbf{1}_{\tilde{f}_s(V)}(r)}{\det\left(J\tilde{f}_s\right)\circ(\tilde{f}_s)^{-1}(r)}, \tag{3.2}$$

where λ_V is m-dimensional Lebesgue on V. By Point 3, $\tilde{f}_s(V) = \mathrm{co}\,\nabla F(\mathbb{R}^d)$, and so the above density never vanishes on this set, from which we conclude $\tilde{f}_s(\tilde{\gamma}) \sim \lambda_V|_{\mathrm{co}\,\nabla F(\mathbb{R}^d)}$. The claim follows easily from this. □

3.3. **The two-dimensional case.** Our first main result for $d = 2$ is a characterization of the structure of sBm, providing a significantly strengthened version of Theorem 1.11 in the introduction.

Theorem 3.12. *Let* $\mu \preceq_c \nu$ *probability measures in* $\mathbb{R}^2$ *with finite second moments. Suppose* $\nu \ll \lambda^2$, *and let* M^* *be the unique optimizer for* (MBB). *Denote* $\pi^t = law\,(M^*_0, M^*_t)$ *with* $0 < t < 1$. *Then stretched Brownian motion* M^* *is on each cell of* C_{π^t} *a standard stretched Brownian motion.*

The second main result of this part is the optimality of s^2Bm whenever we are able to build them with respect to the coarser $C_{\mu,\nu}$ convex paving. Our proof of such result relies on the simplifying assumption 3.7, which as seen in Lemma 3.9 is *verified in dimension two under the further requirement that ν be absolutely continuous*. We therefore place this result here, although in principle it is a result valid in arbitrary dimensions.

Theorem 3.13. *Under Assumption 3.7, if M is a standard stretched Brownian motion on each cell*[4] *of the convex paving $C_{\mu,\nu}$, then it is optimal for (MBB) (i.e. it is a sBm).*

Remark 3.14. The difference between Theorem 1.10 and Theorem 3.13 is as follows: the first one says that standard stretched Brownian motion is optimal in its own, whereas the second statement allows for more freedom in that we are allowed to choose the convex function in the definition of stretched Brownian motion dependent on the cells of $C_{\mu,\nu}$. Therefore this result is a strengthened version of Theorem 1.10.

Remark 3.15. For dimension one ($d = 1$), Theorem 3.1 establishes the existence of standard stretched Brownian motion, and characterize it as the sole optimizer. Both existence and optimality are understood with respect to the same (countable) convex paving. For two dimensions ($d = 2$), Theorems 3.12 and 3.13 and Lemma 3.9 establish, under the assumption that $\nu \ll \lambda^2$, the existence and optimality characterization of standrard stretched Brownian motion. In this case however, existence and optimality are understood with respect to potentially different convex pavings.

The proofs of these results are deferred to Section 5. Theorem 3.12 relies crucially in a monotonicity principle which we now establish and we think is of independent interest.

4. A monotonicity principle for generalized transport costs

For this part only, we adopt a more general setting. Let X, Y be Polish spaces and $C : \mathsf{X} \times \mathcal{P}(\mathsf{Y}) \to \mathbb{R} \cup \{+\infty\}$ Borel measurable. Consider for $\mu \in \mathcal{P}(\mathsf{X}), \nu \in \mathcal{P}(\mathsf{Y})$ the optimization problem

$$\inf_{\pi \in \Pi(\mu,\nu)} \int_X \mu(dx) C(x, \pi_x). \tag{4.1}$$

This is a generalized (non-linear) transport problem in the sense of [28, 27] and the references therein. We now obtain a "monotonicity principle" for this problem, i.e. a finitistic "zeroth-order" necessary optimality condition.

Proposition 4.1. *Suppose that*

- *Problem* (4.1) *is finite with optimizer π;*
- *$\mu(dx)$-a.e. the function $C(x, \cdot)$ is convex.*

Then there exists a Borel set $\Gamma \subseteq X$ with $\mu(\Gamma) = 1$ and the following property

$$\text{if } x, x' \in \Gamma \text{ and } m_x, m_{x'} \in \mathcal{P}(Y) \text{ satisfy } m_x + m_{x'} = \pi_x + \pi_{x'}, \text{ then}$$
$$C(x, \pi_x) + C(x', \pi_{x'}) \leq C(x, m_x) + C(x', m_{x'}).$$

Proof. Let

$$\mathcal{D} := \left\{ \big((x, x'), (m_1, m_2) \big) \in X^2 \times \mathcal{P}(Y)^2 : \begin{array}{l} m_1 + m_2 = \pi_x + \pi_{x'} \text{ , and} \\ C(x, \pi_x) + C(x', \pi_{x'}) > C(x, m_1) + C(x', m_2) \end{array} \right\},$$

which is an analytic set. By Jankov-von Neumann uniformization theorem there is [38, Theorem 18.1] an analytically measurable function

$$D := \mathrm{proj}_{X^2}(\mathcal{D}) \ni (x, x') \mapsto (m_1^{(x,x')}, m_2^{(x,x')}) \in \mathcal{P}(Y)^2,$$

so that $(x, x', m_1^{(x,x')}, m_2^{(x,x')}) \in \mathcal{D}$. Since $(x, x', m_1, m_2) \in \mathcal{D} \iff (x', x, m_2, m_1) \in \mathcal{D}$, it is possible to prove that we may actually assume that

$$(m_1^{(x',x)}, m_2^{(x',x)}) = (m_2^{(x,x')}, m_1^{(x,x')}). \tag{4.2}$$

[4] This means that for $C_{\mu,\nu}(\mu)$-a.e. K, the conditioning of M^* to $M_0^* \in K$ is a stretched Brownian motion between the marginals $\mu(\cdot\,|K)$ and $\nu(\cdot\,|K)$ introduced in Proposition 3.10

Of course the set D is likewise analytic. Thus extending $(m_1^{(\cdot,\cdot)}, m_2^{(\cdot,\cdot)})$ to $(x, x') \notin D$ by setting it to $(\pi_x, \pi_{x'})$, analytic-measurability and the symmetry property (4.2) are preserved.

Assume that there exists $Q \in \Pi(\mu, \mu)$ such that $Q(D) > 0$. We now show that this is in conflict with the optimality of π. By considering $\frac{Q+e(Q)}{2}$, where $e(x, x') := (x', x)$, we may assume that Q is symmetric. We first define

$$\tilde{\pi}(dx, dy) := \mu(dx) \int_{x'} Q_x(dx') m_1^{(x,x')}(dy), \tag{4.3}$$

which is legitimate owing to all measurability precautions we have taken. We will prove

(1) $\tilde{\pi} \in \Pi(\mu, \nu)$,
(2) $\int \mu(dx) C(x, \pi_x) > \int \mu(dx) C(x, \tilde{\pi}_x)$.

For (1): Evidently the first marginal of $\tilde{\pi}$ is μ. On the other hand

$$\int_x \mu(dx) \tilde{\pi}_x(dy) = \int_x \mu(dx) \int_{x'} Q_x(dx') m_1^{(x,x')}(dy) = \int_{x,x'} Q(dx, dx') m_1^{(x,x')}(dy).$$

The last quantity is equal to $\int_{x,x'} Q(dx, dx') m_2^{(x,x')}(dy)$ by symmetry of Q and (4.2). So

$$\int_x \mu(dx) \tilde{\pi}_x(dy) = \int_{x,x'} Q(dx, dx') \tfrac{m_1^{(x,x')} + m_2^{(x,x')}}{2}(dy) = \int_{x,x'} Q(dx, dx') \tfrac{\pi_{x'} + \pi_x}{2}(dy) = \nu(dy),$$

by definition of $m^{(x,x')}$ and Q. Thus $\tilde{\pi}$ has second marginal ν.

For (2): By convexity of $C(x, \cdot)$, the symmetry of Q and (4.2), and by the assumption that on the Q-non negligible set D we have $C(x, \pi_x) + C(x', \pi_{x'}) > C(x, m_1^{(x,x')}) + C(x', m_2^{(x,x')})$, we conclude

$$\begin{aligned}
\int_x \mu(dx) C(x, \tilde{\pi}_x) &= \int_x \mu(dx) C\left(x, \int_{x'} Q_x(dx') m_1^{(x,x')}\right) \\
&\leq \int_x \mu(dx) \int_{x'} Q_x(dx') C\left(x, m_1^{(x,x')}\right) \\
&= \int_{x,x'} Q(dx, dx') C\left(x, m_1^{(x,x')}\right) \\
&= \int_{x,x'} Q(dx, dx') \tfrac{C\left(x, m_1^{(x,x')}\right) + C\left(x, m_2^{(x,x')}\right)}{2} \\
&< \int_{x,x'} Q(dx, dx') \tfrac{C(x, \pi_x) + C(x', \pi_{x'})}{2} \\
&= \int_x \mu(dx) C(x, \pi_x).
\end{aligned}$$

As expected, we have contradicted the optimality of π.

We conclude that no measure Q with the stated properties exists. By "Kellerer's lemma" [9, Proposition 2.1], which is also true for analytic sets, we obtain that D is contained in a set of the form $N \times N$ where $\mu(N) = 0$. Letting $\Gamma := N^c$, so $\Gamma \times \Gamma \subseteq D^c$, we easily conclude. □

We now go back to the main framework in this article. The monotonicity principle will be crucially used, under the following guise, in order to prove the results in Section 3.3. For a kernel $\pi_x(dy)$ and $\tilde{\mu}(d\tilde{x}) = \frac{1}{2}(\delta_x(d\tilde{x}) + \delta_{x'}(d\tilde{x}))$ we write $\pi_{\tilde{x}}(dy)\tilde{\mu}(d\tilde{x}) = \frac{1}{2}(\delta_x \pi_x + \delta_{x'} \pi_{x'})$.

Corollary 4.2. *Let π be optimal for (GOT). Then there exists $\Gamma \subseteq \mathbb{R}^d$ with $\mu(\Gamma) = 1$ such that*

$$\textit{if } x, x' \in \Gamma, \textit{ then the measure } \tfrac{\delta_x \pi_x + \delta_{x'} \pi_{x'}}{2} \textit{ is optimal for}$$

$$\inf_{\substack{mean(m_{x'}) = x',\, mean(m_x) = x \\ (m_x + m_{x'})/2 = (\pi_x + \pi_{x'})/2}} \left\{ \mathcal{W}_2(m_x, \gamma^d)^2 + \mathcal{W}_2(m_{x'}, \gamma^d)^2 \right\}. \tag{4.4}$$

Proof. Consider Proposition 4.1, taking $X = Y = \mathbb{R}^d$ and setting

$$C(x, m) = \mathcal{W}_2(m, \gamma^d)^2,$$

if $mean(m) = x$ and $+\infty$ otherwise. Observe that $C(x, \cdot)$ is convex by [55, Theorem 4.8]. Taking Γ to be the μ-full set given by Proposition 4.1, the result is immediate. □

Observe that Problem (4.4) is of the same kind as (*GOT*), with first marginal $\frac{\delta_x+\delta_{x'}}{2}$ and final marginal $\frac{\pi_x+\pi_{x'}}{2}$. It follows that (4.4) has a continuous-time analogue, which enjoys the dynamic programming principle, and whose optimizer is a strong Markov martingale. This fact will be repeatedly used in the next part.

Remark 4.3. Of course there are versions of the results in this section for general n-tuples instead of pairs. Since we only use the version with pairs we did not state the result in its most general form.

5. Pending proofs

5.1. **Proof of Theorems 1.10 and 3.13.**

Proof of Theorem 1.10. Let $A : \mathbb{R}^d \to \mathbb{R}^d$ be in $L^2(\mu)$ and $\varphi, \psi : \mathbb{R}^d \to \mathbb{R}$ be conjugate convex functions. We start by proving that

$$GT \leq \int \varphi \, d\nu - \int x \cdot A(x) \, d\mu + \int \mu(dx) \int \gamma^{(A(x))}(db)\psi(b), \tag{5.1}$$

where $\gamma^{(a)} := \delta_a * \gamma^d$. First observe that

$$\sup_{q\in\Pi(\pi,\gamma)} \int q(dm, db)m \cdot b = \sup_{q\in\Pi(\pi,\gamma^{(a)})} \int q(dm, db)m \cdot [b - a].$$

Let us write $\Sigma := \{ \{\pi_x\}_x : \text{mean}(\pi_x) = x \text{ and } \int \mu(dx)\pi_x(dy) = \nu(dy) \}$. From here,

$$\begin{aligned}
GT &= \sup_{\{\pi_x\}_x\in\Sigma} \int \mu(dx) \sup_{q\in\Pi(\pi_x,\gamma^{(A(x))})} \int q(dm, db)m \cdot [b - A(x)] \\
&= \sup_{\{\pi_x\}_x\in\Sigma} \int \mu(dx) \left[- \int \pi_x(dm)m \cdot A(x) + \sup_{q\in\Pi(\pi_x,\gamma^{(A(x))})} \int q(dm, db)m \cdot b \right] \\
&= \sup_{\{\pi_x\}_x\in\Sigma} \int \mu(dx) \left[-x \cdot A(x) + \sup_{q\in\Pi(\pi_x,\gamma^{(A(x))})} \int q(dm, db)m \cdot b \right] \\
&\leq - \int x \cdot A(x) \, d\mu + \sup_{\{\pi_x\}_x\in\Sigma} \int \mu(dx) \sup_{q\in\Pi(\pi_x,\gamma^{(A(x))})} \int q(dm, db)[\varphi(m) + \psi(b)] \\
&= - \int x \cdot A(x) \, d\mu + \int \varphi \, d\nu + \int \mu(dx) \int \gamma^{(A(x))}(db) \, \psi(b)
\end{aligned}$$

by the conjugacy relationship $m \cdot b \leq \varphi(m) + \psi(b)$ and the defining property of Σ. Hence, (5.1) follows.

Let now M be standard stretched Brownian motion from μ to ν in the notation of Definition 1.8 and equation (1.9). By classical convex analysis arguments, or optimal transport theory, there exists φ, ψ convex conjugate functions such that λ^d-a.e. (γ^d-a.e.)

$$\nabla F(b) \cdot b = \varphi(\nabla F(b)) + \psi(b).$$

We also choose $A(x) = f_0^{-1}(x)$, which is well defined on $\text{supp}(\mu)$ by Lemma 3.11.

By definition $\mu \sim M_0 = f_0(B_0) \sim f_0(\alpha)$, so

$$\int x \cdot A(x) \, d\mu(x) = \int x \cdot A(x) df_0(\alpha)(x) = \int f_0(x) \cdot x \, d\alpha(x) = \mathbb{E}[f_0(B_0)B_0] = \mathbb{E}[M_0 \cdot B_0].$$

On the other hand

$$\begin{aligned}
\int \varphi \, d\nu + \int \mu(dx) \int \gamma^{(A(x))}(db)\psi(b) &= \mathbb{E}[\varphi(M_1)] + \int A(\mu)(dx) \int \gamma^{(x)}(db)\psi(b) \\
&= \mathbb{E}[\varphi(M_1)] + \int A(\mu)(dx)\mathbb{E}[\psi(B_1)|B_0 = x] \\
&= \mathbb{E}[\varphi(M_1)] + \mathbb{E}[\ \mathbb{E}[\psi(B_1)|B_0 = x]\ ], \\
&= \mathbb{E}[\varphi(M_1) + \psi(B_1)],
\end{aligned}$$

since $A(\mu) = \alpha = law\,(B_0)$. So the r.h.s. of (5.1) becomes in this case

$$\begin{aligned}
\mathbb{E}[\varphi(M_1) + \psi(B_1) - M_0 \cdot B_0] &= \mathbb{E}[\varphi(\nabla F(B_1)) + \psi(B_1) - M_0 \cdot B_0] \\
&= \mathbb{E}[\nabla F(B_1) \cdot B_1 - M_0 \cdot B_0] = \mathbb{E}[M_1 \cdot B_1 - M_0 \cdot B_0] = \mathbb{E}\left[\int_0^1 \text{tr}(\sigma_t) dt \right].
\end{aligned}$$

Hence, Theorem 2.2 implies the optimality of M.

□

We now work under Assumption 3.7, still in arbitrary dimension d.

Proof of Theorem 3.13. We observe from Proposition 3.10 that the optimization problem (GOT) can be decomposed / disintegrated along the cells of $C_{\mu,\nu}$. Therefore, optimality must only hold for $C_{\mu,\nu}(\mu)$-a.e. K for the corresponding transport problems with first and second marginals $\mu(\,\cdot\,|K)$ and $\nu(\,\cdot\,|K)$ respectively. This reduces the argument to the previous case of Theorem 1.10, and we conclude. □

5.2. **Proof of Theorem 3.12.** Although this is eventually a two-dimensional result, for the arguments we do not fix the dimension d to two unless we explicitly say so.

Let M be the unique optimizer of (MBB), where we drop the superscript $*$ for simplicity. By Theorem 2.2 this continuous-time martingale is associated to the unique two-step martingale π optimizing (GOT). Let ∇F^x be the optimal transport map pushing γ^d to π_x.

By Remark 2.3, we know that conditioning on $M_0 = x$ the martingale M is given by

$$M_t^x := f_t^x(B_t), \quad \text{where} \quad f_t^x(\cdot) := \int \nabla F^x(b+\cdot)\gamma^d_{1-t}(db). \tag{5.2}$$

Remark 5.1. By Lemma 3.11, it would be more appropriate to write $M_t^x = \tilde{f}_t^x(PB_t)$, where P is the orthogonal projection onto $\mathrm{aff}(\mathrm{supp}\,\pi_x)$, and $\tilde{f}_t^x(\cdot) := \int \nabla \tilde{F}^x(b+\cdot)P\left(\gamma^d_{1-t}\right)(db)$ with $\nabla\tilde{F}^x(P(\gamma^d_{1-t})) := \pi_x$. This would make the notation unnecessarily heavy. In what follows, we chose to make use of the simpler notation in (5.2), and ask the reader to keep in mind that the Brownian motion B may actually mean a Brownian motion in a subspace (namely PB) and similarly f^x, F^x and γ^d should be understood as living in $\mathrm{aff}(\mathrm{supp}\,\pi_x)$.

We fix $0 < t < 1$ throughout. By Lemma 3.11 we find $B_t = (f_t^x)^{-1}(M_t^x)$. We denote

$$\pi_{x,y} := law\,(M_1|M_0 = x, M_t = y) = \nabla F^x(\delta_{(f_t^x)^{-1}(y)} * \gamma^d_{1-t}). \tag{5.3}$$

Important convention: For the rest of this section we make the convention that x, y, z denote possible values of the random variables M_0, M_t, M_1 respectively.

Lemma 5.2. *Let g be the unique gradient of a convex function such that $g(\gamma^d_{1-t}) = \pi_{x,y}$. Then $\nabla F^x(\cdot) = g(-(f_t^x)^{-1}(y) + \cdot)$. In particular ∇F^x is uniquely determined by the family of translates of g, which we denote by*

$$type(\pi_{x,y}) := \{a \mapsto g(r+a) :\ r \in \mathbb{R}^d\}.$$

Proof. For $r \in \mathbb{R}^d$ write $g_r(\cdot) = g(\cdot - r)$. Then, we have $\pi_{x,y} = g(\gamma^d_{1-t}) = g_{(f_t^x)^{-1}(y)}(\delta_{(f_t^x)^{-1}(y)} * \gamma^d_{1-t})$. Hence, both ∇F^x and $g_{(f_t^x)^{-1}(y)}$ push forward $\delta_{(f_t^x)^{-1}(y)} * \gamma^d_{1-t}$ into $\pi_{x,y}$, and both are gradients of convex functions. By the uniqueness result in Brenier's theorem, it follows that they are equal. Thus knowing ∇F^x determines g modulo translation. Conversely, knowing $\mathrm{type}(\pi_{x,y})$ (i.e. the translations of g) determines ∇F^x upon finding the vector r such that $\int g(r+a)\gamma^d(da) = x$. □

Let

$$\pi^t := law\,(M_0, M_t) \tag{5.4}$$

and consider

$$C := \{C(x)\}_x := C_{\pi^t},$$

the minimal π^t-invariant measurable convex paving of Lemma 3.6. We need to show that on each cell of C, M is a standard stretched Brownian motion, i.e. on each cell $C(x)$, we need to find a convex function $F = F_{C(x)}$ such that

$$M_1^x = \nabla F((f_0^x)^{-1}(x) + B_1), \tag{5.5}$$

where f_0^x and F are related via (5.2). To this end, we introduce

$$A(x) := \mathrm{type}(\pi_x) = \{a \mapsto \nabla F^x(a+r) \ :\ r \in \mathbb{R}^d\}$$

and we need to show that on each cell $A(x)$ is constant. We start by establishing a few preliminary results.

Lemma 5.3. *If $A(x)$ is constant in each cell of C, then M is a standard stretched Brownian motion each of these cells.*

Proof. As in Lemma 5.2. Fix arbitrary $x' \in C(x)$. Then, we have $\nabla F^x(\cdot) = \nabla F^{x'}(r(x) + \cdot)$ as $A(x) = A(x')$. Setting $\nabla F := \nabla F^{x'}$, proves the claim. □

To make use of the previous lemma, we shall study the behaviour of the martingale M for times in $[0, t]$ and $[t, 1]$. Let

$$\tilde{\pi}(dy, dz) := \operatorname{argsup} V(t, 1, law(M_t), \nu),$$

where $\tilde{\pi}$ is understood as the coupling of the initial and terminal marginals of the unique optimizer for $V(t, 1, law(M_t), \nu)$. For π^t from (5.4) we denote its disintegration w.r.t. the second marginal by $(\pi^t_y)_y$. Recall $\pi_{x,y}$ from (5.3).

Lemma 5.4. *For $law(M_t)$-a.e. y and π^t_y-a.e. x, we have*

$$\pi_{x,y}(dz) = \tilde{\pi}_y(dz).$$

Proof. We must have $\tilde{\pi} = law(M_t, M_1)$, by Lemma 2.4 (1). Thus $\tilde{\pi}_y = law(M_1 | M_t = y)$. On the other hand, $\pi_{x,y} = law(M_1 | M_t = y, M_0 = x)$ so by Lemma 2.4 we get $\pi_{x,y}(dz) = \tilde{\pi}_y(dz)$ for $law(M_t)$-a.e. y and π^t_y-a.e. x. □

The previous Lemma shows that the type of $\pi_{x,y}$ only really depends on y. Here is the opposite dependence

Lemma 5.5. *For μ-a.e. x and π^t_x-a.e. y we have*

$$type(\pi_{x,y}) = A(x).$$

Proof. By Lemma 5.2, if $g \in \text{type}(\pi_{x,y})$ then ∇F^x is a translate of g (the translation may depend on x, y). But this means conversely that g is a translate of ∇F^x, i.e. $g \in A(x)$. Reversing the steps gives the equality. □

We finalize the proof of Theorem 3.12. In a nutshell, the key is to deal with the null sets in Lemmas 5.4-5.5. **Only from now on we must assume that** $d = 2$.

Proof of Theorem 3.12. Lemma 5.5 proves that for π^t-a.e. (x, y), $\text{type}(\pi_{x,y}) = A(x)$. On the other hand Lemma 5.4 implies that for π^t-a.e. (x, y), $\text{type}(\pi_{x,y}) = D(y)$, where $D(y)$ is the common almost sure type of all $\pi_{x,y}$ which can be reached from y. By Fubini we have

$$\pi^t(\{(x, y) : A(x) = D(y)\}) = 1.$$

We want to use this to show that $A(\cdot)$ is constant on the cells of C. We first prove this for

$$C^t := \{\operatorname{ri}\operatorname{supp}(\pi^t_x)\}_{x \in \mathbb{R}^d}.$$

By (5.2) and Lemma 3.11 (v) we have $\operatorname{supp}(\pi^t_x) = \operatorname{supp}(law(M_t | M_0 = x)) = \overline{\operatorname{co}}\, \nabla F^x(\mathbb{R}^d)$ is convex. As in the final part of the proof of Lemma 3.6, the martingale property implies $x \in \operatorname{ri}\overline{\operatorname{co}}\operatorname{supp}\pi^t_x = \operatorname{ri}\operatorname{supp}\pi^t_x$, and by [50, Theorem 6.3] we know $\operatorname{ri}\overline{\operatorname{supp}\pi^t_x} = \operatorname{supp}\pi^t_x$. Hence, to show that C^t is a candidate π^t-invariant convex paving, it remains to show that the cells of C^t are pairwise disjoint or equal.

By Proposition 5.6 below, there is a μ-full set of initial positions with the property that, if x, x' satisfy $\operatorname{ri}\operatorname{supp}\pi^t_x \bigcap \operatorname{ri}\operatorname{supp}\pi^t_{x'} \neq \emptyset$, then $A(x) = A(x')$, i.e. the types of π_x and $\pi_{x'}$ coincide. This means that ∇F^x and $\nabla F^{x'}$ are translates of each other, implying that $\overline{\operatorname{co}}\, \nabla F^x(\mathbb{R}^d) = \overline{\operatorname{co}}\, \nabla F^{x'}(\mathbb{R}^d)$. From the previous paragraph, this shows $\operatorname{supp}(\pi^t_x) = \operatorname{supp}(\pi^t_{x'})$ and in particular $\operatorname{ri}\operatorname{supp}\pi^t_x = \operatorname{ri}\operatorname{supp}\pi^t_{x'}$. In one stroke this proves that C^t is a (π^t-invariant) convex paving and that $A(\cdot)$ is constant on its cells.

Since C is finer than C^t, this proves that $A(\cdot)$ is constant in the cells of C as well, and we conclude the proof by Lemma 5.3 □

The crucial Proposition 5.6 below relies on the "monotonicity principle" of Proposition 4.1, and more specifically Corollary 4.2.

For the rest of this section let Γ be the μ-full set of Corollary 4.2.

Proposition 5.6. *There is a μ-full set $S \subseteq \Gamma$ with the following property: If $x, x' \in S$ and satisfy ri supp $\pi^t_x \bigcap$ ri supp $\pi^t_{x'} \neq \emptyset$, then $A(x) = A(x')$.*

Proof. **Step 1:** By Lemma 5.7 below, for $x, x' \in \Gamma$ we have

$$\dim \operatorname{ri\,supp} \pi^t_x = \dim \operatorname{ri\,supp} \pi^t_{x'} = \dim \left(\operatorname{ri\,supp} \pi^t_x \cap \operatorname{ri\,supp} \pi^t_{x'}\right) \Rightarrow$$
$$\pi^t_x(\operatorname{ri\,supp} \pi^t_x \cap \operatorname{ri\,supp} \pi^t_{x'} \cap \{\pi_{x,y} \neq \pi_{x',y}\}) = \pi^t_{x'}(\operatorname{ri\,supp} \pi^t_x \cap \operatorname{ri\,supp} \pi^t_{x'} \cap \{\pi_{x,y} \neq \pi_{x',y}\}) = 0. \tag{5.6}$$

The goal is now to prove that for pairs $x, x' \in \Gamma$ the l.h.s. of (5.6) holds. As we will see in the final step of the proof, the r.h.s. of (5.6) is a strengthening of the dynamic programming principle that allows to deal with the null sets in Lemmas 5.4 and 5.5 more effectively.

Step 2: By Lemma 5.8 we know that if $x, x' \in \Gamma$, then

$$\dim \operatorname{ri\,supp} \pi^t_x = \dim \operatorname{ri\,supp} \pi^t_{x'} = 1 \text{ and } \operatorname{ri\,supp} \pi^t_x \bigcap \operatorname{ri\,supp} \pi^t_{x'} \neq \emptyset$$
$$\Rightarrow \dim \left(\operatorname{ri\,supp} \pi^t_x \cap \operatorname{ri\,supp} \pi^t_{x'}\right) = 1.$$

Step 3: By Lemma 5.9, we have for $x, x' \in \Gamma$

$$\dim \operatorname{ri\,supp} \pi^t_x = 1 \text{ and } \operatorname{ri\,supp} \pi^t_{x'} = \{x'\} \Rightarrow x' \notin \operatorname{ri\,supp} \pi^t_x \ .$$

Step 4: By Lemma 5.10, we have $x, x' \in \Gamma$

$$\dim \operatorname{ri\,supp} \pi^t_x = 2 \text{ and } \dim \operatorname{ri\,supp} \pi^t_{x'} = 1 \Rightarrow \operatorname{ri\,supp} \pi^t_x \bigcap \operatorname{ri\,supp} \pi^t_{x'} = \emptyset \ .$$

Step 5: By Lemma 5.11, the family $C^t_2 := \{\operatorname{ri\,supp} \pi^t_x : x \in \Gamma, \dim \operatorname{ri\,supp} \pi^t_x = 2\}$ consists of pairwise either disjoint or equal sets. As these are open sets, there can only be countable many different such sets (i.e. $|C^t_2| = |\mathbb{N}|$). If $C \in C^t_2$ is such that $\mu(\{x : \operatorname{ri\,supp} \pi^t_x = C\}) = 0$, then we discard this set C from our convex paving. So we may assume, for $C \in C^t_2$ that $\mu(\{x : \operatorname{ri\,supp} \pi^t_x = C\}) > 0$. By Lemma 5.11 the set $\{x' \in C : \operatorname{ri\,supp} \pi^t_{x'} = \{x'\}\}$ is μ-null under the assumption $\nu \ll \lambda^2$. Hence, for each of the conutably many $C \in C^t_2$ we can discard a μ-null set such that on a possibly smaller but still μ-full subset of Γ, which we keep calling Γ for simplicity, we have

$$x, x' \in \Gamma, \dim \operatorname{ri\,supp} \pi^t_x = 2 \text{ and } \{x'\} = \operatorname{ri\,supp} \pi^t_{x'} \Rightarrow x' \notin \operatorname{ri\,supp} \pi^t_x.$$

Final Step: By Steps 3, 4 and 5, we may assume that, for x, x' in a μ-full set, we have

$$\operatorname{ri\,supp} \pi^t_x \bigcap \operatorname{ri\,supp} \pi^t_{x'} \neq \emptyset \Rightarrow \dim \operatorname{ri\,supp} \pi^t_x = \dim \operatorname{ri\,supp} \pi^t_{x'}.$$

In this situation, if the common dimension in the r.h.s. is equal to one, by Step 2 also the dimension of the intersection in the l.h.s. is equal to one. On the other hand, if the common dimension in the r.h.s. is two, then automatically the dimension of the intersection is two (as an open convex set in $\mathbb{R}^2$). In any case, call $d^{(x,x')}$ this common dimension[5]. We find ourselves in the setting of (5.6), so by Step 1 we must have with $I := \operatorname{ri\,supp} \pi^t_x \cap \operatorname{ri\,supp} \pi^t_{x'}$

$$\pi^t_x(I \cap \{\pi_{x,y} \neq \pi_{x',y}\}) = \pi^t_{x'}(I \cap \{\pi_{x,y} \neq \pi_{x',y}\}) = 0. \tag{5.7}$$

Possibly throwing away another μ-null set, we know by Lemma 5.5 that on a μ-full set $S \subseteq \Gamma$ and for sets Y, Y' with $\pi^t_x(\mathsf{Y}) = \pi^t_{x'}(\mathsf{Y}') = 1$ it holds that

$$\operatorname{type}(\pi_{x,y}) = A(x), \ \forall y \in \mathsf{Y},$$
$$\operatorname{type}(\pi_{x',y}) = A(x'), \ \forall y \in \mathsf{Y}'.$$

[5]Actually the case $d^{(x,x')} = 2$ is settled by Lemma 5.11, but we prefer to give a general argument.

By Lemma 3.11, π^t_x is equivalent to $d^{(x,x')}$-dimensional Lebesgue measure on the $d^{(x,x')}$-dimensional open set $\operatorname{ri}\operatorname{supp}\pi^t_x$. Since $\operatorname{ri}\operatorname{supp}\pi^t_x \bigcap \operatorname{ri}\operatorname{supp}\pi^t_{x'}$ is a $d^{(x,x')}$-dimensional open subset it is also of positive $d^{(x,x')}$-Lebesgue measure. Then it is also of positive π^t_x-measure. Thus $\operatorname{ri}\operatorname{supp}\pi^t_x \bigcap \operatorname{ri}\operatorname{supp}\pi^t_{x'} \bigcap \mathsf{Y}$ has positive π^t_x-measure, and positive $d^{(x,x')}$-Lebesgue measure. But then again by Lemma 3.11 this same set must have positive $\pi^t_{x'}$-measure. We conclude that $I \bigcap \mathsf{Y} \bigcap \mathsf{Y}'$ has likewise positive $\pi^t_{x'}$-measure. The symmetric argument shows that the same set has positive π^t_x-measure. But by (5.7) the set $\{y : \pi_{x,y} = \pi_{x',y}\}$ is π^t_x-full in I. It follows that

$$I \bigcap \mathsf{Y} \bigcap \mathsf{Y}' \bigcap \{y : \pi_{x,y} = \pi_{x',y}\},$$

has positive π^t_x-measure, and by the same token it has positive $\pi^t_{x'}$-measure. In particular,

$$\mathsf{Y} \bigcap \{y : \pi_{x,y} = \pi_{x',y}\} \bigcap \mathsf{Y}' \neq \emptyset,$$

and taking y in this intersection we find

$$A(x) = \operatorname{type}(\pi_{x,y}) = \operatorname{type}(\pi_{x',y}) = A(x').$$

□

Lemma 5.7. *We have*

$$\left\{(x,x') : \begin{array}{c} \textit{dim ri supp}\,\pi^t_x = \textit{dim ri supp}\,\pi^t_{x'} = \textit{dim}\left(\textit{ri supp}\,\pi^t_x \bigcap \textit{ri supp}\,\pi^t_{x'}\right), \\ \textit{and either}\ \ \pi^t_x\left(\textit{ri supp}\,\pi^t_x \bigcap \textit{ri supp}\,\pi^t_{x'} \bigcap \{\pi_{x,y} \neq \pi_{x',y}\}\right) > 0, \\ \textit{or}\ \ \pi^t_{x'}\left(\textit{ri supp}\,\pi^t_x \bigcap \textit{ri supp}\,\pi^t_{x'} \bigcap \{\pi_{x,y} \neq \pi_{x',y}\}\right) > 0 \end{array}\right\} \bigcap (\Gamma \times \Gamma) = \emptyset.$$

Proof. Take $x, x' \in \Gamma$. By Corollary 4.2, the two-step martingale $\frac{\delta_x \pi_x + \delta_{x'} \pi_{x'}}{2}$ is optimal for (4.4). Consider its continuous-time analogue, i.e. the martingale which started at x equals M^x and started at x' equals $M^{x'}$ (cf. (5.2)) and both starting points have equal probability. We denote this continuous time martingale by $M^{(x,x')}$. By construction, $\mathit{law}\,(M^{(x,x')}_t | M^{(x,x')}_0 = x) = \pi^t_x$ and likewise for x'. Similarly $\mathit{law}\,(M^{(x,x')}_1 | M^{(x,x')}_0 = x,\ M^{(x,x')}_t = y) = \pi_{x,y}$ and the same holds for x' instead of x. By optimality of $\frac{\delta_x \pi_x + \delta_{x'} \pi_{x'}}{2}$ also $M^{(x,x')}$ is optimal for the continuous-time analogue of (4.4), then by dynamic programming (Lemma 2.4), we obtain sets Y, Y' such that

$$\begin{aligned} &\pi_{x,y} = \mathit{law}\,(M^{(x,x')}_1 | M^{(x,x')}_t = y) \text{ for } y \in \mathsf{Y} \text{ with } \pi^t_x(\mathsf{Y}) = 1, \\ &\pi_{x',y} = \mathit{law}\,(M^{(x,x')}_1 | M^{(x,x')}_t = y) \text{ for } y \in \mathsf{Y}' \text{ with } \pi^t_{x'}(\mathsf{Y}') = 1. \end{aligned}$$

The important point is that this is "pointwise" in $M^{(x,x')}_0 \in \{x, x'\}$.

Now assume further that $\dim \operatorname{ri}\operatorname{supp}\pi^t_x = \dim \operatorname{ri}\operatorname{supp}\pi^t_{x'} = \dim\left(\operatorname{ri}\operatorname{supp}\pi^t_x \bigcap \operatorname{ri}\operatorname{supp}\pi^t_{x'}\right)$, and call $d^{(x,x')}$ this common dimension. By Lemma 3.11 we have that π^t_x and $\pi^t_{x'}$ restricted to $\operatorname{ri}\operatorname{supp}\pi^t_x \bigcap \operatorname{ri}\operatorname{supp}\pi^t_{x'}$ are equivalent to $d^{(x,x')}$-dimensional Lebesgue measure restricted to this same set. We write

$$I := \operatorname{ri}\operatorname{supp}\pi^t_x \bigcap \operatorname{ri}\operatorname{supp}\pi^t_{x'}.$$

Necessarily $\mathsf{Y} \bigcap I$ is π^t_x-full in I, and therefore also $\pi^t_{x'}$-full in I. But then $\mathsf{Y} \bigcap I \bigcap \mathsf{Y}'$ is $\pi^t_{x'}$-full in I too. Inverting the roles of x, x' this set must also be π^t_x-full in I. We conclude

$$\pi^t_x(I \backslash (\mathsf{Y} \bigcap \mathsf{Y}')) = \pi^t_{x'}(I \backslash (\mathsf{Y} \bigcap \mathsf{Y}')) = 0.$$

But on $\mathsf{Y} \bigcap \mathsf{Y}'$ we have $\pi_{x,y} = \mathit{law}\,(M^{(x,x')}_1 | M^{(x,x')}_t = y) = \pi_{x',y}$, so

$$\pi^t_x(I \bigcap \{\pi_{x,y} \neq \pi_{x',y}\}) = \pi^t_{x'}(I \bigcap \{\pi_{x,y} \neq \pi_{x',y}\}) = 0.$$

This finishes the proof. □

Lemma 5.8. *Put*

$$\mathcal{V} := \{(x,x') : \textit{ri supp}\,\pi^t_x \textit{ and ri supp}\,\pi^t_{x'} \textit{ have dimension 1 and intersect in a singleton}\}.$$

Then,

$$\mathcal{V} \cap (\Gamma \times \Gamma) = \emptyset.$$

Proof. We use the same notation as in the proof of Lemma 5.7, and assume $(x, x') \in \mathcal{V}$.

By construction, $law\,(M_t^{(x,x')} | M_0^{(x,x')} = x) = \pi_x^t$ and likewise for x'. So $M_t^{(x,x')}$ conditioned to start at x, or at x', live respectively in line segments exactly intersecting in a single point $p \in \mathbb{R}^2$. By Lemma 3.11, the paths of these martingales (restricted to times in $[0, t]$) evolve in different "space-time" strips that only intersect along the line $L := \{(p, s) : s \geq 0\}$. Let $\tau := \inf\{s : (M_s^{(x,x')}, s) \in L\}$. It follows that $0 < \tau < t$ on a non-negligible set. The law of τ conditioned on the starting point of $M^{(x,x')}$ is equivalent to Lebesgue measure on $(0, 1)$. The reason is that this is true for 1-dimensional Brownian motion, and thanks to Lemma 3.11 the martingale $M^{(x,x')}$ conditioned to start say in x, is a one-dimensional Brownian motion after a continuous strictly increasing time change. Hence for any set $E \subseteq (0, 1)$ of positive Lebesgue measure we have $\mathbb{P}(\tau \in E \bigcap (0, t) \,|\, M_0^{(x,x')} = x) > 0$ and $\mathbb{P}(\tau \in E \bigcap (0, t) \,|\, M_0^{(x,x')} = x') > 0$. Thus we observe that the law of $M_t^{(x,x')}$ given $\{M_s^{(x,x')} : s \leq \tau \wedge t\}$ is different from the law of $M_t^{(x,x')}$ given $M_{\tau\wedge t}^{(x,x')}$. Indeed, when $\tau < t$ (equivalently when $M_{\tau\wedge t}^{(x,x')} = p$) and $\tau \in E$, one cannot for sure say in which of the aforementioned strips the martingale will continue to evolve. On the contrary, by observing $\{M_s^{(x,x')} : s \leq \tau \wedge t\}$ and on $\{\tau < t\} \bigcap \{\tau \in E\}$, such a strip is completely determined. Therefore $M^{(x,x')}$ fails to have the strong Markov property. But then it cannot be optimal between its marginals, by Lemma 2.5, and so neither can be $\frac{\delta_x \pi_x + \delta_{x'} \pi_{x'}}{2}$ optimal for (4.4). We conclude by Corollary 4.2 that $(x, x') \notin \Gamma \times \Gamma$. □

Lemma 5.9. *We have*

$$\{(x, x') : dim\ ri\ supp\, \pi_x^t = 1,\ ri\ supp\, \pi_{x'}^t = x',\ x' \in ri\ supp\, \pi_x^t\} \bigcap (\Gamma \times \Gamma) = \emptyset.$$

Proof. The proof is very similar to that of Lemma 5.8. Let (x, x') belong to the leftmost set. Using the same notation, $M^{(x,x')}$ is a martingale which evolves in a space-time strip if started at x, and otherwise is a constant equal to x'. We denote τ the first hitting time of $\{(x', s) : s \geq 0\}$. Since the martingale lives in a strip, we have that $\tau < t$ has positive probability. The strong Markov property of $M^{(x,x')}$ is destroyed at $\tau \wedge t$, since the knowledge of the past up to $\tau \wedge t$ reveals whether the martingale is constant or not thereafter. As before, by Lemma 2.5 and Corollary 4.2, $M^{(x,x')}$ cannot be optimal and $(x, x') \notin \Gamma \times \Gamma$. □

Lemma 5.10. *We have*

$$\{(x, x') : dim\ ri\ supp\, \pi_x^t = 2,\ dim\ ri\ supp\, \pi_{x'}^t = 1,\ ri\ supp\, \pi_x^t \bigcap ri\ supp\, \pi_{x'}^t \neq \emptyset\} \bigcap (\Gamma \times \Gamma) = \emptyset.$$

Proof. As in the previous proofs, with $M^{(x,x')}$ we associate $\tau = \inf\{s : M_s^{(x,x')} \in \mathrm{ri\,supp}\, \pi_{x'}^t\}$. Taking (x, x') in the leftmost set, it is tedious but not difficult to see that

$$law\left((M_\tau^{(x,x')}, \tau) \,|\, \tau \leq t, M_0^{(x,x')} = x\right), \quad \text{and} \quad law\left((M_U^{(x,x')}, U) \,|\, M_0^{(x,x')} = x'\right),$$

are equivalent to Lebesgue measure on $\mathrm{ri\,supp}\, \pi_{x'}^t \times [0, t]$, where U is uniformly distributed on $[0, t]$ and independent of everything. The point is that there is a common "space-time" set E charged by the two aforementioned laws. But the behaviour of $M_t^{(x,x')}$ conditioned on its past up to $\tau \wedge t$ is drastically depending on its starting position (e.g. whether it will evolve in a one- or two- dimensional set), whereas if for example we knew $(M_\tau^{(x,x')}, \tau) \in E$ then this does not reveal the dimension of the set where the martingale will continue to evolve. This contradicts the strong Markov property and we conclude as before. □

Lemma 5.11. *The family*

$$C_2^t := \{ri\ supp\, \pi_x^t :\ x \in \Gamma,\ dim\ ri\ supp\, \pi_x^t = 2\},$$

consists of open sets which are pairwise disjoint or equal. Assuming $\nu \ll \lambda^2$, we have

$$C \in C_2^t\ and\ \mu(\{x : ri\ supp\, \pi_x^t = C\}) > 0 \Rightarrow \mu(\{x' \in C : ri\ supp\, \pi_{x'}^t = \{x'\}\}) = 0.$$

Proof. Let Λ consist of all (x, x') such that

$$\{\dim \mathrm{ri\,supp}\, \pi_x^t = 2 = \dim \mathrm{ri\,supp}\, \pi_{x'}^t,\ \mathrm{ri\,supp}\, \pi_x^t \neq \mathrm{ri\,supp}\, \pi_{x'}^t, \mathrm{ri\,supp}\, \pi_x^t \bigcap \mathrm{ri\,supp}\, \pi_{x'}^t \neq \emptyset\}.$$

As before we can show that Λ cannot intersect $\Gamma\times\Gamma$. We do not give the argument, to avoid repetition, but mention that instead of contradicting the strong Markov property it suffices to contradict the regular Markov property. We conclude the first assertion.

Now let $C \in C_2^t$ such that $\mu(\{x : \operatorname{ri}\operatorname{supp}\pi_x^t = C\}) > 0$, $K := \{x' \in C : \operatorname{ri}\operatorname{supp}\pi_{x'}^t = \{x'\}\}$, and suppose $\mu(K) > 0$. We think of K as a non-negligible cloud of dots where the martingale M stays frozen. Since $M_0 \in K \Rightarrow M_1 \in K$, we have $\nu(K) > 0$ and by assumption $\lambda^2(K) > 0$. It follows that $\{M_t \in K\}$ is non-negligible, no matter if M has started on K or on $\{x : \operatorname{ri}\operatorname{supp}\pi_x^t = C\}$ at time zero (in the latter case, by Lemma 3.11). Since both sets of initial conditions are non-negligible, we contradict the regular Markov property of M. Indeed, on $\{M_t \in K\}$ the behaviour of M after t is drastically different depending on the starting condition at time zero being in K or $\{x : \operatorname{ri}\operatorname{supp}\pi_x^t = C\}$. This contradicts the optimality of M, and we conclude $\mu(K) = 0$. □

6. Further optimality properties

Let $\mathbb{T} := \{0 = t_0 \le t_1 \le \cdots \le t_{n-1} \le t_n = 1\} \subseteq [0,1]$ be a finite subgrid. Suppose M is a standard stretched Brownian motion from μ to ν, so $M_t = f_t(B_t)$ for a Brownian motion starting with some distribution α; see (1.9). Then

$$M_{t_0} = f_0(B_{t_0}),\ M_{t_1} = f_1(B_{t_1}),\ \dots\ ,\ M_{t_n} = f_n(B_{t_n}).$$

Denote $\nu^{\mathbb{T}} := law\,(M_{t_0}, M_{t_1}, \dots, M_{t_n})$, the projection of ν onto the time indices in $\mathbb{T}$, and $\gamma^{\mathbb{T}} := law\,(B_{t_0}, B_{t_1}, \dots, B_{t_n})$. Finally, consider the *adapted map*

$$[\mathbb{R}^d]^{n+1} \ni (b_0, \dots, b_n) \mapsto f^{\mathbb{T}}(b_0, \dots, b_n) := (f_0(b_0), f_1(b_0), \dots, f_n(b_n)) \in [\mathbb{R}^d]^{n+1}.$$

It follows that

$$f^{\mathbb{T}}(\gamma^{\mathbb{T}}) = \nu^{\mathbb{T}}.$$

Each component of $f^{\mathbb{T}}$ in *increasing* in the sense that it is the gradient of a convex function. Such a map is an example of an "increasing triangular transformation," as in [16]. It can also be understood in terms of increasingness w.r.t. lexicographical order in case $\nu^{\mathbb{T}}$ has a density. In a sense properly explained in [4], $f^{\mathbb{T}}$ sends $\gamma^{\mathbb{T}}$ into $\nu^{\mathbb{T}}$ in a canonical respect. optimal way: see respect. Proposition 5.6 and Corollary 2.10 therein.

Since this is true no matter the subgrid $\mathbb{T}$, we are entitled to think of M as an *adapted increasing rearrangement* of the Brownian motion into a martingale with given initial an final laws. Also, the aforementioned canonical/optimal character of such rearrangements should translate into the optimality of M as obtained in the previous section, and vice-versa. We now make this heuristics rigorous.

Problem (MBB) is equivalent to

$$\inf_{M_t = M_0 + \int_0^t \sigma_s\, dB_s, M_0 \sim \mu, M_1 \sim \nu} \mathbb{E}\left[\langle M - B\rangle_1\right], \tag{6.1}$$

since $\mathbb{E}\left[\langle M-B\rangle_1\right] = \mathbb{E}\left[\langle M\rangle_1\right] + \mathbb{E}\left[\langle B\rangle_1\right] - 2\mathbb{E}\left[\int_0^1 \operatorname{tr}(\sigma_t)dt\right]$, and the first two quantities in the r.h.s. do not depend on the concrete coupling (M, B). This also proves that for (6.1) it is irrelevant where B is started. Taking (6.1) as our starting point we want to formulate a transport problem between laws of stochastic processes which is compatible with this. For ease of notation we denote

$$\Omega := C([0,1]; \mathbb{R}^d).$$

In analogy to the observation in Remark 5.1, the Brownian motion above need not be d-dimensional. Keeping track of this adds notation while not providing further insight. We thus assume that we are in the full-dimensional setting; the changes needed to cover the general case are evident.

Definition 6.1. *A causal coupling between $\mathbb{P}$ and $\mathbb{Q}$ is a probability measure π on $\Omega\times\Omega$ with first and second marginals $\mathbb{P}$ and $\mathbb{Q}$ respectively, and satisfying the additional property*

$$\forall t, \forall A \in \mathcal{F}_t :\ (\ \Omega \ni x \mapsto \pi^x(A) \in [0,1]\ )\ \textit{is } (\mathbb{P}, \mathcal{F}_t)\textit{-measurable}, \tag{6.2}$$

where $\mathcal{F}$ is the $\mathbb{P}$-completed canonical filtration and π^x is a regular conditional probability of π w.r.t. the first marginal. We denote $\Pi_c(\mathbb{P},\mathbb{Q})$ the set of all such π. We also denote $\Pi_{bc}(\mathbb{P},\mathbb{Q}) = \{\pi \in \Pi_c(\mathbb{P},\mathbb{Q}) : e(\pi) \in \Pi_c(\mathbb{Q},\mathbb{P})\}$ for $e(x,y) = (y,x)$, the set of bicausal couplings.

We refer to [41, 4, 1, 3] for more on this definition. In what follows, we write $(\omega,\bar\omega)$ for a generic element in $\Omega\times\Omega$.

Lemma 6.2. *Let $\mathbb{P}$ and $\mathbb{Q}$ be martingale laws, and $\pi \in \Pi_{bc}(\mathbb{P},\mathbb{Q})$. Then the canonical process on $\Omega\times\Omega$ is a π-martingale in its own filtration.*

Proof. One can easily see that under π we have

$\{\bar\omega_s : 0 \le s \le t\}$ is π-conditionally independent from $\{\omega_s : 0 \le s \le 1\}$ given $\{\omega_s : 0 \le s \le t\}$,

$\{\omega_s : 0 \le s \le t\}$ is π-conditionally independent from $\{\bar\omega_s : 0 \le s \le 1\}$ given $\{\bar\omega_s : 0 \le s \le t\}$,

by bicausality. The first property above, for $T > t$, implies

$$\mathbb{E}^\pi[\omega_T|\,\{\omega_s,\bar\omega_s, s\le t\}\,] = \mathbb{E}^\pi[\omega_T|\,\{\omega_s, s\le t\}\,] = \mathbb{E}^{\mathbb{P}}[\omega_T|\,\{\omega_s, s\le t\}\,] = \omega_t.$$

The second property implies similarly $\mathbb{E}^\pi[\bar\omega_T|\,\{\omega_s,\bar\omega_s, s\le t\}\,] = \bar\omega_t$, so we conclude. □

Let us denote by $\mathbb{W}$ Wiener measure (started at zero) on Ω. Now the crucial connection between standard stretched Brownian motion and the present *causal transport* setting.

Lemma 6.3. *Let M be standard stretched Brownian motion from μ to ν, with $M_t = M_0 + \int_0^t \sigma_s dB_s$. Then*

$$law\,(B - B_0, M) \in \Pi_{bc}(\mathbb{W}, law\,(M)).$$

More generally, if M is stretched Brownian motion and B is as in Remark 2.3, the same conclusion holds.

Proof. We know $M_t = f_t(B_t)$, see (1.9). By Lemma 3.11, for $t < 1$ we have $B_t = (f_t)^{-1}(M_t)$. We conclude that $\{B_t - B_0 : t < 1\}$ is M-adapted. Thus $law\,(B_t - B_0|M_s, 0 \le s \le 1) = law\,(B_t - B_0|M_s, 0 \le s \le t)$. This establishes the causality property (6.2) from $law\,(M)$ to $law\,(B - B_0)$. For causality in the opposite direction, observe that $\{B_{t+h} - B_t : h \ge 0\}$ is independent from $\{(M_s, B_s) : s \le t\}$, since M is B-adapted and B is a Brownian motion. In particular, given $\{B_s - B_0 : s \le t\}$, we have that $\{M_s : s \le t\}$ and $\{B_{t+h} - B_t : h \ge 0\}$ are independent. Equivalently: given $\{B_s - B_0 : s \le t\}$, we have that $\{M_s : s \le t\}$ and $\{B_s - B_0 : s \le 1\}$ are independent. This is precisely (6.2) from $law\,(B - B_0)$ to $law\,(M)$.

In case M and B are constructed as in Remark 2.3, we may take $B_0 = 0$ and $M_0 = X$ independent of B with $X \sim \mu$. Then all arguments above can be repeated, with $M_t = f_t^X(B_t)$. Indeed, $\{B_t : t \le 1\}$ is still M adapted and B is a Brownian motion w.r.t. M (since the filtration of M is the enlargement of that of B by an independent random variable). □

We can now put the pieces together to obtain optimality of (standard) stretched Brownian motion in the sense of trajectorial laws. Let us fix a refining sequence of partition P_n of $[0,1]$ in order to define the quadratic variation $\langle\cdot\rangle$ pathwise on $C([0,1];\mathbb{R}^d)$ in the usual manner, namely

$$\omega \mapsto \langle\omega\rangle_1^{i,j} := \lim_{n\to\infty} \sum_{t_m\in P_n} (\omega^i_{t_{m+1}} - \omega^i_{t_m})(\omega^j_{t_{m+1}} - \omega^j_{t_m}),$$

when the limit exist, and otherwise $+\infty$. We then consider

$$\inf_{\substack{\mathbb{Q}\in \mathsf{M}^c(\mu,\nu)\\ \pi\in\Pi_{bc}(\mathbb{W},\mathbb{Q})}} \mathbb{E}^\pi[\langle\omega - \bar\omega\rangle_1], \tag{6.3}$$

where $\mathsf{M}^c(\mu,\nu)$ denotes the set of laws of continuous martingales indexed by $[0,1]$ starting in μ and terminating in ν.

Proposition 6.4. *Problems* (6.1) *and* (6.3) *are equivalent. In particular, let M^* be the optimizer of the former, i.e. stretched Brownian motion. Then $\mathbb{Q}^* := law(M^*)$ is optimal for the latter.*

Proof. Let $\mathbb{Q}, \pi$ be feasible for (6.3). Since

$$\begin{aligned}\mathbb{E}^\pi[\langle \omega - \bar\omega\rangle_1] &= \mathbb{E}^{\mathbb{W}}[\langle \omega\rangle_1] + \mathbb{E}^{\mathbb{Q}}[\langle \bar\omega\rangle_1] - 2\mathbb{E}^\pi[\langle \omega, \bar\omega\rangle_1]\\ &= \mathbb{E}^{\mathbb{W}}[|\omega_1|^2 - |\omega_0|^2] + \mathbb{E}^{\mathbb{Q}}[|\bar\omega_1|^2 - |\bar\omega_0|^2] - 2\mathbb{E}^\pi[\langle \omega, \bar\omega\rangle_1],\end{aligned}$$

it is enough to think of maximizing $\mathbb{E}^\pi[\langle \omega, \bar\omega\rangle_1]$ in (6.3), rather than minimizing $\mathbb{E}^\pi[\langle \omega - \bar\omega\rangle_1]$. But by Lemma 6.2 the canonical process is a π-martingale so

$$\mathbb{E}^\pi[\langle \omega, \bar\omega\rangle_1] = \mathbb{E}^\pi[\omega_1 \cdot \bar\omega_1] = \mathbb{E}^\pi[\,\mathbb{E}^\pi[\omega_1 \cdot \bar\omega_1 \,|\, \bar\omega_0]\,],$$

by the product formula and as $\omega_0 = 0$ under π. Denoting $\pi_x = \mathbb{Q} - law(\bar\omega_1|\bar\omega_0 = x)$ and $q_x = \pi - law((\bar\omega_1, \omega_1)|\bar\omega_0 = x)$ we have that the first marginal of q_x is π_x and the second one is γ^d. Indeed, by bicausality $\pi - law(\omega_1|\omega_0, \bar\omega_0) = \pi - law(\omega_1|\omega_0) = \gamma^d$, so in particular $\pi - law(\omega_1|\bar\omega_0) = \gamma^d$. Therefore

$$\mathbb{E}^\pi[\langle \omega, \bar\omega\rangle_1] = \int \mu(dx) \int q^x(dm, db)\, m \cdot b \le \int \mu(dx) \sup_{q \in \Pi(\pi_x, \gamma^d)} \int q(dm, db)\, m \cdot b. \quad (6.4)$$

By Theorem 2.2 we conclude that the value of (6.1) is greater or equal than that of (6.3). Let M^* be the optimizer of (6.1) (equiv. of (*MBB*)). By Remark 2.3 M^* is precisely built via attaining the r.h.s. of (6.4) when maximizing over Kernels π^x. By the final part of Lemma 6.3 we may build a bicausal coupling π so that in (6.4) we have equality. This proves that Problems (6.1) and (6.3) have the same value and that $law(M^*)$ is optimal for the latter. □

Remark 6.5. The discrete-time version of Problem (6.3) would have shown, in light of [4], that the optimal way to send a Gaussian random walk into a martingale is through the Knothe-Rosenblatt rearrangement (the unique increasing bicausal triangular transformation between its marginals). This is in tandem with the first paragraphs of the present part. Via Proposition 6.4 we know that stretched Brownian motion attains Problem (6.3). Hence, one can arguably describe stretched Brownian motion as the canonical/optimal Knothe-Rosenblatt rearrangement of Brownian motion with prescribed initial and final marginals.

References

[1] B. Acciaio, J. Backhoff Veraguas, and A. Zalashko. Causal optimal transport and its links to enlargement of filtrations and continuous-time stochastic optimization. *ArXiv e-prints*, 2016.

[2] L. Ambrosio and N. Gigli. A user's guide to optimal transport. In *Modelling and optimisation of flows on networks*, volume 2062 of *Lecture Notes in Math.*, pages 1–155. Springer, Heidelberg, 2013.

[3] J. Backhoff Veraguas, M. Beiglböck, M. Eder, and A. Pichler. Fine properties of process distances. *ArXiv e-prints*, 2017.

[4] J. Backhoff Veraguas, M. Beiglböck, Y. Lin, and A. Zalashko. Causal transport in discrete time and applications. *ArXiv e-prints*, June 2016.

[5] Erik J. Balder. Lectures on Young measure theory and its applications in economics. *Rend. Istit. Mat. Univ. Trieste*, 31(suppl. 1):1–69, 2000. Workshop on Measure Theory and Real Analysis (Italian) (Grado, 1997).

[6] R. Bass. Skorokhod imbedding via stochastic integrals. In Jacques Azéma and Marc Yor, editors, *Séminaire de Probabilités XVII 1981/82*, number 986 in Lecture Notes in Mathematics, pages 221–224. 1983.

[7] M. Beiglböck, A. Cox, and M. Huesmann. Optimal transport and Skorokhod embedding. *Invent. Math.*, 208(2):327–400, 2017.

[8] M. Beiglböck, A. Cox, M. Huesmann, and S. Kallblad. Measure-valued martingales and optimality of bass-type solutions to the skorokhod embedding problem. *preprint available at request*, 2017.

[9] M. Beiglböck, M. Goldstern, G. Maresch, and W. Schachermayer. Optimal and better transport plans. *J. Funct. Anal.*, 256(6):1907–1927, 2009.

[10] M. Beiglböck and C. Griessler. A land of monotone plenty. *Annali della SNS, to appear*, April 2016.

[11] 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.

[12] M. Beiglböck, M. Huesmann, and F. Stebegg. Root to Kellerer. *Séminaire de Probabilités XLVIII*, pages 1–12, 2016.

[13] M. Beiglböck and N. Juillet. On a problem of optimal transport under marginal martingale constraints. *Ann. Probab.*, 44(1):42–106, 2016.
[14] M. Beiglboeck, M. Eder, C. Elgert, and U. Schmock. Geometry of Distribution-Constrained Optimal Stopping Problems. *ArXiv e-prints*, December 2016.
[15] J. Benamou and Y. Brenier. A numerical method for the optimal time-continuous mass transport problem and related problems. In *Monge Ampère equation: applications to geometry and optimization (Deerfield Beach, FL, 1997)*, volume 226 of *Contemp. Math.*, pages 1–11. Amer. Math. Soc., Providence, RI, 1999.
[16] V. I. Bogachev, A. V. Kolesnikov, and K. V. Medvedev. Triangular transformations of measures. *Mat. Sb.*, 196(3):3–30, 2005.
[17] B. Bouchard and M. Nutz. Arbitrage and duality in nondominated discrete-time models. *The Annals of Applied Probability*, 25(2):823–859, 2015.
[18] Y. Brenier. Décomposition polaire et réarrangement monotone des champs de vecteurs. *C. R. Acad. Sci. Paris Sér. I Math.*, 305(19):805–808, 1987.
[19] Yann Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Comm. Pure Appl. Math.*, 44(4):375–417, 1991.
[20] L. Campi, I. Laachir, and C. Martini. Change of numeraire in the two-marginals martingale transport problem. *Finance Stoch.*, 21(2):471–486, June 2017.
[21] A. De March and N. Touzi. Irreducible convex paving for decomposition of multi-dimensional martingale transport plans. arXiv:1702.08298, 2017.
[22] Y. Dolinsky and H. M. Soner. Martingale optimal transport and robust hedging in continuous time. *Probab. Theory Relat. Fields*, 160(1-2):391–427, 2014.
[23] A. Galichon, P. Henry-Labordère, and N. Touzi. A stochastic control approach to no-arbitrage bounds given marginals, with an application to lookback options. *Ann. Appl. Probab.*, 24(1):312–336, 2014.
[24] W. Gangbo and R. McCann. The geometry of optimal transportation. *Acta Math.*, 177(2):113–161, 1996.
[25] N. Ghoussoub, Y.-H. Kim, and T. Lim. Optimal skorokhod embedding for radially symmetric marginals in general dimension. *preprint*, 2016.
[26] N. Ghoussoub, Y.-H. Kim, and T. Lim. Structure of optimal martingale transport plans in general dimensions. arXiv:1508.01806, 2016.
[27] N. Gozlan, C. Roberto, P-M. Samson, Y. Shu, and P. Tetali. Characterization of a class of weak transport-entropy inequalities on the line. Preprint arXiv:1509.04202v2, 2015.
[28] N. Gozlan, C. Roberto, P.-M. Samson, and P. Tetali. Kantorovich duality for general transport costs and applications. *ArXiv e-prints*, December 2014.
[29] G. Guo, X. Tan, and N. Touzi. On the monotonicity principle of optimal Skorokhod embedding problem. *ArXiv e-prints*, June 2015.
[30] P. Henry-Labordère, J. Obłój, P. Spoida, and N. Touzi. The maximum maximum of a martingale with given *n* marginals. *Ann. Appl. Probab.*, 26(1):1–44, 2016.
[31] P. Henry-Labord'ere and N. Touzi. An Explicit Martingale Version of Brenier's Theorem. *Finance Stoch., to appear*, 2016.
[32] Francis Hirsch, Christophe Profeta, Bernard Roynette, and Marc Yor. *Peacocks and associated martingales, with explicit constructions*, volume 3 of *Bocconi & Springer Series*. Springer, Milan, 2011.
[33] D. Hobson and A. Neuberger. Robust bounds for forward start options. *Math. Finance*, 22(1):31–56, 2012.
[34] M. Huesmann and D. Trevisan. A Benamou–Brenier formulation of martingale optimal transport. *arXiv:1707.01493*, 2017.
[35] S. Kallblad, X. Tan, and N. Touzi. Optimal Skorokhod embedding given full marginals and Azema-Yor peacocks. *ArXiv e-prints*, March 2015.
[36] L. Kantorovich. On the translocation of masses. *C. R. (Doklady) Acad. Sci. URSS (N.S.)*, 37:199–201, 1942.
[37] Ioannis Karatzas and Steven E. Shreve. *Brownian motion and stochastic calculus*, volume 113 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 1991.
[38] A. S. Kechris. *Classical descriptive set theory*, volume 156 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.
[39] Hans G. Kellerer. Integraldarstellung von Dilationen. In *Transactions of the Sixth Prague Conference on Information Theory, Statistical Decision Functions, Random Processes (Tech. Univ., Prague, 1971; dedicated to the memory of Antonín Špaček)*, pages 341–374. Academia, Prague, 1973.
[40] Y.-H. Kim and B. Pass. A general condition for Monge solutions in the multi-marginal optimal transport problem. *SIAM J. Math. Anal., to appear*, July 2013.
[41] R. Lassalle. Causal transference plans and their Monge-Kantorovich problems. *ArXiv e-prints*, 2013.
[42] Christian Léonard. A survey of the Schrödinger problem and some of its connections with optimal transport. *Discrete Contin. Dyn. Syst.*, 34(4):1533–1574, 2014.
[43] G. Lowther. Fitting martingales to given marginals. *ArXiv e-prints*, August 2008.
[44] G. Lowther. Limits of one-dimensional diffusions. *Ann. Probab.*, 37(1):78–106, 2009.
[45] R. McCann. A convexity theory for interacting gases and equilibrum crystals. *PhD thesis, Princeton University*, 1994.
[46] G. Monge. Memoire sur la theorie des deblais et des remblais. *Histoire de l'académie Royale des Sciences de Paris*, 1781.
[47] M. Nutz and F. Stebegg. Canonical Supermartingale Couplings. *ArXiv e-prints*, September 2016.

[48] J. Obłój and P. Siorpaes. Structure of martingale transports in finite dimensions. arXiv:1702.08433, 2017.
[49] J. Obłój, P. Spoida, and N. Touzi. Martingale inequalities for the maximum via pathwise arguments. In *In Memoriam Marc Yor-Séminaire de Probabilités XLVII*, pages 227–247. Springer, 2015.
[50] R. Tyrrell Rockafellar. *Convex analysis*. Princeton Landmarks in Mathematics. Princeton University Press, Princeton, NJ, 1997. Reprint of the 1970 original, Princeton Paperbacks.
[51] Filippo Santambrogio. *Optimal transport for applied mathematicians*, volume 87 of *Progress in Nonlinear Differential Equations and their Applications*. Birkhäuser/Springer, Cham, 2015. Calculus of variations, PDEs, and modeling.
[52] V. Strassen. The existence of probability measures with given marginals. *Ann. Math. Statist.*, 36:423–439, 1965.
[53] X. Tan and N. Touzi. Optimal transportation under controlled stochastic dynamics. *Ann. Probab.*, 41(5):3201–3240, 2013.
[54] C. Villani. *Topics in optimal transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2003.
[55] C. Villani. *Optimal Transport. Old and New*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2009.
[56] D. Zaev. On the Monge-Kantorovich problem with additional linear constraints. *ArXiv e-prints*, 2014.