Vol. 80, No. 2 DUKE MATHEMATICAL JOURNAL © November 1995

EXISTENCE AND UNIQUENESS OF MONOTONE MEASURE-PRESERVING MAPS

ROBERT J. McCANN

Introduction. Given a pair of Borel probability measures μ and ν on $\mathbf{R}^d$, it is natural to ask whether ν can be obtained from μ by redistributing its mass in a canonical way. In the case of the line $d = 1$ the answer is clear: as long as both measures are free from atoms—$\mu[\{x\}] = \nu[\{x\}] = 0$—there is a map $y(x)$ of the line to itself for which

$$\int_{-\infty}^{x} d\mu = \int_{-\infty}^{y(x)} d\nu.$$

Uniquely determined μ-almost everywhere, this map may be taken to be nondecreasing by a suitable choice of $y(x) \in \mathbf{R} \cup \{\pm\infty\}$ at the remaining points. Interpreting μ and ν as the initial and final distribution of a one-dimensional fluid, the transformation $y(x)$ gives a rearrangement of fluid particles yielding final state ν from the initial state μ; this rearrangement is characterized by the fact that it preserves particle ordering, obviating any need for two particles to cross. Although the generalization of this construction to higher dimensions is the focus of this paper, the one-dimensional case will be pursued slightly further: when the measures are absolutely continuous with respect to Lebesgue—$d\mu(x) = f(x)dx$ and $d\nu(y) = g(y)dy$—then, formally at least (neglecting regularity issues), the fundamental theorem of calculus yields

$$y'(x)g(y(x)) = f(x). \tag{1}$$

When μ and ν measure $\mathbf{R}^d$ rather than $\mathbf{R}$, the properties of y one might hope to preserve are not clear. An answer to this question has been provided by a theorem of Brenier [1], [2], which shows under restrictions on μ and ν, that a measure-preserving transformation $y: (\mathbf{R}^d, \mu) \to (\mathbf{R}^d, \nu)$ can be realized as the gradient of a convex function. In particular, y will be irrotational and will not involve crossings: $(1 - t)x + ty(x) = (1 - t)x' + ty(x')$ implies $x = x'$ if $t \in [0, 1)$. Moreover, this theorem turns out to have striking applications which will shortly be indicated. Motivated by these applications, our present purpose is to extend

Received 12 December 1994. Revision received 9 March 1995.
The author gratefully acknowledges support provided by a postdoctoral fellowship of the Natural Sciences and Engineering Research Council of Canada.

the theorem to a larger class of measures μ and ν. The improvement is achieved through a simple, geometrical proof relying on Rockafellar's characterization [3] of the gradient of a convex function.

A precise statement requires a bit of notation. Let $\mathscr{P}(\mathbf{R}^d)$ be the space of Borel probability measures on $\mathbf{R}^d$, so $\mu \in \mathscr{P}(\mathbf{R}^d)$ is a positive measure with $\mu[\mathbf{R}^d] = 1$.

Definition 1. A measure $\mu \in \mathscr{P}(\mathbf{R}^d)$ together with a Borel transformation $y: \mathbf{R}^d \to \mathbf{R}^n$ defined μ-almost everywhere, induces a measure $y_{\#}\mu$ on $\mathbf{R}^n$ defined by $y_{\#}\mu[M] := \mu[y^{-1}(M)]$ for Borel $M \subset \mathbf{R}^n$. A Borel probability measure itself, $y_{\#}\mu$ is called the *push-forward* of μ through y.

The map y is said to be *measure-preserving* between μ and $y_{\#}\mu$, or to *push μ forward to $y_{\#}\mu$*. For a Borel function f on $\mathbf{R}^n$, the change of variables theorem states that when either integral is defined,

$$\int_{\mathbf{R}^n} f \, d(y_{\#}\mu) = \int_{\mathbf{R}^d} f(y(x)) \, d\mu(x).$$

A *convex* function ψ on $\mathbf{R}^d$ refers to a function $\psi: \mathbf{R}^d \to \mathbf{R} \cup \{+\infty\}$ for which $\psi((1-t)x + tx') \leqslant (1-t)\psi(x) + t\psi(x')$ when the latter is finite and $t \in (0, 1)$. Such a function will be continuous on the interior of the convex set $\operatorname{dom} \psi := \{\psi(x) < \infty\}$, and differentiable except on a set of Hausdorff dimension $d - 1$ in $\operatorname{dom} \psi$ [4]. The *monotone map* of our title refers to the gradient $\nabla\psi$ of such a function.

MAIN THEOREM. *Let $\mu, \nu \in \mathscr{P}(\mathbf{R}^d)$ and suppose μ vanishes on (Borel) subsets of $\mathbf{R}^d$ having Hausdorff dimension $d - 1$. Then there exists a convex function ψ on $\mathbf{R}^d$ whose gradient $\nabla\psi$ pushes μ forward to ν. Although ψ may not be unique, the map $\nabla\psi$ is uniquely determined μ-almost everywhere.*

In the special case $d\nu(y) = k_d^{-1}(1 + |y|^2)^{-(d+2)/2} dy$, this theorem is equivalent to a classical result of Aleksandrov [5] asserting the existence of a unique convex function with prescribed scalar curvature: if $d\mu(x) = f(x)dx$, the graph of ψ will have Gauss curvature $k_d f(x)$ at the point $(x, \psi(x))$—see (3) below. The constant k_d which ensures that $\nu[\mathbf{R}^d] = 1$ is the volume of the unit ball in $\mathbf{R}^d$, while the image of graph(ψ) under the Gauss map covers half of the unit sphere $\mathbf{S}^d \subset \mathbf{R}^{d+1}$.

That this theorem might be true for more general ν was foreshadowed in work of Dowson and Landau [6] and Knott and Smith [7] from the early 1980s. Both studied the joint probability measure which is maximally correlated among those $\gamma \in \mathscr{P}(\mathbf{R}^d \times \mathbf{R}^d)$ having μ and ν as their *marginals*: $M \subset \mathbf{R}^d$ Borel implies $\mu[M] = \gamma[M \times \mathbf{R}^d]$ and $\gamma[\mathbf{R}^d \times M] = \nu[M]$. Letting $\Gamma(\mu, \nu)$ denote the set of such γ, and $\langle x, y \rangle$ denote the Euclidean inner product on $\mathbf{R}^d$, to be *maximally correlated* means that $\gamma' \in \Gamma(\mu, \nu)$ is the optimal solution to the (infinite-dimensional) linear program

$$\sup_{\gamma \in \Gamma(\mu, \nu)} \int \langle x, y \rangle \, d\gamma. \tag{2}$$

Knott and Smith showed that when ψ satisfies the conclusions of the main theorem, then γ' is given by the push-forward of μ through the map $\mathrm{id} \times \nabla\psi$ taking x to $(x, \nabla\psi(x))$. Dowson and Landau gave heuristics indicating that if $(\mathrm{id} \times y)_{\#}\mu$ is maximally correlated, then the transformation y of $\mathbf{R}^d$ must be irrotational. Both assumed regularity. Finally, Brenier established our main theorem and related assertions, under the following restrictions on μ and ν [2]: the first moment $\int |y| \, d\nu(y)$ of ν should be finite, while μ should be absolutely continuous with respect to Lebesgue—$d\mu(x) = f(x)\,dx$—with $f(x)$ vanishing outside some bounded, smooth, connected domain $\Omega \subset \mathbf{R}^d$ and bounded away from zero and infinity on Ω. In this case ψ is unique, and lies in the Sobolev space $W^{1,1}(\Omega, \mu)$. His proof proceeded through an analysis of the variational problem (2) and its dual, both of which attain finite extrema under the additional hypotheses. He went on to remark that the theorem was expected to hold for unbounded Ω, as long as $\int |x|^p \, d\mu(x)$ and $\int |y|^q \, d\nu(y)$ are both finite for some Hölder conjugates $p^{-1} + q^{-1} = 1$. Simpler proofs, based on the dual problem alone, were later discovered by Caffarelli and Varadhan [8] and Gangbo [9].

From our point of view, the content of this theorem is geometrical rather than analytic, which suggests that $W^{1,1}(\Omega, \mu)$ is not its natural setting. The results established here reflect this bias: they go beyond the earlier work in several respects, to include cases in which (2) will not be finite. Proving uniqueness for these cases presents the most serious obstacle, since the variational formulation is not of use. This difficulty is circumvented through a geometrical approach inspired by [5]. The existence assertion, on the other hand, may in principle be extracted from Brenier's result via a continuity argument. Instead, a new proof is provided, which has the advantage of being largely self-contained. Moreover, the idea which underlies it is transparent in a pairing problem involving the elements from two finite subsets of $\mathbf{R}^d$. The solution to this problem is a finite optimization, after which a continuity argument yields Theorem 6: for $\mu, \nu \in \mathscr{P}(\mathbf{R}^d)$, there exists $\gamma \in \Gamma(\mu, \nu)$ whose support in $\mathbf{R}^d \times \mathbf{R}^d$ enjoys a geometrical property known as cyclical monotonicity. Introduced by Rockafellar in [3], this property is characteristic of the gradient of a convex function: a theorem of his allows the function ψ to be recovered. As long as ψ is differentiable μ-almost everywhere, then $\nabla\psi_{\#}\mu = \nu$. This will certainly be the case if none of the mass of μ concentrates on a set of dimension $d - 1$. An example shows that existence and uniqueness both may fail if this condition should be violated.

Finally, some applications of the theorem must be mentioned. When $d\mu(x) = f(x)\,dx$ and $d\nu(y) = g(y)\,dy$, Caffarelli [10], [8], [11] has shown that conditions on f and g ensure the regularity of ψ. In this case the higher-dimensional change of variables formula analogous to (1) holds—

$$\det\left[\frac{\partial^2\psi}{\partial x_i \partial x_j}(x)\right] g(\nabla\psi(x)) = f(x). \tag{3}$$

This method represents a state of the art technique for constructing convex solutions ψ to the Monge-Ampère equation (3).

Brenier's investigations were motivated from a different direction: he was interested in establishing a rearrangement theorem for $\mathbf{R}^d$-valued functions y on a measure space $(\mathbf{R}^d, \mu)$, and he succeeded in expressing y uniquely as the composition $\nabla\psi \circ s$ of a monotone map with a measure-preserving transformation $s_{\#}\mu = \mu$. Here ψ is chosen so that $\nabla\psi_{\#}\mu = y_{\#}\mu$. This theorem generalizes, among other things, the polar factorization of matrices and the Hodge decomposition of vector fields; it is discussed along with further remarks and references in [2].

Our interest in these results was stimulated from yet a third direction: a study of the equilibrium states for an attracting gas led us to discover a new set of convexity relations. Formulated via an alternative convex structure on $\mathscr{P}(\mathbf{R}^d)$, these inequalities include generalizations of the Brunn-Minkowski theorem from sets to measures, and yield uniqueness results not only for attracting gases but also for equilibrium crystals [12], [13]. They are based on a construction which requires an irrotational measure-preserving mapping, free from crossings, between any pair of probability measures on $\mathbf{R}^d$; the monotone map $\nabla\psi$ provides the smoothest possible of such mappings.

Without further ado, the existence assertion of the main theorem will be established; its proof, and a more general existence result, form the next part of the present manuscript. The final section is comprised of the uniqueness arguments, along with the example in which existence and uniqueness fail. An appendix establishes a version of the implicit function theorem used in the uniqueness proof; it applies to functions which, though not continuously differentiable, are differences of two convex functions.

Some results discussed herein appear in a preliminary form as an appendix to my Princeton thesis [12]. It is a pleasure to express my gratitude to my advisor, Elliott Lieb, who drew my attention to the work of Brenier, as well as to Michael Aizenman, Andrew Browder, Walter Craig, Michael Loss, Andrew Mayer, Millie Niss, and Jan-Philip Solovej, who provided fruitful discussions concerning these ideas.

Existence of monotone measure-preserving maps. Although they will not be required until Proposition 10, it is useful to review some facts of life concerning convex functions $\psi\colon \mathbf{R}^d \to \mathbf{R} \cup \{+\infty\}$; the case $\psi(x) = \infty$ is excluded by convention. Rockafellar's text [14, especially §23.2–4, §24.1, §24.4–5, §24.8–9, §25.1, and §7.4] provides the standard reference.

The gradient of a convex function may exist only almost everywhere, but whenever ψ is finite in a neighbourhood of x, its graph admits a supporting hyperplane: there is some $y \in \mathbf{R}^d$ such that

$$\psi(z) \geqslant \langle y, z - x \rangle + \psi(x) \tag{4}$$

for all $z \in \mathbf{R}^d$. In this case y is called a *subgradient* of ψ at x, motivating the following definition.

Definition 2. The *subdifferential* of a convex function ψ on $\mathbf{R}^d$ refers to the subset $\partial\psi \subset \mathbf{R}^d \times \mathbf{R}^d$ of pairs (x, y) satisfying (4) for all $z \in \mathbf{R}^d$.

The subgradients of ψ at x will form a closed and convex set $\partial\psi(x) := \{y | (x, y) \in \partial\psi\}$. Since ψ is finite somewhere, $\partial\psi(x)$ will be empty unless $\psi(x) < \infty$; it will be bounded and nonempty precisely when x is from the interior int dom ψ of dom ψ. Differentiability of ψ at x turns out to be equivalent to the existence of a unique subgradient $y \in \partial\psi(x)$, in which case $y = \nabla\psi(x)$. Moreover, y_n converges to $\nabla\psi(x)$ whenever the latter exists, and $y_n \in \partial\psi(x_n)$ with $x_n \to x$; this fact is often used in lieu of continuity for $\nabla\psi$. Modifying the values of $\psi(x)$ on the boundary of dom ψ (if necessary) ensures ψ lower semicontinuous, and can only enlarge its subdifferential $\partial\psi$; the modification forces $\partial\psi$ to be a closed subset of $\mathbf{R}^d \times \mathbf{R}^d$.

Summing inequality (4) with $z_i = x_{i+1}$ around a cycle in dom ψ shows that the subdifferential $\partial\psi$ of a convex function enjoys a property known as cyclical monotonicity.

Definition 3. A subset $S \subset \mathbf{R}^d \times \mathbf{R}^d$ is said to be *cyclically monotone* if for any finite number of points $(x_i, y_i) \in S$, $i = 1 \cdots k$,

$$\langle y_1, x_2 - x_1 \rangle + \langle y_2, x_3 - x_2 \rangle + \cdots + \langle y_k, x_1 - x_k \rangle \leqslant 0. \tag{5}$$

Rockafellar's theorem [3]—which is ingenious but not hard to prove—asserts a converse: any cyclically monotone set $S \subset \mathbf{R}^d \times \mathbf{R}^d$ is contained in the subdifferential of some convex function on $\mathbf{R}^d$. Note that this is an integrability result: if $S = \mathrm{graph}(\nabla\psi)$ were known to hold for a differentiable function ψ, then the two-point ($k = 2$) inequalities alone would guarantee convexity of ψ.

Although Definition 3 is not transparent, for finite S cyclical monotonicity means that the x and y occurring in $(x, y) \in S$ have been paired so as to maximize the correlation $\sum_S \langle x, y \rangle$, or equivalently, to minimize the sum of the squared distances $|x - y|^2$. The next proposition exploits this observation. Together with Rockafellar's theorem, it constitutes the essence of our existence argument.

PROPOSITION 4 (Cyclical monotonicity of correlated pairs). *Choose n points of origin $x_\alpha \in \mathbf{R}^d$ and n destinations $y_\alpha \in \mathbf{R}^d$, where $\alpha \in \{1, \ldots, n\}$ and neither the x_α nor the y_α need be distinct. For some reordering of the x_α, the set $S = \{(x_\alpha, y_\alpha)\}_\alpha$ will be cyclically monotone.*

Proof. Order the x_α so that $C(\sigma) := \sum_\alpha \langle x_{\sigma(\alpha)}, y_\alpha \rangle$ is maximized by the identity permutation id among permutations σ on n letters. Any k (distinct) points in S may be specified via their labels $\alpha(1), \ldots, \alpha(k)$. Now let $\sigma = (\alpha(1)\alpha(2) \cdots \alpha(k))$ permute the $\alpha(i)$ cyclically (while fixing the remaining $n - k$ letters). The condition (5) for cyclical monotonicity of S is recovered from $C(\sigma) \leqslant C(\mathrm{id})$ and the observation

$$C(\sigma) - C(\mathrm{id}) = \sum_{i=1}^{k} \langle x_{\sigma(\alpha(i))} - x_{\alpha(i)}, y_{\alpha(i)} \rangle. \qquad \square$$

The next theorem follows immediately from this proposition when the measures $\mu, \nu \in \mathscr{P}(\mathbf{R}^d)$ are sums of point masses.

Definition 5. The *support* spt μ of a measure $\mu \in \mathscr{P}(\mathbf{R}^d)$ is the smallest closed subset $M \subset \mathbf{R}^d$ having total mass $\mu[M] = 1$.

THEOREM 6 (Existence of monotonically correlated measures). *Let $\mu, \nu \in \mathscr{P}(\mathbf{R}^d)$. There exists a Borel probability measure γ on $\mathbf{R}^d \times \mathbf{R}^d$ with cyclically monotone support having μ and ν as its marginals.*

The proof will be postponed to state two elementary lemmas. They provide enough analysis to prove the theorem in full generality.

By the Riesz-Markov theorem, $\mathscr{P}(\mathbf{R}^d)$ lies in the Banach space dual $C_\infty(\mathbf{R}^d)^*$ of the continuous functions vanishing at infinity under the sup norm. In fact, $\mathscr{P}(\mathbf{R}^d)$ is the intersection of the unit sphere with the cone of positive measures. The first lemma indicates that $\mathscr{P}(\mathbf{R}^d)$ should be topologized by the weak-$*$ topology it inherits from $C_\infty(\mathbf{R}^d)^*$—that is, by convergence against $C_\infty(\mathbf{R}^d)$ test functions. Dirac's symbol δ_x denotes the measure assigning unit mass to M whenever $x \in M \subset \mathbf{R}^d$.

LEMMA 7 (Weak-$*$ density of point mass sums). *The following measures form a weak-$*$ dense subset of $\mathscr{P}(\mathbf{R}^d) \subset C_\infty(\mathbf{R}^d)^*$:*

$$\left\{ \frac{1}{n} \sum^{n} \delta_{x_i} \,\middle|\, n \in \mathbf{N}, x_1, \ldots, x_n \in \mathbf{R}^d \right\}. \tag{6}$$

Proof. Consider the set of positive measures from the unit ball in $C_\infty(\mathbf{R}^d)^*$. This set is obviously convex, and is weak-$*$ compact by the Banach-Alaoglu theorem. Its extreme points are the zero measure and the Dirac point masses δ_x. The Krein-Milman theorem then ensures that convex combinations $\sum^m t_i \delta_{x_i}$ with $t_i \geqslant 0$ and $\sum^m t_i \leqslant 1$ are dense in this set (and in $\mathscr{P}(\mathbf{R}^d)$ a fortiori). When approximating measures in $\mathscr{P}(\mathbf{R}^d)$, it costs no generality to assume $\sum^m t_i = 1$, nor to restrict the t_i to be rational. In this case each convex combination is a measure from (6): n represents a common denominator for the t_i. □

The second lemma shows that weak-$*$ limits preserve the desired properties —specified marginals and cyclical monotonicity of support—of measures in $\mathscr{P}(\mathbf{R}^d \times \mathbf{R}^d)$.

Definition 8. A measure $\gamma \in \mathscr{P}(\mathbf{R}^d \times \mathbf{R}^d)$ is said to have $\mu, \nu \in \mathscr{P}(\mathbf{R}^d)$ as its (left and right) *marginals* if $M \subset \mathbf{R}^d$ Borel implies $\mu[M] = \gamma[M \times \mathbf{R}^d]$ and $\gamma[\mathbf{R}^d \times M] = \nu[M]$. The set of such γ is denoted by $\Gamma(\mu, \nu)$.

LEMMA 9 (Weak-$*$ limits preserve monotone correlations). *Let a sequence $\gamma_n \in \mathscr{P}(\mathbf{R}^d \times \mathbf{R}^d)$ converge weak-$*$ to $\gamma \in C_\infty(\mathbf{R}^d \times \mathbf{R}^d)^*$. Then*

(i) *γ will have cyclically monotone support if γ_n does for each n;*

(ii) *if the left and right marginals of $\gamma_n \in \Gamma(\mu_n, \nu_n)$ converge weak-$*$ in $C_\infty(\mathbf{R}^d)^*$ to limits $\mu, \nu \in \mathscr{P}(\mathbf{R}^d)$, then $\gamma \in \Gamma(\mu, \nu)$.*

Proof. (i) If spt γ fails to be cyclically monotone, then it contains k points (x_i, y_i) which violate (5). Choosing a sufficiently small neighbourhood $U_i \subset \mathbf{R}^d \times \mathbf{R}^d$ for each of these k points, the inequality continues to be violated whenever one or more (x_i, y_i) is replaced by a point $(u_i, v_i) \in U_i$. Since (x_i, y_i) lies in spt γ, $\gamma[U_i] > 0$. Therefore $\gamma_n[U_i] > 0$ holds for large n and all i. But spt γ_n fails to be cyclically monotone if it intersects all of the U_i.

(ii) Any $f \in C_\infty(\mathbf{R}^d)$ may be viewed as a function on $\mathbf{R}^d \times \mathbf{R}^d$ which is independent of its second argument. The assertion would be immediate if $f \in C_\infty(\mathbf{R}^d \times \mathbf{R}^d)$, but this will not be the case unless $f = 0$. What is true is that f extends naturally to a continuous function on $\overline{\mathbf{R}}^d \times \overline{\mathbf{R}}^d$, where $\overline{\mathbf{R}}^d$ denotes the one-point compactification of $\mathbf{R}^d$. The desired result follows by showing that γ coincides with a weak-$*$ limit of the γ_n in $C(\overline{\mathbf{R}}^d \times \overline{\mathbf{R}}^d)^*$. Such a limit point γ' may be extracted from the γ_n via the Banach-Alaoglu theorem; since $C(\overline{\mathbf{R}}^d \times \overline{\mathbf{R}}^d)$ includes the constant functions, $\gamma'[\overline{\mathbf{R}}^d \times \overline{\mathbf{R}}^d] = 1$. Moreover, the actions of γ and γ' coincide on $C_\infty(\mathbf{R}^d \times \mathbf{R}^d)$, so that $\gamma' - \gamma$ can only be supported "at ∞". If $f \in C(\overline{\mathbf{R}}^d \times \overline{\mathbf{R}}^d)$ is independent of its second argument and $f(\infty, y) = 0$, then using $\gamma_n \in \Gamma(\mu_n, \nu_n)$,

$$\int f\, d\gamma' = \lim_n \int f\, d\gamma_n = \lim_n \int f\, d\mu_n = \int f\, d\mu.$$

Taking a supremum over $\|f\|_\infty \leqslant 1$ implies $\gamma'[\mathbf{R}^d \times \overline{\mathbf{R}}^d] \geqslant \mu[\mathbf{R}^d] = 1$, so μ is in fact the left marginal of γ'. A similar argument shows that ν is the right marginal of γ'. Finally, $\gamma'[\mathbf{R}^d \times \mathbf{R}^d] = 1$ forces $\gamma' - \gamma = 0$. □

Proof of Theorem 6. Use Lemma 7 to choose a sequence $\mu_k \in \mathscr{P}(\mathbf{R}^d)$ converging weak-$*$ to μ in which each μ_k is a (normalized) sum of point masses. Approximate ν similarly by ν_k. Fixing k, there are $x_i, y_i \in \mathbf{R}^d$ such that:

$$\mu_k = \frac{1}{n}\sum_{i=1}^{n} \delta_{x_i}, \qquad \nu_k = \frac{1}{m}\sum_{i=1}^{m} \delta_{y_i}.$$

It is no loss of generality to assume $n = m$: both may be replaced by nm provided the multiplicity of each x_i (resp. y_i) occurring in the sum is increased by a factor of m (resp. n). Reorder the x_i by Proposition 4, so that $\{(x_i, y_i)\}_i$ is cyclically monotone. The joint measure

$$\gamma_k = \frac{1}{n}\sum_{i=1}^{n} \delta_{(x_i, y_i)}$$

on $\mathbf{R}^d \times \mathbf{R}^d$ then has cyclically monotone support, with μ_k and ν_k as its marginals. Now the γ_k lie in the unit ball of $C_\infty(\mathbf{R}^d \times \mathbf{R}^d)^*$, so a weak-$*$ limit point γ may be extracted by the Banach-Alaoglu theorem. γ is a positive measure, which by Lemma 9 must have cyclically monotone support and μ and ν as its marginals. In particular, γ has unit mass. □

A final proposition recovers the connection between measures of cyclically monotone support and convex functions. The map $x \to (x, \nabla\psi(x))$ will be denoted by $\mathrm{id} \times \nabla\psi$.

Proposition 10 (Monotone maps from monotone correlations). *Suppose a measure $\gamma \in \mathscr{P}(\mathbf{R}^d \times \mathbf{R}^d)$ is supported on the subgradient $\partial\psi \supset \operatorname{spt}\gamma$ of a convex function ψ on $\mathbf{R}^d$. Let μ and ν denote the marginals of $\gamma \in \Gamma(\mu, \nu)$. If μ vanishes on (Borel) sets of Hausdorff dimension $d - 1$, then the gradient $\nabla\psi$ pushes μ forward to ν; in fact, $\gamma = (\mathrm{id} \times \nabla\psi)_{\#}\mu$.*

Proof. Modifying ψ on the boundary of the convex set $\operatorname{dom}\psi := \{\psi < \infty\}$ ensures $\partial\psi$ as large as possible and closed; the set $\operatorname{dom}\nabla\psi$ of differentiability for ψ lies in the interior of $\operatorname{dom}\psi$, so $\nabla\psi$ is unaffected, while $\operatorname{spt}\gamma \subset \partial\psi$ continues to hold. Since $\partial\psi \subset \operatorname{dom}\psi \times \mathbf{R}^d$, the left marginal μ of γ is supported on the closure of $\operatorname{dom}\psi$. To begin, one needs to know the transformation $\nabla\psi$, and therefore $\mathrm{id} \times \nabla\psi$, is defined μ-almost everywhere and Borel measurable. Differentiability of ψ fails only on a Borel set of dimension $d - 1$ in $\operatorname{dom}\psi$ [4]. Since $\operatorname{dom}\psi$ forms a convex (a fortiori locally Lipschitz) domain, its boundary is also $d - 1$ dimensional, and $\mu[\operatorname{dom}\nabla\psi] = 1$ follows by hypothesis on μ. Measurability of $\nabla\psi$ is manifest since it coincides with the pointwise limit of a sequence of continuous approximants

$$\langle \nabla\psi(x), z \rangle = \lim_{n\to\infty} n(\psi(x + z/n) - \psi(x)).$$

If $(\mathrm{id} \times \nabla\psi)_{\#}\mu = \gamma$ is now verified, $\nabla\psi_{\#}\mu = \nu$ will be immediate from Definitions 1 and 8.

To complete the proof, it suffices to show that the measure $(\mathrm{id} \times \nabla\psi)_{\#}\mu$ coincides with γ on products $M \times N$ of Borel sets $M, N \subset \mathbf{R}^d$; the semialgebra of such products generates the Borel sets in $\mathbf{R}^d \times \mathbf{R}^d$. Define $S := \{(x, \nabla\psi(x)) \mid x \in \operatorname{dom}\nabla\psi\}$. Since $\partial\psi$ is closed and contains $\operatorname{spt}\gamma$, $S = (\operatorname{dom}\nabla\psi \times \mathbf{R}^d) \cap \partial\psi$ is Borel and contains a set $(\operatorname{dom}\nabla\psi \times \mathbf{R}^d) \cap \operatorname{spt}\gamma$ of full measure for γ. Thus $\gamma[Z \cap S] = \gamma[Z]$ for $Z \subset \mathbf{R}^d \times \mathbf{R}^d$. Applied to

$$(M \times N) \cap S = ((M \cap (\nabla\psi)^{-1}N) \times \mathbf{R}^d) \cap S,$$

this yields

$$\begin{aligned} \gamma[M \times N] &= \gamma[(M \cap (\nabla\psi)^{-1}N) \times \mathbf{R}^d] \\ &= \mu[M \cap (\nabla\psi)^{-1}N] \\ &= (\mathrm{id} \times \nabla\psi)_{\#}\mu[M \times N]. \end{aligned}$$

$\gamma \in \Gamma(\mu, \nu)$ implies the second equation; Definition 1 implies the third. □

Proof of existence in Main Theorem. Theorem 6 guarantees a joint measure $\gamma \in \Gamma(\mu, \nu)$ having cyclically monotone support. Rockafellar's theorem provides a convex function ψ whose subgradient contains $\operatorname{spt} \gamma$. By Proposition 10, $\nabla\psi$ pushes forward μ to ν. □

Remark 11 (A generalized existence result). Before proceeding to the uniqueness question, it seems worthwhile to point out the generality of the arguments which lead to Theorem 6. Let X and Y be locally compact, σ-compact Hausdorff spaces, so that the Riesz-Markov theorem holds, and the Borel probability measures $\mathscr{P}(X)$ on X are regular. If $c(x, y) \in \mathbf{R}$ is a jointly continuous *cost function* on $X \times Y$, then the same proof, mutatis mutandis, yields the following theorem.

THEOREM 12. *If $\mu \in \mathscr{P}(X)$ and $\nu \in \mathscr{P}(Y)$, then there is a measure $\gamma \in \mathscr{P}(X \times Y)$ with μ and ν as its marginals, whose support satisfies the following condition: for any finite number of points $(x_i, y_i) \in \operatorname{spt} \gamma$,*

$$\sum_{i=1}^{k} c(x_i, y_i) \leqslant \sum_{i=1}^{k} c(x_{i+1}, y_i). \tag{7}$$

The second sum is cyclical, $x_{k+1} := x_1$.

Although it will not be proven here, we believe that $\int_{X \times Y} c(x, y)\, d(\gamma' - \gamma) \geqslant 0$ whenever γ' shares the marginals of γ. In other words, condition (7) on its support should be necessary and sufficient to ensure that γ be an optimal solution of the associated Monge-Kantorovich mass transport problem (see, e.g., [15]).

Note added in revision. Since the submission of this work, I have learned of recent investigations which provide substantial links between (7) and the Monge-Kantorovich problem. Smith and Knott [16]—who refer to the condition (7) as *c-cyclical monotonicity*—noted that its sufficiency for optimal mass transport follows from a theorem of Rüschendorf [17]. Abdellaoui and Heinich [18] obtained necessity results. For the cost function $c(x, y) = |x - y|^2$ on $X = Y = \mathbf{R}^d$, or equivalently, $c(x, y) = -2\langle x, y\rangle$, such results are of course present in Brenier's work [1], [2], but also in a paper by Rüschendorf and Rachev [19] and a preprint of Cuesta-Albertos, Matrán, and Tuero-Díaz [20]. Some developments in these last two papers parallel those of the present manuscript, although they are restricted to cases where the measures μ and ν have finite second moments. Under this assumption, an analog to Theorem 6 is proved in [19], and to our main theorem in [20]. Both papers consider the possibility that X and Y are infinite-dimensional topological dual spaces. I thank J. A. Cuesta-Albertos, M. Gelbrich, S. T. Rachev, and L. Rüschendorf for bringing this literature to my attention.

Uniqueness of monotone measure-preserving maps. It remains to establish uniqueness of the monotone map $\nabla\psi$ which pushes μ forward to ν. The proof employs a refinement of Aleksandrov's argument for the uniqueness of convex surfaces with prescribed integral curvature. The idea is that two monotone maps

$\nabla\psi$ and $\nabla\phi$ differing at $p \in \operatorname{spt}\mu$, push μ forward to measures which disagree on some $Y \subset \mathbf{R}^d$ (9). Adding a constant to the convex function ϕ, so that $\phi(p) = \psi(p)$, the set Y may be chosen to consist of the subgradients for ϕ on $M := \{\phi > \psi\}$; it is denoted $\partial\phi(M) := \bigcup_{x \in M} \partial\phi(x)$. Any hyperplane which supports ϕ on M must cut off a part of graph(ψ). A parallel hyperplane supporting ψ must therefore do so over M. Thus $\nabla\psi^{-1}(Y) \subset M$. This idea forms the essence the following lemma, asserted in [5] and used here to prove (9). As before, $\nabla\psi$ is defined where ψ is differentiable: on dom $\nabla\psi$.

LEMMA 13 (Aleksandrov). *Let ϕ and ψ be convex on $\mathbf{R}^d$, differentiable at p with $\phi(p) = \psi(p) = 0$, and $\nabla\phi(p) \neq \nabla\psi(p) = 0$. Define $M := \{\phi > \psi\} \subset \operatorname{dom}\psi$ and $X := \nabla\psi^{-1}(\partial\phi(M))$. Then $X \subset M$, while p lies a positive distance from X.*

Proof. To obtain the inclusion, let $x \in X$ and $y := \nabla\psi(x)$. Then there exists $m \in M$ such that $y \in \partial\phi(m)$. For any $z \in \mathbf{R}^d$

$$\phi(z) \geqslant \langle y, z - m\rangle + \phi(m) \quad \text{and}$$

$$\psi(m) \geqslant \langle y, m - x\rangle + \psi(x).$$

Noting that $\phi(m) > \psi(m)$, these inequalities combine to yield

$$\phi(z) > \langle y, z - x\rangle + \psi(x). \tag{8}$$

Taking $z = x$ shows $x \in M$.

Next, suppose a sequence $x_n \in X$ converges to p. Again, there exist $m_n \in M$ such that $\nabla\psi(x_n) \in \partial\phi(m_n)$. Now $\nabla\psi(p) = 0$ implies $\psi \geqslant 0$ and $\nabla\psi(x_n) \to 0$ by convexity of ψ and continuity of $\partial\psi$; on the other hand, $\nabla\phi(p) \neq 0$ implies $\phi(z) < 0$ for some z near p. Making use of (8) once more yields

$$\begin{aligned} 0 > \phi(z) &> \langle \nabla\psi(x_n), z - x_n\rangle + \psi(x_n) \\ &\geqslant -|\nabla\psi(x_n)|\,|z - x_n|. \end{aligned}$$

Since $x_n \to p$ and $\nabla\psi(x_n) \to 0$, a contradiction is obtained. The conclusion is that p cannot lie in the closure of X. □

Proof of uniqueness in Main Theorem. Let ψ and ϕ be convex on $\mathbf{R}^d$ with $\nabla\phi_{\#}\mu = \nabla\psi_{\#}\mu = \nu$, and suppose that $\nabla\phi = \nabla\psi$ fails to hold μ-almost everywhere. Then there is a point $p \in \operatorname{spt}\mu$ at which ψ and ϕ are differentiable, but their gradients differ $\nabla\psi(p) \neq \nabla\phi(p)$. Both ψ and ϕ may be chosen lower semi-continuous, so that $\partial\phi$ is closed, and constants may be subtracted from each to ensure $\psi(p) = \phi(p) = 0$; neither modification affects the gradients.

Because $p \in \operatorname{spt}\mu$, each of its neighbourhoods must have positive measure under μ. Since the gradient maps differ at p, Corollary 19 of our implicit function theorem provides a neighbourhood of p in which $\psi = \phi$ occurs only on a set of Hausdorff dimension $d - 1$. This set has measure zero for μ. Exchanging ψ with ϕ if necessary, all neighbourhoods of p must intersect $M := \{x \in \operatorname{int} \operatorname{dom} x \mid \psi(x) < \phi(x)\}$ in a set of positive measure for μ. As a final normalization for ψ and ϕ, the same linear function may be subtracted from each to yield $\nabla\psi(p) = 0$; this corresponds to a translation of ν.

Now M is open, while $\partial\phi$ is closed by the lower semicontinuity of ϕ. Thus $Y := \partial\phi(M)$ is Borel, in fact σ-compact: it's the projection of a σ-compact set $(M \times \mathbf{R}^d) \cap \partial\phi$ onto $\mathbf{R}^d$. A contradiction will be derived by showing that the two push-forwards of μ cannot agree on Y:

$$\mu[\nabla\psi^{-1}(Y)] < \mu[M] \leqslant \mu[\nabla\phi^{-1}(Y)]. \tag{9}$$

The second inequality is obvious from the fact that $\mu[\operatorname{dom} \nabla\phi] = \nu[\mathbf{R}^d] = 1$, and the inclusion

$$M \cap \operatorname{dom} \nabla\phi \subseteq \nabla\phi^{-1}(Y).$$

The first inequality follows from Lemma 13, which shows that $\nabla\psi^{-1}(Y) \subset M$ and excludes some neighbourhood U of p; since $\mu[U \cap M] > 0$, the inequality is strict and the proof is complete. □

COROLLARY 14 (Uniqueness of monotone correlations). *Suppose $\mu, \nu \in \mathscr{P}(\mathbf{R}^d)$, and that one of these measures vanishes on all sets of Hausdorff dimension $d - 1$. Then the joint measure $\gamma \in \Gamma(\mu, \nu)$ with cyclically monotone support is unique.*

Proof. From its definition, cyclical monotonicity of $S \subset \mathbf{R}^d \times \mathbf{R}^d$ is equivalent to that of $S^* := \{(y, x) \mid (x, y) \in S\}$. Thus the measure which vanishes on sets of dimension $d - 1$ may be taken to be μ without loss. Now suppose that both γ, $\gamma' \in \Gamma(\mu, \nu)$ have cyclically monotone support. Rockafellar's theorem and Proposition 10 guarantee convex functions ψ and ϕ on $\mathbf{R}^d$ for which $\gamma = (\mathrm{id} \times \nabla\psi)_{\#}\mu$ and $\gamma' = (\mathrm{id} \times \nabla\phi)_{\#}\mu$. Since both $\nabla\psi$ and $\nabla\phi$ push μ forward to ν, the uniqueness part of the Main Theorem shows that they must agree μ-almost everywhere, whence $\gamma = \gamma'$. □

Remark 15 (Nonuniqueness of monotone correlations). In dimension $d \geqslant 2$, the joint distribution $\gamma \in \Gamma(\mu, \nu)$ with cyclically monotone support need not be unique when neither μ nor ν vanishes on sets of codimension 1. For example, let $d = 2$ and suppose that $\operatorname{spt} \mu \subset [-1, 1] \times \{0\}$ while $\operatorname{spt} \nu \subset \{0\} \times [-1, 1]$. Since it has the correct marginals, *any* $\gamma \in \Gamma(\mu, \nu)$ will be supported on the subgradient of the convex function $\psi(x_1, x_2) = |x_2|$.

Remark 16 (The Legendre-Fenchel transform). A parting remark exposes the role of the Legendre-Fenchel transform of a convex function ψ on $\mathbf{R}^d$. Applied to the closure of the cyclically monotone set $\partial\psi^* := \{(y, x) \mid (x, y) \in \partial\psi\}$, Rockafellar's theorem provides a convex function ψ^* on $\mathbf{R}^d$ which is dual to ψ. When μ, $\nu \in \mathscr{P}(\mathbf{R}^d)$ and $\nabla\psi$ pushes μ forward to ν, then $(\nabla\psi \times \mathrm{id})_{\#}\mu$ is supported on the subgradient of ψ^*. If ν vanishes on sets of codimension 1, then Proposition 10 shows that $\nabla\psi^*$ pushes ν forward to μ. Known as the *Legendre-Fenchel transform* of ψ, the function ψ^* is more commonly defined by $\psi^*(y) := \sup_x \langle y, x\rangle - \psi(x)$.

Appendix

A nonsmooth implicit function theorem. To prove the uniqueness result of the previous section, it is necessary to know that when two convex functions ψ and ϕ coincide at a point $p \in \mathbf{R}^d$ where both are differentiable but $\nabla\psi(p) \neq \nabla\phi(p)$, then, locally at least, $\psi = \phi$ occurs only on a set of Hausdorff dimension $d - 1$. One would like to conclude this from the implicit function theorem, but since ψ and ϕ need not be continuously differentiable in any neighbourhood of p, the usual hypotheses of that theorem are not satisfied. This appendix is devoted to establishing a version of the theorem which applies to $\psi - \phi$. A counterexample shows that the theorem is *false* when ψ and ϕ are both Lipschitz, but fail to be convex.

THEOREM 17 (An implicit function theorem). *Let ψ and ϕ be convex functions on $\mathbf{R}^d$, differentiable at p with $\psi(p) = \phi(p)$ but $\nabla\psi(p) \neq \nabla\phi(p)$. Take $\nabla\psi(p) - \nabla\phi(p)$ to be directed along the x_1-axis (without losing generality). Then there is a function $f: \mathbf{R}^{d-1} \to \mathbf{R}$ which is Lipschitz with constant 1, and a neighbourhood U of p on which $\psi(x) = \phi(x)$ if and only if $x_1 = f(x_2, \ldots, x_d)$.*

As might be expected, the proof exploits the continuity of subdifferentials in lieu of continuous differentiability: for x near p, taking $y \in \partial\psi(x)$ and $v \in \partial\phi(x)$ forces $\langle y - v, \hat{x}_1\rangle > 0$; here $\hat{x}_1$ denotes a unit vector in the x_1-direction. As a result, $\psi - \phi$ will be increasing along lines parallel to the x_1-axis in U, implying the existence of f. Failure of the Lipschitz estimate between two points in $\mathbf{R}^{d-1}$, together with a mean-value theorem for $\psi - \phi$, forces subgradients y and v for which $y - v$ is directed far away from $\hat{x}_1$; again, this cannot happen when x is close to p.

Remark 18. An application of the theorem to $[\psi - \phi](\lambda x_1, x_2, \ldots, x_d)$ instead of $\psi - \phi$ shows that in a small enough neighbourhood of p, the Lipschitz constant of f will be arbitrarily small. As a consequence, f is differentiable at $(p_2, \ldots, p_d)$ and its gradient vanishes there.

Before proving this implicit function theorem, the desired result is extracted as a corollary.

COROLLARY 19. *Let ψ and ϕ be convex functions on $\mathbf{R}^d$, differentiable at p with $\psi(p) = \phi(p)$, but $\nabla\psi(p) \neq \nabla\phi(p)$. On a small neighbourhood U of p, the $d - 1$ dimensional Hausdorff measure of $\{x \in U \mid \psi(x) = \phi(x)\}$ is finite.*

Proof. An application of the theorem yields the neighbourhood U, which may be taken to be bounded, and a Lipschitz function f. A standard estimate of geometric measure theory bounds the $d - 1$ dimensional Hausdorff measure $\mathscr{H}^{d-1}$ of the image $g(M)$ of a set under a Lipschitz mapping g by

$$\mathscr{H}^{d-1}[g(M)] \leqslant k^{d-1}\mathscr{H}^{d-1}[M],$$

where $g: \mathbf{R}^n \to \mathbf{R}^m$ satisfies $|g(w) - g(z)| \leqslant k|w - z|$ and $M \subset \mathbf{R}^n$. Taking $g(w) := (f(w), w)$ on $\mathbf{R}^{d-1}$ and the bounded set $M := \{w | g(w) \in U\}$, this estimate yields the desired result: $\mathscr{H}^{d-1}[g(M)] < \infty$. □

The proof of Theorem 17 requires a pair of standard lemmas, suitably adapted to the nonsmooth case. The left and right derivatives of a function f defined in a neighbourhood of $t \in \mathbf{R}$ will be denoted by

$$f'_{\pm}(t) := \lim_{h \to 0\pm} \frac{f(t+h) - f(t)}{h}$$

(when they exist). When ψ is convex on $\mathbf{R}^d$ and finite near p, the right derivative of $f(t) := \psi(p + tx)$ is defined at $t = 0$ for $x \in \mathbf{R}^d$ and given [14, §23.4] by

$$f'_{+}(0) = \sup_{y \in \partial\psi(p)} \langle y, x \rangle; \tag{10}$$

the left derivative is given by an infimum instead of the supremum.

LEMMA 20 (Extremal conditions). *Let f be a function defined in a neighbourhood of $t \in \mathbf{R}$ and assuming a local minimum at t. If f has directional derivatives at t, then $f'_{-}(t) \leqslant 0 \leqslant f'_{+}(t)$.*

Proof. For small h, $f(t + h) - f(x) \geqslant 0$ so $f'_{\pm}(t)$ takes the sign of h. □

LEMMA 21 (A mean value theorem). *Let ψ and ϕ be convex functions on $\mathbf{R}^d$, finite in a neighbourhood of p, $q \in \mathbf{R}^d$ with $\psi(p) - \phi(p) = \psi(q) - \phi(q)$. For some $x = (1 - t)p + tq$ with $0 < t < 1$, there exists $y \in \partial\psi(x)$ and $v \in \partial\phi(x)$ such that $\langle y - v, q - p \rangle = 0$.*

Proof. By convexity, ψ and ϕ are finite in a neighbourhood of the segment joining p to q, thus the function $f(t) := [\psi - \phi]((1 - t)p + tq)$ is continuous on $[0, 1]$. Since $f(0) = f(1)$, for some $t \in (0, 1)$ a maximum or a minimum must be assumed; fix this t and set $x = (1 - t)p + tq$. Now ψ is finite near x so $\partial\psi(x)$ is compact [14, §23.2–4]. The maximum of $\langle y, q - p \rangle$ on $\partial\psi(x)$ is therefore attained at some $y_+ \in \partial\psi(x)$, and the minimum at $y_- \in \partial\psi(x)$. Similarly, $\langle y, q - p \rangle$ attains its extrema on $\partial\phi(x)$ at $y = v_{\pm}$. Thus $f'_{\pm}(t) = \langle y_{\pm} - v_{\pm}, q - p \rangle$ by (10). Lemma 20 shows that $\langle y_- - v_-, q - p \rangle \leqslant 0 \leqslant \langle y_+ - v_+, q - p \rangle$ if f is minimized at t; the inequalities will be reversed if f is maximized. Choose $\lambda \in [0, 1]$ so that $\langle (1 - \lambda)(y_+ - v_+) + \lambda(y_- - v_-), q - p \rangle = 0$, and set $y := (1 - \lambda)y_+ + \lambda y_-$ and $v := (1 - \lambda)v_+ + \lambda v_-$. Noting convexity of the sets $\partial\psi(x)$ and $\partial\phi(x)$, the lemma is proved. □

Proof of Theorem 17. Take $p = 0$ without loss. The convexity of ψ implies $y_n \to \nabla\psi(p)$ when $(x_n, y_n) \in \partial\psi$ and $x_n \to p$; the same is true for ϕ. Since $\nabla\psi(p) - \nabla\phi(p) = \lambda\hat{x}_1$ for $\lambda > 0$, there is a neighbourhood U of p such that

$$y_1 - v_1 > |\pi(y) - \pi(v)| \tag{11}$$

whenever $y \in \partial\psi(x)$, $v \in \partial\phi(x)$ and $x \in U$; here $\pi(y) := (y_2, \ldots, y_d)$, and U may be taken convex with ψ and ϕ continuous on U. In particular, $y_1 - v_1 > 0$, so Lemma 21 shows that $h := \psi - \phi$ must be strictly monotone along lines parallel to the x_1-axis in U. Taking $p_\pm \in U$ to lie on this axis on opposite sides of $p = 0$, continuity of h ensures that $h(x) > h(p) = 0$ whenever $|x - p_+| < r$ for r sufficiently small; the opposite inequality holds when x lies in the ball $B_r(p_-)$. Taking r smaller if necessary forces both $B_r(p_\pm) \subset U$, after which U may be replaced by its intersection with the cylinder $|\pi(x)| < r$ around the x_1-axis. Fix $w \in B_r(0) \subset \mathbf{R}^{d-1}$. Then h takes both positive and negative values on the line $\pi(x) = w$ inside U; since h is also continuous and strictly monotone, it vanishes uniquely at some x on this line. Define $f(w) := x_1$.

It remains to show that $f(w)$ satisfies the Lipschitz bound in $B_r(0)$; no assertion is made concerning f outside this ball (except that it be Lipschitz), and f can certainly be extended to $\mathbf{R}^{d-1}$ (e.g., by inversion) without changing its Lipschitz constant. Therefore, take $w, z \in B_r(0) \subset \mathbf{R}^{d-1}$. Since h vanishes at both $(f(w), w)$ and $(f(z), z)$, Lemma 21 guarantees some $x \in \mathbf{R}^d$ on the line segment joining them, along with $y \in \partial\psi(x)$ and $v \in \partial\phi(x)$, such that

$$\begin{aligned} |y_1 - v_1||f(w) - f(z)| &= |\langle \pi(y) - \pi(v), z - w\rangle| \\ &\leqslant |\pi(y) - \pi(v)||z - w|. \end{aligned}$$

Since U was convex, $x \in U$ and (11) implies the Lipschitz bound. □

Remark 22 (A Lipschitz counterexample). Both the theorem and its corollary fail if ϕ and ψ are not convex, but merely Lipschitz. A counterexample may be constructed by taking $\phi = 0$ and obtaining ψ as the Lipschitz extension of

$$\psi(x_1, x_2) := \begin{cases} x_1 & \text{where } |x_1| > x_2^2 \\ 0 & \text{where } |x_1| < x_2^2/2 \end{cases}$$

from a neighbourhood of the origin to the plane. In this case, $\nabla\psi(0) = (1, 0)$, but $\{\psi = 0\}$ intersects each neighbourhood of the origin in a set of positive area.

References

[1] Y. Brenier, *Décomposition polaire et réarrangement monotone des champs de vecteurs,* C. R. Acad. Sci. Paris Sér. I Math. **305** (1987), 805–808.

[2] ———, *Polar factorization and monotone rearrangement of vector-valued functions,* Comm. Pure Appl. Math. **44** (1991), 375–417.

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

[4] R. D. Anderson and V. L. Klee, Jr., *Convex functions and upper semicontinuous collections,* Duke Math. J. **19** (1952), 349–357.

[5] A. D. Aleksandrov, *Existence and uniqueness of a convex surface with a given integral curvature,* C. R. (Doklady) Acad. Sci. URSS (N.S.) **35** (1942), 131–134.

[6] D. C. Dowson and B. V. Landau, *The Fréchet distance between multivariate normal distributions*, J. Multivariate Anal. **12** (1982), 450–455.
[7] M. Knott and C. S. Smith, *On the optimal mapping of distributions*, J. Optim. Theory Appl. **43** (1984), 39–49.
[8] L. Caffarelli, *Boundary regularity of maps with convex potentials*, Comm. Pure Appl. Math. **45** (1992), 1141–1151.
[9] W. Gangbo, *An elementary proof of the polar factorization of vector-valued functions*, Arch. Rational Mech. Anal. **128** (1994), 381–399.
[10] L. Caffarelli, *The regularity of mappings with a convex potential*, J. Amer. Math. Soc. **5** (1992), 99–104.
[11] ———, *Boundary regularity of maps with convex potentials—2*, preprint.
[12] R. J. McCann, *A Convexity Theory for Interacting Gases and Equilibrium Crystals*, PhD thesis, Princeton University, 1994.
[13] ———, *A convexity principle for attracting gases*, to appear in Adv. Math.
[14] R. T. Rockafellar, *Convex Analysis*, Princeton University Press, Princeton, 1972.
[15] S. T. Rachev, *The Monge-Kantorovich mass transference problem and its stochastic applications*, Theory Probab. Appl. **29** (1984), 647–676.
[16] C. Smith and M. Knott, *On Hoeffding-Fréchet bounds and cyclic monotone relations*, J. Multivariate Anal. **40** (1992), 328–334.
[17] L. Rüschendorf, "Fréchet bounds and their applications" in *Advances in Probability Distributions with Given Marginals (Rome 1990)*, ed. by G. Dall'Aglio, S. Kotz, and G. Salinetti, Math. Appl. **67**, Kluwer Academic Publishers, Dordrecht, 1991, 151–187.
[18] T. Abdellaoui and H. Heinich, *Sur la distance de deux lois dans le cas vectoriel*, C. R. Acad. Sci. Paris Sér I Math. **319**, (1994), 397–400.
[19] L. Rüschendorf and S. T. Rachev, *A characterization of random variables with minimum L^2-distance*, J. Multivariate Anal. **32** (1990), 48–54.
[20] J. A. Cuesta-Albertos, C. Matrán, and A. Tuero-Díaz, *Optimal maps for the L^2-Wasserstein distance*, preprint.

Department of Mathematics, Brown University, Providence, Rhode Island 02912, USA; mccann@math.brown.edu