

1

An Introduction to the Monge Problem

We begin this chapter by introducing two formulations of the Monge problem. The first one is closer to the original one proposed by Monge himself, and is based on a hardly manageable pointwise transport condition (Section 1.1). The second one exploits concepts of modern measure theory to formulate a more flexible transport condition and is the one adopted in the rest of the book (Section 1.2). We then close the chapter by building some intuition on the notions of the transport map and of the *optimal* transport map. This is achieved by looking into a simple “duality-based optimality criterion” (Section 1.3) and by exploiting monotonicity in the construction of transport maps, first in dimension 1 (Section 1.4) and then in higher dimensions (Section 1.5).

1.1 The Original Monge Problem

A modern-language proxy for Monge’s original formulation of his eponymous transport problem can be introduced as follows. Given two smooth, nonnegative functions $\rho, \sigma : \mathbb{R}^n \rightarrow [0, \infty)$ with the dimensions of mass per unit volume, and assuming that the mass distributions $\rho(x) dx$ and $\sigma(y) dy$ have the same (unit) total mass, i.e.,

$$\int_{\mathbb{R}^n} \rho(x) dx = \int_{\mathbb{R}^n} \sigma(y) dy = 1,$$

we consider smooth, injective maps $T : \mathbb{R}^n \rightarrow \mathbb{R}^n$ that *transport* $\rho(x) dx$ to $\sigma(y) dy$, in the sense that the total infinitesimal volume of the origin mass distribution $\rho(x) dx$ at x is required to be equal to the total infinitesimal volume of the final mass distribution $\sigma(y) dy$ at $y = T(x)$. Since, by the change of variables formulae for smooth injective maps, we have $dy|_{y=T(x)} = |\det \nabla T(x)| dx$, the transport constraint takes the form

$$|\det \nabla T(x)| \sigma(T(x)) = \rho(x), \quad \forall x \in \{\rho > 0\}. \quad (1.1)$$

We call (1.1) the **pointwise transport condition** from $\rho(x) dx$ to $\sigma(y) dy$. Taking $|y - x|$ as the transport cost¹ to move a unit mass from x to y , the **original Monge problem** (from $\rho(x) dx$ to $\sigma(y) dy$) is the minimization problem

$$M = \inf \left\{ \int_{\mathbb{R}^n} |T(x) - x| \rho(x) dx : T \text{ is smooth, injective and (1.1) holds} \right\}. \quad (1.2)$$

Since work has the dimensions of force times length, if λ denotes the amount of force per unit mass at our disposal to implement the “instructions” of transport maps, then λM is the **minimal amount of work** needed to transport $\rho(x) dx$ into $\sigma(y) dy$. Since work is a form of (mechanical) energy, (1.2) is, in precise physical terms, an “energy minimization problem.”

From a mathematical viewpoint – even from a modern mathematical viewpoint that takes advantage of all sorts of compactness and closure theorems discovered since Monge’s time – (1.2) is a *very* challenging minimization problem. Let us consider, for example, the problem of showing the mere existence of a minimizer. The baseline, modern strategy to approach this kind of question, the so-called *Direct Method of the Calculus of Variations*, works as follows. Consider an abstract minimization problem, $m = \inf\{f(x) : x \in X\}$, defined by a function $f : X \rightarrow \mathbb{R}$ such that $m \in \mathbb{R}$. By definition of infimum of a set of real numbers, we can consider a *minimizing sequence*² for m , that is, a sequence $\{x_j\}_j$ in X such that $f(x_j) \rightarrow m$ as $j \rightarrow \infty$. Assuming that: (i) there is a notion of convergence in X such that “ $\{f(x_j)\}_j$ bounded in \mathbb{R} implies, up to subsequences, that $x_j \rightarrow x \in X$,” and (ii) “ $f(x) \leq \liminf_j f(x_j)$ whenever $x_j \rightarrow x$,” we conclude that any subsequential limit x of $\{x_j\}_j$ is a minimizer of m , since, using in the order, $x \in X$, properties (i) and (ii), and the minimizing sequence property, we find

$$m \leq f(x) \leq \liminf_j f(x_j) = m.$$

With this method in mind, and back to the original Monge problem (1.2), we assume that M is finite (i.e., we assume the existence of at least one transport map with finite transport cost) and consider a minimizing sequence $\{T_j\}_j$ for (1.9). Thus, $\{T_j\}_j$ is a sequence of smooth and injective maps with

$$\sigma(T_j(x)) |\det \nabla T_j(x)| = \rho(x), \quad (1.3)$$

for all $x \in \{\rho > 0\}$ and $j \in \mathbb{N}$, and such that

$$\lim_{j \rightarrow \infty} \int_{\mathbb{R}^n} |T_j - x| \rho = M < \infty. \quad (1.4)$$

¹ The transport cost $|x - y|$ is commonly named the “linear cost,” although evidently $(x, y) \mapsto |x - y|$ is not linear.

² Notice that a subsequence of a minimizing sequence is still a minimizing sequence.

Trying to check assumption (i) of the Direct Method, we ask if (1.4) implies the compactness of $\{T_j\}_j$, say, in the sense of pointwise (a.e.) convergence. Compactness criteria enforcing this kind of convergence, like the Ascoli–Arzelà criterion, or the compactness theorem of Sobolev spaces, would require some form of uniform control on the gradients (or on some sort of incremental ratio) of the maps T_j . It is, however, clear that no control of that sort is contained in (1.4). It is natural to think about pointwise convergence here, because should the maps T_j converge pointwise to some limit T , then by Fatou's lemma, we would find

$$\int_{\mathbb{R}^n} |T - x| \rho \leq \lim_{j \rightarrow \infty} \int_{\mathbb{R}^n} |T_j - x| \rho = M,$$

thus verifying assumption (ii) of the Direct Method. Finally, even if pointwise convergence could somehow be obtained, we would still face the issue of showing that the limit map T belongs to the competition class (i.e., T is smooth and injective, and it satisfies the transport constraint (1.1)) in order to infer $\int_{\mathbb{R}^n} |T - x| \rho \geq M$ and close the Direct Method argument. Deducing all these properties on T definitely requires some form of convergence of ∇T_j toward ∇T (as is evident from the problem of passing to the limit the nonlinear constraint (1.3)) – a task that is even more out of reach than proving the pointwise convergence of T_j in the first place! Thus, even from a modern perspective, establishing the mere existence of minimizers in the original Monge problem is a formidable task.

1.2 A Modern Formulation of the Monge Problem

We now introduce the modern formulation of the Monge problem that will be used in the rest of this book. The first difference with respect to Monge's original formulation is that we extend the class of distributions of mass to be transported to the whole family $\mathcal{P}(\mathbb{R}^n)$ of probability measures on \mathbb{R}^n . We usually denote by μ the origin distribution of mass, and by ν the final one, thus going back to Monge's original formulation by setting $\mu = \rho d\mathcal{L}^n$ and $\nu = \sigma d\mathcal{L}^n$, where \mathcal{L}^n is the Lebesgue measure on \mathbb{R}^n . This first change demands a second one, namely, we need to reformulate the pointwise transport condition (1.1) in a way that makes sense even when μ and ν are not absolutely continuous with respect to \mathcal{L}^n . This is done by resorting to the notion of push-forward (or direct image) of a measure through a map, which we now recall (see also Appendix A.4).

We say that T **transports** μ if there exists a Borel set $F \subset \mathbb{R}^n$ such that

$$\begin{aligned} T : F &\rightarrow \mathbb{R}^n \text{ is a Borel map,} \\ \text{and } \mu &\text{ is concentrated on } F \text{ (i.e., } \mu(\mathbb{R}^n \setminus F) = 0). \end{aligned} \tag{1.5}$$

Whenever $T : F \rightarrow \mathbb{R}^n$ transports μ , we can define a Borel measure $T_{\#}\mu$ (the push-forward of μ through T) by setting, for every Borel set $E \subset \mathbb{R}^n$,

$$(T_{\#}\mu)(E) = \mu(T^{-1}(E)), \quad (1.6)$$

where $T^{-1}(E) = \{x \in F : T(x) \in E\}$.

Notice that, according to this definition $(T_{\#}\mu)(\mathbb{R}^n) = \mu(F)$; therefore, the requirement that μ is concentrated on F in (1.5) is necessary to ensure that $T_{\#}\mu \in \mathcal{P}(\mathbb{R}^n)$ if $\mu \in \mathcal{P}(\mathbb{R}^n)$. Finally, we say that T is a **transport map from μ to ν** if

$$T_{\#}\mu = \nu.$$

Clearly, the transport condition (1.6) does not require T to be differentiable, nor injective; moreover, it boils down to the pointwise transport condition (1.1) whenever the latter makes sense, as illustrated in the following proposition.

Proposition 1.1 *Let $\mu = \rho d\mathcal{L}^n$ and $\nu = \sigma d\mathcal{L}^n$ belong to $\mathcal{P}(\mathbb{R}^n)$, μ be concentrated on a Borel set F , and $T : F \rightarrow \mathbb{R}^n$ be an injective Lipschitz map. Then, $T_{\#}\mu = \nu$ if and only if*

$$\sigma(T(x)) |\det \nabla T(x)| = \rho(x), \quad \text{for } \mathcal{L}^n\text{-a.e. } x \in F. \quad (1.7)$$

Proof By the injectivity and Lipschitz continuity of T , the area formula,

$$\int_{T(F)} \varphi(y) \sigma(y) dy = \int_F \varphi(T(x)) |\det \nabla T(x)| \sigma(T(x)) dx, \quad (1.8)$$

holds for every Borel function $\varphi : T(F) \rightarrow [0, \infty]$ (see Appendix A.10). Since T is injective, for every Borel set $G \subset F$, we have $G = T^{-1}(T(G))$. Therefore, by definition of $T_{\#}\mu$ and by (1.8) with $\varphi = 1_{T(G)}$, we find

$$\begin{aligned} (T_{\#}\mu)(T(G)) &= \mu(T^{-1}(T(G))) = \int_G \rho, \\ \nu(T(G)) &= \int_{T(G)} \sigma(y) dy = \int_G |\det \nabla T(x)| \sigma(T(x)) dx. \end{aligned}$$

By arbitrariness of $G \subset F$, we find that $T_{\#}\mu = \nu$ if and only if (1.7) holds. \square

Based on these considerations, given $\mu, \nu \in \mathcal{P}(\mathbb{R}^n)$, we formally introduce the **Monge problem** from μ to ν by letting

$$\mathbf{M}_1(\mu, \nu) = \inf \left\{ \int_{\mathbb{R}^n} |T(x) - x| d\mu(x) : T_{\#}\mu = \nu \right\}. \quad (1.9)$$

Problem (1.9) is, in principle, more tractable than (1.2). Transport maps are no longer required to be smooth and injective, as reflected in the new transport condition (1.6). It is still unclear, however, if, given two arbitrary $\mu, \nu \in \mathcal{P}(\mathbb{R}^n)$,

there always exists at least one transport map from μ to ν , and if such transport map can be found with finite transport cost; whenever this is not the case, we have $\mathbf{M}_1(\mu, \nu) = +\infty$, and the Monge problem is ill posed. A more fundamental issue is that, even in a situation where we know from the onset that $\mathbf{M}_1(\mu, \nu) < \infty$, it is still very much unclear how to verify assumption (i) in the Direct Method: what notion of subsequential convergence (for minimizing sequences $\{T_j\}_j$) is needed for passing the transport condition $(T_j)_\# \mu = \nu$ to a limit map T ? This difficulty will eventually be solved by working with the Kantorovich formulation of the transport condition, which requires extending competition classes for transport problems from the family of transport *maps* to that of transport *plans*. From this viewpoint, the modern formulation of the Monge problem, as much as the original one, is still somehow untractable by a direct approach.

We can, of course, formulate the Monge problem with respect to a general³ **transport cost** $c : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$. Interpreting $c(x, y)$ as the cost needed to transport a unit mass from x to y (notice that c does not need to be symmetric in (x, y) !), we define the **Monge problem with transport cost** c by setting

$$\mathbf{M}_c(\mu, \nu) = \inf \left\{ \int_{\mathbb{R}^n} c(x, T(x)) d\mu(x) : T_\# \mu = \nu \right\}. \quad (1.10)$$

Our focus will be largely (but not completely) specific to the cases of the linear cost $c(x, y) = |x - y|$ and of the quadratic cost $c(x, y) = |x - y|^2$. In the following, when talking about “the Monge problem,” we shall either assume that the transport cost under consideration is evident from the context or otherwise add the specification “with general cost,” “with linear cost,” or “with quadratic cost.” From the historical viewpoint, of course, only the Monge problem with linear cost should be called “the Monge problem.”

1.3 Optimality via Duality and Transport Rays

We now anticipate an observation that we will formally reintroduce later on⁴ in our study of the Kantorovich duality theory and that provides a simple and effective criterion to check the optimality of a transport map in the Monge problem. The remark is that, if $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is a Lipschitz function with $\text{Lip}(f) \leq 1$ (briefly, a 1-Lipschitz function), and if T is a transport map from μ to ν , then

$$\int_{\mathbb{R}^n} f d\nu - \int_{\mathbb{R}^n} f d\mu = \int_{\mathbb{R}^n} [f(T(x)) - f(x)] d\mu(x) \leq \int_{\mathbb{R}^n} |T(x) - x| d\mu(x),$$

³ In practice, we shall work with transport costs that are at least lower semicontinuous, thus guaranteeing the Borel measurability of $x \mapsto c(x, T(x))$.

⁴ See Section 3.7.

so that one always has

$$\sup_{\text{Lip}(f) \leq 1} \int_{\mathbb{R}^n} f \, d\nu - \int_{\mathbb{R}^n} f \, d\mu \leq \inf_{T_{\#}\mu = \nu} \int_{\mathbb{R}^n} |T(x) - x| \, d\mu(x). \quad (1.11)$$

In particular, if for a given transport map T from μ and ν , we can find a 1-Lipschitz function f such that

$$f(T(x)) - f(x) = |T(x) - x|, \quad \text{for } \mu\text{-a.e. } x \in \mathbb{R}^n, \quad (1.12)$$

then, integrating (1.12) with respect to $d\mu$ and exploiting $T_{\#}\mu = \nu$, we find

$$\int_{\mathbb{R}^n} f \, d\nu - \int_{\mathbb{R}^n} f \, d\mu = \int_{\mathbb{R}^n} |T(x) - x| \, d\mu(x),$$

thus deducing from (1.11) that T is a minimizer in the Monge problem $\mathbf{M}_1(\mu, \nu)$ (and, symmetrically, that f is a maximizer in the “dual” maximization problem appearing on the left-hand side of (1.11)). We illustrate this idea with the so-called “book-shifting example.” Given $N \geq 2$, let us consider the Monge problem from μ to ν with

$$\mu = \frac{1_{[0, N]}}{N} \, d\mathcal{L}^1, \quad \nu = \frac{1_{[1, N+1]}}{N} \, d\mathcal{L}^1.$$

We can think of μ as a collection of N books of mass $1/N$ that we want to shift to the right (not necessarily in their original order) by a unit length. The map $T(t) = t + 1$ for $t \in \mathbb{R}$ (corresponding to shifting each book to the right by a unit length) is a minimizer in $\mathbf{M}_1(\mu, \nu)$ since it satisfies (1.12) with $f(t) = t$. By computing the transport cost of T , we see that $\mathbf{M}_1(\mu, \nu) = 1$. We easily check that transport map S defined by $S(t) = t$, for $t \in [1, N]$, and $S(t) = t + N$, for $t \in [0, 1)$, which corresponds to moving *only* the left-most book to the right by a length equal to N , has also a unit transport cost and thus is also optimal in $\mathbf{M}_1(\mu, \nu)$. This shows, in particular, that the Monge problem can admit multiple minimizers.

It is interesting to notice that the connection between optimal transport maps and 1-Lipschitz “potential functions” expressed in (1.12) was also clear to Monge, who rather focused on the more expressive identity

$$\nabla f(x) = \frac{T(x) - x}{|T(x) - x|}. \quad (1.13)$$

The relation between (1.13) and (1.12) is clarified by noticing that $\text{Lip}(f) \leq 1$ and (1.12) imply that f is affine with unit slope along the oriented segment from x to $T(x)$, that is,

$$f(x + t(T(x) - x)) = f(x) + t|T(x) - x|, \quad \forall t \in [0, 1], \quad (1.14)$$

from which (1.13) follows⁵ if f is differentiable at x . Such oriented segments are called *transport rays*, and their study plays a central role in the solution to the Monge problem (presented in Part IV). Notice that, by (1.14), the graph of f above the union of such segments is a *developable surface*; this connection seems to be the reason why Monge started (independently from Euler) the systematic study of developable surfaces.

1.4 Monotone Transport Maps

In dimension $n = 1$, it is particularly easy to construct *optimal* transport maps by looking at *monotone* transport maps. Here we just informally discuss this important idea, which will be addressed rigorously in Chapter 16.

It is quite intuitive that, in dimension 1, moving mass by monotone increasing maps must be a good transport strategy (the book-shifting example from Section 1.3 confirming that intuition). Considering the case when $\mu = \rho d\mathcal{L}^1$ and $\nu = \sigma d\mathcal{L}^1$, for an increasing map to be a transport map, we only need to check that the “rate of mass transfer,” i.e., the derivative of the transport map, is compatible with the transport condition (1.1), and this can be achieved quite easily by defining $T(x)$ through the formula

$$\int_{-\infty}^x \rho = \int_{-\infty}^{T(x)} \sigma, \quad \forall x \in \mathbb{R}. \quad (1.15)$$

In more geometric terms, we are prescribing that the mass stored by μ to the left of x corresponds to the mass stored by ν to the left of $T(x)$, i.e., we are setting $\mu((-\infty, x)) = \nu((-\infty, T(x)))$; see Figure 1.1. Indeed, an informal differentiation in x of (1.15) gives

$$\rho(x) = T'(x) \sigma(T(x)),$$

which (thanks to $T' \geq 0$) is the pointwise transport condition (1.1). The map T defined in (1.15) is called the **monotone rearrangement** of μ into ν and provides a minimizer in the Monge problem. We present here a simple argument in support of this assertion, which works under the assumption that $\{T \geq \text{id}\} = \{x : T(x) \geq x\}$ and $\{T < \text{id}\} = \{x : T(x) < x\}$ are equal, respectively, to complementary half-lines $[a, \infty)$ and $(-\infty, a)$ for some $a \in \mathbb{R}$: indeed, in this case, we can consider a 1-Lipschitz function $f : \mathbb{R} \rightarrow \mathbb{R}$ with

$$f'(x) = 1_{\{T \geq \text{id}\}}(x) - 1_{\{T < \text{id}\}}(x)$$

and see that if $x \in \{T \geq \text{id}\}$, and thus $(x, T(x)) \subset \{T \geq \text{id}\} = \{f' = 1\}$, we have

⁵ See Proposition 18.6 for a formal discussion.




Figure 1.1 The amount of mass stored by $\mu = \rho d\mathcal{L}^1$ to the left of x corresponds to the amount of mass stored by $\nu = \sigma d\mathcal{L}^1$ to the left of $T(x)$; see (1.15).

$$f(T(x)) - f(x) = \int_x^{T(x)} f' = T(x) - x = |T(x) - x|;$$

while if $x \in \{T < \text{id}\}$, and thus $(T(x), x) \subset \{T < \text{id}\} = \{f' = -1\}$, we have

$$f(T(x)) - f(x) = \int_x^{T(x)} f' = x - T(x) = |T(x) - x|.$$

Hence (1.12) holds, and T is optimal in the Monge problem.

1.5 Knothe Maps

Monotone rearrangements can be used to define transport maps in higher dimensions. Given $\mu, \nu \in \mathcal{P}(\mathbb{R}^n)$, with $\mu, \nu \ll \mathcal{L}^n$, and an orthonormal basis $\tau = \{\tau_i\}_{i=1}^n$ of \mathbb{R}^n (with coordinates $x^i = x \cdot \tau_i$), one can define **the Knothe map T from μ to ν (relative to the orthonormal basis τ)** by the following procedure (which, for the sake of simplicity, is discussed only informally here and in the case $n = 2$; an expert reader should have little difficulty in formalizing and extending to higher dimensions the following sketch). Writing $\mu = \rho d\mathcal{L}^2$ and $\nu = \sigma d\mathcal{L}^2$, we define the first component of T by a monotone rearrangement depending on the coordinate x^1 only, i.e., we set $T^1(x) = T^1(x^1)$ with

$$\int_{-\infty}^{x^1} ds \int_{\mathbb{R}} \rho(s, t) dt = \int_{-\infty}^{T^1(x^1)} ds \int_{\mathbb{R}} \sigma(s, t) dt, \quad \forall x \in \mathbb{R}^2. \quad (1.16)$$

In this way, the total mass stored by μ inside the half-plane $\{z^1 < x^1\} = \{z \in \mathbb{R}^2 : z^1 < x^1\}$ is set to be equal to the total mass stored by ν inside the half-plane $\{z^1 < T^1(x^1)\}$. This choice of T^1 implies that points in the vertical line $\{z^1 = x^1\}$ are to be mapped by T inside the vertical line $\{z^1 = T^1(x^1)\}$. Thus, it is just natural to do this by a monotone rearrangement of $\rho(x^1, t) dt$ into $\sigma(T^1(x^1), t) dt$. Since these two measures have not the same total mass, we first normalize them into probability measures, and then we define $T^2(x) = T^2(x^1, x^2)$ by setting

$$\frac{\int_{-\infty}^{x^2} \rho(x^1, t) dt}{\int_{\mathbb{R}} \rho(x^1, t) dt} = \frac{\int_{-\infty}^{T^2(x^1, x^2)} \sigma(T^1(x^1), t) dt}{\int_{\mathbb{R}} \sigma(T^1(x^1), t) dt} \quad \forall x \in \mathbb{R}^2. \quad (1.17)$$

Informal differentiations of (1.16) in x^1 and (1.17) in x^2 give

$$\begin{aligned} \int_{\mathbb{R}} \rho(x^1, t) dt &= \frac{\partial T^1}{\partial x^1}(x) \int_{\mathbb{R}} \sigma(T^1(x^1), t) dt \\ \frac{\rho(x)}{\int_{\mathbb{R}} \rho(x^1, t) dt} &= \frac{\partial T^2}{\partial x^2}(x) \frac{\sigma(T(x))}{\int_{\mathbb{R}} \sigma(T^1(x^1), t) dt}, \end{aligned}$$

while, evidently, $\partial T^1 / \partial x^2 = 0$; therefore,

$$\det \nabla T(x) = \frac{\partial T^1}{\partial x^1}(x) \frac{\partial T^2}{\partial x^2}(x) = \frac{\rho(x)}{\sigma(T(x))}.$$

Therefore (1.1) holds (notice that $\det \nabla T(x) \geq 0$), and T transports μ into ν .

The (formal) construction of Knothe maps proves the important point that there are always transport maps between two \mathcal{L}^n -absolutely continuous probability measures. Moreover, because of their componentwise monotonicity, Knothe maps can be used in place of optimal transport maps in certain arguments. For example, the proofs of the sharp Euclidean isoperimetric and Sobolev inequalities presented in Chapter 9 can be rigorously carried over using Knothe maps rather than (as done in that chapter) optimal transport maps in the Monge problem with quadratic transport cost (known, more briefly, as Brenier maps).

This said, when $n \geq 2$, we do not expect Knothe maps to be *optimal* transport maps for the linear and quadratic transport costs, as explained (only informally) in the following remarks.

Remark 1.2 (Knothe maps, in general, fail the noncrossing condition) In general, we do not expect Knothe maps to be optimal in the Monge problem with linear cost. To explain this point, we informally notice that for a transport map T from μ to ν to be optimal in \mathbf{M}_1 , necessary condition is:⁶ for every $x_1, x_2 \in \text{spt } \mu$, it holds

$$|T(x_1) - x_1| + |T(x_2) - x_2| \leq |T(x_1) - x_2| + |T(x_2) - x_1|. \quad (1.18)$$

Indeed, should (1.18) fail at $x_1 \neq x_2$, then one should be able to define a new transport map by sending small neighborhoods of x_1 and x_2 , respectively, to

⁶ In the terminology and notation to be introduced in the next two chapters condition (1.18) can be seen as a particular case of c -cyclical monotonicity condition (with respect to the linear cost $c(x, y) = |x - y|$) applied to $\text{spt } \gamma_T$, $\gamma_T = (\text{id} \times T)_{\#} \mu$; see Remark 3.10.




Figure 1.2 In this example $\mu = 1_E d\mathcal{L}^2$ and $\nu = 1_F d\mathcal{L}^2$, where E and F are two unit area regions, with each obtained as the union of a family of three rectangles, denoted respectively by $\{E_i\}_{i=1}^3$ and $\{F_i\}_{i=1}^3$. There are horizontal vectors v_1 , v_2 , and v_3 such that, if A denotes the diagonal matrix with entries $(1/2, 2)$, then $F_1 = E_1 + v_1$, $F_2 = A[E_2] + v_2$, and $F_3 = E_3 + v_3$. The Knothe map T from μ to ν (in the coordinate system (z^1, z^2)) is such that $T(x) = x + v_1$ if $x \in E_1$, $T(x) = A[x] + v_2$ if $x \in E_2$, and $T(x) = x + v_3$ if $x \in E_3$. (Notice that T is discontinuous on the segments separating E_1 from E_2 and E_2 from E_3 .) In the figure, we have selected points $x_1, x_2 \in E$ such that the corresponding segments $[x_1, T(x_1)]$ and $[x_2, T(x_2)]$ intersect. Correspondingly, condition (1.18) does not hold, and T is not optimal in the Monge problem with linear cost.

small neighborhoods of $T(x_2)$ and $T(x_1)$, thus lowering the total transport cost. This said, it is easily seen that (1.18) can be violated by a Knothe map; see, for example, Figure 1.2.

Remark 1.3 (Knothe maps (in general) are not Brenier maps) In general, we do not expect Knothe maps to be optimal in the Monge problem with quadratic cost. Indeed, looking at Theorem 4.2 in Chapter 4 and keeping in mind the informal character of this remark, this would mean that $\nabla T = \nabla^2 f$ for a convex function $f : \mathbb{R}^n \rightarrow \mathbb{R}$. In particular, ∇T would be symmetric, whereas gradients of Knothe maps are usually represented by triangular matrices with nontrivial off-diagonal entries. Incidentally, this method for excluding optimality in the Monge problem with quadratic cost does not apply to the example of the Knothe map depicted in Figure 1.2, since, in that case, ∇T takes only two symmetric values,

$$A_1 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad A_2 = \begin{pmatrix} 1/2 & 0 \\ 0 & 2 \end{pmatrix}.$$

However, should $\nabla^2 f = A_i$ hold on E_i for $i = 1, 2$, we could then find $a, b \in \mathbb{R}^2$ and $c \in \mathbb{R}$ such that $f(z) = (|z|^2/2) + a \cdot z$ for $z \in E_1$ and $f(z) = (z^1)^2/4 + (z^2)^2 + b \cdot z + c$ for $z \in E_2$. Notice, however, that there is no way to adjust a , b , and c so that f is continuous on the vertical segment at the interface between E_1 and E_2 (f being convex on \mathbb{R}^n , it must be continuous on E); indeed, the $(z^2)^2$ -coefficients of the two polynomials describing f on E_1 and E_2 are different.