Pontifícia Universidade Católica
do Rio de Janeiro

Paulo Najberg Orenstein

Optimal Transport and the Wasserstein Metric

Dissertação de Mestrado

Thesis presented to the Postgraduate Program in Applied Mathematics of the Departamento de Matemática, PUC–Rio as partial fulfillment of the requirements for the degree of Mestre em Matemática Aplicada

Advisor: Prof. Jairo Bochi
Co–Advisor: Prof. Carlos Tomei

Rio de Janeiro
January 2014

Pontifícia Universidade Católica
do Rio de Janeiro

Paulo Najberg Orenstein

Optimal Transport and the Wasserstein Metric

Thesis presented to the Postgraduate Program in Applied Mathematics of the Departamento de Matemática, PUC–Rio as partial fulfillment of the requirements for the degree of Mestre em Matemática Aplicada. Approved by the following commission:

Prof. Jairo Bochi
Advisor
Departamento de Matemática — PUC–Rio

Prof. Carlos Tomei
Co–Advisor
Departamento de Matemática — PUC–Rio

Prof. Nicolau Corção Saldanha
Departamento de Matemática — PUC–Rio

Prof. Milton da Costa Lopes Filho
Instituto de Matemática — UFRJ

Prof. Helena Judith Nussenzveig Lopes
Instituto de Matemática — UFRJ

Rio de Janeiro — January 23, 2014

Paulo Najberg Orenstein

BSc. in Economics from PUC–Rio.

Bibliographic data

Orenstein, Paulo Najberg

Optimal Transport and the Wasserstein Metric / Paulo Najberg Orenstein ; advisor: Jairo Bochi; co–advisor: Carlos Tomei. — 2014.
89 f. : il. ; 30 cm

Dissertação (Mestrado em Matemática Aplicada)-Pontifícia Universidade Católica do Rio de Janeiro, Rio de Janeiro, 2014.
Inclui bibliografia

1. Matemática – Teses. 2. Transporte Ótimo. 3. Problema de Monge–Kantorovich. 4. Mapa de Transporte. 5. Dualidade. 6. Métrica de Wasserstein. 7. Interpolação de Medidas. I. Bochi, Jairo. II. Tomei, Carlos. III. Pontifícia Universidade Católica do Rio de Janeiro. Departamento de Matemática. IV. Título.

CDD: 510

Acknowledgments

To my advisor, Jairo Bochi, for his help, teachings and intense dedication to this project. For showing me what mathematics really is, and for leading me into this journey.

To my co-advisor, Carlos Tomei, for guiding me through the years and making mathematics even more fun, interesting and challenging. For his constant support, energy, honesty and friendship.

To all professors at the Mathematics Department, particularly Nicolau Corção Saldanha and Ricardo de Sá Earp, for their hard work and inspiration.

To the members of staff at the Mathematics Department, particularly Creuza, for their constant help and usual cheerfulness.

To CNPQ, FAPERJ and PUC-Rio, for the financial support.

To my girlfriend, for her love, kindness, and being the best company ever.

To my sister, Leila, for always making me laugh, always picking me up, and always being by my side.

To my parents, Luiz and Sheila, for being my role models. For their strength, love, encouragement, and wisdom throughout the years.

Abstract

Orenstein, Paulo Najberg; Bochi, Jairo; Tomei, Carlos. **Optimal Transport and the Wasserstein Metric**. Rio de Janeiro, 2014. 89p. Dissertação de Mestrado — Departamento de Matemática, Pontifícia Universidade Católica do Rio de Janeiro.

In this work, we study the so called Optimal Transport theory: given two probability measures and a cost function, what is the best way to take one measure to another, minimizing costs? We analyze particular cases and establish necessary and sufficient conditions for optimality. We show that, in general, the transport takes the form of a generalized gradient, and that the problem, which can be viewed as a generalization of Linear Programming to infinite dimensions, admits a rich underlying duality theory.

Moreover, using the aforementioned results, it is possible to define a distance in the space of probability measures, called the Wasserstein metric, as the optimal transport cost between two given measures. This gives rise to a metric space of probability measures, known as the Wasserstein space, which has several interesting properties that are studied in the text, both in terms of its topology (e.g. completeness) and in terms of its geometry (e.g. geodesics and curvature).

Keywords

Optimal Transport. Monge–Kantorovich Problem. Transport Map. Duality. Wasserstein Metric. Measure Interpolation.

Resumo

Orenstein, Paulo Najberg; Bochi, Jairo; Tomei, Carlos. **Transporte Ótimo e a Métrica de Wasserstein**. Rio de Janeiro, 2014. 89p. Dissertação de Mestrado — Departamento de Matemática, Pontifícia Universidade Católica do Rio de Janeiro.

Neste trabalho, estudamos a chamada Teoria de Transporte Ótimo: dadas duas medidas de probabilidade e uma função de custo, qual é a melhor maneira de levar uma medida na outra, minimizando custos? Analisamos alguns casos particulares e estabelecemos critérios de otimalidade para o problema. Mostramos que o transporte toma em geral a forma de um gradiente generalizado e que há uma rica teoria de dualidade subjacente ao problema, que, de fato, pode ser encarado como uma generalização de programação linear para dimensão infinita.

Além disso, através dos resultados obtidos, é possível definir uma distância no espaço de probabilidades, chamada de métrica de Wasserstein, como o custo de transporte ótimo entre duas medidas. Isto permite considerar um espaço métrico de medidas de probabilidade, conhecido como espaço de Wasserstein, que traz consigo diversas propriedades interessantes, tanto de caráter topológico (e.g. completude) quanto de caráter geométrico (e.g. curvatura), que são investigadas no texto.

Palavras–chave

Transporte Ótimo. Problema de Monge–Kantorovich. Mapa de Transporte. Dualidade. Métrica de Wasserstein. Interpolação de Medidas.

Contents

1
Introduction

Though Optimal Transport theory can trace its roots back to the French mathematician Gaspard Monge in 1781, it has recently been the stage of spectacular developments, both in terms of new and deep theorems and also unsuspected applications. The field, once dormant, was first revived by the Russian mathematician and economist Leonid Kantorovich, in 1923, and, arguably, a second time by Yann Brenier, in 1987. In its modern incarnation, Optimal Transport stands at the intersection of measure theory, functional analysis, probability theory and optimization.

The fundamental question is deceivingly simple: given two lumps of masses, what is the best way to transport one of them to the other, minimizing some cost? Of course, the problem can be understood in increasing generality, depending on which spaces the masses are on, what shapes we allow the masses to have, what exactly we mean by 'transporting', and which function we take to be the cost of moving. Though one would have guessed at first that too great a generality would imply in unavailing theorems, we show in these notes that several important results of this theory work for surprisingly general scenarios.

This explains, at least in part, why Optimal Transport has found several applications in recent years. It pervades fields like engineering (where the masses could represent piles of sand and holes), economics (where the masses could be the density of consumers and producers) and even image processing (where the masses could be color histograms). Indeed, this theory has been applied in problems as varied as oceanography [17], the design of antennas [33], matching theory [14], and city planning [13]. It also has profound connections with the theories of partial differential equations [19, 31], gradient flows [4], dynamical systems [5], and stochastic differential equations [23], among others. This, in and on itself, is sufficient to make Optimal Transport an indispensable tool for both the pure and the applied mathematician.

However, even more astounding is the intricate and beautiful mathematics that underlie this field. Indeed, when looked through the right lenses, one can rewrite the fundamental question of transporting masses as a Linear Programming problem in infinite dimensions, and from then try to generalize already established theorems. Surprisingly, several results still hold in this setting (including a duality theory), but many new ones give the theory a new and distinct flavor.

Furthermore, this rich mathematical structure allows one to use the problem of transporting masses to suggest a useful metric in the space of probability measures. It is called the Wasserstein metric, and it lets one look at the space of probability measures (in fact, a slight restriction of it) as if it were a metric space. Though this could have been done in other ways, the particular space that arises, called the Wasserstein space, is full of interesting properties. In these notes, we care about two, in particular. First, we investigate topological properties, such as whether the space is separable or complete, and what kind of topology it possesses. Second, we look at geometric properties, which is possible when one studies the Optimal Transport dynamically, as a sequence of moving measures. It is then reasonable to ask what constitutes the 'shortest path' between two given probability measure, and how 'curved' the Wasserstein space is.

The structure of this work is as follows. In Chapter 2, we study the Optimal Transport problem. Section 2.1 presents an informal introduction to the theory, explaining in mathematical terms what we are trying to answer and establishing a standard notation. It also distinguishes between what is called the Monge Problem, a situation in which we model the solution of the Optimal Transport problem by a function, and the Kantorovich Problem, a relaxation of the Monge Problem where the solution is taken to be a measure. Section 2.2 gives some simple examples to illustrate the theory and motivate some fundamental questions. Section 2.3 lays the mathematical groundwork by defining precisely the basic ingredients of the Optimal Transport problem, as well as some important theorems (for instance, the existence of solutions). Section 2.4 provides the necessary and sufficient conditions for optimality. In particular, it proves that it is possible to verify whether a candidate measure is a solution to the Kantorovich Problem merely by looking at its support; also, it shows that there is a function somehow associated to each Optimal Transport problem. In Section 2.5, we further study the relationship between this newfound function and the Kantorovich Problem, only to conclude that it is an indication of an underlying duality theory, which we then set out to explore.

In chapter 3, we study the Wasserstein metric. Section 3.1 defines what we mean by a Wasserstein space, and deals with some of its basic properties (for example, proving that the proposed Wasserstein metric satisfies the axioms for a metric). Section 3.2 examines the topological properties of such a space; in particular, it investigates what kind of topology the Wasserstein metric induces and whether the Wasserstein space is complete. Finally, Section 3.2 studies two geometrical properties of Wasserstein spaces: it tries to settle what constitutes a 'shortest path', or minimal geodesic, in that space; and it defines and explores a suitable notion of curvature.

2
Optimal Transport

2.1
The Monge and the Kantorovich Problems

Let us start with a concrete example. Suppose we are at war and our city has just been under a fierce bombardment. As a consequence, there are several independent fire spots across town that must be extinguished. All of them are reported to the firemen central, which must act swiftly to deploy fire trucks. Luckily, there are fire stations spread throughout the city, each with a single fire truck full of water (for simplicity, let us assume that the stations are as many as the fire spots).

The firemen central then faces a challenge: what is the best way to transport the fire trucks from the several stations to the fire spots? Note that since there is only one truck per fire station, each station must select a single fire to extinguish. A possible alternative is drawn in figure 2.1, where circles represent fire stations and squares represent the flames.

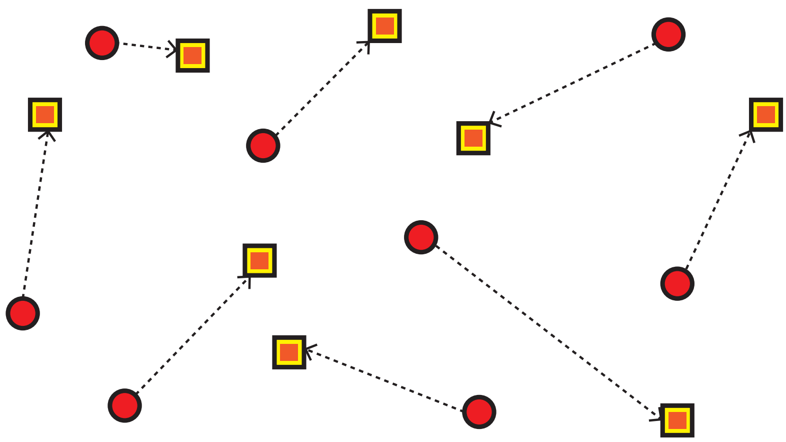

Figure 2.1: *Transporting fire trucks to flames.*

How can we model this problem mathematically? One way is to consider fire stations as points $x_1, \ldots, x_N$ in $\mathbb{R}^2$ and fire spots as points $y_1, \ldots, y_N$ in $\mathbb{R}^2$. The collection of stations can then be seen as a measure in $\mathbb{R}^2$, given by $\mu = \frac{1}{N}\sum_{i=1}^{N}\delta_{x_i}$ (where δ_{x_i} is the standard **Dirac delta measure**, defined by $\delta_{x_i}(A) = 1$ if $x_i \in A$, and 0 otherwise) and the collection of fire spots can be seen as another measure in $\mathbb{R}^2$, given by $\nu = \frac{1}{N}\sum_{j=1}^{N}\delta_{y_j}$. Notice we have normalized μ and ν so that $\mu(\mathbb{R}^2) = \nu(\mathbb{R}^2) = 1$, thereby making them probability measures.

For the problem to make sense we must have a cost associated with moving, given by a function $c : \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}$. The value $c(x, y)$ represents how costly it is to move the mass from point x to point y. In our example, it would be natural to consider the usual euclidean distance as the cost function, $c(x, y) = |x - y|$.

It is clear the problem is to transport measure μ to ν, but what should constitute a candidate for the solution? It seems reasonable to think of it as a function $T : \mathbb{R}^2 \to \mathbb{R}^2$, where $T(x_i) = y_j$ means that the i-th fire truck should proceed to the j-th fire spot. Note that because each fire station must deploy only one truck, the function is well-defined (it even allows for multiple trucks to extinguish a single fire). Thus, the firemen central's problem becomes to pick T so as to minimize $\sum_{i=1}^{N} c(x_i, T(x_i))$, conditional on T taking all the firetrucks $x_1, \ldots, x_N$ to the fire spots $y_1, \ldots, y_N$.

This is a particular instance of a general class of problems known as Optimal Transport. The field started in 1781 when the French mathematician Gaspard Monge tried to solve a question that is the continuous analog of the problem posed above. Essentially, he wanted to find the best way to transport a certain amount of sand to a construction site, where both the shape of the hole and the shape of the construction had already been decided. As before, the transport is costly, and so it should be minimized. Figure 2.2 illustrates a possible scenario.

Calling X the space where the sand is and Y the space where it should be transported to, we can model the mass of sand in X by a measure μ and the mass where the sand should be put in by another measure ν on Y. Of course, the value of the measures of the whole spaces X and Y must both be the same, or else it would be impossible to transport the mass from one place to the other: there would either be too much sand in the hole, or not enough to fill the proposed shape. Since both measures should be finite, we can normalize them and assume for simplicity $\mu(X) = 1$, $\nu(Y) = 1$.

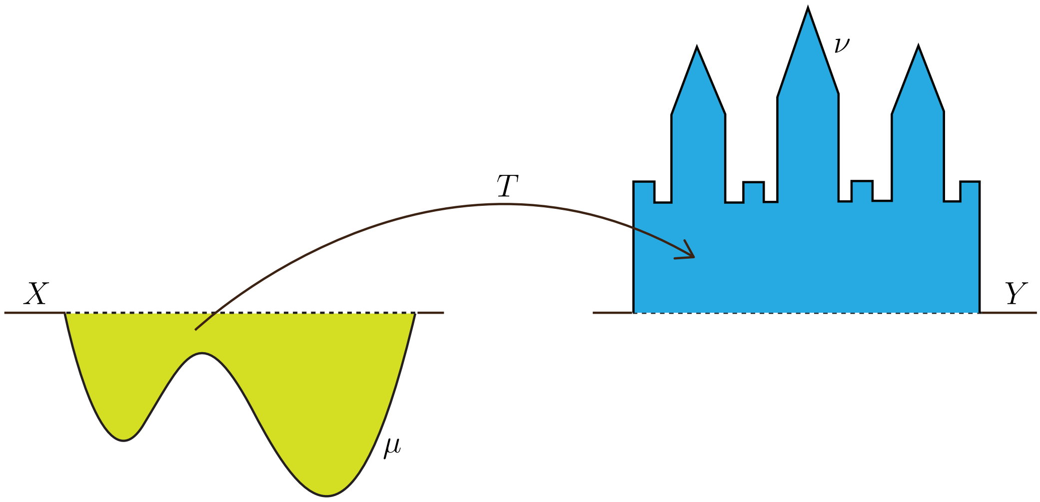

Figure 2.2: *Moving sand.*

Furthermore, as before, there must be a measurable function cost $c : X \times Y \to \mathbb{R}$ associated with transporting.

As candidates for solutions, Monge considered all measurable maps $T : X \to Y$, indicating that the mass of each point $x \in X$ should be transported to point $T(x) \in Y$. Notice this naturally imposes a restriction: since for T to be a map $T(x)$ needs to be single-valued, the mass of a point x cannot be split. Put differently, all mass taken from $T^{-1}(B) \subset X$ must go to $B \subset Y$, and so the measures μ and ν should agree on these sets. Hence, the following condition must hold for T to be considered a candidate for solution:

$$\nu(B) = \mu(T^{-1}(B)), \text{ for any measurable set } B \subset Y. \tag{2.1}$$

When this condition holds, we call ν the **push-forward** of μ by T, and denote this by $\nu = T_{\#}\mu$.

There is another useful way to characterize the push-forward. Denote by χ_B the **characteristic function** on B, that is $\chi_B(x) = 1$ if $x \in B$ and 0 otherwise. Then,

$$\chi_{T^{-1}(B)}(x) = 1 \iff x \in T^{-1}(B) \iff T(x) \in B \iff \chi_B(T(x)) = 1,$$

and we have, for $\psi = \chi_B$,

$$\begin{aligned}\int_Y \psi(y)d\nu(y) &= \int_Y \psi(y)dT_{\#}\mu(y) = T_{\#}\mu(B) = \mu(T^{-1}(B)) \\ &= \int_X \chi_{T^{-1}(B)}(x)d\mu(x) = \int_X \chi_B(T(x))d\mu(x) = \int_X \psi \circ T(x)d\mu(x).\end{aligned}$$

By the linearity of the integral, the formula above must hold for all simple functions $\psi = \sum_{j=1}^{n} c_j \chi_{B_j}$, where $c_j \in \mathbb{R}, B_j \subset Y$ measurable. We can further extend the formula to all $\psi \in L^1(T_{\#}\mu)$ such that $\psi \geq 0$ since, by a well-known Measure Theory argument, there exists a sequence of simple function $\{\psi_n\}$ such that $\psi = \sup_n \psi_n$. Then, by the Monotone Convergence Theorem,

$$\int \psi dT_{\#}\mu = \sup_n \int \psi_n dT_{\#}\mu = \sup_n \int \psi_n \circ T d\mu = \int \psi \circ T d\mu.$$

And, finally, it is easy to see the formula is true for arbitrary $\psi \in L^1(T_{\#}\mu)$: it is enough to write $\psi = \psi^+ - \psi^-$, where $\psi^+, \psi^- \geq 0$ denote the positive and negative parts of ψ, respectively, and $\psi^+, \psi^- \in L^1(T_{\#}\mu)$. So, for any $\psi \in L^1(T_{\#}\mu)$, we have that $\nu = T_{\#}\mu$ holds if and only if the following is true:

$$\int_Y \psi d\nu = \int_X (\psi \circ T) d\mu. \tag{2.2}$$

This alternative characterization of the push-forward will be useful on a number of occasions in the future.

Thus, **Monge's Problem** is to minimize the total cost associated with moving measures from X to Y, choosing among all transportation maps that preserve masses. We state it mathematically as

$$\text{minimize} \int_X c(x, T(x))\ d\mu(x) \tag{2.3}$$

among measurable maps $T : X \to Y$ such that $T_{\#}\mu = \nu$. If (2.3) admits a minimizer T^* we call it an **optimal transport map**.

Simple as the problem seems, several mathematicians tried without success to make progress on it. With enough time, it became clear why: the Monge Problem is in fact ill-posed. Indeed, there are two glaring issues.

First, no admissible T might exist: consider $\mu = \delta_x$ a Dirac delta and ν anything but. It is clear that $T(x)$ must cover at least two points in Y, which is impossible if T is a well-defined function — in our first example, this is equivalent to asking a single fire station to extinguish multiple fires (which cannot happen because, by assumption, each fire station only has one truck at its disposal).

Second, even in situations where there is an admissible map T, it is very hard to find suitable methods to attack this problem. Indeed, the typical approach through the Calculus of Variations does not yield promising results due to the highly non-linear aspect of the constraint $T_{\#}\mu = \nu$. As an example, if we assume a "good scenario" where μ and ν are defined on $\mathbb{R}^n$ and are both absolutely continuous with respect to the Lebesgue measure (i.e. $d\mu(x) = f(x)dx$, $d\nu(y) = g(y)dy$, with f, g Lebesgue-integrable) and if we guess that T must be a C^1 diffeomorphism (which is in general too optimistic a guess), then it is possible to use (2.2) and the change of variables formula to obtain

$$\int_X \psi(T(x))f(x)dx = \int_Y \psi(y)g(y)dy = \int_{T(X)} \psi(y)g(y)dy$$
$$= \int_X \psi(T(x))g(T(x))|\det \nabla T(x)|dx,$$

for all $\psi \in L^1(T_{\#}\mu)$, and so we get the constraint

$$f(x) = g(T(x))|\det \nabla T(x)|,$$

which is by all counts a highly non-linear restriction[1].

Finally, after more than a century, the Russian mathematician Leonid Kantorovich was able to make headways on the problem. He realized that modeling the solutions of Monge's Problem as functions was too stringent: it imposed the restrictive condition that each point in X needed to be taken to only one point in Y. Instead, he allowed for the possibility that a single $x \in X$ might be split and transported to several different places in Y. In our first example, this would be akin to letting one fire station have multiple trucks, each of which might be sent to a different fire spot.

With this generalization, a function would not model a solution appropriately, since functions require that $T(x)$ be unique for each x. Instead, Kantorovich considered the space of Borel probability measures on $X \times Y$, denoted by $\mathcal{P}(X \times Y)$, and thought of a solution as a measure $\pi \in \mathcal{P}(X \times Y)$. Informally, we can interpret $\pi(x, y)$ as the amount of mass that is transferred from point $x \in X$ to $y \in Y$.

Still, there is a clear restriction on the admissible measures. If we are transporting

[1]This is, in fact, the famous *Monge-Ampère* partial differential equation in disguise.

μ to ν, the total mass that π transfers from a subset $A \subset X$ to the whole space Y, namely $\pi(A \times Y)$, must coincide with the given mass of that subset, namely $\mu(A)$; the same is true for ν. Thus, we get the restrictions

$$\pi(A \times Y) = \mu(A), \qquad \pi(X \times B) = \nu(B), \tag{2.4}$$

for all measurable subsets $A \subset X$ and $B \subset Y$. It is easy to see this is equivalent to

$$\int_{X\times Y} [\varphi(x) + \psi(y)] d\pi(x, y) = \int_X \varphi(x) d\mu(x) + \int_Y \psi(y) d\nu(y) \tag{2.5}$$

for all $(\varphi, \psi) \in L^1(\mu) \times L^1(\nu)$, or even for all $(\varphi, \psi) \in L^\infty(\mu) \times L^\infty(\nu)$. Alternatively, this condition is also equivalent to

$$\mathrm{proj}_{X\#}\, \pi = \mu, \ \mathrm{proj}_{Y\#}\, \pi = \nu,$$

where $\mathrm{proj}_X : X \times Y \to X$ and $\mathrm{proj}_Y : X \times Y \to Y$ denote the **standard projections** on X and Y, respectively.

We then define the admissible set for Kantorovich's reformulation, or simply the **admissible set**, by

$$\Pi(\mu, \nu) = \{\pi \in \mathcal{P}(X \times Y) : \pi(A \times Y) = \mu(A), \pi(X \times B) = \nu(B)\}\,. \tag{2.6}$$

By considering measures in $X \times Y$ instead of functions, we arrive at the **Kantorovich Problem**. It can be stated as

$$\text{minimize} \int_{X\times Y} c(x, y) d\pi(x, y) \tag{2.7}$$

over $\pi \in \Pi(\mu, \nu)$. If (2.7) admits a minimizer, π^*, we call it an **optimal transport plan** (in contrast to an optimal transport map, which is a solution to the Monge Problem).

Though this might not be immediately obvious, Kantorovich's reformulation of Monge's Problem has several benefits. Indeed, soon after Kantorovich rephrased Monge's Problem mathematicians were able to achieve remarkable results. We list here four advantages.

First, the admissible set of measures is never empty. For instance, it is easy to see that the **product measure** $\mu \otimes \nu$, given by $(\mu \otimes \nu)(A \times B) = \mu(A) \cdot \nu(B)$ is always available, and satisfies condition (2.4) for all measurable subsets $A \subset X$, $B \subset Y$. Thus, whereas the set of optimal transport maps can be empty, the set of optimal transport plans always has at least one element.

Second, the Kantorovich Problem is really a generalization of Monge's. If there is a solution to Kantorovich's Problem that does not involve the splitting of masses, then it is easy to see it must also be a solution to Monge's. Indeed, if T solves Monge's Problem, it satisfies $T_{\#}\mu = \nu$, and thus $\pi = (\mathrm{Id} \times T)_{\#}\mu$ is an admissible measure for the Kantorovich Problem. This follows because, for all measurable subsets $A \subset X, B \subset Y$,

$$\pi(A \times Y) = (\mathrm{Id} \times T)_{\#}\mu(A \times Y) = \mu((\mathrm{Id} \times T)^{-1}(A \times Y)) = \mu(A),$$
$$\pi(X \times B) = \mu((\mathrm{Id} \times T)^{-1}(X \times B)) = \mu(T^{-1}(B)) = \nu(B),$$

where the last equality is a consequence of $T_{\#}\mu = \nu$.

Third, we can guarantee the existence of a solution to the Kantorovich Problem under quite general conditions. This will be the content of Theorem 10 below. The existence of a minimizer is not really surprising when one understands that Kantorovich's reformulation makes the objective functional linear (in the candidate measure) and the admissible set convex and compact with respect to a natural topology. This should become clearer in section 2.3.

Fourth, and more importantly, Kantorovich's relaxation transformed Monge's Problem into an infinite-dimensional Linear Programming problem. Thus, as one might expect, it admits a useful dual formulation to be explored. This is the content of Theorem 22 below.

Finally, we mention in passing the strong connection the Kantorovich Problem has with Probability Theory. Though this is immediate from the fact that we are taking μ, ν, π to be probability measures, the correspondence goes deeper. By definition, a random variable U in X is a measurable function $U : \Omega \to X$, where Ω is a space endowed with a probability measure $\mathbb{P}$. Also, the law of U, given by a probability measure μ on X, is defined by $\mathrm{law}(A) = \mathbb{P}(U^{-1}(A))$, and the expected value of a random variable U is simply its integral with respect to $\mathbb{P}$. Given these definitions, it is easy to see Kantorovich's Problem can be restated as: given two probability measures μ and ν, minimize the expected value of $\mathbb{E}[c(U, V)]$ considering all pairs of random variables $U : \Omega \to X$, $V : \Omega \to Y$, such that $\text{law}(U) = \mu$, $\text{law}(V) = \nu$.

2.2 Elementary Examples

In order to make the problem less abstract, let us consider some examples.

Example 1 (Dirac mass)**.** Let μ be any measure on X and $\nu = \delta_a$ a Dirac delta measure on Y, with $a \in Y$. Then it is quite clear that the admissible set $\Pi(\mu, \nu)$ has only one element, the measure π that transports all mass from X to a, which must therefore be optimal. The optimal cost is

$$\int_{X\times Y} c(x, y) d\pi(x, y) = \int_X c(x, a) d\mu(x).$$

Thus, both the Monge and Kantorovich Problems admit the same solution.

Example 2 (Non-existence)**.** Let us give an example where there is no optimal map that solves the problem, though there is an optimal plan. Consider $X = Y = \mathbb{R}^2$ with quadratic cost $c(x, y) = |x - y|^2$ and let λ be the Lebesgue measure. Define the measure μ_{-1} to be supported on $\{-1\} \times [0, 1]$ and equal to $\mu_{-1}(\{-1\} \times A) = \frac{1}{2}\lambda(A)$ for any $A \subset [0, 1]$ Borel-measurable; likewise define the measure μ_1 to be supported on $\{1\}\times[0, 1]$ and equal to $\mu_1(\{1\} \times A) = \frac{1}{2}\lambda(A)$ for any $A \subset [0, 1]$ Borel-measurable. Finally, define the measure μ_0 to be supported on $\{0\} \times [0, 1]$ and equal to $\mu_0(\{0\} \times A) = \lambda(A)$ for any $A \subset [0, 1]$ Borel-measurable. We shall take $\mu = \mu_0$ and $\nu = \mu_{-1} + \mu_1$, and again we must transport μ to ν.

It is straightforward to see that the Kantorovich Problem admits a solution: there is an optimal plan π^* that simply takes half the mass of a point $(0, a)$ to $(-1, a)$ and the other half to $(1, a)$, where $a \in [0, 1]$. Analytically, we have

$$\pi^*((0, a), (1, \tilde{a})) = \begin{cases} 1/2, & \text{if } a = \tilde{a} \\ 0, & \text{otherwise,} \end{cases}$$

and

$$\pi^*((0,a),(-1,\tilde{a})) = \begin{cases} 1/2, & \text{if } a = \tilde{a} \\ 0, & \text{otherwise.} \end{cases}$$

Indeed, it is easy to see π^* as defined above must be optimal, since $\int_{X\times Y}|x-y|^2 d\pi^*(x,y) = 1$ and for any candidate plan π the optimal cost will be equal or bigger than one, as each point must necessarily be transported by a distance of at least 1 (thus, $\int_{X\times Y}|x-y|^2 d\pi \geq \int_{X\times Y} 1 d\pi = 1$). See Figure 2.3$a$ below.

On the other hand, there is no solution to the Monge Problem, i.e. there is no map that solves the problem. We can create a sequence of maps that approximate the solution, but the limit of the sequence turns out not to be a map. For instance, partition the interval $\{0\}\times[0,1]$ into n equal-sized intervals, and consider the map T suggested in Figure 2.3b. Each partition interval will be transported to a side with a cost of

$$\int_0^{\frac{1}{n}} |(1,2x)-(0,x)|^2 dx = \frac{1}{n} + \frac{1}{3n^3}.$$

Since there are n intervals to be transported, the total cost will be $1+\frac{1}{3n^2} \to 1$ as $n\to\infty$.

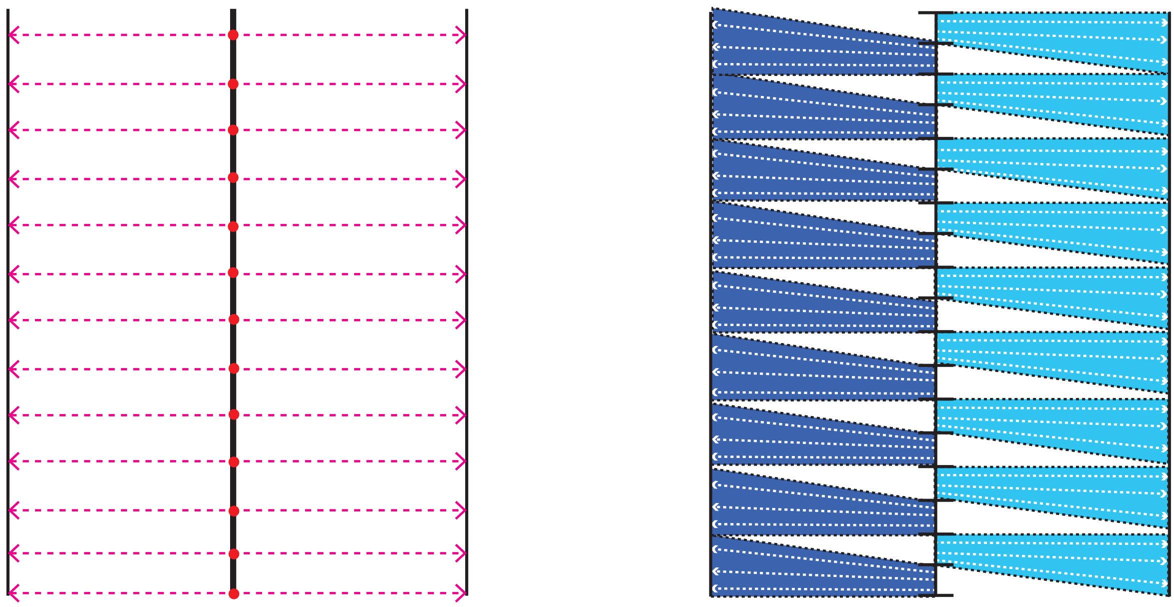

(a) (b)

Figure 2.3: *On the left, the optimal plan; on the right, an attempt to find an optimal map, which gets increasingly close to the optimal plan.*

Example 3 (Discrete case)**.** We shall consider the case where μ and ν are discrete, uniform probability measures (just as in the initial fire trucks example). Take X and Y to be discrete spaces, such that $X = \{x_1, \dots, x_N\}$, $Y = \{y_1, \dots, y_N\}$, and let

$$\mu = \frac{1}{N}\sum_{i=1}^{N} \delta_{x_i}, \qquad \nu = \frac{1}{N}\sum_{j=1}^{N} \delta_{y_j}$$

so both measures are probability measures that assign the same weight to all points. Since the problem is discrete and finite, a measure $\pi \in \Pi(\mu, \nu)$ can be thought of as an $N \times N$ matrix, $\pi = [\pi_{ij}]_{i,j}$, where π_{ij} indicates the fraction of mass that goes from point x_i to point y_j. Note that all mass from x_i goes to Y, so $\sum_{j=1}^{N} \pi_{ij} = 1$, and also all mass that gets to y_j comes from somewhere in X, so $\sum_{i=1}^{N} \pi_{ij} = 1$. If we define the **set of bistochastic** $n \times n$ **matrices** by

$$\mathcal{M}_{\mathcal{B}} = \left\{ M = (m_{ij}) \in \mathbb{R}^{N\times N} \;\middle|\; \sum_{i=1}^{N} m_{ij} = 1 \;\forall i,\; \sum_{j=1}^{N} m_{ij} = 1 \;\forall j, \text{ and } m_{ij} \geq 0 \;\forall i, j \right\}, \tag{2.8}$$

then the Kantorovich Problem in this setting becomes:

$$\min_{\pi} \left\{ \frac{1}{N} \sum_{i,j} \pi_{ij} c(x_i, y_j) \;\middle|\; \pi \in \mathcal{M}_{\mathcal{B}} \right\}. \tag{2.9}$$

This is a simple Linear Programming problem, since the objective function $\sum_{i,j} \pi_{ij} c(x_i, y_j)$ is linear and the admissible set $\mathcal{M}_{\mathcal{B}}$ is both convex and bounded. By Choquet's Theorem [16], we know that at least some solutions to (2.9) must be **extremal points** of $\mathcal{M}_{\mathcal{B}}$, i.e. points that cannot be expressed as a convex combination of other points in $\mathcal{M}_{\mathcal{B}}$. In turn, by Birkhoff's Theorem [7], we know all extremal points of $\mathcal{M}_{\mathcal{B}}$ are **permutation matrices**, i.e. matrices such that, for some permutation σ, $\pi_{ij} = 1$, if $j = \sigma(i)$, and 0 otherwise.

This means that there are solutions to Kantorovich's Problem that are also solutions to Monge's Problem. Indeed, as some optimal plans are permutations (which take all the mass from a point x_i to a point y_j), the solution can be thought of as a mapping. If we denote by S_N the set of permutations of $\{1, \dots, N\}$, the Monge Problem can be recast as

$$\min_{\sigma} \left\{ \frac{1}{N} \sum_{i=1}^{N} c(x_i, y_{\sigma(i)}) \;\middle|\; \sigma \in S_n \right\}, \tag{2.10}$$

and it is clear (2.9) and (2.10) are the same.

It is then easy to see that the Optimal Transportation problem reduces to an **optimal matching problem** in the discrete case, since we are simply 'matching' the points in X to points in Y so as to minimize a given matrix cost, $c = [c_{ij}]_{i,j}$. This particular case has received a lot of attention, and much theory has been developed lately[2].

Example 4 (Book-shifting)**.** Consider the Lebesgue measure on $\mathbb{R}$, denoted by λ, the spaces $X = Y = \mathbb{R}$, the cost function $c(x,y) = |x-y|^p$, where $p > 0$, and the uniform probability measures $\mu = \frac{1}{n}\lambda\big|_{[0,n]}$, and $\nu = \frac{1}{n}\lambda\big|_{[1,n+1]}$, where $n > 1$.

Define two transport maps

$$T_1(x) = x + 1 \qquad \text{and} \qquad T_2(x) = \begin{cases} x & \text{if } 1 < x \leq n \\ x + n & \text{if } 0 \leq x \leq 1. \end{cases}$$

Note T_1 simply "shifts" the measure μ by one, whereas T_2 keeps the common mass in place and moves the remaining mass. It is easy to see from (2.3) that the total costs associated to maps T_1 and T_2 are, respectively

$$\int_X |x - T_1(x)|^p d\mu(x) = 1, \qquad \int_X |x - T_2(x)|^p d\mu(x) = n^{p-1}.$$

If $p = 1$ then both transports cost the same, as our intuition would have made us believe. If $p > 1$, though, T_1 induces a smaller cost, whereas if $p < 1$, T_2 is better. This has to do with the fact that $f(x) = x^p$ is a convex function when $p > 1$ (so two medium-length shifts are better than one big shift and one small shift) and concave when $p < 1$ (so two medium-length shifts are worse than one big shift and one small shift).

Furthermore, if $p = 1$, it is possible to prove that both T_1 and T_2 are optimal transport maps by elementary methods. To see this, first notice the infimum of the Monge Problem is always greater than

$$\sup_{\varphi} \left\{ \int_X \varphi d(\mu - \nu) \;\Big|\; \varphi \in \mathrm{Lip}_1(X) \right\}, \tag{2.11}$$

[2]Though in this discrete setting the Kantorovich Problem becomes a traditional finite-dimensional Linear Programming problem, one should not be tempted to just use the Simplex algorithm to solve it. Indeed, there are better algorithms, such as the *Hungarian algorithm* [25], that exploit the additional structure this formulation imposes.

where $\mathrm{Lip}_1(X)$ denotes the space of **1-Lipschitz functions** on X (i.e. those functions φ that satisfy the condition $|\varphi(x) - \varphi(y)| \le |x - y|$, $\forall x, y \in X$), because, by (2.2),

$$\int_X \varphi d(\mu - \nu) = \int_X (\varphi(x) - \varphi(T(x)))d\mu(x) \le \int_X |x - T(x)|d\mu(x).$$

Thus, if we can find a 1-Lipschitz function φ and a transport map so that equality holds, we will simultaneously have that T solves the Monge Problem and φ solves (2.11). Fortunately, it is easy to see $\varphi(x) = -x$ fits the bill:

$$\begin{aligned}\int_X \varphi d(\mu - \nu) = \int_X -xd(\mu - \nu) &= \int_X xd\nu(x) - \int_X x\mu(x) \\ &= \frac{1}{n}\int_1^{n+1} xdx - \frac{1}{n}\int_0^n xdx = \frac{(n+1)^2 - 1^2 - n^2}{2n} = 1.\end{aligned}$$

Hence, the optimal value of the Monge Problem is 1 and both T_1 and T_2 are optimal maps.

Though the method above solves the Monge Problem, it is not clear that T_1 and T_2 also induce proper solutions to the Kantorovich Problem. As we shall see later, however, the Lipschitz condition devised can be extended to transference plans, so $(\mathrm{Id} \times T_1)_\#$ and $(\mathrm{Id} \times T_2)_\#$ do solve the Kantorovich Problem. In fact, under suitable conditions, the Kantorovich Problem always admits a solution that leaves all common mass in place (see example 32).

Example 5 (Non-uniqueness)**.** An interesting property arises if $n = 1$ in the preceding example. Then $\mu = \lambda\big|_{[0,1]}$, $\nu = \lambda\big|_{[1,2]}$ and we still have $c(x, y) = |x - y|^p$, with $p > 0$. Both T_1 and T_2 now define the same function, so let us consider the following two maps:

$$T_1(x) = x + 1 \qquad \text{and} \qquad T_3(x) = 2 - x.$$

The costs associated with T_1 and T_3 are, respectively,

$$\int_0^1 1dx = 1 \qquad \text{and} \qquad \int_0^1 (2 - 2x)^p dx = \frac{2^p}{p+1}.$$

Again, if $p > 1$, the cost function is convex, and so it is better to move all points by a medium amount than move half by a large amount and half by a small amount; thus, T_1 is optimal. If, on the other hand, $p < 1$, the cost function is concave, so it is better to move all points by a medium amount than half by a large amount and half by a small amount; thus, T_2 is optimal because it reverses the orientation.

A curious property shows up when $p = 1$. Then any transference plan $\pi \in \Pi(\mu, \nu)$ is optimal because the cost function reduces to $c(x, y) = y - x$, and so, by (2.5),

$$\int_{X\times Y} c(x,y)d\pi(x,y) = \int_{X\times Y} (y-x)d\pi(x,y) = \int_Y ydy - \int_X xdx$$
$$= \frac{2^2 - 1^2 - 1^2}{2} = 1.$$

Hence, the solution to the Kantorovich Problem is not unique. As an obvious consequence, the solution to the Monge Problem (which also attains the value of 1) cannot be unique as well.

2.3
Basic Definitions and Results

Though we were able to find interesting characterizations for the examples posed in the last section, in order to derive general results in the Optimal Transport theory, we must restrict the scope of some of our objects. The conditions we shall impose are general enough to encompass all the previous examples and more.

In what follows, we will take X and Y to be complete and separable metric spaces (also called **Polish spaces**), and denote by $\mathcal{P}(X), \mathcal{P}(Y)$ the sets of Borel probability measures on X and Y, respectively (we recall that a **Borel probability measure** is a measure defined on the σ-algebra generated by the open sets of some topological space). Also, we define the **support** of a measure $\mu \in \mathcal{P}$, denoted by $\mathrm{supp}(\mu)$, as the smallest closed set on which μ is concentrated. We will consider $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$, $\pi \in \mathcal{P}(X \times Y)$, and restrict our cost functions to continuous functions $c : X \times Y \to \mathbb{R}$ bounded from below.

Also, given a topological space X, it will be useful to define a function $f : X \to \mathbb{R} \cup \{-\infty, \infty\}$ to be **lower semi-continuous** if, for all $x \in X$ and sequence $\{x_n\}_{n\in\mathbb{N}}$ such that $x_n \to x$, it holds that $f(x) \leq \liminf_{n\to\infty} f(x_n)$. Symmetrically, we define a function $f : X \to \mathbb{R}$ to be **upper semi-continuous** if, for all $x \in X$ and sequence $\{x_n\}_{n\in\mathbb{N}}$ such that $x_n \to x$, we have $f(x) \geq \limsup_{n\to\infty} f(x_n)$.

The **Monge Problem** can then be stated as: given two measures $\mu \in \mathcal{P}(X)$ and $\nu \in \mathcal{P}(Y)$ and a continuous cost function $c : X \times Y \to \mathbb{R}$ bounded below, solve

$$\inf_T \left\{ \int_X c(x, T(x)) d\mu(x) \mid T : X \to Y \text{ measurable}, T_{\#}\mu = \nu \right\}. \tag{2.12}$$

As before, an optimal solution T^* to the Monge Problem will be called an optimal transference map.

The **Kantorovich Problem** can be stated as: given two measures $\mu \in \mathcal{P}(X)$ and $\nu \in \mathcal{P}(Y)$ and a continuous cost function $c : X \times Y \to \mathbb{R}$ bounded below, solve

$$\inf_\pi \left\{ \int_{X\times Y} c(x, y) d\pi(x, y) \mid \pi \in \Pi(\mu, \nu) \right\}, \tag{2.13}$$

where $\Pi(\mu, \nu)$ is defined as in (2.6). An optimal solution π^* to the Kantorovich Problem will be called an optimal transference plan.

Finally, denote by $\Pi^*(\mu, \nu)$ the **set of optimal plans**. The main purpose of this section will be to prove that $\Pi^*(\mu, \nu)$ is always non-empty; i.e. minimizers to the Kantorovich Problem always exist (see Theorem 10).

We first recall basic results concerning analysis on Polish spaces. Let us take as the standard topology on $\mathcal{P}(X)$ the so-called **narrow topology**[3], induced by the convergence against continuous bounded functions. It is possible to show that the narrow topology is metrizable[4]. We say a sequence of measures $\{\mu_n\}_{n\in\mathbb{N}} \subset \mathcal{P}(X)$ **narrowly converges** to another measure μ if

$$\int_X \varphi d\mu_n \to \int_X \varphi d\mu, \qquad \forall \varphi \in C_b(X).$$

Having specified a topology, it will be useful time and again to know when a measure admits a compact approximation for the space X. In fact, this yields the following definition.

Definition 6. Given a Polish space X, a set $\mathcal{S} \subset \mathcal{P}(X)$ is called **tight** if for all $\varepsilon > 0$ there exists a compact set $K_\varepsilon \subset X$ such that $\mu(X \setminus K_\varepsilon) \leq \varepsilon$, for all $\mu \in \mathcal{S}$.

[3]The narrow topology is also referred to as the weak topology of measures; from the point of view of Functional Analysis, by considering $C_b(X)$ to be the base space, it should be called the weak-$*$ topology.

[4]We shall pick up on this issue on Corollary 38. For now it suffices to say that it is metrizable, for instance, by the Levy-Prokhorov distance, defined in (3.10).

Tight sets are important because they have an extremely useful equivalence, as the next theorem shows.

Theorem 7 (Prokhorov's Theorem)**.** *If X is a Polish space, then a subset $\mathcal{S} \subset \mathcal{P}(X)$ is tight if and only if it is pre-compact for the narrow topology (i.e. its closure is a compact set).*

Prokhorov's Theorem is a very important tool in Polish spaces, and shall be employed frequently in this text. For a proof, refer to [30]. As a straightforward corollary, we obtain Ulam's Lemma.

Corollary 8 (Ulam's Lemma)**.** *A Borel probability measure μ on a Polish space X is automatically tight.*

And, lastly, a reassuring result: the completeness and separability of the space X are inherited by the space of probability measures on that set, $\mathcal{P}(X)$. We shall state the theorem here, though we will later have the opportunity to prove it in Theorem 39, after we come across a suitable metric for the space $\mathcal{P}(X)$.

Theorem 9. *If X is a Polish space, then $\mathcal{P}(X)$ is also a Polish space.*

It is now easy to prove the existence of minimizers for the Kantorovich Problem.

Theorem 10. *Let X, Y be Polish spaces and let $c : X \times Y \to \mathbb{R}$ be a continuous function bounded from below. Given two measures $\mu \in \mathcal{P}(X), \nu \in \mathcal{P}(Y)$, there exists a measure $\pi^* \in \Pi(\mu, \nu)$ such that*

$$\int_{X\times Y} c(x,y)d\pi^*(x,y) = \inf_{\pi} \left\{ \int_{X\times Y} c(x,y)d\pi(x,y) \;\middle|\; \pi \in \Pi(\mu,\nu) \right\},$$

i.e. π^* is a minimizing measure for the Kantorovich Problem.

Proof. We proceed as follows: first, we show the set $\Pi(\mu, \nu)$ is tight, hence pre-compact; because we can prove $\Pi(\mu, \nu)$ is actually closed, it is compact; then, as a consequence, we are able to find a convergent minimizing subsequence for the Kantorovich Problem; finally, we show that the continuity of the cost function c guarantees that the limit of the subsequence is indeed an optimal solution.

To show $\Pi(\mu, \nu)$ is tight, notice that, by Ulam's Lemma, both μ and ν are tight. Therefore, given $\varepsilon > 0$, we can find $K_1 \subset X$ and $K_2 \subset Y$ compact subsets such that $\mu(X \setminus K_1) < \varepsilon/2$ and $\nu(Y \setminus K_2) < \varepsilon/2$. Thus, for any $\pi \in \Pi(\mu, \nu)$, we get

$$\pi(X \times Y \setminus K_1 \times K_2) \leq \pi((X \setminus K_1) \times Y) + \pi(X \times (Y \setminus K_2)) = \mu(X \setminus K_1) + \nu(Y \setminus K_2) < \varepsilon,$$

and so $\Pi(\mu, \nu)$ is tight in $\mathcal{P}(X \times Y)$.

By Prokhorov's Theorem, we get that $\Pi(\mu, \nu)$ is pre-compact in $\mathcal{P}(X \times Y)$. To show that it is in fact (narrowly) compact, we must see that $\Pi(\mu, \nu)$ is closed. Let $\{\pi_n\} \subset \Pi(\mu, \nu)$ be such that $\pi_n \to \pi^*$ narrowly. We shall have $\pi^* \in \Pi(\mu, \nu)$ if $\mathrm{proj}_{X\#}\, \pi^* = \mu$ and $\mathrm{proj}_{Y\#}\, \pi^* = \nu$.

Since $\mathrm{proj}_X \in C_b(X)$ and $\mathrm{proj}_Y \in C_b(Y)$, we get, for $\varphi \in C_b(X)$,

$$\begin{aligned}
\int_X \varphi d(\mathrm{proj}_{X\#}\, \pi_*) &= \int_{X \times Y} \varphi(\mathrm{proj}_X(x, y)) d\pi_*(x, y) \\
&= \lim_{n \to \infty} \int_{X \times Y} \varphi(\mathrm{proj}_X(x, y)) d\pi_n(x, y) \\
&= \lim_{n \to \infty} \int_X \varphi d(\mathrm{proj}_{X\#}\, \pi_n) = \int_X \varphi d\mu,
\end{aligned}$$

where the first equality follows by (2.2), the second by narrow convergence of $\{\pi_n\}$, the third again by (2.2), and the fourth because $\pi_n \in \Pi(\mu, \nu)$. Thus, since the equality above holds for any $\varphi \in C_b(X)$, we have that $\mathrm{proj}_{X\#}\, \pi^* = \mu$. An analogous calculation shows that $\mathrm{proj}_{Y\#}\, \pi^* = \nu$, and we get (narrow) compactness of $\Pi(\mu, \nu)$.

Finally, since c is continuous, there exists an increasing sequence of continuous and bounded functions $c_n : X \times Y \to \mathbb{R}$ such that $c(x, y) = \lim_{n \to \infty} c_n(x, y)$. Take a minimizing sequence $\{\pi_k\}_{k \in \mathbb{N}} \subset \Pi(\mu, \nu)$ and let π^* be an accumulation point. Then:

$$\begin{aligned}
\int_{X \times Y} c d\pi^* &= \lim_{n \to \infty} \int_{X \times Y} c_n d\pi^* \leq \lim_{n \to \infty} \limsup_{k \to \infty} \int_{X \times Y} c_n d\pi_k \\
&\leq \limsup_{k \to \infty} \lim_{n \to \infty} \int_{X \times Y} c_n d\pi_k = \limsup_{k \to \infty} \int_{X \times Y} c d\pi_k = \inf_{\pi \in \Pi(\mu, \nu)} \int_{X \times Y} c d\pi,
\end{aligned} \tag{2.14}$$

where the first and second equalities are a consequence of the Monotone Convergence Theorem, and the first inequality follows because $\{\pi_k\}$ is a minimizing sequence.

Hence, as the measure π^* is in the admissible set and attains the infimum in the Kantorovich Problem, it is an optimal solution. ■

Remark 2.3.1. Taking the cost function c to be continuous was more restrictive than needed. Note that any positive lower semi-continuous function $f : X \to \mathbb{R}$ can be approximated by a sequence of continuous bounded functions $\tilde{f}_n : X \to \mathbb{R}$ (for instance, consider $\tilde{f}_n(x) = \inf_{z \in X}\{f(z) + nd(z, x)\}$, where d is any metric on X). Thus, the proof above would hardly need a modification if we considered c to be lower semi-continuous instead of continuous.

Now that we have guaranteed the existence of a minimizer for the Kantorovich Problem, several questions arise: (i) under what circumstances are these minimizers unique? (ii) can we find good criteria for optimality? (iii) are there properties that any optimal solution must satisfy (for instance, do they preserve orietantion)? (iv) when will a solution to the Kantorovich Problem induce a solution to the Monge Problem? (v) can we estimate how close the solution of the Kantorovich Problem is to the solution of the Monge Problem?

In the next sections we try to answer some of these questions.

2.4
Necessary and Sufficient Conditions for Optimality

The first important result of this section is the perhaps surprising condition that optimality in the Kantorovich Problem is a property that can depend solely on the *support* of the candidate measure $\pi \in \Pi(\mu, \nu)$. In other words, one can determine whether a given transference plan is optimal only by looking at its support; there is no need to see how the mass is distributed. The second main result is that the support of an optimal measure is always contained in a generalization of the gradient of a function φ, and, usually, the transport does in fact take the form $\nabla\varphi$. Both of these observations are part of Theorem 19.

To build intuition for the necessary and sufficient conditions for optimality, it is useful to consider a particular case of the Kantorovich Problem. In some sense, we shall be looking at a "best-case scenario": we take $X = \mathbb{R}^n$, $Y = \mathbb{R}^n$, $\mu = \frac{1}{N}\sum_{i=1}^{N} \delta_{x_i}$ and $\nu = \frac{1}{N}\sum_{j=1}^{N} \delta_{y_j}$ and the quadratic cost $c(x, y) = |x - y|^2/2$. Note this is precisely the setting of example 3.

The following three classical concepts in convex analysis are crucial. First, we say a set $\Gamma \subset \mathbb{R}^n \times \mathbb{R}^n$ is **cyclically monotone** if, for any $N \in \mathbb{N}$ and any family of points in Γ, $(x_1, y_1), \ldots, (x_N, y_N)$, we have

$$\sum_{i=1}^{N} \langle x_i, y_i \rangle \geq \sum_{i=1}^{N} \langle x_i, y_{\sigma(i)} \rangle,$$

for all possible permutations σ of $\{1, \ldots, N\}$. It is "cyclical" because it suffices to check the permutation $y_1 \to y_2 \to \cdots \to y_N \to y_1$; it is "monotone" because when $N = 2$ monotonicity and cyclical monotonicity are equivalent.

Second, we define $\varphi : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$ to be a (proper) **convex function** if $\varphi \not\equiv \infty$ and, for any $t \in (0, 1)$ and $x_1, x_2 \in \mathbb{R}^n$, it holds that

$$\varphi((1-t)x_1 + tx_2) \leq (1-t)\varphi(x_1) + t\varphi(x_2).$$

A function is said to be (proper) **concave** if $-\varphi$ is (proper) convex.

Third, we deal with a generalization of the idea of differentiability for convex functions. Given a (proper) convex function $\varphi : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$, the **subdifferential** of φ, denoted by $\partial_-\varphi$, is a set-valued mapping given by the relation

$$y \in \partial_-\varphi(x) \iff \forall z \in \mathbb{R}^n,\ \varphi(z) \geq \varphi(x) + \langle y, z - x \rangle. \tag{2.15}$$

We usually identify the subdifferential $\partial_-\varphi$ with its graph. It is possible to prove that φ is differentiable at a point x if and only if $\partial_-\varphi(x)$ contains a single element, which is, naturally, $\nabla f(x)$. Note that if $\bar{x}$ is such that $\varphi(\bar{x}) = \infty$, then $\partial_-\varphi(\bar{x}) = \varnothing$. Figure 2.4 gives some intuition to definition (2.15).

If we consider φ to be concave, then we define its **superdifferential**, denoted by $\partial_+\varphi$, by the relation

$$y \in \partial_+\varphi(x) \iff \forall z \in \mathbb{R}^n,\ \varphi(z) \leq \varphi(x) + \langle y, z - x \rangle. \tag{2.16}$$

We now proceed to relate these concepts with the optimality conditions of the Kantorovich Problem. It is clear that a plan $\pi \in \Pi(\mu, \nu)$ is optimal if and only if

$$\sum_{i=1}^{N} \frac{|x_i - y_i|^2}{2} \leq \sum_{i=1}^{N} \frac{|x_i - y_{\sigma(i)}|^2}{2},$$

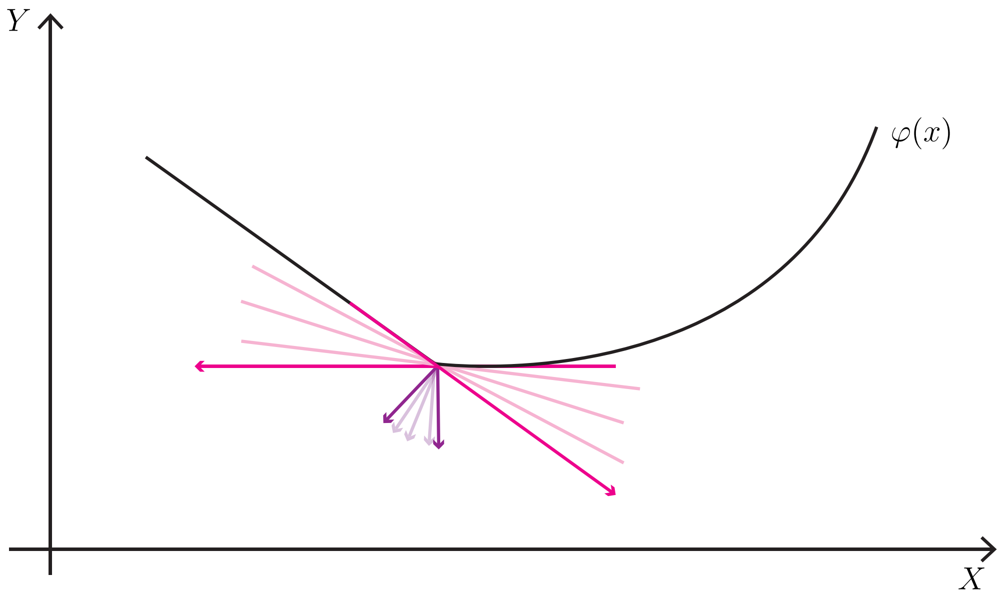

Figure 2.4: *The subdifferential of a convex function $\varphi(x)$, given by the cone of purple normal vectors.*

for all $N \in \mathbb{N}$, $(x_i, y_i) \in \mathrm{supp}(\pi)$, $i = 1, \ldots, N$ and σ permutation of $\{1, \ldots, N\}$. By expanding and using inner products, we get the equivalent condition:

$$\sum_{i=1}^{N} \langle x_i, y_i \rangle \geq \sum_{i=1}^{N} \langle x_i, y_{\sigma(i)} \rangle. \tag{2.17}$$

Thus, we get in a straightforward way that π is optimal if and only if its support $\mathrm{supp}(\pi)$ is cyclically monotone.

We recall an important theorem about cyclically monotone sets.

Theorem 11 (Rockafellar)**.** *A nonempty set $\Gamma \subset \mathbb{R}^n \times \mathbb{R}^n$ is cyclically monotone if and only if it is included in the subdifferential of a lower semi-continuous convex function $\varphi : \mathbb{R}^n \to \mathbb{R} \cup \{\infty\}$, with $\varphi \not\equiv \infty$.*

Proof. ($\Leftarrow$) Let $\varphi : \mathbb{R}^n \to \mathbb{R}$ be a convex function, and $(x_1, y_1), \ldots, (x_N, y_N) \in \mathrm{supp}(\partial_- \varphi)$, so that $y_i \in \partial_- \varphi(x_i)$ for all $i = 1, \ldots, N$. By (2.15), this means that

$$\varphi(z) \geq \varphi(x_i) + \langle y_i, z - x_i \rangle,$$

for all $z \in \mathbb{R}^n$.

And so we find that

$$\begin{aligned}
\varphi(x_2) &\geq \varphi(x_1) + \langle y_1, x_2 - x_1 \rangle, \\
\varphi(x_3) &\geq \varphi(x_2) + \langle y_2, x_3 - x_2 \rangle, \\
&\vdots \\
\varphi(x_N) &\geq \varphi(x_{N-1}) + \langle y_{N-1}, x_N - x_{N-1} \rangle, \\
\varphi(x_1) &\geq \varphi(x_N) + \langle y_N, x_1 - x_N \rangle.
\end{aligned} \tag{2.18}$$

By adding all these inequalities, we have that

$$\sum_{i=1}^{N} \langle y_i, x_{i+1} - x_i \rangle \leq 0. \tag{2.19}$$

(with the convention that $x_{N+1} = x_1$). Since (2.19) holds for all sets of points $(x_1, y_1), \ldots, (x_N, y_N)$, it is clear that the subdifferential of φ is cyclically monotone.

($\Rightarrow$) Let $\Gamma \subset \mathbb{R}^n \times \mathbb{R}^n$ be a cyclically monotone set, we must find a lower semi-continuous convex function φ (not identically infinite) such that $\Gamma \subset \partial_- \varphi$. By (2.18), it seems reasonable to pick some $(x_0, y_0) \in \Gamma$ and from that define

$$\varphi(x) = \sup \left\{ \langle y_N, x - x_N \rangle + \ldots + \langle y_0, x_1 - x_0 \rangle \;\middle|\; N \in \mathbb{N}; (x_1, y_1), \ldots, (x_N, y_N) \in \Gamma \right\}. \tag{2.20}$$

Note that φ is a convex lower semi-continuous function because it is the supremum of affine functions. Also, φ is not identically infinite since $\varphi(x_0) \leq 0$, by (2.19).

Having defined φ, all that is left to check is that Γ is indeed in the subdifferential of φ. Take $(x, y) \in \Gamma$ and let $z \in \mathbb{R}^n$. Since we want to prove that $\varphi(z) \geq \varphi(x) + \langle y, z - x \rangle$, it will suffice to have that, for all $\alpha < \varphi(x)$

$$\varphi(z) \geq \alpha + \langle y, z - x \rangle. \tag{2.21}$$

But if $\alpha < \varphi(x)$, by (2.20), we can find $(x_1, y_1), \ldots, (x_N, y_N) \in \Gamma$, for some $N \in \mathbb{N}$, such that

$$\alpha \leq \langle y_N, x - x_N \rangle + \cdots + \langle y_0, x_1 - x_0 \rangle,$$

and so

$$\alpha + \langle y, z - x \rangle \leq \langle y, z - x \rangle + \langle y_N, x - x_N \rangle + \cdots + \langle y_0, x_1 - x_0 \rangle.$$

By taking $x = x_{N+1}, y = y_{N+1}$, we get (2.21) and conclude the proof. ∎

Remark 2.4.1. Note that although the above theorem might seem a bit mysterious at first, its continuous counterpart is relatively clear: equation (2.17) is equivalent to $\sum_{i=1}^{N}\langle y_i, x_i - x_{i-1}\rangle \leq 0$ (with the understanding that $x_0 = x_N$), which translates in a continuous setting to $\oint \overline{y}(x)dx \leq 0$. This in turn implies that the vector field $\overline{y}(x)$ is conservative, so $\overline{y}(x) = \nabla\varphi(x)$ for some function $\varphi(x)$.

We thus get an equivalence between three key concepts: a measure $\pi \in \Pi(\mu, \nu)$ is optimal if and only if $\text{supp}(\pi)$ is cyclically monotone if and only if there exist a convex lower semi-continuous function φ such that π is concentrated on the subdifferential of φ.

This relationship between optimality, cyclical monotonicity and subdifferentials of convex lower semi-continuous functions is indeed a remarkable one. We can attest a measure's optimality for the Kantorovich Problem simply by looking at its support. Moreover, each solution to the Kantorovich Problem has somehow a convex lower semi-continuous function associated to it.

Unfortunately, the connections made above relied intrinsically on the format of the cost function, assumed to be $c(x, y) = |x - y|^2/2$. Hence, *a priori*, one should have no reason to expect them to hold for more general cost functions. Still, fortunately, we are able to regain these equivalences to any continuous cost function c by appropriately generalizing the notions of cyclical monotonicity, convexity and subdifferential or, as we shall prefer, by generalizing the notions of cyclical monotonicity, concavity and superdifferential. We now turn to this direction.

Definition 12. A set $\Gamma \subset X \times Y$ is said to be c**-cyclically monotone** if, for any $N \in \mathbb{N}$ and any family of points in Γ, $(x_1, y_1), \ldots, (x_N, y_N)$, we have

$$\sum_{i=1}^{N} c(x_i, y_i) \leq \sum_{i=1}^{N} c(x_i, y_{\sigma(i)}),$$

for all possible permutations σ of $\{1, \ldots, N\}$ (or simply for the permutation $\sigma(i) = i+1$ for $i = 1, \ldots, N-1$ and $\sigma(N) = 1$).

Remark 2.4.2. Though it may be hard to characterize c-cyclically monotone sets, it is not difficult to see why they are important. Indeed, they provide a criteria for 'critical points': if $\{(x_1, y_1), \ldots, (x_N, y_N)\}$ is not a c-cyclically monotone set, then they cannot

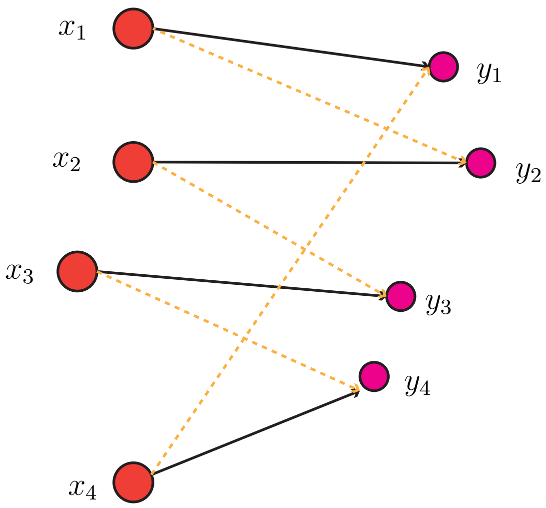

Figure 2.5: *Two possible cycles.*

be contained in the support of an optimal measure, since we can find a permutation σ such that the optimal measure can be improved by sending some mass from x_1 to $y_{\sigma(1)}$, from x_2 to $y_{\sigma(2)}$, and so on. By proceeding this way, we can keep improving a candidate measure until it is supported on a set that is c-cyclically monotone.

The important question is: would that 'improved' measure necessarily be an optimal one? Or could this process just yield a point of 'local minimum'? Though we will only prove the answer in Theorem 19 below, here is a reasonable argument: since the admissible set of the Kantorovich Problem is convex and compact and the objective function linear, we are in essence dealing with an infinite-dimensional Linear Programming problem, as alluded to before. Hence, minimizing the cost functional means walking along the extreme points of our convex domain and all local minimum points should be global minimum points. In particular, the Kantorovich Problem has either one solution (i.e. an 'edge') or infinitely many solutions (i.e. a 'face').

Definition 13. Let $\varphi : X \to \mathbb{R} \cup \{\pm\infty\}$ be a function, $\varphi \not\equiv -\infty$. We define its **$c$-transform** $\varphi^c : Y \to \mathbb{R} \cup \{-\infty\}$ by

$$\varphi^c(y) = \inf_{x \in X} c(x, y) - \varphi(x).$$

Analogously, given a function $\psi : Y \to \mathbb{R} \cup \{\pm\infty\}$, $\psi \not\equiv -\infty$, we define its c-transform $\psi^c : X \to \mathbb{R} \cup \{-\infty\}$ as

$$\psi^c(x) = \inf_{y \in Y} c(x, y) - \psi(y).$$

We are now ready to generalize the idea of concavity, based on the concepts of c-transforms.

Definition 14. A function $\varphi : X \to \mathbb{R} \cup \{-\infty\}$ is said to be **c-concave** if there exists a function $\psi : Y \to \mathbb{R} \cup \{-\infty\}$ such that $\varphi = \psi^c$. Likewise, a function $\psi : Y \to \mathbb{R} \cup \{-\infty\}$ is **c-concave** if there exists $\varphi : X \to \mathbb{R} \cup \{-\infty\}$ such that $\psi = \varphi^c$.

Since c is continuous, it is clear that the c-transforms are automatically upper semi-continuous, so we do not have to worry about the measurability of those functions.

Denoting $(\psi^c)^c$ by ψ^{cc}, we have the following simple proposition.

Proposition 15. *Given a function $\psi : Y \to \mathbb{R} \cup \{\pm\infty\}$, it holds that $\psi^c = \psi^{ccc}$.*

Proof. We simply use the definition of a c-concave function repeatedly.

$$\begin{aligned}
\psi^c(x) &= \inf_{y \in Y} c(x,y) - \psi(y) \\
\psi^{cc}(\tilde{y}) &= \inf_{\tilde{x} \in X} c(\tilde{x},\tilde{y}) - \psi^c(\tilde{x}) = \inf_{\tilde{x} \in X} \left(c(\tilde{x},\tilde{y}) - \inf_{y \in Y} c(\tilde{x},y) - \psi(y) \right) \\
&= \inf_{\tilde{x} \in X} \left(- \inf_{y \in Y} \left(c(\tilde{x},y) - c(\tilde{x},\tilde{y}) - \psi(y) \right) \right) = - \sup_{\tilde{x} \in X} \inf_{y \in Y} c(\tilde{x},y) - c(\tilde{x},\tilde{y}) - \psi(y) \\
\psi^{ccc}(x) &= \inf_{\tilde{y} \in Y} \left(c(x,\tilde{y}) + \sup_{\tilde{x} \in X} \inf_{y \in Y} \left(c(\tilde{x},y) - c(\tilde{x},\tilde{y}) - \psi(y) \right) \right) \\
&= \inf_{\tilde{y} \in Y} \sup_{\tilde{x} \in X} \inf_{y \in Y} c(x,\tilde{y}) - c(\tilde{x},\tilde{y}) + c(\tilde{x},y) - \psi(y).
\end{aligned}$$

By taking $\tilde{x} = x$, we have that $\psi^{ccc} \geq \psi^c$, and by taking $y = \tilde{y}$, we have that $\psi^{ccc} \leq \psi^c$. ■

Using the proposition above, we can find another useful characterization for c-concave functions.

Corollary 16. *A function $\varphi : X \to \mathbb{R} \cup \{-\infty\}$ is c-concave if and only if it holds that $\varphi^{cc} = \varphi$.*

Proof. Suppose $\varphi = \varphi^{cc}$. Considering the function $\psi = \varphi^c$, we see that $\varphi = \psi^c$, so φ is c-concave. Conversely, if φ is c-concave, there exists ψ such that $\psi^c = \varphi$. By Proposition 15, we get $\varphi = \psi^c = \psi^{ccc} = (\psi^c)^{cc} = \varphi^{cc}$. ■

Remark 2.4.3. The definition of a c-concave function generalizes the notion of concavity in the sense that they can be written as the pointwise infima of functions of the form $c(x, y) - \psi(y)$ for some ψ, whereas concave functions can be defined as the infimum of a family of affine (and upper semicontinuous) functions. It should be noted that the definition adopted in this text is not universal, but it is the one more appropriate for the topics that will be studied subsequently. Some authors also prefer to work with the notion of c-convex functions.

Finally, we come to our last crucial definition.

Definition 17. Given a c-concave function $\varphi : X \to \mathbb{R} \cup \{-\infty\}$, we define its **$c$-superdifferential** $\partial^c\varphi \subset X \times Y$ as

$$\partial^c\varphi = \{(x, y) \in X \times Y \mid \varphi(x) + \varphi^c(y) = c(x, y)\}\,.$$

As before, in the above definition we identify the c-superdifferential of a function with its graph. We also define the c-superdifferential of φ at a point x, denoted $\partial^c\varphi(x)$, by $\partial^c\varphi(x) = \{y \in Y \mid (x, y) \in \partial^c\varphi\}$. A symmetric definition holds for c-concave functions of the form $\psi : Y \to \mathbb{R} \cup \{-\infty\}$.

There is another simple characterization of the c-superdifferential, which is more closely related to our earlier definition of subdifferential.

Proposition 18. *It holds that* $y \in \partial^c\varphi(x)$ *if and only if* $\varphi(x) - c(x, y) \geq \varphi(z) - c(z, y)$, $\forall z \in X$.

Proof. By definition of c-superdifferential and c-transform, we have, respectively

$$\begin{aligned}\varphi(x) &= c(x, y) - \varphi^c(y),\\ \varphi(z) &\leq c(z, y) - \varphi^c(y), \quad \forall z \in X.\end{aligned}$$

By adding up these equations we get the proposition. ■

Remark 2.4.4. It is straightforward to see that the definitions above are generalizations of cyclical monotonicity, concavity and superdifferential. Indeed, if we consider the case $X = \mathbb{R}^n$, $Y = \mathbb{R}^n$ and $c(x, y) = -\langle x, y\rangle$, then it holds that: (i) a set is c-cyclically

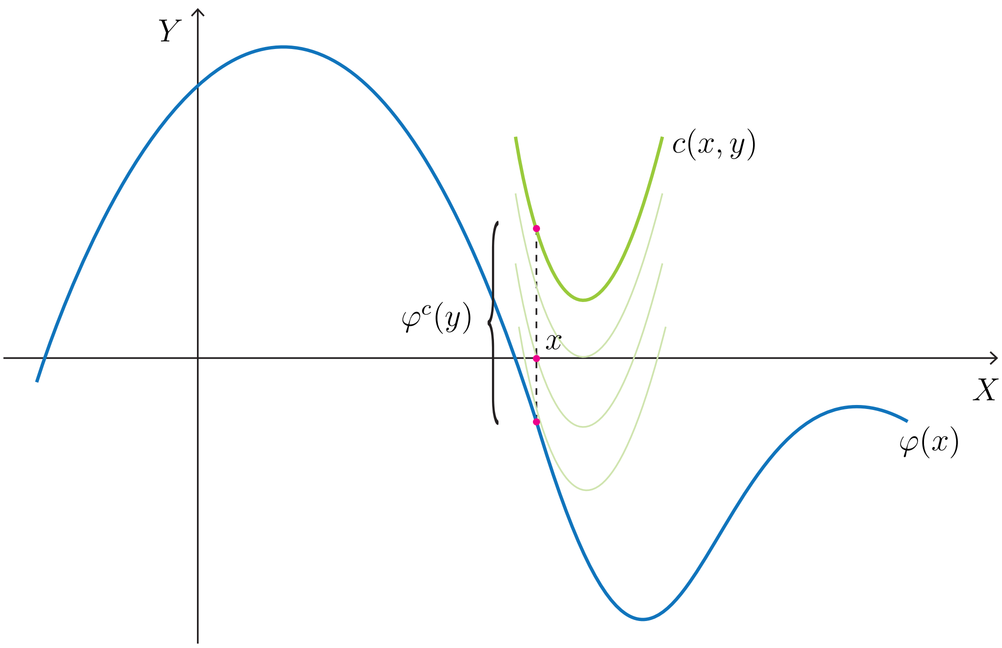

Figure 2.6: *An illustration of* $(x, y) \in \partial^c\varphi$.

monotone if and only it is cyclically monotone; (ii) a function is c-concave if and only if it is concave (and upper semicontinuous); and (iii) the c-superdifferential of a c-concave function is the negative of the classical superdifferential. For this reason taking the quadratic cost in the beginning of this section was useful: we were able to expand the inner product and essentially work with a cost function of the form $c(x, y) = -\langle x, y\rangle$, thereby working with conventional tools of convex analysis.

An interesting consequence of Proposition 18 is that, just as the subdifferential of a convex function is always a cyclically monotone set, the c-superdifferential of a c-concave function is also always a c-cyclically monotone set. This follows directly from the fact that, if $(x_i, y_i) \in \partial^c\varphi$, then for any permutation σ of $\{1, \dots, N\}$,

$$\sum_{i=1}^{N} c(x_i, y_i) = \sum_{i=1}^{N} \varphi(x_i) + \varphi^c(y_i) = \sum_{i=1}^{N} \varphi(x_i) + \varphi^c(y_{\sigma(i)}) \leq \sum_{i=1}^{N} c(x_i, y_{\sigma(i)}).$$

What is more surprising, however, is that the converse holds: *every* c-cyclically monotone set is the c-superdifferential of a c-concave function. Not only that, but one can find conditions for optimality in the Kantorovich Problem through these relations. That is the content of the next theorem, which is the main result of this section.

Theorem 19 (Fundamental Theorem of Optimal Transport)**.** *Let* $c : X \times Y \to \mathbb{R}$ *be a continuous function bounded from below, and* $\mu \in \mathcal{P}(X), \nu \in \mathcal{P}(Y)$ *be such that* $c(x,y) \leq a(x) + b(y)$ *for some* $a \in L^1(\mu), b \in L^1(\nu)$*. Also, let* $\pi \in \Pi(\mu,\nu)$*. Then the following three conditions are equivalent:*

(i) the transference plan π *is optimal;*

(ii) the set $\operatorname{supp}(\pi)$ *is c-cyclically monotone;*

(iii) there exists a c-concave function φ *such that* $\max\{\varphi, 0\} \in L^1(\mu)$ *and* $\operatorname{supp}(\pi) \subset \partial^c\varphi$*.*

Proof. We first note that whatever admissible $\tilde{\pi} \in \Pi(\mu,\nu)$ we take, we can guarantee that $c \in L^1(\tilde{\pi})$, because c is bounded below and, using (2.5),

$$\int c(x,y)d\tilde{\pi}(x,y) \leq \int a(x) + b(y)d\tilde{\pi}(x,y) = \int a(x)d\mu(x) + \int b(y)d\nu(y) < \infty.$$

$(i) \Rightarrow (ii)$ We prove by contradiction: if the set $\operatorname{supp}(\pi)$ is not c-cyclically monotone, there exists an $N \in \mathbb{N}$ and a set of points $\{(x_i,y_i)\}_{i=1}^N \subset \operatorname{supp}(\pi)$, along with a permutation σ of $\{1,\dots,N\}$ such that

$$\sum_{i=1}^N c(x_i,y_i) > \sum_{i=1}^N c(x_i,y_{\sigma(i)}).$$

Since we are taking c to be continuous, we can find neighborhoods U_i and V_i, with $x_i \in U_i$, $y_i \in V_i$, such that it still holds

$$\sum_{i=1}^N c(u_i,v_i) > \sum_{i=1}^N c(u_i,v_{\sigma(i)}),$$

with $u_i \in U_i, v_i \in V_i, 1 \leq i \leq N$.

The idea will be to construct a new measure $\tilde{\pi}$ that contradicts the minimality of π. Intuitively, we will try to create a measure that sends U_i to $V_{\sigma(i)}$ instead of V_i, and in all other respects is just the same as π. We shall take $\tilde{\pi}$ to be a variation of π, so we pick $\tilde{\pi} = \pi + \eta$, and must determine the (signed) measure η. The following three conditions are to be satisfied:

(1) $\eta^- \leq \pi$;

(2) $\text{proj}_{X\#}\,\eta = 0, \text{proj}_{Y\#}\,\eta = 0$;

(3) $\int c d\eta < 0$.

Indeed, the first and second conditions ensure $\tilde{\pi} \in \Pi(\mu, \nu)$ while the third condition defies the optimality of the measure π.

Let $\Omega = \Pi_{i=1}^N U_i \times V_i$, $m_i = \pi(U_i \times V_i)$ and define $P \in \mathcal{P}(\Omega)$ to be the product of the measures $\frac{1}{m_i}\pi\big|_{U_i \times V_i}$. Then, define

$$\eta := \frac{\min_i m_i}{N} \sum_{i=1}^N \left((\text{proj}_{U_i}, \text{proj}_{V_{\sigma(i)}})_\# P - (\text{proj}_{U_i}, \text{proj}_{V_i}) \# P \right),$$

where $(\text{proj}_{U_i}, \text{proj}_{V_i}) : \Omega \to U_i \times V_i$ is the composition of the usual projections. Note $\eta = 0$ for any set outside Ω and, for a given $A \times B \subset U_i \times V_i$,

$$\begin{aligned}
(\text{proj}_{U_i}, \text{proj}_{V_i})_\# P(A \times B) &= P\left(\left(A \cup \bigcup_{j \neq i} U_j \right) \times \left(B \cup \bigcup_{j \neq i} V_j \right) \right) \\
&= 1 \cdots 1 \cdot \frac{1}{\pi(U_i \times V_i)} \pi(A \times B) \cdot 1 \cdots 1 \\
&= \frac{1}{\pi(U_i \times V_i)} \pi(A \times B).
\end{aligned}$$

We must now show η as defined above satisfies (1), (2), (3) above.

(1): Clearly $\eta^- = \frac{\min_i m_i}{N} \sum_{i=1}^N (\text{proj}_{U_i}, \text{proj}_{V_i})_\# P$, and so we would like to prove that

$$\pi - \frac{\min_i m_i}{N} \sum_{i=1}^N (\text{proj}_{U_i}, \text{proj}_{V_i})_\# P \geq 0,$$

and it is sufficient that

$$\frac{1}{N}\pi > \frac{\min_i m_i}{N} (\text{proj}_{U_i}, \text{proj}_{V_i})_\# P, \quad \forall i = 1, \ldots, N.$$

But indeed we have

$$\frac{\min_i m_i}{N} (\text{proj}_{U_i}, \text{proj}_{V_i})_\# P = \frac{\min_i m_i}{N} \cdot \frac{1}{\pi(U_i \times V_i)} \pi\Big|_{U_i \times V_i} < \frac{1}{N}\pi\Big|_{U_i \times V_i}.$$

(2): Using the definitions, and taking $A \subset X, B \subset Y$, we find that

$$\begin{aligned}\tilde{\pi}(A \times Y) &= \pi(A \times Y) + \frac{\min_i m_i}{N} \sum_{i=1}^{N} \big((\mathrm{proj}_{U_i}, \mathrm{proj}_{V_{\sigma(i)}})_{\#} P(A \times Y) \\ &\qquad - (\mathrm{proj}_{U_i}, \mathrm{proj}_{V_i})_{\#} P(A \times Y)\big) \\ &= \mu(A) + \frac{\min_i m_i}{N} \sum_{i=1}^{N} \Big(\frac{1}{\pi(U_i \times V_{\sigma(i)})} \pi\Big|_{U_i \times V_{\sigma(i)}} (A \times Y) \\ &\qquad - \frac{1}{\pi(U_i \times V_i)} \pi\Big|_{U_i \times V_i} (A \times Y)\Big) \\ &= \mu(A) + \frac{\min_{m_i}}{N} \sum_{i=1}^{N} \left(\frac{1}{\mu(U_i)} \mu(A \cap U_i) - \frac{1}{\mu(U_i)} \mu(A \cap U_i)\right) \\ &= \mu(A),\end{aligned}$$

and, similarly,

$$\tilde{\pi}(X \times B) = \nu(B) + \frac{\min_{m_i}}{N} \sum_{i=1}^{N} \left(\frac{1}{\nu(V_{\sigma(i)})} \nu(B \cap V_{\sigma(i)}) - \frac{1}{\nu(V_i)} \nu(B \cap V_i)\right) = \nu(B),$$

so the marginals agree.

(3): By the continuity of c, choosing appropriately small neighborhoods U_i, V_i, there exists a sufficiently small $\varepsilon > 0$ such that

$$\begin{aligned}\int_{X \times Y} c d\pi - \int_{X \times Y} c d\tilde{\pi} &= \frac{\min_i m_i}{N} \sum_{i=1}^{N} \Big(\int c d(\mathrm{proj}_{U_i}, \mathrm{proj}_{V_i})_{\#} P \\ &\qquad - \int c d(\mathrm{proj}_{U_i}, \mathrm{proj}_{V_{\sigma(i)}})_{\#} P\Big) \\ &\geq \frac{\min_i m_i}{N} \Big(\sum_{i=1}^{N} (c(x_i, y_i) - \varepsilon) \int d(\mathrm{proj}_{U_i}, \mathrm{proj}_{V_i})_{\#} P \\ &\qquad - (c(x_i, y_{\sigma(i)}) + \varepsilon) \int d(\mathrm{proj}_{U_i}, \mathrm{proj}_{V_{\sigma(i)}})_{\#} P\Big) \\ &= \frac{\min_i m_i}{N} \sum_{i=1}^{N} (c(x_i, y_i) - \varepsilon) - (c(x_i, y_{\sigma(i)}) + \varepsilon) \\ &= \frac{\min_i m_i}{N} \left(\sum_{i=1}^{N} \big(c(x_i, y_i) - c(x_i, y_{\sigma(i)})\big) - 2N\varepsilon \right) > 0,\end{aligned}$$

which contradicts the optimality of π.

(ii) ⇒ *(iii)* Assuming there exists a set $\Gamma \subset X \times Y$ that is c-cyclically monotone, we need to find a c-concave function φ such that $\Gamma \subset \partial^c\varphi$, with $\max\{\varphi, 0\} \in L^1(\mu)$. If we fix $(\bar{x}, \bar{y}) \in \Gamma$, we must have, for any $(x_i, y_i) \in \Gamma, i = 1, \ldots, N$,

$$\begin{aligned}
\varphi(x) &\leq c(x, y_1) - \varphi^c(y_1) = c(x, y_1) - c(x_1, y_1) + \varphi(x_1) \\
&\leq (c(x, y_1) - c(x_1, y_1)) + c(x_1, y_2) - \varphi^c(y_2) \\
&= (c(x, y_1) - c(x_1, y_1)) + (c(x_1, y_2) - c(x_2, y_2)) + \varphi(x_2) \\
&\vdots \\
&\leq (c(x, y_1) - c(x_1, y_1)) + (c(x_1, y_2) - c(x_2, y_2)) + \cdots + (c(x_N, \bar{y}) - c(\bar{x}, \bar{y})) + \varphi(\bar{x}).
\end{aligned}$$

Since this is our only restriction on φ, and in analogy to Theorem 11, it is reasonable to define

$$\varphi(x) := \inf_{N \in \mathbb{N}} \inf_{\{(x_i, y_i)\}_{i=1}^N \in \Gamma} (c(x, y_1) - c(x_1, y_1)) + \cdots + (c(x_N, \bar{y}) - c(\bar{x}, \bar{y})) \,.$$

Note that this way we have implicitly defined $\varphi(\bar{x}) = 0$ since, on one side, taking $N = 1$ and choosing $(x_1, y_1) = (\bar{x}, \bar{y})$ we get $\varphi(\bar{x}) \leq 0$; on the other, because Γ is c-cyclical monotone, $\varphi(\bar{x}) \geq 0$.

Having defined φ, we now show it is c-concave, $\max\{\varphi, 0\} \in L^1(\mu)$, and then that $\Gamma \subset \partial^c\varphi$. For c-concavity, simply rewrite φ as

$$\varphi(x) = \inf_{y \in Y} \inf_{N \in \mathbb{N}} \inf_{\{(x_i, y_i)\}_{i=1}^N \in \Gamma} (c(x, y) - c(x_1, y)) + \cdots + (c(x_N, \bar{y}) - c(\bar{x}, \bar{y})) \,,$$

and thus, defining

$$\zeta(y) = -\left(\inf_{N \in \mathbb{N}} \inf_{\{(x_i, y_i)\}_{i=1}^N \in \Gamma} -c(x_1, y) + (c(x_1, y_2) - c(x_2, y_2)) + \cdots + (c(x_N, \bar{y}) - c(\bar{x}, \bar{y})) \right),$$

we get

$$\varphi(x) = \inf_{y \in Y} c(x, y) - \zeta(y) = \zeta^c(x),$$

so φ is c-concave.

To see $\max\{\varphi, 0\} \in L^1(\mu)$, pick $N = 1$, $(x_1, y_1) = (\bar{x}, \bar{y})$ and since we assumed $c(x, y) \leq a(x) + b(y)$ with $a \in L^1(\mu)$, using Proposition 18, we get

$$\varphi(x) \leq c(x, \bar{y}) - c(\bar{x}, \bar{y}) < a(x) + b(\bar{y}) - c(\bar{x}, \bar{y}),$$

which proves $\max\{\varphi, 0\} \in L^1(\mu)$.

Finally, we have to check that $\Gamma \subset \partial^c\varphi$. Take any $(\tilde{x}, \tilde{y}) \in \Gamma$, and let $(x_1, y_1) = (\tilde{x}, \tilde{y})$. By definition of φ, it must hold that

$$\begin{aligned}\varphi(x) &\leq c(x, \tilde{y}) - c(\tilde{x}, \tilde{y}) + \inf_{N\in\mathbb{N}} \inf_{\{(x_i,y_i)\}_{i=1}^N \in\Gamma} (c(\tilde{x}, y_2) - c(x_2, y_2)) + \cdots + (c(x_N, \bar{y}) - c(\bar{x}, \bar{y})) \\ &= c(x, \tilde{y}) - c(\tilde{x}, \tilde{y}) + \varphi(\tilde{x}),\end{aligned}$$

which, by Proposition 18, implies that $(\tilde{x}, \tilde{y}) \in \partial^c\varphi$.

$(iii) \Rightarrow (i)$ Take any admissible transport plan $\tilde{\pi} \in \Pi(\mu, \nu)$, we must show $\int c d\pi \leq \int c d\tilde{\pi}$. By the definition of c-transform, we have

$$\varphi(x) + \varphi^c(y) \leq c(x, y), \quad \forall x \in X, \forall y \in Y,$$

and by the definition of c-superdifferential we have

$$\varphi(x) + \varphi^c(y) = c(x, y), \quad \forall (x, y) \in \mathrm{supp}(\pi).$$

Thus, using (2.5),

$$\begin{aligned}\int c(x, y) d\pi(x, y) &= \int \varphi(x) + \varphi^c(y) d\pi(x, y) = \int \varphi(x) d\mu + \int \varphi^c(y) d\nu(y) \\ &= \int \varphi(x) + \varphi^c(y) d\tilde{\pi} \leq \int c(x, y) d\tilde{\pi}(x, y),\end{aligned}$$

and the theorem is established. ■

Several remarks are in place.

Remark 2.4.5. Note that the proof above relies on the continuity hypothesis of the cost function. The theorem is still valid (though it becomes more cumbersome to prove it) if we take the cost function to be only lower semi-continuous (and bounded below). This is done, as before, by considering a sequence of continuous bounded functions approximating it.

Remark 2.4.6. Though the proof above was developed for optimal transference plans, it still holds true for optimal transference maps. That is, assume $T : X \to Y$ is a map such that there exists a c-concave function φ with $T(x) \in \partial^c \varphi(x)$ for all $x \in X$. Then, if $\mu \in \mathcal{P}(X)$ is such that the measure $T_\# \mu$ satisfies the necessary condition $c(x,y) \leq a(x) + b(y)$, for some $a \in L^1(\mu), b \in L^1(T_\# \mu)$, then T must be an optimal map between μ and $T_\# \mu$. This is simply a restatement of the fact that if a map induces the optimal transference plan, then the map itself must be an optimal transference map.

Remark 2.4.7. In what sense is Theorem 19 "fundamental"?

First, it enables us to solve, at least in some cases, the Monge Problem. Since the support of an optimal measure is contained in the c-superdifferential of a c-concave function, if the c-superdifferential is single-valued, then the optimal measure induces map (which must therefore solve the Monge Problem). It is not hard to show that if a function is differentiable at a point, then its c-superdifferential must contain solely the gradient vector. This gives at least a partial answer to the question "when will the Monge Problem admit a solution?" and also provides effective candidates for the Monge Problem: they are gradients of convex functions (this result is generally known as Brenier's Theorem, [10]). Indeed, to show whether we can find a solution to the Monge Problem it then suffices to study how the non-differentiability points of the c-concave function φ behave (this is usually done through some form of Rademacher's Theorem (see, for example, [32, p. 162]).

Second, it already hints at a duality result. Indeed, note that associated to each Kantorovich Problem there seems to be a c-concave function φ. By exploring this relationship further, we will see that the problem of finding an optimal π can be translated to a problem of finding an optimal function φ. This is the main result of Theorem 22 below.

One straightforward consequence of Theorem 19 is that optimality is inherited by restriction, i.e. the restriction of an optimal measure to a certain subset must be optimal in these subsets. Indeed, if π_R is a restriction of $\pi^* \in \Pi^*(\mu, \nu)$ to a certain (measurable) subset of $X \times Y$, then $\mathrm{supp}(\pi_R) \subset \mathrm{supp}(\pi^*)$, so it is contained in the the c-superdifferential of a c-concave function, and thus it is optimal (among its marginals).

The intuitive idea is that if a transport plan is optimal, it must be optimal between its parts because if we could improve the transport in a certain subset, we would be able to improve it when considering the whole space.

Another, perhaps more surprising, corollary is the following.

Corollary 20. *Under the hypotheses of the Fundamental Theorem of Optimal Transport, let* $\pi^* \in \Pi^*(\mu,\nu)$ *be an optimal transference plan, so that there is a c-concave function* φ *such that* $\mathrm{supp}(\pi^*) \subset \partial^c\varphi$*. Then, for any other optimal plan* π'*, it must hold that* $\mathrm{supp}(\pi') \subset \partial^c\varphi$*.*

Proof. First, we check that $\max\{\varphi^c, 0\} \in L^1(\nu)$. Notice that if $(\tilde{x},\tilde{y}) \in \partial^c\varphi$ then, by Proposition 18, it holds that

$$\varphi(\tilde{x}) - c(\tilde{x},\tilde{y}) \geq \varphi(x) - c(x,\tilde{y}),\ \forall x \in X.$$

In particular, by the way we defined φ, we can pick an $\bar{x}$ such that $\varphi(\bar{x}) = 0$, and then $-\varphi(\tilde{x}) \leq c(\bar{x},\tilde{y}) - c(\tilde{x},\tilde{y})$, and so

$$\varphi^c(y) \leq c(\tilde{x},y) - \varphi(\tilde{x}) \leq c(\tilde{x},y) + c(\bar{x},\tilde{y}) - c(\tilde{x},\tilde{y}) \leq c(\tilde{x},y) + M,$$

where we have renamed $M = c(\bar{x},\tilde{y}) - c(\tilde{x},\tilde{y}) \in \mathbb{R}$. By hypothesis, $c(x,y) \leq a(x) + b(y)$, with $a \in L^1(\mu), b \in L^1(\nu)$ so we find that $\varphi^c(y) \leq a(\tilde{x}) + b(y) + M$. Thus, we get $\max\{\varphi^c, 0\} \in L^1(\nu)$.

Hence, for any optimal $\pi' \in \Pi^*(\mu,\nu)$,

$$\begin{aligned}\int \varphi d\mu + \int \varphi^c d\nu = \int \varphi(x) + \varphi^c(y) d\pi'(x,y) \leq \int c(x,y) d\pi'(x,y) = \int c(x,y) d\pi^*(x,y) \\ = \int \varphi(x) + \varphi^c(y) d\pi^*(x,y) = \int \varphi d\mu + \int \varphi^c d\nu.\end{aligned}$$

As the first and last expressions above are the same, the inequality must be an equality, and that is true if and only if $(x,y) \in \partial^c\varphi$, π'-a.e. By the continuity of c, this must be true for all $(x,y) \in \mathrm{supp}(\pi')$, not only almost everywhere. Hence $\mathrm{supp}(\pi') \subset \partial^c\varphi$. ∎

Lastly, let us discuss the question of stability of optimality. That is, given a sequence $\{\mu_k\}_{k\in\mathbb{N}} \subset \mathcal{P}(X)$ narrowly converging to $\mu \in \mathcal{P}(X)$, and another sequence $\{\nu_k\}_{k\in\mathbb{N}} \subset \mathcal{P}(Y)$ narrowly converging to $\nu \in \mathcal{P}(Y)$, what can we say about a sequence of measures made up of $\pi_k \in \Pi^*(\mu_k, \nu_k)$? Is it necessarily a convergent sequence? Is the limit optimal between its marginals?

Proposition 21. *Let X and Y be Polish spaces, and let $c : X \times Y \to \mathbb{R}$ be a cost function bounded below. Given two narrowly convergent sequence of measures, $\mu_k \to \mu, \nu_k \to \nu$, with $\mu_k \in \mathcal{P}(X), \nu_k \in \mathcal{P}(Y), \forall k$, a sequence of optimal maps $\{\pi_k\}_{k\in\mathbb{N}}$ with $\pi_k \in \Pi^*(\mu_k, \nu_k)$, then it must be that, up to a subsequence, $\pi_k \to \pi$ narrowly, and $\pi \in \Pi^*(\mu, \nu)$.*

Proof. Since $\{\mu_k\}$ and $\{\nu_k\}$ are convergent sequences, by Prokhorov's Theorem, the sets $S_1 = \{\mu_k\}$ and $S_2 = \{\nu_k\}$ must be tight. Hence, for any $\varepsilon > 0$, there exist compact sets $K_1, K_2 \subset X$ such that

$$\mu_k(X \setminus K_1) \leq \varepsilon/2 \text{ and } \nu(Y \setminus K_2) \leq \varepsilon/2,\ \forall k.$$

Then, for any element of the set $\pi \in T = \{\pi \in \mathcal{P}(X \times Y : \text{proj}_{X\#}\, \pi \in S_1, \text{proj}_{Y\#}\, \pi \in S_2\}$, we must have

$$\pi(X \times Y \setminus K_1 \times K_2) \leq \mu_k(X \setminus K_1) + \nu_k(Y \setminus K_2) \leq \varepsilon.$$

Hence, by definition, the set T is tight. Thus, given a sequence $\pi_k \in \Pi^*(\mu_k, \nu_k)$ it must be possible to extract a subsequence (still denoted by $\{\pi_k\}$, for simplicity) such that $\pi_k \to \pi$ narrowly, and $\pi \in \Pi(\mu, \nu)$.

We must still prove the optimality of π, i.e. that $\pi \in \Pi^*(\mu, \nu)$. By Theorem 19, we know that the optimality of π_k is equivalent to c-cyclical monotonicity of its support. Fix any $N \in \mathbb{N}$ and take $(x_i, y_i) \in \text{supp}(\pi), i = 1, \ldots, N$. Because $\pi_k \to \pi$ narrowly, there must exist points $(x_i^k, y_i^k) \in \text{supp}(\pi_k), i = 1, \ldots, N$ such that $(x_i^k, y_i^k) \to (x_i, y_i)$ as $k \to \infty$. The c-cyclical monotonicity of $\text{supp}(\pi_k)$ implies

$$\sum_{i=1}^{N} c(x_i^k, y_i^k) \leq \sum_{i=1}^{N} c(x_i^k, y_{i+1}^k),$$

with the usual convention $y_{N+1} = y_1$, and then the c-cyclical monotonicity of $\text{supp}(\pi)$ follows from the continuity of c, which implies $\pi \in \Pi^*(\mu, \nu)$. ∎

2.5
Duality

As we pointed out, the main reason why Kantorovich generalized Monge's Problem was not to merely guarantee the existence of solutions. Kantorovich is considered to be one of the fathers of Linear Programming, where he established very useful duality results. The point of extending Monge's Problem was so that he could generalize his duality theory and prove results in the more general setting usually considered in Optimal Transport. Indeed, this duality formulation is fundamental in several ways, as we shall see later.

What exactly do we mean by duality? It is possible to rewrite Kantorovich's Problem, into a related *dual problem*, so that if we can solve the dual problem then we will have solved the original problem, and vice-versa. Kantorovich's Problem is usually stated as a constrained minimization problem, while its dual is a constrained maximization. In many cases, solving the dual problem is easier than solving the original one, and this allows for novel techniques to be tried.

The main goal of this section is to prove the two following formulations are equivalent, in the sense that the values of (2.22) and (2.23) below are the same.

Problem 1 (Kantorovich Problem). *Let $\mathcal{P}(X), \mathcal{P}(Y)$ be Borel probability spaces on the Polish spaces X and Y, respectively, and let $c : X \times Y \to \mathbb{R}$ be a continuous function, bounded below. Given $\mu \in \mathcal{P}(X)$ and $\nu \in \mathcal{P}(Y)$ we would like to find $\pi \in \mathcal{P}(X \times Y)$ so as to solve*

$$\inf_{\pi} \left\{ \int_{X \times Y} c(x, y) d\pi(x, y) \mid \pi \in \Pi(\mu, \nu) \right\}, \tag{2.22}$$

where $\Pi(\mu, \nu) = \left\{ \pi \in \mathcal{P}(X \times Y) : \operatorname{proj}_{X\#} \pi = \mu, \operatorname{proj}_{Y\#} \pi = \nu \right\}$.

Problem 2 (Dual Problem). *Let $\mathcal{P}(X), \mathcal{P}(Y)$ be Borel probability spaces on the Polish spaces X and Y, respectively, and let $c : X \times Y \to \mathbb{R}$ be a continuous function, bounded below. Given $\mu \in \mathcal{P}(X)$ and $\nu \in \mathcal{P}(Y)$ we would like to find functions $\varphi \in L^1(\mu)$ and $\psi \in L^1(\nu)$ so as to attain*

$$\sup_{\varphi, \psi} \left\{ \int_X \varphi(x) d\mu(x) + \int_Y \psi(y) d\nu(y) \mid \varphi(x) + \psi(y) \leq c(x, y), \ \forall x \in X, y \in Y \right\}. \tag{2.23}$$

As we have noted before, the Kantorovich's Problem is really a linear minimization problem with convex constraints. We want to minimize the functional $\pi \mapsto \int_{X\times Y} c(x,y)d\pi(x,y)$, subject to the constraints $\text{proj}_{X\#}\pi = \mu$, $\text{proj}_{Y\#}\pi = \nu$, and $\pi \geq 0$. Problems of this type generally admit a dual formulation, a fact well-known in Linear Programming. With this in mind, one can try to somewhat reproduce the proof established for duality in finite dimensions. And, by finding a suitable 'minimax theorem', this is possible. To give some intuition, below we give a sketch as to how one might go about in this direction, but we shall bypass some details. We will rigorously prove the theorem afterwards, using the Fundamental Theorem of Optimal Transport we have just proved, so as to make the proof more succinct.

To begin with, notice that

$$\inf_{\pi\in\Pi(\mu,\nu)} \int_{X\times Y} c(x,y)d\pi(x,y) = \inf_{\pi\in M_+(X\times Y)} \left(\int_{X\times Y} c(x,y)d\pi(x,y) + \begin{Bmatrix} 0 \text{ if } \pi \in \Pi(\mu,\nu) \\ \infty \text{ else} \end{Bmatrix} \right),$$

where $M_+(X \times Y)$ is the set of non-negative Borel measures on $X \times Y$. Because the constraints defining $\Pi(\mu,\nu)$ in (2.5) are linear, it is possible to rewrite the indicator function in brackets as a supremum of linear functionals:

$$\begin{Bmatrix} 0 \text{ if } \pi \in \Pi(\mu,\nu) \\ \infty \text{ else} \end{Bmatrix} = \sup_{(\varphi,\psi)\in C_b(X)\times C_b(Y)} \int \varphi d\mu + \int \psi d\nu - \int [\varphi(x) + \psi(y)]d\pi(x,y),$$

since if $\pi \in \Pi(\mu,\nu)$, then (2.5) holds for $(\varphi,\psi) \in C_b(X) \times C_b(Y)$ (because X and Y are Polish), and so the right hand side is zero; on the other hand if there are (φ,ψ) such that the right hand side is not zero, then we can multiply the functions by real numbers so as to render the above supremum infinite.

Hence, we get

$$\begin{aligned} &\inf_{\pi\in\Pi(\mu,\nu)} \int_{X\times Y} c(x,y)d\pi(x,y) = \\ &\qquad = \inf_{\pi\in M_+(X\times Y)} \sup_{(\varphi,\psi)} \int_{X\times Y} c d\pi + \int_X \varphi d\mu + \int_Y \psi d\nu - \int_{X\times Y} [\varphi(x) + \psi(y)]d\pi(x,y). \end{aligned}$$

If we assume that we can find a minimax theorem to enable us to interchange the 'inf' and 'sup' operators, we can rewrite this as

$$\sup_{(\varphi,\psi)} \inf_{\pi\in M_+(X\times Y)} \int_{X\times Y} c(x,y)d\pi(x,y) + \int_X \varphi d\mu + \int_Y \psi d\nu - \int_{X\times Y} [\varphi(x)+\psi(y)]d\pi(x,y)$$
$$= \sup_{(\varphi,\psi)} \left\{ \int_X \varphi d\mu + \int_Y \psi d\nu - \sup_{\pi\in M_+(X\times Y)} \int_{X\times Y} [\varphi(x)+\psi(y)-c(x,y)]d\pi(x,y) \right\}.$$

Now, consider the integrand of the inner 'sup' above. If the function $\theta(x,y) = \varphi(x) + \psi(y) - c(x,y)$ is positive at some point $(\bar{x},\bar{y})$, then since we are free to choose $\pi \in M_+(X \times Y)$, we just pick $\pi = \alpha\delta_{(\bar{x},\bar{y})}$, with $\alpha \in \mathbb{R}$ and $\delta_{(\bar{x},\bar{y})}$ a Dirac measure, and let $\alpha \to \infty$, obtaining a value of infinity for the supremum. On other hand, if θ is non-positive, then it is clear that the value for the supremum must be obtained by $\pi = 0$.

Hence,

$$\sup_{\pi\in M_+(X\times Y)} \int_{X\times Y} [\varphi+\psi-c(x,y)]d\pi = \left\{ \begin{matrix} 0 \text{ if } (\varphi,\psi) \text{ satisfies the constraints in (2.23)} \\ \infty \text{ else} \end{matrix} \right\},$$

and, finally, we get

$$\inf_{\pi\in\Pi(\mu,\nu)} \int c(x,y)d\pi(x,y) = \sup_{(\varphi,\psi)} \int \varphi(x)d\mu(x) + \int \psi(y)d\nu(y)$$

as we hoped for, where $(\varphi,\psi) \in L^1(\mu) \times L^1(\nu)$ satisfies the constraints in (2.23).

Remark 2.5.1. To rigorously establish this proof, one can use the Fenchel-Rockafellar Theorem, a basic result in Functional Analysis (see [11, p. 15]), as the missing minimax theorem. Still, this theorem only yields the desired result if we take X and Y to be compact. The general result follows from approximation arguments (see [31, p. 26]).

Let us now prove the Duality Theorem, using tools from the Fundamental Theorem of Optimal Transport. Actually, we shall prove something stronger: not only the result holds, but the maximizing pair (φ,ψ) can always be taken to be of the form (φ,φ^c).

Theorem 22 (Duality Theorem). *Let $\mu \in \mathcal{P}(X), \nu \in \mathcal{P}(Y)$ and $c : X \times Y \to \mathbb{R}$ be a continuous cost function, bounded from below. If c satisfies $c(x,y) \le a(x) + b(y)$ for $a \in L^1(\mu)$ and $b \in L^1(\nu)$, then the minimum value in the Kantorovich Problem, 1, is equal to the supremum value in the Dual Problem, 2.*

In addition, the supremum of the dual problem is attained, and the maximizing pair (φ, ψ) is of the form (φ, φ^c), for some c-concave function φ.

Proof. First, let $\pi \in \Pi(\mu, \nu)$ and take $(\varphi, \psi) \in L^1(\mu) \times L^1(\nu)$ such that $\varphi(x) + \psi(y) \leq c(x, y), \forall x \in X, y \in Y$. It is clear that

$$\int c(x, y) d\pi(x, y) \geq \int_{X \times Y} \varphi(x) + \psi(y) d\pi(x, y) = \int \varphi(x) d\mu(x) + \int \psi(y) d\nu(y),$$

and so the minimum in the Kantorovich Problem is at least as great as the supremum in the Dual Problem.

Conversely, let $\pi^* \in \Pi^*(\mu, \nu)$ be an optimal transport plan. By the Fundamental Theorem of Optimal Transport, there exists a c-concave function φ such that $\text{supp}(\pi^*) \subset \partial^c \varphi$, as well as $\max\{\varphi, 0\} \in L^1(\mu)$ and $\max\{\varphi^c, 0\} \in L^1(\nu)$. Because

$$\int c(x, y) d\pi(x, y) = \int \varphi(x) + \varphi^c(y) d\pi(x, y) = \int \varphi(x) d\mu(x) + \int \varphi^c(y) d\nu(y),$$

and we already know $\int c d\pi \in \mathbb{R}$, we get $\varphi \in L^1(\mu)$ and $\varphi^c \in L^1(\nu)$, so that (φ, φ^c) is an admissible pair and establishes that the supremum of the Dual Problem is at least as great as the minimum in the Kantorovich Problem. Hence, (φ, φ^c) is a maximizing couple for the Dual Problem. ■

Remark 2.5.2. The Duality Theorem holds in greater generality than considered here. For instance, as before, we could have taken c to be lower semi-continuous. Also, it is possible to first prove the Duality Theorem and from that derive the Fundamental Theorem of Optimal Transport (see [31]).

An interesting consequence of the Duality Theorem is that if (φ, ψ) is a maximizing couple for the Dual Problem, with φ c-concave, then (φ, φ^c) must also be a maximizing couple. Note the same process could be again applied so that $(\varphi^{cc}, \varphi^c)$ yields a pair that is not worse than (φ, φ^c), $(\varphi^{cc}, \varphi^{ccc})$ yields a pair that is not worse than $(\varphi^{cc}, \varphi^c)$, and so on. However, we know by Proposition 15 and Proposition 16 that the above process must stop: if φ is c-concave, it stops because $\varphi^{cc} = \varphi$, otherwise it stops because $\varphi^{ccc} = \varphi^c$.

Thus, we can restrict our search of maximizers to those of the form (φ, φ^c), with φ c-concave. This motivates the next definition.

Definition 23. A **Kantorovich potential** for the measures μ, ν is a couple of functions (φ, φ^c), with φ c-concave, that maximizes the Dual Problem.

If φ is not c-concave, we can consider $(\varphi^{cc}, \varphi^c)$, so that the first function in the couple is in fact c-concave.

Remark 2.5.3. The way we proved optimality in example 4 above was to consider the supremum in (2.11) and prove it must equal the value of the Kantorovich Problem. Notice this was a particular instance of duality at work: in the example considered, the cost function was a distance on the Polish space $X = Y$, given by $c(x, y) = |x - y|$, so the constraint in the dual reads $\varphi(x) + \psi(y) \leq |x - y|$. But we have seen that in the dual problem we can consider, with no loss of generality, $\varphi(x) = \inf_{y \in X}(|x - y| - \psi(y))$, so φ satisfies the 1-Lipschitz condition. But then

$$\psi(y) = \inf_{x \in X} |x - y| - \varphi(x) = -\varphi(x),$$

since the infimum is actually achieved. Hence, the maximizing pair for the dual must be of the form $(\varphi, -\varphi)$, with $\varphi \in \mathrm{Lip}_1(X)$. This implies the dual problem can be rewritten as in (2.11), which is precisely the condition used to solve the example.

3
The Wasserstein Metric

3.1
Wasserstein Spaces

From our previous chapter, we know that associated to each pair of Borel probability measures μ, ν, there is a unique real number

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

that quantifies in a very precise way how 'hard' it is to transport μ to ν. On this direction, assuming both μ, ν are in the same space $\mathcal{P}(X)$, it is natural to ask: is it possible to use $C(\mu, \nu)$ as a metric for the measures?

Though for general costs $c : X \times X \to \mathbb{R}$ the answer is readily seen to be 'no', if we restrict our possible cost functions to *distances* in $X \times X$, then the answer is 'yes'. This is the starting point for the concept of the Wasserstein metric that shall be investigated in this section. It not only provides a metric for the somewhat complicated space of Borel probability measures, but it also has several interesting and useful properties, as we shall see ahead.

First, since a metric cannot attain the value ∞, it would be hopeless to try to find a metric based on $C(\mu, \nu)$ for the entire space $\mathcal{P}(X)$. Instead, we shall restrict our attention to a slightly smaller space, called the Wasserstein space.

Definition 24. Let (X, d) be a Polish metric space, and take $p \in [1, \infty)$. For any two Borel probability measure μ on X, the **Wasserstein space of order** p is defined to be

$$P_p(X) = \left\{ \mu \in \mathcal{P}(X) : \int_X d(x_0, x)^p d\mu(x) < \infty \right\}, \tag{3.1}$$

where $x_0 \in X$ is arbitrary.

Note that the triangle inequality ensures that the space $P_p(X)$ is completely independent from the choice of x_0. Also, if d is a bounded function, then clearly $P_p(X) = \mathcal{P}(X)$.

We proceed to define a reasonable candidate for a metric in this space.

Definition 25. Let (X, d) be a Polish metric space, and take $p \in [1, \infty)$. For any two measures $\mu, \nu \in P_p(X)$, the **Wasserstein metric of order** p from μ to ν is defined by

$$W_p(\mu, \nu) = \left(\inf_{\pi \in \Pi(\mu,\nu)} \int_{X \times X} d^p(x, y) d\pi(x, y) \right)^{1/p}.$$

Before proving that the Wasserstein metric indeed satisfies the conditions for a metric, let us quickly check that $W_p(\mu, \nu)$ is always finite on $P_p(X)$. Indeed, if $\mu, \nu \in P_p(X)$, then, since $d^p(x, y) \leq 2^{p-1}[d^p(x, x_0) + d^p(x_0, y)]$ and using (2.5) it is clear that $d^p(\cdot, \cdot)$ is π-integrable whenever $d^p(\cdot, x_0)$ is μ-integrable and $d^p(x_0, \cdot)$ is ν-integrable.

Now, in order to prove that $W_p(\mu, \nu)$ satisfies the triangle inequality, we shall need the following basic lemma. It essentially states that it is possible to 'glue' two probability measures with a common marginal together to form a third probability measure.

Lemma 26 (Gluing Lemma)**.** *Let X_1, X_2, X_3 be three Polish spaces, and let $\mu_1 \in \mathcal{P}(X_1), \mu_2 \in \mathcal{P}(X_2), \mu_3 \in \mathcal{P}(X_3)$. Take $\pi_{1,2} \in \Pi(\mu_1, \mu_2)$ and $\pi_{2,3} \in \Pi(\mu_2, \mu_3)$ to be two transference plans with a common marginal μ_2. Then there exists a Borel probability measure $\pi \in \mathcal{P}(X_1 \times X_2 \times X_3)$ with marginals $\pi_{1,2}$ on $X_1 \times X_2$ and $\pi_{2,3}$ on $X_2 \times X_3$.*

Proof. We provide a sketch of the proof. The Disintegration of Measure Theorem (see, for instance, [30, p.23] or [31, p. 209]) states that, if X, Y are Polish spaces, then any probability measure $\pi \in \mathcal{P}(X \times Y)$ can be written as the average of probability measures on $\{x\} \times Y$ with $x \in X$. Particularly, if π has a marginal μ on X, then there exists a measurable application $x \mapsto \pi_{Y|x}$ from X to $\mathcal{P}(Y)$ such that

$$\pi = \int_X (\delta_x \otimes \pi_{Y|x}) d\mu(x),$$

where this identity should be taken to mean that for any measurable set $A \subset X \times Y$,

$$\pi(A) = \int_X (\delta_x \otimes \pi_{Y|x})(A) d\mu(x) = \int_X \pi_{Y|x}(A_x) d\mu(x)$$

with $A_{Y|x} = \{y \in Y : (x, y) \in A\}$.

Now, take $\pi_{1,2} \in \Pi(\mu_1, \mu_2)$ and $\pi_{2,3} \in \Pi(\mu_2, \mu_3)$, and disintegrate them with respect to their marginal μ_2. Then we get two measurable applications $\hat{\pi}_{X_1|x_2}$ and $\hat{\pi}_{X_3|x_2}$, from X_2 to $\mathcal{P}(X_1), \mathcal{P}(X_3)$, respectively, such that

$$\pi_{1,2} = \int_{X_2} \hat{\pi}_{X_1|x_2} \otimes \delta_{x_2} d\mu_2(x_2),$$
$$\pi_{2,3} = \int_{X_2} \delta_{x_2} \otimes \hat{\pi}_{X_3|x_2} d\mu_2(x_2).$$

Finally, we define $\pi \in \mathcal{P}(X_1 \times X_2 \times X_3)$ by

$$\pi = \int_{X_2} (\hat{\pi}_{X_1|x_2} \otimes \delta_{x_2} \otimes \hat{\pi}_{X_3|x_2}) d\mu_2(x_2). \tag{3.2}$$

To check that π has the correct marginals, just take an arbitrary measurable subset $A \times B \subset X_1 \times X_2$ so that

$$\begin{aligned}\pi(A \times B \times X_3) &= \int_{X_2} \hat{\pi}_{X_1|x_2}(A)\ \delta_{x_2}(B)\ \hat{\pi}_{X_3|x_2}(X_3) d\mu_2(x_2) \\ &= \int_{X_2} \hat{\pi}_{X_1|x_2}(A)\ \delta_{x_2}(B) d\mu_2(x_2) = \pi_{1,2}(A \times B).\end{aligned}$$

This establishes that π has the correct marginal $\pi_{1,2}$ on $X_1 \times X_2$. The proof for the other marginal is analogous. ■

Definition 27. Given two measures $\mu_{1,2} \in \Pi(\mu_1, \mu_2)$, $\mu_{2,3} \in \Pi(\mu_2, \mu_3)$ with a common marginal μ_2, and letting π be as in the above lemma, the measure $\pi_{1,3} = (\text{proj}_1, \text{proj}_3)_{\#}\pi \in \Pi(\mu_1, \mu_3)$ will be called the **composition of two measures** $\mu_{1,2}$ and $\mu_{2,3}$, and shall be denoted by

$$\pi_{1,3} = \pi_{2,3} \circ \pi_{1,2}.$$

We are finally in condition to prove that $W_p(\mu, \nu)$ is a distance.

Theorem 28. *For all $p \in [1, \infty)$, given a Polish space X and two measures $\mu, \nu \in P_p(X)$, the Wasserstein metric $W_p(\mu, \nu)$ does indeed define a metric.*

Proof. We have already checked that W_p is finite on $P_p(X)$. It is straightforward to check it is nonnegative and symmetric.

Clearly, $W_p(\mu, \mu) = 0$. Now, consider a pair of probability measures such that $W_p(\mu, \nu) = 0$. We must prove $\mu = \nu$. Take $\pi^* \in \Pi^*$ to be an optimal transport plan between μ and ν. Since $\int_{X\times Y} c(x,y)d\pi^*(x,y) = 0$, the measure π^* must be concentrated on the diagonal of $X \times Y$ (given by $y = x$). This implies, for all $\varphi \in C_b(X)$,

$$\int_X \varphi(x)d\mu(x) = \int_{X\times Y} \varphi(x)d\pi^*(x,y) = \int_{X\times Y} \varphi(y)d\pi^*(x,y) = \int_X \varphi(x)d\nu(x),$$

and so $\mu = \nu$.

Lastly, it remains to prove the triangle inequality. Consider $\mu_1, \mu_2, \mu_3 \in P_p(X)$, and take $\pi_{1,2} \in \Pi^*(\mu_1, \mu_2)$ and $\pi_{2,3} \in \Pi^*(\mu_2, \mu_3)$ to be optimal transfer plans between μ_1, μ_2 and μ_2, μ_3, respectively. Define X_i to be the support of measure μ_i. Let π be a measure such as the one in the Gluing Lemma, in (3.2), and call its marginal on $X_1 \times X_3$ by $\pi_{1,3}$, so that $\pi_{1,3} \in \Pi(\mu_1, \mu_3)$. Then:

$$\begin{aligned}
W_p(\mu_1, \mu_3) &\leq \left(\int_{X_1\times X_3} d^p(x_1, x_3)d\pi_{1,3}(x_1, x_3)\right)^{1/p} \\
&= \left(\int_{X_1\times X_2\times X_3} d^p(x_1, x_3)d\pi(x_1, x_2, x_3)\right)^{1/p} \\
&\leq \left(\int_{X_1\times X_2\times X_3} [d(x_1, x_2) + d(x_2, x_3)]^p d\pi(x_1, x_2, x_3)\right)^{1/p} \\
&\leq \left(\int_{X_1\times X_2\times X_3} d^p(x_1, x_2)d\pi(x_1, x_2, x_3)\right)^{1/p} \\
&\qquad\qquad + \left(\int_{X_1\times X_2\times X_3} d^p(x_2, x_3)d\pi(x_1, x_2, x_3)\right)^{1/p} \\
&= \left(\int_{X_1\times X_2} d^p(x_1, x_2)d\pi_{1,2}(x_1, x_2)\right)^{1/p} \\
&\qquad\qquad + \left(\int_{X_2\times X_3} d^p(x_2, x_3)d\pi_{2,3}(x_2, x_3)\right)^{1/p} \\
&= W_p(\mu_1, \mu_2) + W_p(\mu_2, \mu_3).
\end{aligned}$$

Note the first inequality follows because $\pi_{1,3}$ is not guaranteed to be optimal; the first equality follows because $\pi_{1,3}$ is just the marginal of π on $X_1 \times X_3$; the second inequality follows from the usual triangle inequality of d; and the third inequality follows from the Minkowski's inequality for $L^p(X \times X \times X, \pi)$ functions. This establishes that W_p is a metric, and finishes the theorem. ■

Let us explore a few examples of the Wasserstein metric in order to familiarize ourselves with the concept.

Example 29 (Trivial)**.** It is readily seen that $W_p(\delta_x, \delta_y) = d(x, y)$ for any $p \in [1, \infty)$.

Example 30. If $X = \mathbb{R}$ (or, more generally, a Hilbert space), μ a Borel probability measure, and $a \in X$, then

$$W_2^2(\mu, \delta_a) = \int_X |x - a|^2 d\mu(x).$$

We can then find the mean of μ, defined to be $m_\mu = \int_X x d\mu(x)$, as the solution to the optimization problem $\inf_{a \in X} W_2^2(\mu, \delta_a)$. Also, the variance of μ, defined to be $\int_X |x - m_\mu|^2 d\mu(x)$, is the minimum value obtained.

Remark 3.1.1. An important consequence of the way the Wasserstein metric is defined is that it can be ordered in the following sense: if $1 \leq p \leq q$, then $W_p(\mu, \nu) \leq W_q(\mu, \nu)$ for any μ, ν. Indeed, given $1 \leq p \leq q$ and any $\pi \in \Pi(\mu, \nu)$, a simple application of Hölder's inequality yields $\left(\int_{X \times Y} d^p(x, y) d\pi(x, y)\right)^{1/p} \leq \left(\int_{X \times Y} d^q(x, y) d\pi(x, y)\right)^{1/q}$. It thus follows that $W_p(\mu, \nu) \leq W_q(\mu, \nu)$.

When we take the cost function of the Optimal Transport problem to be a metric d, it is possible to use the Duality Theorem to find an alternative formula for the 1-Wasserstein Metric. The following theorem is simply a generalization of remark 2.5.3.

Theorem 31 (Kantorovich-Rubinstein)**.** *Let X be a Polish space, and d a metric on X. Take* Lip(X) *to be the space of all Lipschitz functions on X, and define*

$$||\varphi||_{\text{Lip}} = \sup_{x \neq y} \frac{|\varphi(x) - \varphi(y)|}{d(x, y)}.$$

If $c(x, y) = d(x, y)$, then

$$W_1(\mu, \nu) = \inf_{\pi \in \Pi(\mu, \nu)} \left\{ \int_{X \times X} d(x, y) d\pi(x, y) \right\} = \sup_{\varphi \in L^1(d|\mu - \nu|)} \left\{ \int_X \varphi d(\mu - \nu) \mid ||\varphi||_{\text{Lip}} \leq 1 \right\},$$

where we define $|\mu - \nu| = (\mu - \nu)_+ + (\mu - \nu)_-$.

Proof. First, notice that all Lipschitz functions φ are integrable with respect to μ, ν, so that the right hand side of the equation above makes sense. This holds because, without loss of generality, we can take d to be bounded (otherwise replace it by a sequence of bounded distances $d_n = d/(1+n^{-1}d)$, which satisfy $d_n \leq d$ and $d_n(x,y) \to d(x,y)$ as $n \to \infty$), and so any Lipschitz function will be bounded, thus integrable.

Now, by the Duality Theorem, we now the following holds:

$$\inf_{\pi \in \Pi(\mu,\nu)} \left\{ \int_{X \times X} d(x,y) d\pi(x,y) \right\} = \sup_{\varphi,\psi} \left\{ \int \varphi d\mu + \int \psi d\nu \mid \varphi(x) + \psi(y) \leq d(x,y) \right\}$$
$$= \sup_{\varphi} \left\{ \int_X \varphi(x) d\mu(x) + \int_X \varphi^c(y) d\nu(y) \right\}.$$

Also, our definition for c-concave functions gives $\varphi^c(y) = \inf_{x \in X} d(x,y) - \varphi(x)$ and $\varphi^{cc}(x) = \inf_{y \in X} d(x,y) - \varphi^c(y)$. Since φ^c is 1-Lipschitz (because it is the infimum of 1-Lipschitz functions, bounded from below) we get

$$-\varphi^c(x) \leq \inf_{y \in X} d(x,y) - \varphi^c(y) \leq -\varphi^c(x), \tag{3.3}$$

where the first inequality follows because of the 1-Lipschitz property, and the second by considering $x = y$ in the infimum. This shows that $\varphi^c(x) = -\varphi^{cc}(x) = -\varphi(x)$. Thus

$$W_1(\mu,\nu) = \sup_{\varphi \in L^1(d|\mu-\nu|)} \left\{ \int_X \varphi d\mu - \int_X \varphi d\nu \mid ||\varphi||_{\text{Lip}} \leq 1 \right\},$$

which establishes the theorem. ■

One might ask where precisely the hypothesis that the cost is given by a metric enters in the above result. The fact is that the inequalities (3.3) are only true if the cost function satisfies the triangle inequality.

The Kantorovich-Rubinstein Theorem is very useful in many situations. We illustrate this with the following example.

Example 32. One of the most common metrics used in probability spaces is the so-called **total variation distance**, defined as

$$||\mu - \nu||_{TV} = \frac{1}{2}|\mu - \nu|(X) = \sup_A |\mu(A) - \nu(A)|,$$

where A is a (Borel) subset of X, and $\mu, \nu \in \mathcal{P}(X)$. In fact, this is just the 1-Wasserstein metric considered with the particular cost $c(x,y) = \chi_{x \neq y}$ (sometimes called the 0-

Wasserstein metric, W_0). Using the Kantorovich-Rubinstein Theorem, we get

$$W_0(\mu,\nu) = \inf_{\pi\in\Pi(\mu,\nu)} \int_{X\times X} \chi_{x\neq y} d\pi(x,y) = \inf_{\pi\in\Pi(\mu,\nu)} \pi(\{x\neq y\}) = \sup_{0\leq f\leq 1} \int_X f d(\mu-\nu).$$

By the Jordan Decomposition Theorem, we can split $(\mu-\nu) = (\mu-\nu)_+ - (\mu-\nu)_-$, where both $(\mu-\nu)_+, (\mu-\nu)_-$ are positive measures and singular to each other. It is then easy to see that

$$\sup_{0\leq f\leq 1} \int_X f d(\mu-\nu) = (\mu-\nu)_+(X) = (\mu-\nu)_-(X) = ||\mu-\nu||_{TV}.$$

Of course, both metrics $W_0(\mu,\nu)$ and $||\mu-\nu||_{TV}$ are essentially measuring the amount of mass that needs to be moved if we are required to leave as much mass in its place as possible. Thus, it is not very surprising that they are, in fact, the same.

3.2
Topological Properties

Having defined the Wasserstein distance, it is only natural to ask what kind of topology it induces. Is it the same as some other known topology, or is it something completely different? In any case, what kind of properties does this topology possess?

Before setting out to answer these questions, however, let us note that it is already expected that the Wasserstein distance has some nice properties. Indeed, it is possible to consider X as a 'subset' of $P_p(X)$ via the map $x \mapsto \delta_x$, and also $W_p(\delta_x, \delta_y) = d(x,y)$, as we noted in example 29. Therefore, there is a straightforward isometric immersion of (X,d) into $(P_p(X), W_p)$, and so 'good' properties of (X,d) are somewhat expected to appear in $(P_p(X), W_p)$.

First, we define a convergence mode suitable in $P_p(X)$, which we shall call narrow convergence in $P_p(X)$, that is essentially the usual narrow convergence plus a moment condition. We then prove that the Wasserstein metric can metrize the narrow convergence in $P_p(X)$. In particular, by tweaking with the distance function, we shall see that the Wasserstein metric can in fact metrize the weak convergence.

We recall that a sequence of measures $\{\mu_k\}_{k\in\mathbb{N}}$ is said to converge narrowly to μ if $\int_X \varphi(x)d\mu_k \to \int_X \varphi(x)d\mu$ for all $\varphi \in C_b(X)$, where $C_b(X)$ denote the space of bounded continuous functions on X.

Definition 33. Let (X, d) be a Polish metric space, and $p \in [1, \infty)$. Take $\{\mu_k\}_{k\in\mathbb{N}}$ to be a sequence of measures in $P_p(X)$ and let μ be another element of $P_p(X)$. Then $\{\mu_k\}$ is said to **converge narrowly in** $P_p(X)$ if, for any $x_0 \in X$, $\mu_k \to \mu$ narrowly and $\int_X d^p(x_0, x)d\mu_k(x) \to \int_X d^p(x_0, x)d\mu(x)$.

We now consider alternative ways to get narrow convergence in $P_p(X)$.

Proposition 34. *Let (X, d) be a Polish metric space, $x_0 \in X$, and $p \in [1, \infty)$. Take $\{\mu_k\}_{k\in\mathbb{N}}$ to be a sequence of measures in $P_p(X)$ and let μ be another element of $P_p(X)$. Then the following are equivalent:*

(i) $\mu_k \to \mu$ narrowly in $P_p(X)$, i.e. $\mu_k \to \mu$ narrowly and $\int_X d^p(x_0, x)d\mu_k(x) \to \int_X d^p(x_0, x)d\mu(x)$;

(ii) $\mu_k \to \mu$ narrowly and $\lim_{R\to\infty} \limsup_{k\to\infty} \int_{d(x_0,x)\geq R} d^p(x_0, x)d\mu_k(x) = 0$;

(iii) For all continuous functions φ with $|\varphi(x)| \leq K(1 + d^p(x_0, x))$, $K \in \mathbb{R}$ a constant, it holds that $\int_X \varphi(x)d\mu_k(x) \to \int_X \varphi(x)d\mu(x)$.

(iv) $\mu_k \to \mu$ narrowly and $\limsup_{k\to\infty} \int_X d^p(x_0, x)d\mu_k(x) \leq \int_X d^p(x_0, x)d\mu(x)$;

Proof. We will prove that (iii) $\to$ (i) $\to$ (ii) $\to$ (iii), and then establish that (i) is equivalent to (iv).

First, it is clear that (iii) implies (i). Let us show that (i) implies (ii). We denote $f \wedge g = \inf\{f, g\}$, so that, by narrow convergence,

$$\int_X [d(x_0, x) \wedge R]^p d\mu_k(x) \xrightarrow{k\to\infty} \int_X [d(x_0, x) \wedge R]^p d\mu(x);$$

also, the Monotone Convergence Theorem implies

$$\lim_{R\to\infty} \int_X [d(x_0, x) \wedge R]^p d\mu_k(x) = \int_X d^p(x_0, x)d\mu(x);$$

and, finally, by the hypothesis in (i),

$$\int_X d^p(x_0, x)d\mu_k(x) \xrightarrow{k\to\infty} \int_X d^p(x_0, x)d\mu(x).$$

Thus, we get that

$$\lim_{R\to\infty}\lim_{k\to\infty}\int_X \left(d^p(x_0,x)-[d(x_0,x)\wedge R]^p\right)d\mu_k(x)=0.$$

To conclude, we just investigate what happens when $d(x_0,x)\geq 2R$. A simple rearrangement yields $d^p(x_0,x)-R^p\geq(1-2^{-p})d^p(x_0,x)$, and then

$$\lim_{R\to\infty}\limsup_{k\to\infty}\int_{d(x_0,x)\geq R} d^p(x_0,x)d\mu_k(x)=0,$$

which is (ii).

Next, we prove that (ii) implies (iii). Take any function φ satisfying $|\varphi(x)|\leq K(1+d^p(x_0,x))$, $K\in\mathbb{R}$, and let $R>1$. We can decompose $\varphi(x)=\varphi_1(x)+\varphi_2(x)$, where $\varphi_1(x)=[\varphi(x)\wedge K(1+R^p)]$, and $\varphi_2(x)=\varphi(x)-\varphi_1(x)$. Then, as $\varphi_2(x)$ is bounded pointwise by $Kd^p(x_0,x)\chi_{d(x_0,x)\geq R}$, we find

$$\begin{aligned}\left|\int_X\varphi(x)d\mu_k(x)-\int_X\varphi(x)d\mu(x)\right|&\leq\left|\int_X\varphi_1(x)d(\mu_k-\mu)\right|\\&+K\int_{d(x_0,x)\geq R}d^p(x_0,x)d\mu_k(x)+K\int_{d(x_0,x)\geq R}d^p(x_0,x)d\mu(x).\end{aligned}$$

Because $\mu_k\to\mu$ narrowly, we get

$$\limsup_{k\to\infty}\left|\int_X\varphi(x)d\mu_k(x)-\int_X\varphi(x)d\mu(x)\right|\leq K\limsup_{k\to\infty}\int_{d(x_0,x)\geq R}d^p(x_0,x)d(\mu_k+\mu).$$

And, as $R\to\infty$, using the added hypothesis in (ii), we see that the right hand side above must go to zero. This establishes (iii).

Finally, let us show that (i) is equivalent to (iv). Because $\mu_k\to\mu$ narrowly, we get

$$\int_X d^p(x_0,x)d\mu(x)=\lim_{R\to\infty}\lim_{k\to\infty}\int_X[d(x_0,x)\wedge R]^p d\mu_k(x)\leq\liminf_{k\to\infty}\int_X d^p(x_0,x)d\mu_k(x).$$

Therefore, the convergence in (i) is true if and only if

$$\limsup_{k\to\infty}\int_X d^p(x_0,x)d\mu_k(x)\leq\int_X d^p(x_0,x)d\mu(x),$$

that is, the inequality in (iv) is true. ■

We now want to prove that weak convergence in $P_p(X)$ is actually the same as convergence in the Wasserstein metric.

For the proof we shall need the following important but technical lemma, that essentially states that Cauchy sequences in W_p are tight; that is, given a sequence $\{\mu_k\}_{k\in\mathbb{N}}$ such that $W_p(\mu_k, \mu_j) \xrightarrow{k,j\to\infty} 0$, then, for all $\varepsilon > 0$, one can always find a compact set $K_\varepsilon \subset X$ such that $\mu_k(X \setminus K_\varepsilon) \leq \varepsilon, \forall k \in \mathbb{N}$.

Lemma 35. *Let X be a Polish space, $p \geq 1$, and take $\{\mu_k\}_{k\in\mathbb{N}}$ to be a Cauchy sequence in $(P_p(X), W_p)$. Then $\{\mu_k\}$ is tight.*

Proof. We shall prove the lemma by adapting a canonical proof of Ulam's Lemma. The idea will be to split the sequence $\{\mu_k\}_{k\in\mathbb{N}}$ into a finite part, which will be tight by Ulam's Lemma, and an infinite part, which can be controlled because it is a Cauchy sequence in W_p. Some work must be done to find a compact S that, given ε, satisfies $\mu_k(X \setminus S) \leq \varepsilon$ for all the measures μ_k in the sequence, out of a compact K that satisfies $\mu_j(X \setminus K) \leq \varepsilon$ for just some finite j.

First, from remark 3.1.1 we know that $W_p(\mu, \nu) \geq W_1(\mu, \nu)$, so that if $\{\mu_k\}_{k\in\mathbb{N}}$ is Cauchy in W_p then it must also be Cauchy in W_1. Thus, given $\varepsilon > 0$, there must exist $N \in \mathbb{N}$ such that $W_1(\mu_k, \mu_N) \leq \varepsilon^2$ for $n \geq N$. As a consequence, for any $k \in \mathbb{N}$, there exists $j \in \{1, \ldots, N\}$ such that $W_1(\mu_k, \mu_j) \leq \varepsilon^2$ (indeed, either $k \geq N$, and so $j = N$ suffices, or $k < N$, in which case just take $j = k$).

Now, the finite family $\{\mu_1, \ldots, \mu_N\}$ is tight by Ulam's Lemma, so there exists a compact set $K \subset X$ such that $\mu_j(K) \geq 1 - \varepsilon$ for $j = 1, \ldots, N$. Because K is compact, it can be covered by a finite subcover, i.e. there exists q points $x_1, \ldots, x_q \in X$ such that $K \subset B(x_1, \varepsilon) \cup \cdots B(x_q, \varepsilon)$. Call $U = \bigcup_{n=1}^{q} B(x_n, \varepsilon)$, so that $\mu_j(U) \geq 1 - \varepsilon$ for any $j = 1, \ldots, N$.

Consider the set $U_\varepsilon = \{x \in X : d(x, U) < \varepsilon\} \subset B(x_1, 2\varepsilon) \cup \cdots \cup B(x_q, 2\varepsilon)$ as an enlargement of ϵ around U, still contained in X. Also, define $\varphi(x) = \max\left\{1 - \frac{d(x,U)}{\varepsilon}, 0\right\}$, and note that φ is $(1/\varepsilon)$-Lipschitz.

Since $\chi_U \leq \varphi(x) \leq \chi_{U_\varepsilon}$, we get that, for any $k \in \mathbb{N}$ and $j \leq N$,

$$\mu_k(U_\varepsilon) \geq \int_X \varphi(x) d\mu_k(x) = \int_X \varphi(x) d\mu_j(x) + \left(\int_X \varphi(x) d\mu_k(x) - \int_X \varphi(x) d\mu_j(x) \right),$$

but, because φ is $(1/\varepsilon)$-Lipschitz, we have, for $\pi \in \Pi(\mu_j, \mu_k)$,

$$\int_X \varphi(x)d\mu_j(x) - \int_X \varphi(x)d\mu_k(x) = \int_{X\times X} (\varphi(x) - \varphi(y))d\pi(x,y)$$
$$\leq \frac{1}{\varepsilon}\int_{X\times X} d(x,y)d\pi(x,y) = \frac{W_1(\mu_k, \mu_j)}{\varepsilon},$$

and so

$$\mu_k(U_\varepsilon) \geq \int_X \varphi(x)d\mu_j(x) - \frac{W_1(\mu_k,\mu_j)}{\varepsilon} \geq \mu_j(U) - \frac{W_1(\mu_k,\mu_j)}{\varepsilon}.$$

Since $\mu_j(U) \geq \mu_j(K) \geq 1-\varepsilon$ if $j \leq N$ and for each $k \in \mathbb{N}$ we can find j such that $W_1(\mu_k, \mu_j) \leq \varepsilon^2$, we have

$$\mu_k(U_\varepsilon) \geq 1 - \varepsilon - \frac{\varepsilon^2}{\varepsilon} = 1 - 2\varepsilon.$$

Thus, from $U_\varepsilon \subset \bigcup_{i=1}^q B(x_i, 2\varepsilon)$, we get

$$\mu_k\left(X \setminus \bigcup_{i=1}^q B(x_i, 2\varepsilon)\right) \leq 2\varepsilon.$$

Note this is almost what we need, except that $\bigcup_{i=1}^q B(x_i, 2\varepsilon)$ might not be compact.

To fix this issue, first replace ε by $\varepsilon 2^{-m-1}$, where m is an integer, and find $q(m)$ points $x_1^m, \ldots, x_{1(m)}^m$ in X in order to have

$$\mu_k\left(X \setminus \bigcup_{i=1}^{q(m)} B(x_i^m, \varepsilon 2^{-m})\right) \leq \varepsilon 2^{-m}$$

for any $k \in \mathbb{N}$.

It suffices to find a compact set S such that $S \subset \bigcup_{i=1}^{q(m)} B(x_i^m, \varepsilon 2^{-m})$ to finish this lemma. To this end, consider

$$S = \bigcap_{m=1}^{\infty} \bigcup_{k=1}^{q(m)} \overline{B(x_k^m, \varepsilon 2^{-m})}.$$

It is easy to see that S must be compact because on the one hand it is totally bounded (that is, S can be covered by a finite number of balls of arbitrarily small radius δ; indeed, just take l so that $2^{-l}\varepsilon < \delta$ and then $\overline{B(x_i, 2^{-l}\varepsilon)} \subset B(x_i, \delta)$); on the hand hand, S is closed (as an infinite intersection of a finite union of closed sets). Because X is complete, it holds that $S = \overline{S}$ is compact.

Finally, because we picked S to satisfy $S \subset \bigcup_{i=1}^{q(m)} B(x_i^m, \varepsilon 2^{-m})$, we have

$$\mu_k(X \setminus S) \leq \sum_{m=1}^{\infty} \mu_k\left(X \setminus \bigcup_{i=1}^{q(m)} B(x_i^m, \varepsilon 2^{-m})\right) \leq \sum_{m=1}^{\infty} \varepsilon 2^{-m} = \varepsilon.$$

Therefore, given $\varepsilon > 0$, we found a compact set S such that $\mu_k(X \setminus S) \leq \varepsilon$ for any $k \in \mathbb{N}$. This means that the Cauchy sequence $\{\mu_k\}_{k\in\mathbb{N}}$ is tight, as we wished to prove. ■

We are finally ready to prove that the Wasserstein metric metrizes the narrow convergence in $P_p(X)$.

Theorem 36. *Let (X, d) be a Polish metric space, and $p \in [1, \infty)$. Then the narrow convergence in $P_p(X)$ is the same as convergence in the Wasserstein metric. That is, given a sequence $\{\mu_k\}_{k\in\mathbb{N}} \subset P_p(X)$ and a measure $\mu \in \mathcal{P}(X)$, then μ_k converges narrowly in $P_p(X)$ if and only if $W_p(\mu_k, \mu) \to 0$.*

Proof. Take $\{\mu_k\}_{k\in\mathbb{N}}$ with $W_p(\mu_k, \mu) \to 0$, and let us prove that $\mu_k \to \mu$ narrowly in $P_p(X)$. We must show two things: (i) $\mu_k \xrightarrow{k\to\infty} \mu$ narrowly; and (ii) $\limsup_{k\to\infty} \int_X d^p(x_0, x) d\mu_k(x) \leq \int_X d^p(x_0, x) d\mu(x)$. By Proposition 34, we will then have the narrow convergence in $P_p(X)$.

To prove (i), we first use Lemma 35 to see that $\{\mu_k\}_{k\in\mathbb{N}}$ is tight, and then, by Prokhorov's Theorem, we get a subsequence $\{\tilde{\mu}_k\}_{k\in\mathbb{N}}$ that is narrowly convergent to a probability measure $\tilde{\mu} \in \mathcal{P}(X)$. Since $S_1 = \{\tilde{\mu}_k\}_{k\in\mathbb{N}}$ and $S_2 = \tilde{\mu}$ are tight, we already know by Proposition 21 that there exists a sequence given by $\tilde{\pi}_k \in \Pi^*(\tilde{\mu}_k, \mu)$ that is narrowly convergent to $\pi \in \Pi^*(\tilde{\mu}, \mu)$. Then, proceeding as in (2.14), take $\{d_n^p\}_{n\in\mathbb{N}}$ to be a sequence of bounded, continuous function converging pointwise to d^p, so

$$\begin{aligned}
\inf_{\tilde{\pi}\in\Pi(\tilde{\mu},\mu)} \int_{X\times X} d^p(x,y) d\tilde{\pi}(x,y) &= \inf_{\tilde{\pi}\in\Pi(\tilde{\mu},\mu)} \lim_{n\to\infty} \int_{X\times X} d_n^p(x,y) d\tilde{\pi} \\
&= \inf_{\tilde{\pi}_k\in\Pi(\tilde{\mu}_k,\mu)} \lim_{n\to\infty} \lim_{k\to\infty} \int_{X\times X} d_n^p(x,y) d\tilde{\pi}_k \\
&\leq \liminf_{k\to\infty} \inf_{\tilde{\pi}_k\in\Pi(\tilde{\mu}_k,\mu)} \int_{X\times X} d^p(x,y) d\tilde{\pi}_k.
\end{aligned}$$

As a consequence, we find that

$$W_p(\tilde{\mu}, \mu) \leq \liminf_{k\to\infty} W_p(\tilde{\mu}_k, \mu) = 0.$$

So it must be that $\tilde{\mu} = \mu$ and therefore we must have $\mu_k \to \mu$ narrowly.

To prove (ii), we first use a classical inequality: for any $\varepsilon > 0$, there exists a constant C such that for $a, b \in \mathbb{R}_+$, it holds

$$(a+b)^p \leq (1+\varepsilon)a^p + Cb^p. \tag{3.4}$$

Using the triangle inequality, as well as the inequality above, we find that, for $x_0, x, y \in X$,

$$d^p(x_0, x) \leq (1+\varepsilon)d^p(x_0, y) + Cd^p(x, y). \tag{3.5}$$

Now, take a sequence $\{\mu_k\}_{k\in\mathbb{N}} \subset P_p(X)$ with $W_p(\mu_k, \mu) \to 0$. Let $\pi_k \in \Pi^*(\mu_k, \mu)$ and integrate (3.5) with respect to π_k to get

$$\int_X d^p(x_0, x) d\mu_k(x) \leq (1+\varepsilon) \int_X d^p(x_0, y) d\mu(y) + C \int_{X\times X} d^p(x, y) d\pi_k(x, y).$$

Since $\int_{X\times X} d^p(x, y) d\pi_k(x, y) = W_p^p(\mu_k, \mu) \to 0$ when $k \to \infty$,

$$\limsup_{k\to\infty} \int_X d^p(x_0, x) d\mu_k(x) \leq (1+\varepsilon) \int_X d^p(x_0, x) d\mu(x).$$

By making $\varepsilon \to 0$, we get property (iv) of Proposition 34. This proves that $\mu_k \to \mu$ narrowly in $P_p(X)$.

Conversely, we assume $\mu_k \to \mu$ narrowly in $P_p(X)$ and must show that $\lim_{k\to\infty} W_p(\mu_k, \mu) = 0$. Take as before $\pi_k \in \Pi^*(\mu_k, \mu)$ so that

$$\int_{X\times X} d^p(x, y) d\pi_k(x, y) \to 0.$$

By Prokhorov's Theorem, the sequence $\{\mu_k\}_{k\in\mathbb{N}}$ is tight, and $\{\mu\}$ is also trivially tight. Then, by Proposition 21, there must exist a subsequence (which we still denote by $\{\pi_k\}$) such that $\pi_k \to \pi$ narrowly in $\mathcal{P}(X \times X)$ as $k \to \infty$. Also, since each π_k is optimal, Proposition 21 tells us that π must be an optimal transport plan with marginals μ and μ. Hence, it clearly must be that $\pi = (\mathrm{Id}, \mathrm{Id})_{\#}\mu$. Because the limit does not depend on the subsequence chosen, π needs to be the limit of the whole sequence $\{\pi_k\}$.

Note that, given $x_0 \in X$ and $R > 0$, we have following inequality:

$$d(x,y) \leq d(x,y) \wedge R + 2d(x,x_0)\chi_{d(x,x_0)\geq R/2} + 2d(x_0,y)\chi_{d(x_0,y)\geq R/2},$$

which is simply stating that if $d(x,y) > R$ then the largest of $d(x,x_0)$ and $d(x_0,y)$ needs to be greater than $R/2$, and not less than $d(x,y)/2$. Then, taking the p-th power and using (3.4), it must hold that

$$d^p(x,y) \leq C_p\left([d(x,y)\wedge R]^p + d^p(x,x_0)\chi_{d(x,x_0)\geq R/2} + d^p(x_0,y)\chi_{d(x_0,y)\geq R/2}\right), \tag{3.6}$$

for some constant $C_p > 0$ that only depends on p.

Finally, let $\pi_k \in \Pi^*(\mu_k,\mu)$. From (3.6) above, with $R \geq 1$, we have

$$\begin{aligned}
W_p^p(\mu_k,\mu) &= \int_{X\times X} d^p(x,y)d\pi_k(x,y) &(3.7)\\
&\leq C_p\int_{X\times X}[d(x,y)\wedge R]^p d\pi_k(x,y) + C_p\int_{d(x,x_0)\geq R/2} d^p(x,x_0)d\pi_k(x,y)\\
&\qquad + C_p\int_{d(x_0,y)\geq R/2} d^p(x_0,y)d\pi_k(x,y) &(3.8)\\
&\leq C_p\int_{X\times X}[d(x,y)\wedge R]^p d\pi_k(x,y) + C_p\int_{d(x,x_0)\geq R/2} d^p(x,x_0)d\mu_k(x)\\
&\qquad + C_p\int_{d(x_0,y)\geq R/2} d^p(x_0,y)d\mu_k(y). &(3.9)
\end{aligned}$$

Since $\pi_k \to \pi$ narrowly and $\pi \in \Pi^*(\mu,\mu)$,

$$\begin{aligned}
\lim_{R\to\infty}\lim_{k\to\infty}\int_{X\times X}[d(x,y)\wedge R]^p d\pi_k(x,y) &= \lim_{R\to\infty}\int_{X\times X}[d(x,y)\wedge R]^p d\pi(x,y)\\
&= \int_{X\times X} d^p(x,y)d\pi_k(x,y),
\end{aligned}$$

and thus, when taking limits in R and k, (3.9) simply becomes

$$\begin{aligned}
\limsup_{k\to\infty} W_p^p(\mu_k,\mu) &\leq \lim_{R\to\infty} C_p \limsup_{k\to\infty}\int_{d(x,x_0)\geq R/2} d^p(x,x_0)d\mu_k(x)\\
&\quad + \lim_{R\to\infty} C_p \limsup_{k\to\infty}\int_{d(x_0,y)\geq R/2} d^p(x_0,y)d\mu_k(y)\\
&= 0,
\end{aligned}$$

where above we have used the equivalence between converging narrowly in $P_p(X)$ and property (ii) in Proposition 3.9. This concludes the theorem. ■

Two important corollaries are straightforward from Theorem 36. The first attests the continuity of W_p, while the second shows that, in fact, the Wasserstein metric can properly metrize the narrow topology in $\mathcal{P}(X)$.

Corollary 37. *Let (X, d) be a Polish metric space, and $p \in [1, \infty)$, then, W_p is continuous on $P_p(X)$. That is, if both $\{\mu_k\}_{k\in\mathbb{N}}$ and $\{\nu_k\}_{k\in\mathbb{N}}$ converge narrowly in $P_p(X)$ to μ and ν as $k \to \infty$, respectively, then*

$$W_p(\mu_k, \nu_k) \xrightarrow{k\to\infty} W_p(\mu, \nu).$$

Proof. Because the sequences $\{\mu_k\}, \{\nu_k\}$ converge narrowly in $P_p(X)$, by Theorem 36, we have that $W_p(\mu_k, \mu) \to 0$ and $W_p(\nu_k, \nu) \to 0$ as $k \to \infty$. By the triangle inequality, we get

$$W_p(\mu, \nu) - (W_p(\mu, \mu_k) + W_p(\nu, \nu_k)) \leq W_p(\mu_k, \nu_k) \leq W_p(\mu, \nu) + (W_p(\mu, \mu_k) + W_p(\nu, \nu_k)),$$

and so $W_p(\mu_k, \nu_k) \to W_p(\mu, \nu)$ as $k \to \infty$. ■

Remark 3.2.1. If, in the theorem above, we only had that $\mu_k \to \mu, \nu_k \to \nu$ narrowly (as opposed to narrowly in $P_p(X)$), then the Wasserstein metric would only be *lower-semicontinuous*. That is, it would only hold that $W_p(\mu, \nu) \leq \liminf_{k\to\infty} W_p(\mu_k, \nu_k)$.

Corollary 38. *Let (X, d) be a Polish space. It is always possible to find a distance $\tilde{d}$ such that the convergence in the Wasserstein sense for the distance d is equivalent to the narrow convergence of probability measures in $\mathcal{P}(X)$.*

Proof. We have already observed that if d is a bounded function then $P_p(X) = \mathcal{P}(X)$ and all the extra conditions in Proposition 34 are trivially satisfied, so that narrow convergence is equivalent to narrow convergence in $P_p(X)$ or, by Theorem 36, convergence in the Wasserstein metric.

To find a bounded distance inducing the same topology as d, consider $\tilde{d} = d/(1 + d)$. ■

Now is a good time to stop and discuss why the Wasserstein metric has any importance. Though it is very appealing that it can metrize the narrow topology, it is certainly not the only to do so. For instance, this holds for both the **Lévy-Prokhorov** distance, defined by

$$d_p(\mu, \nu) = \inf \{\varepsilon > 0 : \mu(B) \leq \nu(B^\varepsilon) + \varepsilon, \text{ for all Borel sets } B\}, \tag{3.10}$$

where $B^\varepsilon = \{x : \inf_{y \in B} d(x, y) \leq \varepsilon\}$; or the **bounded Lipschitz distance**, defined by

$$d_{bL}(\mu, \nu) = \sup \left\{ \int \varphi d\mu - \int \varphi d\nu : ||\varphi||_\infty + ||\varphi||_{Lip} \leq 1 \right\},$$

where $||\varphi||_\infty = \sup |\varphi|$.

Still, there are quite a few reasons why the Wasserstein metric is of interest. First, as we have argued before, there is an isometric embedding of (X, d) in $(P_p(X), W_p)$, that suggests an interesting geometric structure for Wasserstein spaces. Second, there is the direct connection with Optimal Transport that can sometimes lead to interesting insights. Third, and as a consequence of the Optimal Transport theory, we know there is an interesting duality theorem that might lead to a new and rich set of tools. Fourth, from Proposition 21, we get that the Wasserstein distance is reasonably 'stable' under perturbations. Finally, because W_p is defined as an infimum, it is relatively easy to find upper bounds for it.

Let us now make due on our promise that the space $(P_p(X), W_p)$ has a rich structure with the next theorem. We shall prove that if (X, d) is a Polish space, then $(P_p(X), W_p)$ is itself a Polish space, that is, it is a complete and separable space.

Theorem 39. *Let (X, d) be a metric Polish space with $p \in [1, \infty)$; then $(P_p(X), W_p)$ is also a Polish space.*

Proof. We must prove that $(P_p(X), W_p)$ is a complete and separable metric space. We have already proved that it is a metric space in Theorem 28; now we prove separability and completeness.

For separability, we must find a sequence of points $x_n \in X$ such that the countable set of measures $M = \{\sum_{n=1}^N b_n \delta_{x_n} : N \in \mathbb{N}, b_n \in \mathbb{Q}\}$ is dense in $(P_p(X), W_p)$.

Let $\mu \in P_p(X)$ and $\varepsilon > 0$. First, we construct a measure $\mu_1 = \sum_{n=1}^N a_n \delta_{x_n}$, with $a_n \in \mathbb{R}_+$ and $\sum_{n=1}^\infty a_n = 1$, such that $W_p(\mu, \mu_1) \leq \varepsilon$. Since μ_1 does not satisfy all the criteria we need, we then find another measure $\mu_2 \in M$ such that $W_p(\mu_1, \mu_2) \leq 2\varepsilon$, and then, by the triangle inequality, we shall have $W_p(\mu, \mu_2) \leq 3\varepsilon$ for $\mu_2 \in M$, as we wanted.

Because X itself is separable, we can find a sequence of points $\{x_n\}_{n\in\mathbb{N}}$ such that balls of the form $B(x_n, \varepsilon^{\max\{1,1/p\}})$ cover the entire space X. Moreover, we can construct a partition of X by considering the sets

$$\tilde{B}_n = \left(B(x_n, \varepsilon^{\max\{1,1/p\}}) \setminus \bigcup_{k=1}^{n-1} B(x_k, \varepsilon^{\max\{1,1/p\}}) \right) \cap X.$$

The idea will be to concentrate each point of $\tilde{B}_n$ in x_n, and consider how costly the transport is. In this case, take $a_n = \mu(\tilde{B}_n)$, so that $\sum_{n=1}^{\infty} a_n = 1$, and define $\mu_1 = \sum_{n=1}^{\infty} a_n \delta_{x_n}$. The optimal cost of transporting μ to μ_1 is

$$\sum_{n=1}^{\infty} d^p(x, x_n) d\mu(x) \leq \sum_{n=1}^{\infty} a_n \varepsilon^{p\max\{1,1/p\}} = \varepsilon^{\max\{p,1\}}$$

and we get $W_p(\mu, \mu_1)$, as we wanted. Now we need to find a suitable μ_2 to prove separability.

Because $W_p(\mu, \mu_1)$, from (3.1) and property (ii) in Proposition 34, it is clear that $\mu \in P_p(X)$. Since trivially $\delta_{x_1} \in P_p(X)$, it must be that

$$\sum_{n=1}^{\infty} a_n |x_n - x|^p = (W_p(\mu_1, \delta_{x_1}))^{\max\{p,1\}} < \infty.$$

Since the sequence $\sum_{n=1}^{\infty} a_n |x_n - x|^p$ is convergent, there must exist $N \in \mathbb{N}$ so that

$$\sum_{n=N+1}^{\infty} a_n |x_n - x|^p \leq \varepsilon^{\max\{p,1\}}.$$

For $n = 2, \ldots, N$, it is then possible to pick nonnegative rational numbers b_n sufficiently close to a_n to ensure that

$$0 \leq a_n - b_n \leq a_n \frac{\varepsilon^{\max\{p,1\}}}{\left(\sum_{n=1}^{N} a_n |x_n - x_1|^p \right)}$$

with

$$b_1 = a_1 + \sum_{n=2}^{N} (a_n - b_n) + \sum_{n=N+1}^{\infty} a_n.$$

Note that, by the way we defined b_1, it holds that $\sum_{n=1}^{N} b_n = 1$. The idea now is transport μ_1 to $\mu_2 = \sum_{n=1}^{N} b_n \delta_{x_n}$ in the following way: for $n = 1, \ldots N$ keep a b_n mass at each x_n and transport the remaining mass $a_n - b_n$ back to x_1 (keeping the b_n mass in place costs

0, and since $a_n - b_n$ is really small, it should not cost too much); for $n > N$, transport all the mass a_n from x_n to x_1 (because the masses a_n are small, this should not cost much either). In any case, the total cost of transport will be

$$\sum_{n=1}^{N}(a_n - b_n)|x_n - x_1|^p + \sum_{n=N+1}^{\infty} a_n|x_n - x_1|^p \leq 2\varepsilon^{\max\{p,1\}}.$$

Thus, $W_p(\mu_1, \mu_2) \leq 2\varepsilon$ and, finally we obtain, $W_p(\mu, \mu_2) \leq 3\varepsilon$. Since ε is arbitrary, we proved separability.

To show that $(P_p(X), W_p)$ is complete, take $\{\mu_k\}_{k\in\mathbb{N}}$ to be a Cauchy sequence, and let us prove convergence. By Lemma 35, we already know that $\{\mu_k\}$ must be tight and, hence, it has a subsequence (still denoted by $\{\mu_k\}$ for simplicity) that converges narrowly to a measure μ. We must then have, for an arbitrary $x_0 \in X$,

$$\int_X d^p(x_0, x)d\mu(x) \leq \liminf_{k\to\infty} \int_X d^p(x_0, x)d\mu_k(x) < \infty,$$

which shows that $\mu \in P_p(X)$. Now, consider a subsequence of $\{\mu_k\}$, which we shall denote $\{mu_{k'}\}$ Since $\mu_k \to \mu$ narrowly, by remark 3.2.1, we have

$$W_p(\mu, \mu_{k'}) \leq \liminf_{k\to\infty} W_p(\mu_k, \mu_{k'})$$

which yields

$$\limsup_{k'\to\infty} W_p(\mu, \mu_{k'}) \leq \limsup_{k,k'\to\infty} W_p(\mu_k, \mu_{k'}) = 0.$$

Thus, $\mu_{k'} \to \mu$ narrowly in $P_p(X)$. Since $\{\mu_{k'}\}$ is an arbitrary subsequence of $\{\mu_k\}$, and $\{\mu_k\}$ itself is a Cauchy sequence, the entire sequence must be convergent in W_p. ■

3.3 Geometric Properties

We now investigate the geometric structure of the Wasserstein space $(P_p(X), W_p)$. We try to answer two important questions: (i) given two points in $P_p(X)$ what is the shape of the 'shortest path' between the points?; and (ii) how 'curved' is the space $P_p(X)$? The first question leads us to the concept of geodesics, and the second to the notion of curvature. To answer them, we will restrict our attention to the case where the base space X is a separable Hilbert space so as to simplify several proofs, but we shall state the theorem in full generality when suitable.

Geodesics

Why should we care for the behavior of shortest paths in Wasserstein spaces? Linear structures provide very useful aids when trying to understand a space, and several definitions are based upon on it (for instance, convexity). Yet, in arbitrary structures, the key concept of a line segment does not exist, and some other notion must take its place. Since part of the reason why line segments are so important is because they give the shortest path between two points, it seems useful to understand how shortest paths behave in arbitrary metric spaces. Indeed, it is the basic motivation for defining geodesic curves, of which shortest paths are a particular case.

First, we will need some preliminary concepts.

Definition 40. Given a topological space X, a **curve** γ is a continuous map $\gamma : I \to X$, where $I \subset \mathbb{R}$ is an interval.

Of course, two different curves can have the same image if they have different parametrizations. We would like to identify when that is the case.

Definition 41. A curve $\tilde{\gamma} : I \to X$ is said to be a **reparametrization** of another curve $\gamma : J \to X$ if there exists a nondecreasing and continuous function $\theta : I \to J$ such that $\tilde{\gamma} = \gamma \circ \theta$.

For simplicity and without any loss of generality, we shall take $I = [0, 1]$ from now on.

Another important concept regarding curves is its length, which we define next.

Definition 42. Given a metric space (X, d) and a curve $\gamma : [0, 1] \to X$, the **length** of γ is given by

$$\text{len}(\gamma) = \sup_{0=t_0<t_1<\cdots<t_n=1} \sum_{i=1}^{n} d(\gamma(t_i - 1), \gamma(t_i)),$$

where the supremum is taken over all the partitions $0 = t_1 < t_2 < \cdots < t_n = 1$, and $n \in \mathbb{N}$ is arbitrary.

We say a curve is **rectifiable** if $\text{len}(\gamma) < \infty$. In this text, we shall not bother with non-rectifiable curves.

It will be useful to introduce the following notation: we let $\gamma_{[s,t]}$ be the restriction of a curve $\gamma : [0, 1] \to X$ to the interval $[s, t]$, with $0 \leq s \leq t \leq 1$.

Now, in order to avoid working with curves that are simply reparametrizations of each other, let us define a somewhat 'natural' parametrization that will be standard from this point on.

Definition 43. A parametrization of a rectifiable curve $\gamma : [0, 1] \to X$ is called a **constant speed parametrization** if there exists $v > 0$ such that $\text{len}(\gamma_{[s,t]}) = v(t - s)$, for all $0 \leq s \leq t \leq 1$.

To verify that a parametrization is constant speed, one can let $s = a$ be fixed and just prove $\text{len}(\gamma_{[a,t]}) = v(t - a)$ for all $t \in [a, 1]$, since $\text{len}(\gamma_{[s,t]}) = \text{len}(\gamma_{[a,t]}) - \text{len}(\gamma_{[a,s]})$. Indeed, this point of view shows why this parametrization is called 'constant speed': we have that $\frac{d}{dt} \text{len}(\gamma_{[a,t]}) = v$.

Of course, it remains to be shown that every (rectifiable) curve does indeed admit such a constant speed parametrization.

Proposition 44. *Given $v > 0$, any rectifiable curve $\gamma : [0, 1] \to X$ can be rewritten as $\gamma = \tilde{\gamma} \circ \zeta$, where $\tilde{\gamma} : [0, \frac{1}{v} \text{len}(\gamma)] \to X$ is a constant speed parametrization, and $\zeta : [0, 1] \to [0, \frac{1}{v} \text{len}(\gamma)]$ is a nondecreasing continuous function.*

Proof. The idea will be to find a constant speed parametrization $\tilde{\gamma}$ by taking $\zeta : [0, 1] \to [0, \frac{1}{v} \text{len}(\gamma)]$ to be length of the curve $\gamma_{[0,t]}$ divided by v.

First, then, we define $\zeta(t) = \frac{1}{v} \text{len}(\gamma_{[0,t]})$, $\forall t \in [0, 1]$, and it is immediate to see that $\zeta(t)$ is nondecreasing. It is also continuous: for a fixed t and any $\varepsilon > 0$, consider a partition of the interval $[0, t]$ given by $0 = s_0 < s_1 < \cdots < s_{n-1} < s_n = t$ with $d(\gamma(s_{n-1}), \gamma(t)) < v\varepsilon/2$ and such that $\text{len}(\gamma_{[0,t]}) - \sum_{i=1}^{n} d(\gamma(s_{i-1}), \gamma(s_i)) < v\varepsilon/2$. Then we have

$$\text{len}(\gamma_{[s_{n-1},t]}) - d(\gamma(s_{n-1}), \gamma(t)) < \frac{v\varepsilon}{2},$$

and so $\text{len}(\gamma_{[s_{n-1},t]}) < v\varepsilon$. Thus, given $\varepsilon > 0$, for any t' such that $s_{n-1} \leq t' \leq t$ we have that

$$\zeta(t) - \zeta(t') = \frac{1}{v} \left(\text{len}(\gamma_{[0,t]}) - \text{len}(\gamma_{[0,t']}) \right) = \frac{\text{len}(\gamma_{[t',t]})}{v} < \varepsilon.$$

Since $\varepsilon > 0$ was arbitrary, we get the continuity of $\zeta(t)$ from the left. Analogously, one proves it is continuous from the right; hence we have that $\zeta(t)$ is continuous.

Now, we must define $\tilde{\gamma} : [0, \frac{1}{v} \text{len}(\gamma)] \to X$. To do so, pick $\tau \in [0, \frac{1}{v} \text{len}(\gamma)]$ and $t \in [0, 1]$ such that $\tau = \zeta(t)$; since we want $\gamma(t) = \tilde{\gamma} \circ \zeta(t) = \tilde{\gamma}(\tau)$, simply define

$\tilde{\gamma}(\tau) = \gamma(t)$. By construction, it must hold that $\gamma = \tilde{\gamma} \circ \zeta$. We are left to show that $\tilde{\gamma}$ is continuous and has constant speed parametrization.

To see that $\tilde{\gamma}$ is indeed a constant speed parametrization, note that $\tilde{\gamma}$ is just a reparametrization of γ, and so

$$\mathrm{len}(\tilde{\gamma}_{[\tau_1,\tau_2]}) = \mathrm{len}(\gamma_{[t_1,t_2]}) = \mathrm{len}(\gamma_{[0,t_2]}) - \mathrm{len}(\gamma_{[0,t_1]}) = v\tau_2 - v\tau_1 = v(\tau_2 - \tau_1). \tag{3.11}$$

For continuity, take $\tau_1 = \zeta(t_1), \tau_2 = \zeta(t_2)$, so that $\tilde{\gamma}(\tau_1) = \gamma(t_1), \tilde{\gamma}(\tau_2) = \gamma(t_2)$ and the endpoints of $\gamma_{[t_1,t_2]}$ are $\tilde{\gamma}(\tau_1)$ and $\tilde{\gamma}(\tau_2)$. From the triangle inequality we have that $d(\tilde{\gamma}(\tau_1), \tilde{\gamma}(\tau_2)) \leq \mathrm{len}(\tilde{\gamma}_{[\tau_1,\tau_2]})$, and thus using (3.11) above we get that

$$d(\tilde{\gamma}(\tau_1), \tilde{\gamma}(\tau_2)) \leq \mathrm{len}(\gamma_{[t_1,t_2]}) = v(\tau_2 - \tau_1),$$

which proves continuity.

■

Having defined curves, and given them a 'natural' constant speed parametrization, we must now ensure that the space we are working with does indeed admit a shortest path. In this sense, we would like to consider as a metric the infimum of the lengths of all curves connecting two given points.

Definition 45. Given a metric space (X, d), its **intrinsic metric**, $d_I : X \times X \to \mathbb{R}_+$, is given by

$$d_I(x, y) = \inf_{\gamma \in \mathcal{C}_{xy}} \mathrm{len}(\gamma),$$

with $x, y \in X$, where $\mathcal{C}_{xy}$ denote the set of curves in X connecting x to y. If there is no path in X of finite length between x and y, we define $d_I(x, y) = \infty$.

The intrinsic metric is simply a notion of distance based on the lengths of connecting curves. It is not hard to see that (X, d_I) is a metric space as long as $d_I < \infty$. Note that, in general,

$$d_I(x, y) \geq d(x, y), \tag{3.12}$$

since, as we have noted, $\mathrm{len}(\gamma) \geq d(\gamma(0), \gamma(1))$ follows from the triangular inequality. The next example illustrates a case where the inequality may be strict.

Example 46. Consider the circle $S^1 \subset \mathbb{R}^2$, with the euclidean distance $d(x, y) = |x - y|$. Then no matter how two distinct points are connected by a curve, the curve's length needs to be strictly greater than the euclidean distance between the two points. See Figure 3.1.

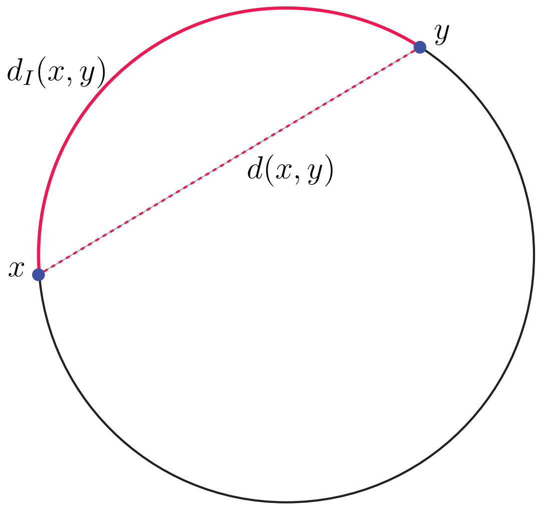

Figure 3.1: *A situation where* $d_I(x, y) > d(x, y)$.

In light of inequality (3.12) above, we would hope that the spaces we work with are such that the distance between two points can be approximated by the length of the curves connecting them. We call these length spaces.

Definition 47. Given a metric space (X, d), we call it a **length space** if, for every $x, y \in X$,

$$d(x, y) = d_I(x, y),$$

where $d_I(x, y)$ denotes the intrinsic metric of the space.

Example 48. From example 46, we know that S^1 (with the euclidean distance) is not a length space, though $\mathbb{R}^d$ (with the euclidean distance) is. More generally, any convex, connected subset of $\mathbb{R}^d$ is a length space, and any non-convex set cannot be one.

Example 49. The set $\mathbb{R}^2 \setminus \{(0, 0)\}$ with the euclidean distance is still a length space. Even if the line segment connecting two points goes through the origin, we can still consider a sequence of curves whose lengths are arbitrarily close to the distance between the two points. See Figure 3.2.

Finally, we define the notion of geodesic.

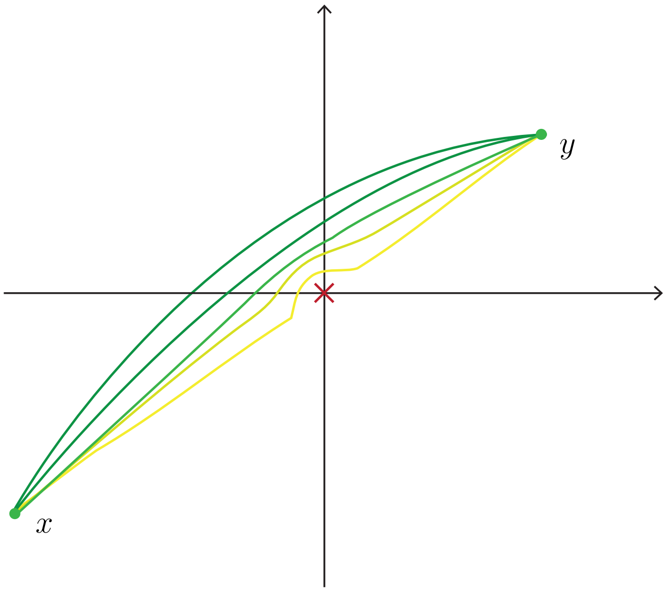

Figure 3.2: *A case where $d_I(x,y) = d(x,y)$, though the infimum is not achieved.*

Definition 50. Given a length space X, a curve $\gamma : [0,1] \to X$ is said to be a **(constant speed) geodesic** if there exists a constant $v > 0$ such that for all $t \in [0,1]$ there exists a neighborhood J of t in $[0,1]$ such that, for any $t_1, t_2 \in J$ with $t_1 > t_2$,

$$d(\gamma(t_1), \gamma(t_2)) = v(t_1 - t_2).$$

Geodesics, as defined above, are curves that are everywhere *locally* distance minimizers (because of (3.12)). Since shortest paths are curves that are *globally* distance minimizers, they are a particular case of geodesics.

Definition 51. Given a length space X, a curve $\gamma : [0,1] \to X$ is said to be a **(constant speed) minimal geodesic** or **shortest path** if there exists a constant $v > 0$ such that for all $t_1, t_2 \in [0,1]$ with $t_1 > t_2$,

$$d(\gamma(t_1), \gamma(t_2)) = v(t_1 - t_2).$$

The definition of minimal geodesic above simply means that for a curve to be a shortest path, it needs to be a shortest path between all points $\gamma(t_1), \gamma(t_2)$ with $t_1, t_2 \in [0,1]$, whereas a geodesic might only be a shortest path on small neighborhoods. As we have noted, every shortest path between two points, when it exists, must be a geodesic. The converse, however, is not true, as the following example shows.

Example 52. Consider the sphere $S^2 \subset \mathbb{R}^3$, and take $x, y \subset S^2$, as in Figure 3.3. Since x and y are not antipodal, the shortest path between them is precisely a segment of the great circle passing through them, colored in purple in the Figure. However, the other segment constituting the geodesic, colored in orange, must also be a geodesic, since it locally minimizes the distance between the two points.

Besides, note that if x and y are antipodal (as the green points in Figure 3.3), then there are infinitely many geodesics connecting x and y. Indeed there are infinitely many great circles passing through x and y in this case, and any of the two segments of a great circle connecting x and y constitutes a geodesic.

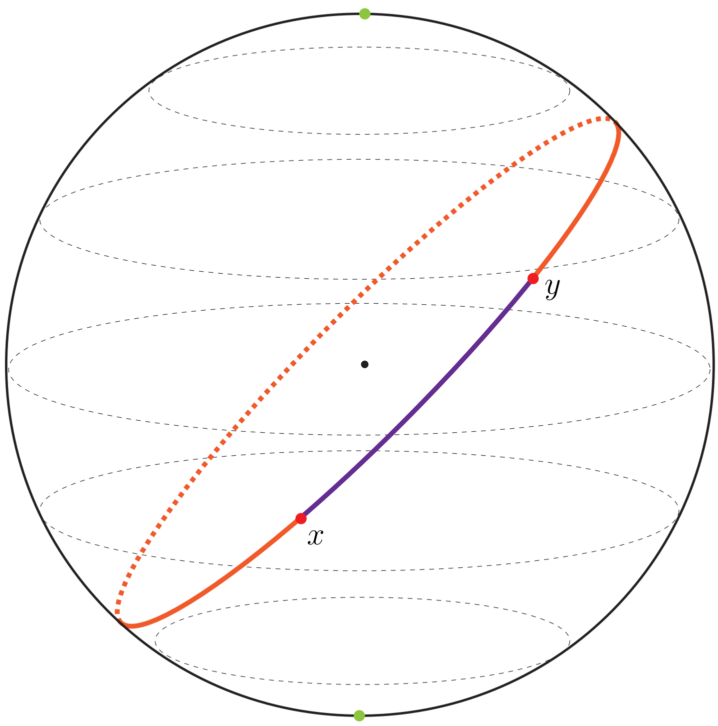

Figure 3.3: *Geodesics on a sphere.*

Essentially all the geodesics that will be considered from now on are minimal geodesics.

As example 49 showed, in arbitrary spaces it is possible that no geodesic exists between two given points. We would like to avoid this situation, and so the next definition comes naturally.

Definition 53. A metric space (X, d) is called **geodesic** if for every $x, y \in X$ there exists a minimal geodesic connecting them.

Now, in a geodesic space the notion of a shortest path makes sense, and we are guaranteed to find a curve that realizes it. The main aim of this section will be to prove that if the base space is geodesic, then its Wasserstein space is also geodesic. To simplify several results and avoid measure theoretical complications, we shall take our base space to be one that, given two points, admits only one geodesic.

Definition 54. A metric space (X, d) is called **uniquely geodesic** if for every $x, y \in X$ there exists a single minimal geodesic connecting them.

From example 52 above, we have already learned not to expect geodesics to be unique. Yet, several important spaces are in fact uniquely geodesic.

Example 55. It is easy to see that any normed vector space V must be geodesic, since given two points $x, y \in V$, the minimal geodesic connecting x and y must be given by $\gamma(t) = tx + (1-t)y$, with $t \in [0, 1]$. The more interesting question is: are all normed vector spaces uniquely geodesic? The general answer must clearly be 'no': just take $\mathbb{R}^2$ with the l^1 norm, and notice there are infinitely many geodesics connecting two arbitrary points x, y, as in Figure 3.4.

Still, if we restrict our attention to **strictly convex normed vector spaces** (i.e. a space E such that, for all $x_1, x_2 \in E, x_1 \neq x_2, ||x_1|| = ||x_2|| = 1$, we have $||(1-t)x_1 + tx_2|| < 1$), then the answer is 'yes'. Indeed, a normed vector space is uniquely geodesic if and only if it is strictly convex (see [28, p. 180]). As a consequence, any vector space equipped with a norm that comes from an inner product must be uniquely geodesic. In particular, any Hilbert space is uniquely geodesic.

Finally, let us return to our study of the Wasserstein space. From now on, we assume the base space X is a separable Hilbert space, with $p > 1$; from example 55 above we already know this is a uniquely geodesic Polish space.

A curve in the Wasserstein space can be thought of as a family of probability measures μ_t, one for each $t \in [0, 1]$. According to Definition 51, such a curve is a minimal geodesic if it satisfies

$$W_p(\mu_s, \mu_t) = (t-s)W_p(\mu_0, \mu_1), \quad \forall 0 \leq s \leq t \leq 1. \tag{3.13}$$

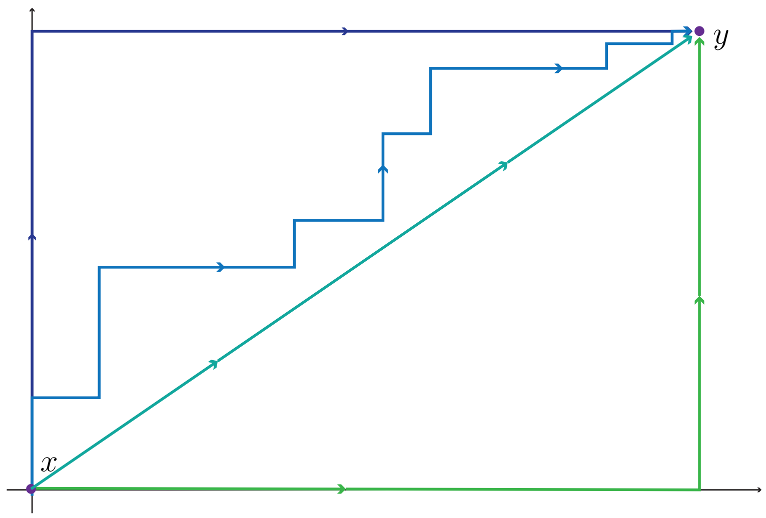

Figure 3.4: $(\mathbb{R}^2, ||\cdot||_1)$ *is not uniquely geodesic.*

Now, given two probability measures, how can we characterize the minimal geodesic between them? As a starting point, let us first consider the case $\mu = \delta_x, \nu = \delta_y$, with $x, y \in X$.

Perhaps a naive guess, in analogy with the fact that $\gamma(t) = (1-t)x + ty$ is the shortest path between $x, y \in X$, is to conjecture that a minimal geodesic between measures δ_x, δ_y, should be of the form $\mu_t = t\delta_x + (1-t)\delta_y$. However, if we take $p = 2$ and $0 \leq s \leq t \leq 1$, we readily see that

$$W_2(\mu_s, \mu_t) = \sqrt{t-s}\; d(x,t)$$

which of course does not satisfy equation 3.13 – indeed, $\{\mu_t\}_{t\in[0,1]}$ would have infinite length! This strange curve amounts to the somewhat artificial idea of sending mass from δ_x to δ_y at a distance, as Figure 3.5 shows.

A better guess would be to transport the mass from δ_x to δ_y optimally along geodesics in the base space X, as Figure 3.6 suggests. In other words, geodesics in $P_p(X)$ should be given by optimally *interpolating* two measures. This will be the content of Theorem 58 below.

Before doing so, we need to introduce some notation and an important lemma.

To make the notation more homogeneous, we let proj_i denote the standard projection in the $(i+1)$-th variable (so that proj_0 projects in the first variable, proj_1 in the second, etc.) and $\mathrm{proj}_{i,k}$ denote the standard projection on both the $(i+1)$-th and

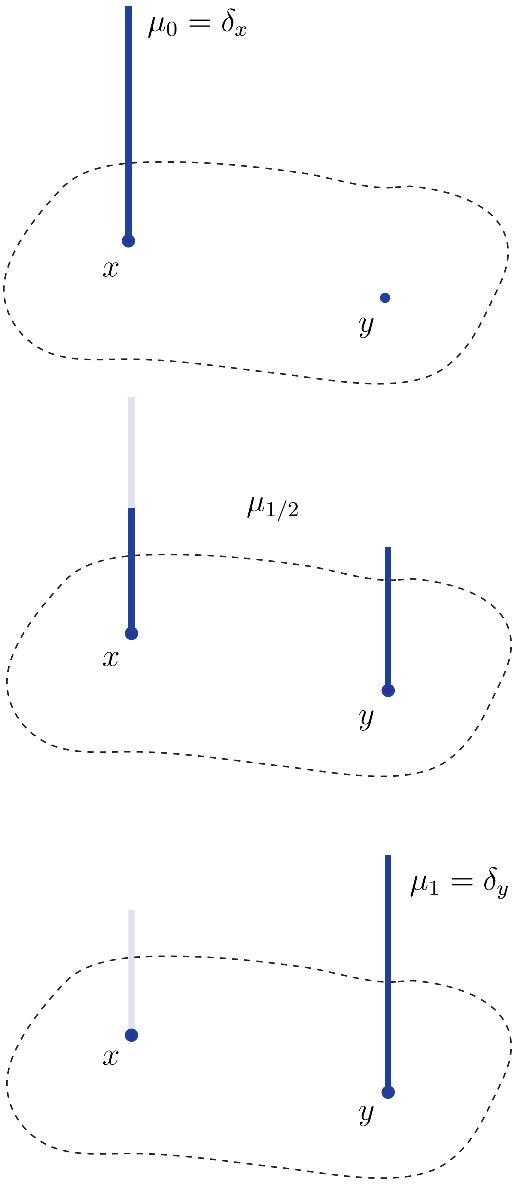

Figure 3.5: *The curve* $\mu_t = (1-t)\delta_x + t\delta_y$.

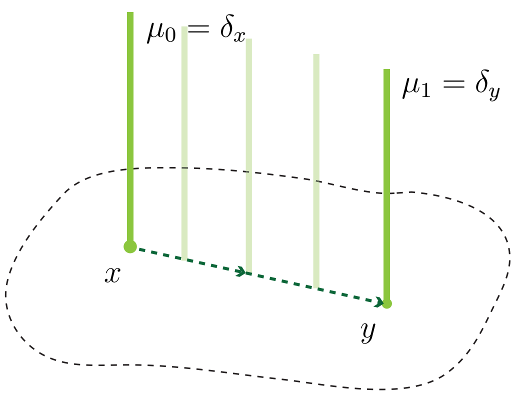

Figure 3.6: *Interpolating measures* δ_x *and* δ_y *provides a minimal geodesic curve.*

the $(k+1)$-th variables. Also, take $N \geq 2$, $0 \leq i, j, k \leq N-1$, and $t \in [0,1]$.

Now, given $\bar{\boldsymbol{\mu}} \in \mathcal{P}(X^N)$, define

- $\text{proj}_t^{i\to j} : X^N \to X$; with $\text{proj}_t^{i\to j} = (1-t)\,\text{proj}_i + t\,\text{proj}_j$;
- $\text{proj}_t^{i\to j,k} : X^N \to X^2$; with $\text{proj}_t^{i\to j,k} = (1-t)\,\text{proj}_{i,k} + t\,\text{proj}_{j,k}$;
- $\mu_t^{i\to j} \in \mathcal{P}(X)$; with $\mu_t^{i\to j} = (\text{proj}_t^{i\to j})_\# \bar{\boldsymbol{\mu}}$;
- $\mu_t^{i\to j,k} \in \mathcal{P}(X^2)$; with $\mu_t^{i\to j,k} = (\text{proj}_t^{i\to j,k})_\# \bar{\boldsymbol{\mu}}$.

The idea behind this notation is the following. Let $N = 2$, and take $\bar{\boldsymbol{\mu}} = \pi^* \in \Pi^*(\mu_0, \mu_1)$ to be an optimal coupling between μ_0 and μ_1 (greek boldface letters will denote optimal couplings, so that $\bar{\boldsymbol{\mu}} \in \mathcal{P}(X^N)$ but $\mu \in \mathcal{P}(X)$). Then we will prove that $\mu_t = \mu_t^{0\to 1} = (\text{proj}_t^{0\to 1})_\# \bar{\boldsymbol{\mu}}$ is a geodesic from $\mu_0 = (\text{proj}_0)_\# \bar{\boldsymbol{\mu}}$ to $\mu_1 = (\text{proj}_1)_\# \bar{\boldsymbol{\mu}}$.

Notice that, for any measurable $U \subset X$,

$$\mu_t(U) = (\text{proj}_t^{0\to 1})_\# \pi^*(U) = \pi^* \left\{(x,y) \in X \times X : tx + (1-t)y \in U\right\},$$

so that if $\pi^*(U_0 \times U_1) > 0$, then the set

$$U_t = \{(1-t)x + ty : x \in U_0, y \in U_1\},$$

illustrated in Figure 3.7, has positive measure μ_t. Indeed,

$$\mu_t(U_t) = \pi^*\{(x,y) \in X \times X : (1-t)x + ty \in U_t\} \geq \pi^*(U_0 \times U_1).$$

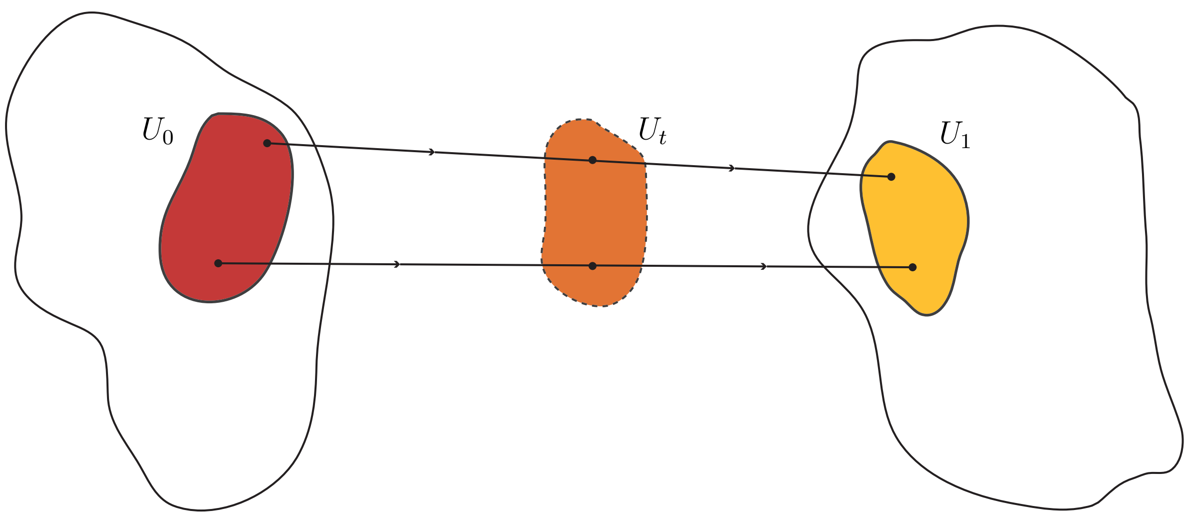

Figure 3.7: *Understanding the behavior of the geodesic* $\{\mu_t\}_{t\in[0,1]}$.

We will need the following lemma, which essentially states that though there could be several geodesics from μ_0 to μ_1, any point μ_t along a geodesic between these measures is such that $\Pi^*(\mu_0, \mu_t)$ and $\Pi^*(\mu_t, \mu_1)$ have only one element, and the optimal transport from μ_0 to μ_t and the one from μ_t to μ_1 are induced by a map.

Lemma 56. *Given a minimal geodesic $\{\mu_t\}_{t\in[0,1]}$ in $P_p(X)$, for each $t \in (0,1)$ the sets $\Pi^*(\mu_0, \mu_t)$ and $\Pi^*(\mu_t, \mu_1)$ have a unique optimal plan, which we respectively call $\bar{\boldsymbol{\mu}}^{0\to t}$ and $\bar{\boldsymbol{\mu}}^{t\to 1}$. Furthermore, there exists $\bar{\boldsymbol{\mu}}$ such that both $\bar{\boldsymbol{\mu}}^{0\to t}$ and $\bar{\boldsymbol{\mu}}^{t\to 1}$ are induced by transport maps:*

$$\bar{\boldsymbol{\mu}}^{0\to t} = (\mathrm{proj}_t^{0,0\to 1})_{\#}\bar{\boldsymbol{\mu}} \text{ and } \bar{\boldsymbol{\mu}}^{t\to 1} = (\mathrm{proj}_t^{0\to 1,1})_{\#}\bar{\boldsymbol{\mu}}$$

and $\bar{\boldsymbol{\mu}} = \bar{\boldsymbol{\mu}}^{t\to 1} \circ \bar{\boldsymbol{\mu}}^{0\to t}$ (in the sense of remark 27),

Proof. Fix $t \in (0,1)$. We first let $\boldsymbol{\alpha} \in \Pi^*(\mu_0, \mu_t)$ and $\boldsymbol{\beta} \in \Pi^*(\mu_t, \mu_1)$, so that both $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ constitute optimal plans. To make the proof easier to understand, we consider three distinct copies of X, and name them X_1, X_2, X_3 so that $\mu_0 \in \mathcal{P}(X_1), \mu_t \in \mathcal{P}(X_2)$ and $\mu_1 \in \mathcal{P}(X_3)$.

We now use Lemma 26 (the Gluing Lemma) to 'glue' the measures $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ with respect to their common variable x_2, obtaining a third measure $\boldsymbol{\gamma} \in \mathcal{P}(X_1 \times X_2 \times X_3)$. Indeed, we first disintegrate the measures $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ as

$$\boldsymbol{\alpha} = \int_{X_2} \alpha_{X_1|x_2} \otimes \delta_{x_2} d\mu_t(x_2)$$
$$\boldsymbol{\beta} = \int_{X_2} \beta_{X_3|x_2} \otimes \delta_{x_2} d\mu_t(x_2),$$

and then 'glue' them together, obtaining

$$\boldsymbol{\gamma} = \int_{X_2} \alpha_{X_1|x_2} \otimes \delta_{x_2} \otimes \beta_{X_3|x_2} d\mu_t(x_2),$$

which is essentially equation (3.2) rephrased.

Thus, using the notation of remark 27, we define

$$\bar{\boldsymbol{\mu}} = \boldsymbol{\alpha} \circ \boldsymbol{\beta} = (\mathrm{proj}_{1,3})_{\#}\boldsymbol{\gamma} \in \Pi(\mu_0, \mu_1), \tag{3.14}$$

and obtain

$$\begin{aligned}
W_p(\mu_0,\mu_1) &\le \left(\int_{X_1\times X_3} d^p(x_1,x_3)d\bar{\boldsymbol{\mu}}\right)^{\frac{1}{p}} = \left(\int_{X_1\times X_2\times X_3} d^p(x_1,x_3)d\boldsymbol{\gamma}\right)^{\frac{1}{p}}\\
&\le \left(\int_{X_1\times X_2\times X_3} d^p(x_1,x_2)d\boldsymbol{\gamma}\right)^{\frac{1}{p}} + \left(\int_{X_1\times X_2\times X_3} d^p(x_2,x_3)d\boldsymbol{\gamma}\right)^{\frac{1}{p}}\\
&= \left(\int_{X_1\times X_2} d^p(x_1,x_2)d\boldsymbol{\alpha}\right)^{\frac{1}{p}} + \left(\int_{X_2\times X_3} d^p(x_2,x_3)d\boldsymbol{\beta}\right)^{\frac{1}{p}}\\
&= W_p(\mu_0,\mu_t) + W_p(\mu_t,\mu_1) = W_p(\mu_0,\mu_1).
\end{aligned}$$

The first inequality follows from the definition of the Wasserstein metric; the first equality follows (3.14) above; the second inequality is simply the usual triangular inequality in L^p; the second equality comes from the marginal identities of $\boldsymbol{\gamma}$; and the third equality comes from $\boldsymbol{\alpha} \in \Pi^*(\mu_0,\mu_t)$, $\boldsymbol{\beta} \in \Pi^*(\mu_t,\mu_1)$, and the fact that $\{\mu_t\}_{t\in[0,1]}$ is geodesic.

Since we must then have equality everywhere, this calculation shows that $\bar{\boldsymbol{\mu}}$ is an optimal measure between μ_0 and μ_1, i.e. $\bar{\boldsymbol{\mu}} \in \Pi^*(\mu_0,\mu_1)$.

Also, we get that

$$||x_1 - x_3||_{L^p(X_1\times X_2\times X_3,\gamma)} = ||x_1 - x_2||_{L^p(X_1\times X_2\times X_3,\gamma)} + ||x_2 - x_3||_{L^p(X_1\times X_2\times X_3,\gamma)},$$

and so, by the strict convexity of the L^p-norm, $\boldsymbol{\gamma}$-a.e. $x_2 - x_1$ and $x_3 - x_1$ should be collinear; that is, there exists $k > 0$ such that

$$x_2 - x_1 = k(x_3 - x_1)$$

for $\boldsymbol{\gamma}$-a.e. triples (x_1,x_2,x_3).

Because $\{\mu_t\}$ is a geodesic in $P_p(X)$, we have that $W_p(\mu_0,\mu_t) = tW_p(\mu_0,\mu_1)$, which implies

$$\begin{aligned}
\int_{X_1\times X_2\times X_3} d^p(x_1,x_2)d\boldsymbol{\gamma} &= \int_{X_1\times X_2} d^p(x_1,x_2)d\boldsymbol{\alpha} = t^p \int_{X_1\times X_3} d^p(x_1,x_3)d\bar{\boldsymbol{\mu}}\\
&= \int_{X_1\times X_2\times X_3} (td(x_1,x_3))^p d\boldsymbol{\gamma}
\end{aligned}$$

and so $k = t$. That is,

$$x_2 - x_1 = t(x_3 - x_1) \tag{3.15}$$

for $\boldsymbol{\gamma}$-a.e. triples (x_1, x_2, x_3).

Let $z(x_2) = \int_{X_1} x_1 d\alpha_{X_1|x_2}$, and note that $\alpha_{X_1|x_2} = \gamma_{X_1|x_2,x_3}$ by the way we defined $\boldsymbol{\gamma}$. Then, integrating (3.15) with respect to the measure $\alpha_{X_1|x_2}$, we get

$$x_2 - z(x_2) = t(x_3 - z(x_2)),$$

for β-a.e. pairs (x_2, x_3).

This, in turn, implies the transport map $r_t : X_2 \to X_3$ given by

$$r_t(x_2) = \frac{x_2}{t} - (1-t)\frac{z(x_2)}{t}$$

induces the measure $\boldsymbol{\beta}$, i.e. $\boldsymbol{\beta} = (\mathrm{Id} \times r_t)_{\#}\mu_t$. Note that $\boldsymbol{\alpha}$ determines z, which depends on r_t, which in turn depends on $\boldsymbol{\beta}$. Still, as $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ were chosen independently, $\boldsymbol{\beta}$ is properly determined by the transport map r_t; thus, $\boldsymbol{\beta}$ must be unique.

Finally, by taking $\bar{\boldsymbol{\mu}}^{t\to 1} = \boldsymbol{\beta}$, we get that $\bar{\boldsymbol{\mu}}^{t\to 1}$ is indeed unique and induced by the appropriate transport map. Analogously, one sees that $\bar{\boldsymbol{\mu}}^{0\to t} = \boldsymbol{\alpha}$ also fits the bill, and so $\bar{\boldsymbol{\mu}} = \bar{\boldsymbol{\mu}}^{t\to 1} \circ \bar{\boldsymbol{\mu}}^{0\to t} \in \Pi^*(\mu_0, \mu_1)$. This finishes the lemma. ■

Finally, we prove the main result of this subsection: if we optimally interpolate two measures μ_0 and μ_1 using the formula for $\mu_t^{0\to 1}$ above, we get that $\{\mu_t\}_{t\in[0,1]}$ is a geodesic. Conversely, given any geodesic $\{\mu_t\}_{t\in[0,1]}$ we can find an optimal map $\bar{\boldsymbol{\mu}} \in \Pi^*(\mu_0, \mu_1)$ such that this geodesic satisfies $\mu_t = \mu_t^{0\to 1} = (\mathrm{proj}_t^{0\to 1})_{\#}\bar{\boldsymbol{\mu}}$.

Theorem 57. *Given an optimal plan $\bar{\boldsymbol{\mu}} \in \Pi^*(\mu_0, \mu_1)$, the curve $\{\mu_t\}_{t\in[0,1]}$ with $\mu_t = \mu_t^{0\to 1}$ is a minimal geodesic from μ_0 to μ_1. Conversely, all geodesics are of this form.*

Proof. ($\Rightarrow$) Note that, for $0 \le s \le t \le 1$, since μ_t, μ_s are induced by transport maps,

$$\begin{aligned} W_p^p(\mu_t, \mu_s) &= W_p^p((\mathrm{proj}_t)_{\#}\bar{\boldsymbol{\mu}}, (\mathrm{proj}_s)_{\#}\bar{\boldsymbol{\mu}}) \le \int_{X\times X} d^p(x,y) d\left((\mathrm{proj}_t, \mathrm{proj}_s)_{\#}\bar{\boldsymbol{\mu}}\right)(x,y) \\ &= \int_{X\times X} d^p(\mathrm{proj}_t(x,y), \mathrm{proj}_s(x,y)) d\bar{\boldsymbol{\mu}}(x,y) = \int_{X\times X} (t-s)^p d^p(x,y) d\bar{\boldsymbol{\mu}}(x,y) \\ &= (t-s)^p \int_{X\times X} d^p(x,y) d\bar{\boldsymbol{\mu}} = (t-s)^p W_p^p(\mu_0, \mu_1). \end{aligned}$$

The inequality comes form the fact that $(f,g)_{\#}\bar{\boldsymbol{\mu}} \in \Pi(f_{\#}\bar{\boldsymbol{\mu}}, g_{\#}\bar{\boldsymbol{\mu}})$. The first equality comes from equation (2.2); and the second equality comes from the fact that $|(1-t)x + ty - (1-s)x - sy| = (t-s)|x-y|$.

Thus, $W_p(\mu_t, \mu_s) \leq (t-s)W_p(\mu_0, \mu_1)$. However, if the inequality is strict we can find a contradiction by simply applying the triangular inequality to μ_0, μ_s, μ_t and μ_1. Therefore, equality must hold, and, by definition, $\{\mu_t\}$ is a minimal geodesic.

($\Leftarrow$) Conversely, take $\{\mu_t\}_{t\in[0,1]}$ to be a minimal geodesic. Fix t, and consider the geodesic restricted to the interval $[0,t]$; that is, consider the minimal geodesic $\{\mu_s\}_{s\in[0,t]}$, which can of course be thought of as $\{\mu_{st}\}_{s\in[0,1]}$. From Lemma 56, we know that μ_{st} is uniquely determined by a transport map, so

$$\mu_{st} = (\mathrm{proj}_s^{0\to 1})_{\#}\bar{\boldsymbol{\mu}}^{0\to t} = (\mathrm{proj}_s^{0\to 1} \circ \mathrm{proj}_t^{0,0\to 1})_{\#}\bar{\boldsymbol{\mu}} = (\mathrm{proj}_{st}^{0\to 1})_{\#}\bar{\boldsymbol{\mu}},$$

and the geodesic going from 0 to t does indeed take the form $(\mathrm{proj}_{st}^{0\to 1})_{\#}\bar{\boldsymbol{\mu}}, s \in [0,1]$. The argument is analogous for the other part of the geodesic, though we use the measure $\bar{\boldsymbol{\mu}}^{t\to 1}$ instead of $\bar{\boldsymbol{\mu}}^{0\to t}$. In any case, since $\bar{\boldsymbol{\mu}} = \bar{\boldsymbol{\mu}}^{t\to 1} \circ \bar{\boldsymbol{\mu}}^{0\to t}$, we get that the whole geodesic must satisfy

$$\mu_t = (\mathrm{proj}_t^{0\to 1})_{\#}\bar{\boldsymbol{\mu}},$$

which finishes the theorem. ■

Though the proof above relied heavily on properties of the separable Hilbert space X, Theorem 57 can be generalized in several directions. For instance, one could have let the base space X to be any geodesic space, and still obtain that $(P_p(X), W_p)$ is a geodesic space for $p > 1$. Indeed, let $\mathrm{Geod}(X)$ denote the metric space of all minimal geodesics on X, endowed with the supremum norm. Also, define the evaluation maps $e_t : \mathrm{Geod}(X) \to X$ by $e_t(\gamma) = \gamma(t)$ for $t \in [0,1]$. Then the following holds.

Theorem 58. *If (X,d) is a Polish geodesic space, then $(P_p(X), W_p)$ is also Polish and geodesic; the converse also holds. Furthermore, for every minimal geodesic $\{\mu_t\}_{t\in[0,1]}$, there exists $\tilde{\boldsymbol{\mu}} \in P_p(\mathrm{Geod}(X))$ such that $(e_0, e_1)_{\#}\tilde{\boldsymbol{\mu}} \in \Pi^*(\mu_0, \mu_1)$ and $\mu_t = (e_t)_{\#}\tilde{\boldsymbol{\mu}}$; conversely, any such curve $\{\mu_t\}_{t\in[0,1]}$ must be a minimal geodesic.*

Proof. See [3, p. 31]. ■

Curvature

We would like to somehow measure how 'curved' the Wasserstein space is. To make matters simpler, we restrict attention to the case where $p = 2$ (and, as before, X is a separable Hilbert space). Now, we must first define a notion of curvature that is suitable for metric spaces. Intuitively, we want it to distinguish between Figures 3.8(a), 3.8(b) and 3.8(c).

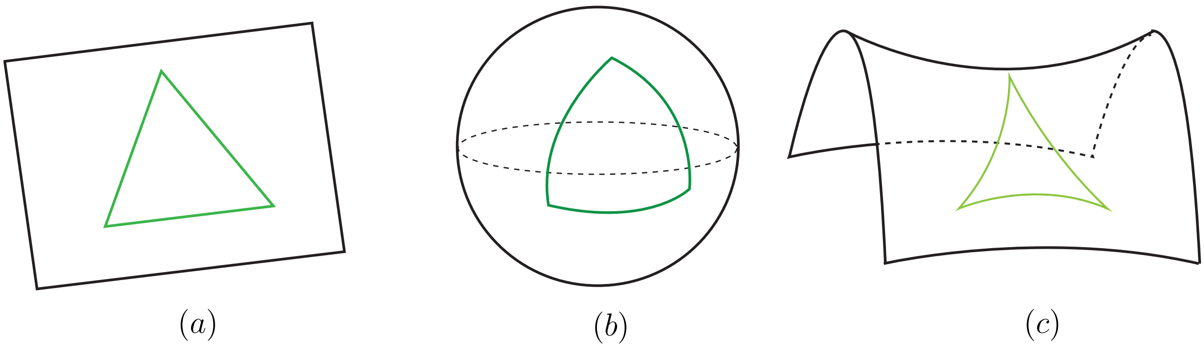

Figure 3.8: *A space like (a) should have zero curvature; a space like (b) should have positive curvature; and a space like (c) should have negative curvature.*

The idea will be as follows: given three points $x_1, x_2, x_3 \in X$, and geodesics $x^{1\to 2}, x^{2\to 3}, x^{3\to 1}$ among them, compare the triangle formed in X with an euclidean triangle in $\mathbb{R}^2$ with sides of length $d(x_1, x_2), d(x_2, x_3)$, and $d(x_3, x_1)$. If the distance between a vertice and any point in the opposing geodesic is bigger or equal to the distance that between the corresponding two points in the $\mathbb{R}^2$ triangle, we will say that the original space is non-negatively curved, or NNC. If, on the other hand, the distance is smaller or equal to the corresponding distance in the $\mathbb{R}^2$ triangle, we will say that the original space is non-positively curved, or NPC.

Note that this notion agrees with our intuition for Figure 3.8, and it would make the space in 3.8(b) NNC, the one in 3.8(c) NPC, and the space in 3.8(a) both NNC and NPC.

Since in $\mathbb{R}^n$, or, more generally, a Hilbert space, geodesics take the form $x_t^{1\to 2} = (1-t)x_1 + tx_2$, we have the equality

$$|x_t^{1\to 2} - x_3| = (1-t)|x_1 - x_3|^2 + t|x_2 - x_3|^2 - t(1-t)|x_1 - x_2|^2.$$

Hence, it makes sense to define NNC spaces as follows.

Definition 59. A geodesic metric space (X, d) is said to be **non-negatively curved** or NNC if for every $x_3 \in X$ and every minimal geodesic $\{x_t^{1\to 2}\}_{t\in[0,1]}$ connecting $x_1, x_2 \in X$,

$$d^2(x_t^{1\to 2}, x_3) \geq (1-t)d^2(x_1, x_3) + td^2(x_2, x_3) - t(1-t)d^2(x_1, x_2). \tag{3.16}$$

The space is said to be **non-positively curved** or NPC if the reverse inequality holds.

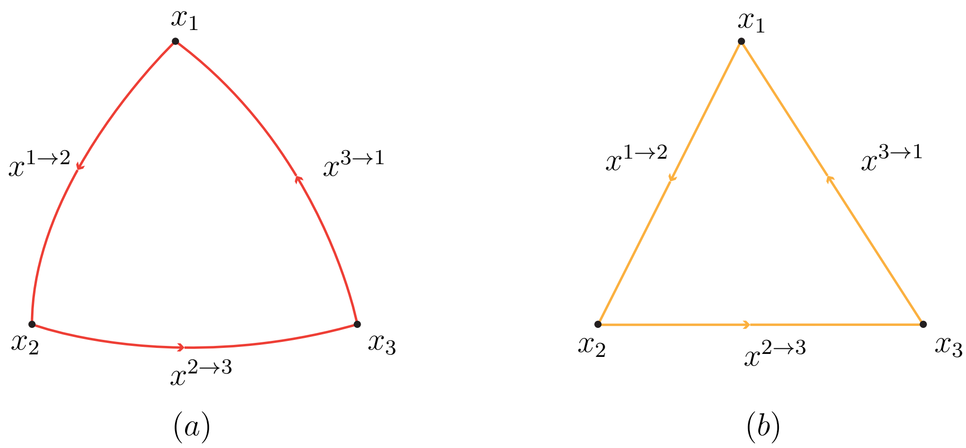

Figure 3.9: *Comparing triangles.*

It is not hard to show that the definition above does indeed agree with our intuition regarding triangles. In any case, for our purposes, the Wasserstein space $(P_2(X), W_2)$ is said to be NNC if for each choice of $\mu_1, \mu_2, \mu_3 \in P_2(X)$ it holds that

$$W_2^2(\mu_t^{1\to 2}, \mu_3) \geq (1-t)W_2^2(\mu_1, \mu_3) + tW_2^2(\mu_2, \mu_3) - t(1-t)W_2^2(\mu_1, \mu_2).$$

In this section, our aim will be to prove that if X is a NNC space, then so will be $P_2(X)$. First, we need two technical results; we now let proj_i denote the standard projection in the i-th coordinate, again.

Lemma 60. *Given* $\bar{\boldsymbol{\mu}} \in \Pi(\mu_1, \mu_2, \mu_3) \subset P_2(X \times X \times X)$*,* $i, j, k \in \{1, 2, 3\}$ *and* $t \in [0, 1]$*, if we define*

$$C^2_{\bar{\boldsymbol{\mu}}}(\mu_t^{i\to j}, \mu_k) = \int_{X\times X\times X} |(1-t)x_i + tx_j - x_k|^2 d\bar{\boldsymbol{\mu}}(x_1, x_2, x_3),$$

then it holds that

$$W_2^2(\mu_t^{1\to 2}, \mu_3) \leq (1-t)C^2_{\bar{\boldsymbol{\mu}}}(\mu_1, \mu_3) + tC^2_{\bar{\boldsymbol{\mu}}}(\mu_2, \mu_3) - t(1-t)C^2_{\bar{\boldsymbol{\mu}}}(\mu_1, \mu_2).$$

Proof. First, recall that for any measure $\mu \in \mathcal{P}(X)$ and any pair of Borel measurable maps $f, g : X \to X$, we must have

$$W_2^2(f_\#\mu, g_\#\mu) \leq \int_X d^2(f(x), g(x)) d\mu(x),$$

since $(f, g)_\#\mu \in \Pi(f_\#\mu, g_\#\mu)$ and the Wasserstein metric is defined as the infimum cost among all admissible measures.

Furthermore, note that $C^2_{\bar{\boldsymbol{\mu}}}(\mu_t^{i\to j}, \mu_k)$ is simply the squared cost of the push-forward of $\bar{\boldsymbol{\mu}}$ by the projections $f(x_1, x_2, x_3) = (1-t)x_i + tx_j$ and $g(x_1, x_2, x_3) = x_k$, so we get

$$W_2^2(\mu_t^{1\to 2}, \mu_3) \leq C^2_{\bar{\boldsymbol{\mu}}}(\mu_t^{1\to 2}, \mu_3). \tag{3.17}$$

Now, the identity

$$|(1-t)a + tb - c|^2 = (1-t)|a-c|^2 + t|b-c|^2 - t(1-t)|b-a|^2$$

yields

$$\begin{aligned} C^2_{\bar{\boldsymbol{\mu}}}(\mu_t^{1\to 2}, \mu_3) &= \int_{X\times X\times X} |(1-t)x_1 + tx_2 - x_3|^2 d\bar{\boldsymbol{\mu}}(x_1, x_2, x_3) \\ &= (1-t)\int |x_1 - x_3|^2 d\bar{\boldsymbol{\mu}} + t\int |x_2 - x_3|^2 d\bar{\boldsymbol{\mu}} - t(1-t)\int |x_2 - x_1|^2 d\bar{\boldsymbol{\mu}} = \\ &= (1-t)C^2_{\bar{\boldsymbol{\mu}}}(\mu_1, \mu_3) + tC^2_{\bar{\boldsymbol{\mu}}}(\mu_2, \mu_3) - t(1-t)C^2_{\bar{\boldsymbol{\mu}}}(\mu_1, \mu_2), \end{aligned}$$

and so we have

$$W_2^2(\mu_t^{1\to 2}, \mu_3) \leq (1-t)C^2_{\bar{\boldsymbol{\mu}}}(\mu_1, \mu_3) + tC^2_{\bar{\boldsymbol{\mu}}}(\mu_2, \mu_3) - t(1-t)C^2_{\bar{\boldsymbol{\mu}}}(\mu_1, \mu_2),$$

as we wished. ■

Proposition 61. *Let $\boldsymbol{\mu}^{1,2} \in \Pi(\mu_1, \mu_2)$ and $\boldsymbol{\mu}^{t,3} \in \Pi^*(\mu_t^{1\to 2}, \mu_3)$, with $t \in (0, 1)$. Then there exists a plan $\boldsymbol{\mu}_t \in \Pi(\mu_1, \mu_2, \mu_3)$ such that $(\text{proj}_t^{1\to 2,3})_\#\boldsymbol{\mu}_t = \boldsymbol{\mu}^{t,3}$ and thus*

$$W_2^2(\mu_t^{1\to 2}, \mu_3) = (1-t)C^2_{\boldsymbol{\mu}_t}(\mu_1, \mu_3) + tC^2_{\boldsymbol{\mu}_t}(\mu_2, \mu_3) - t(1-t)C^2_{\boldsymbol{\mu}_t}(\mu_1, \mu_2), \tag{3.18}$$

Moreover, the plan $\boldsymbol{\mu}_t$ is unique if $\boldsymbol{\mu}^{1,2} \in \Pi^*(\mu_1, \mu_2)$.

Proof. Define the homeomorphisms $Q_t : X \times X \to X \times X$ and $R_t : X \times X \times X \to X \times X \times X$ by

$$Q_t(x_1, x_2) = ((1-t)x_1 + tx_2, x_2)$$
$$R_t(x_1, x_2, x_3) = ((1-t)x_1 + tx_2, x_2, x_3)\,,$$

and consider the measure $\eta = (R_t)_\#\mu$. Note that the properties $\boldsymbol{\mu}_t \in \Pi(\mu_1, \mu_2, \mu_3)$ and $(\mathrm{proj}_t^{1\to 2,3})_\#\boldsymbol{\mu}_t = \boldsymbol{\mu}^{t,3}$ are equivalent to asking $(\mathrm{proj}_{1,2})_\#\eta = (Q_t)_\#\boldsymbol{\mu}^{1,2}$ and $(\mathrm{proj}_{1,3})_\#\eta = \boldsymbol{\mu}^{t,3}$, since

- $(\mathrm{proj}_i)_\#\eta = (\mathrm{proj}_i)_\#((R_t)_\#\boldsymbol{\mu}) = (\mathrm{proj}_i)_\#((1-t)\mu_1 + t\mu_2, \mu_2, \mu_3)$
- $(Q_t)_\#\boldsymbol{\mu}^{1,2} = ((1-t)\mu_1 + t\mu_2, \mu_2)$
- $\boldsymbol{\mu}^{t,3} = ((1-t)\mu_1 + t\mu_2, \mu_3)$.

Because $(\mathrm{proj}_{1,2})_\#\eta$ and $(\mathrm{proj}_{1,3})_\#\eta$ have a common marginal, Lemma 26 (the Gluing Lemma) guarantees the existence of a measure η such as the one we want. Then, the fact that R_t is a homeomorphism is enough to ensure the existence of $\boldsymbol{\mu}_t$.

From $(\mathrm{proj}_t^{1\to 2,3})_\#\boldsymbol{\mu}_t = \boldsymbol{\mu}^{t,3} \in \Pi^*(\mu_t^{1\to 2}, \mu_3)$ and Lemma 60, we get

$$\begin{aligned}
W_2^2(\mu_t^{1\to 2}, \mu_3) &= \int_{X\times X} d^2(x,y)\; d\left((\mathrm{proj}_t^{1\to 2}, \mathrm{proj}_3)_\#\mu_t)\right) \\
&= \int_{X\times X\times X} d^2\left(\mathrm{proj}_t^{1\to 2}, \mathrm{proj}_3\right) d\boldsymbol{\mu}_t(x_1, x_2, x_3) \\
&= \int_{X\times X\times X} |(1-t)x_1 + tx_2 - x_3|^2 d\boldsymbol{\mu_t}(x_1, x_2, x_3) = C^2_{\boldsymbol{\mu}_t}(\mu_t^{1\to 2}, \mu_3) \\
&= (1-t)C^2_{\boldsymbol{\mu}_t}(\mu_1, \mu_3) + tC^2_{\boldsymbol{\mu}_t}(\mu_2, \mu_3) - t(1-t)C^2_{\boldsymbol{\mu}_t}(\mu_1, \mu_2).
\end{aligned}$$

Finally, if $\boldsymbol{\mu}^{1,2} \in \Pi^*(\mu_1, \mu_2)$, then $(Q_t)_\#\boldsymbol{\mu}^{1,2} \in \Pi^*(\mu_t^{1\to 2}, \mu_2)$ and from Lemma 56 we have that $(Q_t)_\#\boldsymbol{\mu}^{1,2}$ is unique because it is induced by a transport map. In turn, since Q_t is a homeomorphism, we have $\boldsymbol{\mu}_t$ is uniquely determined. ■

We thus obtain the main result of this subsection.

Theorem 62. *For any measures $\mu_1, \mu_2, \mu_3 \in P_2(X)$, the following inequality holds:*

$$W_2^2(\mu_t^{1\to 2}, \mu_3) \geq (1-t)W_2^2(\mu_1, \mu_3) + tW_2^2(\mu_2, \mu_3) - t(1-t)W_2^2(\mu_1, \mu_2);$$

that is, $P_2(X)$ is a NNC space.

Proof. Given $\mu_1, \mu_2, \mu_3 \in P_2(X)$, $\boldsymbol{\mu}^{1,2} \in \Pi(\mu_1, \mu_2)$ and using (3.17), (3.18) we obtain

$$W_2^2(\mu_t^{1\to 2}, \mu_3) \geq (1-t)W_2^2(\mu_1, \mu_3) + tW_2^2(\mu_2, \mu_3) - t(1-t)C^2_{\boldsymbol{\mu}^{1,2}}(\mu_1, \mu_2). \tag{3.19}$$

Now, since we are considering a geodesic between μ_1 and μ_2, because of Theorem 57 it must be that $\boldsymbol{\mu}^{1,2} \in \Pi^*(\mu_1, \mu_2)$. This, in turn, means that

$$C^2_{\boldsymbol{\mu}^{1,2}}(\mu_1, \mu_2) = W_2^2(\mu_1, \mu_2). \tag{3.20}$$

Putting equations (3.19) and (3.20) together yields

$$W_2^2(\mu_t^{1\to 2}, \mu_3) \geq (1-t)W_2^2(\mu_1, \mu_3) + tW_2^2(\mu_2, \mu_3) - t(1-t)W_2^2(\mu_1, \mu_2),$$

which shows that $P_2(X)$ is indeed a NNC space. ■

As in the last subsection, the proofs above relied on the underlying Hilbert structure of the space X. Still, Theorem 62 holds in far greater generality: we could have taken X to be just a geodesic space.

Theorem 63. *If (X, d) is a geodesic space that is NNC, then the space $(P_2(X), W_2)$ is also NNC.*

Proof. See [3, p. 39]. ■

Though the non-negative curvature seems to be inherited by the Wasserstein space, the same is not true for non-positive curvature. That is, it is possible for (X, d) to be a NPC space while $(P_2(X), W_2)$ is not, as the next example shows.

Example 64. Consider the space $X = \mathbb{R}^2$ with the usual euclidean distance. it is straightforward to see that since in this case (3.16) holds in equality, $\mathbb{R}^2$ must be NPC. Now, let us show that $(P_2(\mathbb{R}^2), W_2)$ is not NPC. Define

$$\begin{aligned}
\mu_0 &= \frac{1}{2}\left(\delta_{(1,1)} + \delta_{(5,3)}\right), \\
\mu_1 &= \frac{1}{2}\left(\delta_{(-1,1)} + \delta_{(-5,3)}\right), \\
\nu &= \frac{1}{2}\left(\delta_{(0,0)} + \delta_{(0,-4)}\right),
\end{aligned}$$

as in Figure 3.10.

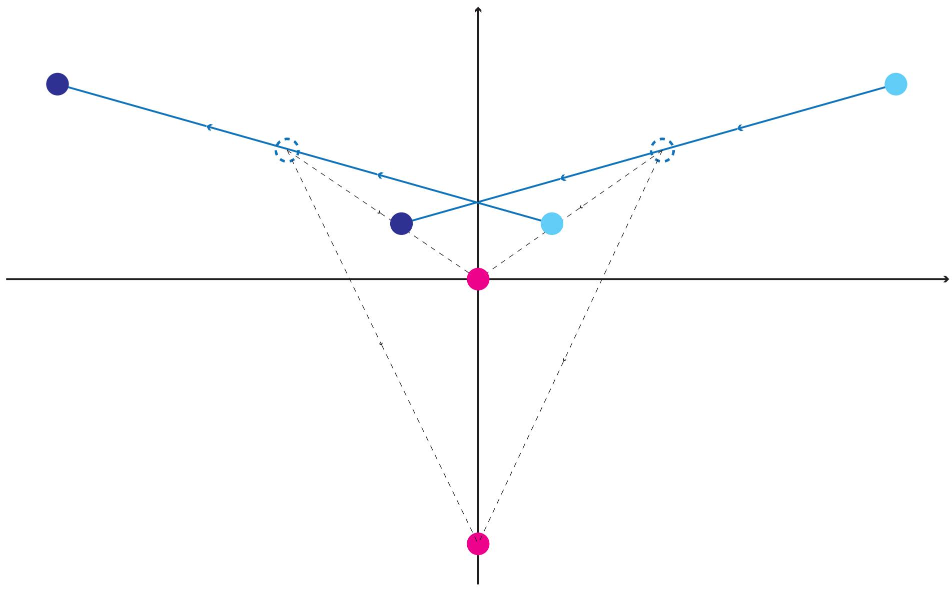

Figure 3.10: *Illustration of μ_0 (light blue), μ_1 (dark blue), ν (pink), and their transports.*

Since the cost function is convex, it will be better to transport the masses along two medium length paths than along one big and one small path, as suggested by the blue path in Figure 3.10. It is then straightforward to see that $W_2^2(\mu_0, \mu_1) = 40$ and $W_2^2(\mu_0, \nu) = W_2^2(\mu_1, \nu) = 30$.

Besides, we already know that $\mathbb{R}^2$ with the euclidean distance is uniquely geodesic, and in this particular example the geodesic between μ_1 and μ_2 should follow along the blue path in Figure 3.10. Therefore, the geodesic between μ_1 and μ_2 should be

$$\mu_t^{1\to 2} = \frac{1}{2}\left(\delta_{(1-6t,1+2t)} + \delta_{(5-6t,3-2t)}\right).$$

Then, the transport from $\mu_{1/2}^{1\to 2}$ to ν should follow along the dotted black line in Figure 3.10, and easy calculations show that

$$W_2^2(\mu_{1/2}^{1\to 2}, \nu) = 24 > \frac{30}{2} + \frac{30}{2} - \frac{40}{4} = tW_2^2(\mu_0, \nu) + (1-t)W_2^2(\mu_1, \nu) - t(1-t)W_2^2(\mu_0, \mu_1).$$

Hence, $(P_2(\mathbb{R}^2), W_2)$ violates the NPC condition.

This example also illustrates two additional interesting facts: first, the inequality (3.16) is oftentimes strict; put another way, the triangles in the Wasserstein space are usually more like the ones in Figure 3.3(a) then the one in Figure 3.3(b). Second, although geodesics in $\mathbb{R}^2$ behave very well (they are line segments, after all), the geodesics in $(P_2(\mathbb{R}^2), W_2)$ can have both intersecting paths and splitting of masses.

Bibliography

[1] AHMAD, N.; KIM, H. K. ; MCCANN, R. J. **Bulletin of Mathematical Sciences**. Optimal transportation, topology and uniqueness, journal, v.1, n.1, p. 13–32, 2011.

[2] AMBROSIO, L. **Lecture notes on optimal transport problems**. Springer, 2003.

[3] AMBROSIO, L.; GIGLI, N. **A user's guide to optimal transport**. In: MODELLING AND OPTIMISATION OF FLOWS ON NETWORKS, p. 1–155. Springer, 2013.

[4] AMBROSIO, L.; GIGLI, N. ; SAVARÉ, G. **Gradient flows**. Springer, 2005.

[5] BERNARD, P.; BUFFONI, B. **arXiv preprint math/0412299**. Optimal mass transportation and mather theory, journal, 2004.

[6] BERTRAND, J.; KLOECKNER, B. **Journal of Topology and Analysis**. A geometric study of wasserstein spaces: Hadamard spaces, journal, v.4, n.04, p. 515–542, 2012.

[7] BIRKHOFF, G. **Revi. Univ. Nac. Tucuman, ser A**. Three observations on linear algebra, journal, v.5, p. 147–151, 1946.

[8] BLANCHET, A.; CARLIER, G. **arXiv preprint arXiv:1206.6571**. Optimal transport and cournot-nash equilibria, journal, 2012.

[9] BOLLEY, F. **Separability and completeness for the wasserstein distance**. In: SÉMINAIRE DE PROBABILITÉS XLI, p. 371–377. Springer, 2008.

[10] BRENIER, Y. **Communications on pure and applied mathematics**. Polar factorization and monotone rearrangement of vector-valued functions, journal, v.44, n.4, p. 375–417, 1991.

[11] BREZIS, H. **Functional analysis, Sobolev spaces and partial differential equations**. Springer, 2010.

[12] BURAGO, D.; BURAGO, Y. ; IVANOV, S. **A course in metric geometry**, volume 33. American Mathematical Society Providence, 2001.

[13] BUTTAZZO, G.; SANTAMBROGIO, F. **SIAM journal on mathematical analysis**. A model for the optimal planning of an urban area, journal, v.37, n.2, p. 514–530, 2005.

[14] CARLIER, G.; EKELAND, I. **Economic Theory**. Matching for teams, journal, v.42, n.2, p. 397–418, 2010.

[15] CARLIER, G.; JIMENEZ, C. ; SANTAMBROGIO, F. **SIAM Journal on Control and Optimization**. Optimal transportation with traffic congestion and wardrop equilibria, journal, v.47, n.3, p. 1330–1350, 2008.

[16] CHOQUET, G. **Séminaire Bourbaki**. Existence et unicité des représentations intégrales au moyen des points extrémaux dans les cônes convexes, journal, v.4, p. 33–47, 1956.

[17] CULLEN, M.; DOUGLAS, R. **Contemporary Mathematics**. Applications of the monge-ampere equation and monge transport problem to meteorology and oceanography, journal, v.226, p. 33–54, 1999.

[18] DANERI, S.; SAVARÉ, G. **arXiv preprint arXiv:1009.3737**. Lecture notes on gradient flows and optimal transport, journal, 2010.

[19] EVANS, L. C. **Current developments in mathematics**. Partial differential equations and monge-kantorovich mass transfer, journal, p. 65–126, 1997.

[20] GANGBO, W. **Contemporary Mathematics**. The monge mass transfer problem and its applications, journal, v.226, p. 79–104, 1999.

[21] GANGBO, W.; MCCANN, R. J. **Acta Mathematica**. The geometry of optimal transportation, journal, v.177, n.2, p. 113–161, 1996.

[22] GUILLEN, N.; MCCANN, R. **arXiv preprint arXiv:1011.2911**. Five lectures on optimal transportation: Geometry, regularity and applications, journal, 2010.

[23] JORDAN, R.; KINDERLEHRER, D. ; OTTO, F. **SIAM journal on mathematical analysis**. The variational formulation of the fokker–planck equation, journal, v.29, n.1, p. 1–17, 1998.

[24] KLOECKNER, B. **arXiv preprint arXiv:0804.3505**. A geometric study of wasserstein spaces: Euclidean spaces, journal, 2008.

[25] KUHN, H. W. **Naval research logistics quarterly**. The hungarian method for the assignment problem, journal, v.2, n.1-2, p. 83–97, 1955.

[26] LOTT, J. **Communications in Mathematical Physics**. Some geometric calculations on wasserstein space, journal, v.277, n.2, p. 423–437, 2008.

[27] MCCANN, R. J. **arXiv preprint arXiv:1207.1867**. A glimpse into the differential topology and geometry of optimal transport, journal, 2012.

[28] PAPADOPOULOS, A. **Metric spaces, convexity and nonpositive curvature**, volume 6. European Mathematical Society, 2005.

[29] ROCKAFELLAR, R. T. **Pacific J. Math**. Characterization of the subdifferentials of convex functions, journal, v.17, n.3, p. 497–510, 1966.

[30] VAN GAANS, O. **Gradient flows in measure spaces**. Notes for "Topics in Analysis 2011", Leiden University, 2011.

[31] VILLANI, C. **Topics in optimal transportation**, volume 58. AMS Bookstore, 2003.

[32] VILLANI, C. **Optimal transport: old and new**, volume 338. Springer, 2009.

[33] WANG, X.-J. **Calculus of Variations and Partial Differential Equations**. On the design of a reflector antenna ii, journal, v.20, n.3, p. 329–341, 2004.