



INSTITUT
POLYTECHNIQUE
DE PARIS

IP Paris M2 Optimisation Internship, completed at the Centre de
Mathématiques Appliquées de l'École Polytechnique

Quantitative Stability in Discrete Optimal Transport

Author:

William Ford

Directors:

Michael Goldman et Cyril Letrouit

September 18, 2025

To my grandma Josie, the kindest woman I have ever known.

Acknowledgements

Tout d'abord, je présente mes remerciements à mes encadrants, Michael Goldman et Cyril Letrouit. Ça a été un plaisir de travailler avec vous pendant cette année, et j'ai hâte de continuer ce voyage avec vous durant mon doctorat.

Thanks to my parents, for their love, support and counsel, and who always pushed me to achieve my best. I do not comment on who gave me my “maths gene”. Thank you to Toby, with whom I first really fell in love with mathematics through our countless discussions.

I also want to thank all my friends who made my first year in Paris during this M2 so enjoyable. In particular, I would like to thank Aryna, Enzo, Ester, Francesca, Guido, Jacob, Jacopo, Jules, Michele, Pablo, Petra, Raphaël, Timur, Thomas, and Yegor. I probably would have done a lot more maths but had much less fun if it weren't for all of you.

Contents

1	Introduction	1
1.1	Overview	1
1.2	Background on optimal transport	1
1.3	Motivation for quantitative stability	3
1.4	Stability of transport plans	5
2	Stability with one measure fixed absolutely continuous	9
2.1	The Kantorovich functional	9
2.2	Glueing strong convexity inequalities	11
3	The discrete transport problem	14
3.1	Primal and dual polyhedra	14
3.2	The sets of optimal plans and potentials	18
4	Perturbation of support positions	23
4.1	Discussion on how to use linear programming connections	23
4.2	Structure of primal optimisers under deformations	25
5	Uniqueness of continuous Kantorovich potentials	31
5.1	Literature review	31
5.2	Uniqueness	32
6	Conclusions	39
	Bibliography	40
A	Appendix	42

1 Introduction

1.1 Overview

The optimal transport problem provides a rigorous framework for comparing probability measures, and has become a central tool in modern analysis and probability, as well as having found wide applications ranging from fluid mechanics and economics to data science and imaging. This report is concerned with the stability properties of optimal transport. While qualitative stability is well understood, quantitative estimates remain a subject of active research. Such estimates are crucial in applied contexts where data are inevitably discretised or noisy. The structure of the report is as follows:

In Chapter 1, we provide a brief summary of the optimal transport theory, then provide some motivation for the quantitative stability problem. Notably, we emphasise the necessity of studying the stability of *optimal plans* rather than maps. We formulate the main problem, providing some preliminary results in this direction.

In Chapter 2, we review some of the current literature on quantitative stability with a source fixed absolutely continuous, as well as providing some extensions of a certain result to a larger class of measures.

In Chapter 3, we shift towards the main theme of this report. Our approach is to begin to study the fully discrete setting, where the transport problem admits a linear programming formulation with a rich combinatorial structure.

In Chapter 4, we provide a (partial) analysis of quantitative stability properties of discrete optimal plans under perturbations.

In Chapter 5, we use the intuition given by the discrete transport problem to provide a sharper understanding of optimal transport in the general (non-discrete) case, significantly refining known uniqueness criteria for Kantorovich potentials (optimisers of the dual problem).

1.2 Background on optimal transport

Let ρ and μ be two probability measures supported inside $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^d$. Let $c : \mathcal{X} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ be lower semi-continuous and integrable in the sense that there exist $a \in L^1(\rho), b \in L^1(\mu)$ such that $c(x, y) \leq a(x) + b(y)$ pointwise. The optimal transport problem is the minimisation problem

$$\min_{\gamma \in \Pi(\rho, \mu)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\gamma(x, y), \quad (1.1)$$

where $\Pi(\rho, \mu)$ is the set of probabilities over $\mathbb{R}^d \times \mathbb{R}^d$ with first marginal ρ and second marginal μ . Under our hypotheses on c , the existence of solutions to (1.1) is classical, see [33, Theorem 1.3]. As a linear optimisation problem with convex constraints, problem (1.1) admits the following dual formulation

$$\max_{\substack{\phi, \psi \\ \phi \oplus \psi \leq c}} \int_{\mathcal{X}} \phi(x) d\rho(x) + \int_{\mathcal{Y}} \psi(y) d\mu(y). \quad (1.2)$$

Here the supremum is taken over all $\phi \in L^1(\mathcal{X}, \rho)$ and $\psi \in L^1(\mathcal{Y}, \mu)$ such that

$$\phi(x) + \psi(y) \leq c(x, y) \text{ for all } (x, y) \in \mathcal{X} \times \mathcal{Y}. \quad (1.3)$$

By a density argument (See Lusin's theorem [29, Theorem 2.24] and its corollary), the value of (1.2) is equal to if we took the supremum only over $\phi \in C_b(\mathcal{X})$ and $\psi \in C_b(\mathcal{Y})$. Under our hypotheses on c we have existence of optimisers for (1.2) in $L^1(\rho), L^1(\mu)$ (but not necessarily C_b), see [34, Theorem 5.10]. Given ϕ, ψ satisfying (1.3), we define their c -transforms as

$$\phi^c(y) := \inf_{x \in \mathcal{X}} c(x, y) - \phi(x); \quad \psi^c(x) := \inf_{y \in \mathcal{Y}} c(x, y) - \psi(y). \quad (1.4)$$

Functions which can be written as the c -transform of some function are referred to as c -concave. Take any (ϕ, ψ) admissible for (1.2), then

$$\phi(x) \leq \psi^c(x) \quad \forall x \in \mathcal{X} \quad (1.5)$$

pointwise as a consequence of (1.3). Hence, the pair (ψ^c, ψ) always performs at least as well while still respecting the constraint. Applying another c -transform yields $(\psi^c, (\psi^c)^c)$ which also performs no worse. However, this process cannot be iterated indefinitely, as $((\psi^c)^c)^c = \psi^c$. Now take any (ϕ, ψ) optimal for (1.2). Taking a c transform cannot increase the value of the objective in (1.2), thus

$$\int_{\mathcal{X}} \phi(x) d\rho(x) = \int_{\mathcal{X}} \psi^c(x) d\rho(x). \quad (1.6)$$

Combining (1.5) and (1.6), we deduce $\phi = \psi^c$ ρ -a.e. and hence without loss of generality, we can always choose representatives of any optimal potentials which are c -concave. The dual problem (1.2) is also often reformulated as the following so-called *semi-dual* problem

$$\sup_{\psi \in C_b(\mathcal{Y})} \int_{\mathcal{X}} \psi^c(x) d\rho(x) + \int_{\mathcal{Y}} \psi(y) d\mu(y). \quad (1.7)$$

Convex duality (said to be ‘‘Kantorovich Duality’’ in this case) tells us that the values of the primal and dual problems (1.1) and (1.2) are equal. Furthermore, for any $\gamma \in \Pi(\rho, \mu)$ optimal for (1.1) and any pair (ϕ, ϕ^c) optimal for (1.2),

$$\text{Supp } \gamma \subset \{(x, y) \in \mathcal{X} \times \mathcal{Y} : \phi(x) + \phi^c(y) = c(x, y)\}. \quad (1.8)$$

The set on the right-hand side is referred to as the graph of the c -superdifferential of ϕ , denoted $\partial^c \phi$, with a symmetric definition given for ψ . The c -superdifferential can equivalently be characterised by $y \in \partial^c \phi(x)$ if and only if

$$c(x, y) - \phi(x) \leq c(z, y) - \phi(z) \quad \forall z \in \mathcal{X}. \quad (1.9)$$

Furthermore, if for some admissible γ and (ϕ, ψ) , (1.8) holds, both are optimal for their respective problems.

A set $W \subset \mathcal{X} \times \mathcal{Y}$ is said to be c -cyclically monotone if for every finite sequence of points $(x_1, y_1), \dots, (x_n, y_n) \in W$ and any permutation σ of $\{1, \dots, n\}$,

$$\sum_{i=1}^n c(x_i, y_i) \leq \sum_{i=1}^n c(x_i, y_{\sigma(i)}) \quad (1.10)$$

When the cost is continuous $\gamma \in \Pi(\rho, \mu)$ is optimal for (1.1) if and only if $\text{Supp } \gamma$ is c -cyclically monotone, see [31, Corollary 1].

In the specific case $c(x, y) = \|x - y\|^p$ for some $p > 0$ and $\mathcal{X} = \mathcal{Y} = \mathbb{R}^d$, the p -th root of the value of the optimal transport problem between ρ and μ defines the p -Wasserstein distance

$$W_p(\rho, \mu) := \left(\inf_{\gamma \in \Pi(\rho, \mu)} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^p d\gamma(x, y) \right)^{1/p}.$$

When $p \geq 1$, W_p defines a metric on the space of probability measures with bounded p th moment. When $p = 2$, observing that

$$\int_{\mathcal{X} \times \mathcal{Y}} \|x - y\|^2 d\gamma = \int_{\mathcal{X}} \|x\|^2 d\rho + \int_{\mathcal{Y}} \|y\|^2 d\mu - 2 \int_{\mathcal{X} \times \mathcal{Y}} \langle x | y \rangle d\gamma. \quad (1.11)$$

The first two terms are independent of the choice of $\gamma \in \Pi(\rho, \mu)$ and so (1.1) is equivalent to maximising the correlation $\int \langle x | y \rangle d\gamma$, or using the cost $c(x, y) = -\langle x | y \rangle$. For this problem, Kantorovich duality reads

$$\max_{\gamma \in \Pi(\rho, \mu)} \int_{\mathcal{X} \times \mathcal{Y}} \langle x | y \rangle d\gamma = \min_{\psi \in \mathcal{C}^0(\mathcal{Y})} \int_{\mathcal{X}} \psi^* d\rho + \int_{\mathcal{Y}} \psi d\mu, \quad (1.12)$$

where here the $c(x, y) = -\langle x | y \rangle$ transform acts as convex conjugation, defined for $\psi \in \mathcal{C}^0(\mathcal{Y})$ by

$$\psi^*(x) = \sup_{y \in \mathcal{Y}} \langle x | y \rangle - \psi(y). \quad (1.13)$$

This is equivalent to extending ψ by $+\infty$ outside of $\mathcal{Y}$, then taking the usual Legendre transform with supremum over all of $\mathbb{R}^d$. Dual potentials φ for the cost $c(x, y) = \frac{1}{2}\|x - y\|^2$ and dual potentials ϕ for $c(x, y) = -\langle x | y \rangle$ are thus in bijection, related by

$$\phi(x) = \frac{1}{2}\|x\|^2 - \varphi(x).$$

In the case that ρ is absolutely continuous and both measures have bounded second moment, a theorem of Y. Brenier [9] asserts that the optimal transport plan $\gamma \in \Pi(\rho, \mu)$ is unique and induced by the map $T = \nabla \psi^*$, where ψ is a solution to the dual problem on the right hand side of (1.12). Here, induced means that $\gamma = (\text{Id}, T)_\# \rho$. Such an optimal ψ^* is henceforth referred to as a Brenier potential. If we assume further that $\text{Supp } \rho$ is the closure of a connected open set and $\text{Supp } \mu$ is bounded, the potentials ψ, ψ^* are respectively μ -a.e./ ρ -a.e. uniquely defined up to a constant [30, Proposition 7.18]. We present sharper conditions for uniqueness in Chapter 5.

1.3 Motivation for quantitative stability

Stability in optimal transport is important for applications, as well as being significant from a purely mathematical viewpoint. *Qualitative* stability in optimal transport is well understood, established as a consequence of compactness arguments. In particular, from [30, Section 1.6.4]

Theorem 1.1. *Let $\mathcal{X}$ and $\mathcal{Y}$ be compact metric spaces, let $\rho_n \in \mathcal{P}(\mathcal{X})$ and $\mu_n \in \mathcal{P}(\mathcal{Y})$ be sequences of probability measures converging weakly (in duality with continuous bounded functions) to limits $\rho_n \rightharpoonup \rho$ and $\mu_n \rightharpoonup \mu$. Then, up to a subsequence*

- *Optimal plans γ_n between ρ_n and μ_n converge weakly to an optimal plan between ρ and μ .*
- *If they exist, for $\rho_n \equiv \rho$ fixed, optimal maps T_n from ρ to μ_n converge strongly in $L^2(\rho)$ to the optimal map T between ρ and μ .*

- Optimal dual potentials between ρ_n and μ_n converge uniformly to an optimal pair for ρ and μ .
- $W_p(\rho_n, \mu_n) \rightarrow W_p(\rho, \mu)$.

The question of *quantitative* stability in optimal transport has been the subject of numerous recent works. The case where one measure is absolutely continuous has received a lot of interest recently, for example [15, 13, 22, 18]. Here, the source measure ρ is fixed as an absolutely continuous measure with support $\mathcal{X} \subset \mathbb{R}^d$, and the target measure μ is perturbed. The authors derive results of the form

$$\|T_{\rho \rightarrow \mu_0} - T_{\rho \rightarrow \mu_1}\|_{L^2(\rho)} \leq C_{\rho, \mathcal{Y}} W_1(\mu_0, \mu_1)^q \text{ for all } \mu_0, \mu_1 \in \mathcal{P}(\mathcal{Y}), \quad (1.14)$$

for some $q \leq 1$. Here $T_{\rho \rightarrow \mu_i}$ represents the unique quadratic optimal transport map from ρ to μ_i . The best exponent currently proven here is $q = 1/6$ in [13]. The following example illustrates that we cannot hope for better than $q = 1/2$ in general.

Example 1.2. [12, Lemma 5.2] Let $\rho(x) = \frac{1}{\pi} \chi_{B_1}(x)$ as the uniform probability measure on the unit disc in $\mathbb{R}^2$. Let $x_\theta = (\sin \theta, \cos \theta)$, then define $\mu_\theta = \frac{1}{2} \delta_{x_\theta} + \frac{1}{2} \delta_{x_{\theta+\pi}}$. In this case, the optimal transport between ρ and μ_θ is given by

$$T_{\rho \rightarrow \mu_\theta}(x) = \begin{cases} x_\theta & \text{if } \langle x | x_\theta \rangle \geq 0 \\ x_{\theta+\pi} & \text{if } \langle x | x_\theta \rangle < 0 \end{cases}$$

see Figure 1.1. We can calculate explicitly the quantities involved. For small θ ,




Figure 1.1: The optimal map $T_{\rho \rightarrow \mu_\theta}$. Taken from [21, Figure 1].

$$W_2(\mu_0, \mu_\theta) = 2 \sin(\theta/2) \sim |\theta|$$

for any $p \geq 1$. Also,

$$\|T_{\rho \rightarrow \mu_\theta} - T_{\rho \rightarrow \mu_0}\|_{L^2(\rho)}^2 = \left(1 - \frac{|\theta|}{\pi}\right) |2 \sin(\theta/2)|^2 + \frac{|\theta|}{\pi} |2 \sin((\pi - \theta)/2)|^2 \sim \frac{4|\theta|}{\pi}.$$

Hence, we cannot hope for better than $q = 1/2$ in (1.14).

Current proof techniques actually do not establish (1.14) directly. Instead, they first establish a similar statement for optimisers of the dual problem (1.2), of the form

$$\|\phi_{\rho \rightarrow \mu_0} - \phi_{\rho \rightarrow \mu_1}\|_{L^2(\rho)} \leq C_{\rho, \mathcal{Y}} W_1(\mu_0, \mu_1)^q \text{ for all } \mu_0, \mu_1 \in \mathcal{P}(\mathcal{Y}). \quad (1.15)$$

Here $\phi_{\rho \rightarrow \mu_i}$ denotes the Brenier potential between ρ and μ_i , satisfying $(\int \phi d\rho = 0)$. This constraint is imposed to guarantee the uniqueness of potentials (they are unique up to a constant here without this additional constraint). We give a slightly more detailed overview of this theory in Chapter 2, for an in depth exposition we refer the reader to [21].

1.4 Stability of transport plans

Stability beyond when one measure is absolutely continuous has received comparatively little study. [24] treats stability in both marginals in a neighbourhood of measures whose optimal map between them is Lipschitz, slightly extending the result of [15]. Such regularity assumptions on optimal maps are hard to verify in applications. Moreover, they are often not true - Brenier's theorem does not apply, so in many cases we may not even have the existence of optimal maps, let alone regular ones. It is usually more reasonable to have some regularity assumptions on one of the measures.

Problem 1. Let ρ satisfy some regularity assumptions, for example, of the form

- (i) *Continuous*: ρ has density comparable to Lebesgue measure of a bounded compact set.
- (ii) *Discrete*: ρ is a fully discrete measure, given by the projection of an absolutely continuous measure onto a uniform grid mesh placed over its support.
- (iii) *Rough*: ρ is comparable to Lebesgue measure “up to resolution ε ” in the sense of [17, Assumption 1], in that there exist constants $0 < \lambda \leq \Lambda$ such that for any ball $B_r \subset \text{Supp } \rho$ with $r \geq \varepsilon$, we have

$$\lambda|B_r| \leq \rho(B_r) \leq \Lambda|B_r|.$$

For what constants C and exponents q are the optimal transport plans stable in the following sense: There exists a “small” error $\mathcal{E}$ depending on (i) the datum of ρ (ii) the mesh grid size or (iii) the resolution ε , such that

$$W_p(\gamma_{\rho \rightarrow \mu_0}, \gamma_{\rho \rightarrow \mu_1}) \leq CW_p(\mu_0, \mu_1)^q + \mathcal{E}? \quad (1.16)$$

Note that on the left-hand side, W_p is taken with respect to the Euclidean distance on the product space $\mathbb{R}^d \times \mathbb{R}^d$, while those on the right-hand side are for the Euclidean distance on $\mathbb{R}^d$. We do not distinguish between these cases in notation, as it is clear from the context.

In case (iii), $\mathcal{E}$ should depend on ε in a way which vanishes as $\varepsilon \rightarrow 0$. Understanding the sharp behaviour in this case should simultaneously explain (ii), and this in turn should induce a result for (i) as a consequence of qualitative stability arguments. It is also possible that a thorough study of (ii) could lead to (iii) by qualitative limits. The advantage of studying the discrete formulation is that the Laguerre cells induced on the target domain by choosing a dual potential for ρ have a very nice structure generally here, as will be further discussed in Chapter 3. One can also perturb both marginals at the same time, with an additional $W_p(\rho_0, \rho_1)$ term on the right-hand side of (1.16).

We make three immediate observations regarding the stability of optimal plans. The first is that a form of “reverse stability” of plans always holds.

Proposition 1.3. *Let $p \in (1, \infty)$. Let ρ_0, ρ_1, μ_0 and μ_1 be probabilities on $\mathbb{R}^d$ with finite p -moments. Let $\gamma_i \in \Pi(\rho_i, \mu_i)$ be optimal transport plans between their marginals for the p -cost. Then*

$$W_p^p(\gamma_0, \gamma_1) \geq c_p(W_p^p(\rho_0, \rho_1) + W_p^p(\mu_0, \mu_1)).$$

where $c_p = \min(1, 2^{p/2-1})$.

Proof. Set $\eta \in \Pi(\gamma_0, \gamma_1) \subset \Pi(\rho_0, \mu_0, \rho_1, \mu_1)$ as an optimal p -cost transport on $\mathbb{R}^d \times \mathbb{R}^d$ between γ_0 and γ_1 . The projections onto first and third or second and fourth coordinates give couplings between ρ_0 and ρ_1 or μ_0 and μ_1 respectively. These are candidates for the respective transport

problems, so

$$\begin{aligned}
W_p^p(\gamma_0, \gamma_1) &= \int_{\mathbb{R}^{2d} \times \mathbb{R}^{2d}} \|(x_0, y_0) - (x_1, y_1)\|^p d\eta(x_0, y_0, x_1, y_1) \\
&= \int_{\mathbb{R}^{2d} \times \mathbb{R}^{2d}} \left(\sum_{i=1}^d (x_{0,i} - x_{1,i})^2 + \sum_{i=1}^d (y_{0,i} - y_{1,i})^2 \right)^{p/2} d\eta(x_0, y_0, x_1, y_1) \\
&\geq c_p \int_{\mathbb{R}^{2d}} \|x_0 - x_1\|^p d\pi_{13\#}\eta + c_p \int_{\mathbb{R}^{2d}} \|y_0 - y_1\|^p d\pi_{24\#}\eta \geq W_p^p(\rho_0, \rho_1) + W_p^p(\mu_0, \mu_1).
\end{aligned}$$

where we used the inequality $(a+b)^{p/2} \geq c_p(a^{p/2} + b^{p/2})$ with $c_p = \min(1, 2^{\frac{p}{2}-1})$. $\square$

The second observation is that “plans are always at least as stable as maps”. Fix ρ and two targets μ_0, μ_1 . Assume the existence of maps T_i for the p -transport between ρ and μ_i and set $\gamma_{\rho \rightarrow \mu_i} = (\text{Id}, T_{\rho \rightarrow \mu_i})_{\#}\rho \in \Pi(\rho, \mu_i)$ as the optimal plans induced by them. Then

$$W_p(\gamma_{\rho \rightarrow \mu_0}, \gamma_{\rho \rightarrow \mu_1}) \leq \|T_0 - T_1\|_{L^p(\rho)}. \quad (1.17)$$

To see this, it suffices to observe that $(\text{Id}, T_{\rho \rightarrow \mu_0}, \text{Id}, T_{\rho \rightarrow \mu_1})_{\#}\rho$ is an admissible coupling between $\gamma_{\rho \rightarrow \mu_0}$ and $\gamma_{\rho \rightarrow \mu_1}$, and its transport cost is precisely $\|T_0 - T_1\|_{L^p(\rho)}^p$. Consequently, all the results of [13, 22] pass directly to plans, so that for the continuous case (i), $\mathcal{E} = 0$ with $q = 1/6$, under suitable assumptions on ρ and the target support.

The third observation is related to the following natural question: Is the worst-case behaviour of plans any better than that of maps? Consider again Example 1.2, except now we are interested in the Wasserstein distance between the optimal plans $\gamma_{\rho \rightarrow \mu_\theta} = (\text{Id}, T_{\rho \rightarrow \mu_\theta})_{\#}\rho$. Consider the following coupling $\eta \in \Pi(\gamma_{\rho \rightarrow \mu_0}, \gamma_{\rho \rightarrow \mu_1})$, defined by

$$\eta = (\text{Id}, T_{\rho \rightarrow \mu_0}, R_\theta, R_\theta T_{\rho \rightarrow \mu_0}) \in \Pi(\gamma_{\rho \rightarrow \mu_0}, \gamma_{\rho \rightarrow \mu_\theta}),$$

where R_θ is the matrix representing rotation anticlockwise by θ in $\mathbb{R}^2$. This is a valid coupling since $(R_\theta, R_\theta)_{\#}\gamma_{\rho \rightarrow \mu_0} = \gamma_{\rho \rightarrow \mu_\theta}$. The coupling η gives an upper bound on the W_2 distance between the two plans. In particular,

$$\begin{aligned}
W_2^2(\gamma_{\rho \rightarrow \mu_0}, \gamma_{\rho \rightarrow \mu_\theta}) &\leq \int_{\mathbb{R}^2 \times \mathbb{R}^2} \|x_0 - x_1\|^2 + \|y_0 - y_1\|^2 d\eta(x_0, y_0, x_1, y_1) \\
&= \int_{\mathbb{R}^2} \|x - R_\theta x\|^2 + \|T_{\rho \rightarrow \mu_0}(x) - R_\theta T_{\rho \rightarrow \mu_0}(x)\|^2 d\rho(x) \\
&= \frac{1}{\pi} \int_0^1 \int_0^{2\pi} \left(2r \sin(\theta/2)\right)^2 r d\phi dr + \left(2 \sin(\theta/2)\right)^2 \\
&= \left(\frac{1}{2} + 1\right) \left(2 \sin(\theta/2)\right)^2 = \frac{3}{2} W_2^2(\mu_0, \mu_\theta)
\end{aligned}$$

since $\|T_{\rho \rightarrow \mu_0}(x) - R_\theta T_{\rho \rightarrow \mu_0}(x)\| = 2 \sin(\theta/2)$ ρ -a.e. Hence, in this case, the plans exhibit Lipschitz behaviour, better than the Hölder $1/2$ of the maps. The argument below demonstrates that plans cannot be Lipschitz in general, but leaves the door open to exponents in $(1/2, 1)$.

Lemma 1.4. *For fixed ρ and two targets μ_0, μ_1 , all compactly supported. Set $\gamma_i \in \Pi(\rho, \mu_i)$ as the optimal plans induced by optimal maps T_i . Let*

$$c_\varepsilon(x_0, y_0, x_1, y_1) = \|x_0 - x_1\|^2 + \varepsilon \|y_0 - y_1\|^2,$$

and denote

$$\tau_{c_\varepsilon}(\gamma_0, \gamma_1) = \inf_{\eta \in \Pi(\gamma_0, \gamma_1)} \int_{\mathbb{R}^4} c_\varepsilon(x_0, y_0, x_1, y_1) d\eta(x_0, y_0, x_1, y_1) \quad (1.18)$$

as the transport cost between plans with respect to c_ε . Then

$$\frac{1}{\varepsilon} \tau_{c_\varepsilon}(\gamma_0, \gamma_1) \rightarrow \|T_0 - T_1\|_{L^2(\rho)}^2,$$

and optimal plans η_ε for (1.18) converge weakly to $(\text{Id}, T_0, \text{Id}, T_1)_\# \rho$.

Proof. This is simply a consequence of the stability of plans with respect to uniform convergence of costs. We present here a self-contained proof of this fact. Consider optimisers $\eta_\varepsilon \in \Pi(\gamma_0, \gamma_1)$ for the c_ε problem. By compactness, we can assume these converge weakly up to a subsequence to some $\eta_0 \in \Pi(\gamma_0, \gamma_1)$. Then strong-weak convergence gives

$$\int_{\mathbb{R}^{2d}} c_\varepsilon d\eta_\varepsilon \rightarrow \int_{\mathbb{R}^{2d}} c_0 d\eta_0.$$

But we also have

$$|\tau_{c_0}(\gamma_0, \gamma_1) - \tau_{c_\varepsilon}(\gamma_0, \gamma_1)| \leq \|c_0 - c_\varepsilon\|_\infty,$$

and hence $\int c_0 d\eta_0 = \tau_{c_0}(\gamma_0, \gamma_1)$. In other words, the limit η_0 is optimal. But this optimiser must precisely be

$$\eta_0 = (\text{Id}, T_0, \text{Id}, T_1)_\# \rho$$

since it has zero cost for $c_0 = \|x_0 - x_1\|^2$, and this is the only way to glue through the identity (since we have existence of maps). Hence

$$\int_{\mathbb{R}^{4d}} \|y_0 - y_1\|^2 d\eta_\varepsilon(x_0, y_0, x_1, y_1) \leq \frac{1}{\varepsilon} \tau_{c_\varepsilon}(\gamma_0, \gamma_1) \leq \|T_0 - T_1\|_{L^2(\rho)}^2,$$

where the first inequality is by monotonicity of the cost $\|y_0 - y_1\|^2 \leq \frac{1}{\varepsilon} c_\varepsilon(x_0, y_0, x_1, y_1)$ and the fact that η_ε is optimal for c_ε , and the second inequality is a consequence of the plan $(\text{Id}, T_0, \text{Id}, T_1)_\# \rho$ as a competitor for the c_ε problem. Passing to the limit $\varepsilon \rightarrow 0$, the left hand side converges to $\|T_0 - T_1\|_{L^2(\rho)}^2$ by weak convergence, and hence so does $\frac{1}{\varepsilon} \tau_{c_\varepsilon}(\gamma_0, \gamma_1)$ by squeezing. $\square$

Proposition 1.5 (Optimal plans cannot be Lipschitz). *Fix ρ a probability density uniform on a compact set with non-empty interior in $\mathbb{R}^d$. Then, given a compact set $\mathcal{Y}$ in $\mathbb{R}^d$ with non empty interior, there **does not exist** $C_{\rho, \mathcal{Y}} > 0$ such that for all μ_0, μ_1 probabilities on $\mathcal{Y}$,*

$$W_2(\gamma_{\rho \rightarrow \mu_0}, \gamma_{\rho \rightarrow \mu_1}) \leq C W_2(\mu_0, \mu_1). \quad (1.19)$$

where here $\gamma_{\rho \rightarrow \mu_i} = (\text{Id}, T_{\rho \rightarrow \mu_i})_\# \rho$ are the plans induced by the optimal maps from $\rho \rightarrow \mu_i$ for quadratic cost.

Proof. Since $\mathcal{Y}$ has non-empty interior, it contains a closed ball. By invariance of optimisers to the quadratic OT problem under dilations and translations, we may assume $\mathcal{Y} = \overline{B}(0, 1) =: B_1$. Assume that (1.19) holds on B_1 with constant $C > 0$. Given $\mu_0, \mu_1 \in \mathcal{P}(B_1)$, we define the contractions

$$\mu_i^\varepsilon := p_\varepsilon \# \mu_i \quad \text{where } p_\varepsilon(y) = \varepsilon y.$$

We have $\mu_i^\varepsilon \in \mathcal{P}(B_1)$ and hence

$$W_2(\gamma_{\rho \rightarrow \mu_0^\varepsilon}, \gamma_{\rho \rightarrow \mu_1^\varepsilon}) \leq C W_2(\mu_0^\varepsilon, \mu_1^\varepsilon). \quad (1.20)$$

A change of variables shows that

$$W_2(\mu_0^\varepsilon, \mu_1^\varepsilon) = \varepsilon W_2(\mu_0, \mu_1)$$

since $\pi \in \Pi(\mu_0, \mu_1)$ if and only if $(\varepsilon y_0, \varepsilon y_1)_{\#} \pi \in \Pi(\mu_0^\varepsilon, \mu_1^\varepsilon)$. Furthermore,

$$\eta \in \Pi(\gamma_{\rho \rightarrow \mu_0}, \gamma_{\rho \rightarrow \mu_1}) \iff (x_0, \varepsilon y_0, x_1, \varepsilon y_1)_{\#} \eta \in \Pi(\gamma_{\rho \rightarrow \mu_0^\varepsilon}, \gamma_{\rho \rightarrow \mu_1^\varepsilon})$$

and so

$$\begin{aligned} W_2^2(\gamma_{\rho \rightarrow \mu_0^\varepsilon}, \gamma_{\rho \rightarrow \mu_1^\varepsilon}) &= \inf_{\eta \in \Pi(\gamma_{\rho \rightarrow \mu_0}, \gamma_{\rho \rightarrow \mu_1})} \int \|x_0 - x_1\|^2 + \|\varepsilon y_0 - \varepsilon y_1\|^2 d\eta(x_0, y_0, x_1, y_1) \\ &= \tau_{c_{\varepsilon^2}}(\gamma_{\rho \rightarrow \mu_0}, \gamma_{\rho \rightarrow \mu_1}). \end{aligned}$$

Consequently, (1.20) becomes

$$\tau_{c_{\varepsilon^2}}(\gamma_{\rho \rightarrow \mu_0}, \gamma_{\rho \rightarrow \mu_1}) \leq C^2 \varepsilon^2 W_2^2(\mu_0, \mu_1).$$

Dividing by ε^2 and taking $\varepsilon \rightarrow 0$ using Lemma 1.4 followed by a square root implies

$$\|T_{\rho \rightarrow \mu_0} - T_{\rho \rightarrow \mu_1}\|_{L^2(\rho)} \leq C W_2(\mu_0, \mu_1) \text{ for all } \mu_0, \mu_1 \in \mathcal{P}(B_1).$$

But this is known to be impossible: Example 1.2 contradicts this. Thus, there can be no $C_{\rho, \mathcal{Y}} > 0$ such that (1.19) holds. $\square$

Remark 1.6. If one could quantify the convergence in Lemma 1.4, then the cost c_ε could potentially be used to invert inequalities of the form (1.17) and, given a stability result about plans, infer back to stability about the maps.

The principal aim of this report is to begin to treat quantitative stability beyond the case where one measure is fixed absolutely continuous. This is important for applications, as one might often work with a source that is a discretisation of an absolutely continuous measure, using a very fine mesh. To the best of our knowledge, there are no known techniques to treat stability with this type of source measure. Our approach is, in leaving behind absolutely continuous measures, to go to the other extreme. We study the quantitative stability properties of the fully discrete transport problem, aiming to recover statements for general target measures as the result of a qualitative limit. In the fully discrete case, the optimal transport problem becomes a linear program, which gives us another lens through which to understand the stability.

2 Stability with one measure fixed absolutely continuous

In the first half of this chapter, we review some of the literature on the quantitative stability of the quadratic optimal transport problem with one measure fixed and absolutely continuous. The first result of this form without regularity hypotheses on the maps was attained in [7]. In [13], techniques were developed to attain stability exponents in (1.14) that are independent of the dimension d . These have been further used to extend the same results to a larger class of source measures in [22], and to Riemannian manifolds in [18]. In general, these techniques have been designed to treat the quadratic cost, with some extensions to p -costs in [26]. For a more complete account of this theory, see [21].

In the second half of the chapter, we provide an extension of a certain strong convexity inequality to a larger class of measures, using techniques developed in [22]. This extends the results on the quantitative stability of Wasserstein barycentres in [11] to a much larger class of absolutely continuous measures.

2.1 The Kantorovich functional

Fix ρ to be an absolutely continuous probability measure with support $\mathcal{X}$, which is the closure of a connected open bounded set in $\mathbb{R}^d$. Fix also $\mathcal{Y} \subset \mathbb{R}^d$ compact. To study the stability of Brenier potentials, we consider the semi-dual problem

$$\min_{\psi \in C_b(\mathcal{Y})} \int_{\mathcal{X}} \psi^*(x) d\rho(x) + \int \psi(y) d\mu(y), \quad (2.1)$$

where ψ^* represents the Legendre transform defined in (1.13). This problem is a convex optimisation problem over $C_b(\mathcal{Y})$, and hence (informally) stability of minimisers is equivalent to strong convexity of the objective (2.1). The second term is linear, so we can focus on the first term, which motivates the following definition. The Kantorovich functional is defined as

$$\mathcal{K}_\rho : C_b(\mathcal{Y}) \rightarrow \mathbb{R} \quad \psi \mapsto \int_{\mathcal{X}} \psi^*(x) d\rho(x),$$

where the Legendre transform is taken as in (1.13). First introduced in [19] to study the convergence of algorithms for semi-discrete optimal transport, the Kantorovich functional is the central tool for the stability results presented in this chapter. The following proposition is a minor generalisation of [12, Lemma 1.8], providing a characterisation of the subdifferential of $\mathcal{K}_\rho$.

Proposition 2.1 (Characterisation of the Subdifferential). *$\mathcal{K}_\rho$ is a convex, 1-Lipschitz functional on $C_b(\mathcal{Y})$ equipped with $\|\cdot\|_\infty$ norm. The subdifferential of $\mathcal{K}_\rho$ is given by*

$$\partial \mathcal{K}_\rho(\psi) = \{-p_{Y\#}\gamma \mid \gamma \in \mathcal{P}(\mathcal{X} \times \mathcal{Y}) \text{ with } p_{X\#}\gamma = \rho \text{ and } \text{Supp } \gamma \subset \partial\psi\}. \quad (2.2)$$

where $p_Y : \mathcal{X} \times \mathcal{Y} \rightarrow \mathcal{Y}$ denotes the canonical projection onto $\mathcal{Y}$ and $\partial\psi \subset \mathcal{X} \times \mathcal{Y}$ denotes the graph on the subdifferential of ψ^* .

Proof. The subdifferential is a priori valued in the topological dual $C_b(\mathcal{Y})^*$, the space of finite signed regular Borel measures. Since the Legendre transform is a 1-Lipschitz operator on $C_b(\mathcal{Y})$, and it is pointwise convex, $\mathcal{K}_\rho$ is a 1-Lipschitz convex functional on $C_b(\mathcal{Y})$. Hence, by subdifferential rules for convex functions (see [2, Corollary 2.1]),

$$-\mu \in \partial\mathcal{K}_\rho(\psi) \iff \psi \in \arg \min \int_{\mathcal{X}} \psi^* d\rho + \int_{\mathcal{Y}} \psi d\mu. \quad (2.3)$$

We now show that any such μ should be a positive probability measure. Assume first $\langle 1 | \mu \rangle = m \neq 1$. Then for any $\psi \in C_b(\mathcal{Y})$, replacing with $\psi + \lambda$ for some $\lambda \in \mathbb{R}$ gives

$$\int_{\mathcal{X}} (\psi + \lambda)^* d\rho + \int_{\mathcal{Y}} \psi + \lambda d\mu = \int_{\mathcal{X}} \psi^* d\rho + \int_{\mathcal{Y}} \psi d\mu + \lambda(m - 1)$$

and so taking $\lambda \rightarrow \pm\infty$ depending on the sign of $m - 1$ gives that no argmin can exist. Thus, any element of $\partial\mathcal{K}_\rho$ should have total mass -1 . Similarly, assume there exists some measurable $A \subset \mathcal{Y}$ for which $\langle \chi_A | \mu \rangle = m < 0$. As mentioned in the introduction, by a density argument (see [29, Theorem 2.24]) (2.1) has the same value as if we instead took the infimum over all $\psi \in L^1(\mu)$. Then, if $\psi \in C_b(\mathcal{Y})$ is optimal, it is also optimal over all $L^1(\mu)$. Considering $\psi + \lambda\chi_A$ for $\lambda > 0$, by monotonicity of the Legendre transform

$$\int_{\mathcal{X}} (\psi + \lambda\chi_A)^* d\rho + \int_{\mathcal{Y}} \psi + \lambda\chi_A d\mu \leq \int_{\mathcal{X}} \psi^* d\rho + \int_{\mathcal{Y}} \psi d\mu + \lambda m$$

from which again ψ cannot belong to the argmin. Hence, any subdifferential element is a *negative* measure of mass -1 .

Now knowing that the only measures in $\partial\mathcal{K}_\rho$ are negative probability measures, by (2.3), combined with the compatibility condition (1.8), we arrive at (2.2). $\square$

Remark 2.2. In general, $\mathcal{K}_\rho$ is never strongly convex, as it is affine along the path $\psi + t$ for $t \in \mathbb{R}$. Paths corresponding to the same dual potential up to a constant. To avoid this, we fix the constants of the potentials so that the Brenier potential is mean free ($\int \psi^* d\rho = 0$). Since $\text{Supp } \rho$ is the closure of a connected open set, [30, Proposition 7.18] tells us that mean free potentials are unique in this setting.

We now turn to strong convexity estimates for $\mathcal{K}_\rho$, which in turn will imply quantitative stability of Brenier potentials.

Theorem 2.3. [13, Theorem 2.1][22, Theorem 2.1] *Let $Q \subset \mathbb{R}^d$ be a compact convex set with non-empty interior, and let ρ be a probability density over Q satisfying $M_\rho \geq \rho \geq m_\rho > 0$. Then for all $\psi_0, \psi_1 \in C_b(\mathcal{Y})$ we have*

$$\text{Var}_\rho(\psi_1^* - \psi_0^*) \leq C_\psi \langle \psi_1^* - \psi_0^* | (\nabla \psi_0^*)_\# \rho - (\nabla \psi_1^*)_\# \rho \rangle$$

where $C_\psi = e^{\frac{M_\rho}{m_\rho}}(M_\phi - m_\phi)$ for $m_\phi \leq \psi_i^* \leq M_\phi$.

Here, we make explicit the dependence of the above constant on ψ , in terms of the oscillation of the Brenier potentials $(M_\phi - m_\phi)$. When both $\mathcal{X}$ and $\mathcal{Y}$ are bounded, we can uniformly bound this, and also when $\mathcal{X}$ is bounded and all target measures satisfy a uniform p th moment bound for $p > d$, as demonstrated in [7, 13]. The above strong convexity directly implies the following quantitative stability estimates for both maps and potentials.

Corollary 2.3.1. [13] *Let $\mathcal{X} \subset \mathbb{R}^d$ be a compact convex set with non-empty interior, and let ρ be a probability density over $\mathcal{X}$ satisfying $M_\rho \geq \rho \geq m_\rho > 0$. Then for any $\mathcal{Y} \subset \mathbb{R}^d$ compact, there exists $C_{\rho, \mathcal{Y}} > 0$ such that*

$$\|\psi_1^* - \psi_0^*\|_{L^2(\rho)} \leq CW_1(\mu_0, \mu_1)^{1/2} \quad \forall \mu_0, \mu_1 \in \mathcal{P}(\mathcal{Y}), \quad (2.4)$$

where ψ_i are optimal for (2.1) between ρ and μ_i .

Proof. We observe that in the preceding theorem, the constant C depended on ψ only by the oscillation $M_\psi - m_\psi$. By definition of the Legendre transform restricted to $\mathcal{Y}$,

$$\psi^*(x) = \sup_{y \in \mathcal{Y}} \langle x | y \rangle - \psi(y)$$

is a supremum of $R_{\mathcal{Y}}$ Lipschitz functions $x \mapsto \langle x | y \rangle - \psi(y)$ where $R_{\mathcal{Y}}$ is the radius of the smallest ball containing $\mathcal{Y}$. Hence, ψ^* is $R_{\mathcal{Y}}$ Lipschitz, and its oscillation over $\mathcal{X}$ is bounded by

$$M_\psi - m_\psi \leq (\text{Diam } \mathcal{X}) R_{\mathcal{Y}}.$$

Furthermore, due to the invariance of the quadratic optimal transport problem to translations, translating $\mathcal{Y}$ allows us to replace $R_{\mathcal{Y}}$ with $\text{Diam } \mathcal{Y}$. Thus C_ψ in (2.6) is uniformly bounded in ψ , depending only on ρ and $\mathcal{Y}$.

Secondly, by Brenier's theorem, $(\nabla \psi_i^*)_{\#} \rho = \mu_i$ so that applying Theorem 2.3 gives

$$\|\psi_1^* - \psi_0^*\|_{L^2(\rho)}^2 = \text{Var}_\rho(\psi_1^* - \psi_0^*) \leq C \langle \psi_0 - \psi_1 | \mu_1 - \mu_0 \rangle,$$

where the first equality is due to our choice of mean free potentials. Finally, since $\psi_1^* - \psi_0^*$ is $2 \text{Diam } \mathcal{Y}$ Lipschitz continuous, Kantorovich Rubinstein duality gives

$$\|\psi_1^* - \psi_0^*\|_{L^2(\rho)} \leq CW_1(\mu_0, \mu_1)^{1/2}$$

as required. □

Corollary 2.3.2. [13] *Let $\mathcal{X} \subset \mathbb{R}^d$ be a compact convex set with non-empty interior and $\mathcal{H}^{d-1}$ rectifiable boundary, and let ρ be a probability density over $\mathcal{X}$ satisfying $M_\rho \geq \rho \geq m_\rho > 0$. Then for any $\mathcal{Y} \subset \mathbb{R}^d$ compact, there exists $C_{\rho, \mathcal{Y}} > 0$ such that*

$$\|T_1 - T_0\|_{L^2(\rho)} \leq CW_1(\mu_0, \mu_1)^{\frac{1}{6}} \quad \forall \mu_0, \mu_1 \in \mathcal{P}(\mathcal{Y}), \quad (2.5)$$

where T_i is the unique optimal transport map between ρ and μ_i .

Proof. One combines [13, Proposition 4.1] with Corollary 2.3.1, using again that the Lipschitz constants of ψ_0^* and ψ_1^* are uniformly bounded. □

2.2 Glueing strong convexity inequalities

The proof of Theorem 2.3 allows for a slightly stronger inequality presented below, which has found other usage, for example, in studying the stability of Wasserstein Barycentres in [11] and for the optimal matching problem in [16]. In particular, we have

Theorem 2.4. [11, Appendix B], [13] *Under the same hypotheses as Theorem 2.3, for any $\mathcal{Y} \subset \mathbb{R}^d$ compact there exists $C_{\rho, \mathcal{Y}} > 0$ such that for all $\psi_0, \psi_1 \in C_b(\mathcal{Y})$,*

$$\text{Var}_\rho(\psi_1^* - \psi_0^*) \leq C \left(\mathcal{K}_\rho(\psi_1) - \mathcal{K}_\rho(\psi_0) + \langle \psi_1 - \psi_0 | (\nabla \psi_0^*)_{\#} \rho \rangle \right). \quad (2.6)$$

In [22, 18], the authors treat stability for a much larger class of absolutely continuous source measures ρ . One decomposes the source domain into cubes, upon each of which we can apply Theorem 2.4. Then, under regularity assumptions on ρ , it is shown that the sum of stability errors locally over each cube can be combined to give a stability result globally across the domain. These techniques were applied to deduce the stability of maps and potentials for more general source measures. However, the techniques here did not directly establish (2.6) for the same larger class of ρ . Using the same glueing techniques, we obtain a strong convexity inequality (2.6) for John domains, which, to our knowledge, is new in this generality. We first review the glueing theory used in [22].

Definition 2.5 (Boman Chain condition). [8] *A probability measure ρ on an open set $\mathcal{X} \subset \mathbb{R}^d$ satisfies the Boman chain condition with parameters $A, B, C > 1$ if there exists a covering $\mathcal{F}$ of $\mathcal{X}$ of open cubes $Q \in \mathcal{F}$ such that*

- for $\sigma = \min(10/9, (d+1)/d)$ and for any $x \in \mathbb{R}^d$,

$$\sum_{Q \in \mathcal{F}} \chi_{\sigma Q}(x) \leq A \chi_{\mathcal{X}}(x). \quad (2.7)$$

- For some fixed cube Q_0 in $\mathcal{F}$, called the central cube, and for every $Q \in \mathcal{F}$, there exists a chain $Q_0, Q_1, \dots, Q_N = Q$ of distinct cubes from $\mathcal{F}$ such that for any $j \in \{0, \dots, N-1\}$,

$$Q \subset BQ_j \quad (2.8)$$

Where BQ_j is the cube with the same centre as Q_j and side length multiplied by B .

- Consecutive cubes of the above chain rho-verlap quantitatively: for any $j \in \{0, \dots, N-1\}$,

$$\rho(Q_j \cap Q_{j+1}) \geq C^{-1} \max(\rho(Q_j), \rho(Q_{j+1})). \quad (2.9)$$

Note that the length N of the chain can depend on $Q \in \mathcal{F}$ and need not be uniformly bounded over $Q \in \mathcal{F}$. The following lemma is the crux of the argument, providing a way to combine local stability on each cube of a Boman decomposition into a global stability result.

Lemma 2.6. [22, Lemma 3.3] *Let ρ be a probability density over a domain $\mathcal{X}$ satisfying the Boman chain condition for some covering $\mathcal{F}$ and some $A, B, C > 1$. Assume moreover that there exists $D > 0$ such that*

$$\forall Q \in \mathcal{F}, \quad \rho(5B\sqrt{d}) \leq D\rho(Q). \quad (2.10)$$

Then setting $\tilde{\rho}_Q = \frac{1}{\rho(Q)}\rho|_Q$, for any continuous function f on $\mathcal{X}$, it holds

$$\text{Var}_\rho(f) \leq 200A^2CD^3 \sum_{Q \in \mathcal{F}} \rho(Q) \text{Var}_{\tilde{\rho}_Q}(f).$$

Theorem 2.7 (Strong convexity for more general source). *Let ρ be a probability density over a domain $\mathcal{X} \subset \mathbb{R}^d$ satisfying the Boman chain condition for some covering $\mathcal{F}$ and $A, B, C > 1$. Assume moreover that (2.10) holds for some $D > 0$ (w.l.o.g. we assume $D > 1$) and there exists $E > 0$ such that*

$$\sup_{Q \in \mathcal{F}} \frac{M_{\rho_Q}}{m_{\rho_Q}} \leq E < +\infty. \quad (2.11)$$

Then for any $\mathcal{Y} \subset \mathbb{R}^d$ compact, there exists $C_{\rho, \mathcal{Y}} > 0$ such that for all $\psi_0, \psi_1 \in C_b(\mathcal{Y})$,

$$\text{Var}_\rho(\psi_1^* - \psi_0^*) \leq C \left(\mathcal{K}_\rho(\psi_1) - \mathcal{K}_\rho(\psi_0) + \langle \psi_1 - \psi_0 | (\nabla \psi_0^*)_{\#} \rho \rangle \right). \quad (2.12)$$

Proof. For each $Q \in \mathcal{F}$ set $\tilde{\rho}_Q = \frac{1}{\rho(Q)}\rho|_Q$, then Theorem 2.6 tells us

$$\text{Var}_{\tilde{\rho}_Q}(\psi_1^* - \psi_0^*) \leq C \left(\mathcal{K}_{\tilde{\rho}_Q}(\psi_1) - \mathcal{K}_{\tilde{\rho}_Q}(\psi_0) + \langle \psi_1 - \psi_0 | (\nabla \psi_0^*)_{\#} \tilde{\rho}_Q \rangle \right). \quad (2.13)$$

where $C_1 = e(d+1)2^{d+1}E^2 \text{Diam}(\mathcal{Y}) \text{Diam}(\mathcal{X})$. Combining Lemma 2.6 with (2.13), we get

$$\begin{aligned} \text{Var}_{\rho}(\psi_1^* - \psi_0^*) &\leq 200A^2CD^3 \sum_{Q \in \mathcal{F}} \rho(Q) \text{Var}_{\tilde{\rho}_Q}(\psi_1^* - \psi_0^*) \\ &\leq 200A^2CD^3C_1 \sum_{Q \in \mathcal{F}} \left(\mathcal{K}_{\rho|_Q}(\psi_1) - \mathcal{K}_{\rho|_Q}(\psi_0) + \langle \psi_1 - \psi_0 | (\nabla \psi_0^*)_{\#} \rho|_Q \rangle \right). \end{aligned}$$

Now define a partition $\mathcal{F}'$ of $\mathcal{X}$ such that x, x' belong to the same element $P \in \mathcal{F}'$ if and only if they belong to the same elements of $\mathcal{F}$. Proposition 2.1 applied to $\mathcal{K}_{\rho|_P}$ tells us that

$$\mathcal{K}_{\rho|_P}(\psi_1) - \mathcal{K}_{\rho|_P}(\psi_0) + \langle \psi_1 - \psi_0 | (\nabla \psi_0^*)_{\#} \rho|_P \rangle \geq 0 \quad \forall P \in \mathcal{F}'. \quad (2.14)$$

It follows that

$$\begin{aligned} \text{Var}_{\rho}(\psi_1^* - \psi_0^*) &\leq 200A^3CD^3C_1 \sum_{P \in \mathcal{F}'} (\mathcal{K}_{\rho|_P}(\psi_1) - \mathcal{K}_{\rho|_P}(\psi_0) + \langle \psi_1 - \psi_0 | (\nabla \psi_0^*)_{\#} \rho|_P \rangle) \\ &= 200A^3CD^3C_1 \left(\mathcal{K}_{\rho}(\psi_1) - \mathcal{K}_{\rho}(\psi_0) + \langle \psi_1 - \psi_0 | (\nabla \psi_0^*)_{\#} \rho \rangle \right). \end{aligned}$$

Where the first inequality is due to the chain condition of no more than A cubes simultaneously overlapping by (2.7) and the second equality is due to P being a partition of $\mathcal{X}$. $\square$

In particular, since the uniform measure on a John domain satisfies the Bowman chain condition (see [8, Lemma 2.1]), we have the following.

Theorem 2.8 (Strong convexity of the Kantorovich function on John domains). *Let $\mathcal{X} \subset \mathbb{R}^d$ be a John domain with non-empty interior. Let ρ be a probability density on $\mathcal{X}$ bounded above and below by positive constants. Then for every compact set $\mathcal{Y}$ there exists a constant $C_{\rho, \mathcal{Y}} > 0$ such that for all $\psi_0, \psi_1 \in C_b(\mathcal{Y})$,*

$$\text{Var}_{\rho}(\psi_1^* - \psi_0^*) \leq C \left(\mathcal{K}_{\rho}(\psi_1) - \mathcal{K}_{\rho}(\psi_0) + \langle \psi_1 - \psi_0 | (\nabla \psi_0^*)_{\#} \rho \rangle \right).$$

3 The discrete transport problem

We now shift to the main focus of this report, the fully discrete optimal transport problem. Such problems have a natural interpretation as a linear programming problem, which we will exploit heavily.

3.1 Primal and dual polyhedra

In the special case of fully discrete optimal transport, both dual and primal problems are linear programming problems. Specifically, fix weights $\alpha \in \mathbb{R}^M$, $\beta \in \mathbb{R}^N$ strictly positive probability vectors. Given $X = (x_1, \dots, x_M) \in (\mathbb{R}^d)^M$ and $Y = (y_1, \dots, y_N) \in (\mathbb{R}^d)^N$, we set

$$\rho_X = \sum_{i=1}^M \alpha_i \delta_{x_i} \quad \text{and} \quad \mu_Y = \sum_{j=1}^N \beta_j \delta_{y_j}. \quad (3.1)$$

We define a cost matrix $C(X, Y) \in \mathbb{R}^{M \times N}$ by $C(X, Y)_{ij} = c(x_i, y_j)$, and also the convex polyhedron of couplings

$$\Pi(\alpha, \beta) := \left\{ \hat{\gamma} \in \mathbb{R}^{M \times N} : \hat{\gamma}_{ij} \geq 0, \sum_{i=1}^M \hat{\gamma}_{ij} = \beta_j \text{ and } \sum_{j=1}^N \hat{\gamma}_{ij} = \alpha_i \text{ for } i = 1, \dots, M; j = 1, \dots, N \right\}.$$

Note here that we use $\hat{\gamma}$ to refer to couplings living in the polyhedron $\Pi(\alpha, \beta) \subset \mathbb{R}^{M \times N}$, while we will reserve γ for plans as probability measures on $\mathbb{R}^d \times \mathbb{R}^d$. The primal transport problem (1.1) between ρ_X and μ_Y is equivalent to

$$\inf_{\hat{\gamma} \in \Pi(\alpha, \beta)} \langle C(X, Y) | \hat{\gamma} \rangle. \quad (3.2)$$

Given $C \in \mathbb{R}^{M \times N}$, we define the dual polyhedron

$$D(C) := \{(\phi, \psi) \in \mathbb{R}^{M+N} : \phi_i + \psi_j \leq C_{ij} \text{ for all } i = 1, \dots, M, j = 1, \dots, N\}.$$

Then the dual transport problem (1.2) between ρ_X and μ_Y is equivalent to

$$\sup_{\phi, \psi \in D(C(X, Y))} \langle \phi | \alpha \rangle + \langle \psi | \beta \rangle. \quad (3.3)$$

Uniqueness of solutions in linear programming occurs when the cost vector to be maximised (here $-C$ for the primal or (α, β) for the dual) does not point towards a face of the polyhedron to which it is orthogonal. The unique solution is then one of the vertices (extreme points) of the polyhedron. This is illustrated in Figure 3.1. Thus, in the discrete case, uniqueness of the primal and dual problems is understood by characterising the extreme points of the polyhedra $\Pi(\alpha, \beta)$ and $D(C)$.

The following lemma, due to [20], characterises the extreme points of the primal transport polyhedron $\Pi(\alpha, \beta) \subset \mathbb{R}^{M \times N}$. This serves as a generalisation of the Birkhoff theorem for bi-stochastic matrices, which corresponds to the case $N = M$ and α, β uniform probability vectors. The second half of the proof presented here can be found in [10, Chapter 8], the first half is formulated with the results of Chapter 4 in mind.




Figure 3.1: Uniqueness vs non-uniqueness of the minimisers of $\langle \hat{\gamma} | C \rangle$ depending on the direction of $-C$.

Lemma 3.1. [20] *Given $\hat{\gamma} \in \Pi(\alpha, \beta)$, define the bipartite graph $G(\hat{\gamma}) = (V, E(\hat{\gamma}))$ of its support, with $V = (x_i)_{i=1}^M \cup (y_j)_{j=1}^N$ and*

$$(x_i, y_j) \in E(\hat{\gamma}) \iff \hat{\gamma}_{ij} > 0. \quad (3.4)$$

Then $\hat{\gamma} \in \Pi(\alpha, \beta)$ is an extreme point if and only if $G(\hat{\gamma})$ contains no cycles (it is a disjoint union of trees, often referred to as a forest).

Proof. Assume $\hat{\gamma}$ is such that $G(\hat{\gamma})$ has a cycle. Since no two x 's or y 's are connected, this is of the form $x_{i_1}, y_{j_1}, \dots, x_{i_n}, y_{j_n}, x_{i_1}$, such that $\hat{\gamma}_{i_k j_k} > 0$ and $\hat{\gamma}_{i_1 j_n} > 0$ and $\hat{\gamma}_{i_{k+1} j_k} > 0$. Here we choose $n \in \mathbb{N}$ to minimise the length of the chain

$$i_1, j_1, \dots, i_n, j_n, i_1, \quad (3.5)$$

so the cycle is simple (it does not retrace itself). Set

$$\varepsilon := \min \left(\hat{\gamma}_{i_1 j_n}, \min_{k=1}^n \hat{\gamma}_{i_k j_k}, \min_{k=1}^{n-1} \hat{\gamma}_{i_{k+1} j_k} \right).$$

Then $\hat{\gamma} = \frac{1}{2}\hat{\gamma}_\varepsilon + \frac{1}{2}\hat{\gamma}_{-\varepsilon}$ where

$$(\hat{\gamma}_\varepsilon)_{ij} := \begin{cases} \hat{\gamma}_{ij} + \varepsilon & i, j \text{ appears in the chain (3.5),} \\ \hat{\gamma}_{ij} - \varepsilon & j, i \text{ appears in the chain (3.5),} \\ \hat{\gamma}_{ij} & \text{otherwise.} \end{cases}$$

By construction, $\hat{\gamma}_\varepsilon, \hat{\gamma}_{-\varepsilon} \in \Pi(\alpha, \beta)$. Hence $\hat{\gamma}$ is not an extreme point, since $\hat{\gamma}_\varepsilon$ and $\hat{\gamma}_{-\varepsilon}$ are distinct.

Now assume the induced bipartite graph by $\hat{\gamma} \in \Pi(\alpha, \beta)$ is a forest. We show that there exists no other $\hat{\gamma}_0 \in \Pi(\alpha, \beta)$ whose induced graph is a subgraph of that of $\hat{\gamma}$ (recall that we are considering our graph as an unweighted graph, only the pattern is important here). Since α and β are strictly positive probability vectors, each vertex has degree at least 1. Since $G(\hat{\gamma})$ is a forest, there exists a vertex of degree 1 on the graph, without loss of generality, let it be y_1 , and let (x_1, y_1) be the single edge connecting to y_1 . So that the right amount of mass arrives at y_1 , it must be that $\hat{\gamma}_{11} = \beta_1$. Repeatedly applying this observation to degree 1 vertices on the tree as we remove them, we see that a tree graph uniquely specifies the element of $\Pi(\alpha, \beta)$.

Thus, there can be no other $\hat{\gamma}_0 \in \Pi(\alpha, \beta)$ whose graph is a subgraph of $G(\hat{\gamma})$. Consequently, if $\hat{\gamma} = \frac{1}{2}\hat{\gamma}^- + \frac{1}{2}\hat{\gamma}^+$, then $\hat{\gamma} = \hat{\gamma}^- = \hat{\gamma}^+$ since $G(\hat{\gamma}^-), G(\hat{\gamma}^+)$ are subgraphs of $G(\hat{\gamma})$. Hence $\hat{\gamma}$ is an extreme point of $\Pi(\alpha, \beta)$. $\square$

In [5, 4, 3], extreme points for the dual transport polyhedron $D(C) \cap \{\phi_1 = 0\}$ are also characterised in terms of a bipartite graph with node sets corresponding to the support points X and Y . (Here we fix the first coordinate arbitrarily, as otherwise there are no extreme points, with the lowest-dimensional faces being lines.)

Lemma 3.2. [4, Lemma 1] *Given $(\phi, \psi) \in D(C) \cap \{\phi_1 = 0\}$, define the bipartite graph $G(\phi, \psi) = (V, E(\phi, \psi))$ with $V = (x_i)_{i=1}^M \cup (y_j)_{j=1}^N$ and*

$$(x_i, y_j) \in E(\phi, \psi) \iff \phi_i + \psi_j = C_{ij}. \quad (3.6)$$

Then (ϕ, ψ) is an extreme point of $D(C) \cap \{\phi_1 = 0\}$ if and only if $G(\phi, \psi)$ is connected.

Proof. Assume first that (ϕ, ψ) are such that $G(\phi, \psi)$ is connected. Then there exists a path from each vertex starting from x_1 (which represents $\phi_1 = 0$). By definition of $G(\phi, \psi)$, iterating along any path in the graph $1 = i_1, j_1, i_2, j_2, \dots, j_n$ in $G(\phi, \psi)$ we must have that $\phi_{i_{k+1}} = C_{i_k j_k} - \psi_{j_k}$, and $\psi_{j_k} = C_{i_k j_k} - \phi_{i_k}$. Hence, for a connected graph, the values of (ϕ, ψ) are uniquely determined. Let $(\phi^-, \psi^-), (\phi^+, \psi^+)$ be such that $(\phi, \psi) = \frac{1}{2}(\phi^-, \psi^-) + \frac{1}{2}(\phi^+, \psi^+)$. By definition of $D(C)$, for each i, j

$$\phi_i^- + \psi_j^- \leq C_{ij} \quad \text{and} \quad \phi_i^+ + \psi_j^+ \leq C_{ij}.$$

Hence $(x_i, y_j) \in E(\phi, \psi)$ if and only if it is in both $E(\phi^-, \psi^-)$ and $E(\phi^+, \psi^+)$; in other words $G(\phi, \psi) = G(\phi^-, \psi^-) \cap G(\phi^+, \psi^+)$. Consequently, if $G(\phi, \psi)$ is connected, both $G(\phi^-, \psi^-)$ and $G(\phi^+, \psi^+)$ are, and the same arcs are sufficient to connect them. It follows from the unique determinacy property of connected graphs shown above that $(\phi, \psi) = (\phi^-, \psi^-) = (\phi^+, \psi^+)$, and hence (ϕ, ψ) is an extreme point of $D(C) \cap \{\phi_1 = 0\}$.

Assume now that (ϕ, ψ) are such that $G(\phi, \psi)$ are disconnected. The above argument establishes that on each connected component of $G(\phi, \psi)$, the corresponding coordinates of ϕ and ψ are fixed up to a constant. We will show that, pairwise, between different connected components of $G(\phi, \psi)$, there is some flexibility (we can add a small constant to all the coordinates in one component, while keeping the other components unchanged). This maximum amount of flexibility will define a sort of distance between two components. Since the graph is finite, choosing some component D not containing ϕ_1 , and taking the minimum distance between all other components, we will be able to perturb all the values on this component slightly in each direction whilst respecting the constraints of $D(C) \cap \{\phi_1 = 0\}$. This will allow us to write (ϕ, ψ) as a non-trivial convex combination, proving that they do not constitute an extreme point.

Forgetting briefly the constraint $\{\phi_1 = 0\}$, decompose $G(\phi, \psi) = \sqcup_{k=1}^r G_k$, where $G_k = (V_k, E_k)$ are the connected components of $G(\phi, \psi)$. Consider two components G_0, G_1 . We consider all the edges missing from $G(\phi, \psi)$ that could join G_0 to G_1 , and see how much slack they have in their constraint

$$C_{ij} - \phi_i - \psi_j > 0.$$

First, we consider all the vertices starting from the x side of G_0 to the y side of G_1 . For $k_0 \neq k_1$, set

$$\varepsilon_{\phi, \psi}(G_k, G_l) := \min_{\substack{x_i \in V_k \\ y_j \in V_l}} C_{ij} - \phi_i - \psi_j > 0, \quad (3.7)$$

(Note this is not at all symmetric in G_{k_0}, G_{k_1}) where the strict inequality comes from the fact that we are taking two vertices from different components, so none of these edges can be in the

graph. Then for each coordinate (i_0, j_1) in (3.7), adding $\varepsilon_{\phi, \psi}(G_0, G_1)$ to each ψ_{j_1} coordinate and subtracting it from each ϕ_{i_0} coordinate preserves all constraints within each component, as well as (G_0, x) to (G_1, y) edge constraints. Symmetrically, for each coordinate (i_1, j_0) in, subtracting $\varepsilon_{\phi, \psi}(G_1, G_0)$ from each ψ_{j_1} coordinate and adding it to each ϕ_{i_0} coordinate preserves all constraints within each component, as well as (G_0, y) to (G_1, x) edge constraints. Returning now to our setting $(\phi, \psi) \in D(C) \cap \{\phi_1 = 0\}$ with disconnected graph choose any connected component G_l not containing ϕ_1 , then set

$$\varepsilon_{\phi, \psi}^+ = \min_{\substack{k=1, \dots, r \\ k \neq l}} \varepsilon_{\phi, \psi}(G_k, G_l) \quad \text{and} \quad \varepsilon_{\phi, \psi}^- = \min_{\substack{k=1, \dots, r \\ k \neq l}} \varepsilon_{\phi, \psi}(G_k, G_l).$$

The above computations show that, keeping all the values of (ϕ, ψ) fixed on every component apart from G_l , we can add ε^+ to each ψ and subtract from each ϕ value while staying inside $D(C) \cap \{\phi_1 = 0\}$. Similarly, we can subtract ε^- from each ψ and add it to each ϕ . Hence, defining (ϕ^+, ψ^+) as the potential obtained by adding $\varepsilon := \min(\varepsilon^+, \varepsilon^-)$ to each ψ coordinate in G_l and subtracting from each ϕ coordinate in G_l , and defining (ϕ^-, ψ^-) in the opposite manner with respect to addition and subtraction by ε , we have

$$(\phi, \psi) = \frac{1}{2}(\phi^+, \psi^+) + \frac{1}{2}(\phi^-, \psi^-).$$

Hence, we have written (ϕ, ψ) as a non-trivial linear combination of two elements of $D(C) \cap \{\phi_1 = 0\}$, and so it is not an extreme point. $\square$

Remark 3.3. The second half of the above proof involves quantitative calculations that are essentially identical to those of [1, Theorem 3.9]. We will return to this connection in later sections.

In [1], the authors also study the uniqueness of solutions to the fully discrete dual transport problem. They also characterise uniqueness in terms of similar bipartite graphs as in Lemma 3.2. Here, they arrive at what is (usually) the same graph, but instead, it is constructed implicitly using the set of primal optimisers, combined with the compatibility condition (1.8).

Proposition 3.4. [1, Proposition 3.5 (ii)/Corollary 3.14] *For $\hat{\gamma} \in \Pi(\alpha, \beta)$, define the bipartite graph as in Lemma 3.1. Define the superposition of optimal graphs $G_\Gamma = (V, E_\Gamma)$ where*

$$E_\Gamma := \bigcup_{\substack{\hat{\gamma} \in \Pi(\alpha, \beta) \\ \hat{\gamma} \text{ optimal for (3.3)}}} E(\hat{\gamma}).$$

Then the solutions of (3.3) are unique up to a constant if and only if G_Γ is connected.

Condition (1.8) forces that for any dual optimisers (ϕ, ψ) , $G_\Gamma \subset G(\phi, \psi)$. [1, Proposition 3.6] establishes the inverse direction $G_\Gamma = G(\phi, \psi)$ when dual potentials are unique up to a constant. Effectively, since in this case $G(\phi, \psi)$ is connected, if there were some edge in $G(\phi, \psi) \setminus G_\Gamma$, it must belong to a cycle of $G(\phi, \psi)$. The reason for this is that the tree-structured sections of $G(\phi, \psi)$ force all optimal plans supported there to have the same weights, so all edges of $G(\phi, \psi)$ not contained in a cycle must also lie in G_Γ . But if there is some edge lying in a cycle of $G(\phi, \psi)$ we can always take another optimal plan and permute some of the mass round the cycle as in the proof of Lemma 3.1, to create an optimal plan giving mass to this edge, and hence the edge is in G_Γ . Thus, in this case $G_\Gamma = G(\phi, \psi)$. In general, G_Γ can be a strict subset of $G(\phi, \psi)$. To construct such a case, take any primal and dual problem for which $G(\phi, \psi)$ has at least two connected components, then add one edge to the graph joining the two components, by reducing a C_{ij} connecting them until it is equal to $\phi_i + \psi_j$. Then no new optimal plans are created so G_Γ remains unchanged, but $G(\phi, \psi)$ just gained an edge.

We note that in [4], the authors distinguish between two types of vertices of dual transport polyhedra, referred to as *non-degenerate* when $G(\phi, \psi)$ is a spanning tree, and *degenerate* when $G(\phi, \psi)$ contains a cycle. The degenerate case occurs precisely when we have non-uniqueness of the primal problem; this fact can be seen as a consequence of Lemma 3.1.

To synthesise the discrete case: uniqueness of the primal problem occurs if and only if the corresponding graph G_Γ has no cycles, and uniqueness of the dual if and only if G_Γ is connected.

3.2 The sets of optimal plans and potentials

We now view the discrete transport problem again through continuous rather than linear programming language. As discussed, in general, we do not have uniqueness, so we are interested in the solution set of minimisers. We define

$$\begin{aligned} \Gamma_c : (\mathbb{R}^d)^M \times (\mathbb{R}^d)^N &\rightrightarrows \mathcal{P}(\mathbb{R}^{2d}); \\ \Gamma_c(X, Y) &:= \arg \min_{\gamma \in \Pi(\rho_X, \mu_Y)} \int_{\mathbb{R}^d \times \mathbb{R}^d} c(x, y) d\gamma(x, y), \end{aligned} \quad (3.8)$$

where ρ_X and μ_Y are defined as in (3.1). We denote

$$\text{Supp } \Gamma(X, Y) = \bigcup_{\gamma \in \Gamma(X, Y)} \text{Supp } \gamma,$$

sometimes viewing $\text{Supp } \Gamma$ as a graph and sometimes as a set, as it should be clear from the context to which we refer. We know from the theory of the previous section that the optimal set should correspond to a (potentially zero-dimensional) face of the polyhedron $\Pi(\alpha, \beta)$. Consider sets of points of the form

$$P = \{(x_i, y_i)\}_{i=1}^n \subset \text{Supp } \Gamma;$$

which are *distinct* in that

$$x_i \neq x_j \quad \text{and} \quad y_i \neq y_j \quad \text{for all } i \neq j. \quad (3.9)$$

Under these hypotheses,

$$\{(x_{i+1}, y_i)\}_{i=1}^n \cap P = \emptyset, \quad (3.10)$$

where $x_{n+1} := x_1$. Given such a P , we define the cyclical monotonicity gap

$$\Delta_P = \sum_{i=1}^n c(x_{i+1}, y_i) - c(x_i, y_i).$$

We can then say the following.

Proposition 3.5. *Let $X \in (\mathbb{R}^d)^M$ and $Y \in (\mathbb{R}^d)^N$ and consider the transport problem (3.8) for some $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$. Then*

- (i) $\gamma \in \Gamma(X, Y)$ if and only if $\Delta_P \geq 0$ for all distinct (in the sense of (3.9)) P in $\text{Supp } \gamma$.
- (ii) Problem (3.8) has uniqueness of solutions if and only if there exists $\gamma \in \Pi(\rho_X, \mu_Y)$, such that $\Delta_P > 0$ for all distinct (in the sense of (3.9)) $P \subset \text{Supp } \gamma$ of length ≥ 2 (and this γ is the unique solution).

Proof. Claim (i) is effectively a rephrasing of the classical cyclical monotonicity in the discrete case, as stated in the introduction (1.10). We know optimality is equivalent to testing $\Delta_P \geq 0$ for all $P = \{(x_i, y_i)\}_{i=1}^n \subset \text{Supp } \gamma$ without distinctness assumptions. We write the cycle

$$x_1, y_1, x_2, y_2, \dots, x_n, y_n, x_1,$$

and decompose into cycles $\{P_k\}_{k=1}^m$ of maximal length in which no x or y point is equal to another in the cycle. Then each $P_k = \{(x_i, y_i)\}_{i=1}^{n_k}$ is distinct, and $\Delta_P = \sum_{k=1}^m \Delta_{P_k}$, so that it is sufficient to test only distinct P for optimality as required.

Claim (ii) can be seen as follows. Let $\gamma \in \Gamma(X, Y)$, assume that there exists $P = (x_i, y_i)_{i=1}^n \subset \text{Supp } \gamma$ distinct in the sense of (3.9), with $\Delta_P = 0$. Define $\tilde{\gamma}$ by

$$\tilde{\gamma} := \gamma + \varepsilon \sum_{i=1}^n (\delta_{x_{i+1}, y_i} - \delta_{x_i, y_i}),$$

where $\varepsilon = \min_{i=1, \dots, n} \gamma((x_i, y_i)) > 0$. We have $\tilde{\gamma} \in \Pi(\rho_X, \mu_Y)$ by construction. By (3.10), γ and $\tilde{\gamma}$ are distinct transport plans, and by $\Delta_P = 0$, $\int c d\gamma = \int c d\tilde{\gamma}$, so $\tilde{\gamma}$ is also an optimal transport plan for c cost, and the problem does not have a unique solution.

Conversely, assume that we have uniqueness, take $\gamma_0, \gamma_1 \in \Pi(\rho_X, \mu_Y)$ as distinct optimal plans. Then the measure $\gamma_0 - \gamma_1$ is a non-zero signed zero mass discrete probability measure on $\mathbb{R}^d \times \mathbb{R}^d$ with zero first and second marginals. Consider the bipartite graph of $\text{Supp}(\gamma_0 - \gamma_1)$ defined as (V, E) where

$$V = p_X(\text{Supp}(\gamma_0 - \gamma_1)) \cup p_Y(\text{Supp}(\gamma_0 - \gamma_1)),$$

and $E = E^+ \sqcup E^-$ where

$$E^+ = \{(x, y) : (\gamma_0 - \gamma_1)(x, y) > 0\} \quad \text{and} \quad E^- = \{(x, y) : (\gamma_0 - \gamma_1)(x, y) < 0\}.$$

We have $E^+ \subset \text{Supp } \gamma_0$ and $E^- \subset \text{Supp } \gamma_1$. Every vertex has at least one positive and one negative edge leaving from it (since $\gamma_0 - \gamma_1$ has zero marginals). And the graph is simple, with at most one edge joining any two points, either positive or negative. We will construct a simple cycle in the graph. Start from some x_{i_1} , there exists a positive edge to some y_{j_1} . Then there exists some negative edge to some x_{i_2} in the graph, necessarily $i_2 \neq i_1$ since the graph is simple and the signs of the edges are different. We repeat this process, always choosing to visit a new vertex, until eventually we are forced to revisit a vertex for the first time. Let y_{j_k} the last distinct point before we subsequently return to vertex x_{i_m} for some $m < k$. (We cannot have $i_m = i_k$ since the signs of their edges in the graph must be different.) Up to removing the points before i_m and relabelling, it follows that there exists a nontrivial simple cycle of the form

$$\underbrace{x_1, y_1}_{\in E^+}, \underbrace{x_2, y_2}_{\in E^+} \cdots \underbrace{x_n, y_n}_{\in E^+}, \underbrace{x_1}_{\in E^-}. \quad (3.11)$$

Take $P = (x_i, y_i)_{i=1}^n$, then since the cycle is simple, P is distinct in the sense of (3.9). Moreover, $P \subset \text{Supp } \gamma_0$ since these edges are all in E^+ . For any optimal potential pair $\phi, \psi \in \mathbb{R}^{M \times N}$ for (3.3), the compatibility condition with both γ_0 and γ_1 implies $\phi_i + \psi_i = c(x_i, y_i)$ and $\phi_{i+1} + \psi_i = c(x_{i+1}, y_i)$, and consequently

$$\Delta_P = \sum_{i=1}^n c(x_{i+1}, y_i) - c(x_i, y_i) = \sum_{i=1}^n (\phi_{i+1} + \psi_i - \phi_i - \psi_i) = 0.$$

Hence, we have shown that non-uniqueness implies that for any optimal γ_0 , there exists a distinct $P \subset \text{Supp } \gamma_0$ with $\Delta_P = 0$. The characterisation ((ii)) follows. $\square$

We define similar optimal sets for the Kantorovich potentials. We will use the semi-dual formulation (1.7), which takes a special formulation as soon as one of the measures is discrete. In general, we are only interested in the values of the potentials on the support points X and Y , and so we can reduce to a finite-dimensional problem. We define

$$\Phi_c : (\mathbb{R}^d)^M \times (\mathbb{R}^d)^N \rightrightarrows \mathbb{R}^M, \quad \Psi_c : (\mathbb{R}^d)^M \times (\mathbb{R}^d)^N \rightrightarrows \mathbb{R}^N;$$

$$\Phi_c(X, Y) := \arg \max_{\phi \in \mathbb{R}^M} \langle \phi | \alpha \rangle_{\mathbb{R}^M} + \int_{\mathcal{X}} \phi^c(y|X) d\mu_Y(y), \quad (3.12)$$

$$\Psi_c(X, Y) := \arg \max_{\psi \in \mathbb{R}^N} \int_{\mathcal{X}} \psi^c(x|Y) d\rho_X(x) + \langle \psi | \beta \rangle_{\mathbb{R}^N}, \quad (3.13)$$

where the c -transforms are taken with respect to points in the support of the target X and Y :

$$\psi^c(x|Y) := \inf_{j=1, \dots, N} c(x, y_j) - \psi_j, \quad \text{and} \quad \phi^c(y|X) := \inf_{i=1, \dots, M} c(x_i, y) - \phi_i.$$

Problems (3.12) and (3.13) are referred to as the *semi-discrete* transport problem, which has been the subject of numerous recent works, see [19, 23, 14, 25]. Due to symmetry, we can define both Φ and Ψ as above. It is not clear which is more useful, or if both are necessary (recall sometimes we want to fix X and perturb Y so our perturbations are not necessarily symmetric). To avoid developing a dual theory unnecessarily, we will concentrate on problem (3.12). Symmetric calculations apply to (3.13).

On one side of (3.12), we have a vector in $\mathbb{R}^M$, corresponding to the dual in the linear programming formulation, and on the other side we have $\phi^c \in C_b(\mathcal{Y})$, corresponding to the measure theoretic formulation. In this way, (3.12) will serve as a bridge between these two formulations, allowing us to pass stability results from linear programming to functional analysis language. Then, qualitative stability, Theorem 1.1, should allow us to take a limit of discrete measures to say statements about more general optimal transport problems beyond the discrete case. For ease of exposition, we fix the cost to be quadratic, $c(x, y) = -\langle x|y \rangle$, and state the generalisations to general c after. Here we simply denote Γ and Φ for the sets of optimisers. In this case, (3.12) becomes

$$\inf_{\phi \in \mathbb{R}^M} \langle \phi | \alpha \rangle_{\mathbb{R}^M} + \int_{\mathbb{R}^d} \phi^*(y) d\mu(y). \quad (3.14)$$

with Legendre transform $\psi^* : \mathbb{R}^d \rightarrow \mathbb{R}$ defined by

$$\phi^*(y|X) := \sup_{i=1, \dots, M} \langle y|x_i \rangle - \phi_i. \quad (3.15)$$

To keep notation light, where clear, we omit the dependence on X in the Legendre transform. The function $\phi^* : \mathbb{R}^d \rightarrow \mathbb{R}$ is a finite supremum of hyperplanes, with one hyperplane $y \mapsto \langle y|x_i \rangle - \phi_i$ of gradient of x_i corresponding to each point in the support of ρ_X . Thus ϕ^* is piecewise flat, and the region for which the hyperplane i is active in the sup (3.15) is called the i th Laguerre cell of $\phi \in \mathbb{R}^M$, denoted

$$\text{Lag}_i(\phi|X) = \left\{ y \in \mathbb{R}^d : \langle y|x_i \rangle - \phi_i \geq \langle y|x_j \rangle - \phi_j \text{ for all } j = 1, \dots, M \right\} = \left\{ y \in \mathbb{R}^d : x_i \in \partial \phi^*(y) \right\}.$$

These cells have pairwise disjoint interiors, providing a decomposition of $\mathbb{R}^d$. This definition *does not* depend on μ , only on the points $X \in (\mathbb{R}^d)^M$. As (3.14) is a convex minimisation problem, for each $X, Y \in \mathbb{R}^N$, $\Phi(X, Y)$ is a convex set. We will give an explicit characterisation of Φ as a finite intersection of certain half-spaces. This will allow us to study how Φ changes as we perturb X and Y .

Lemma 3.6. For any vector of target positions $Y \in (\mathbb{R}^d)^N$, set

$$S_i(X, Y) := \{y \in \{y_j\}_{j=1}^N : \text{there exists } \gamma \in \Gamma(X, Y) \text{ such that } (x_i, y) \in \text{Supp } \gamma\}. \quad (3.16)$$

Then

$$\Phi(X, Y) = \{\phi \in \mathbb{R}^N : S_i(X, Y) \subset \text{Lag}_i(\phi|X) \text{ for all } i = 1, \dots, M\}.$$

Proof. By definition $y \in \text{Lag}_i(\phi|X)$ if and only if $x_i \in \partial\phi^*(y)$. Hence $S_i(X, Y) \subset \text{Lag}_i(\phi|X)$ for all $i = 1, \dots, M$ if and only if for any $\gamma \in \Pi(\rho_X, \mu_Y)$ optimal and any $(x_i, y) \in \text{Supp } \gamma$, we have $x_i \in \partial\phi^*(y)$. But this is precisely saying that $\text{Supp } \gamma \subset \partial\phi^*$. By the compatibility condition (1.8), this is equivalent to $\phi \in \Phi(X, Y)$. $\square$

Proposition 3.7 (Characterisation of the set of dual potentials). *Let ρ_X and μ_Y be discrete probability measures on $\mathbb{R}^d$ for some $X \in (\mathbb{R}^d)^M$ and $Y \in (\mathbb{R}^d)^N$. Then $\Phi(X, Y)$ is a convex polyhedron in $\mathbb{R}^M$ characterised by a finite intersection of half spaces, one for each pairing of distinct points $x, x' \in \{x_i\}_{i=1}^M$. Explicitly,*

$$\Phi(X, Y) = \bigcap_{i,j} H_{ij}(X, Y), \quad (3.17)$$

where

$$H_{ij}(X, Y) := \{\phi \in \mathbb{R}^M : \langle \phi | e_i - e_j \rangle \leq h_{ij}(X, Y)\}, \quad (3.18)$$

e_i are basis elements of $\mathbb{R}^M$, and

$$h_{ij} : (\mathbb{R}^d)^M \times (\mathbb{R}^d)^N \rightarrow \mathbb{R}; \quad X, Y \mapsto \inf_{y \in S_i(X, Y)} \langle x_i - x_j | y \rangle.$$

Proof. We will show that

$$\phi \in \bigcap_{i,j} H_{ij} \iff S_i \subset \text{Lag}_i(\phi) \text{ for all } i,$$

which implies the result by Lemma 3.6. We have $S_i \subset \text{Lag}_i(\phi)$ if and only if for all $y \in S_i$ and all $j = 1, \dots, M$,

$$\langle x_i | y \rangle - \phi_i \geq \langle x_j | y \rangle - \phi_j. \quad (3.19)$$

This is equivalent to asking that for all $j = 1, \dots, M$,

$$\langle \phi | e_i - e_j \rangle_{\mathbb{R}^M} \leq \inf_{y \in S_i} \langle x_i - x_j | y \rangle_{\mathbb{R}^d} = h_{ij}.$$

Hence $\phi \in H_{ij}$ for all $j \neq i$, and so

$$\phi \in \bigcap_{j=1}^M H_{ij} \iff S_i \subset \text{Lag}_i(\phi).$$

Finally, we want $S_i \subset \text{Lag}_i(\phi)$ for any $i = 1, \dots, M$. Hence, we take the intersection over i , giving the characterisation (3.17). $\square$

Remark 3.8. We have two parallel half-spaces with gradient $e_i - e_j$ in $\mathbb{R}^M$. Their intersection $\{\phi : \phi_i - \phi_j \in [-h_{ji}, h_{ij}]\}$ is all the admissible intervals of the relative vertical displacements between the hyperplane associated to x_i and that associated to x_j , which preserve $S_i \subset \text{Lag}_i(\phi)$ and $S_j \subset \text{Lag}_j(\phi)$. Each of the hyperplanes defining H_{ij} is orthogonal to the vector $(1, 1, \dots, 1)$, and so the optimal set is symmetrical in this direction. This is consistent with the fact that potentials are only ever defined up to a constant. What matters are the relative heights of




Figure 3.2: Laguerre cell decompositions of the source (left) and target (middle) domains induced by an optimal dual vector, and corresponding graph G_Γ (right).

the hyperplanes defining ϕ^* , as these decide the cell boundaries. If two hyperplanes are both translated vertically by the same amount, their intersection set remains the same, so if their cells are adjacent, the boundary is preserved.

The compatibility condition (1.8) forces that at $y \in \text{Supp } \mu_Y$ the subgradient of any optimal $\phi^* : \mathcal{Y} \rightarrow \mathbb{R}$ must contain all x such that $(x, y) \in \text{Supp } \gamma$ for any $\gamma \in \Gamma(X, Y)$. We then have the liberty to choose any convex function whose subdifferential contains these values (it can always also contain more values too; (1.8) demands containment, not equality). If μ_Y is a very fine mesh, then the subgradients of ϕ^* are specified on a very dense set, so there is less flexibility in choice. If ρ_X is a very fine mesh compared to μ_Y , then for $\phi^{**} : \mathcal{X} \rightarrow \mathbb{R}$, many y will lie in the subdifferentials of multiple x_i . This also leaves little choice, as we are forced to interpolate affinely between any such pair of points so that the value can be in the subgradient of both. Many points $x_i \in \text{Supp } \rho$ will send their mass split to one y , and so for any potential ϕ^* , these y should lie at the intersection of many cells rather than just inside a single cell.

Remark 3.9 (Extension to cost c). The extension of this to cost c simply uses generalised Laguerre cells

$$\text{Lag}_c(\phi) = \left\{ y \in \mathbb{R}^d : c(x_i, y) - \phi_i \leq c(x_j, y) - \phi_j \quad \forall j = 1, \dots, M \right\}.$$

The compatibility condition demands $S_i^c \subset \text{Lag}_i^c(\phi)$ where S_i^c is defined similarly to the quadratic case, using Γ_c instead. The hyperplane constraints are given explicitly using h_{ij}^c defined by

$$h_{ij}^c : (\mathbb{R}^d)^M \times (\mathbb{R}^d)^N \rightarrow \mathbb{R}; \quad X, Y \mapsto \sup_{y \in S_i^c(X, Y)} c(x_i, y) - c(x_j, y).$$

$\phi \in \mathbb{R}^M$ is optimal if and only if for all $i = 1, \dots, M$ and $j = 1, \dots, N$,

$$\langle \phi | e_i - e_j \rangle \geq h_{ij}^c(X, Y).$$

Example 3.10. In $\mathbb{R}^2$, let $\rho = \sum_{i=1}^4 \frac{1}{4} \delta_{x_i}$ and $\mu = \sum_{j=1}^3 \frac{1}{3} \delta_{y_j}$, where $\{x_i\}_{i=1}^4$ and $\{y_j\}_{j=1}^3$ are as in Figure 3.2. The points on the left picture of Figure 3.2 correspond to the cells (or hyperplanes) in the centre picture, and vice versa. Non-uniqueness of the primal problem occurs precisely when there is a non-trivial cycle of the form point 1, hyperplane 1, point 2, hyperplane 2, $\dots$ hyperplane n , point 1, where each point lies on the boundary of the hyperplanes it is next to in the cycle. Here, there exist three such cycles in G_Γ , given by

$$(x_1, y_2, x_2, y_1, x_1), (x_1, y_2, x_3, y_3, x_1), \text{ and } (x_1, y_3, x_3, y_2, x_1).$$

4 Perturbation of support positions

We now investigate how solutions to discrete optimal transport problems change when the support points X and Y are perturbed. This leads naturally to questions of stability for both primal and dual formulations.

4.1 Discussion on how to use linear programming connections

The following could lead to a better understanding of Problem 1:

Problem 2. How does $\Gamma(X, Y)$ change as we perturb X and Y ? To what extent can we say that if X, X' and Y, Y' are close (and hence $W_2(\rho_X, \rho_{X'})$ and $W_2(\mu_Y, \mu_{Y'})$ are small), then $\Gamma_c(X, Y)$ and $\Gamma_c(X', Y')$ are quantitatively close with respect to an appropriate metric? We would like results of the form

$$W_2(\gamma_{X_0, Y_0}, \gamma_{X_1, Y_1}) \leq C \left(W_2(\rho_{X_0}, \rho_{X_1}) + W_2(\mu_{Y_0}, \mu_{Y_1}) \right)^q + \mathcal{E}(X, Y, \alpha, \beta) \quad (4.1)$$

for all $\gamma_{X_i, Y_i} \in \Gamma(X_i, Y_i)$, where $\mathcal{E}(X, Y, \alpha, \beta)$ is an error term to be understood depending on the structure of the measures.

If we fix X and just vary Y with all points staying inside the same compact set, can we arrive at an error $\mathcal{E}(X, \alpha)$ not depending on Y , which is, for example, very small when ρ_X represents a very fine mesh discrete approximation of a regular absolutely continuous measure? Note here that in the spirit of known results in the absolutely continuous case, we should consider some constraint that all the coordinates of Y live inside the same compact set $\mathcal{Y}$, and the error should be allowed to depend on $\mathcal{Y}$ also.

In general, quantitative stability cannot hold without some form of error term $\mathcal{E}$. Sometimes this error can be large, as the following example demonstrates.

Example 4.1. Let $\rho_\varepsilon = \frac{1}{2}\delta_{(-1, \varepsilon)} + \frac{1}{2}\delta_{(1, -\varepsilon)}$ and $\mu = \frac{1}{2}\delta_{(0, 1)} + \frac{1}{2}\delta_{(0, -1)}$, see Figure 4.1. For $\varepsilon \neq 0$, there is a unique optimal plan γ_ε , but at $\varepsilon = 0$ we have non-uniqueness, with any convex combination of the two limiting cases giving an optimal plan. For $\varepsilon > 0$, $W_2(\gamma_\varepsilon, \gamma_{-\varepsilon}) = 2$




Figure 4.1: Optimal plans between ρ_ε and μ .

while $W_2(\rho_\varepsilon, \rho_{-\varepsilon}) = 2\varepsilon$. The size of the jump discontinuity is $\lim_{\varepsilon \rightarrow 0} W_2(\gamma_\varepsilon, \gamma_{-\varepsilon}) = 2$, which is precisely the Wasserstein diameter of the set of optimal plans between ρ_0 and μ . This is the worst it can be, as a consequence of qualitative stability.

Problem 2 is related to the stability of the linear programs (3.2) and (3.3), with respect to perturbations of $C(X, Y)$. However, the geometry of linear programming; i.e. $\Pi(\alpha, \beta) \subset \mathbb{R}^{M \times N}$ equipped with some norm, and the geometry of $(\mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d), W_2)$ are very different. In particular, in (3.2), all the geometric information X and Y concerning the ambient space $\mathbb{R}^d$ is embedded directly into the cost vector $C(X, Y)$, and bears no relation to the geometry of the polyhedron $\Pi(\alpha, \beta)$. Consider the mapping

$$\gamma : (\mathbb{R}^d)^M \times (\mathbb{R}^d)^N \times \Pi(\alpha, \beta) \rightarrow \mathcal{P}_2(\mathbb{R}^d \times \mathbb{R}^d); \quad (X, Y, \hat{\gamma}) \mapsto \gamma(X, Y),$$

where we denote

$$\gamma(X, Y) := \sum_{i=1}^M \sum_{j=1}^N \hat{\gamma}_{ij} \delta_{x_i, y_j} \in \Pi(\rho_X, \mu_Y).$$

First, assume we are comparing two different transport plans in $\Pi(\rho_X, \mu_Y)$. Two distinct couplings $\hat{\gamma}_0, \hat{\gamma}_1 \in \Pi(\alpha, \beta)$ have a fixed Euclidean distance $\|\hat{\gamma}_0 - \hat{\gamma}_1\|_2 > 0$ in the polyhedron, whilst $W_2(\hat{\gamma}_0(X, Y), \hat{\gamma}_1(X, Y))$ can be made arbitrarily large or small by varying X and Y . We note that the l^1 norm on $\Pi(\alpha, \beta) \subset \mathbb{R}^{M \times N}$ corresponds to the total variation:

$$\|\hat{\gamma}_0 - \hat{\gamma}_1\|_1 = \|\gamma_0(X, Y) - \gamma_1(X, Y)\|_{TV}$$

In general, however, we want to compare two plans, one in $\Pi(\rho_X, \mu_Y)$ and one in $\Pi(\rho_{X'}, \mu_{Y'})$, for different X, X' and Y, Y' . In this case, neither $W_2(\hat{\gamma}_0(X, Y), \hat{\gamma}_1(X', Y'))$ nor $\|\gamma_0(X, Y) - \gamma_1(X', Y')\|_{TV}$ are represented by the geometry of $\Pi(\alpha, \beta)$. It is unclear, then, that the quantitative stability results in linear programming can be directly applied to the primal problem to deduce those of a functional analysis flavour. Regardless, the linear programming formulation will give us some qualitative insight into how optimal plans deform as we vary X and Y continuously.

It is natural to ask an analogue of Problem 2 for the stability of the set of dual optimisers Φ . This is relevant both for independent interest as well as that it may hopefully allow better understanding of Problem 2.

Problem 3. How do $\Phi_c(X, Y)$ and $\Psi_c(X, Y)$ change as we perturb X and Y ? If the supports are close so that $\|X - X'\|, \|Y - Y'\| \ll 1$ (and hence $W_2(\rho_X, \rho_{X'})$ and $W_2(\mu_Y, \mu_{Y'})$ are small), are $\Phi_c(X, Y)$ and $\Phi_c(X', Y')$ quantitatively close in some sense? We would like results of the form

$$d_{H, \|\cdot\|_\phi}(\Phi_c(X, Y), \Phi_c(X', Y')) \leq C \left(W_2(\rho_{X_0}, \rho_{X_1}) + W_2(\mu_{Y_0}, \mu_{Y_1}) \right)^q + \mathcal{E} \quad (4.2)$$

where $d_{H, \|\cdot\|_\phi}$ is the Hausdorff distance between sets on $\mathbb{R}^N$ with respect to a choice of norm $\|\cdot\|_\phi$, and $\mathcal{E}$ is an error term to be understood based on the structure of the measures.

A direct comparison between vectors $\phi \in \Phi(X, Y)$ and $\phi' \in \Phi(X', Y')$ for some norm on $\mathbb{R}^M$ has no direct interpretation if $X \neq X'$. If, however, $X = X'$, then norm distances do have a useful interpretation. For example

$$\|\phi\|_{\mathbb{R}^M, \alpha} := \sum_{i=1}^M \alpha_i |\phi_i|^2,$$

so that if $X = X'$, then for $\phi \in \Phi(X, Y)$ and $\phi' \in \Phi(X, Y')$ we have

$$\|\phi - \phi'\|_{\mathbb{R}^M, \alpha} = \|\phi^{**} - \phi'^{**}\|_{L^2(\rho)},$$

where here $\phi^{**} : \mathcal{X} \rightarrow \mathbb{R}$ corresponds to the double Legendre transform with respect to X then Y , so that ϕ^{**} restricted to X returns the vector ϕ . (This is only true when ϕ was already

assumed to be optimal; in general, the restriction of the double Legendre transform is just the restriction of the largest convex function which is less than ϕ_i at each x_i , which may not return ϕ_i itself.) Thus, a norm on $\mathbb{R}^N$ corresponds directly to an $L^2(\rho)$ distance between continuous potentials which is amenable to qualitative limits.

Problem 4. Use the understanding gained from Problem 2 and 3 to take qualitative limits of (4.1) and (4.2) to say things about general transport problems beyond the discrete case. In this way, a sequence of linear programs which may a priori seem unrelated can correspond to weakly converging marginal sequences of discrete measures for which (4.1) may hopefully be stable.

Given that the norm structure on $D(C)$ is much more useful than that of $\Pi(\alpha, \beta)$, to approach Problem 2 it seems reasonable to first understand component wise effects on $\phi \in \Phi(X, Y)$ of perturbing Y through the explicit characterisation given in Proposition 3.7. The ϕ_i values correspond to the translations of hyperplanes with gradient x_i , casting this over the target domain we could hope to understand how the Laguerre cells induced on the target domain evolve. Since there should be the same amount of mass in each cell, we could use this information to construct a coupling between two different transport plans with the same first marginal, giving an estimate on $W_2(\gamma_{X,Y}, \gamma_{X,Y'})$.

4.2 Structure of primal optimisers under deformations

In this section, we provide a qualitative explanation of how the set of optimisers $\Gamma_c(X, Y)$ defined in (3.8) deforms under Perturbations of X and Y . Intuitively, the idea is as follows: if we perturb the positions of each of the masses continuously such that at each point in time, the transport problem has a **unique solution**, then the structure of the optimal plan remains the same. For the structure of the optimal plan to change, it must be that at some point, one way of transporting becomes better than another, and at the point of this change, the before plan and after plan must have the same cost, contradicting uniqueness.

We first present a direct proof in the language of measure-theoretic optimal transport using the c -cyclical monotonicity (1.10). We then describe how this connects to uniqueness in linear programming through the characterisation of extreme points of $\Pi(\alpha, \beta)$, Lemma 3.1.

Definition 4.2 (Glue-composition of transport plans). *Let ρ, μ_0, μ_1 be probability measures on $\mathbb{R}^d$ and consider transport plans $\gamma_0 \in \Pi(\rho, \mu_0)$ and $\pi \in \Pi(\mu_0, \mu_1)$. Then $\gamma_1 \in \Pi(\rho, \mu_1)$ is said to be a glue-composition of γ_0 and π if there exists a measure $\eta \in \Pi(\rho, \mu_0, \mu_1)$ with $p_{12\#}\eta = \gamma_0$, $p_{23\#}\eta = \pi$ and $p_{13\#}\eta = \gamma_1$. Here p_{ij} represents the canonical projection onto the i th and j th coordinates. We denote by*

$$\gamma_0 \circ \pi \subset \Pi(\rho, \mu_1)$$

the set of all possible glue-compositions of γ_0 and π . There always exists at least one way to glue transport plans with a common marginal, since one can always take $\gamma_{0,y} \otimes \pi_y d\mu_0(y)$ where $\gamma_{0,y}$ and π_y are disintegrations of γ_0 and π in the second/first coordinates respectively, see [30, Lemma 5.5] for more details. Unlike the composition of maps, glue-compositions need not be unique, see Example 4.3.

Example 4.3 (Non-uniqueness of glueing). Consider the transport problem illustrated in Figure 4.2 between ρ_X and $\mu_{Y(t)}$ where

$$\alpha = (3/4, 1/4), \quad \beta = (1/2, 1/4, 1/4),$$

$$X = (A, B), \quad Y_0 = (C, D, D), \quad Y_1 = (E, F, G), \quad Y(t) = (1-t)Y_0 + tY_1.$$

There are multiple glue-compositions of the optimal plan $\gamma \in \Pi(\rho_X, \mu_{Y_0})$ with the optimal plan $\pi_t \in \Pi(\mu_{Y_0}, \mu_{Y(t)})$. Only one of them is the optimal plan for quadratic cost between ρ_X and $\mu_{Y(t)}$ (the middle diagram).




Figure 4.2: Measures ρ_X , μ_{Y_0} , $\mu_{Y(t)}$ and μ_{Y_1} and couplings between them (left); and two potential glue-compositions of γ_0 and π_t (middle and right).

Theorem 4.4 (Structure of discrete plans along trajectories preserving uniqueness). *Let $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ be continuous. Let $X : [0, 1] \rightarrow (\mathbb{R}^d)^M$ and $Y : [0, 1] \rightarrow (\mathbb{R}^d)^N$ be continuous curves such that*

- (i) *For each $t \in (0, 1)$ and any distinct indices $i \neq j$, $x_i(t) \neq x_j(t)$ and $y_i(t) \neq y_j(t)$.*
- (ii) *For each $t \in [0, 1]$ solutions to the c -cost transport between $\rho_{X(t)}$ and $\mu_{Y(t)}$ are unique.*

Then there exists $\hat{\gamma} \in \Pi(\alpha, \beta)$ such that

$$\Gamma_c(X(t), Y(t)) = \left\{ \sum_{i=1}^M \sum_{j=1}^N \hat{\gamma}_{ij} \delta_{x_i(t), y_j(t)} \right\} \text{ for all } t \in [0, 1],$$

so that the “same” plan is the unique optimal plan for all $t \in [0, 1]$. In other words, given the optimal plan for some $s \in (0, 1)$, the unique optimal plan for each other $t \in [0, 1]$ is a glue-composition (note there can be multiple glue compositions here, but only one can be optimal since we assume uniqueness) of $\pi_{ts}^\rho \circ \gamma(s) \circ \pi_{st}^\mu$ where

$$\pi_{ts}^\rho := \sum_{i=1}^M \alpha_i \delta_{x_i(t), x_i(s)} \in \Pi(\rho_{X(t)}, \rho_{X(s)}) \text{ and } \pi_{st}^\mu := \sum_{i=1}^M \beta_i \delta_{y_i(s), y_i(t)} \in \Pi(\mu_{Y(s)}, \mu_{Y(t)}).$$

Proof. Fix some $s \in (0, 1)$ and let $\gamma(s)$ be the unique optimal plan between $\rho_{X(s)}$ and $\mu_{Y(s)}$ for cost c , written in the form

$$\gamma(s) = \sum_{i=1}^M \sum_{j=1}^N \hat{\gamma}_{ij}(s) \delta_{x_i(s), y_j(s)}. \quad (\text{Here some } \hat{\gamma}_{ij}(s) \text{ may be zero.}) \quad (4.3)$$

Here a priori, the $\hat{\gamma}_{ij}(s)$ depend on the choice of s . We want to show that they are independent of s , so that the same $\gamma(t)$ is optimal for all $t \in [0, 1]$, rather than just $t = s$. The ways to choose a finite number of ways to choose distinct $\hat{P} \subset \hat{\gamma}(s)$ and distinct $P \subset \gamma(t)$ are in bijection, and there are a finite number of possibilities r . We denote these by $\{P_k(t)\}_{k=1}^r$, where each is of the form

$$P_k(t) = \{(x_{i_l}(t), y_{j_l}(t))\}_{l=1}^{L_k} \quad \text{with} \quad x_{i_l}(t) \neq x_{i_{l'}}(t), y_{j_l}(t) \neq y_{j_{l'}}(t) \quad \forall l \neq l', t \in (0, 1).$$

By Proposition 3.5, $\Delta_{P_k(s)} > 0$, and since c , $X(t)$ and $Y(t)$ are continuous, so is $t \mapsto \Delta_{P_k(t)}$. As there are a finite number of distinct support selections,

$$\underline{\Delta}(t) := \inf_{k=1, \dots, r} \Delta_{P_k(t)}$$

is continuous and $\underline{\Delta}(s) > 0$. Set

$$t^+ = \inf \{t \geq s : \underline{\Delta}(t) = 0\} \quad \text{and} \quad t^- = \sup \{t \leq s : \underline{\Delta}(t) = 0\}.$$

Assume $t^- \in (0, s)$, then there exists $P \subset \text{Supp } \gamma(t^-)$ such that $\Delta_P = 0$, and by continuity and definition of t^- as a supremum, $\Delta_{P'} \geq 0$ for all other $P' \subset \text{Supp } \gamma(t^-)$. Proposition 3.5 implies both that $\gamma(t^-)$ is optimal and that it is not the only optimiser, contradicting the non-uniqueness assumption. Hence, it must be that $\underline{\Delta}(t) > 0$ for all $t \in (0, s)$. Arguing symmetrically for t^+ , we deduce that $\underline{\Delta}(t) > 0$ for all $t \in (0, 1)$, and hence is the unique optimiser. This extends to $t \in [0, 1]$ by qualitative stability.

To explicit the glueing structure, for $s \in (0, 1)$ and $t \in [0, 1]$ consider the measure

$$\eta(s, t) := \sum_{i=1}^M \sum_{j=1}^N \hat{\gamma}_{ij} \delta_{x_i(t), x_i(s), y_j(s), y_j(t)} \in \Pi(\rho_{X(t)}, \rho_{X(s)}, \mu_{Y(s)}, \mu_{Y(t)}). \quad (4.4)$$

It follows that

$$p_{12\#}\eta(s, t) = \pi_{ts}^\rho; \quad p_{23\#}\eta(s, t) = \gamma(s); \quad \text{and} \quad p_{34\#}\eta(s, t) = \pi_{st}^\mu$$

and

$$p_{14\#}\eta(s, t) = \gamma(t) \quad \forall t \in [0, 1].$$

Thus, we have shown that the optimal plan over $\Pi(\rho_t, \mu_t)$ is given by glue composition, supported by the measure $\eta(s, t)$. $\square$

Remark 4.5. The connection between Theorem 4.4 and the linear programming formulation is made clear by the characterisation of extreme points of $\Pi(\alpha, \beta)$, Lemma 3.1. Since the weight vectors $\alpha \in \mathbb{R}^N$ and $\beta \in \mathbb{R}^M$ are fixed, varying X and Y corresponds to varying the direction of the cost vector $C(X, Y)$ inside $\Pi(\alpha, \beta)$. So long as we move $C(X, Y)$ in a way which preserves non-orthogonality to the face of $\Pi(\alpha, \beta)$ in the direction $-C$, we retain uniqueness, and the same vertex (and hence “same” plan) is optimal. If at some point $-C(X, Y)$ is orthogonal to a face, this entire face will be optimal (non-uniqueness). One of the vertices on this face must necessarily be the unique optimal vertex from which we arrived (this is qualitative stability/upper semicontinuity). When we are on the interior of this face, the graph $G(\hat{\gamma})$ has cycles, which correspond to multiple possible permutations of the support being optimal (we can construct multiple optimal plans à la c -cyclical monotonicity). One cannot perturb C continuously such that we have uniqueness all along the trajectory, with a different vertex $\hat{\gamma}_1 \in \Pi(\alpha, \beta)$ optimal at the end compared to $\hat{\gamma}_0$ at the beginning: If

$$\langle \hat{\gamma}_0 | C(0) \rangle > \langle \hat{\gamma}_1 | C(0) \rangle \quad \text{and} \quad \langle \hat{\gamma}_1 | C(1) \rangle > \langle \hat{\gamma}_0 | C(1) \rangle,$$

then by continuity of $C(t)$, the intermediate value theorem tells us there exists $t \in [0, 1]$ with

$$\langle \hat{\gamma}_0 | C(t) \rangle = \langle \hat{\gamma}_1 | C(t) \rangle \implies \langle \hat{\gamma}_0 - \hat{\gamma}_1 | C(t) \rangle = 0.$$

Thus $C(t)$ is orthogonal to the 1 dimensional edge connecting $\hat{\gamma}_0$ and $\hat{\gamma}_1$.

There is, however, no direct equivalence between the uniqueness of the measure-theoretic and linear programming formulations. If for some i_0, i_1 or j_0, j_1 we have $x_{i_0} = x_{i_1}$ or $y_{j_0} = y_{j_1}$, then we have non-uniqueness of the Linear program (3.2) whilst the measure theoretic problem (3.8) could retain uniqueness. Thus, we can pass from one vertex to another being optimal in $\Pi(\alpha, \beta)$ whilst $\Gamma(t)$ remains a singleton. This is illustrated in the example below.

Example 4.6 (Intersecting particles). For interpolations $X(t) = (1 - t)X_0 + tX_1$ and $Y(t) = (1 - t)Y_0 + tY_1$ that correspond to quadratic geodesic interpolations in the sense of Remark 4.9, the non-intersecting particle assumption (i) is guaranteed. Without this assumption, consider

$$\alpha = \beta = (1/2, 1/2),$$

$$X = (A, B), Y_0 = (C, D), Y_1 = (F, E),$$

as illustrated in Figure 4.3. We retain uniqueness all along the trajectory, but after $t = 1/2$ when $y_1(t) = y_2(t)$, the glueing no longer corresponds to the optimal coupling.




Figure 4.3: Optimal couplings between ρ_X and $\mu_{Y(t)}$ for $t = 0.2$ and $t = 0.7$.

Lemma 4.7. Under the hypotheses of Theorem 4.4, let $X(t)$ and $Y(t)$ be of the form $X(t) = (1 - t)X_1 + tX_2$, $Y(t) = (1 - t)Y_1 + tY_2$ for $X_i \in (\mathbb{R}^d)^M$ and $Y_i \in (\mathbb{R}^d)^N$. Then there exists a glueing measure $\eta(0, 1) \in \Pi(\rho_{X_1}, \rho_{X_0}, \mu_{Y_0}, \mu_{Y_1})$ which describes all optimal plans between $\rho_{X(t)}$ and $\mu_{Y(t)}$ simultaneously.

Proof. Recalling (4.4) which defines $\eta(s, 1)$ for $s \in (0, 1)$, we set

$$\eta(0, 1) := \lim_{s \rightarrow 0^+} \eta(s, 1) \in \Pi(\rho_{X_1}, \rho_{X_0}, \mu_{Y_0}, \mu_{Y_1})$$

where the limit is in the sense of weak convergence since $X(s) \rightarrow X_0$ and $Y(s) \rightarrow Y_0$ are continuous. Explicitly, (4.4) tells us that

$$\eta(0, 1) = \sum_{i=1}^M \sum_{j=1}^N \hat{\gamma}_{ij} \delta_{x_i^1, x_i^0, y_j^0, y_j^1} \in \Pi(\rho_{X(t)}, \rho_{X(s)}, \mu_{Y(s)}, \mu_{Y(t)})$$

where $\hat{\gamma} \in \Pi(\alpha, \beta)$ is the vertex corresponding to optimality for $s \in (0, 1)$ in Theorem 4.4. Then since $x_i(t) = (1 - t)x_i^0 + tx_i^1$ and similarly for $Y(t)$, the optimal plan $\gamma(t) \in \Gamma(X(t), Y(t))$ is

$$\gamma(t) = (tx^1 + (1 - t)x^0, (1 - t)y^0 + ty^1)_{\#} \eta(0, 1) \text{ for all } t \in [0, 1].$$

□

Remark 4.8. Another way of formulating the above is as follows. For $X(t), Y(t) : [0, 1] \rightarrow (\mathbb{R}^d)^M \times (\mathbb{R}^d)^N$ continuous, the measures $\rho_{X(t)}, \mu_{Y(t)}$ are absolutely continuous curves in $(\mathcal{P}_p(\mathbb{R}^d), W_p)$ for $p > 1$. By [30, Chapter 5] there exist vector fields v_t^X, v_t^Y convecting these curves, in that

$$\partial_t \rho_{X(t)} + \nabla \cdot (v_t^X \rho_{X(t)}) = 0, \quad \partial_t \mu_{Y(t)} + \nabla \cdot (v_t^Y \mu_{Y(t)}) = 0.$$

(Here one can explicitly just take vector fields agreeing a.e. with the coordinates of $X'(t)$ and $Y'(t)$.) Under the Hypotheses of Theorem 4.4, for γ_t the unique optimal plan, we have

$$\partial_t \gamma_t + \nabla \cdot ((v_t^X, v_t^Y) \gamma_t) = 0.$$

Remark 4.9. The reason why we must define $\eta(0, 1)$ delicately is that in (i) we do not assume the coordinates are distinct at $s = 0$. Thus, there may be multiple ways to glue, only one of which will give the optimal plan by our uniqueness assumption. Taking the limit down from $s > 0$ ensures we select the right glueing by qualitative stability. The reason we prove the result with these slightly more general hypotheses (particles can intersect at $t = 0, 1$) is explained below.

Take couplings in $\pi^\rho \in \Pi(\rho_{X_0}, \rho_{X_1})$ and $\pi^\mu \in \Pi(\rho_{Y_0}, \rho_{Y_1})$. Let $M = \#\text{Supp } \pi^\rho$ and $N = \#\text{Supp } \pi^\mu$, then let each coordinate of X_0, X_1 correspond to a point in $\text{Supp } \pi^\rho$ with

$$\alpha_i = \pi^\rho((x_i^0, x_i^1)) \quad \text{and} \quad \beta_j = \pi^\mu((y_j^0, y_j^1))$$

for each distinct pair of points $x^0, x^1 \in \text{Supp } \pi^\rho$ and $y^0, y^1 \in \text{Supp } \pi^\mu$. If π^ρ and π^μ are chosen to be optimal for p cost with $p > 1$, then the interpolations $X(t) = (1 - t)X_0 + tX_1$ and $Y(t) = (1 - t)Y_0 + tY_1$ correspond to W_p geodesic interpolations, and in fact all W_p geodesics are of this form, see [30, Chapter 5].

$$\rho_t = ((1 - t)x_0 + tx_1)_\# \pi^\rho \quad \text{and} \quad \mu_t = ((1 - t)y_0 + ty_1)_\# \pi^\mu.$$

Thus, we retain the slight generality of allowing particles to intersect at $t \in \{0, 1\}$ in the statement of Theorem 4.4: in general transport plans will split mass on departure and arrival, and we would like these interpolations to be covered by the theorem.

Lemma 4.10. *Let $c(x, y) = \|x - y\|^p$ with $p > 1$. Assume that $X(t) = (1 - t)X_0 + tX_1$ corresponds to a W_p geodesic in the sense of Remark 4.9. Then $X(t)$ satisfies the non-crossing particle assumption (i).*

Proof. This is a straightforward consequence of c -monotonicity and strict convexity, following an argument very similar in flavour to [30, Theorem 2.9]. Assume that $(x_0, x_1), (z_0, z_1) \in \text{Supp } \pi^\rho$, where $\pi^\rho \in \Pi(\rho_{X_0}, \rho_{X_1})$ is the p cost optimal plan along which $X(t)$ is supported, in the sense of Remark 4.9. Let

$$x_t = (1 - t)x_0 + tx_1 = x_0 + tv_x \quad \text{and} \quad y_t = (1 - t)z_0 + tz_1 = z_0 + tv_z,$$

where $v_x = x_1 - x_0$ and $v_z = z_1 - z_0$. By p -cost cyclical monotonicity,

$$\|x_0 - x_1\|^p + \|z_0 - z_1\|^p \leq \|z_1 - x_0\|^p + \|z_0 - x_1\|^p. \quad (4.5)$$

Assume there exists $t \in (0, 1)$ such that

$$x(t) = (1 - t)x_0 + tx_1 = (1 - t)z_0 + tz_1 = z(t). \quad (4.6)$$

Then for this t ,

$$z_1 - x_0 = z_1 - z(t) + x(t) - x_0 = (1 - t)v_z + tv_x.$$

and similarly $z_0 - x_1 = tv_z + (1 - t)v_x$. If $v_z = v_x$, then (4.6) forces that $x_0 = z_0$ and $x_1 = z_1$. Thus, assuming $v_z \neq v_x$, By strict convexity of $\|\cdot\|^p$,

$$\|(1 - t)v_z + tv_x\|^p + \|tv_z + (1 - t)v_x\|^p < (1 - t)\|v_z\|^p + t\|v_x\|^p + t\|v_z\|^p + (1 - t)\|v_x\|^p \quad (4.7)$$

$$= \|x_0 - x_1\|^p + \|z_0 - z_1\|^p. \quad (4.8)$$

Combining the strict inequality above with (4.5) gives a contradiction, and so there can exist no such $t \in (0, 1)$ intersection point. $\square$

The qualitative structural result Theorem 4.4 allows for the following quantitative result.

Theorem 4.11 (Stability of plans under uniqueness). *For the p cost with $1 < p < \infty$, let ρ_0, ρ_1, μ_0 and μ_1 be discrete measures in $\mathbb{R}^d$. Let $\gamma_0 \in \Pi(\rho_0, \mu_0)$ and $\gamma_1 \in \Pi(\rho_1, \mu_1)$ be optimal plans. Set $(\rho_t)_{t \in [0,1]}$ and $(\mu_t)_{t \in [0,1]}$ as W_p geodesics connecting ρ_0 to ρ_1 and μ_0 to μ_1 . Assume for each $t \in [0, 1]$ that the W_p transport problem over $\Pi(\rho_t, \mu_t)$ has a unique solution. Then*

$$W_p^p(\gamma_0, \gamma_1) \leq C_p(W_p^p(\rho_0, \rho_1) + W_p^p(\mu_0, \mu_1)).$$

where $C_p = \max(1, 2^{\frac{p}{2}-1})$.

Proof. Let π^ρ, π^μ be the plans supporting the geodesics ρ_t and μ_t , in the sense of Remark 4.9. Since π^ρ, π^μ are optimal, the non-intersecting particle assumption (i) is satisfied due to Lemma 4.10, so apply Theorem 4.4. Hence γ_1 has the structure

$$\gamma_1 = \pi^\rho \circ \gamma_0 \circ \pi^\mu.$$

Let $\eta \in \Pi(\rho_1, \rho_0, \mu_0, \mu_1)$ be the corresponding glueing measure, so that $p_{12\#}\eta = \pi^\rho$, $p_{23\#}\eta = \gamma_0$, $p_{34\#}\eta = \pi^\mu$ and $p_{14\#} = \gamma_1$. Then η defines a coupling between γ_0 and γ_1 so that

$$\begin{aligned} W_p^p(\gamma_0, \gamma_1) &\leq \int_{\mathbb{R}^{4d}} \|(x_0, y_0) - (x_1, y_1)\|^p d\eta(x_1, x_0, y_0, y_1) \\ &\leq C_p \int_{\mathbb{R}^{4d}} \|x_0 - x_1\|^p + \|y_0 - y_1\|^p d\eta(x_1, x_0, y_0, y_1) \\ &= C_p \left(\int_{\mathbb{R}^{2d}} \|x_1 - x_0\|^p d\pi^\rho(x_1, x_0) + \int_{\mathbb{R}^{2d}} \|y_0 - y_1\|^p d\pi^\mu(y_0, y_1) \right) \\ &= C_p(W_p^p(\rho_1, \rho_0) + W_p^p(\mu_0, \mu_1)). \end{aligned}$$

□

Remark 4.12. The distances between plans and between marginals are actually comparable under these assumptions, since we always have the reverse inequality as a consequence of Proposition 1.3. In the case $p = 2$, $c_p = C_p = 1$ so that $W_2^2(\gamma_0, \gamma_1) = W_2^2(\rho_0, \rho_1) + W_2^2(\mu_0, \mu_1)$.

5 Uniqueness of continuous Kantorovich potentials

In this chapter, we study the uniqueness properties of the continuous dual problem (1.2). In general, optimisers to (1.2) or (3.3) are at best unique up to a constant, since given any (ϕ, ψ) admissible, $(\phi + \lambda, \psi - \lambda)$ is also admissible for any $\lambda \in \mathbb{R}$ and has the same cost. We provide what is, to our knowledge, the first dual uniqueness criterion applicable when both measures are concentrated on lower-dimensional subsets of $\mathbb{R}^d$. We then use this characterisation to give a combinatorial, graph-based characterisation of uniqueness akin to Lemma 3.2 and Proposition 3.4.

5.1 Literature review

We begin by reviewing the literature on the uniqueness of the dual problem. As discussed in Chapter 3, the uniqueness of the discrete primal and dual problems is understood by characterising the extreme points of the respective polyhedra. Lemma 3.1 and Lemma 3.2 describe these in terms of a bipartite graph between source and target support points due to [6, 1]. Uniqueness of the primal problem occurs if and only if the corresponding graph G_Γ has no cycles, and uniqueness of the dual occurs if and only if G_Γ is connected. We also mention [27], which arrives at an essentially equivalent graph theoretic criterion for dual potentials in the quadratically regularised discrete problem, as well as giving a quantitative description akin to [1, Theorem 3.9].

We now turn to the uniqueness of the continuous dual problem (1.2). The following is the classically dual uniqueness criterion.

Proposition 5.1. [30, Proposition 7.18] *Let ρ and μ be probability measures supported on compact subsets $\mathcal{X}$ and $\mathcal{Y}$ of $\mathbb{R}^d$, let $c \in C^1(\mathcal{X} \times \mathcal{Y})$. Assume that $\text{Supp } \rho$ is the closure of a connected open set. Then optimisers to the dual problem (1.2) are unique up to a constant.*

Two recent papers [32, 35] have approached the uniqueness of the continuous dual problem. Neither explicitly states a bipartite graph theoretic formulation, but both give results which are of this flavour. We also mention [27], which establishes an analogue of Proposition 5.1 in the quadratically regularised case. In [35, Theorem 4.5], the authors characterise uniqueness under the assumption that $\text{Supp } \rho$ is the closure of a connected open set. Here they use what amounts to a bipartite graph-like connectedness condition à la Proposition 3.4, which is asymmetric in ρ and μ . One set of nodes is chosen to be the interiors of the connected components of ρ , while the other is simply the connected components of μ . In [32], the authors also consider the decomposition of $\text{Supp } \rho$ and $\text{Supp } \mu$ into connected components, presenting results under the assumption that the restriction of Kantorovich potentials to each component is unique, as well as presenting some results under the “closure of a connected open set” support hypothesis.

5.2 Uniqueness

We begin by extending Proposition 5.1. The assumption “closure of a connected open set” is inherently “full dimensional” in that each point in the interior should contain an open neighbourhood, so $\text{Supp } \rho$ cannot have any lower dimensional structure. Our result provides what is to our knowledge, the first dual uniqueness result where *both* measures are supported on sets of dimension strictly smaller than the ambient space. We will require the following assumptions on the cost

Assumption 1.

- (i) c is continuous and bounded below.
- (ii) For fixed $y \in \mathcal{Y}$, $x \mapsto c(x, y)$ is differentiable for all $x \in \mathcal{X}$.
- (iii) The y parametrised family of functions $x \mapsto c(x, y)$ are equilocally Lipschitz - given a compact $K \subset \mathcal{Y}$ we can choose a Lipschitz constant L_K which works for all $y \in K$.

In particular, $c(x, y) = \|x - y\|^p$ satisfies Assumption 1 for all $p > 1$ (for $p \in (0, 1]$, differentiability of $x \mapsto \|x - y\|^p$ fails at $x = y$). Continuity of c forces that the maps $x \mapsto c(x, y)$ are locally bounded in x and y . These hypotheses induced the following regularity on c -concave (and hence optimal) potentials. We leave the proof to the Appendix, Lemma A.1.

Lemma 5.2. *Assume c satisfies Assumption 1, and $\mathcal{Y}$ is compact. Then given any $\psi \in C_b(\mathcal{Y})$, ψ^c defined as in (1.4) is Locally Lipschitz on $\mathcal{X}$, and does not take the value $-\infty$ anywhere.*

Definition 5.3. *We say a set Ω is Lipschitz-path connected if for any $x_0, x_1 \in \Omega$, there exists a finite length curve $\omega : [0, 1] \rightarrow \Omega$ such that $\omega(0) = x_0$ and $\omega(1) = x_1$. Without loss of generality, we take ω to be parametrised with constant speed, hence Lipschitz.*




Figure 5.1: Two Lipschitz path-connected sets. Their interiors are not connected, so known uniqueness results do not apply to measures with this support. Uniqueness with bounded target, Theorem 5.4 applies to both domains, whilst for unbounded target, Theorem 5.10 only applies to the left picture.

In particular, Lipschitz-path connected spaces include many lower-dimensional subspaces of $\mathbb{R}^d$ which are not covered by the closure of the interior style assumption, which is inherently d -dimensional. Measures with such support, include, for example, measures with densities with respect to some Hausdorff measure $\mathcal{H}^k$, restricted to some submanifold of $\mathbb{R}^d$.

Theorem 5.4. *Let $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ satisfy Assumption 1, let ρ and μ be probability measures on $\mathbb{R}^d$. Assume that*

- (i) $\text{Supp } \rho$ is Lipschitz-path connected,
- (ii) $\text{Supp } \mu$ is bounded.

Then optimisers of the dual problem (1.2) are unique up to a constant.

Proof. Let $\phi_0, \phi_1 \in L^1(\rho)$ be two optimal dual potentials for (1.2). Take any $x_0, x_1 \in \text{Supp } \rho$, let $\omega : [0, 1] \rightarrow \text{Supp } \rho$ be a Lipschitz curve with $\omega(0) = x_0$ and $\omega(1) = x_1$. We set

$$\tilde{\phi}_i : [0, 1] \rightarrow \mathbb{R}; \quad \tilde{\phi}_i(t) := \phi_i(\omega(t)).$$

By Lemma 5.2, each ϕ_i is locally Lipschitz and hence each $\tilde{\phi}_i$ is Lipschitz as a composition of Lipschitz functions. By Rademacher's theorem, both $\tilde{\phi}_i$ and ω are differentiable t -a.e., so the intersection of these points has full Lebesgue measure on $[0, 1]$. Let t be such a differentiability point. Fix any optimal $\gamma \in \Gamma(\rho, \mu)$. Since $\text{Supp } \mu$ is bounded $p_X(\text{Supp } \gamma) = \text{Supp } \rho$ where p_X denotes projection onto the x coordinate, see Appendix, Lemma A.2. Hence for each $t \in [0, 1]$, there exists at least one y such that $(\omega(t), y) \in \text{Supp } \gamma$. By the compatibility condition (1.8), $y \in \partial^c \phi_i(\omega(t))$. Applying a c -superdifferential inequality (1.9) at $\omega(t)$ for $t \in (0, 1)$ gives

$$c(\omega(t), y) - \phi_i(\omega(t)) \leq c(\omega(t+h), y) - \phi_i(\omega(t+h)).$$

Considering small $h > 0$ and $h < 0$, dividing by h and passing to the limit $h \rightarrow 0$ which exists by differentiability, we deduce

$$\tilde{\phi}'_i(t) = \langle \nabla_x c(\omega(t), y) | \omega'(t) \rangle.$$

The above shows that for a.e. t , the projection of $\nabla_x c(\omega(t), \partial^c \phi_i)$ in the direction $\omega'(t)$ is unique, and specified to be the same for any dual potential. It follows that

$$\tilde{\phi}'_1(t) - \tilde{\phi}'_0(t) = 0 \text{ for a.e. } t \in [0, 1],$$

and consequently

$$\phi_0(x_1) - \phi_0(x_0) = \int_0^1 \tilde{\phi}'_0(t) dt = \int_0^1 \tilde{\phi}'_1(t) dt = \phi_1(x_1) - \phi_1(x_0)$$

so that $\phi_1(x_0) - \phi_0(x_0) = \phi_1(x_1) - \phi_0(x_1)$ for all $x_0, x_1 \in \text{Supp } \rho$. In other words, ϕ_0 and ϕ_1 are equal up to a constant. $\square$

Example 5.5. For quadratic cost $c(x, y) = -\langle x | y \rangle$, it is the projection of $\partial \phi_i$ in the direction $\omega'(t)$ which is a.e. unique. In general, however, the set of $t \in [0, 1]$ along a given path for which there are multiple $y \in \partial \phi(\omega(t))$ can have non-zero Lebesgue measure. The above proof




Figure 5.2: Transport between $\rho = \mathcal{H}^1|_{\{0\} \times [0,1]}$ and $\mu = \frac{1}{2}(\mathcal{H}^1|_{\{-1\} \times [0,1]} + \mathcal{H}^1|_{\{1\} \times [0,1]})$.

establishes that the set of possible projections of all such y in the direction $\omega'(t)$ is a.e. a singleton. Consider the following example given in [30, Section 1.4] and illustrated in Figure 5.2. On $\mathbb{R}^2$, let

$$\rho = \mathcal{H}^1|_{\{0\} \times [0,1]} \quad \text{and} \quad \mu = \frac{1}{2}\mathcal{H}^1|_{\{-1\} \times [0,1]} + \frac{1}{2}\mathcal{H}^1|_{\{1\} \times [0,1]}.$$

Here, Theorem 5.4 applies so we have uniqueness. Despite this, at every point in $\text{Supp } \rho$, any optimal potential's subgradient is multivalued; it must contain at least two points corresponding to the horizontal splitting of the mass of the optimal plan (and by convexity it will also contain the full line between them). But the projection of these gradient values vertically (directions which stay in $\text{Supp } \rho$) gives a single value.

Example 5.6. There is no reason why we cannot have uniqueness of the dual problem when the primal problem has non-uniqueness. On $\mathbb{R}^2$, let

$$\rho = \mathcal{H}^1|_{\{0\} \times [-1,1]} \quad \text{and} \quad \mu = \mathcal{H}^1|_{[1,3] \times \{0\}}.$$

Here every transport plan $\gamma \in \Pi(\rho, \mu)$ is optimal, since $\langle x_0 - x_1 | y \rangle = 0$ for any $x_0, x_1 \in \text{Supp } \rho$ and $y \in \text{Supp } \mu$, and so all plans have cyclically monotone support.




Figure 5.3: Transport between $\mathcal{H}^1|_{\{0\} \times [-1,1]}$ and $\mu = \mathcal{H}^1|_{[1,3] \times \{0\}}$.

Any plan is optimal, so the primal has non-uniqueness while the dual solution is unique up to a constant.

Remark 5.7. Assume we have two Lipschitz-path connected sets A_0, A_1 in $\text{Supp } \rho$, and consider a “distance” between each defined by

$$d(A_0, A_1) = \inf_{\substack{x_0 \in A_0 \\ x_1 \in A_1}} \|x_0 - x_1\|.$$

(Of course, this is not a proper distance, as it satisfies neither positive definiteness nor the triangle inequality.) If $d(A_0, A_1) = 0$, then in the quadratic case, (still assuming $\text{Supp } \mu$ is bounded so that ϕ is Lipschitz) ϕ is unique up to a constant on $A_0 \cup A_1$. To see this, let ϕ_0 and ϕ_1 be two optimal potentials. Theorem 5.4 tells us that when restricted to each A_i , $\phi_1 = \phi_0 + a_i$ for some $a_i \in \mathbb{R}$. Take $\{x_i^n\}_{k=1}^\infty \subset A_i$ points such that $\|x_0^n - x_1^n\| \rightarrow 0$ as $n \rightarrow \infty$. Then

$$|a_0 - a_1| = |\phi_1(x_0^n) - \phi_0(x_0^n) - (\phi_1(x_1^n) - \phi_0(x_1^n))| \leq 2L\|x_0^n - x_1^n\| \rightarrow 0.$$

Where L is the Lipschitz constant of the ϕ_i (which can be taken to be $R_{\text{Supp } \mu}$, the radius of the smallest ball containing $\text{Supp } \mu$). Hence, $a_0 = a_1$ and potentials are specified up to a constant on $A_0 \cup A_1$ as required. More generally, the above can easily be extended to ρ with support which is connected and is the union of finitely many Lipschitz-path connected components. Here, one can decompose ρ into its Lipschitz-path components, observe that the graph with vertices as these components and an edge if and only if the two components are distance zero apart is connected, then argue as above to deduce that the potential is unique up to a constant across the entire support. The limit of applicability for this argument seems to be to any $\text{Supp } \rho$ such that for any $x_0, x_1 \in \text{Supp } \rho$,

$$\inf_{\substack{\omega: [0,1] \rightarrow \text{Supp } \rho \\ \omega(0)=x_0 \ \omega(1)=x_1}} \left\{ \sum_{k=1}^{\infty} \|\omega(t_k^+) - \omega(t_k^-)\| : \omega \text{ piecewise Lipschitz with jumps at } t_0, t_1, \dots \right\} = 0.$$

However, this technique does not seem amenable to a more general connected $\text{Supp } \rho$, or even just a path-connected $\text{Supp } \rho$ (without the hypothesis that the path should be Lipschitz). In particular, assume $f : [0, 1] \rightarrow [0, 1]$ is a sample path of a Brownian motion. Consider $\rho = (\text{Id}, f)_\# \text{Leb}|_{[0,1]}$. Then $\text{Supp } \rho$ is path connected, but every path connecting two points is locally of infinite length and nowhere differentiable, so the above argument cannot propagate uniqueness along such a curve.

For $c(x, y) = \|x - y\|^2$, in certain cases, the above argument also allows the removal of the assumption that $\text{Supp } \mu$ is bounded.

Definition 5.8. *The relative interior of a set, denoted ri is the interior taken with respect to the affine hull of the set rather than the full space. Here, the affine hull of $\mathcal{X} \subset \mathbb{R}^d$ is defined as*

$$\text{aff}(\mathcal{X}) := \left\{ \sum_{i=1}^k \lambda_i x_i \mid k > 0, x_i \in \mathcal{X}, \lambda_i \in \mathbb{R} \text{ with } \sum_{i=1}^k \lambda_i = 1 \right\}.$$

Theorem 5.9. [28, Theorem 10.4]. *Let f be a proper convex function, then f is locally Lipschitz on its restriction to $\text{ri}(\text{dom } f)$.*

Theorem 5.10. *Let $c(x, y) = \|x - y\|^2$, let ρ and μ be probability measures on $\mathbb{R}^d$. Assume that $\text{Supp } \rho$ is Lipschitz-path connected such that for any $x_0, x_1 \in \text{Supp } \rho$ the path can be chosen with*

$$\omega(t) \in \text{int}(\text{conv } \text{Supp } \rho) \text{ for all } t \in [0, 1]. \quad (5.1)$$

Then optimisers of the dual problem (1.2) are unique up to a constant.

Proof. This follows almost directly from the proof of Theorem 5.4. Fix a t where (5.1) holds, and take $\varepsilon > 0$ with $B_\varepsilon(\omega(t)) \subset \text{int } \text{conv}(\text{Supp } \rho)$. By definition of the support, $\rho(B_\varepsilon(\omega(t))) > 0$, and so given some optimal plan $\gamma \in \Pi(\rho, \mu)$, we define

$$\tilde{\rho} = \frac{1}{\rho(B_\varepsilon(\omega(t)))} \rho|_{B_\varepsilon(\omega(t))}$$

and its image measure under γ denoted $\tilde{\mu}_\gamma$, defined by

$$\tilde{\mu}_\gamma(A) = \frac{1}{\rho(B_\varepsilon(\omega(t)))} \gamma(B_\varepsilon(\omega(t))) \times A \text{ for all measurable } A \subset \mathbb{R}^d.$$

By Theorem 5.9, any optimal potential ϕ is Lipschitz when restricted to $B_\varepsilon(\omega(t))$, and so it follows that $\text{Supp } \tilde{\mu}_\gamma$ is bounded. Consequently, we can apply Theorem 5.4 to the dual problem between $\tilde{\rho}$ and $\tilde{\mu}_\gamma$, which gives that the restrictions of all potentials for the original problem are unique up to a constant on $\text{Supp } \rho \cap B_\varepsilon(\omega(t))$. We conclude that potentials are unique globally up to a constant on $\text{Supp } \rho$ as required, by choosing overlapping neighbourhoods covering the path ω . $\square$

Remark 5.11. One could attempt to extend the above to only requiring

$$\omega(t) \in \text{ri}(\text{conv } \text{Supp } \rho) \text{ for all } t \in [0, 1],$$

following the same proof only instead using balls $B_\varepsilon^{\text{ri}}$ with respect to $\text{aff } \text{Supp } \rho$ rather than all of $\mathbb{R}^d$. There are two places in the proof of Theorem 5.4 uses that $\text{Supp } \mu$ is bounded. Once is for the potentials to be Lipschitz, this extends directly to a relative interior assumption by Theorem 5.9. The second is to apply Lemma A.2 and deduce $\text{Supp } \rho \subset p_X(\text{Supp } \gamma)$. This requires some work, and one should prove that for an optimal quadratic cost γ ,

$$(\text{ri } \text{conv } \text{Supp } \rho) \cap \text{Supp } \rho \subset p_X(\text{Supp } \gamma).$$

We now present an analogue of the bipartite graph condition Proposition 3.4 for the continuous problem.

Theorem 5.12. Let ρ and μ be probability measures on compact $\mathcal{X}$ and $\mathcal{Y}$ subsets of $\mathbb{R}^d$. Let $c \in C^1(\mathcal{X} \times \mathcal{Y})$. Assume that each support is a finite union of Lipschitz-path connected components $(A_i)_{i \in I}$ and $(B_j)_{j \in J}$ respectively. Define the bipartite graph $G = (V, E)$ with $V = (A_i)_{i \in I} \cup (B_j)_{j \in J}$ and

$$(A_i, B_j) \in E \iff \text{there exists an optimal } \gamma \in \Pi(\rho, \mu) \text{ with } \gamma(A_i \times B_j) > 0.$$

If G is connected, then solutions to the dual problem (1.2) are unique up to a constant.

Proof. Let $\phi_0, \phi_1 \in C_b(\mathcal{X})$ be two Brenier potentials, let $\psi_0, \psi_1 \in C_b(\mathcal{Y})$ be their corresponding dual potentials. In general, if $\gamma \in \Pi(\rho, \mu)$ is optimal, then any $\tilde{\gamma} \leq \gamma$ is also optimal for the c cost transport between its marginals, see Appendix, Lemma A.3. Consequently, the compatibility condition (1.8) implies that if $\phi \in L^1(\rho)$ is optimal for the dual problem over $\Pi(\rho, \mu)$, then it is optimal for the dual problem over $\Pi(\tilde{\rho}, \tilde{\mu})$, since $\text{Supp } \tilde{\gamma} \subset \text{Supp } \gamma \subset \partial\phi$.

Both $\mathcal{X}$ and $\mathcal{Y}$ are compact and $c \in C^1$ so c is also Lipschitz, and hence satisfies Assumption 1 in both variables. Thus by Theorem 5.4, there exist $a_i, b_j \in \mathbb{R}$ such that on each component A_i , $\phi_0 = \phi_1 + a_i$ and on each component B_j , $\psi_0 = \psi_1 + b_j$. We want to show that all a_i are equal to some $a \in \mathbb{R}$, and consequently each $b_j = -a$ also.

Fix some i, i' , let $x \in A_i$ and $x' \in A_{i'}$. We will show that $a_i = \phi_0(x) - \phi_1(x) = \phi_0(x') - \phi_1(x') = a_{i'}$. Let $i = i_1, j_1, i_2, j_2, \dots, j_{n-1}, i_n = i'$ be a path between A_i and $A_{i'}$ in G . Then there exist

$$x = x_1, \hat{x}_1, y_1, \hat{y}_1, x_2, \hat{x}_2, \dots, y_{n-1}, \hat{y}_{n-1}, x_n, \hat{x}_n = x'$$

with $x_k, \hat{x}_k \in A_{i_k}$ and $y_k, \hat{y}_k \in B_{j_k}$ such that for each $k = 1, \dots, n-1$ there exists $\gamma \in \Pi(\rho, \mu)$ optimal with $(\hat{x}_k, y_k) \in \text{Supp } \gamma \cap (A_{i_k} \times B_{j_k})$ and there exists $\gamma \in \Pi(\rho, \mu)$ optimal with $(x_{k+1}, \hat{y}_k) \in \text{Supp } \gamma \cap (A_{i_{k+1}} \times B_{j_k})$. Consequently, $(\hat{x}_k, y_k), (x_{k+1}, \hat{y}_k) \in \partial\phi_i$ and so for $l = 0, 1$,

$$\phi_l(\hat{x}_k) + \psi_l(y_k) = \langle \hat{x}_k | y_k \rangle \quad \text{and} \quad \phi_l(x_{k+1}) + \psi_l(\hat{y}_k) = \langle x_{k+1} | \hat{y}_k \rangle. \quad (5.2)$$

It follows that

$$\begin{aligned} a_i &= \phi_0(x) - \phi_1(x) = \phi_0(\hat{x}_1) - \phi_1(\hat{x}_1) = \psi_1(y_1) - \psi_0(y_1) \\ &= \psi_1(\hat{y}_1) - \psi_0(\hat{y}_1) = \phi_0(x_2) - \phi_1(x_2) = \dots = \phi_0(\hat{x}_n) - \phi_1(\hat{x}_n) = a_{i'}, \end{aligned}$$

and so $a_i = a_{i'} = -b_j = -b_{j'}$ for all $i, j \in I \times J$. In other words, the dual potentials are unique up to a constant as required. $\square$

Remark 5.13. In [32], the authors establish uniqueness under a hypothesis they refer to as *non-degeneracy*; that there exist no subsets of indices $I' \subset I$ and $J' \subset J$ such that

$$0 < \rho(\cup_{i \in I'} A_i) = \mu(\cup_{j \in J'} B_j) < 1. \quad (5.3)$$

This hypothesis is much stronger; it forces that *any* $\gamma \in \Pi(\rho, \mu)$, its support graph G_γ is connected. To see this, take any $\gamma \in \Pi(\rho, \mu)$ and assume there exist at least two distinct connected components of G_γ , choose one, and set I' and J' to be all the indices of vertices corresponding to this component. By definition of G_γ , for any $i \in I'$ and $j \notin J'$, $\gamma(A_i \times B_j) = 0$ and similarly for $i \notin I'$ and $j \in J'$. Consequently

$$\rho(\cup_{i \in I'} A_i) = \gamma(\cup_{i \in I'} A_i \times \mathcal{Y}) = \gamma(\cup_{i \in I'} A_i \times \cup_{j \in J'} B_j) = \gamma(\mathcal{X} \times \cup_{j \in J'} B_j) = \mu(\cup_{j \in J'} B_j).$$

By the assumption that G_γ has more than one component, each of I' and J' correspond to a proper subset of the sets of support components of ρ and μ respectively, and hence the $0 < \rho(\cup_{i \in I'} A_i) = \mu(\cup_{j \in J'} B_j) < 1$.

Remark 5.14. It would be interesting to consider if the above can be extended to the multi-marginal transport problem, considering the connectivity of an n -partite hypergraph.

When there are an infinite number of components, there can exist A_i with $\rho(A_i) = 0$ and B_j with $\mu(B_j) = 0$. This can occur when these components are limit sets of points from other components, see Example 5.15.

Example 5.15 (Uniqueness of potentials despite disconnected graph). In $\mathbb{R}^2$, set $\rho = \frac{1}{2}\mathcal{H}_{\{-1\} \times [0,1]}^1 + \frac{1}{2}\mathcal{H}_{\{1\} \times [0,1]}^1$. Define the sequence of points $\{x_n\}_{n \in \mathbb{N}} \subset \mathbb{R}^2$ by

$$x_{2n-1} := \left(-\frac{1}{2n}, \frac{1}{n}\right); \quad x_{2n} := \left(\frac{1}{2n}, \frac{1}{n}\right),$$

then set

$$\mu = \sum_{n=1}^{\infty} 2^{-n} (\delta_{x_{2n-1}} + \delta_{x_{2n}}),$$

see Figure 5.4. The connected components of $\text{Supp } \rho$ are $A_1 = \{-1\} \times [0, 1]$ and $A_2 = \{1\} \times [0, 1]$.




Figure 5.4: Transport between ρ and μ (Left), and the Bipartite support graph (Right). The graph is disconnected, but the uniqueness of potentials holds as the components are “asymptotically connected” via the point $(0, 0)$.

The connected components of $\text{Supp } \mu$ are $B_0 = \{(0, 0)\}$ and $B_n = \{x_n\}$ for $n \in \mathbb{N}$. Here, the component B_0 receives no mass from μ , but belongs to the support nonetheless as a limit point of points from the other components. The graph for the transport has three disjoint components, one from the left pairings of points and one from the right. Theorem 5.12 gives that the potentials are defined up to a constant on each component of the graph. Then continuity of any Brenier potential at $(0, 0)$ forces uniqueness up to a constant globally, as a consequence of the arguments presented in Remark 5.7.

One can also use calculations similar to those in the proof of Lemma 3.2 to ascertain quantitative uniqueness results, i.e. bounds on the diameter of the optimal set. In particular, we have the following

Proposition 5.16. *Let ρ and μ be probability measures on $\mathbb{R}^d$, let $c(x, y) = \|x - y\|^2$. Assume that each support is a finite union of Lipschitz-path connected components $(A_i)_{i \in I}$ and $(B_j)_{j \in J}$ respectively. Assume that the graph G defined in Theorem 5.12 consists of two connected components $G = G_0 \sqcup G_1$. Then for any two Brenier potentials ϕ_0, ϕ_1 , we have*

$$\inf_{\lambda \in \mathbb{R}} \|\phi_0 - \phi_1 - \lambda\|_{\infty} \leq \frac{1}{2} \inf_{\substack{(x_0, y_0) \in \text{Supp } \Gamma \cap G_0 \\ (x_1, y_1) \in \text{Supp } \Gamma \cap G_1}} \langle x_1 - x_0 | y_1 - y_0 \rangle. \quad (5.4)$$

Proof. Assume we know the values of a Brenier potential ϕ on the G_0 component of G . For any $(x_0, y_0) \in \text{Supp } \Gamma \cap G_0$ and $(x_1, y_1) \in \text{Supp } \Gamma \cap G_1$, we have $y_0 \in \partial\phi(x_0)$ and $y_1 \in \partial\phi(x_1)$. By

the subdifferential inequality applied at each (x_i, y_i) , the smallest and largest ϕ could be at x_1 , in terms of $\phi(x_0)$, is

$$\phi(x_1) \in \left[\phi(x_0) + \langle y_0 | x_1 - x_0 \rangle, \phi(x_0) + \langle y_1 | x_1 - x_0 \rangle \right]. \quad (5.5)$$

Hence for two potentials ϕ_0, ϕ_1 with $\phi_0(x_0) = \phi_1(x_0)$, rearranging the above gives

$$|\phi_0(x_1) - \phi_1(x_1)| \leq \langle x_1 - x_0 | y_1 - y_0 \rangle.$$

Since on each component of G the potentials are unique up to a constant,

$$|\phi_0(x_1) - \phi_1(x_1)| = |\phi_0(x'_1) - \phi_1(x'_1)|$$

for all $x_1, x'_1 \in G_1$, and so if the difference between two potentials is bounded at one place in G_1 , the same bound holds everywhere. (Note here by $x \in G_1$ we mean x belongs to the union of the A_i vertices in the component G_1 .) Thus, all differences globally on G_1 are bounded by the tightest possible bound at any one point in the component, so we can pass to the infimum. Hence (5.4) follows after dividing by 2 since we can optimise λ so the displacement is centred, rather than $\phi_0 = \phi_1$ on G_0 . $\square$

Remark 5.17. The right-hand side of (5.4) can be bounded using Cauchy-Schwarz to give

$$\inf_{\lambda \in \mathbb{R}} \|\phi_0 - \phi_1 - \lambda\|_\infty \leq \frac{1}{2} \min(\varepsilon_{\mathcal{X}} R_{\mathcal{Y}}, \varepsilon_{\mathcal{Y}} R_{\mathcal{X}}),$$

where $R_{\mathcal{X}}$ and $R_{\mathcal{Y}}$ are the smallest radii of balls centred at the origin containing $\text{Supp } \rho$ and $\text{Supp } \mu$ respectively, and $\varepsilon_{\mathcal{X}}$ is the largest pairwise distance attainable by partitioning (A_i) into two components,

$$\varepsilon_{\mathcal{X}} = \inf_{\mathbb{N}_1 \sqcup \mathbb{N}_2 = \{A_i\}_{i \in I}} d\left(\bigcup_{A \in \mathbb{N}_1} A, \bigcup_{A \in \mathbb{N}_2} A\right)$$

with d denoting the closest distance between sets defined in Remark 5.7, and $\varepsilon_{\mathcal{Y}}$ is defined symmetrically.

For example, suppose $\text{Supp } \rho$ has two Lipschitz-path connected components A_0 and A_1 , then we have the following bound on the uniform norm diameter of the optimal set

$$\sup_{\substack{\phi_0, \phi_1 \\ \text{optimal}}} \inf_{\lambda \in \mathbb{R}} \|\phi_0 - \phi_1 - \lambda\|_\infty \leq \frac{1}{2} d(A_0, A_1) \text{Diam}(\text{Supp } \mu).$$

Where here we replaced $R_{\mathcal{Y}}$ with $\text{Diam } \mathcal{Y}$ due to the translation invariance of quadratic optimal transport, as in Chapter 2.

When there are more than two connected components of G , one should instead consider the graph of connected components of G given the above distance quantity, and the diameter of the optimal set becomes less clear. One can see parallels between this calculation and those of Proposition 3.7. A better understanding of this type of graph structure is the subject of ongoing work.

6 Conclusions

The contributions of this work are:

Chapter 1: The observations and preliminary results on the stability of optimal transport plans.

Chapter 2: The extension of the strong convexity inequality to a larger class of measures using glueing, which implies stability of Wasserstein barycentres for a large class of measures.

Chapter 3: Unifying known literature on the discrete primal and dual problems. Explicit characterisation of the set of dual optimisers for the discrete transport problem. (Some unfinished calculations have been omitted here, which seem to suggest this formulation could be very helpful for studying the quantitative stability of Φ and Ψ .)

Chapter 4: A qualitative understanding of how fully discrete transport plans behave under perturbation of source support points - quantitatively stable as long as uniqueness is preserved, at which case the optimal plans may have jump discontinuities corresponding to changing vertex on the polyhedron $\Pi(\alpha, \beta)$. This is the first step towards a resolution of Problem 1 under the discrete hypothesis on ρ .

Chapter 5: Significant weakening of source hypotheses needed for uniqueness of optimisers for the dual problem. Lipschitz-path connected spaces are a huge class of spaces, including any connected differentiable submanifold of $\mathbb{R}^d$. Formalised graph-based uniqueness criteria in the continuous setting. First calculations in the direction of Quantitative uniqueness results (bounds on the diameter of the optimal set).

The main takeaway of this report is that the connections between linear programming and optimal transport should not be overlooked. After a period of intense interaction in the 20th century, modern optimal transport has, in recent decades, moved closer to analysis and PDEs. As shown in Chapter 3, the combinatorial characterisations of extreme points provide valuable intuition for the discrete transport problem, motivating the novel results of Chapters 4 and 5. This demonstrates that linear programming and combinatorial perspectives remain relevant for advancing our theoretical understanding of optimal transport. Problems 1, 2, and 3 remain wide open, and I look forward to tackling them during the course of my PhD...

Bibliography

- [1] Beatrice Acciaio and Berenice Anne Neumann. Characterization of transport optimizers via graphs and applications to Stackelberg–Cournot–Nash equilibria. *Mathematics and Financial Economics*, 19(1):67–99, 2025.
- [2] Hédÿ Attouch and Haïm Brezis. Duality for the sum of convex functions in general Banach spaces. In *North-Holland Mathematical Library*, volume 34, pages 125–133. Elsevier, 1986.
- [3] Michel L Balinski. Signatures des points extremes du polyedre dual du probleme de transport. 1983.
- [4] Michel L Balinski. Faces of dual transportation polyhedra. *Mathematical Programming Study*, 22:1–8, 1984.
- [5] Michel L Balinski. The Hirsch conjecture for dual transportation polyhedra. *Mathematics of Operations Research*, 9(4):629–633, 1984.
- [6] Michel L Balinski and Fred J Rispoli. Signature classes of transportation polytopes. *Mathematical programming*, 60(1):127–144, 1993.
- [7] Robert J Berman. Convergence rates for discretized Monge–Ampère equations and quantitative stability of optimal transport. *Foundations of Computational Mathematics*, 21(4):1099–1140, 2021.
- [8] Jan Boman. *Lp-estimates for very strongly elliptic systems*. Department of Mathematics, University of Stockholm, Sweden, 1982.
- [9] Yann Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on pure and applied mathematics*, 44(4):375–417, 1991.
- [10] Richard A Brualdi. *Combinatorial matrix classes*, volume 13. Cambridge University Press, 2006.
- [11] Guillaume Carlier, Alex Delalande, and Quentin Merigot. Quantitative stability of barycenters in the Wasserstein space. *Probability Theory and Related Fields*, 188(3):1257–1286, 2024.
- [12] Alex Delalande. *Quantitative Stability in Quadratic Optimal Transport*. PhD thesis, Université Paris-Saclay, 2022.
- [13] Alex Delalande and Quentin Merigot. Quantitative stability of optimal transport maps under variations of the target measure. *Duke Mathematical Journal*, 172(17):3321–3357, 2023.
- [14] Thomas O Gallouët and Quentin Mérigot. A lagrangian scheme à la Brenier for the incompressible Euler equations. *Foundations of Computational Mathematics*, 18(4):835–865, 2018.
- [15] Nicola Gigli. On Hölder continuity-in-time of the optimal transport map towards measures along a curve. *Proceedings of the Edinburgh Mathematical Society*, 54(2):401–409, 2011.

- [16] Martin Huesmann, Michael Goldman, and Dario Trevisan. Asymptotics for random quadratic transportation costs. *arXiv preprint arXiv:2409.08612*, 2024.
- [17] Pierre-Emmanuel Jabin and Antoine Mellet. Hölder and Sobolev regularity of optimal transportation potentials with rough measures. *arXiv preprint arXiv:2308.10804*, 2023.
- [18] Jun Kitagawa, Cyril Letrouit, and Quentin Mérigot. Stability of optimal transport maps on Riemannian manifolds. *arXiv preprint arXiv:2504.05412*, 2025.
- [19] Jun Kitagawa, Quentin Mérigot, and Boris Thibert. Convergence of a Newton algorithm for semi-discrete optimal transport. *Journal of the European Mathematical Society*, 21(9):2603–2651, 2019.
- [20] Victor Klee and Christoph Witzgall. Facets and vertices of transportation polytopes. Technical report, 1967.
- [21] Cyril Letrouit. Lectures on quantitative stability of optimal transport. 2025.
- [22] Cyril Letrouit and Quentin Mérigot. Gluing methods for quantitative stability of optimal transport maps. *arXiv preprint arXiv:2411.04908*, 2024.
- [23] Bruno Lévy. A numerical algorithm for $l_{\{2\}}$ semi-discrete optimal transport in 3D. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1693–1715, 2015.
- [24] Wenbo Li and Ricardo H Nochetto. Quantitative stability and error estimates for optimal transport plans. *IMA Journal of Numerical Analysis*, 41(3):1941–1965, 2021.
- [25] Quentin Mérigot. A multiscale approach to optimal transport. In *Computer graphics forum*, volume 30, pages 1583–1592. Wiley Online Library, 2011.
- [26] Octave Mischler and Dario Trevisan. Quantitative stability in optimal transport for general power costs. *arXiv preprint arXiv:2407.19337*, 2024.
- [27] Marcel Nutz. Quadratically regularized optimal transport: Existence and multiplicity of potentials. *SIAM Journal on Mathematical Analysis*, 57(3):2622–2649, 2025.
- [28] Ralph Tyrell Rockafellar. *Convex analysis:(pms-28)*. Princeton University Press, 1992.
- [29] Walter Rudin. *Real and complex analysis*. McGraw-Hill, Inc., 1987.
- [30] Filippo Santambrogio. *Optimal transport for applied mathematicians*. Springer, Birkhäuser, 2015.
- [31] Walter Schachermayer and Josef Teichmann. Characterization of optimal transport plans for the monge-kantorovich problem. *Proceedings of the American Mathematical Society*, 137(2):519–529, 2009.
- [32] Thomas Staudt, Shayan Hundrieser, and Axel Munk. On the uniqueness of Kantorovich potentials. *SIAM Journal on Mathematical Analysis*, 57(2):1452–1482, 2025.
- [33] Cédric Villani. *Topics in optimal transportation*, volume 58. American Mathematical Soc., 2021.
- [34] Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2008.
- [35] Yunan Yang, Levon Nurbekyan, Elisa Negrini, Robert Martin, and Mirjeta Pasha. Optimal transport for parameter identification of chaotic dynamics via invariant measures. *SIAM Journal on Applied Dynamical Systems*, 22(1):269–310, 2023.

A Appendix

Lemma A.1. *Assume c satisfies Assumption 1, and $\mathcal{Y}$ is compact. Then given any $\psi \in C_b(\mathcal{Y})$, ψ^c defined as in (1.4) is Locally Lipschitz on $\mathcal{X}$, and does not take the value $-\infty$ anywhere.*

Proof. Since $\mathcal{Y}$ is compact, the functions

$$x \mapsto c(x, y)$$

are locally uniformly Lipschitz by assumption. To see this, Fix some $K_{\mathcal{X}} \subset \mathcal{X}$ compact, and some $x_0 \in K_{\mathcal{X}}$, then by continuity of c in y ,

$$y \mapsto c(x_0, y)$$

Attains its maximum and minimum over $\mathcal{Y}$ by compactness. Since all $x \mapsto c(x, y)$ share the same Lipschitz constant on $K_{\mathcal{X}}$, it follows that the $x \mapsto c(x, y)$ are uniformly bounded on $K_{\mathcal{X}}$. Consequently,

$$|\psi^c(x)| = \left| \inf_{y \in \mathcal{Y}} c(x, y) - \psi(y) \right| \leq \|c(x_0, \cdot)\|_{\infty} + L_{K_{\mathcal{X}}} \text{Diam } \mathcal{Y} + \|\psi\|_{\infty} \text{ for all } x \in K_{\mathcal{X}},$$

where $L_{K_{\mathcal{X}}}$ is the Lipschitz constant of the equicontinuous family $c(\cdot, y)$ on $K_{\mathcal{X}}$. It follows that ψ^c is bounded on each compact set in $\mathcal{X}$, and $L_{K_{\mathcal{X}}}$ Lipschitz by classical envelope arguments as an infimum of uniformly Lipschitz functions, see [30, Box 1.8]. $\square$

Lemma A.2. *Let ρ and μ be probabilities on $\mathbb{R}^d$, and let $\gamma \in \Pi(\rho, \mu)$ with $\text{Supp } \mu$ bounded. Then $p_X(\text{Supp } \gamma) = \text{Supp } \rho$, so that for each $x \in \text{Supp } \rho$ there exists y with $(x, y) \in \text{Supp } \gamma$.*

Proof. One inclusion always holds without any hypotheses: Since $\text{Supp } \rho \times \mathbb{R}^d$ is a closed set of full measure, $\text{Supp } \gamma \subset \text{Supp } \rho \times \mathbb{R}^d$, and hence $p_X(\text{Supp } \gamma) \subset \text{Supp } \rho$. The inverse inclusion does not always hold in general, as it could be that $p_X(\text{Supp } \gamma)$ is not closed. For example, $\mathcal{X} = \mathcal{Y} = \mathbb{R}$ and γ supported on the graph of $\{(x, 1/x) : x > 0\}$, then $p_X(\text{Supp } \gamma) = (0, \infty)$, whereas the support of the $\mathcal{X}$ marginal is $[0, \infty)$.

Assume $\text{Supp } \mu$ is bounded. Necessarily $p_X(\text{Supp } \gamma)$ has full ρ mass, and so it holds that

$$\text{Supp } \rho \subset \overline{p_X(\text{Supp } \gamma)}.$$

Thus, it suffices to show that $p_X(\text{Supp } \gamma)$ is closed. Take $x_n \in p_X(\text{Supp } \gamma)$, with $x_n \rightarrow x$ then there exist y_n with $(x_n, y_n) \in \text{Supp } \gamma$, by compactness of $\text{Supp } \mu$ we can assume that y_n converges to some y , and hence $(x, y) \in \text{Supp } \gamma$ by closedness of $\text{Supp } \gamma$. In other words, $x \in p_X(\text{Supp } \gamma)$, establishing the result. $\square$

Lemma A.3. *If $\gamma \in \Pi(\rho, \mu)$ is an optimal transport plan for cost $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ integrable and $\pi \in \mathcal{M}^+(\mathbb{R}^d \times \mathbb{R}^d)$ such that $\pi \leq \gamma$, then*

$$\hat{\pi} := \frac{\pi}{\pi(\mathbb{R}^d \times \mathbb{R}^d)}$$

is an optimal transport plan between its marginals $\hat{\rho} = p_{X\#}\hat{\pi}$ and $\hat{\mu} = p_{Y\#}\hat{\pi}$.

Proof. Assume $\hat{\pi}$ is not optimal, then there exists a transport plan w with marginals $\hat{\rho}, \hat{\mu}$ such that $\int c dw < \int c d\hat{\pi}$. Thus the transport plan between marginals ρ and μ defined by

$$\hat{\gamma} = \hat{\pi}(\mathbb{R}^d \times \mathbb{R}^d)w + (\gamma - \pi)$$

has cost

$$\begin{aligned} \int c(x, y) d\hat{\gamma} &= \hat{\pi}(\mathbb{R}^d \times \mathbb{R}^d) \int c(x, y) dw + \int c(x, y) d(\gamma - \pi) \\ &< \hat{\pi}(\mathbb{R}^d \times \mathbb{R}^d) \int c(x, y) d\hat{\pi} + \int c(x, y) d(\gamma - \pi) \\ &= \int c(x, y) d\gamma, \end{aligned}$$

contradicting the optimality of γ . (For the well definition of $\hat{\gamma}$ as a transport plan it was crucial $\pi \leq \gamma$ to ensure $\hat{\gamma} \geq 0$.) □