

Quadratic-form Optimal Transport

Ruodu Wang*

Zhenyuan Zhang†

8th September 2025

Abstract

We introduce the framework of quadratic-form optimal transport (QOT), whose transport cost has the form $\iint c d\pi \otimes d\pi$ for some coupling π between two marginals. Interesting examples of quadratic-form transport cost and their optimization include inequality measurement, the variance of a bivariate function, covariance, Kendall's tau, the Gromov–Wasserstein distance, quadratic assignment problems, and quadratic regularization of classic optimal transport. QOT leads to substantially different mathematical structures compared to classic transport problems and many technical challenges. We illustrate the fundamental properties of QOT and provide several cases where explicit solutions are obtained. For a wide class of cost functions, including the rectangular cost functions, the QOT problem is solved by a new coupling called the diamond transport, whose copula is supported on a diamond in the unit square.

Keywords: Quadratic programming, diamond transport, quadratic assignment problem, Gromov–Wasserstein distance, regularization, submodularity

MSC classification 2020: 49Q22; 62H05; 91B70; 62H20

1 Introduction

Given probability measures μ on a space \mathfrak{X} and ν on a space \mathfrak{Y} , a transport plan, also called a coupling, is a joint distribution on $\mathfrak{X} \times \mathfrak{Y}$ with marginals μ and ν . It does not hurt to think of $\mathfrak{X} = \mathfrak{Y} = \mathbb{R}$ in this section, and \mathfrak{X} and \mathfrak{Y} will be general Polish spaces in the formal theory. The set of all such transport plans is denoted by $\Pi(\mu, \nu)$. The classic Kantorovich optimal transport (OT) problem is

$$\begin{aligned} & \text{to minimize} && \int c(x, y) d\pi(x, y) \\ & \text{subject to} && \pi \in \Pi(\mu, \nu), \end{aligned}$$

where $c : \mathfrak{X} \times \mathfrak{Y} \rightarrow \mathbb{R}$ is a fixed cost function. This problem can be written in a probabilistic form:

$$\text{to minimize} \quad \mathbb{E}[c(X, Y)] \quad \text{subject to} \quad X \stackrel{\text{law}}{\sim} \mu; Y \stackrel{\text{law}}{\sim} \nu,$$

where $X \stackrel{\text{law}}{\sim} \mu$ means the distribution of the random variable X is μ . We also call (X, Y) a coupling. The objective $\int c d\pi$ of the OT problem is called the transport cost. The OT problem and its numerous extensions have wide applications in various fields including statistics, machine learning, operations research, mathematical finance, and economics. We refer to Villani [2003, 2009], Santambrogio [2015] for the theory of OT, and Galichon [2018], Peyré and Cuturi [2019] for applied perspectives.

*Department of Statistics and Actuarial Science, University of Waterloo, Canada. Email: wang@uwaterloo.ca.

†Department of Mathematics, Stanford University, USA. Email: zzy@stanford.edu.

A classic application of OT in economics and operations research concerns the problems of assignment and matching. While the transport cost $\mathbb{E}[c(X, Y)]$ includes many quantities of interest in these problems, such as total production or total cost, a larger framework is needed when a notion of equality needs to be optimized, as we illustrate below.

Inequality minimization. Suppose that there are two types of non-transferable and indivisible resources A and B with their distributions given by μ and ν on \mathbb{R} , and a social planner needs to choose a coupling $\pi \in \Pi(\mu, \nu)$ to allocate pairs of resources to (possibly a continuous spectrum of) individuals. For two individuals with vectors of assigned resources (x, y) and (x', y') , their *discrepancy* is defined as a squared weighted sum of differences in each resource, that is,

$$(\theta_1|x - x'| + \theta_2|y - y'|)^2,$$

where the weights $\theta_1, \theta_2 \geq 0$ are given. If we consider general spaces $\mathfrak{X}, \mathfrak{Y}$ instead of \mathbb{R} , then the discrepancy is $(\theta_1 d_{\mathfrak{X}}(x, x') + \theta_2 d_{\mathfrak{Y}}(y, y'))^2$, where $d_{\mathfrak{X}}$ and $d_{\mathfrak{Y}}$ are some distances on \mathfrak{X} and \mathfrak{Y} .

The social planner would like to minimize the average (or total) discrepancy between two randomly selected individuals in the population, that is

$$\text{to minimize } \mathbb{E}[(\theta_1|X - X'| + \theta_2|Y - Y'|)^2] \quad \text{over } \pi \in \Pi(\mu, \nu), \quad (1.1)$$

where (X, Y) and (X', Y') are independently drawn from the distribution π .¹

If ν is degenerate (i.e., all individuals are assigned the same resource B) or $\theta_2 = 0$ (i.e., discrepancy in resource B does not matter), the objective in (1.1) becomes $2\theta_1^2$ times the variance of X , which is a natural measure of distributional inequality. Hence, the expectation in (1.1) can be seen as a generalization of the variance when two different types of quantities are compared simultaneously. In addition to the variance, the idea of computing the expected difference between two randomly selected individuals is used to define other classic measures of inequality, such as the Gini deviation and the Gini coefficient in economics; see Example 3.5.

For a concrete example, suppose that a government agency aims to enhance equality in a population by distributing a menu of economic benefits according to a given distribution ν (items in the menu cannot be combined or divided). The current wealth level of the population is described by a distribution μ . After the financial policy, the wealth and benefits of the population are distributed as $\pi \in \Pi(\mu, \nu)$, determined by the agency. Discrepancy between two individuals occurs when either their wealth levels or their benefits differ (or both). The problem (1.1) is to minimize such discrepancy according to a certain policy π .

To analyze (1.1), one may first look at some commonly encountered couplings. A quick observation is that if $\mu = \nu$ and π is the comonotone coupling π_{com} (formal definition in Section 2), then $Y = X$ and $Y' = X'$ (almost surely), and thus $|X - X'| = |Y - Y'|$. Expanding (1.1), we see that the transport cost (the expectation in (1.1)) consists of terms that do not depend on π , plus $2\theta_1\theta_2\text{Cov}(|X - X'|, |Y - Y'|)$. Since the distributions of $|X - X'|$ and $|Y - Y'|$ are determined solely by μ and ν , the covariance and hence the transport cost is maximized when $|X - X'| = |Y - Y'|$. From there, one may then wonder whether the antimonotone coupling π_{ant} minimizes (1.1). However, by choosing $\mu = \nu$ as the standard uniform distribution, π_{ant} also yields $|X - X'| = |Y - Y'|$, thus also maximizing transport cost. Therefore, intuitively, the optimal coupling must lie somewhere between the most positive coupling π_{com} and the most negative coupling π_{ant} . It turns out that the optimal coupling, solved in full generality in Section 6, has a special structure that is different from both independence and mixtures of π_{com} and π_{ant} .

Inspired by the above assignment problem, we propose a new formulation of OT, called the *quadratic-form optimal transport* (QOT). Given a cost function $c : (\mathfrak{X} \times \mathfrak{Y})^2 \rightarrow \mathbb{R}$, we define the

¹This problem has a Kantorovich formulation. For a practical application with finitely many individuals, one may further require π to be induced by a permutation (that is, the Monge formulation), but we will mainly focus on the Kantorovich formulation in our study, which is technically tractable and offers approximations for optimizers in the Monge setting (see Section 6).

QOT problem as:

$$\begin{aligned} & \text{to minimize} && \iint c(x, y, x', y') \, d\pi(x, y) \, d\pi(x', y') && (1.2) \\ & \text{subject to} && \pi \in \Pi(\mu, \nu). \end{aligned}$$

The term ‘‘quadratic-form’’ reflects that a discrete formulation of the problem (1.2) can be written into the minimization of a quadratic form (see Appendix A). This term also distinguishes (1.2) from optimal transport with the quadratic cost function $c(x, y) = (x - y)^2$, which has been widely studied in the classic OT literature.² Compared to the classic OT problem, the transport cost (1.2) in QOT is defined as a linear function of $\pi \otimes \pi$ instead of π itself, and hence the problem is non-linear. Clearly, (1.1) is a special case of the QOT problem.

Certain cost functions, such as those that only involve (x, x') or (y, y') , have a transport cost determined by the marginals μ, ν , and they are called *QOT-irrelevant*.

The flexibility of the 4-variate cost function c allows for a rich spectrum of interesting instances of QOT. We pay special attention to two general sub-classes of cost functions, the class of *type-XX* cost functions,³

$$c(x, y, x', y') = h(f(x, x'), g(y, y')) \quad \text{for some real-valued bivariate functions } f, g, h, \quad (1.3)$$

which includes (1.1), and the class of *type-XY* cost functions

$$c(x, y, x', y') = h(f(x, y), g(x', y')) \quad \text{for some real-valued bivariate functions } f, g, h. \quad (1.4)$$

Throughout, equations for the form of cost functions, such as (1.3) and (1.4), are meant to hold for all (x, y, x', y') . The terms XX and XY reflect the idea that the cost functions in (1.3) aggregate some costs (e.g., distances) between x, x' and between y, y' , and the cost functions in (1.4) aggregate costs between x, y and between x', y' . It is possible that a non-constant cost function belongs to both types, such as $|x + y + x' + y'|$. The two types of cost functions lead to very different mathematical structures, which will be explored in this paper.

Although the QOT framework is much more general than the two classes above, these two types cover many commonly encountered problems. We give a few examples here with $\mathfrak{X} = \mathfrak{Y} = \mathbb{R}$, with detailed definitions and discussions in Section 3. The QOT problem becomes a classic OT problem by choosing the type-XY cost function

$$c(x, y, x', y') = f(x, y) + g(x', y') \quad \text{for some bivariate functions } f, g;$$

the transport cost is the variance of $f(X, Y)$ by choosing the type-XY cost function

$$c(x, y, x', y') = \frac{1}{2}(f(x, y) - f(x', y'))^2 \quad \text{for some bivariate function } f;$$

the transport cost is the covariance of (X, Y) by choosing the cost function of both types (up to QOT-irrelevant terms)

$$c(x, y, x', y') = \frac{1}{2} \underbrace{(x - x')(y - y')}_{\text{type-XX}} = \frac{1}{2} \underbrace{(xy + x'y')}_{\text{type-XY}} - \frac{1}{2} \underbrace{(xy' + x'y)}_{\text{QOT-irrelevant}};$$

the transport cost is Kendall’s tau of (X, Y) by choosing the type-XX cost function

$$c(x, y, x', y') = \text{sgn}(x - x') \text{sgn}(y - y');$$

²As a side note, the abbreviation QOT was also used for *quadratically regularized optimal transport* in González-Sanz and Nutz [2024], Nutz [2025] and Wiesel and Xu [2024]; see Example 3.8 below.

³In (1.3) and (1.4), it should be clear that f maps either \mathfrak{X}^2 or $\mathfrak{X} \times \mathfrak{Y}$ to \mathbb{R} , g maps either \mathfrak{Y}^2 or $\mathfrak{X} \times \mathfrak{Y}$ to \mathbb{R} , and h maps \mathbb{R}^2 to \mathbb{R} .

the optimal transport cost is the p -th power of a Gromov–Wasserstein (GW) distance by choosing the type-XX cost function

$$c(x, y, x', y') = ||x - x'| - |y - y'|\|^p, \quad p \geq 1;$$

QOT also includes the quadratic regularization of classic optimal transport as a special case. Moreover, a specific Monge formulation of QOT includes the quadratic assignment problems (QAP) of [Koopmans and Beckmann \[1957\]](#) by choosing the type-XX cost function

$$c(x, y, x', y') = f(x, x')g(y, y') \quad \text{for some bivariate functions } f, g.$$

The Koopmans–Beckmann QAP solves

$$\min_{\sigma \in S_n} \sum_{i=1}^n \sum_{j=1}^n a_{ij} b_{\sigma_i, \sigma_j}, \tag{1.5}$$

where $\{a_{ij}\}_{1 \leq i, j \leq n}$, $\{b_{ij}\}_{1 \leq i, j \leq n}$ are given $n \times n$ matrices and S_n is the set of all permutations of $[n] = \{1, \dots, n\}$. This includes many prominent examples such as the traveling salesman problem. Our theory of QOT in the discrete case includes but is much more general than QAP. For instance, if the Monge assumption is relaxed, the optimizer may not even be a (deterministic) assignment; see [Example 1.1](#) below.

The non-linearity of QOT induces many difficulties and peculiarities. Since the problem is neither convex nor concave, duality is not generally available, and computational methods are also quite limited. In addition, explicit solutions are rare, while many peculiar examples exist due to the non-linearity. We illustrate a simple example below.

Example 1.1. Let $\mu = \nu$ be the two-point uniform distribution on $\{0, 1\}$, that is, Bernoulli(1/2), and consider the (type-XX) rectangular cost function $c(x, y, x', y') = |x - x'| |y - y'|$, which is equivalent to [\(1.1\)](#) up to QOT-irrelevant terms. Any $\pi \in \Pi(\mu, \nu)$ can be parameterized by $p = 2\pi(\{(1, 1)\}) \in [0, 1]$, and thus we may write $\Pi(\mu, \nu) = \{\pi_p\}_{p \in [0, 1]}$. By direct computation,

$$\iint c \, d\pi_p \otimes d\pi_p = 2\pi_p \otimes \pi_p(\{(0, 0, 1, 1), (0, 1, 1, 0)\}) = 2 \left(\frac{p^2}{4} + \frac{(1-p)^2}{4} \right) = p^2 - p + \frac{1}{2}.$$

Therefore, the transport cost $\iint c \, d\pi_p \otimes d\pi_p$ is a quadratic function in p that is uniquely minimized by $p = 1/2$. In other words, the independent coupling is the unique minimizer. On the contrary, it is well-known that for classic OT, if μ, ν are both uniformly distributed on the same number of points, a bijective optimizer exists, which follows from Birkhoff’s theorem ([Birkhoff \[1946\]](#)).

Given the richness of possible special cases and applications and the mathematical novelty of the new framework, we dedicate this paper to a systematic study of QOT. Our main contributions are summarized below.

In [Section 2](#), we formally present the framework of QOT on general Polish spaces. [Section 3](#) collects many relevant examples of QOT in optimization, economics, computational OT, and statistics.

[Section 4](#) provides general results on QOT, including properties of the optimizers and general lower bounds of the QOT cost, some of which extend known results on the GW distance. [Section 5](#) examines several explicit solvable cases. We show that the comonotone coupling π_{com} and the anti-monotone coupling π_{ant} , as well as their mixtures, form solutions to many classes of cost functions, including quadratic costs, jointly submodular/supermodular costs, and Gromov–Wasserstein-type costs.

Due to the fundamental differences from the classic OT setting, many new optimal transport plans emerge besides the comonotone and antimonotone ones. [Figure 1](#) illustrates the support of




Figure 1: Illustration of the support of the comonotone transport π_{com} , the antimonotone transport π_{ant} , the X-transport π_x^λ with $\lambda \in (0, 1)$, the V-transport π_v , and the diamond transport π_{dia} , which appear as QOT minimizers with marginals normalized to uniform distributions on $[0, 1]$; for their corresponding cost functions and marginals, see Table 1 below. Blue lines indicate the support of the transport plans. The transport plans in (c), (d), and (e) appear new compared to classic OT. Precise definitions are given in Section 2 and Definitions 5.9 and 6.1.

some QOT minimizers. The most interesting one is arguably the diamond transport π_{dia} , for the following reasons: first, it does not appear in other contexts of OT or QAP; second, it is perfectly symmetric but is not Monge; third, it serves as a universal minimizer of a large class of QOT problems with some assumptions on the marginals. A simple example with a diamond minimizer is given by the rectangular cost function $c(x, y, x', y') = |(x - x')(y - y')|$ in Example 1.1, equivalent to the one in (1.1). Section 6 focuses on the diamond transport.

Section 7 concludes with a few open questions. The appendices contain further discussions, additional results, and omitted proofs. Appendix A gives a quadratic programming formulation of QOT. Appendix B shows that the independent coupling is rarely, but possibly, an optimizer of the QOT, and gives some interesting examples. In Appendix C, we consider the class of linear-exponential distance cost functions of the form $c(x, y, x', y') = |y - y'|e^{-\gamma|x-x'|}$, $\gamma > 0$. This class of cost functions are minimized by the comonotone coupling, but its maximizers have interesting limiting behavior as γ goes to 0 or ∞ , such as the diamond and independent couplings, in some special senses. In particular, the limiting case $\gamma \rightarrow \infty$ is connected to recently studied measures of association (Chatterjee [2021], Deb et al. [2020]). Appendix D contains a more detailed discussion of the open questions from Section 7. Appendices E–G collect omitted proofs of the results from Sections 4–6.

Before moving on to the formal analysis, we summarize in Table 1 the QOT problems with known explicit optimizers obtained in this paper.

2 Framework

As in the Introduction, let \mathfrak{X} and \mathfrak{Y} be two Polish spaces, μ and ν be two probability measures on \mathfrak{X} and \mathfrak{Y} respectively, and $\Pi(\mu, \nu)$ be the set of all distributions on $\mathfrak{X} \times \mathfrak{Y}$ with marginals μ, ν . In many explicit results and examples, we will take $\mathfrak{X} = \mathfrak{Y} = \mathbb{R}$, but we also present some results on more general spaces. For a function $c : (\mathfrak{X} \times \mathfrak{Y})^2 \rightarrow \mathbb{R}$, called a cost function, and a coupling $\pi \in \Pi(\mu, \nu)$, define the *quadratic-form transport cost* as⁴

$$\iint c \, d\pi \otimes d\pi = \iint c(x, y, x', y') \, d\pi(x, y) \, d\pi(x', y') \quad (2.1)$$

and the (Kantorovich) QOT problem is to minimize (and occasionally, to maximize) this transport cost over all $\pi \in \Pi(\mu, \nu)$ such that the integral (2.1) is well-defined (taking possibly infinite values).

⁴Throughout, we tacitly assume suitable measurability of the cost function c so that (2.1) is meaningful.

Cost function c	Marginals μ, ν	Minimizer	Location
$f(x, y)g(x', y')$ f, g quadratic	general	$\pi_x^\lambda = \lambda\pi_{\text{com}} + (1 - \lambda)\pi_{\text{ant}}$ for some $\lambda \in [0, 1]$	Proposition 5.1
$(\alpha_1 f(x, y) + C_1)(\alpha_2 f(x', y') + C_2)$ f submodular, $C_1, C_2, \alpha_1, \alpha_2 \in \mathbb{R}$	general	$\pi_x^\lambda = \lambda\pi_{\text{com}} + (1 - \lambda)\pi_{\text{ant}}$ for some $\lambda \in [0, 1]$	Proposition 5.3
$(x', y') \mapsto c(x, y, x', y')$ and $(x, y) \mapsto c(x, y, x', y')$ both submodular	general	π_{com}	Theorem 5.4
$h(x - x' , y - y')$ h submodular	$\nu = \mu \circ \ell^{-1}$ ℓ linear map	π_{com} (ℓ increasing) π_{ant} (ℓ decreasing)	Theorem 5.6
$f(x - x')g(y, y')$ f nonnegative increasing g increasing supermodular and some regularity conditions	μ is uniform on an interval	π_ν	Theorem 5.10
$ (x - x')(y - y') $ or $(\theta_1 x - x' + \theta_2 y - y')^2$ with $\theta_1, \theta_2 \geq 0$	general	π_{dia}	Theorem 6.2
$\phi((x - x')^2)\phi((y - y')^2)$ ϕ completely monotone $\phi'(u) + 2u\phi''(u) \leq 0$	μ, ν symmetric	π_{dia}	Theorem 6.5
$ (x - x')(y - y') ^q, q \in (1, 2]$	μ, ν symmetric	π_{dia}	Theorem 6.9

Table 1: A selected list of explicitly solved examples of QOT problems on the real line. The minimizers may not be unique. Certain moment assumptions on the marginals are omitted (compactness of support is sufficient). For definitions of the couplings, see Section 2 and Definitions 5.9 and 6.1. In Examples 3.1–5.5, many more explicit cost functions, some of which belong to the above general classes, are presented. The conclusions remain the same if QOT-irrelevant terms like $w_1(x, x') + w_2(y, y') + w_3(x, y') + w_4(x', y)$ are added to the cost function c (Fact 2.1).

We omit “Kantorovich” in the sequel. The probabilistic formulation of the quadratic-form transport cost (2.1) is $\mathbb{E}[c(\mathbf{Z}, \mathbf{Z}')]]$, where $\mathbf{Z}, \mathbf{Z}' \stackrel{\text{law}}{\sim} \pi$ are iid. A quadratic program formulation of discrete QOT is presented in Appendix A.

Fact 2.1. The QOT problem remains equivalent (that is, with the same set of minimizers) if QOT-irrelevant terms are added to the cost function. For instance, the cost functions c and

$$(x, y, x', y') \mapsto c(x, y, x', y') + w_1(x, x') + w_2(y, y') + w_3(x, y') + w_4(x', y)$$

lead to equivalent QOT problems. All results in this paper automatically hold when QOT-irrelevant terms are added to the cost functions.

In certain applications, one may restrict to the Monge setting, where the set of couplings is induced by functions. Denote by $\mathcal{T}(\mu, \nu)$ the set of measurable maps $T : \mathfrak{X} \rightarrow \mathfrak{Y}$ satisfying $\mu \circ T^{-1} = \nu$, also known as the set of transport maps (or Monge maps) from μ to ν . The *Monge QOT* problem is to minimize

$$\iint c(x, T(x), x', T(x')) d\mu(x) d\mu(x'),$$

over the set $T \in \mathcal{T}(\mu, \nu)$.

The QOT problem can be realized as a variation of the multi-marginal OT problem under independence and marginal constraints. Consider the set $\Pi(\mu, \nu, \mu, \nu)$ of all probability measures on $(\mathfrak{X} \times \mathfrak{Y})^2$ with the four marginals given respectively by μ, ν, μ, ν . The *multi-marginal optimal transport* problem minimizes the transport cost

$$\int c(x, y, x', y') d\tilde{\pi}(x, y, x', y')$$

over $\tilde{\pi} \in \Pi(\mu, \nu, \mu, \nu)$; see [Pass \[2015\]](#), [Pass and Vargas-Jiménez \[2024\]](#) for surveys. Let $\Pi_{\text{ind}}(\mu, \nu, \mu, \nu)$ be the couplings (X, Y, X', Y') of μ, ν, μ, ν such that (X, Y) and (X', Y') are independent and have the same distribution. We then arrive at the equivalence

$$\inf_{\tilde{\pi} \in \Pi_{\text{ind}}(\mu, \nu, \mu, \nu)} \int c(x, y, x', y') d\tilde{\pi}(x, y, x', y') = \inf_{\pi \in \Pi(\mu, \nu)} \iint c(x, y, x', y') d\pi(x, y) d\pi(x', y').$$

In the rest of this section, we recall some fundamental results in classic OT and set up the necessary notation. Recall that a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ is called *submodular* if for any $x < x'$ and $y < y'$,

$$f(x, y) + f(x', y') \leq f(x, y') + f(x', y), \quad (2.2)$$

and f is called *supermodular* if for any $x < x'$ and $y < y'$,

$$f(x, y) + f(x', y') \geq f(x, y') + f(x', y). \quad (2.3)$$

In case the inequalities in (2.2) and (2.3) are strict, we say f is strictly submodular (or supermodular). Assuming $f \in C^2(\mathbb{R}^2)$, the cross partial derivative f_{xy} is nonnegative (resp. nonpositive) if and only if f is supermodular (resp. submodular).

Two classic couplings are fundamental to classic OT on $\mathfrak{X} \times \mathfrak{Y} = \mathbb{R}^2$, which we define below. Denote by $\mathcal{P}(\mathbb{R})$ the set of probability measures on \mathbb{R} . For a probability measure $\mu \in \mathcal{P}(\mathbb{R})$, let Q_μ be the left quantile function of μ , that is, $Q_\mu(t) = \inf\{x \in \mathbb{R} : \mu((-\infty, x]) \geq t\}$ for $t \in [0, 1]$ with $\inf \emptyset = \infty$. A coupling (X, Y) with marginals $\mu, \nu \in \mathcal{P}(\mathbb{R})$, or its joint distribution, is *comonotone* if $(X, Y) \stackrel{\text{law}}{=} (Q_\mu(U), Q_\nu(U))$, where U is uniformly distributed on $[0, 1]$; it is *antimonotone* if $(X, Y) \stackrel{\text{law}}{=} (Q_\mu(U), Q_\nu(1-U))$. It is well known that these two couplings either maximize or minimize classic OT problems when the cost function is submodular or supermodular (e.g., Theorem 2.9 of [Santambrogio \[2015\]](#)). We let $\pi_{\text{com}} \in \Pi(\mu, \nu)$ denote the comonotone coupling, $\pi_{\text{ant}} \in \Pi(\mu, \nu)$ denote the antimonotone coupling, and $\pi_{\text{ind}} = \mu \otimes \nu \in \Pi(\mu, \nu)$ denote the independent coupling. Couplings such as π_{com} , π_{ant} , and π_{ind} depend on the marginals μ, ν , which should be clear from context. In addition, for $\lambda \in [0, 1]$, let $\pi_x^\lambda = \lambda\pi_{\text{com}} + (1-\lambda)\pi_{\text{ant}} \in \Pi(\mu, \nu)$; the coupling π_x^λ for $\lambda \in (0, 1)$ is called an *X-transport* because its support has an X-shape; see [Figure 1](#), panel (c).

Some further notation and terminologies will be used throughout the paper. We say that a measure $\mu \in \mathcal{P}(\mathbb{R})$ is *symmetric* if there exists $m \in \mathbb{R}$ such that $\mu(A) = \mu(m-A)$ for all Borel sets $A \subseteq \mathbb{R}$. This means $X \stackrel{\text{law}}{\sim} \mu \iff m-X \stackrel{\text{law}}{\sim} \mu$. Otherwise, we say μ is *asymmetric*. For a probability measure μ on \mathbb{R}^d with $d \in \mathbb{N}$, we denote by F_μ the cdf of μ . For $p \geq 1$, we let $\mathcal{P}_p(\mathbb{R})$ denote the set of probability measures $\mu \in \mathcal{P}(\mathbb{R})$ with a finite p -th absolute moment, i.e., $\int_{\mathbb{R}} |x|^p d\mu(x) < \infty$. For $a < b$, we denote by $U(a, b)$ the uniform distribution on $[a, b]$. We also write $U = U(0, 1)$. Denote by \mathbb{R}_+ the set of nonnegative real numbers.

3 Examples of QOT

In this section, we discuss several examples of QOT that appear in different fields and how they connect to the results in this work.

Example 3.1 (Sum of bivariate functions). Suppose that the cost function c can be written as the sum of several bivariate functions, that is,

$$c(x, y, x', y') = f(x, y) + g(x', y') + w_1(x, x') + w_2(y, y') + w_3(x, y') + w_4(x', y).$$

By Fact 2.1, the QOT problem with the above cost is equivalent to the type-XY cost function $g(x, y) + h(x', y')$. In this case,

$$\int (f(x, y) + g(x', y')) d\pi(x, y)d\pi(x', y') = \int (f + g) d\pi.$$

In other words, the QOT problem reduces to a classic OT problem with cost function $f + g$. In particular, if $f + g$ is submodular, then the comonotone coupling is a minimizer. On the other hand, if the cost function c has a multiplicative form $c(x, y, x', y') = f(x, y)g(x', y')$, then $\iint c d\pi \otimes d\pi = (\int f d\pi)(\int g d\pi)$, which is the product of two transport costs in classic OT. If f and g are both submodular and nonnegative, a minimizer is the comonotone coupling.

Example 3.2 (Variance minimization with given marginals). Let $f : \mathfrak{X} \times \mathfrak{Y} \rightarrow \mathbb{R}$ be a measurable function. Suppose that the goal is to minimize the variance of $f(X, Y)$ subject to $X \stackrel{\text{law}}{\sim} \mu$ and $Y \stackrel{\text{law}}{\sim} \nu$. This problem is QOT with the type-XY nonnegative cost function c given by $c(x, y, x', y') = (f(x, y) - f(x', y'))^2/2$ because, for $(X, Y) \stackrel{\text{law}}{\sim} \pi$,

$$\begin{aligned} \iint c d\pi \otimes d\pi &= \begin{cases} \mathbb{E}[f(X, Y)^2] - \mathbb{E}[f(X, Y)]^2 & \text{if } \mathbb{E}[f(X, Y)^2] < \infty \\ \infty & \text{otherwise} \end{cases} \\ &= \text{Var}(f(X, Y)). \end{aligned} \tag{3.1}$$

This QOT problem is well-posed even when $f(X, Y)$ does not have a finite variance for some coupling π . The minimization of (3.1) is not a classic OT problem, because transport costs in classic OT are linear in π , whereas (3.1) is not.

Example 3.3 (Covariance). Assume that $\mu, \nu \in \mathcal{P}_2(\mathbb{R})$. Consider the QOT problem with the type-XX (and also type-XY, up to QOT-irrelevant terms) cost function c given by

$$c(x, y, x', y') = \frac{1}{2}(x - x')(y - y') = \frac{1}{2}(xy + x'y') - \frac{1}{2}(xy' + x'y).$$

For $(X, Y) \stackrel{\text{law}}{\sim} \pi$, we can verify

$$\begin{aligned} \iint c d\pi \otimes d\pi &= \frac{1}{2} (\mathbb{E}[XY] + \mathbb{E}[X'Y'] - \mathbb{E}[XY'] - \mathbb{E}[X'Y]) \\ &= \mathbb{E}[XY] - \mathbb{E}[X]\mathbb{E}[Y] = \text{Cov}(X, Y). \end{aligned}$$

Therefore, the transport cost is the covariance of (X, Y) . It is well-known that the unique minimizer of covariance is the antimonotone coupling and the unique maximizer is the comonotone coupling, which is also a consequence of Theorem 5.4.

Example 3.4 (Kendall's tau). Kendall's tau, also called Kendall's rank correlation coefficient, is one of the most popular measures of bivariate rank correlation, widely used in statistics and stochastic modeling; see e.g., Nelsen [2006, Chapter 5] and McNeil et al. [2015, Chapter 7]. For a random vector (X, Y) taking values in \mathbb{R}^2 , its Kendall's tau is defined as

$$\tau = \mathbb{E}[\text{sgn}((X - X')(Y - Y'))],$$

where (X', Y') is an independent copy of (X, Y) . Intuitively, it equals the probability of concordance minus that of discordance between (X, Y) and (X', Y') . Clearly, τ of $(X, Y) \stackrel{\text{law}}{\sim} \pi$ can be written as

the quadratic-form transport cost with the type-XX cost function $c(x, y, x', y') = \text{sgn}(x - x') \text{sgn}(y - y')$. For given marginals μ, ν , it is well-known that $\tau(\pi)$ over $\pi \in \Pi(\mu, \nu)$ is maximized by the comonotone coupling with maximum value 1 and minimized by the antimonotone coupling with minimum value -1 (this can also be checked by Theorem 5.4; see Example 5.5). Another equivalent formulation is $\tau(\pi) = 4 \int \pi d\pi - 1$, from which the quadratic form in π is visible.

Example 3.5 (Gini deviation and Gini coefficient). Let L^1 be the set of integrable random variables and $L^1_+ = \{Z \in L^1 : Z \geq 0; \mathbb{E}[Z] > 0\}$. Define the mappings GD and GC on L^1_+ by

$$\text{GD}(Z) = \frac{1}{2} \mathbb{E}[|Z - Z'|] \quad \text{and} \quad \text{GC}(Z) = \frac{\text{GD}(Z)}{\mathbb{E}[Z]} = \frac{\mathbb{E}[|Z - Z'|]}{\mathbb{E}[Z + Z']},$$

where Z' is an independent copy of Z . The value $\text{GD}(Z)$ is called the Gini deviation of Z , and $\text{GC}(Z)$ is called the Gini coefficient of Z , both of which are commonly used as measures of distributional variability or inequality in economics and risk management; see e.g., Gastwirth [1971] and Furman et al. [2017]. Similarly to the variance in (3.1), minimization of $\text{GD}(f(X, Y))$ for some measurable $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ over $X \stackrel{\text{law}}{\sim} \mu$ and $Y \stackrel{\text{law}}{\sim} \nu$ can be written as the QOT problem with the type-XY cost function $|f(x, y) - f(x', y')|$. Moreover, the minimization of $\text{GC}(X + Y)$ can be written as

$$\min_{\pi \in \Pi(\mu, \nu)} \frac{\iint |x + y - x' - y'| d\pi(x, y) d\pi(x', y')}{2(\int x d\mu(x) + \int y d\nu(y))}, \quad (3.2)$$

which is equivalent to a QOT problem with cost function $c(x, y, x', y') = |x + y - x' - y'|$, noting that the denominator of (3.2) does not involve π . A minimizer of (3.2) is given by the antimonotone coupling π_{ant} ; see Example 5.5(v) below. This cost function is both type-XX and type-XY.

Example 3.6 (Gromov–Wasserstein (GW) distance). A special case of QOT is the GW distance, a measure of the distance (or similarity) between two metric measure spaces introduced and studied by Mémoli [2007, 2011a]. Suppose that $(\mathfrak{X}, d_{\mathfrak{X}}, \mu)$ and $(\mathfrak{Y}, d_{\mathfrak{Y}}, \nu)$ are metric measure spaces. For $p, q \geq 1$, the (p, q) -GW distance is defined as

$$\text{GW}_{p,q}(\mathfrak{X}, \mathfrak{Y}) := \left(\inf_{\pi \in \Pi(\mu, \nu)} \iint |d_{\mathfrak{X}}(x, x')^q - d_{\mathfrak{Y}}(y, y')^q|^p d\pi(x, y) d\pi(x', y') \right)^{1/p}. \quad (3.3)$$

This is an increasing transform of the minimum transport cost of QOT with a type-XX cost function. The GW distance satisfies the triangle inequality and defines a pseudo-metric on metric measure spaces (and a metric on isomorphism classes of metric measure spaces; see Sturm [2023, Corollary 9.3]). The GW distance is a widely used technique in data science, machine learning, computer vision, and computer graphics to align heterogeneous data sets or images (Mémoli [2011a], Peyré and Cuturi [2019]). However, in general, solving for the GW distance is a challenging task. In Theorems 6.2 and 6.9, we explicitly characterize the *maximizers* of the transport cost that appears in (3.3) for $p = 2, q \in [1, 2], \mathfrak{X} = \mathfrak{Y} = \mathbb{R}$ equipped with the Euclidean distance, and symmetric marginals, where the symmetry is not needed for $q = 1$. The GW literature also incorporates some earlier ideas discussed in this paper, such as existence of minimizers and Monge–Kantorovich equivalence (Proposition 4.4), lower bounds on the transport cost (Propositions 4.6 and 4.7), and connections to QAP. These ideas also manifest in studies of extensions of the GW distance (Arya et al. [2024], Bauer et al. [2024], Chowdhury and Mémoli [2019], Mémoli [2011b, 2012], Mémoli et al. [2023]), sometimes leading to closed-form solutions.

Example 3.7 (Quadratic Assignment Problem (QAP)). If the probability measures μ, ν are each uniformly distributed on N points, the Monge QOT problem reduces to QAP, which was first introduced by Koopmans and Beckmann [1957] as a model for the allocation problem of indivisible economic activities. This connection is well known in the GW literature; see, for instance, Remark

4.6 of [Mémoli \[2011a\]](#). Recall the Koopmans–Beckmann problem from [\(1.5\)](#). The work of [Lawler \[1963\]](#) proposed a generalization of the Koopmans–Beckmann QAP, where one solves

$$\min_{\sigma \in \mathcal{S}_n} \sum_{i=1}^n \sum_{j=1}^n d_{ij\sigma_i\sigma_j},$$

where $D = \{d_{ijkl}\}_{1 \leq i,j,k,\ell \leq n}$ is a given 4-index cost array. Since any Monge transport map between μ and ν matches the N elements bijectively, the Lawler QAP coincides with Monge QOT with discrete uniform marginals, and the Koopmans–Beckmann QAP further constrains that the cost function is of the form $c(x, y, x', y') = c_1(x, x')c_2(y, y')$, thus type-XX. The QAP class includes many well-known combinatorial optimization problems, such as the traveling salesman problem (Section 7.1.2 of [Burkard et al. \[1998\]](#)) and the campus planning model ([Dickey and Hopkins \[1972\]](#)). However, in general, even approximating QAP is NP-hard ([Loiola et al. \[2007\]](#), [Queyranne \[1986\]](#)). For instance, QAP of size $n > 30$ cannot be solved in a reasonable amount of time. We refer to [Burkard et al. \[2012\]](#) and [Çela \[2013\]](#) for comprehensive surveys on QAP and various extensions of the problem.

Example 3.8 (Quadratic regularization of discrete optimal transport). Regularization is a modern technique in OT that facilitates computation by introducing strong convexity to the linear OT problem. The most prevalent choice is arguably the entropic regularized OT (EOT), which enables the Sinkhorn’s algorithm and generates smoothness properties of the solution ([Nutz \[2021\]](#)). An alternate of EOT is given by the following quadratically regularized OT:

$$\begin{aligned} \text{to minimize} \quad & \int c(x, y) d\pi(x, y) + \frac{\varepsilon}{2} \int \left(\frac{d\pi}{d\pi_{\text{ind}}}(x, y) \right)^2 d\pi_{\text{ind}}(x, y) \\ \text{subject to} \quad & \pi \in \Pi(\mu, \nu), \end{aligned} \tag{3.4}$$

where by convention, the objective value is ∞ if $\pi \not\ll \pi_{\text{ind}}$. The quadratically regularized OT was introduced by [Blondel et al. \[2018\]](#) and [Essid and Solomon \[2018\]](#) in the discrete setting, and rigorously studied by [Lorenz et al. \[2021\]](#) in the continuous case. The authors highlighted that quadratically regularized OT gives rise to sparse couplings, a desirable property when the OT itself is of interest. Another advantage of quadratically regularized OT over EOT is the allowance of small regularization parameters, as the computation for EOT is difficult for a small ε ([Nutz \[2025\]](#)). In the discrete setting, suppose that μ has mass $\{p_i\}$ on points $\{x_i\}$ and ν has mass $\{q_j\}$ on points $\{y_j\}$. Denote also by π_{ij} the mass of π on (x_i, y_j) . The regularization term of [\(3.4\)](#) can be written as

$$\begin{aligned} \int \left(\frac{d\pi}{d\pi_{\text{ind}}}(x, y) \right)^2 d\pi_{\text{ind}}(x, y) &= \sum_i \sum_j \frac{\pi_{ij}^2}{p_i q_j} = \sum_{i,j} \sum_{k,\ell} \frac{1}{p_i q_j} \mathbb{1}_{\{i=k, j=\ell\}} \pi_{ij} \pi_{k\ell} \\ &= \iint \frac{\mathbb{1}_{\{x=x', y=y'\}}}{\mu(\{x\})\nu(\{y\})} d\pi(x, y) d\pi(x', y'), \end{aligned}$$

which is a transport cost in QOT (type-XX) with the unique minimizer given by π_{ind} , a consequence of Jensen’s inequality. Since the classic OT is also a special case of QOT ([Example 3.1](#)), the quadratically regularized OT in the discrete case belongs to the QOT class.

4 General properties

We now proceed to study general properties of QOT. The results presented in this section hold for probability measures on general Polish spaces, except for the stability of QOT ([Proposition 4.5](#)), where we require $\mathfrak{X} = \mathfrak{Y} = \mathbb{R}$.

Convexity is an essential issue in classic OT theory, giving rise to many useful techniques such as duality and c -cyclical monotonicity. We start with a simple result that gives a sufficient condition

for QOT to be convex. For a non-empty set S , we say that a symmetric kernel $\phi : S \times S \rightarrow \mathbb{R}$ is a *positive definite kernel* if for any $n \in \mathbb{N}$, $s_1, \dots, s_n \in S$, and $c_1, \dots, c_n \in \mathbb{R}$, it holds

$$\sum_{i=1}^n \sum_{j=1}^n c_i c_j \phi(s_i, s_j) \geq 0,$$

which is a generalization of positive semi-definite matrices. A closely related concept is that of *positive definite functions*, that is, functions $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ satisfying that $(x, y) \mapsto \varphi(x - y)$ is a positive definite kernel on \mathbb{R}^2 .

Proposition 4.1. *If the cost function c satisfies $c(x, y, x', y') = \phi((x, y), (x', y'))$ for some $\phi : (\mathfrak{X} \times \mathfrak{Y})^2 \rightarrow \mathbb{R}$ that is bounded, continuous, and positive definite, then QOT is a convex optimization problem. In other words, the quadratic-form transport cost (2.1) is a convex function of π .*

Example 4.2. Denote by $\|\cdot\|$ the Euclidean norm. For $\mathbf{z}, \mathbf{z}' \in \mathbb{R}^d$, the kernels $(\mathbf{z}, \mathbf{z}') \mapsto e^{-\alpha\|\mathbf{z}-\mathbf{z}'\|}$ and $(\mathbf{z}, \mathbf{z}') \mapsto e^{-\alpha\|\mathbf{z}-\mathbf{z}'\|^2}$ are both positive definite for $\alpha > 0$. Therefore, the QOT problem with the type-XX cost function $e^{-\alpha\|(x,y)-(x',y')\|}$ or $e^{-\alpha\|(x,y)-(x',y')\|^2}$ is convex. We will see more examples in Example 6.7, where we also show that QOT with cost function $e^{-\alpha\|(x,y)-(x',y')\|^2}$, $\alpha \in (0, 1/2]$ is minimized by the diamond transport.

Despite the above result, the majority of QOT problems are not convex in π . This non-convex structure of QOT prohibits the use of classic tools such as duality.⁵ In the rest of this section, we study the fundamental properties of QOT, which may not be convex. Specifically, we show that under certain assumptions, minimizers exist, the Monge optimal transport cost is equivalent to the Kantorovich one, QOT on \mathbb{R} is stable, and the independent coupling is rarely an optimizer.

To discuss the finiteness of the transport cost in QOT, denote by $\mathcal{C}(\mu, \nu)$ the set

$$\mathcal{C}(\mu, \nu) = \left\{ c : (\mathfrak{X} \times \mathfrak{Y})^2 \rightarrow \mathbb{R} \mid \begin{array}{l} c(x, y, x', y') \geq f(x, x') + g(y, y') \text{ everywhere} \\ \text{for some } f \in L^1(\mu \otimes \mu) \text{ and } g \in L^1(\nu \otimes \nu) \end{array} \right\}.$$

Note that a cost function c is in $\mathcal{C}(\mu, \nu)$ if it is bounded from below, or if it is lower semi-continuous and μ and ν are compactly supported. The next remark is immediate.

Fact 4.3. If $c \in \mathcal{C}(\mu, \nu)$, then the infimum of the quadratic-form transport cost in (2.1) is well-defined and not $-\infty$.

The next result gives conditions under which minimizers of QOT exist and under which Monge is equivalent to Kantorovich. Similar results have been established for special cases such as the GW distance (Mémoli [2011a, Corollary 10.1] and Mémoli and Needham [2024, Theorem 3.2]; see Example 3.6 for the formulation) and its extensions (e.g., Bauer et al. [2024, Theorem 2]). We also refer to Section 3 of Mémoli and Needham [2024] for further results that compare the Monge and Kantorovich problems in the case of GW costs.

Proposition 4.4. *Suppose that $c \in \mathcal{C}(\mu, \nu)$ is lower semi-continuous. Then a minimizer of (2.1) exists. In particular, if μ is atomless, c is continuous, and $\mathfrak{X}, \mathfrak{Y}$ are compact, then*

$$\min_{\pi \in \Pi(\mu, \nu)} \iint c \, d\pi \otimes d\pi = \inf_{T \in \mathcal{T}(\mu, \nu)} \iint c(x, T(x), x', T(x')) \, d\mu(x) \, d\mu(x').$$

The minimizer in Proposition 4.4 may not be unique even in many non-trivial cases, which we will see later.

We next show that, similar to classic OT, QOT satisfies stability with respect to the marginals.

⁵Notable exceptions include Vayer [2020, Theorem 4.2.5] and Zhang et al. [2024, Theorem 1], and the latter result also analyzes sample complexity of the (2, 2)-GW distance; see Example 3.6.

Proposition 4.5. *Suppose that $\mu, \nu \in \mathcal{P}(\mathbb{R})$ and $\mu_n \rightarrow \mu$, $\nu_n \rightarrow \nu$ weakly. Let $c : \mathbb{R}^4 \rightarrow \mathbb{R}$ be a continuous function satisfying the uniform integrability condition*

$$\sup_{\pi \in \Pi(\mu, \nu)} \iint |c(x, y, x', y')|^{1+\delta} d\pi(x, y) d\pi(x', y') < \infty. \quad (4.1)$$

for some $\delta > 0$. Let $\pi_n \in \Pi(\mu_n, \nu_n)$ be any QOT minimizer with cost function c . Then the sequence $\{\pi_n\}_{n \in \mathbb{N}}$ admits weak limit points in $\Pi(\mu, \nu)$ and every weak limit point of $\{\pi_n\}_{n \in \mathbb{N}}$ is a QOT minimizer with cost function c and marginals μ, ν .

Stability is crucial in classic OT theory as it ensures that numerical algorithms for solving the OT problem converge consistently. On the other hand, numerically computing or approximating the QOT solution remains a difficult task, as the discretized version remains an NP-hard problem (Loiola et al. [2007]), a fact already noted in Mémoli [2011a, Remark 4.6]. The recent work of Kravtsova [2024] also indicates the NP-hardness of the GW distance. As a discrete Monge version of QOT, QAP provides feasible heuristic algorithms; see Burkard et al. [2012] and Çela [2013]. However, one needs to be careful here since discrete QOT may not always have a Monge minimizer even if a bijective transport map exists (see Example 1.1). We leave the computational aspects of QOT for further investigation.

Lower bounds on transport costs in QOT can be easily obtained from classic OT by either a two-step optimization or an optimization over aggregated marginals, which we summarize in the next two simple results. Let $\mathcal{C}_c(\mu, \nu)$ denote the classic optimal transport cost from μ to ν with cost function c , i.e.,

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

Proposition 4.6. *Suppose that $c \in \mathcal{C}(\mu, \nu)$. It holds that*

$$\inf_{\pi \in \Pi(\mu, \nu)} \iint c d\pi \otimes d\pi \geq \mathcal{C}_{\hat{c}}(\mu, \nu),$$

where $\hat{c}(x, y) = \mathcal{C}_{c_{x,y}}(\mu, \nu)$ and $c_{x,y}(x', y') = c(x, y, x', y')$. Moreover, if there exists $\pi_* \in \Pi(\mu, \nu)$ such that π_* is the optimal coupling for both cost functions \hat{c} and $c_{x,y}$ for π_* -a.e. (x, y) , then π_* is a minimizer of the QOT problem with cost function c .

Proposition 4.7. *Suppose that the cost function c is type-XX, i.e., of the form*

$$c(x, y, x', y') = h(f(x, x'), g(y, y'))$$

for some $h : \mathbb{R}^2 \rightarrow \mathbb{R}$, $f : \mathfrak{X}^2 \rightarrow \mathbb{R}$, and $g : \mathfrak{Y}^2 \rightarrow \mathbb{R}$. It holds that

$$\inf_{\pi \in \Pi(\mu, \nu)} \iint c d\pi \otimes d\pi \geq \inf_{\hat{\pi} \in \Pi(\mu_f, \nu_g)} \int h(\xi, \zeta) d\hat{\pi}(\xi, \zeta) = \mathcal{C}_h(\mu_f, \nu_g),$$

where μ_f is the law of $f(X, X')$ for X, X' independent following law μ , and ν_g is the law of $g(Y, Y')$ for Y, Y' independent following law ν .

The lower bounds in Propositions 4.6 and 4.7 are generally not sharp, but they are useful in proving the optimality of some transport plans. The essential ideas of these lower bounds are present in the literature on QAP and GW distances. More precisely, Proposition 4.6 generalizes Lawler’s lower bound on QAP (Burkard et al. [2012]) and is known as the “third lower bound” in the GW setting (Mémoli [2007, 2011a]); Proposition 4.7 is an extension of the “second lower bound” for the GW distance. Examples 3.6 and 3.7 detail the connections between GW distance and QAP to QOT.

5 Explicit solutions between measures on the real line

In this section, we discuss a few instances where the QOT problem for $\mu, \nu \in \mathcal{P}(\mathbb{R})$ allows for an explicit solution. Most of our results will be built on the lower bounds obtained in Section 4. More precisely, the strategy is to show that certain lower bounds are achieved by specific transport plans (such as the comonotone coupling). Further results where the minimizer is attained by the diamond transport will be discussed in Section 6.

5.1 Type-XY product of two quadratic cost functions

A natural class of cost functions to consider in OT theory is the quadratic ones. For instance, martingale optimal transport with a quadratic cost function is trivial. In the QOT framework, we consider a cost function $c(x, y, x', y')$ that is a quadratic function of four variables. After adding terms that do not depend on the coupling, any such cost function c is equivalent to one of the form $c(x, y, x', y') = f(x, y)g(x', y')$, where f, g are quadratic functions of the two variables.

We describe an algorithm that explicitly solves QOT problems whose cost function is of the form $c(x, y, x', y') = f(x, y)g(x', y')$, where f, g are quadratic. This cost function is type-XY. Vayer [2020, Theorem 4.2.4] studies a special case $c(x, y, x', y') = xyx'y'$. Recall the notation $\pi_x^\lambda = \lambda\pi_{\text{com}} + (1 - \lambda)\pi_{\text{ant}}$ for $\lambda \in [0, 1]$, which is called an X-transport if $\lambda \in (0, 1)$.

Proposition 5.1. *Suppose that $\mu, \nu \in \mathcal{P}(\mathbb{R})$. If the cost function c is given by*

$$c(x, y, x', y') = f(x, y)g(x', y'), \quad \text{where } f, g \text{ are quadratic functions,} \quad (5.1)$$

then there exists a QOT minimizer π_x^λ for some $\lambda \in [0, 1]$. Moreover, if μ, ν are not degenerate, then every π_x^λ minimizes the quadratic-form transport cost for some cost function in (5.1) uniquely among the class $(\pi_x^\lambda)_{\lambda \in [0, 1]}$.

The proof of Proposition 5.1 contains an algorithm that explicitly solves such a QOT problem. We illustrate it with the following example.

Example 5.2. Consider μ, ν both distributed as $N(0, 1)$ with the cost function given by $c(x, y, x', y') = -(x + y)^2(2x' - y')^2$. A standard computation yields that for $\pi \in \Pi(\mu, \nu)$,

$$\begin{aligned} \mathbb{E}_{\pi \otimes \pi}[c(X, Y, X', Y')] &= -\mathbb{E}[(X + Y)^2]\mathbb{E}[(2X' - Y')^2] \\ &= -(2 + 2\mathbb{E}[XY])(5 - 4\mathbb{E}[X'Y']) = 8\mathbb{E}[XY]^2 - 2\mathbb{E}[XY] - 10. \end{aligned}$$

Since the quadratic function $z \mapsto 8z^2 - 2z - 10$ is minimized at $z = 1/8$, we see that if $\text{Cov}(X, Y) = 1/8$, the law π of (X, Y) is a QOT minimizer. This is achieved, for example, by $\pi = (9/16)\pi_{\text{com}} + (7/16)\pi_{\text{ant}}$, where $(X, X) \stackrel{\text{law}}{\sim} \pi_{\text{com}}$ and $(X, -X) \stackrel{\text{law}}{\sim} \pi_{\text{ant}}$. On the other hand, since the range of $\mathbb{E}[XY]$ is $[-1, 1]$, the unique QOT maximizer is given by the antimonotone coupling $X = -Y$, where $\mathbb{E}[XY] = -1$.

Proposition 5.1 follows from the following more general result, which replaces the quadratic functions f, g in Proposition 5.1 by submodular functions f, g that are identical up to a constant term.

Proposition 5.3. *Suppose that $\mu, \nu \in \mathcal{P}(\mathbb{R})$ and f is submodular. If the cost function c is given by*

$$c(x, y, x', y') = (\alpha_1 f(x, y) + C_1)(\alpha_2 f(x', y') + C_2), \quad \text{where } C_1, C_2, \alpha_1, \alpha_2 \in \mathbb{R}, \quad (5.2)$$

then there exists a QOT minimizer π_x^λ for some $\lambda \in [0, 1]$. Moreover, if μ, ν are not degenerate, then every π_x^λ minimizes the quadratic-form transport cost for some cost function in (5.2) uniquely among the class $(\pi_x^\lambda)_{\lambda \in [0, 1]}$.

5.2 Jointly submodular cost functions

Similarly to the classic OT problems, submodular and supermodular cost functions lead to explicit optimizers of QOT, which are the comonotone and antimonotone couplings (π_{com} and π_{ant}), as we present in the next result.

Theorem 5.4. *Suppose that $\mu, \nu \in \mathcal{P}(\mathbb{R})$ and the cost function $c \in \mathcal{C}(\mu, \nu)$ satisfies that both $(x', y') \mapsto c(x, y, x', y')$ and $(x, y) \mapsto c(x, y, x', y')$ are submodular (resp. supermodular) for every $x, y, x', y' \in \mathbb{R}$. Then the comonotone (resp. antimonotone) coupling is a minimizer.*

Example 5.5. We give several examples in which the conditions in Theorem 5.4 are satisfied, including both type-XX and type-XY ones.

- (i) If $c(x, y, x', y') = c_1(x, y)c_2(x', y')$ where both c_1, c_2 are nonnegative and submodular, the submodularity condition in Theorem 5.4 is clearly satisfied. In fact, a direct proof that the comonotone coupling is an optimizer follows from $\iint c \, d\pi \otimes d\pi = \int c_1 \, d\pi \int c_2 \, d\pi$.
- (ii) Generalizing (i), suppose that $c(x, y, x', y') = h(c_1(x, y), c_2(x', y'))$ and $c \in \mathcal{C}(\mu, \nu)$, where both c_1, c_2 are submodular and componentwise increasing, and $h : \mathbb{R}^2 \rightarrow \mathbb{R}$ is componentwise increasing and concave. We can check that the submodularity condition in Theorem 5.4 is satisfied. Hence, the comonotone coupling is a minimizer. Similarly, the antimonotone coupling is a maximizer if $-c \in \mathcal{C}(\mu, \nu)$. This includes, for instance, $c(x, y, x', y') = \min\{c_1(x, y), c_2(x', y')\}$, and $c(x, y, x', y') = (c_1(x, y) + c_2(x', y'))^p$ where $p \in (0, 1)$ and c_1, c_2 nonnegative.
- (iii) Suppose that $c(x, y, x', y') = h(|x - y|, |x' - y'|)$ and $c \in \mathcal{C}(\mu, \nu)$, where h is componentwise increasing and convex. It is elementary to check that the function $(x, y) \mapsto h(|x - y|, a)$ for $a \in \mathbb{R}$ is submodular. By Theorem 5.4, the comonotone coupling is a minimizer. This includes, for instance, $c(x, y, x', y') = \max\{|x - y|, |x' - y'|\}$. Note that for $(X, Y) \stackrel{\text{law}}{\sim} \pi$,

$$\iint \max\{|x - y|, |x' - y'|\} \, d\pi(x, y) \, d\pi(x', y') = \int_0^\infty \left(1 - (\mathbb{P}(|X - Y| \leq x))^2\right) \, dx,$$

which is the transport cost in a distorted OT problem (Liu et al. [2025])⁶ with distortion function $\eta : t \mapsto 1 - (1 - t)^2$ and cost function $\tilde{c} : (x, y) \mapsto |x - y|$.

- (iv) Consider $c(x, y, x', y') = c_1(x, x')c_2(y, y')$, where $c_1 \geq 0$ is increasing in both arguments and $c_2 \geq 0$ is decreasing in both arguments. Observe that a function $f(x, y) = a(x)b(y)$ is submodular if a is increasing positive and b is decreasing positive. Therefore, the function $(x', y') \mapsto c(x, y, x', y')$ is submodular by our assumption, and the same holds for $(x, y) \mapsto c(x, y, x', y')$.
- (v) The function $|x + y - x' - y'|$ is related to the Gini coefficient (see Example 3.5) and satisfies the supermodularity condition in Theorem 5.4, since $(x, y) \mapsto |x + y + c|$ is supermodular for every $c \in \mathbb{R}$. More generally, the cost function $|x + y - x' - y'|^p$ (resp. $|x - y + x' - y'|^p$) for $p \geq 1$ induces the antimonotone (resp. comonotone) coupling as an optimizer. Similarly, the antimonotone coupling is an optimizer of QOT with cost function $|x + y + x' + y'|^p$ for $p \geq 1$.
- (vi) The cost function $c(x, y, x', y') = \text{sgn}(x - x') \text{sgn}(y - y')$ defines Kendall's tau; see Example 3.4. It is elementary to verify that $(x', y') \mapsto c(x, y, x', y')$ and $(x, y) \mapsto c(x, y, x', y')$ are supermodular, and hence the transport cost is maximized by the comonotone coupling and minimized by the antimonotone coupling.
- (vii) Let $c(x, y, x', y') = \min\{x - x', y - y'\}$. Using that $(x, y) \mapsto \min\{x, y\}$ is supermodular, we see that the cost function c satisfies the supermodularity condition in Theorem 5.4, and hence a minimizer is given by the comonotone coupling.

⁶For a nonnegative cost function $\tilde{c} : \mathfrak{X} \times \mathfrak{Y} \rightarrow \mathbb{R}$ and a distortion function $\eta : [0, 1] \rightarrow [0, 1]$ increasing with $\eta(0) = 0$ and $\eta(1) = 1$, the distorted OT problem has transport cost formulated by $\int_0^\infty \eta(\mathbb{P}(\tilde{c}(X, Y) > x)) \, dx$, with the classic OT corresponding to $\eta(t) = t$ on $[0, 1]$; see Liu et al. [2025].

5.3 Gromov–Wasserstein-type cost functions

We now consider a family of type-XX cost functions, called the GW-type cost functions (see Example 3.6 for the GW distance). In what follows, we say that ν is an increasing (resp. decreasing) location-scale transform of μ if $\nu = \mu \circ \ell^{-1}$ for some strictly increasing (resp. decreasing) linear map $\ell : \mathbb{R} \rightarrow \mathbb{R}$; that is, $\ell(x) = ax + b$ for some $a > 0$ (resp. $a < 0$) and $b \in \mathbb{R}$.

Theorem 5.6. *Suppose that $\mu \in \mathcal{P}(\mathbb{R})$, ν is an increasing (resp. decreasing) location-scale transform of μ , and $h : \mathbb{R}_+^2 \rightarrow \mathbb{R}$ is a submodular function. Then the comonotone (resp. antimonotone) coupling is a minimizer of the QOT with the cost function c given by*

$$c(x, y, x', y') = h(|x - x'|, |y - y'|).$$

Moreover, such a minimizer is unique if it yields a finite transport cost, h is strictly submodular, and μ is asymmetric. In the same setting except that μ is symmetric, the comonotone and antimonotone couplings are the only minimizers.

In particular, an optimal coupling in Theorem 5.6 is precisely given by the Monge map $x \mapsto \ell(x)$, that is, the linear transform connecting μ and ν .

Example 5.7. We give a few examples of QOT problems that satisfy the conditions in Theorem 5.6. Assume $\mu = \nu \in \mathcal{P}(\mathbb{R})$ in all items.

- (i) Let $h(u, v) = \max\{u, v\}$, which is submodular. Theorem 5.6 then implies that π_{com} is a minimizer for the cost function $c(x, y, x', y') = \max\{|x - x'|, |y - y'|\}$.
- (ii) Let $h(u, v) = -u^{-\alpha}v^{-\alpha}$ for $\alpha \in (0, 1/2)$. If μ has a uniformly bounded density,

$$0 \leq -\mathbb{E}_{\pi_{\text{com}} \otimes \pi_{\text{com}}}[h(|X - X'|, |Y - Y'|)] = \mathbb{E}[|X - X'|^{-2\alpha}] < \infty.$$

In this case, Theorem 5.6 implies that π_{com} (and π_{ant} if μ is symmetric) is the unique minimizer for the cost function $c(x, y, x', y') = -|x - x'|^{-\alpha}|y - y'|^{-\alpha}$.

- (iii) Let $h(u, v) = -u^\beta v^\beta$ where $\beta > 0$. It is easy to verify that h is submodular. If $\mu \in \mathcal{P}_{2\beta}(\mathbb{R})$,

$$0 \leq -\mathbb{E}_{\pi_{\text{com}} \otimes \pi_{\text{com}}}[h(|X - X'|, |Y - Y'|)] = \mathbb{E}[|X - X'|^{2\beta}] < \infty.$$

It follows from Theorem 5.6 that π_{com} (and π_{ant} if μ is symmetric) is the unique minimizer for the cost function $c(x, y, x', y') = -|x - x'|^\beta |y - y'|^\beta$. The same cost function is also investigated in Beinert et al. [2023] in the case $\mu \neq \nu$, where it is shown that the Monge minimizer may be far away from the comonotone coupling.

- (iv) The choice $h(u, v) = |u^q - v^q|^p$ corresponds to the (p, q) -GW transport cost, see (3.3). One can verify that h is submodular if $p \geq 1, q > 0$ on $[0, \infty)^2$. Theorem 5.6 then implies that in the case $\mu = \nu$, π_{com} (and π_{ant} if μ is symmetric) is a minimizer for (3.3). This aligns with the intuition that the GW distance measures distances between metric measure spaces.

Remark 5.8. Theorem 5.6 extends naturally to more general Polish spaces, where $\mathfrak{X} = \mathfrak{Y}$, ν is a lateral shift of μ , and $c(x, y, x', y') = h(d(x, x'), d(y, y'))$. In this case, the lateral shift is always a Monge minimizer, but the uniqueness of the minimizer may depend on the geometry of the Polish space and the measures μ, ν .

5.4 A special class of separable cost functions and the V-transport

In this section, we prove that the V-transport mentioned in the Introduction serves as a minimizer for a special class of separable type-XX cost functions. We first rigorously define the V-transport. Recall that Q_μ is the left quantile function of $\mu \in \mathcal{P}(\mathbb{R})$.

Definition 5.9. A coupling (X, Y) with marginals μ and ν , or its joint distribution, is the V -transport if $(X, Y) \stackrel{\text{law}}{=} (Q_\mu(U), Q_\nu(|2U - 1|))$, where $U \stackrel{\text{law}}{\sim} U$. In this case, we denote the law of (X, Y) by π_ν .

For instance, if $\mu, \nu \stackrel{\text{law}}{\sim} U$, then π_ν is the distribution of $(U, |2U - 1|)$ where $U \stackrel{\text{law}}{\sim} U$; see Figure 1(b). If ν is atomless with median m_ν , then π_ν is the arithmetic average of the antimonotone coupling of μ and $2\nu|_{(-\infty, m_\nu]}$, and the comonotone coupling of μ and $2\nu|_{(m_\nu, \infty)}$.

Theorem 5.10. Suppose $\mu = U(a, b)$ for some $a < b$, $\nu \in \mathcal{P}(\mathbb{R})$, and the cost function c has the form

$$c(x, y, x', y') = f(|x - x'|)g(y, y'), \quad (5.3)$$

where g is right-continuous, increasing in both arguments, supermodular, and satisfies

$$\lim_{y \rightarrow -\infty} g(y, y') = \lim_{y' \rightarrow -\infty} g(y, y') = 0,$$

and h is nonnegative, right-continuous, and increasing. Then the V -transport is a minimizer.

Theorem 5.10 requires that μ is uniformly distributed on a compact interval. For a general atomless μ , we can transform the marginal by using $f(|F_\mu(x) - F_\mu(x')|)$ instead of $f(|x - x'|)$ in (5.3), and the same result applies.

The conditions on g hold, for instance, if $g(y, y') = \phi(y)\psi(y')$ where both ϕ, ψ are increasing, right-continuous, and satisfy $\lim_{y \rightarrow -\infty} \phi(y) = \lim_{y \rightarrow -\infty} \psi(y) = 0$.

The QAP version of Theorem 5.10 is contained in Burkard et al. [1998]. However, the Monge assumptions of QAP cannot be relaxed in general, and hence we cannot directly apply stability (Proposition 4.5) to solve the corresponding QOT.

6 The diamond transport

In this section, we systematically study a new transport, the diamond transport, which turns out to be a minimizer for several classes of type-XX QOT cost functions. Its definition is presented below.

Definition 6.1. Let $D = \{(x, y) \in [0, 1]^2 : |y - 1/2| + |x - 1/2| = 1/2\}$. The *diamond copula* C_{dia} is the cdf of the uniform distribution on D . The *diamond transport* $\pi_{\text{dia}} \in \Pi(\mu, \nu)$ between μ, ν is the law of $(Q_\mu(U), Q_\nu(V))$ where $(U, V) \stackrel{\text{law}}{\sim} C_{\text{dia}}$.

In terms of cdf, $\pi_{\text{dia}} \in \Pi(\mu, \nu)$ can be expressed as

$$F_{\pi_{\text{dia}}}(x, y) = C_{\text{dia}}(F_\mu(x), F_\nu(y)), \quad x, y \in \mathbb{R}. \quad (6.1)$$

Denote by $a \wedge b$ the minimum of a, b and by $a \vee b$ the maximum of a, b . Moreover, let

$$a \diamond b = \frac{a}{2} + \frac{b}{2} - \frac{1}{4}.$$

By direct calculation, the diamond copula has an explicit cdf formula

$$C_{\text{dia}}(u, v) = \begin{cases} (u \diamond v)_+ & (u, v) \in [0, 1/2]^2 \\ (u \diamond v) \wedge v & (u, v) \in (1/2, 1] \times [0, 1/2] \\ (u \diamond v) \wedge u & (u, v) \in [0, 1/2] \times (1/2, 1] \\ (u \diamond v) \vee (u + v - 1) & (u, v) \in (1/2, 1]^2. \end{cases} \quad (6.2)$$

In particular, $C_{\text{dia}}(u, v) = u \diamond v$ when (u, v) is in the area inside D . In the case $\mu = \nu = U$, the diamond copula coincides with the diamond transport; see Figure 2 for an illustration.




Figure 2: The value $C_{\text{dia}}(u, v)$ of the diamond copula, illustrated by distinct values in different regions. The blue shape is the support D of the diamond copula, which also indicates transitions of the cdf across different regions.

6.1 The rectangular cost function

We now consider the *rectangular cost function* $c(x, y, x', y') = |(x - x')(y - y')|$, which is the area of the rectangle formed by the two vertices (x, y) and (x', y') . By Fact 2.1, this cost function is equivalent to the ones in two other problems:

- (a) minimizing the transport cost is equivalent to the problem of inequality minimization in the Introduction with cost function $(\theta_1|x - x'| + \theta_2|y - y'|)^2$ in (1.1) with $\theta_1, \theta_2 > 0$;
- (b) maximizing the transport cost is equivalent to computing the (2, 1)-GW distance in Section 3 defined in (3.3).

In the next theorem, we see that the diamond transport uniquely solves the QOT problem. A family of more general cost functions will be studied in Section 6.3, where we prove analogous results.

Theorem 6.2. *Let $\mu, \nu \in \mathcal{P}_1(\mathbb{R})$. For the rectangular cost function c given by*

$$c(x, y, x', y') = |(x - x')(y - y')|,$$

the unique minimizer of the QOT problem is the diamond transport π_{dia} .

Proof. For $(x, y), (x', y') \in \mathbb{R}^2$, denote by $[(x, y), (x', y')]$ the unique closed rectangle in \mathbb{R}^2 whose sides are parallel to the xy -axes and two of whose corners are given by $(x, y), (x', y')$ if $x \neq x'$ and $y \neq y'$, and the empty set otherwise. It holds that for $\pi \in \Pi(\mu, \nu)$,

$$\begin{aligned} & \iint |(x - x')(y - y')| d\pi(x, y) d\pi(x', y') \\ &= \iiint \mathbb{1}_{\{(u, v) \in [(x, y), (x', y')]\}} d\pi(x, y) d\pi(x', y') du dv \\ &= \iint \pi \otimes \pi(\{(x, y, x', y') \in \mathbb{R}^4 : (u, v) \in [(x, y), (x', y')]\}) du dv. \end{aligned} \quad (6.3)$$

Our goal is to show that the integrand in (6.3) is uniquely minimized for all (u, v) by $\pi = \pi_{\text{dia}}$, which suffices for our purpose: since $\mu, \nu \in \mathcal{P}_1(\mathbb{R})$, the transport cost for the independent coupling

Coupling \ Marginals	π_{com}	π_{ant}	π_{ind}	$\pi_x^{0.5}$	π_{dia}
	U(0, 1), N(0, 1)	(3.296)	(3.296)	2.916	3.051
U(0, 1), U(0, 1)	(0.667)	(0.667)	0.555	0.583	0.547
N(0, 1), N(0, 1)	(8.001)	(8.001)	6.543	7.273	6.439
Exp(1), Exp(1)	(7.998)	6.580	5.998	6.772	5.763
U(0, 1), Exp(1)	3.166	(3.168)	2.832	2.963	2.775
N(0, 1), Exp(1)	7.612	(7.615)	6.255	7.002	6.007

Table 2: Quadratic-form transport costs for different couplings π_{com} , π_{ant} , π_{ind} , $\pi_x^{0.5}$, and π_{dia} , with the cost function $c(x, y, x', y') = (|x - x'| + |y - y'|)^2$. The marginals are chosen from $N(0, 1)$, $U(0, 1)$, or $\text{Exp}(1)$. The smallest transport cost in each row is marked in bold and the largest in brackets. Transport costs are computed using Monte Carlo simulation with sample size 10^7 .

is finite, and hence so is the transport cost for π_{dia} . For a fixed $(u, v) \in \mathbb{R}^2$ such that μ has no atom at u and ν has no atom at v , define $A(u, v) = \pi((-\infty, u] \times (-\infty, v])$. Also, denote by F_μ, F_ν the distribution functions of μ, ν , so F_μ is continuous at u and F_ν is continuous at v . Using $\pi \in \Pi(\mu, \nu)$, we have

$$\begin{aligned}
& \pi \otimes \pi(\{(x, y), (x', y') \in \mathbb{R}^4 : (u, v) \in [(x, y), (x', y')]\}) \\
&= 2(A(u, v)(1 - F_\mu(u) - F_\nu(v) + A(u, v)) + (F_\mu(u) - A(u, v))(F_\nu(v) - A(u, v))) \\
&= 4A(u, v)^2 - (4F_\mu(u) + 4F_\nu(v) - 2)A(u, v) + 2F_\mu(u)F_\nu(v).
\end{aligned}$$

This is a quadratic function in $A(u, v)$, which is uniquely minimized on \mathbb{R} at $A(u, v) = (2F_\mu(u) + 2F_\nu(v) - 1)/4 = F_\mu(u) \diamond F_\nu(v)$. Note that the feasible region for $A(u, v)$ is $(F_\mu(u) + F_\nu(v) - 1)_+ \leq A(u, v) \leq \min\{F_\mu(u), F_\nu(v)\}$. Hence, the unique minimizer is given by

$$A(u, v) = \min \{ \max\{F_\mu(u) \diamond F_\nu(v), F_\mu(u) + F_\nu(v) - 1, 0\}, F_\mu(u), F_\nu(v) \}.$$

By (6.1) and (6.2), this is the cdf of the diamond transport. \square

As explained in item (a) at the beginning of this section, Theorem 6.2 fully solves problem (1.1) described in the Introduction in the Kantorovich setting; that is, the diamond transport yields a minimum inequality quantified by (1.1). For applications in the Monge setting, this result also leads to approximately optimal transport maps by approximating the diamond transport with a permutation map, since the cost function is continuous. A notable feature of this solution, different from the comonotone or antimonotone coupling, is $\mathbb{E}[V | U] = 1/2$ when $U, V \stackrel{\text{law}}{\sim} C_{\text{dia}}$. Hence, in the financial policy application described in the Introduction, if the marginal distributions are uniform, individuals across different wealth levels get the same benefit on average under the optimal policy.

In Table 2, we report some numerical values for the cost function $c(x, y, x', y') = (|x - x'| + |y - y'|)^2$, with different couplings and marginal distributions. We can observe that, while π_{dia} uniquely minimizes the transport cost as shown in Theorem 6.5, the situation is unclear for the maximizers.

The diamond transport is in general not Monge. Obtaining a Monge QOT minimizer for generic discrete marginals seems a technically challenging task, which we do not pursue here.

One may recall from Example 1.1 that for the marginals μ, ν being Bernoulli(1/2), the independent coupling uniquely minimizes the quadratic-form transport cost with the rectangular cost function. It may be useful to note that for this particular pair of discrete marginals, the independent coupling coincides with the diamond coupling because the class $\Pi(\mu, \nu)$ only has one parameter.

Remark 6.3. If μ is an increasing (resp. decreasing) location-scale transform of ν , then a maximizer of the transport cost with the cost function c in Theorem 6.2 is the comonotone (resp. antimonotone) coupling. Indeed, the function $(x, y) \mapsto -xy$ is submodular and the claim follows from Theorem 5.6.

Remark 6.4. The absolute value in the cost function c in Theorem 6.2 is important. If c is specified by $c(x, y, x', y') = (x - x')(y - y')$ (or equivalently, if the objective is $\text{Cov}(X, Y)$; c.f. Example 3.3), then the minimizer is the antimonotone coupling instead of the diamond transport.

6.2 A class of type-XX cost functions with convex QOT

We next provide a general result on the diamond transport as the unique minimizer. An example is the cost function

$$c(x, y, x', y') = e^{-\alpha((x-x')^2 + (y-y')^2)}$$

for $\alpha \in (0, 1/2]$. Specifically, one of our assumptions in this result is that the quadratic-form transport cost is convex in the transport plan π . Schoenberg's theory of complete monotonicity provides a convenient sufficient condition for such convexity. A nonnegative continuous function $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is *completely monotone* if ϕ is C^∞ on $(0, \infty)$ and satisfies $(-1)^n \phi^{(n)}(u) \geq 0$ for $n \geq 0, u > 0$ (see Berg et al. [1984, Section 4.6]). In particular, ϕ is bounded and decreasing. We use the standard calculus notation ϕ' for the first derivative $\phi^{(1)}$, which should not be confused with the apostrophe in x', y' .

Theorem 6.5. *Suppose that both $\mu, \nu \in \mathcal{P}(\mathbb{R})$ are symmetric and the cost function c is given by*

$$c(x, y, x', y') = \phi((x - x')^2) \phi((y - y')^2) \tag{6.4}$$

for some completely monotone function $\phi : \mathbb{R}_+ \rightarrow \mathbb{R}$ satisfying

$$\phi'(u) + 2u\phi''(u) \leq 0, \quad u \in \mathcal{D}, \tag{6.5}$$

where $\mathcal{D} \subseteq \mathbb{R}_+$ is such that $(\mu \otimes \mu)\{|x - x'|^2 \in \mathcal{D}\} = (\nu \otimes \nu)\{|y - y'|^2 \in \mathcal{D}\} = 1$. Then the diamond transport π_{dia} is a minimizer of the QOT problem, and the unique minimizer if ϕ is non-constant.

Remark 6.6. The cost function (6.4) can be written as $c(x, y, x', y') = h(|x - x'|, |y - y'|)$, where $h(u, v) = \phi(u^2)\phi(v^2)$. Since $u, v \geq 0$ and $\phi' \leq 0$, we have $h_{uv} \geq 0$ and hence $-h$ is submodular. By arguing in the same way as Remark 6.3, we see that under the setting of Theorem 6.5, if μ is a location-scale transform of ν , then the comonotone and antimonotone couplings are both maximizers.

Example 6.7. We give three classes of examples of the type-XX cost function c where the conditions in Theorem 6.5 are satisfied and the unique minimizer is given by π_{dia} .

- (i) Let $\phi(u) = e^{-\alpha u}$ where $\alpha \in (0, 1/2]$. In this case,

$$c(x, y, x', y') = e^{-\alpha((x-x')^2 + (y-y')^2)} = e^{-\alpha\|(x,y) - (x',y')\|^2}.$$

The completely monotone condition is evident by definition. To check (6.5), simply note that for all $u \in \mathbb{R}$,

$$\phi'(u) + 2u\phi''(u) = e^{-\alpha u}(-\alpha + 2\alpha^2) \leq 0.$$

- (ii) Suppose that μ, ν are supported on $[-1/2, 1/2]$. Let $\phi(u) = (\beta + u)^{-\gamma}$, where $\gamma > 0$ and $\beta > 2\gamma + 1$. This leads to the cost function

$$c(x, y, x', y') = (\beta + |x - x'|^2)^{-\gamma} (\beta + |y - y'|^2)^{-\gamma}.$$

Again, the completely monotone condition is evident from the definition. To see (6.5), we observe that

$$\phi'(u) + 2u\phi''(u) = \gamma(\beta + u)^{-\gamma-2}(2u(\gamma + 1) - (\beta + u)) \leq 0$$

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

- (iii) Suppose that μ, ν are supported on $[-L, L]$ where $L > 0$. Let $p \in (1/2, 1)$ and $\phi(u) = e^{-\alpha u^p}$, where $0 < \alpha \leq (2p - 1)/(2p(4L^2)^p)$. The cost function is then given by

$$c(x, y, x', y') = e^{-\alpha(|x-x'|^{2p} + |y-y'|^{2p})}. \quad (6.6)$$

The complete monotonicity of ϕ follows from Exercise 55.1 of Sato [1999] (which can be seen from the infinite divisibility of the Weibull distribution with parameter in $(0, 1]$, a result in Steutel [1970]). To verify (6.5), we compute that for $u \in [0, 4L^2]$,

$$\phi'(u) + 2u\phi''(u) = \alpha p u^{p-1} e^{-\alpha u^p} (1 + 2\alpha p u^p - 2p) \leq 0.$$

To prove Theorem 6.5, we introduce a technical lemma. Define

$$\tilde{c}(x, y) := \int c(x, y, x', y') d\pi_{\text{dia}}(x', y'). \quad (6.7)$$

Lemma 6.8. *Assume the same setting as Theorem 6.5. Then \tilde{c} is supermodular on the first and third quadrants, and submodular on the second and fourth quadrants.*

The proof of Lemma 6.8 involves detailed analysis and is deferred to Appendix G.

Proof of Theorem 6.5. We first claim that the resulting QOT problem is convex in π . Since ϕ is completely monotone, the function $\psi(u) := \phi(u^2)$ is a continuous positive definite function on \mathbb{R} by Schoenberg's theorem (Theorem 4.6.13 and Example 5.1.3 of Berg et al. [1984]). In this way, we write $c(x, y, x', y') = \psi(x - x')\psi(y - y')$ where ψ is positive definite, so $(x, x') \mapsto \psi(x - x')$ and $(y, y') \mapsto \psi(y - y')$ are positive definite kernels. By the Schur product theorem, the product of two positive definite kernels is also positive definite (see Theorem 3.1.12 of Berg et al. [1984]). Therefore, $c(x, y, x', y')$ is a positive definite kernel in the two variables $(x, y), (x', y') \in \mathbb{R}^2$. Proposition 4.1 then implies the QOT problem is convex in π .

Since the cost function (6.4) is translation-invariant, we may assume that μ, ν are both symmetric along 0. We next verify

$$\tilde{c}(x, y) = \tilde{c}(-x, y) = \tilde{c}(x, -y) = \tilde{c}(-x, -y). \quad (6.8)$$

By the symmetry of $\pi_{\text{dia}} \in \Pi(\mu, \nu)$, we have

$$\begin{aligned} \tilde{c}(x, y) &= \int c(x, y, x', y') d\pi_{\text{dia}}(x', y') \\ &= \int c(x, y, -x', y') d\pi_{\text{dia}}(x', y') = \int c(-x, y, x', y') d\pi_{\text{dia}}(x', y') = \tilde{c}(-x, y), \end{aligned}$$

where we have used (6.4) in the third equality. Similarly, $\tilde{c}(x, y) = \tilde{c}(x, -y)$, and hence (6.8) holds.

Consider the classic OT problem with cost function \tilde{c} given by (6.7). Denote a minimizer by $\pi_{\tilde{c}}$. By linearity of the classic OT problem and (6.8), another minimizer is given by the symmetrized version $\hat{\pi}_{\tilde{c}}$ of $\pi_{\tilde{c}}$: the law of the uniform mixture of (X, Y) , $(-X, Y)$, $(X, -Y)$, and $(-X, -Y)$ where

$(X, Y) \stackrel{\text{law}}{\simeq} \pi_{\tilde{c}}$. Therefore, $\hat{\pi}_{\tilde{c}}$ is symmetric along the x and y axes. By cyclical monotonicity and Lemma 6.8, any minimizer is antimonotone on the first and third quadrants and comonotone on the second and fourth quadrants. This implies that $\hat{\pi}_{\tilde{c}} = \pi_{\text{dia}}$, and hence π_{dia} is a minimizer of the OT problem with cost function \tilde{c} .

Suppose that $\pi \in \Pi(\mu, \nu)$ satisfies

$$\iint c \, d\pi \otimes d\pi < \iint c \, d\pi_{\text{dia}} \otimes d\pi_{\text{dia}}.$$

For $\delta \in [0, 1]$, let $\pi_\delta = (1 - \delta)\pi_{\text{dia}} + \delta\pi$. It follows by convexity of the QOT that $\delta \mapsto \iint c \, d\pi_\delta \otimes \pi_\delta$ is convex in δ , so there exists $\varepsilon > 0$ such that for all $\delta \in [0, 1]$ small enough,

$$\begin{aligned} & \iint c \, d\pi_{\text{dia}} \otimes d\pi_{\text{dia}} - \delta\varepsilon \\ & > \iint c \, d\pi_\delta \otimes \pi_\delta \\ & = (1 - \delta)^2 \iint c \, d\pi_{\text{dia}} \otimes d\pi_{\text{dia}} + 2\delta(1 - \delta) \iint c \, d\pi_{\text{dia}} \otimes d\pi + \delta^2 \iint c \, d\pi \otimes d\pi \\ & = \iint c \, d\pi_{\text{dia}} \otimes d\pi_{\text{dia}} + 2\delta \left(\iint c \, d\pi_{\text{dia}} \otimes d\pi - \iint c \, d\pi_{\text{dia}} \otimes d\pi_{\text{dia}} \right) + O(\delta^2). \end{aligned}$$

Therefore, letting $\delta \rightarrow 0$ yields

$$\iint c \, d\pi_{\text{dia}} \otimes d\pi < \iint c \, d\pi_{\text{dia}} \otimes d\pi_{\text{dia}},$$

contradicting π_{dia} being a minimizer of the OT problem with cost function \tilde{c} .

The final claim on uniqueness follows from the fact that if ϕ is non-constant, the cost function c is a strictly positive definite kernel in the two variables $(x, y), (x', y') \in \mathbb{R}^2$ (Theorem 3' of Schoenberg [1938]), and hence the transport cost is strictly convex in π with a unique minimizer. \square

6.3 The q -rectangular cost function

Applying Theorem 6.5, we show that the diamond transport solves another class of QOT problems with the q -rectangular cost function $|(x - x')(y - y')|^q$ for $1 < q \leq 2$. Note that in these cases, the transport cost is not convex in π in general, since the map $((x, y), (x', y')) \mapsto |(x - x')(y - y')|^q$ is not a positive definite kernel. QOT with this cost function is equivalent to the one with cost function $(|x - x'|^q + |y - y'|^q)^2$, or equivalently, the maximization of the $(2, q)$ -GW transport cost defined in (3.3).

Theorem 6.9. *Let $q \in (1, 2]$ and $\mu, \nu \in \mathcal{P}_{2+\delta}(\mathbb{R})$ for some $\delta > 0$. Suppose that μ, ν are symmetric and the cost function c is given by*

$$c(x, y, x', y') = |(x - x')(y - y')|^q.$$

Then the diamond transport π_{dia} is a minimizer of the QOT problem.

Proof. Fix $q \in (1, 2]$ and assume without loss of generality that μ, ν are both symmetric around 0. Assume first that μ, ν are supported in $[-L, L]$ for some $L > 0$. For $\alpha > 0$, consider the cost function

$$c_\alpha(x, y, x', y') := \frac{e^{-\alpha(|x-x'|^q + |y-y'|^q)} - 1 + \alpha(|x-x'|^q + |y-y'|^q)}{\alpha^2}.$$

The QOT problem with cost function c_α is equivalent to that with cost function $e^{-\alpha(|x-x'|^q + |y-y'|^q)}$ (c.f. (6.6)), since the term $|x - x'|^q + |y - y'|^q$ is QOT-irrelevant. By Theorem 6.5 and Example

6.7-(iii), for $0 < \alpha \leq (q-1)/(q(4L^2)^{q/2})$, the unique minimizer is given by the diamond transport π_{dia} .

On the other hand, by the Taylor expansion, we have

$$c_\alpha(x, y, x', y') \rightarrow \frac{1}{2}(|x - x'|^q + |y - y'|^q)^2$$

uniformly on $[-L, L]^4$ as $\alpha \rightarrow 0$. Therefore, the QOT problem with cost function $(|x - x'|^q + |y - y'|^q)^2/2$ also has π_{dia} as a minimizer. Removing the QOT-irrelevant terms, we see that this is also the case for the cost function $|(x - x')(y - y')|^q$, as desired.

Now suppose that $\mu, \nu \in \mathcal{P}_{2+\delta}(\mathbb{R})$. We may then apply the stability for QOT (Proposition 4.5) to approximate μ, ν with measures with bounded supports. It remains to show (4.1), i.e.,

$$\sup_{\pi \in \Pi(\mu, \nu)} \iint |(x - x')(y - y')|^{q+\delta/2} d\pi(x, y) d\pi(x', y') < \infty.$$

Indeed, this follows from Cauchy–Schwarz: there exists some constant $C(p, \delta) > 0$ such that uniformly for $\pi \in \Pi(\mu, \nu)$,

$$\begin{aligned} \iint |(x - x')(y - y')|^{q+\delta/2} d\pi(x, y) d\pi(x', y') &= \mathbb{E}_{\pi \otimes \pi}[|X - X'|^{q+\delta/2} |Y - Y'|^{q+\delta/2}] \\ &\leq \mathbb{E}_{\pi \otimes \pi}[|X - X'|^{2q+\delta}]^{1/2} \mathbb{E}_{\pi \otimes \pi}[|Y - Y'|^{2q+\delta}]^{1/2} \\ &\leq C(p, \delta) \mathbb{E}_\mu[|X|^{2q+\delta}]^{1/2} \mathbb{E}_\nu[|Y|^{2q+\delta}]^{1/2}. \end{aligned}$$

This completes the proof. \square

Different from the case of $q = 1$ analyzed in Theorem 6.5, the assumption that both μ, ν are symmetric in Theorem 6.9 is essential for the diamond transport to minimize the QOT problem when $q \in (1, 2]$. In some unreported numerical results, we find that for asymmetric marginals, (approximate) QOT minimizers are quite different from the diamond transport even for the simple case $q = 2$.

For the cost function $c(x, y, x', y') = |(x - x')(y - y')|^q$, $q \in [1, 2]$, a maximizer of the QOT problem is given by the comonotone coupling π_{com} when $\mu = \nu$, as explained in Example 5.7-(iii) (this also holds true for $q > 0$). For this class of QOT problems, we know neither explicit maximizers when μ and ν are not identical nor explicit minimizers when μ and ν are not symmetric.

7 Conclusion

The new framework of quadratic-form optimal transport (QOT) is proposed, with the key feature that the transport cost is linear in $\pi \otimes \pi$. Due to the possible non-convex structure, the QOT problem is difficult to solve numerically. We prove fundamental properties of QOT and highlight cases with explicit solutions, summarized in Table 1. Compared to classic OT, QOT gives rise to two new and special optimal transport plans, the V-transport and the diamond transport. The latter is particularly interesting since it is not Monge, but serves as a universal minimizer of wide classes of QOT problems (Theorems 6.5 and 6.9), some of which are non-convex.

As a new framework, there are many unsolved problems on QOT. We briefly list some promising and important directions below. Details of these directions are further explained in Appendix D.

- (i) In view of Brenier’s theorem (Brenier [1987]), classic OT has a Monge solution under standard conditions. We wonder whether a similar phenomenon exists for QOT. For instance, some transport plans supported on the union of the graphs of two maps, such as π_{dia} and π_x^λ are optimal for some QOT problems, but a general picture is not clear.

- (ii) It is worth exploring how our explicit QOT results, especially the diamond transport, can lead to solutions to QAP.
- (iii) QOT may give rise to various applications to classic OT, especially through regularization that is different from the quadratically regularized OT discussed in Example 3.8.
- (iv) It remains open to solve explicitly the QOT minimizers for many simple cost functions, such as that in Theorem 6.5 without symmetry assumption, $|(x - x')(y - y')|^q$ for $q > 1$, and $\min\{|x - x'|, |y - y'|\}$.

Acknowledgments

We thank Jose Blanchet, Job Boerma, Marcel Nutz, and Johannes Wiesel for helpful discussions, Shengtong Zhang for proposing a simple proof of Theorem 6.2 in the case of uniform marginals, and Facundo Mémoli for bringing up a series of relevant literature on the Gromov–Wasserstein distance. We also thank two referees for their helpful comments that improved this paper. R. Wang acknowledges financial support from the Natural Sciences and Engineering Research Council of Canada (RGPIN-2024-03728, CRC-2022-00141).

Declaration of interests

The authors have no competing interests to declare that are relevant to the content of this article.

References

- L. Ambrosio, E. Brué, and D. Semola. *Lectures on Optimal Transport*, volume 130. Springer, 2021.
- S. Arya, A. Auddy, R. A. Clark, S. Lim, F. Memoli, and D. Packer. The Gromov–Wasserstein distance between spheres. *Foundations of Computational Mathematics*, pages 1–56, 2024.
- J. Backhoff-Veraguas, M. Beiglböck, and G. Pammer. Existence, duality, and cyclical monotonicity for weak transport costs. *Calculus of Variations and Partial Differential Equations*, 58(6):203, 2019.
- M. Bauer, F. Mémoli, T. Needham, and M. Nishino. The Z-Gromov-Wasserstein distance. *arXiv preprint arXiv:2408.08233*, 2024.
- M. Beiglböck and N. Juillet. On a problem of optimal transport under marginal martingale constraints. *The Annals of Probability*, 44(1):42–106, 2016.
- R. Beinert, C. Heiss, and G. Steidl. On assignment problems related to Gromov–Wasserstein distances on the real line. *SIAM Journal on Imaging Sciences*, 16(2):1028–1032, 2023.
- C. Berg, J. P. R. Christensen, and P. Ressel. *Harmonic Analysis on Semigroups: Theory of Positive Definite and Related Functions*, volume 100. Springer, 1984.
- P. Billingsley. *Convergence of Probability Measures*. John Wiley & Sons, 2013.
- G. Birkhoff. Tres observaciones sobre el algebra lineal. *Universidad Nacional de Tucumán Revista Series A*, 5:147–154, 1946.
- M. Blondel, V. Seguy, and A. Rolet. Smooth and sparse optimal transport. In *International Conference on Artificial Intelligence and Statistics*, pages 880–889. PMLR, 2018.

- Y. Brenier. Décomposition polaire et réarrangement monotone des champs de vecteurs. *CR Acad. Sci. Paris Sér. I Math.*, 305:805–808, 1987.
- R. Burkard, M. Dell’Amico, and S. Martello. *Assignment Problems*. SIAM, 2012.
- R. E. Burkard, E. Çela, G. Rote, and G. J. Woeginger. The quadratic assignment problem with a monotone anti-Monge and a symmetric Toeplitz matrix: Easy and hard cases. *Mathematical Programming*, 82:125–158, 1998.
- E. Çela. *The Quadratic Assignment Problem: Theory and Algorithms*, volume 1. Springer Science & Business Media, 2013.
- E. Çela, V. Deineko, and G. J. Woeginger. New special cases of the quadratic assignment problem with diagonally structured coefficient matrices. *European Journal of Operational Research*, 267(3):818–834, 2018.
- S. Chatterjee. A new coefficient of correlation. *Journal of the American Statistical Association*, 116(536):2009–2022, 2021.
- S. Chowdhury and F. Mémoli. The Gromov–Wasserstein distance between networks and stable network invariants. *Information and Inference: A Journal of the IMA*, 8(4):757–787, 2019.
- N. Deb, P. Ghosal, and B. Sen. Measuring association on topological spaces using kernels and geometric graphs. *arXiv preprint arXiv:2010.01768*, 2020.
- N. Deb, P. Ghosal, and B. Sen. Distribution-free measures of association based on optimal transport. *arXiv preprint arXiv:2411.13080*, 2024.
- J. Dickey and J. Hopkins. Campus building arrangement using TOPAZ. *Transportation Research*, 6(1):59–68, 1972.
- T. Dumont, T. Lacombe, and F.-X. Vialard. On the existence of Monge maps for the Gromov–Wasserstein problem. *Foundations of Computational Mathematics*, pages 1–48, 2024.
- M. Essid and J. Solomon. Quadratically regularized optimal transport on graphs. *SIAM Journal on Scientific Computing*, 40(4):A1961–A1986, 2018.
- E. Furman, R. Wang, and R. Zitikis. Gini-type measures of risk and variability: Gini shortfall, capital allocations, and heavy-tailed risks. *Journal of Banking & Finance*, 83:70–84, 2017.
- A. Galichon. *Optimal Transport Methods in Economics*. Princeton University Press, 2018.
- W. Gangbo and R. J. McCann. The geometry of optimal transportation. *Acta Mathematica*, 177(2):113–161, 1996.
- J. L. Gastwirth. A general definition of the Lorenz curve. *Econometrica*, 39(6):1037–1039, 1971.
- A. González-Sanz and M. Nutz. Sparsity of quadratically regularized optimal transport: Scalar case. *arXiv preprint arXiv:2410.03353*, 2024.
- N. Gozlan, C. Roberto, P.-M. Samson, and P. Tetali. Kantorovich duality for general transport costs and applications. *Journal of Functional Analysis*, 273(11):3327–3405, 2017.
- T. C. Koopmans and M. Beckmann. Assignment problems and the location of economic activities. *Econometrica*, 25(1):53–76, 1957.
- N. Kravtsova. The NP-hardness of the Gromov-Wasserstein distance. *arXiv preprint arXiv:2408.06525*, 2024.

- E. L. Lawler. The quadratic assignment problem. *Management Science*, 9(4):586–599, 1963.
- H. Liu, B. Wang, R. Wang, and S. C. Zhuang. Distorted optimal transport. *Mathematics of Operations Research*, 2025.
- E. M. Loiola, N. M. M. De Abreu, P. O. Boaventura-Netto, P. Hahn, and T. Querido. A survey for the quadratic assignment problem. *European Journal of Operational Research*, 176(2):657–690, 2007.
- D. A. Lorenz, P. Manns, and C. Meyer. Quadratically regularized optimal transport. *Applied Mathematics & Optimization*, 83(3):1919–1949, 2021.
- A. J. McNeil, R. Frey, and P. Embrechts. *Quantitative Risk Management: Concepts, Techniques and Tools-Revised Edition*. Princeton University Press, 2015.
- F. Mémoli. On the use of Gromov-Hausdorff distances for shape comparison. In *Proceedings of Point Based Graphics*, 2007.
- F. Mémoli. Gromov-Wasserstein distances and the metric approach to object matching. *Foundations of Computational Mathematics*, 11:417–487, 2011a.
- F. Mémoli. A spectral notion of Gromov-Wasserstein distance and related methods. *Applied and Computational Harmonic Analysis*, 30(3):363–401, 2011b.
- F. Mémoli. Some properties of Gromov-Hausdorff distances. *Discrete & Computational Geometry*, 48:416–440, 2012.
- F. Mémoli and T. Needham. Comparison results for Gromov-Wasserstein and Gromov-Monge distances. *ESAIM: Control, Optimisation and Calculus of Variations*, 30:78, 2024.
- F. Mémoli, A. Munk, Z. Wan, and C. Weitkamp. The ultrametric Gromov-Wasserstein distance. *Discrete & Computational Geometry*, 70(4):1378–1450, 2023.
- R. B. Nelsen. *An Introduction to Copulas*. Springer, 2006.
- M. Nutz. Introduction to entropic optimal transport. *Lecture notes, Columbia University*, 2021.
- M. Nutz. Quadratically regularized optimal transport: Existence and multiplicity of potentials. *SIAM Journal on Mathematical Analysis*, 57(3):2622–2649, 2025.
- M. Nutz and R. Wang. The directional optimal transport. *The Annals of Applied Probability*, 32(2):1400–1420, 2022.
- B. Pass. Multi-marginal optimal transport: theory and applications. *ESAIM: Mathematical Modelling and Numerical Analysis*, 49(6):1771–1790, 2015.
- B. Pass and A. Vargas-Jiménez. A general framework for multi-marginal optimal transport. *Mathematical Programming*, pages 1–36, 2024.
- G. Peyré and M. Cuturi. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- M. Queyranne. Performance ratio of polynomial heuristics for triangle inequality quadratic assignment problems. *Operations Research Letters*, 4(5):231–234, 1986.
- F. Santambrogio. *Optimal Transport for Applied Mathematicians*. Birkhäuser, 2015.
- K.-I. Sato. *Lévy Processes and Infinitely Divisible Distributions*. Cambridge University Press, 1999.

- I. J. Schoenberg. Metric spaces and completely monotone functions. *Annals of Mathematics*, 39(4): 811–841, 1938.
- F. W. Steutel. Preservation of infinite divisibility under mixing and related topics. *Math. Centre Tracts*, 33, 1970.
- K.-T. Sturm. *The space of spaces: curvature bounds and gradient flows on the space of metric measure spaces*, volume 290. American Mathematical Society, 2023.
- A. W. Van der Vaart. *Asymptotic Statistics*, volume 3. Cambridge university press, 2000.
- T. Vayer. A contribution to optimal transport on incomparable spaces. *arXiv preprint arXiv:2011.04447*, 2020.
- C. Villani. *Topics in Optimal Transportation*, volume 58. American Mathematical Society, 2003.
- C. Villani. *Optimal Transport: Old and New*. Springer, 2009.
- J. Wiesel and X. Xu. Sparsity of quadratically regularized optimal transport: Bounds on concentration and bias. *arXiv preprint arXiv:2410.03425*, 2024.
- Z. Zhang, Z. Goldfeld, Y. Mroueh, and B. K. Sriperumbudur. Gromov–Wasserstein distances: Entropic regularization, duality and sample complexity. *The Annals of Statistics*, 52(4):1616–1645, 2024.

Appendices

A Quadratic programming formulation

In the discrete case where μ and ν are supported on N and M points, respectively, QOT can be formulated by a quadratic program. Denote by $\{x_1, \dots, x_N\}$ the support of μ and by $\{y_1, \dots, y_M\}$ the support of ν . Let $\mu_i = \mu(\{x_i\})$ for $i \in [N]$ and $\nu_j = \nu(\{y_j\})$ for $j \in [M]$. Now the measure π can be expressed by a matrix, and we write $\pi = (\pi_{ij})_{i \in [N], j \in [M]}$. QOT can be written as the quadratic program

$$\begin{aligned}
& \text{to minimize} && \sum_{i,k \in [N], j, \ell \in [M]} c(x_i, y_j, x_k, y_\ell) \pi_{ij} \pi_{k\ell} \\
& \text{over} && \pi \in \mathbb{R}_+^{N \times M} \\
& \text{subject to} && \sum_{j \in [M]} \pi_{ij} = \mu_i \quad \text{for all } i \in [N] \\
& && \sum_{i \in [N]} \pi_{ij} = \nu_j \quad \text{for all } j \in [M].
\end{aligned} \tag{A.1}$$

If one considers Monge QOT, there is the extra constraint that $\pi_{ij} \in \{0, \mu_i\}$, and the problem is not a quadratic program.

Further, let us denote by $\boldsymbol{\pi} \in \mathbb{R}^{NM}$ the vectorization of π , which has entries

$$(\boldsymbol{\pi})_{i(M-1)+j} = \pi_{ij}, \quad i \in [N]; j \in [M].$$

Let $\boldsymbol{\mu} \in \mathbb{R}^N$ and $\boldsymbol{\nu} \in \mathbb{R}^M$ be the vectorizations of μ and ν , respectively, and let $\mathbf{1}_n$ be the vector $(1, \dots, 1) \in \mathbb{R}^n$. Moreover, let C be the $NM \times NM$ matrix with entries given by

$$C_{i(M-1)+j, k(M-1)+\ell} = c(x_i, y_j, x_k, y_\ell), \quad i, k \in [N]; j, \ell \in [M].$$

Then (A.1) has the following concise form

$$\begin{aligned} & \text{to minimize } \boldsymbol{\pi}^\top C \boldsymbol{\pi} \\ & \text{over } \boldsymbol{\pi} \in \mathbb{R}_+^{NM} \\ & \text{subject to } \boldsymbol{\pi} \mathbf{1}_N = \boldsymbol{\mu} \text{ and } \boldsymbol{\pi}^\top \mathbf{1}_M = \boldsymbol{\nu}. \end{aligned} \quad (\text{A.2})$$

Note that the constraints in (A.2) are written in matrix form, which can also be written in vector form, but is less concise.

B Independent coupling is rarely a QOT minimizer

In this section, we prove the following result, which states that similarly to classic OT, under mild conditions, the independent coupling cannot be a QOT optimizer. Recall the notation $\pi_{\text{ind}} = \mu \otimes \nu$.

Proposition B.1. *Let $\mu \in \mathcal{P}(\mathfrak{X})$ and $\nu \in \mathcal{P}(\mathfrak{Y})$. Suppose that there exist continuous functions $c_X \in L^1(\mu)$ and $c_Y \in L^1(\nu)$ such that the cost function c is jointly continuous and satisfies*

$$|c(x, y, x', y')| \leq c_X(x) + c_Y(y) + c_X(x') + c_Y(y') \quad (\text{B.1})$$

pointwise. Then for the following statements:

- (i) every $\pi \in \Pi(\mu, \nu)$ is a minimizer of (2.1);
- (ii) π_{ind} is a minimizer or maximizer of (2.1);
- (iii) there exist functions $\varphi : \mathfrak{X} \rightarrow \mathbb{R}$ and $\psi : \mathfrak{Y} \rightarrow \mathbb{R}$ such that

$$\tilde{c}(x, y) := \frac{1}{2} \iint (c(x, y, x', y') + c(x', y', x, y)) d\mu(x') d\nu(y') = \varphi(x) + \psi(y), \quad \pi_{\text{ind}}\text{-a.e.}, \quad (\text{B.2})$$

we have (i) \implies (ii) \implies (iii). Moreover, if the cost function satisfies the conditions of Proposition 4.1, (iii) implies that π_{ind} is a minimizer of (2.1), and hence (ii) and (iii) are equivalent.

Proof. That (i) \implies (ii) is trivial, so we prove (ii) \implies (iii). Let $\pi \in \Pi(\mu, \nu)$ be arbitrary and suppose that π_{ind} is a minimizer. For $\delta \in [0, 1]$, let $\pi_\delta = \delta\pi + (1 - \delta)\pi_{\text{ind}} \in \Pi(\mu, \nu)$. By optimality of π_{ind} and (B.1), for all $\delta \in [0, 1]$,

$$\begin{aligned} \iint c d\pi_{\text{ind}} \otimes d\pi_{\text{ind}} &\leq \iint c d\pi_\delta \otimes d\pi_\delta \\ &= (1 - 2\delta) \iint c d\pi_{\text{ind}} \otimes d\pi_{\text{ind}} + \delta \iint c d(\pi_{\text{ind}} \otimes d\pi_\delta + d\pi_\delta \otimes \pi_{\text{ind}}) + O(\delta^2). \end{aligned}$$

Therefore, we must have, for \tilde{c} defined in (B.2),

$$\frac{1}{2} \int \tilde{c} d\pi_{\text{ind}} = \iint c d\pi_{\text{ind}} \otimes d\pi_{\text{ind}} \leq \frac{1}{2} \iint c d(\pi_{\text{ind}} \otimes d\pi_\delta + d\pi_\delta \otimes \pi_{\text{ind}}) = \frac{1}{2} \int \tilde{c} d\pi_\delta. \quad (\text{B.3})$$

Hence, π_{ind} is a minimizer of the classic OT problem with cost function \tilde{c} . By our assumptions and the dominated convergence theorem, \tilde{c} is continuous. Therefore, classic OT duality (Villani [2009, Theorem 5.10]) yields a \tilde{c} -cyclically monotone set Γ and dual potentials φ, ψ such that $\pi_{\text{ind}}(\Gamma) = 1$ and $\varphi(x) + \psi(y) = \tilde{c}(x, y)$. This implies (B.2). The case of π_{ind} being a maximizer can be similarly established.

For the final statement, suppose that (iii) holds but π has a strictly smaller quadratic-form transport cost than π_{ind} . Denote by $\pi_\delta = \delta\pi + (1 - \delta)\pi_{\text{ind}}$. By the convexity of the quadratic-form transport cost and a similar argument leading to (B.3), it holds that $\int \tilde{c} d\pi_{\text{ind}} > \int \tilde{c} d\pi_\delta$ for $\delta > 0$ small enough. This means that π_{ind} is not a minimizer of the classic OT problem with cost \tilde{c} , contradicting the separability assumption (iii). \square

Example B.2. Suppose that μ, ν are discrete with the cost function given by $c(x, y, x', y') = \mathbb{1}_{\{x=x', y=y'\}} / (\mu(\{x\})\nu(\{y\}))$. This is precisely the regularization term in quadratically regularized OT as discussed in Example 3.8; see (3.4). By Jensen's inequality, the unique minimizer of this QOT problem is given by π_{ind} . In this case, the cost \tilde{c} in (B.2) is constant one.

Condition (B.2) can often be checked explicitly (which often does not hold for common cost functions) and thus offers a neat necessary condition for π_{ind} to be a minimizer and for the QOT problem to be trivial. As a sanity check, for Example 1.1, (B.2) can be easily verified as \tilde{c} is constant on $\{0, 1\}^2$. The same example (as well as Example B.2) also shows that (ii) does not imply (i) in general. We next provide two counter-examples that satisfy (iii) but not (ii).

Example B.3. Let μ, ν be the standard normal distribution, and the cost function c given by $c(x, y, x', y') = ((xy)^2 - 1)(x'y')$. It is easy to verify $\tilde{c} = 0$. If $\pi \in \Pi(\mu, \nu)$ is the joint normal distribution with correlation coefficient $\rho \in [-1, 1]$, then $\iint c d\pi \otimes d\pi = 2\rho^3$. Clearly, this transport cost is neither maximized nor minimized by the independent coupling ($\rho = 0$).

Example B.4. Consider μ uniform on $\mathfrak{X} = \{0, 1\}$ and ν uniform on $\mathfrak{Y} = \{0, 1, 2, 3\}$, with cost function c given by $c(1, 0, 1, 0) = 1$, $c(0, 0, 0, 0) = 1$, $c(1, 1, 1, 1) = -1$, $c(0, 1, 0, 1) = -1$, $c(1, 2, 1, 2) = 1$, $c(0, 2, 0, 2) = 1$, and zero otherwise. Denote the transition probabilities from $0 \in \mathfrak{X}$ to $0, 1, 2, 3 \in \mathfrak{Y}$ respectively by $p, q, r, 1 - (p + q + r)$. It follows that the transition probabilities from $1 \in \mathfrak{X}$ are $1/2 - p, 1/2 - q, 1/2 - r, p + q + r - 1/2$. So, the total transport cost is

$$\left(\frac{1}{2} - p\right)^2 + p^2 - \left(\frac{1}{2} - q\right)^2 + q^2 + \left(\frac{1}{2} - r\right)^2 + r^2 = 2\left(\left(\frac{1}{4} - p\right)^2 - \left(\frac{1}{4} - q\right)^2 + \left(\frac{1}{4} - r\right)^2\right) + \frac{1}{8}.$$

Note that the set of all couplings $\Pi(\mu, \nu)$ can be parameterized by $\{(p, q, r) \in [0, 1/2]^3 : 1/2 \leq p + q + r \leq 1\}$. The minimizers are then given by $(1/4, 0, 1/4)$ and $(1/4, 1/2, 1/4)$, and the maximizers are given by $(0, 1/4, 1/2)$ and $(1/2, 1/4, 0)$. None of them is the independent coupling, which is given by $(1/4, 1/4, 1/4)$. On the other hand, straightforward calculation shows that $\tilde{c}(x, y)$ is a function of y only. For instance, by symmetry of c ,

$$\tilde{c}(x, 0) = \sum_{x' \in \mathfrak{X}} \sum_{y' \in \mathfrak{Y}} c(x, 0, x', y') \mu(\{x'\}) \nu(\{y'\}) = \begin{cases} c(0, 0, 0, 0) \mu(\{0\}) \nu(\{0\}) = 1/8 & \text{if } x = 0; \\ c(1, 0, 1, 0) \mu(\{1\}) \nu(\{0\}) = 1/8 & \text{if } x = 1. \end{cases}$$

C Linear-exponential distance cost functions

C.1 Basic facts

All cost functions in Section 6 are symmetric in $|x - x'|$ and $|y - y'|$. In this section, we consider a special class of type-XX cost function (a sub-class of the one treated in Theorem 5.6) that is not symmetric in $|x - x'|$ and $|y - y'|$, which we call the class of *linear-exponential distance* cost functions, defined by

$$c_\gamma(x, y, x', y') := |y - y'| e^{-\gamma|x - x'|}, \quad \gamma > 0. \quad (\text{C.1})$$

In probabilistic terms, $\iint c_\gamma d\pi \otimes d\pi = \mathbb{E}[|Y - Y'| e^{-\gamma|X - X'|}]$ for $(X, Y), (X', Y') \stackrel{\text{law}}{\sim} \pi$ iid. The intuition is that the cost function (C.1) measures the difference between Y and Y' when X and X' are close, and the parameter γ controls how the distance between X and X' is discounted. The QOT problem is then formulated as

$$\begin{aligned} & \text{to minimize} && \iint |y - y'| e^{-\gamma|x - x'|} d\pi(x, y) d\pi(x', y') \\ & \text{subject to} && \pi \in \Pi(\mu, \nu). \end{aligned} \quad (\text{C.2})$$

We summarize the results for this class of QOT in the following, where we observe that the minimizers and maximizers lead to very different mathematical structures.

- (i) Assume $\mu \in \mathcal{P}_1(\mathbb{R})$ and ν is an increasing (resp. decreasing) location-scale transform of μ . By checking the conditions in Theorem 5.6, the minimizers of (C.2) are (a) the comonotone (resp. antimonotone) coupling when μ is asymmetric, and (b) the comonotone and antimonotone couplings when μ is symmetric.
- (ii) Assume $\mu \in \mathcal{P}_2(\mathbb{R})$ and $\nu \in \mathcal{P}_{2+\delta}(\mathbb{R})$ for some $\delta > 0$. As $\gamma \rightarrow 0^+$, the maximizer of (C.2) converges weakly to the diamond transport.
- (iii) Assume $\mu \in \mathcal{P}(\mathbb{R})$ and $\nu \in \mathcal{P}_1(\mathbb{R})$. As $\gamma \rightarrow \infty$, the unique maximizer of the limit of (C.2) is the independent coupling. The set of minimizers of the limit of (C.2) is given by the set of Monge maps $\mathcal{T}(\mu, \nu)$.

For the minimizers, we only obtain marginals in the same location-scale class as in the first item. For the maximizers, we do not need to assume identical marginals, but we only have asymptotic results as in the second and the third items. For arbitrary μ, ν and fixed $\gamma > 0$, we do not know either the minimizer or the maximizer for the QOT problem in general.

The first item above is rigorously presented in the following simple proposition.

Proposition C.1. *Suppose that $\gamma > 0$, $\mu \in \mathcal{P}_1(\mathbb{R})$, and ν is an increasing (resp. decreasing) location-scale transform of μ . If μ is symmetric, the set of all minimizers of (C.2) is given by the comonotone and antimonotone couplings. If μ is asymmetric, the comonotone (resp. antimonotone) coupling is the unique minimizer.*

Proof. This follows directly from Theorem 5.6 applied with $h(s, t) = se^{-\gamma t}$. To check the conditions, note that the function h is strictly submodular as the product of a strictly increasing function in s and a strictly decreasing function in t . In addition, since $\int |x| d\mu(x) < \infty$, the independent coupling has a finite quadratic-form transport cost, and hence so does the comonotone and antimonotone couplings. \square

In the remainder of this section, we analyze the limit behavior of the optimizers of (C.2) as $\gamma \rightarrow 0^+$ and as $\gamma \rightarrow \infty$, corresponding to the second and third items above.

C.2 The first limit case

Observe that for any $\gamma > 0$, a maximizer of (C.2) also minimizes

$$\iint |y - y'| \left(\frac{1 - e^{-\gamma|x-x'|}}{\gamma} \right) d\pi(x, y) d\pi(x', y')$$

over $\pi \in \Pi(\mu, \nu)$. Formally, as $\gamma \rightarrow 0^+$, we arrive at the limit optimization problem

$$\begin{aligned} & \text{to minimize} && \iint |y - y'| |x - x'| d\pi(x, y) d\pi(x', y') \\ & \text{subject to} && \pi \in \Pi(\mu, \nu). \end{aligned}$$

We have shown in Theorem 6.2 above that if $\mu, \nu \in \mathcal{P}_1(\mathbb{R})$, the unique minimizer is given by the diamond transport π_{dia} in Definition 6.1. For each $\gamma > 0$, let π^γ be a maximizer of (C.2).

Proposition C.2. *Let $\mu \in \mathcal{P}_2(\mathbb{R})$ and $\nu \in \mathcal{P}_{2+\delta}(\mathbb{R})$ for some $\delta > 0$. Then $\lim_{\gamma \rightarrow 0^+} \pi^\gamma = \pi_{\text{dia}}$ weakly.*

Proof. By the definition (C.2) and Fact 2.1, for each $\gamma > 0$, π^γ is also a minimizer of

$$\iint |y - y'| \left(\frac{1 - e^{-\gamma|x-x'|}}{\gamma} \right) d\pi(x, y) d\pi(x', y'). \quad (\text{C.3})$$

Without loss of generality, $\delta \in (0, 1)$. Observe the elementary inequality $|u + e^{-u} - 1| \leq u^{1+\delta}$ for $u \geq 0$.⁷ It follows that uniformly in $\pi \in \Pi(\mu, \nu)$,

$$\begin{aligned} & \left| \iint |y - y'| \left(\frac{1 - e^{-\gamma|x-x'|}}{\gamma} \right) d\pi(x, y) d\pi(x', y') - \iint |y - y'| |x - x'| d\pi(x, y) d\pi(x', y') \right| \\ & \leq \gamma^{\delta/2} \iint |y - y'| |x - x'|^{1+\delta/2} d\pi(x, y) d\pi(x', y') \\ & \leq \gamma^{\delta/2} \left(\iint |y - y'|^2 d\nu(y) d\nu(y') \right)^{1/2} \left(\iint |x - x'|^{2+\delta} d\mu(x) d\mu(x') \right)^{1/2}, \end{aligned}$$

where the right-hand side does not depend on π . Therefore, the functional (C.3) converges uniformly to $\iint |y - y'| |x - x'| d\pi(x, y) d\pi(x', y')$ as $\gamma \rightarrow 0^+$. By Theorem 6.2, $\iint |y - y'| |x - x'| d\pi(x, y) d\pi(x', y')$ is uniquely minimized by π_{dia} . Hence, $\pi^\gamma \rightarrow \pi_{\text{dia}}$ weakly. \square

C.3 The second limit case: Weak OT and measures of association

Next, we study the limit behavior as $\gamma \rightarrow \infty$. Consider the scaled version of (C.2):

$$\frac{\gamma}{2} \iint c_\gamma d\pi \otimes d\pi = \iint |y - y'| \frac{\gamma}{2} e^{-\gamma|x-x'|} d\pi(x, y) d\pi(x', y').$$

As $\gamma \rightarrow \infty$, the double integral has a formal limit of $\mathbb{E}[|Y - Y''|]$, where Y, Y'' are conditionally iid on X and $(X, Y) \stackrel{\text{law}}{\sim} \pi$, which is closely connected to measures of association studied by Deb et al. [2020, 2024].

The limit is verified for well-behaved couplings π in Proposition C.3 below. Stated in probabilistic terms, the following optimization problem arises as $\gamma \rightarrow \infty$:

$$\begin{aligned} & \text{maximize} && \mathbb{E}[|Y - Y''|] \\ & \text{subject to} && Y, Y'' \text{ are conditionally independent given } X; \\ & && (X, Y), (X, Y'') \stackrel{\text{law}}{\sim} \pi; \\ & && \pi \in \Pi(\mu, \nu). \end{aligned} \tag{C.4}$$

This problem does not belong to our QOT framework but is a weak optimal transport problem. For $\pi \in \Pi(\mu, \nu)$, let $\kappa = \{\kappa_x\}_{x \in \mathfrak{X}}$ be a regular disintegration with respect to the first marginal, and we write $\pi = \mu \otimes \kappa$. Given a cost function $c : \mathfrak{X} \times \mathcal{P}(\mathfrak{Y}) \rightarrow \mathbb{R}$, the *weak optimal transport problem* is

$$\begin{aligned} & \text{to minimize} && \int c(x, \kappa_x) d\mu(x) \\ & \text{subject to} && \pi \in \Pi(\mu, \nu). \end{aligned}$$

We refer to Backhoff-Veraguas et al. [2019] and Gozlan et al. [2017] for thorough treatments on this topic.

Proposition C.3. *Let $\mu \in \mathcal{P}(\mathbb{R})$ and $\nu \in \mathcal{P}_1(\mathbb{R})$. Suppose that $\pi = \mu \otimes \kappa \in \Pi(\mu, \nu)$ satisfies the following: either π is absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^2 , or κ_x is continuous in x in the weak topology. Assume further that there exist constants $C, L > 0$ such that for μ -a.e. x , $\int y \kappa_x(dy) \leq C$ and μ has a continuous density $f \leq L$ with respect to the Lebesgue measure. Then*

$$\iint |y - y'| \frac{\gamma}{2} e^{-\gamma|x-x'|} d\pi(x, y) d\pi(x', y') \rightarrow \mathbb{E}[|Y - Y''|],$$

where Y, Y'' are conditionally iid given X .

⁷To see this, observe that $u + e^{-u} - 1 \geq 0$. Let $f(u) = u^{1+\delta} + 1 - u - e^{-u}$. Clearly, $f(u) \geq 0$ for $u \geq 1$. For $u \in (0, 1)$, $f'(u) = (1 + \delta)u^\delta + e^{-u} - 1 \geq (1 + \delta)u^\delta - u \geq 0$. Therefore, $f(u) \geq 0$ for $u \in [0, 1]$.

Proof. Define $g(x, y) := \int |y - y'| \kappa_x(dy')$. We first claim that under our assumptions,

$$\lim_{\gamma \rightarrow \infty} \int \frac{\gamma}{2} e^{-\gamma|u|} g(x+u, y) f(x+u) du = g(x, y) f(x) \quad \pi\text{-a.e.} \quad (\text{C.5})$$

Indeed, if κ_x is continuous in x , $g(x, y)$ would be jointly continuous in (x, y) , and hence (C.5) is valid. On the other hand, by Lebesgue's differentiation theorem, (C.5) holds for x -a.e. for fixed y , and hence for Lebesgue-a.e. (x, y) . If π is absolutely continuous, (C.5) also holds for π -a.e. (x, y) .

Using $\mathbb{P}(Y' \leq x \mid X') = \mathbb{P}(Y' \leq x \mid X, X')$ almost surely for all $x \in \mathbb{R}$, we have by first conditioning on (X, Y) and then conditioning on X' that

$$\begin{aligned} \mathbb{E} \left[\frac{\gamma}{2} |Y - Y'| e^{-\gamma|X - X'|} \right] &= \int \frac{\gamma}{2} \mathbb{E}[|y - Y'| e^{-\gamma|x - X'|}] d\pi(x, y) \\ &= \iint \frac{\gamma}{2} g(x', y) e^{-\gamma|x - x'|} f(x') dx' d\pi(x, y) \end{aligned} \quad (\text{C.6})$$

Next, we apply (C.5) and the dominated convergence theorem to show that as $\gamma \rightarrow \infty$,

$$\iint \frac{\gamma}{2} g(x', y) e^{-\gamma|x - x'|} f(x') dx' d\pi(x, y) \rightarrow \int g(x, y) f(x) d\pi(x, y). \quad (\text{C.7})$$

To see this, it remains to verify

$$\int \sup_{\gamma \geq 0} \int \frac{\gamma}{2} g(x', y) e^{-\gamma|x - x'|} f(x') dx' d\pi(x, y) < \infty. \quad (\text{C.8})$$

By our assumption and the triangle inequality, $g(x, y) \leq |y| + C$ and $f(x') \leq L$. It follows that uniformly in $\gamma \geq 0$,

$$\int \frac{\gamma}{2} g(x', y) e^{-\gamma|x - x'|} f(x') dx' \leq \frac{\gamma L}{2} \int (|y| + C) e^{-\gamma|x - x'|} dx' \leq L(C + |y|).$$

Therefore,

$$\int \sup_{\gamma \geq 0} \int \frac{\gamma}{2} g(x', y) e^{-\gamma|x - x'|} f(x') dx' d\pi(x, y) \leq \int L(C + |y|) d\pi(x, y) < \infty.$$

This proves (C.8). The proof is then complete, by (C.6), (C.7), and the observation that

$$\int g(x, y) f(x) d\pi(x, y) = \mathbb{E}[|Y - Y''|],$$

where Y, Y'' are conditionally iid given X . □

In Proposition C.4 below, we explicitly solve (C.4). The same problem is studied as a special case of Proposition 1.1 of Deb et al. [2020] as a measure of association of (X, Y) , given by

$$\eta(X, Y) = 1 - \frac{\mathbb{E}[|Y - Y''|]}{\mathbb{E}[|Y - Y'|]} \in [0, 1],$$

where Y and Y' are iid and Y and Y'' are conditionally iid given X . Under some additional assumptions on (X, Y) , Deb et al. [2020] showed that if X and Y are non-degenerate, then $\eta(X, Y) = 0$ if and only if X and Y are independent, and $\eta(X, Y) = 1$ if and only if Y is a measurable function of X . Our next result, with a self-contained proof, implies the above conclusion on η . It assumes only the first moment condition on Y , much weaker than the conditions in Deb et al. [2020].

Proposition C.4. *Suppose that $\mu \in \mathcal{P}(\mathbb{R})$ and $\nu \in \mathcal{P}_1(\mathbb{R})$. The unique maximizer π to (C.4) is given by the independent coupling π_{ind} . The set of minimizers of (C.4) is given by the set of Monge maps $\mathcal{T}(\mu, \nu)$.*

Proof. We first analyze the maximizer of (C.4). Since $\nu \in \mathcal{P}_1(\mathbb{R})$, the independent coupling yields a finite transport cost. Let X, U, V be independent, with $X \stackrel{\text{law}}{\sim} \mu$ and $U, V \stackrel{\text{law}}{\sim} U$. Suppose that Y and Y' are conditionally iid on X , and $Y \stackrel{\text{law}}{\sim} \nu$. For any coupling (X, Y) , we note that $(Y, Y', X) \stackrel{\text{law}}{=} (f(X, U), f(X, V), X)$ where $u \mapsto f(x, u)$ is (a regular version of) the conditional quantile function of Y on $X = x$. Write $g(t) = \mathbb{E}[f(X, t)]$ for $t \in [0, 1]$, which implies $g(U) = \mathbb{E}[f(X, U) | U]$. It follows that

$$\begin{aligned} \mathbb{E}[|Y - Y'|] &= \mathbb{E}[f(X, U) \vee f(X, V)] - \mathbb{E}[f(X, U) \wedge f(X, V)] \\ &= \mathbb{E}[f(X, U \vee V)] - \mathbb{E}[f(X, U \wedge V)] \\ &= \int_0^1 g(t) dt^2 - \int_0^1 g(t) d(2t - t^2) \\ &= 2 \int_0^1 g(t)(2t - 1) dt. \end{aligned}$$

For two random variables, we write the convex order relation $Z \leq_{\text{cx}} W$ if $\mathbb{E}[h(Z)] \leq \mathbb{E}[h(W)]$ for all convex functions h such that the two expectations are well-defined. Note that $g(U) \leq_{\text{cx}} f(X, U) \stackrel{\text{law}}{=} Y$. Moreover, $g(U) \stackrel{\text{law}}{=} Y$ when X and Y are independent. By [Furman et al. \[2017, Theorem 4.5\]](#), the functional $X \mapsto \int_0^1 (2t - 1) Q_\mu(t) dt$ is strictly increasing in convex order, where $X \stackrel{\text{law}}{\sim} \mu$. Therefore, $\mathbb{E}[|Y - Y'|]$ is maximized if and only if $g(U) \stackrel{\text{law}}{=} Y$. Therefore, for the maximizer (X, Y) , $\mathbb{E}[f(X, U) | U] = g(U) = f(X, U)$ holds true, implying that X and $f(X, U)$ are independent as X and U are independent. This shows that the independent coupling π_{ind} is the unique maximizer of $\mathbb{E}[|Y - Y'|]$.

Next, we derive the set of minimizers. If the coupling (X, Y) is induced by a Monge map $Y = f(X)$ for some measurable f , the objective is $\mathbb{E}[|Y - Y'|] = \mathbb{E}[|f(X) - f(X)|] = 0$. Conversely, write $\pi = \mu \otimes \kappa$. If $\mathbb{E}[|Y - Y'|] = 0$, then $\mathbb{E}[|Y - Y'| | X] = 0$ almost surely, implying $\mu(\{x : \kappa_x \text{ is degenerate}\}) = 1$, proving that (X, Y) is Monge. \square

The upshot of the above results is that, although [Proposition B.1](#) implies that the independent coupling is never a minimizer for (C.2) with $\gamma > 0$, we expect that the maximizers π^γ behave like the independent coupling as $\gamma \rightarrow \infty$. However, we do not have a proof to guarantee that the maximizer π^γ of (C.2) converges to π_{ind} .

C.4 Numerical approximations for the optimizers

We next present some numerical approximation for QOT optimizers. The goal here is to understand how the QOT minimizers and maximizers for the linear-exponential cost function behave when we cannot compute them explicitly (recall that, for minimizers, we need μ, ν to be in the same location-scale family, and for maximizers, we only have some limiting results).

The QOT problems with cost functions c_γ and $-c_\gamma$ are not convex, indicating that exact solutions may be difficult to compute numerically, and hence we apply heuristic local search algorithms to solve for an optimal transport map. More precisely, we apply the metaheuristic improvement method with pair exchange neighborhood (see Section 3.2 of [Çela \[2013\]](#), or Section 8.2.3 of [Burkard et al. \[2012\]](#)) to a discretized version of the maximization problem with cost function (C.1), which is a quadratic assignment problem. The discretization procedure is justified by [Proposition 4.5](#). The resulting matching may approximate the minimizers and maximizers of $\iint c_\gamma d\pi \otimes d\pi$.

The minimizer of $\iint c_\gamma d\pi \otimes d\pi$ is known to be the comonotone coupling ([Proposition C.1](#)) when μ, ν are in the same location-scale family, and hence we choose uniform and normal marginals, that




Figure 3: Optimal coupling with quadratic-form cost function $c_\gamma(x, y, x', y') = |y - y'|e^{-\gamma|x-x'|}$ for $\gamma \in \{0.3, 2, 6\}$, where μ and ν are the empirical measures of 50 simulated points from $U(0, 1)$ and from $N(0, 1)$, respectively. The optimal coupling is reported in ranks. The support appears close to a λ -shape. The numerical procedure is based on the heuristic improvement method with pair exchange neighborhood with 500 iterations, initiated from the comonotone transport π_{com} .

is, $\mu = U(0, 1)$ and $\nu = N(0, 1)$. In Figure 3, we report the approximate minimizers (normalized by their ranks) with parameters $\gamma \in \{0.3, 2, 6\}$, obtained from the numerical scheme above. Each of them has an interesting “ λ -shaped” support, clearly different from the comonotone coupling, or any other explicit coupling that we studied.

For the maximizers of $\iint c_\gamma d\pi \otimes d\pi$, we do not know explicit forms even for the case $\mu = \nu = U(0, 1)$, so we consider these marginals in the numerical scheme. In Figure 4, we report the approximate maximizers with parameters $\gamma \in \{0.3, 2, 6\}$. Maximizers for smaller γ appear closer to π_{dia} (thus reassuring Proposition C.2), and for larger γ appear closer to π_{ind} (thus reassuring Proposition C.4). The support of the optimizer seems to be contained in a certain symmetric convex shape E_γ in $[0, 1]^2$. As γ increases, the support expands, and less mass is concentrated near the boundary of E_γ but more mass in the interior of E_γ .

D Extended discussions on unsolved questions

We provide details for the promising directions and unsolved problems outlined in Section 7.

- (i) Brenier’s theorem in classic OT states that if $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$, μ is absolutely continuous, and the cost is given by the squared Euclidean distance $\|x - y\|^2$, then the (unique) transport plan is Monge and induced by the gradient of a convex function (Brenier [1987]; see also Santambrogio [2015, Theorem 1.17] for a more general version). The analogous question in the QOT context remains very challenging. Recent studies on the (2,2)-GW cost (3.3) suggest the existence of optimal 2-maps (transport plans supported on the union of the graphs of two maps, such as π_{dia} , π_x^λ , and π_y with x, y flipped) and the non-existence of Monge minimizers under certain assumptions including absolute continuity of μ (Dumont et al. [2024, Theorem 3.6]). Our closed-form results (Propositions 5.1 and 5.3, and Theorems 5.9, 6.5, and 6.9) provide evidence that the existence of optimal 2-maps might be a universal phenomenon for many QOT problems (but not all of them, in view of Figure 4). This phenomenon is also present in many other extensions or special cases of classic OT (mostly on the real line) such as the martingale optimal transport (Beiglböck and Juillet [2016, Corollary 1.6]), the directional optimal transport (Nutz and Wang [2022, Corollary 2.9]), and OT with concave costs (Gangbo and McCann [1996, Theorem 6.4]).




Figure 4: Plots of the QOT maximizers with cost function $c_\gamma(x, y, x', y') = |y - y'|e^{-\gamma|x-x'|}$ for μ, ν both uniformly distributed on 300 equally spaced points on $[0, 1]$ with various parameters γ . The numerical procedure is performed using the heuristic improvement method with pair exchange neighborhood with 10^4 iterations, initiated from the (discretized version of the) diamond transport π_{dia} .

- (ii) Recent works on explicit solutions of QAP (Burkard et al. [2012], Çela et al. [2018]) may hint at certain cost structures leading to further closed-form minimizers of QOT. On the other hand, the analysis of the cost function $c(x, y, x', y') = |(x - x')(y - y')|^q$, $q \in [1, 2]$ seems reminiscent in the QAP literature, and hence our results in Section 6 may potentially inspire new explicitly solvable cases in QAP. In particular, we expect that solving for (1.1) above (equivalent to the case $q = 2$) in the Monge setting may leverage on tools in the QAP literature, where the minimizer is a discrete approximation to the diamond transport in a suitable sense (since the transport cost is continuous in the weak topology).
- (iii) We anticipate that our work will inspire various applications of QOT to classic OT. For example, building on the convex QOT cost functions introduced in Section 6.2, a theory of (convex) quadratic-form regularized OT can be developed. Unlike the quadratically regularized OT discussed in Example 3.8, the quadratic-form approach offers a rich variety of parameterized regularizer classes (Example 6.7) and does not require the solution to be absolutely continuous with respect to the independent coupling. We expect that the quadratic-form regularized OT generally also leads to sparse (or even singular) couplings.
- (iv) There are many simple cost functions for which we do not have an explicit solution to the corresponding QOT problem. We list a few examples below.
 - (a) We wonder whether Theorem 6.5 extends to marginal distributions that are not symmetric. We conjecture that some “diamond-type” coupling is the minimizer of the corresponding QOT problem. Such a coupling is a combination of four comonotone and antitone pieces, and it is numerically supported by Figure 5.
 - (b) In Theorem 6.9, we explicitly solved a class of QOT problems with cost function $|(x - x')(y - y')|^q$, $1 < q \leq 2$ by realizing it as a limit of other solvable classes. We conjecture that the moment condition can be relaxed to $\mu, \nu \in \mathcal{P}_q(\mathbb{R})$ and the minimizer is unique. The case $q > 2$ also deserves future study, as it is equivalent to maximization of the $(2, q)$ -GW transport cost in (3.3).
 - (c) In addition to the results we obtained and the conjectures above, many other cost functions may yield explicit optimizers of the QOT, which need to be further explored. For instance, we do not know the QOT minimizers for the type-XX cost function $\min\{|x - x'|, |y -$

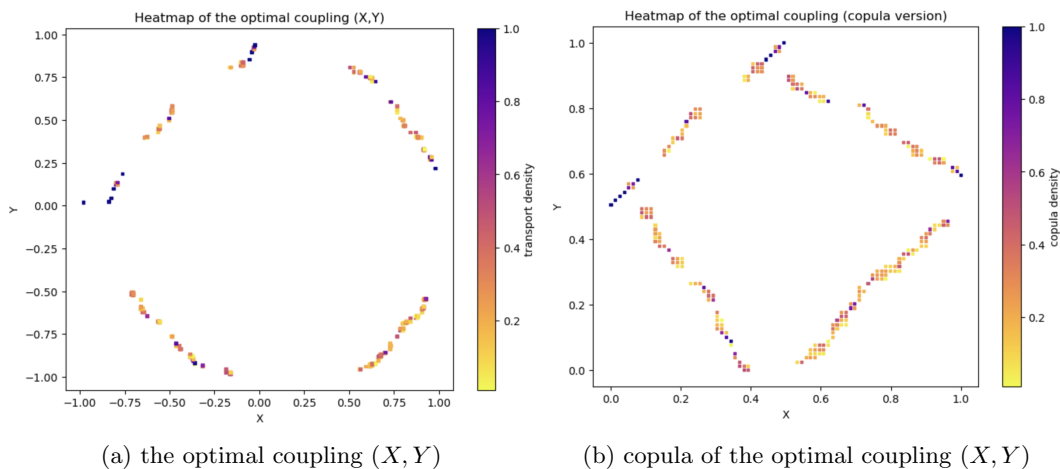


Figure 5: Plots of the quadratic-form optimal coupling (X, Y) and its copula version, with cost function $e^{-((x-x')^2+(y-y')^2)/2}$ and marginals μ uniformly distributed on $[-1, 0] \cup [1/2, 1]$ and ν uniform on $[-1, -1/2] \cup [0, 1]$, both approximated by 80 iid samples. Since the QOT problem is convex (Theorem 6.5), we apply the OSQP solver to find the optimal transport plan, illustrated with the heatmaps. The copula remains of diamond shape, but differs from the diamond copula which is perfectly symmetric.

$y'|$ }, although the QOT minimizers for similar cost functions $\max\{|x - x'|, |y - y'|\}$ and $\min\{x - x', y - y'\}$ are solved in Example 5.5. As another example, we do not know the QOT minimizers for the cost function $|(x - x')(y - y')|^q$, $1 < q \leq 2$ when μ, ν are not symmetric (the symmetric case is solved in Theorem 6.9).

E Omitted proofs of results from Section 4

Proof of Proposition 4.1. Consider distinct transport plans $\pi_0, \pi_1 \in \Pi(\mu, \nu)$ and denote by $\pi_\lambda = (1 - \lambda)\pi_0 + \lambda\pi_1$ their convex combination for $\lambda \in [0, 1]$. By symmetry of ϕ , it holds

$$\begin{aligned}
 & (1 - \lambda) \iint c \, d\pi_0 \otimes d\pi_0 + \lambda \iint c \, d\pi_1 \otimes d\pi_1 - \iint c \, d\pi_\lambda \otimes d\pi_\lambda \\
 & = \lambda(1 - \lambda) \iint c \, d(\pi_0 - \pi_1) \otimes d(\pi_0 - \pi_1).
 \end{aligned} \tag{E.1}$$

Let $\{\pi^{(n)}\}_{n \in \mathbb{N}}$ be a sequence of signed atomic measures converging weakly to the signed measure $\pi_0 - \pi_1$. It follows from Billingsley [2013, Theorem 2.8] that $\pi^{(n)} \otimes \pi^{(n)} \rightarrow (\pi_0 - \pi_1) \otimes (\pi_0 - \pi_1)$ weakly. By the positive definiteness of ϕ , $\iint c \, d\pi^{(n)} \otimes d\pi^{(n)} \geq 0$ for each n . Since c is bounded and continuous,

$$\iint c \, d\pi^{(n)} \otimes d\pi^{(n)} \rightarrow \iint c \, d(\pi_0 - \pi_1) \otimes d(\pi_0 - \pi_1) \quad \text{as } n \rightarrow \infty.$$

This shows that (E.1) is nonnegative, and hence $\pi \mapsto \iint c \, d\pi \otimes d\pi$ is convex. \square

Proof of Proposition 4.4. Suppose that $\pi_n \rightarrow \pi$ in the weak topology on $\Pi(\mu, \nu)$. Theorem 2.8 of Billingsley [2013] yields that $\pi_n \otimes \pi_n \rightarrow \pi \otimes \pi$ weakly in the space of probability measures on $(\mathfrak{X} \times \mathfrak{Y})^2$. Since $c \in \mathcal{C}(\mu, \nu)$ is lower semi-continuous, the map

$$\pi \mapsto \iint c \, d\pi \otimes d\pi$$

is lower semi-continuous by the Portmanteau lemma. Since $\Pi(\mu, \nu)$ is weakly compact, a minimizer of (2.1) exists. The second claim follows immediately since the set $\mathcal{T}(\mu, \nu)$ of Monge transport maps is weakly dense in $\Pi(\mu, \nu)$ for μ atomless and \mathfrak{X} compact (Theorem 1.32 of Santambrogio [2015]). \square

Proof of Proposition 4.5. A standard argument using Prokhorov's theorem shows that $\{\mu_n\}$ and $\{\nu_n\}$ are equi-tight, and hence $\{\pi_n\}$ is relatively compact; see the proof of Theorem 6.8 of Ambrosio et al. [2021]. Let $\pi \in \Pi(\mu, \nu)$ be a limit point of $\{\pi_n\}$. Theorem 2.8 of Billingsley [2013] then implies that $\pi_n \otimes \pi_n \rightarrow \pi \otimes \pi$ weakly. Since c is continuous and satisfies (4.1), we have $\iint c \, d\pi_n \otimes d\pi_n \rightarrow \iint c \, d\pi \otimes d\pi$ (see Van der Vaart [2000, Theorem 2.20] and the example that follows). On the other hand, for any $\hat{\pi} \in \Pi(\mu, \nu)$, Sklar's theorem (McNeil et al. [2015, Theorem 7.3]) implies that there exists a copula C such that the cdf of $\hat{\pi}$ is equal to $C(F_\mu, F_\nu)$, where F_μ is the cdf of μ . Take π'_n specified by its cdf $C(F_{\mu_n}, F_{\nu_n})$. We have $\pi'_n \in \Pi(\mu_n, \nu_n)$ and $\pi'_n \rightarrow \hat{\pi}$ weakly. Therefore,

$$\iint c \, d\pi \otimes d\pi = \lim_{n \rightarrow \infty} \iint c \, d\pi_n \otimes d\pi_n \leq \lim_{n \rightarrow \infty} \iint c \, d\pi'_n \otimes d\pi'_n = \iint c \, d\hat{\pi} \otimes d\hat{\pi}.$$

Altogether, we conclude that π is a QOT minimizer with marginals μ, ν and cost function c . \square

Proof of Proposition 4.6. For each $\pi \in \Pi(\mu, \nu)$, we have by the Fubini–Tonelli theorem,

$$\begin{aligned} \iint c(x, y, x', y') \, d\pi(x, y) \, d\pi(x', y') &= \int \left(\int c(x, y, x', y') \, d\pi(x', y') \right) d\pi(x, y) \\ &\geq \int \mathcal{C}_{c_{x,y}}(\mu, \nu) \, d\pi(x, y) \geq \mathcal{C}_{\hat{c}}(\mu, \nu). \end{aligned} \tag{E.2}$$

This proves the first claim. The second claim follows by noting that, under the given assumptions, both inequalities in (E.2) are equalities. \square

Proof of Proposition 4.7. We extend the domain of the infimum by considering the infimum over a larger class of probability measures on $(\mathfrak{X} \times \mathfrak{Y})^2$ that contains $\pi \otimes \pi$. Define $\Pi_{f,g}$ as the set of probability measures $\tilde{\pi}$ on $(\mathfrak{X} \times \mathfrak{Y})^2$ such that for $(X, Y, X', Y') \stackrel{\text{law}}{\sim} \tilde{\pi}$, we have $f(X, X') \stackrel{\text{law}}{\sim} \mu_f$ and $g(Y, Y') \stackrel{\text{law}}{\sim} \nu_g$. Clearly, $\pi \otimes \pi \in \Pi_{f,g}$. This implies that

$$\begin{aligned} \inf_{\pi \in \Pi(\mu, \nu)} \iint h(f(x, x'), g(y, y')) \, d\pi(x, y) \, d\pi(x', y') &\geq \inf_{\tilde{\pi} \in \Pi_{f,g}} \int h(f(x, x'), g(y, y')) \, d\tilde{\pi}(x, y, x', y') \\ &= \inf_{\hat{\pi} \in \Pi(\mu_f, \nu_g)} \int h(\xi, \zeta) \, d\hat{\pi}(\xi, \zeta), \end{aligned}$$

as desired. \square

F Omitted proofs of results from Section 5

We prove Proposition 5.3 first. The proof of Proposition 5.1 uses Proposition 5.3.

Proof of Proposition 5.3. Using the independence of (X, Y) and (X', Y') , we see that the objective $\mathbb{E}_{\pi \otimes \pi}[c(X, Y, X', Y')]$ is a linear or quadratic function of $\mathbb{E}_\pi[f(X, Y)]$. On the other hand, since f is submodular, the upper and lower bounds for $\mathbb{E}_\pi[f(X, Y)]$ are attained explicitly by π_{com} and π_{ant} , respectively. In addition, for any β in the interval

$$\left[\inf_{\pi \in \Pi(\mu, \nu)} \mathbb{E}_\pi[f(X, Y)], \sup_{\pi \in \Pi(\mu, \nu)} \mathbb{E}_\pi[f(X, Y)] \right], \tag{F.1}$$

there exists $\lambda \in [0, 1]$ such that $\mathbb{E}_{\pi_\lambda}[f(X, Y)] = \beta$ because $\lambda \mapsto \mathbb{E}_{\pi_\lambda}[f(X, Y)]$ is affine. This shows that a QOT minimizer π_λ exists. The last statement follows by noting that, for some choices of

$C_1, C_2, \alpha_1, \alpha_2 \in \mathbb{R}$ (which are arbitrary) in (5.2), the range (F.1) for β is not a singleton and any β in (F.1) can be a unique minimizer. Hence, any $\lambda \in [0, 1]$ can yield a unique minimizer in the class $(\pi_x^\lambda)_{\lambda \in [0, 1]}$. \square

Proof of Proposition 5.1. By independence, we may write

$$\mathbb{E}_{\pi \otimes \pi}[c(X, Y, X', Y')] = \mathbb{E}_\pi[f(X, Y)]\mathbb{E}_\pi[g(X, Y)], \quad \pi \in \Pi(\mu, \nu).$$

The marginal terms of f, g have constant expectations, so there exist constants $C_1, C_2, \alpha_1, \alpha_2 \in \mathbb{R}$ such that

$$\mathbb{E}_\pi[f(X, Y)]\mathbb{E}_\pi[g(X, Y)] = (C_1 + \alpha_1\mathbb{E}_\pi[XY])(C_2 + \alpha_2\mathbb{E}_\pi[XY]). \quad (\text{F.2})$$

The desired result is then a special case of Proposition 5.3 since $(x, y) \mapsto xy$ is supermodular. The last statement follows since any combination of $(C_1, C_2, \alpha_1, \alpha_2)$ can be achieved by certain choices of the quadratic functions f, g . \square

Proof of Theorem 5.4. Define the function $c_{x,y}(x', y') = c(x, y, x', y')$ and recall the notation in Proposition 4.6. It follows from the optimality of π_{com} for submodular cost functions and the Fubini–Tonelli theorem that

$$\iint c \, d\pi_{\text{com}} \otimes d\pi_{\text{com}} = \int \left(\int c(x, y, x', y') \, d\pi_{\text{com}}(x', y') \right) d\pi_{\text{com}}(x, y) = \int \mathcal{C}_{c_{x,y}}(\mu, \nu) \, d\pi_{\text{com}}(x, y).$$

Note that the function $(x, y) \mapsto \mathcal{C}_{c_{x,y}}(\mu, \nu)$ is also submodular, as a weighted combination of submodular functions. As a consequence,

$$\int \mathcal{C}_{c_{x,y}}(\mu, \nu) \, d\pi_{\text{com}}(x, y) = \mathcal{C}_{\hat{c}}(\mu, \nu).$$

In other words, both inequalities in (E.2) are equalities for $\pi = \pi_{\text{com}}$. This implies that π_{com} must be a minimizer because of Proposition 4.6. The supermodular case is analogous. \square

Proof of Theorem 5.6. Without loss of generality, we can assume $\mu = \nu$, as the location-scale transform can be absorbed into h without affecting submodularity or the uniqueness. Let κ be the law of $|X - X'|$, where $X \stackrel{\text{law}}{\sim} \mu$ and X' is an independent copy of X . Let $\hat{\Pi}$ be the set of probability measures $\hat{\pi}$ on \mathbb{R}^4 such that for $(X, Y, X', Y') \stackrel{\text{law}}{\sim} \hat{\pi}$, we have $|X - X'|, |Y - Y'| \stackrel{\text{law}}{\sim} \kappa$. Proposition 4.7 then implies

$$\inf_{\pi \in \Pi(\mu, \nu)} \iint c \, d\pi \otimes d\pi \geq \inf_{\hat{\pi} \in \hat{\Pi}} \int_{\mathbb{R}^4} h(|x - x'|, |y - y'|) \, d\hat{\pi}(x, y, x', y'). \quad (\text{F.3})$$

The integral on the right-hand side of (F.3) depends only on the coupling of $(|Y - Y'|, |X - X'|)$ under the law $\hat{\pi}$. Since the marginals of $|X - X'|$ and $|Y - Y'|$ both follow the law κ under any $\hat{\pi} \in \hat{\Pi}$, the right-hand side of (F.3) coincides with the optimal transport cost between laws κ and κ with cost function h . If h is submodular, the problem is uniquely minimized by π_{com} . This is equivalent to $|X - X'| = |Y - Y'|$ almost surely. To show that the comonotone coupling is a minimizer, observe that under π_{com} , $X = Y$ and $X' = Y'$ hold, and hence $|X - X'| = |Y - Y'|$.

Assume that h is strictly submodular. The right-hand side of (F.3) is then uniquely minimized by the comonotone coupling, or $|X - X'| = |Y - Y'|$ almost surely. It remains to show that π_{com} (and π_{ant} if μ is symmetric) is the unique transport plan that verifies $|X - X'| = |Y - Y'|$. Indeed, this relation implies

$$(X - X' + Y - Y')(X - X' - Y + Y') = (X - X')^2 - (Y - Y')^2 = 0.$$

Hence, either $X + Y = X' + Y'$ or $X - Y = X' - Y'$ almost surely. Since the two sides are independent, we have either $X + Y$ is a constant (only if $(X, Y) \stackrel{\text{law}}{\sim} \pi_{\text{ant}}$) or $X - Y$ is a constant (only if $(X, Y) \stackrel{\text{law}}{\sim} \pi_{\text{com}}$). Since $\mu = \nu$, the comonotone coupling verifies $X - Y = 0$; the antimonotone coupling verifies $X + Y$ is a constant if and only if μ is symmetric. This completes the proof. \square

To prove Theorem 5.10, we need the following lemma.

Lemma F.1. Fix $\alpha, \beta \in (0, 1]$ and $\gamma > 0$. Consider the following optimization problem:

$$\begin{aligned} & \text{maximize} && \mathbb{P}(|X - Y| \leq \gamma) \\ & \text{subject to} && X, Y \text{ are independent with respective densities } f_X, f_Y; \\ & && f_X(x) \leq \alpha^{-1} \mathbb{1}_{[0,1]}(x); \\ & && f_Y(y) \leq \beta^{-1} \mathbb{1}_{[0,1]}(y). \end{aligned}$$

Then an optimizer (X, Y) is given by $X \stackrel{\text{law}}{\simeq} U((1-\alpha)/2, (1+\alpha)/2)$ and $Y \stackrel{\text{law}}{\simeq} U((1-\beta)/2, (1+\beta)/2)$.

The proof of Lemma F.1 is based on a result in Burkard et al. [1998] on discrete assignment. The following result is equivalent to Lemma 2.8 of Burkard et al. [1998], which is the discrete version of Lemma F.1.

Lemma F.2. Let p, q, n be integers satisfying $1 \leq p, q \leq n$. For a fixed $\gamma > 0$, consider the following optimization problem:

$$\begin{aligned} & \text{maximize} && \mathbb{P}(|X - Y| \leq \gamma) \\ & \text{subject to} && X, Y \text{ are independent;} \\ & && X \text{ is uniformly distributed on } p \text{ points in } [n]; \\ & && Y \text{ is uniformly distributed on } q \text{ points in } [n]. \end{aligned}$$

Then an optimizer (X, Y) is given by X being uniformly distributed on the last p points of the finite sequence

$$1, n, 2, n-1, 3, \dots, \tag{F.4}$$

and Y being uniformly distributed on the last q points of (F.4).

Proof of Lemma F.1. Let (X, Y) be given by the claimed optimal solution. Suppose on the contrary that there exist independent random variables \hat{X}, \hat{Y} satisfying the constraints, whose joint law (\hat{X}, \hat{Y}) is different from (X, Y) , and furthermore,

$$\mathbb{P}(|\hat{X} - \hat{Y}| \leq \gamma) > \mathbb{P}(|X - Y| \leq \gamma) \tag{F.5}$$

for some $\gamma > 0$. Note that (\hat{X}, \hat{Y}) is absolutely continuous with bounded density by the constraints. Then there exist a sequence of random variables $(\hat{X}_n, \hat{Y}_n)_{n \geq 1}$ such that:

- for each n , \hat{X}_n and \hat{Y}_n are independent;
- \hat{X}_n (resp. \hat{Y}_n) is supported uniformly on at most $\lfloor \alpha n \rfloor$ (resp. $\lfloor \beta n \rfloor$) points of $\mathbb{Z}/n \cap [0, 1]$;
- $(\hat{X}_n, \hat{Y}_n) \rightarrow (\hat{X}, \hat{Y})$ in distribution.

Similarly, there exist a sequence of random variables $(X_n, Y_n)_{n \geq 1}$ such that:

- for each n , X_n and Y_n are independent;
- X_n (resp. Y_n) is uniformly supported on the last $\lfloor \alpha n \rfloor$ (resp. $\lfloor \beta n \rfloor$) elements of (F.4) scaled by $1/n$.

It follows that $(X_n, Y_n) \rightarrow (X, Y)$ in distribution. By Lemma F.2, we have for each n that

$$\mathbb{P}(|\hat{X}_n - \hat{Y}_n| \leq \gamma) \leq \mathbb{P}(|X_n - Y_n| \leq \gamma).$$

Taking the limit in n and applying the Portmanteau lemma, we have

$$\mathbb{P}(|\hat{X} - \hat{Y}| \leq \gamma) \leq \mathbb{P}(|X - Y| \leq \gamma),$$

leading to a contradiction against (F.5). □

Proof of Theorem 5.10. By absorbing a and b into f , without loss of generality we can assume $a = 0$ and $b = 1$. We first apply the decomposition

$$f(|x - x'|) = \int \mathbb{1}_{\{|x-x'| \geq u\}} d\lambda(u) =: \int c_u(x, x') d\lambda(u),$$

and

$$g(y, y') = \int \mathbb{1}_{[v, \infty) \times [v', \infty)}(y, y') d\eta(v, v') =: \int c_{v, v'}(y, y') d\eta(v, v')$$

where λ, η are some positive measures. For instance, the measure η may be defined via

$$\eta((s, t] \times (s', t']) = g(t, t') + g(s, s') - g(s, t') - g(t, s') \quad \text{for } s < t, s' < t'.$$

By the monotone convergence and Fubini theorems, it remains to show that for every fixed u, u', v, v' π_v is a minimizer of the QOT problem with the cost function

$$c(x, y, x', y') = c_u(x, x')c_{v, v'}(y, y') = \mathbb{1}_{\{|x-x'| \geq u\}} \mathbb{1}_{[v, \infty) \times [v', \infty)}(y, y').$$

Such a problem is equivalent to finding $(X, Y, X', Y') \stackrel{\text{law}}{\sim} \pi$ that minimizes

$$\mathbb{P}(|X - X'| \geq u, Y \geq v, Y' \geq v'), \tag{F.6}$$

subject to $(X, Y) \stackrel{\text{law}}{=} (X', Y')$, the independence of (X, Y) and (X', Y') , and the marginal constraints from π that $X \stackrel{\text{law}}{\sim} \mu$ and $Y \stackrel{\text{law}}{\sim} \nu$. We focus on the case where $\mathbb{P}(Y \geq v) > 0$ and $\mathbb{P}(Y' \geq v') > 0$, otherwise the problem is trivial as (F.6) evaluates to zero. Without loss of generality, we may first remove the constraint that $(X, Y) \stackrel{\text{law}}{=} (X', Y')$ and later show that it is indeed satisfied by the minimizer. Denote by ξ_1 the law of $X \mid Y \geq v$ and ξ_2 the law of $X' \mid Y' \geq v'$. Minimizing (F.6) is then equivalent to minimizing $\mathbb{P}(|\xi_1 - \xi_2| \geq u)$, where ξ_1, ξ_2 are independent. Observe that the marginal constraints on π are equivalent to constraining ξ_1 having density bounded by $1/\mathbb{P}(Y \geq v)$ on $[0, 1]$, and similarly ξ_2 having density bounded by $1/\mathbb{P}(Y' \geq v')$ on $[0, 1]$. Indeed, any such law ξ_1 can be written as the law of $X \mid Y \geq v$ for some coupling (X, Y) satisfying the marginal constraints. In other words, we have reduced to the following problem:

$$\begin{aligned} & \text{to minimize } \mathbb{P}(|\xi_1 - \xi_2| \geq u) \\ & \text{subject to } \xi_1, \xi_2 \text{ are independent r.v.s on } [0, 1] \text{ with respective densities } f_{\xi_1}, f_{\xi_2}; \\ & \quad f_{\xi_1} \leq 1/\mathbb{P}(Y \geq v) \text{ on } [0, 1]; \\ & \quad f_{\xi_2} \leq 1/\mathbb{P}(Y' \geq v') \text{ on } [0, 1]. \end{aligned}$$

By Lemma F.1, a solution is given by

$$\xi_1 \stackrel{\text{law}}{\sim} \text{U}\left(\frac{1 - \mathbb{P}(Y \geq v)}{2}, \frac{1 + \mathbb{P}(Y \geq v)}{2}\right) \quad \text{and} \quad \xi_2 \stackrel{\text{law}}{\sim} \text{U}\left(\frac{1 - \mathbb{P}(Y' \geq v')}{2}, \frac{1 + \mathbb{P}(Y' \geq v')}{2}\right).$$

By Definition 5.9, the V-transport satisfies that for each $v \in \mathbb{R}$,

$$X \mid Y \geq v \stackrel{\text{law}}{\sim} \text{U}\left(\frac{1 - \mathbb{P}(Y \geq v)}{2}, \frac{1 + \mathbb{P}(Y \geq v)}{2}\right).$$

Since π_v does not depend on the choices of u, u', v , the constraint $(X, Y) \stackrel{\text{law}}{=} (X', Y')$ in the minimization problem (F.6) is automatically satisfied. This completes the proof. \square

G Omitted proofs of results from Section 6

Proof of Lemma 6.8. By (6.8), it remains to show that \tilde{c} is supermodular on the support of π_{ind} in $(-\infty, 0]^2$, i.e., with Q_μ, Q_ν denoting the left quantile functions of μ, ν ,

$$\int c_{xy}(Q_\mu(p), Q_\nu(q), Q_\mu(p'), Q_\nu(q')) dC_{\text{dia}}(p', q') \geq 0, \quad p, q \in (0, 1/2), \quad (\text{G.1})$$

where c_{xy} denotes the second-order partial derivative of c with respect to the first two variables. Let $\psi : \mathbb{R}_+ \rightarrow \mathbb{R}$ be the function given by $\psi(u) = \phi(u^2)$. For notational simplicity, let $g : [0, 1]^4 \rightarrow \mathbb{R}$ be given by

$$g(p, q, p', q') := c_{xy}(Q_\mu(p), Q_\nu(q), Q_\mu(p'), Q_\nu(q'))$$

and $\eta_\mu, \eta_\nu : \{(p, p') : 0 \leq p' \leq p \leq 1\} \rightarrow \mathbb{R}$ be given by

$$\eta_\mu(p, p') := \psi'(Q_\mu(p) - Q_\mu(p')) \quad \text{and} \quad \eta_\nu(p, p') := \psi'(Q_\nu(p) - Q_\nu(p')).$$

Using the assumption

$$c(x, y, x', y') = \phi((x - x')^2)\phi((y - y')^2) = \psi(|x - x'|)\psi(|y - y'|)$$

and monotonicity of Q_μ, Q_ν , we have (in the a.e. sense)

$$g(p, q, p', q') = \text{sgn}(p - p') \text{sgn}(q - q') \psi'(|Q_\mu(p) - Q_\mu(p')|) \psi'(|Q_\nu(q) - Q_\nu(q')|). \quad (\text{G.2})$$

We first deal with the case $p + q \leq 1/2$. Using the definition of π_{dia} , we compute the left-hand side of (G.1) as

$$\begin{aligned} & \int g(p, q, p', q') dC_{\text{dia}}(p', q') \\ &= \int_0^{1/2} g\left(p, q, p', \frac{1}{2} + p'\right) dp' + \int_0^{1/2} g\left(p, q, p', \frac{1}{2} - p'\right) dp' \\ & \quad + \int_{1/2}^1 g\left(p, q, p', \frac{3}{2} - p'\right) dp' + \int_{1/2}^1 g\left(p, q, p', p' - \frac{1}{2}\right) dp' \\ &\geq - \int_0^p \eta_\mu(p, p') \eta_\nu\left(\frac{1}{2} + p', q\right) dp' + \int_p^{1/2} \eta_\nu(p', p) \eta_\nu\left(\frac{1}{2} + p', q\right) dp' \\ & \quad - \int_0^{1/2} \psi'(|Q_\mu(p) - Q_\mu(p')|) \psi'\left(\left|Q_\nu\left(\frac{1}{2} - p'\right) - Q_\nu(q)\right|\right) dp' + \int_{1/2}^1 \eta_\nu(p', p) \eta_\nu\left(\frac{3}{2} - p', q\right) dp' \\ & \quad - \int_{1/2}^{1/2+q} \eta_\nu(p', p) \eta_\nu\left(q, p' - \frac{1}{2}\right) dp' + \int_{1/2+q}^1 \eta_\nu(p', p) \eta_\nu\left(p' - \frac{1}{2}, q\right) dp' \\ &= \underbrace{\int_{1-p}^1 \eta_\nu(p', p) \eta_\nu\left(\frac{3}{2} - p', q\right) dp' - \int_0^p \eta_\mu(p, p') \eta_\nu\left(\frac{1}{2} + p', q\right) dp'}_{=: I_1} \\ & \quad + \underbrace{\int_{1-p}^1 \eta_\nu(p', p) \eta_\nu\left(p' - \frac{1}{2}, q\right) dp' - \int_0^p \eta_\mu(p, p') \eta_\nu\left(\frac{1}{2} - p', q\right) dp'}_{=: I_2} \\ & \quad + \underbrace{\int_{1/2+q}^{1-p} \eta_\nu(p', p) \eta_\nu\left(p' - \frac{1}{2}, q\right) dp' - \int_p^{1/2-q} \eta_\nu(p', p) \eta_\nu\left(\frac{1}{2} - p', q\right) dp'}_{=: I_3} \end{aligned}$$

$$\begin{aligned}
& + \underbrace{\int_p^{1/2} \eta_\nu(p', p) \eta_\nu\left(\frac{1}{2} + p', q\right) dp' - \int_{1/2-q}^{1/2} \eta_\nu(p', p) \eta_\nu\left(q, \frac{1}{2} - p'\right) dp'}_{=: I_4} \\
& + \underbrace{\int_{1/2}^{1-p} \eta_\nu(p', p) \eta_\nu\left(\frac{3}{2} - p', q\right) dp' - \int_{1/2}^{1/2+q} \eta_\nu(p', p) \eta_\nu\left(q, p' - \frac{1}{2}\right) dp'}_{=: I_5}.
\end{aligned}$$

By (6.5), we have for all u such that $u^2 \in \mathcal{D}$,

$$\psi''(u) = 2\phi'(u^2) + 4u^2\phi''(u^2) \leq 0. \quad (\text{G.3})$$

As a consequence, ψ' is decreasing on the domain of interest. Since $\phi' \leq 0$, we have $\psi' \leq 0$. Using a change of variable, we obtain

$$I_1 = \int_0^p (\psi'(Q_\mu(1-p') - Q_\mu(p)) - \eta_\mu(p, p')) \eta_\nu\left(\frac{1}{2} + p', q\right) dp' \geq 0.$$

A similar argument using the symmetry properties of Q_μ and Q_ν shows that $I_j \geq 0$ for each $j = 2, 3, 4, 5$. Combining the above yields

$$\int g(p, q, p', q') dC_{\text{dia}}(p', q') \geq 0,$$

proving (G.1) in the case $p + q \leq 1/2$.

Next, we deal with the case $1/2 \leq p + q \leq 1$. It remains to show that

$$\int g\left(\frac{1}{2}, q, p', q'\right) dC_{\text{dia}}(p', q') = 0, \quad \text{for all } q \in [0, 1/2] \quad (\text{G.4})$$

and that for all $p, q \in [0, 1/2]$,

$$\frac{\partial}{\partial x} \int g(p, q, p', q') dC_{\text{dia}}(p', q') = \int c_{xxy}(Q_\mu(p), Q_\nu(q), Q_\mu(p'), Q_\nu(q')) dC_{\text{dia}}(p', q') \leq 0. \quad (\text{G.5})$$

Indeed, integrating (G.5) and using (G.4) imply (G.1). By (6.8) and the smoothness of \tilde{c} , we have $\tilde{c}_x(Q_\mu(1/2), Q_\nu(q)) = 0$ for each $q \in [0, 1]$, and hence (G.4) follows. To prove (G.5), we first differentiate (G.2) to get that in the a.e. sense,

$$\begin{aligned}
& c_{xxy}(Q_\mu(p), Q_\nu(q), Q_\mu(p'), Q_\nu(q')) \\
& = \text{sgn}(p - p')^2 \text{sgn}(q - q') \psi''(|Q_\mu(p) - Q_\mu(p')|) \psi'(|Q_\nu(q) - Q_\nu(q')|) \\
& \quad + 2\delta_{p-p'} \text{sgn}(q - q') \psi'(|Q_\mu(p) - Q_\mu(p')|) \psi'(|Q_\nu(q) - Q_\nu(q')|) \quad (\text{G.6}) \\
& = \text{sgn}(q - q') \psi''(|Q_\mu(p) - Q_\mu(p')|) \psi'(|Q_\nu(q) - Q_\nu(q')|) \\
& \quad + 2\psi'(0) \delta_{p-p'} \text{sgn}(q - q') \psi'(|Q_\nu(q) - Q_\nu(q')|).
\end{aligned}$$

Let $p, q \in [0, 1/2]$ with $p + q \geq 1/2$. We first check that

$$\int \text{sgn}(q - q') \psi''(|Q_\mu(p) - Q_\mu(p')|) \psi'(|Q_\nu(q) - Q_\nu(q')|) dC_{\text{dia}}(p', q') \leq 0. \quad (\text{G.7})$$

Recall that $\psi' \leq 0$ and $\psi'' \leq 0$ by (G.3). For notational simplicity, let $\hat{\eta} : \{(p, p') : 0 \leq p' \leq p \leq 1\} \rightarrow \mathbb{R}$ be given by

$$\hat{\eta}(p, p') := \psi''(|Q_\mu(p) - Q_\mu(p')|).$$

We compute

$$\begin{aligned}
& \int \operatorname{sgn}(q - q') \hat{\eta}(p, p') \psi'(|Q_\nu(q) - Q_\nu(q')|) \, dC_{\text{dia}}(p', q') \\
&= \int_0^{1/2} \operatorname{sgn}\left(q - \frac{1}{2} - p'\right) \hat{\eta}(p, p') \psi' \left(\left| Q_\nu(y) - Q_\nu\left(\frac{1}{2} + p'\right) \right| \right) \, dp' \\
&\quad + \int_0^{1/2} \operatorname{sgn}\left(q - \frac{1}{2} + p'\right) \hat{\eta}(p, p') \psi' \left(\left| Q_\nu(y) - Q_\nu\left(\frac{1}{2} - p'\right) \right| \right) \, dp' \\
&\quad + \int_{1/2}^1 \operatorname{sgn}\left(q - \frac{3}{2} + p'\right) \hat{\eta}(p, p') \psi' \left(\left| Q_\nu(y) - Q_\nu\left(\frac{3}{2} - p'\right) \right| \right) \, dp' \\
&\quad + \int_{1/2}^1 \operatorname{sgn}\left(q + \frac{1}{2} - p'\right) \hat{\eta}(p, p') \psi' \left(\left| Q_\nu(y) - Q_\nu\left(p' - \frac{1}{2}\right) \right| \right) \, dp' \\
&= - \int_0^p \hat{\eta}(p, p') \eta_\nu\left(\frac{1}{2} + p', q\right) \, dp' - \int_p^{1/2} \hat{\eta}(p, p') \eta_\nu\left(\frac{1}{2} + p', q\right) \, dp' \\
&\quad - \int_0^{1/2-q} \hat{\eta}(p, p') \eta_\nu\left(q, \frac{1}{2} - p'\right) \, dp' + \int_{1/2-q}^{1/2} \hat{\eta}(p, p') \eta_\nu\left(q, \frac{1}{2} - p'\right) \, dp' \\
&\quad - \int_{1/2}^{1/2+q} \hat{\eta}(p, p') \eta_\nu\left(\frac{3}{2} - p', q\right) \, dp' - \int_{1/2+q}^1 \hat{\eta}(p, p') \eta_\nu\left(\frac{3}{2} - p', q\right) \, dp' \\
&\quad + \int_{1/2}^{1/2+q} \hat{\eta}(p, p') \eta_\nu\left(q, p' - \frac{1}{2}\right) \, dp' - \int_{1/2+q}^1 \hat{\eta}(p, p') \eta_\nu\left(p' - \frac{1}{2}, q\right) \, dp' \\
&\leq \int_{1/2}^{1/2+q} \hat{\eta}(p, p') \eta_\nu\left(q, p' - \frac{1}{2}\right) \, dp' - \int_{1/2}^{1/2+q} \hat{\eta}(p, p') \eta_\nu\left(\frac{3}{2} - p', q\right) \, dp' \\
&\quad + \int_{1/2-q}^p \hat{\eta}(p, p') \eta_\nu\left(q, \frac{1}{2} - p'\right) \, dp' - \int_{1/2-q}^p \hat{\eta}(p, p') \eta_\nu\left(\frac{1}{2} + p', q\right) \, dp' \\
&\quad + \int_p^{1/2} \hat{\eta}(p, p') \eta_\nu\left(q, \frac{1}{2} - p'\right) \, dp' - \int_p^{1/2} \hat{\eta}(p, p') \eta_\nu\left(\frac{1}{2} + p', q\right) \, dp' \\
&\leq 0,
\end{aligned}$$

where the last step follows from $\hat{\eta}(p, p') \leq 0$ along with the following considerations:

- since $q \leq 1/2$, it holds $2Q_\nu(q) \leq 0 = Q_\nu(p' - \frac{1}{2}) + Q_\nu(\frac{3}{2} - p')$, so that $\eta_\nu(q, p' - \frac{1}{2}) \geq \eta_\nu(\frac{3}{2} - p', q)$;
- again since $q \leq 1/2$, we have $2Q_\nu(q) \leq 0 = Q_\nu(\frac{1}{2} - p') + Q_\nu(\frac{1}{2} + p')$, so $\eta_\nu(q, \frac{1}{2} - p') \geq \eta_\nu(\frac{1}{2} + p', q)$.

This proves (G.7).

In addition, using $p, q \in [0, 1/2]$, (G.3), and the definition of C_{dia} , we obtain

$$\begin{aligned}
& \int 2\delta_{p-p'} \operatorname{sgn}(q - q') \psi'(|Q_\nu(q) - Q_\nu(q')|) \, dC_{\text{dia}}(p', q') \\
&= \operatorname{sgn}\left(q - \frac{1}{2} - p\right) \psi' \left(\left| Q_\nu(q) - Q_\nu\left(\frac{1}{2} + p\right) \right| \right) + \operatorname{sgn}\left(q - \frac{1}{2} + p\right) \psi' \left(\left| Q_\nu(q) - Q_\nu\left(\frac{1}{2} - p\right) \right| \right) \\
&= \psi' \left(Q_\nu(q) - Q_\nu\left(\frac{1}{2} - p\right) \right) - \psi' \left(Q_\nu\left(p + \frac{1}{2}\right) - Q_\nu(q) \right) \geq 0.
\end{aligned}$$

Therefore,

$$\int 2\psi'(0) \delta_{p-p'} \operatorname{sgn}(q - q') \psi'(|Q_\nu(q) - Q_\nu(q')|) \, dC_{\text{dia}}(p', q') \leq 0. \quad (\text{G.8})$$

Combining (G.6), (G.7), and (G.8) yields (G.5) and thus proves (G.1) in the case $p + q \geq 1/2$. \square