

DUALITY FOR OPTIMAL COUPLINGS IN FREE PROBABILITY

WILFRID GANGBO, DAVID JEKEL, KYEONGSIK NAM, AND DIMITRI SHLYAKHTENKO

ABSTRACT. We study the free probabilistic analog of optimal couplings for the quadratic cost, where classical probability spaces are replaced by tracial von Neumann algebras and probability measures on \mathbb{R}^m are replaced by non-commutative laws of m -tuples. We prove an analog of the Monge-Kantorovich duality which characterizes optimal couplings of non-commutative laws with respect to Biane and Voiculescu's non-commutative L^2 -Wasserstein distance using a new type of convex functions. As a consequence, we show that if (X, Y) is a pair of optimally coupled m -tuples of non-commutative random variables in a tracial W^* -algebra \mathcal{A} , then $W^*((1-t)X + tY) = W^*(X, Y)$ for all $t \in (0, 1)$. Finally, we illustrate the subtleties of non-commutative optimal couplings through connections with results in quantum information theory and operator algebras. For instance, two non-commutative laws that can be realized in finite-dimensional algebras may still require an infinite-dimensional algebra to optimally couple. Moreover, the space of non-commutative laws of m -tuples is not separable with respect to the Wasserstein distance when m is sufficiently large.

1. INTRODUCTION

1.1. Context and motivation. Tracial von Neumann algebras have long been viewed as a non-commutative analog of probability spaces, where the elements of the von Neumann algebra play the role of non-commuting random variables, but it was Voiculescu who pointed out that free products of operator algebras provide an analog of probabilistic independence with its own central limit theorem [70, 71], initiating the discipline of free probability theory. Free probability has since had many applications both to random matrix theory e.g. [72] and to von Neumann algebras e.g. [75]. Many developments in free probability theory have been motivated by information geometry (here by “information geometry” we mean the study of the space $\mathcal{P}(M)$ of probability measures on a manifold M , both as a metric space with the Wasserstein distance and as a formal Riemannian manifold, as well as the study of entropy and Fisher's information as functions on $\mathcal{P}(M)$; see [43, 45, 47, 56, 57]). For instance, Voiculescu introduced free entropy and Fisher information [73, 74, 76] and Biane and Voiculescu [11] defined an analog of the L^p Wasserstein distance for non-commutative laws (the analog of probability distributions for m -tuples of non-commuting random variables), which was then used in free Talagrand inequalities [11, 36, 37, 22].

Information-geometric ideas have also been used in quantum information theory, another non-commutative analog of probability theory that is distinct from free probability theory, even though it uses similar concepts and terminology. For a survey of quantum information theory, see [79, 80]. To prevent any confusion, in free probability, operators in a tracial von Neumann algebra are viewed as non-commutative random variables (and there is no known analog of multivariable densities), while in quantum information theory, a positive operator with trace 1 in a tracial von Neumann algebra is viewed as a density.¹ Hence, for example, a random matrix is typically studied in free probability theory, while a matrix-valued density is typically studied in quantum information theory. Our paper is focused on the free probabilistic framework; however, in §5, we will draw a connection between free probabilistic optimal couplings

2020 *Mathematics Subject Classification.* Primary: 46L53; Secondary: 49Q22, 46L52, 81P45.

W.G. was supported by NSF grant DMS-1700202 and Air Force grant FA9550-18-1-0502. D.J. was supported by NSF postdoc grant DMS-2002826. D.S. was supported by NSF grant DMS-1762360. We thank IPAM for the stimulating environment in the online long program on High-dimensional Hamilton-Jacobi PDEs in Spring 2020 where some conversations around this work began. D.S. thanks Alice Guionnet and Yoann Dabrowski for multiple discussions about free transport; D.J. thanks Ben Hayes and Srivatsav Kunnawalkam Elayavalli for discussions about non-commutative laws and model theory.

¹More precisely, a positive operator ρ defines a non-tracial state on the von Neumann algebra, and ρ is the density of this state with respect to the trace. However, tracial von Neumann algebras themselves are difficult to classify, which makes it difficult to classify non-commutative laws in free probability.

and certain aspects of quantum information theory, specifically quantum channels or unital completely positive trace-preserving maps.

In classical information geometry, both the Wasserstein distance and the entropy are intimately related to transport equations (differential equations describing functions which push forward some given probability distribution to another given probability distribution). In the free setting, there has been some success in constructing non-commutative transport of measure for a special type of non-commutative law known as a *free Gibbs law from a convex potential* V in [23, 31, 39, 40, 41]; these ideas have even been generalized beyond the setting of tracial von Neumann algebras [53, 54, 64]. Unfortunately, the transport maps constructed in [23, 39, 40] were not optimal. The transport in [31] was shown to be the gradient of a convex function, hence one would expect it to be optimal in light of the classical Monge-Kantorovich duality, but it was not clear yet how to prove this because there was no known non-commutative Monge-Kantorovich duality. The optimality of these couplings was later verified in [41, Remark 9.11] by studying a Legendre transform for (sufficiently regular, uniformly convex) non-commutative functions [41, Lemma 9.10]. This idea was one of the starting points for our current investigation into non-commutative optimal couplings, Legendre transforms, and Monge-Kantorovich duality with minimal regularity assumptions.

One of the challenges in even formulating a Monge-Kantorovich duality for the free setting is to decide what type of convex functions to use. Operator algebras are often thought of as non-commutative analogs of algebras of functions on a topological space or a measure space, but without a clear analog for points of the underlying space. Our approach is to consider functions that can be evaluated on random variables rather than on points, or more precisely, to study functions $f : L^2(\mathcal{A})_{\text{sa}}^m \rightarrow \mathbb{R}$ where \mathcal{A} is a tracial von Neumann algebra, $L^2(\mathcal{A})$ is the non-commutative L^2 space, and the subscript sa indicates the real subspace of self-adjoint elements. The classical analog would be a function $L^2(\Omega, P; \mathbb{R}^m) \rightarrow \mathbb{R}$ where (Ω, P) is a probability space, rather than a function $\mathbb{R}^m \rightarrow \mathbb{R}$. As we discuss in §1.4, such functions on the space of classical random variables have already found applications to Hamilton-Jacobi equations on the Wasserstein space [29, 26] as well as the master equation on $\mathbb{R}^m \times \mathcal{P}(\mathbb{R}^m)$ in mean field games [14, 27, 28].

As in [31] and [39], we remark that the complexity of classifying von Neumann algebras presents serious obstructions to non-commutative transport theory that simply do not exist in the classical setting. It is a widely used fact in classical probability theory that any two standard Borel probability spaces with no atoms are measurably isomorphic; hence one can always arrange that their random variables are on some canonical probability space. By contrast, McDuff [48] showed that there are uncountably many non-isomorphic tracial von Neumann algebras that are diffuse with trivial center (that is, II_1 factors). This provides a real obstruction to non-commutative transport of measure, because if $X = (X_1, \dots, X_m)$ and $Y = (Y_1, \dots, Y_m)$ are m -tuples of self-adjoint non-commutative random variables such that X is expressed as a “function” of Y and vice versa (for some reasonable notion of non-commutative functions), then X and Y generate the same von Neumann algebra. Hence, non-commutative laws which produce non-isomorphic von Neumann algebras simply cannot be transported to each other in an invertible way. Another result of Ozawa [58] (based on group-theoretic results of Gromov [30] and Olshanskii [55]) shows there is no separable II_1 factor that contains an isomorphic copy of every separable II_1 -factor. Hence, we cannot even expect that there is some non-commutative law μ such that all other non-commutative laws can be expressed as push-forwards of μ .

These obstructions must inform how we go about defining the convex functions for the Monge-Kantorovich duality, as well as the level of regularity that we expect from an optimal coupling. In fact, in §5 we make a more explicit connection between optimal couplings and this result of Gromov, Olshanskii, and Ozawa as well as exploring other pathological properties of the non-commutative Wasserstein distance through connections with quantum information theory.

1.2. Main results. Before stating the non-commutative Monge-Kantorovich duality, we establish following notational conventions; see §2 for background. By *tracial W^* -algebra* we mean a pair $\mathcal{A} = (A, \tau)$ where A is a W^* -algebra (or von Neumann algebra) and $\tau : A \rightarrow \mathbb{C}$ is a faithful normal tracial state. In analogy with classical probability, we will denote the underlying algebra A by $L^\infty(\mathcal{A})$ and the trace by $\tau_{\mathcal{A}}$ when it is convenient to avoid naming A and τ explicitly. We denote by $L^2(\mathcal{A})$ the Hilbert space obtained from the GNS construction of A and τ .

We denote by $L^\infty(\mathcal{A})_{\text{sa}}^m$ the set of m -tuples of self-adjoint elements of $L^\infty(\mathcal{A})$ and for $X = (X_1, \dots, X_m) \in L^\infty(\mathcal{A})_{\text{sa}}^m$, we write $\|X\|_{L^\infty(\mathcal{A})_{\text{sa}}^m} = \max_{j=1, \dots, m} \|X_j\|$. If $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$, then $W^*(X)$ denotes the W^* -algebra generated by X equipped with the appropriate trace.

For each $X = (X_1, \dots, X_m) \in L^\infty(\mathcal{A})_{\text{sa}}^m$, the *non-commutative law* λ_X is the linear map from the non-commutative polynomial algebra $\mathbb{C}\langle x_1, \dots, x_d \rangle$ to \mathbb{C} given by $\lambda_X(p) = \tau_{\mathcal{A}}(p(X))$. The space of non-commutative laws (of self-adjoint m -tuples from any tracial W^* -algebra) is denoted Σ_m . Furthermore, $\Sigma_{m,R}$ denotes the subspace of those laws λ_X where $\|X\|_{L^\infty(\mathcal{A})_{\text{sa}}^m} \leq R$ (where \mathcal{A} is a tracial W^* -algebra and $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$). The *weak- $*$ topology* on $\Sigma_{m,R}$ refers to the topology of pointwise convergence on $\mathbb{C}\langle x_1, \dots, x_m \rangle$.

Following [11], a *coupling* of $\mu, \nu \in \Sigma_m$ is a triple (\mathcal{A}, X, Y) where \mathcal{A} is a tracial W^* -algebra and $X, Y \in L^\infty(\mathcal{A})_{\text{sa}}^m$ such that $\lambda_X = \mu$ and $\lambda_Y = \nu$. The *Wasserstein distance* $d_W^{(2)}(\mu, \nu)$ is the infimum of $\|X - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}$ over all couplings (\mathcal{A}, X, Y) . We denote by $C(\mu, \nu)$ the supremum of $\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$ over all couplings (\mathcal{A}, X, Y) . We say that a coupling is *optimal* if it achieves the infimum of $\|X - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}$ or equivalently if it achieves the supremum of $\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$. The existence of optimal couplings was observed in [11]. That paper also showed that the non-commutative Wasserstein distance agrees with the classical one in the situation that X_1, \dots, X_m commute and Y_1, \dots, Y_m commute [11, Theorem 1.5].

As mentioned before, the functions used in the non-commutative Monge-Kantorovich duality are functions on $L^2(\mathcal{A})_{\text{sa}}^m$ for tracial W^* -algebra \mathcal{A} with separable predual. However, because of Ozawa's result [58], it is not sufficient to fix a single such tracial W^* -algebra, but rather we must consider functions that are defined on $L^2(\mathcal{A})_{\text{sa}}^m$ for every such \mathcal{A} . We give more precise versions of the definitions in §3.

Definition 1.1. A *tracial W^* -function with values in $(-\infty, \infty]$* is a collection of functions $f^{\mathcal{A}} : L^2(\mathcal{A})_{\text{sa}}^m \rightarrow (-\infty, +\infty]$, such that whenever $\iota : \mathcal{A} \rightarrow \mathcal{B}$ is an inclusion map of tracial W^* -algebras, $f^{\mathcal{A}} = f^{\mathcal{B}} \circ \iota$ (here ι is extended to a map $L^2(\mathcal{A})_{\text{sa}}^m \rightarrow L^2(\mathcal{B})_{\text{sa}}^m$). If $\mu \in \Sigma_m$ and f is a tracial W^* -function, then $\mu(f)$ is defined as $f^{\mathcal{A}}(X)$ whenever \mathcal{A} is a tracial W^* -algebra with separable predual and $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$ with $\lambda_X = \mu$; this is well-defined because $W^*(X)$ is determined up to isomorphism by $\lambda_X = \mu$.

One example of a tracial W^* -function would be

$$f^{\mathcal{A}}(X) = \begin{cases} \tau_{\mathcal{A}}(p(X)), & \|X\|_\infty \leq R \\ \infty, & \text{otherwise,} \end{cases}$$

where p is a non-commutative polynomial. Tracial W^* -functions also include scalar-valued tracial non-commutative smooth functions as in [40] and [41] in the following sense. If ϕ is such a tracial non-commutative smooth function, then $\phi^{\mathcal{A}}(X)$ is only a priori when $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$; however, in many cases ϕ is Lipschitz with respect to $\|\cdot\|_{L^2(\mathcal{A})_{\text{sa}}^m}$ and hence can be extended to a function on $L^2(\mathcal{A})_{\text{sa}}^m$ which will be a tracial W^* -function. However, tracial W^* -functions are much more general because they are not assumed to be continuous in any sense.

Definition 1.2. We say that f is *E -convex* if $f^{\mathcal{A}}$ is convex and lower semi-continuous on $L^2(\mathcal{A})_{\text{sa}}^m$ for each \mathcal{A} , and if for every inclusion $\iota : \mathcal{A} \rightarrow \mathcal{B}$, letting $E : \mathcal{B} \rightarrow \mathcal{A}$ be the corresponding trace-preserving conditional expectation, we have $f^{\mathcal{A}}(E[X]) \leq f^{\mathcal{B}}(X)$ for $X \in L^2(\mathcal{B})_{\text{sa}}^m$. Here we use the notation $E[X] = (E[X_1], \dots, E[X_m])$ when $X = (X_1, \dots, X_m)$.

Motivation for the definition of E -convexity will be given in Lemmas 1.10 and 1.17.

Proposition 1.3. $C(\mu, \nu)$ is equal to the infimum of $\mu(f) + \nu(g)$ over pairs (f, g) of E -convex W^* -functions that satisfy $f^{\mathcal{A}}(X) + g^{\mathcal{A}}(Y) \geq \langle X, Y \rangle_{L^2(\mathcal{A})}$ for every tracial W^* -algebra with separable predual and $X, Y \in L^2(\mathcal{A})_{\text{sa}}^m$. There exists an admissible pair of E -convex functions that achieves the infimum. See Propositions 3.22 and 3.23.

Another consequence of the classification-related obstructions to non-commutative transport is that we cannot expect too much regularity in general for the E -convex functions associated to an optimal coupling. For instance, suppose two non-commutative laws μ and ν generate tracial von Neumann algebras that cannot embed into each other. This implies that if (X, Y) is an optimal coupling of these two laws on a tracial W^* -algebra \mathcal{A} , then neither of $W^*(X)$ and $W^*(Y)$ is contained in the other. Thus, even though the non-commutative laws may be diffuse, the situation is similar to when coupling the classical measures

$(1/2)(\delta_{-1} + \delta_1)$ and $(1/3)(\delta_{-1} + \delta_0 + \delta_1)$; in the optimal coupling, neither random variable can be expressed as a function of the other. However, if a pair of E -convex functions associated to an optimal coupling were differentiable, that would imply that X is in the von Neumann algebra generated by Y and vice versa as a consequence of Lemma 3.10.

It is natural to ask how close an arbitrary non-commutative optimal coupling is to a coupling where X and Y generate the same von Neumann algebra. As a first application of duality, we show that every optimal coupling can be decomposed into an optimal coupling where the two variables generate the same W^* -algebra and some additional orthogonal pieces.

Theorem 1.4. *Suppose that (\mathcal{A}, X, Y) is an optimal coupling of $\mu, \nu \in \Sigma_m$. Then there exists a W^* -subalgebra \mathcal{B} such that the following hold. Let $E_{\mathcal{B}} : \mathcal{A} \rightarrow \mathcal{B}$ be the trace-preserving conditional expectation, and let $X' = E[X]$ and $Y' = E[Y]$.*

- (1) X' and Y' each generate \mathcal{B} .
- (2) (\mathcal{B}, X', Y') is an optimal coupling of $\lambda_{X'}$ and $\lambda_{Y'}$.
- (3) $\mathcal{B}, X - X', Y - Y'$ are mutually orthogonal.

See Theorem 3.24.

Our main results in §4 concern the displacement interpolation. If (\mathcal{A}, X, Y) is an optimal coupling of μ and ν , then the *displacement interpolation* refers to the family of random variables $X_t = (1-t)X + tY$ for $t \in [0, 1]$. The associated laws $\mu_t = \lambda_{X_t}$ form a metric geodesic in Σ_m with respect to the Wasserstein distance (see Proposition A.22). With the help of non-commutative Legendre transforms and Hopf-Lax semigroups, we will see that the E -convex functions associated to the couplings (\mathcal{A}, X_s, X_t) for $s, t \in (0, 1)$ have more regularity than the E -convex functions associated to the original coupling (\mathcal{A}, X, Y) (see Proposition 4.12). As a consequence, we obtain the following non-commutative transport result.

Theorem 1.5. *Let (\mathcal{A}, X, Y) be an optimal coupling of $\mu, \nu \in \Sigma_m$. Then $W^*((1-t)X + tY) = W^*(X, Y)$ for all $t \in (0, 1)$. For proof, see §4.3.*

For instance, this theorem entails that for classical optimal couplings, the σ -algebra generated by X_t is the same for all $t \in (0, 1)$, which could be deduced directly from classical optimal transport theory by a similar proof. The reader is encouraged to work out the classical example of $(1/2)(\delta_{-1} + \delta_1)$ and $(1/3)(\delta_{-1} + \delta_0 + \delta_1)$ as motivation.

The results of §5 highlight additional ways in which non-commutative optimal transport theory is significantly more complicated than its classical counterpart. The following proposition is a corollary of the results of [32, 42, 52] obtained by reinterpreting non-commutative optimal couplings in the framework of quantum information theory through Observation 5.5 and Lemma 5.7. Connes-embeddability is a certain type of finite-dimensional approximation; see §5.3 for precise definition.

Proposition 1.6. *Thanks to [32] and [42], for certain $n \in \mathbb{N}$, there exist non-commutative laws μ and ν associated to n^2 -tuples in $M_n(\mathbb{C})$ for which an optimal coupling requires a non-Connes embeddable tracial W^* -algebra; see Corollary 5.14. Furthermore, thanks to [52], for every $n \geq 11$ and $d \in \mathbb{N}$, there exist n^2 -tuples in $M_n(\mathbb{C})$ such that if (\mathcal{A}, X, Y) is a coupling that is optimal among couplings on Connes-embeddable tracial W^* -algebras, then \mathcal{A} must have dimension at least d ; see Corollary 5.8 and Remark 5.15.*

In contrast to classical probability theory, we show that the L^2 -Wasserstein metric does not generate the weak- $*$ topology on $\Sigma_{m,R}$. We call the topology on $\Sigma_{m,R}$ generated by the Wasserstein distance the *Wasserstein topology*. We characterize when the two topologies agree at some μ in terms of the associated tracial W^* -algebra (Proposition 5.21) and hence obtain the following results (relying on the work of Connes [17]).

Proposition 1.7. *The Wasserstein topology on $\Sigma_{m,R}$ is strictly stronger than the weak- $*$ topology; see [11] and Corollary 5.17. Furthermore, let $\Sigma_{m,R}^{\text{fin}}$ denote the set of non-commutative laws λ_X where X comes from $L^2(\mathcal{A})_{\text{sa}}^m$ with \mathcal{A} finite-dimensional. Let μ be a non-commutative law and let \mathcal{A} be a tracial W^* -algebra with a generating m -tuple X such that $\lambda_X = \mu$ and $\|X\|_{L^\infty(\mathcal{A})_{\text{sa}}^m} \leq R$. Then μ is in the weak- $*$ closure of $\Sigma_{m,R}^{\text{fin}}$ if and only if \mathcal{A} is Connes-embeddable; see Lemma 5.12. Moreover, in this case,*

the weak-* and Wasserstein topologies on $\Sigma_{m,R}$ agree at μ if and only if μ is in the Wasserstein closure of $\Sigma_{m,R}^{\text{fin}}$, which is equivalent to \mathcal{A} being approximately finite-dimensional; see Proposition 5.26.

Approximate finite-dimensionality (see §5.4 for definition) is the strongest way that a W^* -algebra can be approximated by finite-dimensional algebras (besides being finite-dimensional itself), and thus the latter condition is quite restrictive when $m > 1$. For instance, there is up to isomorphism only one AFD II_1 factor [67, §XIV.2]. In §6.1, we explain how Propositions 1.6 and 1.7 pose challenges to studying the large- N convergence of Wasserstein distance for random matrix models.

The results of Propositions 1.6 and 1.7 contrast strongly with the classical situation. Some treatments of optimal transport (e.g. [69, p. 75]) take for granted the fact that finitely supported probability measures are weak-* dense in the space of probability measures on a compact set. Such approximation arguments do not work in the non-commutative case for several reasons. Due to the negative resolution of the Connes embedding problem [42], the non-commutative laws that can be realized in finite-dimensional algebras are not weak-* dense. Furthermore, by Proposition 1.7, the weak-* closure of $\Sigma_{m,R}^{\text{fin}}$ is much larger than its Wasserstein closure (assuming $m > 1$). Finally, by Proposition 1.6, even if two laws μ and ν can be realized in finite-dimensional algebras, an optimal coupling need not be weak-* approximable by couplings in finite-dimensional algebras.

Because the weak-* and Wasserstein topologies are different for $m > 1$, one can deduce that $\Sigma_{m,R}$ with the Wasserstein distance is not compact (Corollary 5.27). The following even more startling result is a consequence of Gromov, Olshanskii, and Ozawa's work [58, Theorem 1].

Theorem 1.8. *For $m > 1$ and $R > 0$, the space $\Sigma_{m,R}$ is not separable with respect to $d_W^{(2)}$.*

1.3. Organization. The paper is organized as follows:

- In §1.4 and §1.5, we motivate the definition of E -convex functions and the associated duality result in terms of two toy examples, classical probability spaces and $M_n(\mathbb{C})$.
- In §2, we recall standard background on tracial W^* -algebras and their interpretation as non-commutative probability spaces for the sake of readers who are not specialists in that topic.
- In §3, we describe the properties of E -convex functions and the associated Legendre transform; we prove the non-commutative Monge-Kantorovich duality (Proposition 1.3) and the decomposition theorem for optimal couplings (Theorem 1.4).
- In §4, we study the non-commutative analog of inf-convolution and the regularity properties of E -convex and semi-concave functions; we prove Theorem 1.5 and give further detail about the functions associated to the displacement interpolation in Proposition 4.12.
- In §5, we connect non-commutative optimal couplings with quantum information theory and prove Proposition 1.6. Then we study the differences between the weak-* and the Wasserstein topology using a certain stability property (Proposition 5.21) and hence prove Proposition 1.7. Finally, we show non-separability of the Wasserstein space in §5.5.
- In §6.1, we explain how §5 illustrates the difficulty of studying random matrix optimal transport in the large- N limit. Then §6.2 sketches a different but analogous theory of non-commutative optimal couplings that uses bimodules and UCPT-maps of tracial W^* -algebras.
- In the appendix §A, we define non-commutative laws and optimal couplings for elements of non-commutative L^p spaces, and show the existence of L^p optimal couplings and Wasserstein geodesics.

1.4. Motivation from classical probability. First, we recall the classical Monge-Kantorovich duality. Fix a standard Borel probability space (Ω, P) with no atoms. For probability measures μ and ν compactly supported probability measures on \mathbb{R}^m , a *coupling* of μ and ν is a pair (X, Y) of random variables on Ω with $X \sim \mu$ and $Y \sim \nu$. The classical Wasserstein distance is the infimum of $\|X - Y\|_{L^2(\Omega, P; \mathbb{R}^m)}$ over all such couplings, and a coupling is said to be *optimal* if it achieves this infimum.

Theorem 1.9 (See [69, Theorem 5.10, Particular Case 5.17]). *Let (X, Y) be a coupling of two compactly supported measures μ and ν on \mathbb{R}^m . Then (X, Y) is optimal if and only if there exists a pair of convex functions $f, g : \mathbb{R}^m \rightarrow \mathbb{R}$ satisfying $f(x) + g(y) \geq \langle x, y \rangle$ for $x, y \in \mathbb{R}^m$ and $\mathbb{E}[f(X)] + \mathbb{E}[g(Y)] = \mathbb{E}\langle X, Y \rangle$.*

Furthermore, $\mathbb{E}[f(X)] + \mathbb{E}[g(Y)] = \mathbb{E}\langle X, Y \rangle$ implies that Y is almost surely in the subdifferential of f at X and X is almost surely in the subdifferential of g at Y .

As explained above, E -convex functions will be an analog of functions on $L^2(\Omega, P; \mathbb{R}^m)$ rather than \mathbb{R}^m . Every convex function on \mathbb{R}^m defines a convex function on $L^2(\Omega, P; \mathbb{R}^m)$ as follows.

Lemma 1.10. *Let $f : \mathbb{R}^m \rightarrow (-\infty, \infty]$ be convex and lower semi-continuous. Let (Ω, P) be a non-atomic standard Borel probability space with underlying σ -algebra \mathcal{F} . Define*

$$\tilde{f} : L^2(\Omega, P; \mathbb{R}^m) \rightarrow \mathbb{R}, X \mapsto \mathbb{E}[f(X)],$$

which is well-defined in $(-\infty, \infty]$ thanks to Jensen's inequality. Then

- (1) $\tilde{f}(X)$ only depends on the law (probability distribution) of X .
- (2) \tilde{f} is convex and lower semi-continuous.
- (3) Suppose that $\tilde{f}(X) < \infty$. Then Y is in the subdifferential of \tilde{f} at X if and only if Y is in the subdifferential of f at X almost surely.
- (4) \tilde{f} is monotone under conditional expectations: If \mathcal{G} is a sub- σ -algebra of \mathcal{F} , then

$$\tilde{f}(E[X|\mathcal{G}]) \leq \tilde{f}(X).$$

Sketch of proof. (1) This is immediate.

(2) Convexity of \tilde{f} is immediate from convexity of f . To show lower semi-continuity of \tilde{f} , it suffices to show lower semi-continuity of \tilde{g} where $g(x) = f(x) + |x|^2/2$. But g is bounded below and therefore lower semi-continuity follows from Fatou's lemma for convergence in probability (since convergence in $L^2(\Omega, P; \mathbb{R}^m)$ implies convergence in probability).

(3) If Y is in the subdifferential of f at X almost surely and $Z \in L^2(\Omega, P)$, then $f(Z) \geq f(X) + \langle X, Y \rangle_{\mathbb{R}^m}$ almost surely, and thus by taking expectations $\tilde{f}(Z) \geq \tilde{f}(X) + \langle X, Y \rangle_{L^2(\Omega, P; \mathbb{R}^m)}$. For the converse, let $S = \{x \in \mathbb{R}^m : f(x) < \infty\}$ and fix a countable dense subset Ξ of S . For each $n > 0$ and $\xi \in \Xi$, let $E_{n,\xi}$ be the event

$$E_{n,\xi} = \{f(\xi) \leq f(X) + \langle \xi - X, Y \rangle_{\mathbb{R}^m} - 1/n\}.$$

Because Y is in the subdifferential of \tilde{f} at X , we have

$$\tilde{f}(1_{E_{n,\xi}^c} X + 1_{E_{n,\xi}} \xi) \geq \tilde{f}(X) + \langle 1_{E_{n,\xi}}(\xi - X), Y \rangle_{L^2(\Omega, P; \mathbb{R}^m)}.$$

On the other hand, by definition of $E_{n,\xi}$, we have

$$\tilde{f}(1_{E_{n,\xi}^c} X + 1_{E_{n,\xi}} \xi) \leq \tilde{f}(X) + \langle 1_{E_{n,\xi}}(\xi - X), Y \rangle_{L^2(\Omega, P; \mathbb{R}^m)} + \frac{1}{n} P(E_{n,\xi}).$$

Therefore, $P(E_{n,\xi}) = 0$. Since this holds for all $n \in \mathbb{N}$, we have $f(\xi) \geq f(X) + \langle \xi - X, Y \rangle_{\mathbb{R}^m}$ almost surely for each ξ . Since Ξ is countable, we have this condition every $\xi \in \Xi$ at once almost surely. On this event, if $x \in \mathbb{R}^m$ with $f(x) < \infty$, then f is continuous at x , and therefore by taking sequence of $\xi \in \Xi$ that converges to x we obtain $f(x) \geq f(X) + \langle x - X, Y \rangle_{\mathbb{R}^m}$.

(4) This follows from Jensen's inequality and the existence of regular conditional distributions for standard Borel probability spaces. \square

Remark 1.11. Similar reasoning shows that if g is the Legendre transform of f on \mathbb{R}^m , then \tilde{g} is the Legendre transform of \tilde{f} on $L^2(\Omega, P; \mathbb{R}^m)$.

Let us call a function $F : L^2(\Omega, P; \mathbb{R}^m) \rightarrow (-\infty, \infty]$ *classically E -convex* if

- (1) $F(X)$ depends only on the law of X .
- (2) F is convex and lower semi-continuous.
- (3) We have $F(E[X|\mathcal{G}]) \leq F(X)$ for every sub- σ -algebra \mathcal{G} and every $X \in L^2(\Omega, P; \mathbb{R}^m)$.

Then we have the following version of Monge-Kantorovich duality using classically E -convex functions on $L^2(\Omega, P; \mathbb{R}^m)$.

Corollary 1.12. *Let (X, Y) be a coupling on (Ω, P) of two compactly supported measures μ and ν on \mathbb{R}^m . Then (X, Y) is optimal if and only if there exists a pair of classically E -convex functions $F, G : L^2(\Omega, P; \mathbb{R}^m) \rightarrow (-\infty, \infty]$ such that*

$$F(X') + G(Y') \geq \langle X', Y' \rangle_{L^2(\Omega, P; \mathbb{R}^m)} \text{ for all } X', Y' \in L^2(\Omega, P; \mathbb{R}^m),$$

and

$$F(X) + G(Y) = \langle X, Y \rangle_{L^2(\Omega, P; \mathbb{R}^m)}.$$

Proof. (\implies) By the classical Monge-Kantorovich duality, there are convex functions $f, g : \mathbb{R}^m \rightarrow (-\infty, \infty]$ with $f(x) + g(y) \geq \langle x, y \rangle_{\mathbb{R}^m}$ and $\mathbb{E}f(X) + \mathbb{E}g(Y) = \langle X, Y \rangle_{L^2(\Omega, P; \mathbb{R}^m)}$. Let $F = \tilde{f}$ and $G = \tilde{g}$. By Lemma 1.10, F and G are classical E -convex and clearly $F(X) + G(Y) = \langle X, Y \rangle_{L^2(\Omega, P; \mathbb{R}^m)}$. Also, $F(X') + G(Y') \geq \langle X', Y' \rangle_{L^2(\Omega, P; \mathbb{R}^m)}$ since $f(x) + g(y) \geq \langle x, y \rangle_{\mathbb{R}^m}$.

(\impliedby) Suppose that (X', Y') is another coupling of μ and ν on (Ω, P) . Then

$$\langle X', Y' \rangle_{L^2(\Omega, P; \mathbb{R}^m)} \leq F(X') + G(Y') = F(X) + G(Y) = \langle X, Y \rangle_{L^2(\Omega, P; \mathbb{R}^m)},$$

where in the middle equality we have used that $F(X) = F(X')$ and $G(Y) = G(Y')$ since $X \sim X'$ and $Y \sim Y'$ in law. Therefore, the coupling (X, Y) is optimal. \square

Corollary 1.12 is the statement that we will generalize to the non-commutative setting. We remark that although classically E -convex functions are much less concrete than convex functions on \mathbb{R}^m , Corollary 1.12 still has the power to prove the classical analogs of Theorems 1.4 and 1.5 by exactly the same arguments that we will use in the non-commutative case.

In fact, convex functions on a space of classical random variables have also been used in the theory of mean field games [28]. Mean field games involves the study of the *master equation* [14, 27], a differential equation for a function $u(t, x, \mu)$ depending on a time variable t , a space variable x (representing the position of an individual agent), and a measure μ (representing the distribution of the positions of a continuum of other agents). We can define a function \hat{u} on $[0, \infty) \times \mathbb{R}^m \times L^2(\Omega, P; \mathbb{R}^m)$ by $\hat{u}(t, x, X) = u(t, x, \mu_X)$, where μ_X is the law of X . The first-order regularity conditions needed to solve the master equation are more easily stated in terms of the function \hat{u} on the Hilbert space $\mathbb{R}^m \times L^2(\Omega, P; \mathbb{R}^m)$. Moreover, the proof of existence and uniqueness of solutions to Hamilton-Jacobi equations on Wasserstein space $\mathcal{P}_2(\mathbb{R}^m)$ [26, 29] relies on the theory of viscosity solutions to Hamilton-Jacobi equations on Hilbert spaces [18, 19, 49, 46].

The inf-convolution techniques that we use in §4 are an important special case of this theory of Hamilton-Jacobi equations on Hilbert spaces. In fact, part of our motivation was to understand the non-commutative version of Hamilton-Jacobi equations for functions of a random variable. Recent work has connected random matrix theory to viscosity solutions of Hamilton-Jacobi equations [10] and mean field games [15]. However, these connections are restricted to the setting of a single random matrix because they rely heavily on the description of self-adjoint random matrices in terms of their eigenvalues. It would be of great interest to have a theory of viscosity solutions to partial differential equations in several non-commuting variables as is suggested by the study of heat equations in [23, 38, 41] and the Hamilton-Jacobi-Bellman equation in [21, 38].

1.5. Motivation from matrix tuples. In order to motivate some of the ideas of our paper, we explain a toy model of couplings between tuples of $n \times n$ matrices. Let $M_n(\mathbb{C})$ denote the space of complex $n \times n$ matrices. Let $\text{tr}_n = (1/n) \text{Tr}_n$ be the normalized trace on $M_n(\mathbb{C})$. We define an inner product on $M_n(\mathbb{C})$ by

$$\langle S, T \rangle_{\text{tr}_n} = \text{tr}_n(S^*T).$$

Let $M_n(\mathbb{C})_{\text{sa}}$ denote the real subspace of self-adjoint matrices. Then $\langle X, Y \rangle_{\text{tr}_n} \in \mathbb{R}$ for all $X, Y \in M_n(\mathbb{C})_{\text{sa}}$. Every element of $M_n(\mathbb{C})$ can be uniquely written as $S + iT$ with $S, T \in M_n(\mathbb{C})_{\text{sa}}$, and hence there is a natural identification of the complex inner product space $M_n(\mathbb{C})$ with the complexification of the real inner product space $M_n(\mathbb{C})_{\text{sa}}$.

From a non-commutative probability viewpoint, we can view $M_n(\mathbb{C})$ as an algebra of “random variables” and the normalized trace $\text{tr}_n : M_n(\mathbb{C}) \rightarrow \mathbb{C}$ as the “expectation.” To motivate this, suppose

$X \in M_n(\mathbb{C})_{\text{sa}}$. The *empirical spectral distribution* of X is the measure $\mu = \frac{1}{n} \sum_{j=1}^n \delta_{\lambda_j}$ where $\lambda_1, \dots, \lambda_n$ are the eigenvalues of X listed with multiplicity. We then have for every polynomial p that

$$\text{tr}_n(p(X)) = \int p d\mu.$$

Thus, μ is analogous to the distribution of a random variable.

If $X = (X_1, \dots, X_m) \in M_n(\mathbb{C})_{\text{sa}}^m$, the “joint distribution” of X_1, \dots, X_m is not described by a measure on \mathbb{R}^d , since X_1, \dots, X_m do not commute. Rather we consider the *non-commutative law* λ_X , which is the linear functional on the algebra of m -variable non-commutative polynomials given by

$$p \mapsto \text{tr}_n(p(X_1, \dots, X_m)).$$

It turns out that two tuples X and $Y \in M_n(\mathbb{C})_{\text{sa}}^m$ agree if and only if they are unitarily conjugate (which is equivalent to the fact that any two trace-preserving embeddings of a finite-dimensional tracial $*$ -algebra into $M_n(\mathbb{C})$ are unitarily conjugate).

Lemma 1.13 (see e.g. [25, §2]). *Let $X, Y \in M_n(\mathbb{C})_{\text{sa}}^m$. Then the following are equivalent*

- (1) $\text{tr}_n(p(X)) = \text{tr}_n(p(Y))$ whenever p is a non-commutative polynomial in m variables.
- (2) There exists a unitary U in $M_n(\mathbb{C})$ such that $Y_j = UX_jU^*$ for $j = 1, \dots, m$.

We consider the toy problem of optimally coupling two matrix tuples inside $M_n(\mathbb{C})$ (beware that because of Proposition 1.6 an optimal coupling inside $M_n(\mathbb{C})$ is not necessarily optimal among all couplings in tracial W^* -algebras). Because of Lemma 1.13, the toy problem reduces to the following: Given $X, Y \in M_n(\mathbb{C})$, find a unitary U so that $\|UXU^* - Y\|_{\text{tr}_n}$ is as small as possible, where $UXU^* = (UX_1U^*, \dots, UX_mU^*)$, and where $\|\cdot\|_{\text{tr}_n}$ is the normalized Hilbert-Schmidt norm

$$\|T\|_{\text{tr}_n} = \left(\sum_{j=1}^m \text{tr}_n(T_j^* T_j) \right)^{1/2}.$$

This motivates the following definition: For $X, Y \in M_n(\mathbb{C})_{\text{sa}}^m$, we say that (X, Y) are an *optimal coupling* in $M_n(\mathbb{C})$ if $\|UXU^* - Y\|_{\text{tr}_n} \geq \|X - Y\|_{\text{tr}_n}$ for every unitary U . The next lemma guarantees existence of optimal couplings.

Lemma 1.14. *Let $X, Y \in M_n(\mathbb{C})_{\text{sa}}^m$. Then there exists an $n \times n$ unitary U that minimizes $\|UXU^* - Y\|_{\text{tr}_n}$. Moreover, every such unitary must satisfy*

$$\sum_{j=1}^m [UX_jU^*, Y] = 0,$$

where $[S, T] = ST - TS$ is the commutator.

Proof. Existence of a minimizer follows from the fact that the unitary group is compact and $U \mapsto \|UXU^* - Y\|_{\text{tr}_n}$ is continuous. Now suppose that U is a minimizer and let $Z = UXU^*$. Let A be a self-adjoint matrix, and consider the unitary e^{itA} for $t \in \mathbb{R}$. By minimality, we have $\|e^{itA}Ze^{-itA} - Y\|_{\text{tr}_n}^2 \geq \|Z - Y\|_{\text{tr}_n}^2$. Since $\|e^{itA}Ze^{-itA}\|_{\text{tr}_n}^2 = \|Z\|_{\text{tr}_n}^2$, it follows that $\langle e^{itA}Ze^{-itA}, Y \rangle_{\text{tr}_n}$ is minimized at $t = 0$. Differentiating at $t = 0$ yields

$$\sum_{j=1}^m \text{tr}_n((iAZ_j - iZ_jA)Y_j) = \sum_{j=1}^m \text{tr}_n(Ai(Z_jY_j - Y_jZ_j)) = \text{tr}_n \left(A \sum_{j=1}^m i[Z_j, Y_j] \right).$$

Since this holds for all $A \in M_n(\mathbb{C})_{\text{sa}}^m$, it follows that $\sum_{j=1}^m [Z_j, Y_j] = 0$ as desired. \square

Remark 1.15. In the case $m = 1$, this lemma actually provides an alternative proof the spectral theorem as follows. Let $X \in M_n(\mathbb{C})_{\text{sa}}$. Let Y be a fixed diagonal matrix with distinct diagonal entries y_1, \dots, y_n . Let U be a unitary minimizing $\|UXU^* - Y\|_{\text{tr}_n}$. Then $[UXU^*, Y] = 0$. Any matrix A that commutes with Y must satisfy $a_{i,j}y_j = y_ia_{i,j}$, and hence A must be diagonal. Therefore, UXU^* is diagonal.²

²One might object that the preceding lemma seems to assume the spectral theorem already because it uses functional calculus to define e^{itA} . However, this only requires analytic functional calculus, not continuous functional calculus. One

Next, we describe an analog of the Monge-Kantorovich duality for the setting of matrix tuples.

Lemma 1.16. *Let $X, Y \in M_n(\mathbb{C})_{\text{sa}}^m$. Then (X, Y) is an optimal coupling in $M_n(\mathbb{C})$ if and only if there exist functions $f, g : M_n(\mathbb{C})_{\text{sa}}^m \rightarrow \mathbb{R}$ satisfying the following properties:*

- (1) f and g are convex.
- (2) f and g are unitarily invariant, that is, $f(UX'U^*) = f(X')$ and $g(UY'U^*) = g(Y')$ for $X', Y' \in M_n(\mathbb{C})_{\text{sa}}^m$.
- (3) $f(X') + g(Y') \geq \langle X', Y' \rangle_{\text{tr}_n}$ for all $X', Y' \in M_n(\mathbb{C})_{\text{sa}}^m$.
- (4) $f(X) + g(Y) = \langle X, Y \rangle_{\text{tr}_n}$.

Proof. (\implies). Let $\mathcal{U}(M_n(\mathbb{C}))$ be the unitary group. Let

$$f(X') = \sup_{U \in \mathcal{U}(M_n(\mathbb{C}))} \langle X', UYU^* \rangle_{\text{tr}_n}.$$

Note that f is convex because it is the supremum of a family of affine functions. Moreover, f is unitarily invariant because we took the supremum over all unitaries U .

Let g be the Legendre transform of f , that is,

$$g(Y') = \sup_{X' \in M_n(\mathbb{C})_{\text{sa}}^m} (\langle Y', X' \rangle_{\text{tr}_n} - f(X')).$$

It is immediate that g is convex, g is unitarily invariant because f is unitarily invariant and the inner product is unitarily invariant, and $f(X') + g(Y') \geq \langle X', Y' \rangle_{\text{tr}_n}$ for all $X', Y' \in M_n(\mathbb{C})_{\text{sa}}^m$. In particular, $f(X) + g(Y) \geq \langle X, Y \rangle_{\text{tr}_n}$.

On the other hand, note that the supremum defining $f(X)$ is achieved when $U = 1$ because we assumed that (X, Y) is optimal coupling, hence $\langle X, UYU^* \rangle$ is maximized when $U = 1$. Hence, $f(X) = \langle X, Y \rangle$. Moreover,

$$f(X') \geq \langle X', Y \rangle_{\text{tr}_n},$$

hence

$$g(Y) \leq \sup_{X' \in M_n(\mathbb{C})_{\text{sa}}^m} (\langle X', Y \rangle_{\text{tr}_n} - \langle X', Y \rangle_{\text{tr}_n}) = 0.$$

Thus, $f(X) + g(Y) \leq \langle X, Y \rangle_{\text{tr}_n}$. Hence, $f(X) + g(Y) = \langle X, Y \rangle$ as desired.

(\impliedby) Suppose that f and g satisfy (1)–(4). Let U be a unitary. Then

$$\langle U XU^*, Y \rangle_{\text{tr}_n} \leq f(U XU^*) + g(Y) = f(X) + g(Y) = \langle X, Y \rangle_{\text{tr}_n}.$$

Therefore, (X, Y) is optimal. \square

Unitarily invariant convex functions on $M_n(\mathbb{C})_{\text{sa}}^m$ satisfy a monotonicity property with respect to the non-commutative condition expectation from $M_n(\mathbb{C})$ onto a $*$ -subalgebra A , which is one motivation for our notion of E -convexity in the tracial W^* -setting.

Lemma 1.17. *Let A be a $*$ -subalgebra of $M_m(\mathbb{C})$, and let $E : M_n(\mathbb{C}) \rightarrow A \subseteq M_n(\mathbb{C})$ be the orthogonal projection with respect to the inner product $\langle S, T \rangle_{\text{tr}_n} = \text{tr}_n(S^*T)$. Then $E[ST] = SE[T]$ and $E[TS] = E[T]S$ and $E[T^*] = E[T]^*$ for $T \in M_n(\mathbb{C})$ and $S \in A$. Moreover, if $f : M_n(\mathbb{C})_{\text{sa}}^m \rightarrow \mathbb{R}$ is a convex function that is invariant under unitary conjugation, then for $X = (X_1, \dots, X_m) \in M_n(\mathbb{C})_{\text{sa}}^m$, we have*

$$f(E[X]) \leq f(X).$$

Here $E[X] = (E[X_1], \dots, E[X_m])$.

Proof. First, let us show that E is an A - A -bimodule map, that is, $E[ST] = SE[T]$ and $E[TS] = E[T]S$ for all $T \in M_n(\mathbb{C})$ and $S \in A$. For the second equality, note that $E[T]S \in A$ and for all $Z \in A$, we have

$$\text{tr}_n(Z^*(E[T]S)) = \text{tr}_n(SZ^*E[T]) = \text{tr}_n(SZ^*T) = \text{tr}_n(Z^*(TS)).$$

The proof of the first equality is similar, even easier. To show that $E[T^*] = E[T]^*$, observe that for all $Z \in A$, we have

$$\text{tr}_n(Z^*E[T]^*) = \overline{\text{tr}_n(E[T]Z)} = \overline{\text{tr}_n(ZE[T])} = \overline{\text{tr}_n(ZT)} = \text{tr}_n(T^*Z^*) = \text{tr}_n(Z^*T^*).$$

can use power series to define e^{itA} , show that $e^{i(s+t)A} = e^{isA}e^{itA}$ for $s, t \in \mathbb{R}$, show that $(e^{itA})^* = e^{-itA^*}$, and hence conclude that e^{itA} is unitary when A is self-adjoint.

Therefore, the orthogonal projection of T^* onto \mathcal{A} is $E[T]^*$ as desired.

For a subalgebra $A \subseteq M_n(\mathbb{C})$, we denote by $\mathcal{U}(A)$ the group of unitary matrices that are contained in A . We define the *commutant*

$$A' = \{S \in M_n(\mathbb{C}) \mid [S, T] = 0 \text{ for all } T \in A\}.$$

We recall that $A'' = A$ by von Neumann's bicommutant theorem [65, Theorem II.3.9].

Let μ be the Haar measure on $\mathcal{U}(A')$, and define

$$X' = \int_{\mathcal{U}(A')} UXU^* d\mu(U).$$

We claim that $X' = E[X]$. First, by invariance of Haar measure, $VX'V^* = X'$ for $V \in \mathcal{U}(A')$, hence $[X', V] = 0$. Since A' is a $*$ -algebra, it is linearly spanned by its unitaries, and therefore, $[X', S] = 0$ for all $S \in A'$. So $X' \in A'' = A$. Furthermore, for all $T \in A$, we have

$$\mathrm{tr}_n(T^*X') = \int_{\mathcal{U}(A')} \mathrm{tr}_n(T^*UXU^*) d\mu(U) = \int_{\mathcal{U}(A')} \mathrm{tr}_n(UT^*XU^*) d\mu(U) = \mathrm{tr}_n(T^*X).$$

Thus, $X' = E[X]$ as desired. Since μ is a probability measure, Jensen's inequality and the unitary invariance of f imply that

$$f(E[X]) \leq \int_{\mathcal{U}(A')} f(UXU^*) d\mu(U) = f(X). \quad \square$$

2. BACKGROUND ON TRACIAL W^* -ALGEBRAS

For the sake of readers who are less familiar tracial W^* -algebras, we explain the prerequisites needed for the paper: the definition of a tracial W^* -algebra, its interpretation as a non-commutative generalization of probability spaces, inclusions and trace-preserving conditional expectations of tracial W^* -algebras, free products with amalgamation, and non-commutative laws.

2.1. Tracial W^* -algebras. Historically, von Neumann algebras and W^* -algebras were defined differently, but it turns out that these two definitions give the same objects thanks to work of Sakai; see e.g. [63, Theorem 1.16.7]. Here we follow Sakai's approach that starts with the definition of W^* -algebras as C^* -algebras which are dual Banach spaces [63]. Other background references on von Neumann algebras include [2, 65, 66, 67].

Definition 2.1. A *unital $*$ -algebra* is a (unital) algebra A over \mathbb{C} together with a skew-linear involution $a \mapsto a^*$ such that $(ab)^* = b^*a^*$. If A and B are $*$ -algebras, then a map $\rho : A \rightarrow B$ is said to be a *$*$ -homomorphism* if it is linear and respects multiplication and the $*$ -operation.

Definition 2.2. A *unital C^* -algebra* is a $*$ -algebra A equipped with a norm $\|\cdot\|$ such that

- A is a Banach space with respect to $\|\cdot\|$;
- $\|ab\| \leq \|a\|\|b\|$ for $a, b \in A$;
- $\|a^*a\| = \|a\|^2$ for $a \in A$.

Definition 2.3. A *W^* -algebra* is a C^* -algebra A together with a topology \mathcal{T} , such that A as a Banach space is the dual of some Banach space A_* and \mathcal{T} is the weak- $*$ topology on A .

We remark that A_* can be uniquely recovered from (A, \mathcal{T}) as the subspace of A^{**} consisting of linear functionals that are continuous with respect to \mathcal{T} . In fact, it turns out that the predual of A_* of a W^* -algebra A is uniquely determined by A alone without reference to its weak- $*$ topology [63, Corollary 1.13.3].

Definition 2.4. If A is a W^* -algebra and A_* is a predual of A , then a *faithful normal trace on A* is an element $\tau \in A_*$ satisfying the following properties:

- $\tau(\mathbf{1}) = 1$;
- $\tau(a^*a) \geq 0$ for $a \in A$;
- $\tau(a^*a) = 0$ if and only if $a = 0$;
- $\tau(ab) = \tau(ba)$ for $a, b \in A$.

We remark that in general von Neumann algebra theory, the word “trace” is often used to refer to the semi-finite trace on a semi-finite von Neumann algebra, but in this paper “trace” always means “tracial state.”

Definition 2.5. A *tracial W^* -algebra* is a pair $\mathcal{A} = (A, \tau)$, where A is a W^* -algebra and τ is a faithful normal trace.

Example 2.6. Let (Ω, P) be a probability space. We take $A = L^\infty(\Omega, P)$, with the pointwise addition and multiplication operations. The $*$ -operation is pointwise complex conjugation. The norm is the standard one for $L^\infty(\Omega, P)$, and note that $\|fg\| \leq \|f\|\|g\|$ and $\|f^*f\| = \|f\|^2$. By the Riesz representation theorem, $L^\infty(\Omega, P) = L^1(\Omega, P)^*$, and therefore, we can take $A_* = L^1(\Omega, P)$, and then equip $L^\infty(\Omega, P)$ with the corresponding weak- $*$ topology. We define τ using the element $1 \in L^1(\Omega, P)$, so that $\tau(f) = \int_\Omega f dP$. Since $L^\infty(\Omega, P)$ is commutative, it is immediate that $\tau(fg) = \tau(gf)$. The other properties of τ are straightforward to check from well-known facts in measure theory. Conversely, it turns out that every commutative tracial W^* -algebra is isomorphic to L^∞ of some probability space [63, §1.18], [65, Theorem 1.18].

Example 2.7. Let H be an infinite-dimensional Hilbert space, and let $A = B(H)$ be the algebra of bounded operators on H equipped with the operator norm. Let A_* be the space of trace class operators. Then A can be canonically identified with the dual of A_* by the pairing $(a, T) = \text{Tr}(aT)$ for $a \in A$ and $T \in A_*$. The weak- $*$ topology on $B(H)$ is also known as the σ -weak operator topology. Thus, $B(H)$ is a W^* -algebra. However, it is not a tracial W^* -algebra because Tr is not well-defined on all of $B(H)$ and $\text{Tr}(1) = \infty$. See for instance [63, Theorem 1.15.3].

Theorem 2.8 (GNS construction for tracial W^* -algebras). *Let $\mathcal{A} = (A, \tau)$ be a tracial W^* -algebra. Note that $\langle a, b \rangle_{\mathcal{A}} := \tau(a^*b)$ defines an inner product on A (which is non-degenerate because τ is faithful). This can be completed to a Hilbert space, which we denote by $L^2(\mathcal{A})$. Let us denote the map $A \rightarrow L^2(\mathcal{A})$ by $a \mapsto \widehat{a}$. Then for each $a \in A$, there is a unique operator $\pi_\ell(a), \pi_r(a) \in B(L^2(\mathcal{A}))$ such that $\pi_\ell(a)\widehat{b} = \widehat{ab}$ and $\pi_r(a)\widehat{b} = \widehat{ba}$ for $b \in A$. Moreover, π_ℓ defines a $*$ -homomorphism $A \rightarrow B(L^2(\mathcal{A}))$ which is continuous with respect to the weak- $*$ topologies on A and $B(L^2(\mathcal{A}))$. Similarly, π_r is a $*$ -anti-homomorphism (it preserves $+$ and $*$ but reverses the order of multiplication) that is weak- $*$ continuous. Furthermore, since $\|a^*\|_{L^2(\mathcal{A})} = \|a\|_{L^2(\mathcal{A})}$, there is a unique skew-linear isometry $J : L^2(\mathcal{A}) \rightarrow L^2(\mathcal{A})$ such that $J(\widehat{a}) = \widehat{a^*}$. See [51, §IV] and [2, §7].*

Example 2.9. Let $A = L^\infty(\Omega, P)$ and let τ be integration against P . Then $\langle f, g \rangle_{L^2(\mathcal{A})} = \int_\Omega \overline{f}g dP$. The completion $L^2(\mathcal{A})$ can be canonically identified with $L^2(\Omega, P)$. The map $\widehat{\cdot}$ is the standard inclusion $L^\infty(\Omega, P) \rightarrow L^2(\Omega, P)$. The operator $\pi(f) \in B(L^2(\Omega, P))$ is the operator of multiplication by f .

Remark 2.10. Our examples indicate that if $\mathcal{A} = (A, \tau)$ is a tracial W^* -algebra, then A is an analog of $L^\infty(\Omega, P)$, A_* is an analog of $L^1(\Omega, P)$ and $L^2(\mathcal{A})$ is an analog of $L^2(\Omega, P)$. In fact, there is an even a non-commutative analog of measurable functions on Ω that are finite almost everywhere; this is known as the algebra $\text{Aff}(\mathcal{A})$ of operators *affiliated to \mathcal{A}* , certain closed unbounded operators on the Hilbert space $L^2(\mathcal{A})$. The space $L^2(\mathcal{A})$ can be canonically identified with a subspace of the affiliated operators. Thus, the left and right multiplication operators $\pi_\ell(a)$ and $\pi_r(a)$ for $a \in A$ become instances of multiplying affiliated operators. Moreover, there are subspaces $L^p(\mathcal{A}) \subseteq \text{Aff}(\mathcal{A})$ for $p \in [1, \infty)$ which share many properties of the classical L^p spaces. There is also a natural identification of A_* with $L^1(\mathcal{A})$. See §A.1 and the references therein for details.

2.2. W^* -embeddings, trace-preserving conditional expectations, and W^* -isomorphisms.

Notation 2.11. If $\mathcal{A} = (A, \tau)$ is a tracial W^* -algebra, we will use the notation $L^\infty(\mathcal{A})$ for A and $\tau_{\mathcal{A}}$ for τ when it is convenient to avoid naming A and τ explicitly. In particular, the norm on A will be denoted $\|\cdot\|_{L^\infty(\mathcal{A})}$. Furthermore, we will treat $L^\infty(\mathcal{A})$ as a subspace of $L^2(\mathcal{A})$. We will also write ab rather than $\pi_\ell(a)\widehat{b}$ and ba rather than $\pi_r(a)\widehat{b}$ for $a \in L^\infty(\mathcal{A})$ and $b \in L^2(\mathcal{A})$. Finally, we write a^* instead of $J(a)$ for $a \in L^2(\mathcal{A})$. We denote by $L^2(\mathcal{A})_{\text{sa}}$ the real subspace of $L^2(\mathcal{A})$ consisting of those elements fixed by J .

Definition 2.12. Let \mathcal{A} and \mathcal{B} be tracial W^* -algebras. A linear map $\phi : L^\infty(\mathcal{A}) \rightarrow L^\infty(\mathcal{B})$ is said to be *trace-preserving* if $\tau_{\mathcal{A}} = \tau_{\mathcal{B}} \circ \phi$.

Lemma 2.13 (See [12, Lemma 1.5.11] and [2, §9.1]). *Let \mathcal{A} and \mathcal{B} be tracial W^* -algebras. Let $\phi : L^\infty(\mathcal{A}) \rightarrow L^\infty(\mathcal{B})$ be a trace-preserving unital $*$ -homomorphism. Then*

- (1) ϕ extends to an isometry $L^2(\mathcal{A}) \rightarrow L^2(\mathcal{B})$, and in particular ϕ is injective on \mathcal{A} .
- (2) ϕ is a contraction $L^\infty(\mathcal{A}) \rightarrow L^\infty(\mathcal{B})$.
- (3) The adjoint map $E = \phi^* : L^2(\mathcal{B}) \rightarrow L^2(\mathcal{A})$ restricts to a map $L^\infty(\mathcal{B}) \rightarrow L^\infty(\mathcal{A})$ that is contractive with respect to the L^∞ norm.
- (4) We have $E[b^*] = E[b]^*$ for $b \in L^\infty(\mathcal{B})$, and in fact also for $b \in L^2(\mathcal{B})$.
- (5) E is a bimodule map over $L^\infty(\mathcal{A})$, that is, for $a \in L^\infty(\mathcal{A})$ and $b \in L^2(\mathcal{B})$, we have $E(\phi(a)b) = aE(b)$ and $E(b\phi(a)) = E(b)a$.
- (6) E is unital ($E(1) = 1$) and trace-preserving ($\tau_{\mathcal{A}} \circ E = \tau_{\mathcal{B}}$).

Definition 2.14. In the situation of the previous lemma, we call ϕ a (tracial W^*)-embedding $\mathcal{A} \rightarrow \mathcal{B}$ and E the associated trace-preserving conditional expectation. (Note that both maps are unital by definition and the previous proposition.)

Remark 2.15. It turns out that a trace-preserving $*$ -homomorphism $L^\infty(\mathcal{A}) \rightarrow L^\infty(\mathcal{B})$ is automatically continuous with respect to the weak- $*$ topology, essentially because the weak- $*$ topology can be recovered from the action of $L^\infty(\mathcal{A})$ on $L^2(\mathcal{A})$ by Theorem 2.8; see [24] or [2, Proposition 2.6.4]. For similar reasons, the trace-preserving conditional expectation is also weak- $*$ continuous.

Example 2.16. Suppose that $\mathcal{B} = L^\infty(\Omega, \mathcal{F}, P)$ for some probability space (Ω, \mathcal{F}, P) , where \mathcal{F} is the σ -algebra associated to the measure. Let \mathcal{G} be a σ -subalgebra of \mathcal{F} . Then there is an expectation-preserving inclusion $L^\infty(\Omega, \mathcal{G}, P) \rightarrow L^\infty(\Omega, \mathcal{F}, P)$. This extends to a map on the L^2 spaces, and the adjoint of this map is the conditional expectation $E : L^2(\Omega, \mathcal{F}, P) \rightarrow L^2(\Omega, \mathcal{G}, P)$ sending X to $E[X|\mathcal{G}]$. The properties in Lemma 2.13 then reduce to the well-known classical properties of conditional expectation. For instance, (2) the conditional expectation is contractive on L^∞ , (3) The conditional expectation respects complex conjugation, (4) if $X \in L^2(\Omega, \mathcal{F}, P)$ and $Y \in L^\infty(\Omega, \mathcal{G}, P)$, then $E[XY|\mathcal{G}] = E[X|\mathcal{G}]Y$, (5) the conditional expectation is expectation-preserving: $E[E[X|\mathcal{G}]] = E[X]$.

Notation 2.17. If \mathcal{A} and \mathcal{B} are tracial W^* -algebras, we say that $\mathcal{A} \subseteq \mathcal{B}$ if $L^\infty(\mathcal{A}) \subseteq L^\infty(\mathcal{B})$, the addition, product, $*$ -operation and weak- $*$ topology for $L^\infty(\mathcal{A})$ are the restrictions of those from $L^\infty(\mathcal{B})$, and $\tau_{\mathcal{A}} = \tau_{\mathcal{B}}|_{L^\infty(\mathcal{A})}$. In this case, we denote the conditional expectation $\mathcal{B} \rightarrow \mathcal{A}$ by $E_{\mathcal{A}}$.

As the paper will often deal with m -tuples of self-adjoint elements of L^2 , we

Notation 2.18. If \mathcal{A} and \mathcal{B} are tracial W^* -algebras and $\phi : L^\infty(\mathcal{A}) \rightarrow L^\infty(\mathcal{B})$ is a tracial W^* -embedding or a trace-preserving conditional expectation, then we will use the same letter ϕ to denote the extension of the map to the L^2 spaces. Furthermore, if $X = (X_1, \dots, X_m) \in L^2(\mathcal{A})_{\text{sa}}^m$, then we will write $\phi(X) = (\phi(X_1), \dots, \phi(X_m))$.

Definition 2.19. A tracial W^* -embedding $\phi : \mathcal{A} \rightarrow \mathcal{B}$ is said to be a *tracial W^* -isomorphism* if it is bijective and the inverse map is also a tracial W^* -embedding.

For reasons of mathematical logic, the class of tracial W^* -algebras is not a set. However, it will be convenient for us in §3.2 to have a *set* of isomorphism class representatives of tracial W^* -algebras with separable predual.

Lemma 2.20. *There exists a set \mathbb{W} of tracial W^* -algebras, such that*

- (1) *the elements of \mathbb{W} are pairwise non-isomorphic,*
- (2) *every tracial W^* -algebra with separable predual is isomorphic to some element of \mathbb{W} .*

Sketch of proof. We saw earlier that if $\mathcal{A} = (A, \tau)$ is a tracial W^* -algebra with separable predual, then there is a W^* -embedding $A \rightarrow B(H_{\mathcal{A}})$. Also, $H_{\mathcal{A}} \cong L^2(\mathcal{A})$ is separable and hence isomorphic as a Hilbert space to $\ell^2(\mathbb{N})$. Therefore, A is isomorphic to some W^* -subalgebra of $B(\ell^2(\mathbb{N}))$. Let S_1 be the set of W^* -subalgebras of $B(\ell^2(\mathbb{N}))$ (which is a subset of the power set of $B(\ell^2(\mathbb{N}))$). Let S_2 be the set of pairs $\{(A, \tau) : A \in S_1, \tau : A \rightarrow \mathbb{C} \text{ faithful normal trace}\}$. If $(A, \tau) \in S_2$, then the adjoint of the inclusion map produces a map from the space $B(\ell^2(\mathbb{N}))_*$ of trace class operators to $A_* \cong L^1(A, \tau)$, and hence A_* is separable. Thus, S_2 is a set of tracial W^* -algebras such that every tracial W^* -algebra with separable

predual is isomorphic to some element of S_2 . Finally, observe that tracial W^* -isomorphism defines an equivalence relation on S_2 , and let S_3 be the set of equivalence classes. \square

2.3. Amalgamated free products. Next, we explain the definition of free independence with amalgamation. This is an analog of conditional independence in classical probability theory. For background see for instance [78] or [12, §4.7].

Definition 2.21. Let $\mathcal{A} = (A, \tau)$ be a tracial W^* -algebra. Let B, A_1, \dots, A_N be W^* -algebras of A with $B \subseteq A_j$ for every j . Let $\mathcal{B} = (B, \tau|_B)$ and let $E_B : \mathcal{A} \rightarrow \mathcal{B}$ be the trace-preserving conditional expectation. We say that A_1, \dots, A_N are *freely independent with amalgamation over \mathcal{B}* if the following condition holds: Whenever $\ell \in \mathbb{N}$ and $i_1, \dots, i_\ell \in \{1, \dots, N\}$ with $i_1 \neq i_2, i_2 \neq i_3, \dots, i_{\ell-1} \neq i_\ell$ and $a_j \in A_{i_j}$ with $E_B[a_j] = 0$ $j = 1, \dots, \ell$, then $E_B[a_1 \dots a_\ell] = 0$.

Proposition 2.22. Let $\mathcal{B} = (B, \sigma)$ be a tracial W^* -algebra. For $j = 1, \dots, N$, let $\mathcal{A}_j = (A_j, \tau_j)$ be a tracial W^* -algebra and let $\iota_j : \mathcal{B} \rightarrow \mathcal{A}_j$ be a tracial W^* -embedding. Then there exists a tracial W^* -algebra $\mathcal{A} = (A, \tau)$ and tracial W^* -embeddings $\iota : \mathcal{B} \rightarrow \mathcal{A}$ and $\phi_j : \mathcal{A}_j \rightarrow \mathcal{A}$ such that $\iota = \phi_j \circ \iota_j$ for all j , and such that $\phi_1(A_1), \dots, \phi_N(A_N)$ are freely independent in \mathcal{A} with amalgamation over $\iota(\mathcal{B})$. Moreover, $(\mathcal{A}, \tau, \iota, \phi_1, \dots, \phi_j)$ are unique up to a canonical isomorphism; in other words, if $(\tilde{\mathcal{A}}, \tilde{\tau}, \tilde{\iota}, \tilde{\phi}_1, \dots, \tilde{\phi}_N)$ are another such tuple, then there is a unique tracial W^* -isomorphism $\pi : \mathcal{A} \rightarrow \tilde{\mathcal{A}}$ satisfying $\pi \circ \iota = \tilde{\iota}$ and $\pi \circ \phi_j = \tilde{\phi}_j$ for all j .

Definition 2.23. If $\mathcal{B}, \mathcal{A}_1, \dots, \mathcal{A}_N$, and \mathcal{A} are as above (with the specified maps $\iota, \phi_1, \dots, \phi_N$), then we say that \mathcal{A} is a *free product of $\mathcal{A}_1, \dots, \mathcal{A}_N$ with amalgamation over $\iota_1(\mathcal{B}), \dots, \iota_N(\mathcal{B})$* .

In the case where $\mathcal{B} = \mathbb{C}$, we refer to these concepts simply as *free independence* and *free products*.

2.4. Non-commutative laws and generators. Next, we describe the space of non-commutative laws. A non-commutative law is the analog of a linear functional $\mathbb{C}\langle x_1, \dots, x_m \rangle \rightarrow \mathbb{R}$ given by $f \mapsto \int f d\mu$ for some compactly supported measure on \mathbb{R}^m . Instead of $\mathbb{C}\langle x_1, \dots, x_m \rangle$, we use the non-commutative polynomial algebra in d variables.

Definition 2.24 (Non-commutative polynomial algebra). We denote by $\mathbb{C}\langle x_1, \dots, x_m \rangle$ the universal unital algebra generated by variables x_1, \dots, x_m . As a vector space, $\mathbb{C}\langle x_1, \dots, x_m \rangle$ has a basis consisting of all products $x_{i_1} \dots x_{i_\ell}$ for $\ell \geq 0$ and $i_1, \dots, i_\ell \in \{1, \dots, d\}$. We equip $\mathbb{C}\langle x_1, \dots, x_m \rangle$ with the unique $*$ -operation such that $x_j^* = x_j$; more explicitly, the $*$ -operation is defined on monomials by $(x_{i_1} \dots x_{i_\ell})^* = x_{i_\ell}^* \dots x_{i_1}^*$.

Definition 2.25 (Non-commutative law). A linear functional $\lambda : \mathbb{C}\langle x_1, \dots, x_m \rangle \rightarrow \mathbb{C}$ is said to be *exponentially bounded* if there exists $R > 0$ such that $|\lambda(x_{i_1} \dots x_{i_\ell})| \leq R^\ell$ for all $\ell \in \mathbb{N}_0$ and $i_1, \dots, i_\ell \in \{1, \dots, d\}$, and in this case we say R is an *exponential bound* for λ . A *non-commutative law* is a unital, positive, tracial, exponentially bounded linear functional $\lambda : \mathbb{C}\langle x_1, \dots, x_m \rangle \rightarrow \mathbb{C}$. We denote the space of non-commutative laws by Σ_m , and we equip it with the weak- $*$ topology (that is, the topology of pointwise convergence on $\mathbb{C}\langle x_1, \dots, x_m \rangle$). We denote by $\Sigma_{m,R}$ the subset of Σ_m comprised of non-commutative laws with exponential bound R .

Observation 2.26. The space $\Sigma_{m,R}$ is convex, compact, and metrizable.

Observation 2.27. Let \mathcal{A} be a $*$ -algebra and $X = (X_1, \dots, X_m) \in L^\infty(\mathcal{A})_{\text{sa}}^m$. Then there is a unique $*$ -homomorphism $\pi_X : \mathbb{C}\langle x_1, \dots, x_m \rangle \rightarrow \mathcal{A}$ such that $\pi_X(x_j) = X_j$ for $j = 1, \dots, m$.

Definition 2.28 (Non-commutative law of an m -tuple). Let \mathcal{A} be a tracial W^* -algebra. Let $X = (X_1, \dots, X_m) \in L^\infty(\mathcal{A})_{\text{sa}}^m$. Then we define $\lambda_X : \mathbb{C}\langle x_1, \dots, x_m \rangle \rightarrow \mathbb{C}$ by $\lambda_X = \tau \circ \rho_X$.

Notation 2.29. If \mathcal{A} is a tracial W^* -algebra and $X \in L^\infty(\mathcal{A})^m$, we write

$$\|X\|_{L^\infty(\mathcal{A})^m} := \max(\|X_j\|_{L^\infty(\mathcal{A})} : j = 1, \dots, m).$$

Observation 2.30. If \mathcal{A} and X are as above, then λ_X is a non-commutative law with exponential bound $\|X\|_\infty$. Conversely, if R is an exponential bound for λ_X , then

$$\|X\|_{L^\infty(\mathcal{A})_{\text{sa}}^m} = \max_j \lim_{m \rightarrow \infty} \tau(X_j^{2n})^{1/2n} \leq R.$$

Hence, $\|X\|_\infty$ is the smallest exponential bound for λ_X and in particular it is uniquely determined by λ_X .

In the case of a single operator X , we can apply the spectral theorem to show that there is a unique probability measure μ_X on \mathbb{R} satisfying

$$\int_{\mathbb{R}} f d\mu_X = \tau(f(X)) \text{ for } f \in C_0(\mathbb{R}).$$

Since X is bounded, μ_X is compactly supported and thus makes sense to evaluate on polynomials. If p is a polynomial, then $\lambda_X[p] = \int_{\mathbb{R}} p d\mu_X$. Thus, λ_X is simply the linear functional on polynomials corresponding to the spectral distribution.

We use the notation λ_X in particular when $\mathcal{A} = M_n(\mathbb{C})$. We denote by tr_n the normalized trace $(1/n) \text{Tr}$ on $M_n(\mathbb{C})$; recall that this is the unique (unital) trace on $M_n(\mathbb{C})$. Thus, for any $X \in M_n(\mathbb{C})_{\text{sa}}^m$, a non-commutative law λ_X is unambiguously specified by the previous definition. In the $m = 1$ case, the non-commutative law is given by the empirical spectral distribution. Note that when X is a random m -tuple of matrices, we will use the notation λ_X by default to refer to the empirical non-commutative law, that is, the (random) non-commutative law of X with respect to tr_n .

The next proposition shows that any non-commutative law can be realized by a self-adjoint m -tuple in some tracial W^* -algebra. This is a version of the *Gelfand-Naimark-Segal construction* (or *GNS construction*). A proof can be found in [3, Proposition 5.2.14(d)].

Proposition 2.31 (GNS construction for non-commutative laws). *Let $\lambda \in \Sigma_{m,R}$. Then we may define a semi-inner product on $\mathbb{C}\langle x_1, \dots, x_m \rangle$ by*

$$\langle p, q \rangle_\lambda = \lambda(p^*q).$$

*Let H_λ be the separation-completion of $\mathbb{C}\langle x_1, \dots, x_m \rangle$ with respect to this inner product, that is, the completion of $\mathbb{C}\langle x_1, \dots, x_m \rangle / \{p : \lambda(p^*p) = 0\}$, and let $[p]$ denote the equivalence class of a polynomial p in H_λ .*

There is a unique unital $$ -homomorphism $\pi : \mathbb{C}\langle x_1, \dots, x_m \rangle \rightarrow B(H_\lambda)$ satisfying $\rho(p)[q] = [pq]$ for $p, q \in \mathbb{C}\langle x_1, \dots, x_m \rangle$. Moreover, $\|\pi(x_j)\| \leq R$.*

Let $X_j = \pi(x_j)$, let $X = (X_1, \dots, X_m)$ and let A be the W^ -subalgebra of $B(H_\lambda)$ generated by X_1, \dots, X_m . Define $\tau : A \rightarrow \mathbb{C}$ by $\tau(Y) = \langle [1], Y[1] \rangle_\lambda$. Then τ is a faithful normal trace on A , and hence $\mathcal{A} = (A, \tau)$ is a tracial W^* -algebra.*

Definition 2.32. In the situation of the previous proposition, we call (\mathcal{A}, X) the *GNS realization* of λ .

The tracial W^* -algebra associated to λ is canonical in the sense that any other construction would yield an isomorphic tracial W^* -algebra. The following lemma can be deduced from the well-known properties of the GNS representation associated to a faithful trace τ on a W^* -algebra \mathcal{A} (which gives the so-called standard form of a tracial W^* -algebra).

Lemma 2.33. *Let \mathcal{A} and \mathcal{B} be tracial W^* -algebras. Let $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$ and $Y \in L^\infty(\mathcal{B})_{\text{sa}}^m$ such that $\lambda_X = \lambda_Y$. Let $W^*(X)$ and $W^*(Y)$ be the W^* -subalgebras of \mathcal{A} and \mathcal{B} generated by X and Y respectively. Then there is a unique tracial W^* -isomorphism $\rho : W^*(X) \rightarrow W^*(Y)$ such that $\rho(X_j) = Y_j$.*

Here is a related lemma about generating sets for a tracial W^* -algebra, which relies on the Kaplansky density theorem [65, Theorem II.4.8].

Lemma 2.34. *Let \mathcal{A} be a tracial W^* -algebra. Let $S \subseteq L^\infty(\mathcal{A})$. Let $W^*(S)$ be the smallest W^* -subalgebra of \mathcal{A} containing S , which is equal to the weak- $*$ closure of the unital $*$ -algebra generated by S . Then every $Z \in W^*(S)$ can be approximated in the $L^2(\mathcal{A})$ norm by a sequence Z_n in the unital $*$ -algebra generated by S such that $\|Z_n\|_{L^\infty(\mathcal{A})} \leq \|Z\|_{L^\infty(\mathcal{A})}$. Furthermore, if $\phi : \mathcal{A} \rightarrow \mathcal{B}$ is a W^* -embedding, then $\phi|_{W^*(S)}$ is uniquely determined by $\phi|_S$.*

In fact, the notion of generators for a W^* -algebra extends to elements of $L^2(\mathcal{A})$. For instance, for a self-adjoint tuple $X \in L^2(\mathcal{A})_{\text{sa}}^m$, using the theory of affiliated operators sketched in §A.1, it is valid to apply a bounded Borel function f to X_j through functional calculus, and $f(X_j)$ will be an element of \mathcal{A} . Thus, we may define $W^*(X)$ as (for instance) the W^* -subalgebra generated by $\arctan(X_1), \dots, \arctan(X_m)$, and then, as one would hope, X turns out to be in $L^2(W^*(X))_{\text{sa}}^m$. See also [12, p. 482-483]. We can state a characterization of $W^*(X)$ without reference to affiliated operators as follows.

Lemma 2.35. *Let \mathcal{A} be a tracial W^* -algebra and $X \in L^2(\mathcal{A})_{\text{sa}}^m$. Then there exists a unique smallest subalgebra \mathcal{B} (denoted by $W^*(X)$) of \mathcal{A} such that $X \in L^2(\mathcal{B})_{\text{sa}}^m$.*

3. DUALITY FOR L^2 OPTIMAL COUPLINGS

Our goal is to prove a version of the Monge-Kantorovich duality for the non-commutative version of the L^2 Wasserstein distance defined by Biane and Voiculescu [11]. In §3.1 we recall the definitions of optimal couplings that were stated more succinctly in the introduction. We define E -convex functions in §3.2 and the corresponding Legendre transform in §3.3. Then we prove the non-commutative Monge-Kantorovich duality in §3.4, and as an application we prove a decomposition result for optimal couplings in §3.5.

3.1. Wasserstein distance and optimal couplings.

Definition 3.1 (Biane-Voiculescu [11, §1.1]). Let $\mu, \nu \in \Sigma_m$ be non-commutative laws. A *coupling* of μ and ν is a triple (\mathcal{A}, X, Y) where \mathcal{A} is a tracial W^* -algebra and $X, Y \in L^\infty(\mathcal{A})_{\text{sa}}^m$ such that $\lambda_X = \mu$ and $\lambda_Y = \nu$. For $\mu, \nu \in \Sigma_m$, the (*non-commutative L^2*) *Wasserstein distance* $d_W^{(2)}(\mu, \nu)$ is the infimum of $\|X - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}$ over all couplings (\mathcal{A}, X, Y) for $\mathcal{A} \in \mathbb{W}$.

It is shown in [11, Theorem 1.3] that $d_W^{(2)}$ is a metric on the set Σ_m , and for each $R > 0$, $\Sigma_{m,R}$ is complete in this metric. However, as shown in §5.4, the topology generated by $d_W^{(2)}$ is strictly stronger than the weak- $*$ topology on Σ_m . The notion of optimal couplings corresponding to the Wasserstein distance is as follows.

Definition 3.2. A coupling (\mathcal{A}, X, Y) of two non-commutative laws μ and ν is *optimal* if $\|X - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m} = d_W^{(2)}(\mu, \nu)$.

Remark 3.3. As remarked in [11], for every $\mu, \nu \in \Sigma_m$, some optimal coupling exists. To see this, suppose $R > 0$ is an exponential bound for μ and ν . Note that if (\mathcal{A}, X, Y) is a coupling and γ is the joint law of (X, Y) , then $\|X - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m} = \left(\sum_{j=1}^m \gamma((x_j - x_{m+j})^2) \right)^{1/2}$. The space of joint laws $\gamma \in \Sigma_{2m,R}$ with marginals μ and ν is closed in $\Sigma_{2m,R}$ and therefore compact, and $\gamma \mapsto \left(\sum_{j=1}^m \gamma((x_j - x_{m+j})^2) \right)^{1/2}$ is continuous. Thus, it achieves a minimum at some γ^* , and we obtain an optimal coupling (\mathcal{A}, X, Y) from the GNS construction with γ^* (Proposition 2.31).

Just as in classical optimal transport theory, it is convenient to frame L^2 optimal couplings in terms of inner products rather than L^2 norms in order to relate them with Legendre transforms. If (\mathcal{A}, X, Y) is a coupling of μ and ν , then

$$\|X - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 = \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - 2\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \|Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2.$$

Since $\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2$ and $\|Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2$ are uniquely determined by μ and ν , a coupling minimizes $\|X - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}$ if and only if it maximizes the inner product $\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$. This motivates the following definition.

Definition 3.4. For $\mu, \nu \in \Sigma_m$, we denote by $C(\mu, \nu)$ the maximal value of $\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$ over all probabilistic couplings (\mathcal{A}, X, Y) of μ and ν .

The preceding paragraph shows that

$$d_W^{(2)}(\mu, \nu)^2 = \sum_{j=1}^m \mu(x_j^2) + \sum_{j=1}^m \nu(x_j^2) - 2C(\mu, \nu).$$

The goal of the section is to establish a duality result that $C(\mu, \nu)$ is the infimum of $\mu(f) + \nu(g)$ over certain pairs (f, g) of E -convex functions.

3.2. E -convex functions. Fix a set \mathbb{W} of isomorphism class representatives for tracial W^* -algebras with separable predual, as was given by Lemma 2.20.

Definition 3.5. Let S be a set. A *tracial W^* -function with values in S* is tuple $f = (f^{\mathcal{A}})_{\mathcal{A} \in \mathbb{W}}$, where $f^{\mathcal{A}} : L^2(\mathcal{A})_{\text{sa}}^m \rightarrow S$ if whenever $\iota : \mathcal{A} \rightarrow \mathcal{B}$ is a tracial W^* -embedding, we have $f^{\mathcal{A}} = f^{\mathcal{B}} \circ \iota$. (Here the inclusion ι is understood to extend to a map $L^2(\mathcal{A})_{\text{sa}}^m \rightarrow L^2(\mathcal{B})_{\text{sa}}^m$ per Notation 2.18.)

Thus, roughly speaking, being a W^* -function means that the evaluation of f on some $X \in L^2(\mathcal{A})_{\text{sa}}^m$ is independent of the ambient algebra. Hence, in particular, for bounded operators, $f^{\mathcal{A}}(X)$ only depends on the non-commutative law of X .

Although the definition of f only specifies $f^{\mathcal{A}}$ when \mathcal{A} is in the set \mathbb{W} , it will sometimes be convenient to use the notation $f^{\mathcal{A}}$ for a general tracial W^* -algebra \mathcal{A} with separable predual. This is defined by fixing an isomorphism ϕ from \mathcal{A} to some $\mathcal{B} \in \mathbb{W}$ and letting $f^{\mathcal{A}} = f^{\mathcal{B}} \circ \phi$. This is independent of the choice of ϕ because $f^{\mathcal{B}} \circ \psi = f^{\mathcal{B}}$ for every automorphism ψ of \mathcal{B} ; this in turn follows from the definition of W^* -functions since an automorphism ψ is in particular an inclusion from \mathcal{B} into \mathcal{B} .

Definition 3.6. A tracial W^* -function $f = (f^{\mathcal{A}})_{\mathcal{A} \in \mathbb{W}}$ with values in $[-\infty, +\infty]$ is said to be *E -convex* if either it is identically equal to $-\infty$ or the following conditions hold:

- (1) For each \mathcal{A} , $f^{\mathcal{A}}$ is a convex and lower semi-continuous function $L^2(\mathcal{A})_{\text{sa}}^m \rightarrow (-\infty, +\infty]$.
- (2) If $\iota : \mathcal{A} \rightarrow \mathcal{B}$ is a trace-preserving embedding, and if $E = \iota^* : \mathcal{B} \rightarrow \mathcal{A}$ is the corresponding trace-preserving conditional expectation, then

$$f^{\mathcal{A}}(E[X]) \leq f^{\mathcal{B}}(X)$$

for $X \in \mathcal{B}_{\text{sa}}^m$. (Here E is understood to extend to a map $L^2(\mathcal{B})_{\text{sa}}^m \rightarrow L^2(\mathcal{A})_{\text{sa}}^m$ per Notation 2.18.)

Example 3.7. For $t \in (0, \infty)$, let $q_t^{\mathcal{A}}(X) = (1/2t)\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2$. Then q_t is E -convex. Indeed, it is convex because of the Cauchy-Schwarz and arithmetic-geometric mean inequalities. It is clearly continuous. Finally, it satisfies monotonicity under conditional expectation because conditional expectations are contractive in $\|\cdot\|_{L^2(\mathcal{A})_{\text{sa}}^m}$.

We next explain an equivalent characterization of E -convexity using subgradient vectors.

Definition 3.8. If H is a real Hilbert space and let $f : H \rightarrow (-\infty, \infty]$ is a function, we say that $y \in H$ is a *subgradient* for f at x if

$$f(x') \geq f(x) + \langle y, x - x' \rangle \text{ for all } x' \in H.$$

We define the *subdifferential* $\partial f(x)$ as the set of subgradient vectors at x .

The following facts are well-known in convex analysis.

Lemma 3.9. *Let H be a Hilbert space. If $f : H \rightarrow [-\infty, \infty]$ is convex and lower semi-continuous and $f(x)$ is finite, then $\partial f(x)$ is nonempty, closed, and convex. Conversely, $f : H \rightarrow (-\infty, \infty)$ and ∂f is nonempty for every x , then f is convex.*

Analogously, we will show that E -convex W^* -functions are characterized by the existence of a subgradient vector Y to $f^{\mathcal{A}}$ at X such that $Y \in L^2(W^*(X))_{\text{sa}}^m$ (where $W^*(X)$ is given by Lemma 2.35). In addition, we handle the case where f can take the value $+\infty$.

Lemma 3.10. *Let f be a W^* -function taking values in $(-\infty, \infty)$. Then f is E -convex if and only if for each $\mathcal{A} \in \mathbb{W}$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$, there exists $Y \in L^2(W^*(X))_{\text{sa}}^m$ which is a subgradient vector to $f^{\mathcal{A}}$ at X . Here $W^*(X)$ is given by Lemma 2.35.*

Proof. First, suppose that f is E -convex. Fix $X \in L^2(\mathcal{A})_{\text{sa}}^m$. By a corollary of the Hahn-Banach theorem, there exists some subgradient vector Z to $f^{\mathcal{A}}$ at X . Let $\mathcal{B} = W^*(X)$, let $E : \mathcal{A} \rightarrow \mathcal{B}$ be the trace-preserving conditional expectation, and let $Y = E[Z]$. Then for $X' \in L^2(\mathcal{A})_{\text{sa}}^m$, we have

$$\begin{aligned} f^{\mathcal{A}}(X') &\geq f^{\mathcal{B}}(E_{\mathcal{B}}[X']) = f^{\mathcal{A}}(E[X']) \\ &\geq f^{\mathcal{A}}(X) + \langle Z, E[X'] - X \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \\ &= f^{\mathcal{A}}(X) + \langle Z, E[X' - X] \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \\ &= f^{\mathcal{A}}(X) + \langle Y, X' - X \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}. \end{aligned}$$

Thus, the desired subgradient condition holds.

Conversely, suppose this subgradient condition holds. Lower semi-continuity of $f^{\mathcal{A}}$ follows from the existence of subgradient vectors. For $X_0, X_1 \in L^2(\mathcal{A})_{\text{sa}}^m$ and $t \in (0, 1)$, we have

$$f^{\mathcal{A}}((1-t)X_0 + tX_1) \leq (1-t)f^{\mathcal{A}}(X_0) + tf^{\mathcal{A}}(X_1).$$

because of the existence of a subgradient vector at $(1-t)X_0 + tX_1$. To check the monotonicity under conditional expectation, consider an embedding $\iota : \mathcal{A} \rightarrow \mathcal{B}$ and let $E : \mathcal{B} \rightarrow \mathcal{A}$ be the corresponding conditional expectation. Let $X \in L^2(\mathcal{B})_{\text{sa}}^m$ and let $X' = E[X]$. By (1), there is a subgradient vector Y to $f^{\mathcal{B}}$ at the point X' that is in $L^2(W^*(X'))_{\text{sa}}^m$, and in particular $Y \in L^2(\mathcal{A})_{\text{sa}}^m$. But then

$$\begin{aligned} f^{\mathcal{B}}(X) &\geq f^{\mathcal{B}}(X') + \langle Y, X - X' \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} \\ &= f^{\mathcal{B}}(E[X]) + \langle Y, X - E[X] \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} \\ &= f^{\mathcal{B}}(E[X]). \end{aligned}$$

□

Remark 3.11. The same argument shows that for a W^* -function taking values in $(-\infty, +\infty]$, E -convexity is equivalent to the combination of the following three conditions:

- (1) For each $\mathcal{A} \in \mathbb{W}$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$, if $f^{\mathcal{A}}(X) < \infty$, then there exists $Y \in L^2(W^*(X))$ which is a subgradient vector to $f^{\mathcal{A}}$ at X .
- (2) For each \mathcal{A} , the set $(f^{\mathcal{A}})^{-1}((-\infty, M])$ is closed and convex in $L^2(\mathcal{A})_{\text{sa}}^m$.
- (3) If $\iota : \mathcal{A} \rightarrow \mathcal{B}$ is a tracial W^* -embedding and $E = \iota^* : \mathcal{B} \rightarrow \mathcal{A}$ is the corresponding conditional expectation, then $f^{\mathcal{B}}(X) < +\infty$ implies $f^{\mathcal{A}}(E[X]) < +\infty$.

Remark 3.12. If f is a tracial W^* -function, then $f^{\mathcal{A}}(UXU^*) = f^{\mathcal{A}}(X)$ for every unitary U in $L^\infty(\mathcal{A})$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$; this is because conjugation by U defines an automorphism of \mathcal{A} (hence in particular a tracial W^* -embedding $\mathcal{A} \rightarrow \mathcal{A}$), and f respects tracial W^* -embeddings.

If f is E -convex, then this unitary invariance gives rise to a “sum of commutators” condition on subgradient vectors related to Lemma 1.14. More precisely, suppose f is E -convex, $Y \in \partial f^{\mathcal{A}}(X)$ and U is a unitary in $L^\infty(\mathcal{A})$. Then

$$f^{\mathcal{A}}(X) = f^{\mathcal{A}}(UXU^*) \geq f^{\mathcal{A}}(X) + \langle UXU^* - X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}.$$

As in Lemma 1.14, by taking $U = e^{itA}$ for $A \in L^\infty(\mathcal{A})_{\text{sa}}$ and differentiating at $t = 0$, we obtain $\sum_{j=1}^m [X_j, Y_j] = 0$.

The next lemma describes how the subdifferential interacts with conditional expectations.

Lemma 3.13. *Let f be an E -convex W^* -function. Let $\mathcal{A} \in \mathbb{W}$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$. Let \mathcal{B} be a tracial W^* -subalgebra of \mathcal{A} .*

- (1) *If $f^{\mathcal{B}}(E_{\mathcal{B}}[X]) = f^{\mathcal{A}}(X)$, then*

$$\partial f^{\mathcal{B}}(E_{\mathcal{B}}[X]) = L^2(\mathcal{B})_{\text{sa}}^m \cap \partial f^{\mathcal{A}}(X).$$

- (2) *If $Y \in \partial f^{\mathcal{A}}(X)$, then $E_{W^*(X)}[Y] \in \partial f^{\mathcal{A}}(X)$.*

Proof. (1) First, we show that $\partial f^{\mathcal{B}}(E_{\mathcal{B}}[X]) \subseteq L^2(\mathcal{B})_{\text{sa}}^m \cap \partial f^{\mathcal{A}}(X)$. If $Y \in \partial f^{\mathcal{B}}(E_{\mathcal{B}}[X])$, then clearly $Y \in L^2(\mathcal{B})_{\text{sa}}^m$. Moreover, for all $Z \in L^2(\mathcal{A})$, we have

$$\begin{aligned} f^{\mathcal{A}}(Z) &\geq f^{\mathcal{B}}(E_{\mathcal{B}}[Z]) \\ &\geq f^{\mathcal{B}}(E_{\mathcal{B}}[X]) + \langle Y, E_{\mathcal{B}}[Z] - E_{\mathcal{B}}[X] \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} \\ &= f^{\mathcal{A}}(X) + \langle Y, Z - X \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}, \end{aligned}$$

where we have used the fact that $E_{\mathcal{B}}$ is self-adjoint and $E_{\mathcal{B}}[Y] = Y$. Hence, $Y \in \partial f^{\mathcal{A}}(X)$ as desired.

Conversely, to show that $L^2(\mathcal{B})_{\text{sa}}^m \cap \partial f^{\mathcal{A}}(X) \subseteq \partial f^{\mathcal{B}}(E_{\mathcal{B}}[X])$, suppose that $Y \in L^2(\mathcal{B})_{\text{sa}}^m \cap \partial f^{\mathcal{A}}(X)$. Then for $Z \in L^2(\mathcal{B})_{\text{sa}}^m$,

$$\begin{aligned} f^{\mathcal{B}}(Z) &= f^{\mathcal{A}}(Z) \\ &\geq f^{\mathcal{A}}(X) + \langle Y, Z - X \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \\ &= f^{\mathcal{B}}(E_{\mathcal{B}}[X]) + \langle Y, Z - E_{\mathcal{B}}[X] \rangle_{L^2(\mathcal{B})_{\text{sa}}^m}, \end{aligned}$$

because $E_{\mathcal{B}}$ is a self-adjoint operator on $L^2(\mathcal{A})$ and $Y \in L^2(\mathcal{B})_{\text{sa}}^m$.

(2) Let $\mathcal{B} = W^*(X)$ (where the trace is given by the restriction of $\tau_{\mathcal{A}}$). Let $Z \in L^2(\mathcal{B})_{\text{sa}}^m$. Then

$$\begin{aligned} f^{\mathcal{B}}(Z) &= f^{\mathcal{A}}(Z) \\ &\geq f^{\mathcal{A}}(X) + \langle Y, Z - X \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \\ &= f^{\mathcal{B}}(X) + \langle E_{\mathcal{B}}[Y], Z - X \rangle_{L^2(\mathcal{B})_{\text{sa}}^m}. \end{aligned}$$

Thus, $E_{\mathcal{B}}[Y] \in \partial f^{\mathcal{B}}(X)$, and so by (1), $E_{\mathcal{B}}[Y] \in \partial f^{\mathcal{A}}(X)$. \square

Lemma 3.14. *Let f be an E -convex W^* -function. Let $\mathcal{A} \in \mathbb{W}$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$.*

- (1) *There exists a unique $Y \in \partial f^{\mathcal{A}}(X)$ of minimal L^2 -norm.*
- (2) *$Y \in L^2(W^*(X))_{\text{sa}}^m$.*
- (3) *Let $\mathcal{B} = W^*(Y)$ as described in Lemma 2.35. Then $f^{\mathcal{B}}(E_{\mathcal{B}}[X]) = f^{\mathcal{A}}(X)$ and $\mathcal{B} = W^*(E_{\mathcal{B}}[X])$.*

Proof. (1) Because $\partial f^{\mathcal{A}}(X)$ is a closed convex set, it has a unique element of minimal L^2 -norm.

(2) Let $\mathcal{C} = W^*(X)$. Let $Y' = E_{\mathcal{C}}[Y]$. We claim that $Y' \in \partial f^{\mathcal{C}}(X)$. Let $Z \in L^2(\mathcal{C})_{\text{sa}}^m$. Then

$$\begin{aligned} f^{\mathcal{C}}(Z) &= f^{\mathcal{A}}(Z) \\ &\geq f^{\mathcal{A}}(X) + \langle Y, Z - X \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \\ &= f^{\mathcal{C}}(X) + \langle Y', Z - X \rangle_{L^2(\mathcal{C})_{\text{sa}}^m}. \end{aligned}$$

Thus, $Y' \in \partial f^{\mathcal{C}}(X)$. By the previous lemma, $Y' \in \partial f^{\mathcal{A}}(X)$. But because Y has minimal norm, we have $\|E_{\mathcal{C}}[Y]\|_{L^2(\mathcal{A})_{\text{sa}}^m} = \|Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}$, hence $E_{\mathcal{C}}[Y] = Y$, so $Y \in L^2(\mathcal{C})_{\text{sa}}^m$.

(3) First, we show that $f^{\mathcal{B}}(E_{\mathcal{B}}[X]) = f^{\mathcal{A}}(X)$. By E -convexity, $f^{\mathcal{B}}(E_{\mathcal{B}}[X]) \leq f^{\mathcal{A}}(X)$. Conversely,

$$f^{\mathcal{B}}(E_{\mathcal{B}}[X]) = f^{\mathcal{A}}(E_{\mathcal{B}}[X]) \geq f^{\mathcal{A}}(X) + \langle Y, E_{\mathcal{B}}[X] - X \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = f^{\mathcal{A}}(X).$$

Let $\mathcal{D} = W^*(E_{\mathcal{B}}[X])$ and $\phi = \tau|_{\mathcal{D}}$. By Lemma 3.13 (2), $E_{\mathcal{D}}[Y] \in \partial f^{\mathcal{B}}(E_{\mathcal{B}}[X])$, hence also $E_{\mathcal{D}}[Y] \in \partial f^{\mathcal{A}}(X)$ by Lemma 3.13 (1). Because Y was chosen to have minimal norm, we have $E_{\mathcal{D}}[Y] = Y$, and thus, $\mathcal{D} \supseteq W^*(Y) = \mathcal{B}$ by the characterization of $W^*(Y)$ given in Lemma 2.35. Hence, $\mathcal{B} = \mathcal{D} = W^*(E_{\mathcal{B}}[X])$. \square

3.3. Legendre transforms.

Definition 3.15. We define the *Legendre transform* as the tuple $\mathcal{L}f = (\mathcal{L}f^{\mathcal{A}})_{\mathcal{A} \in \mathbb{W}}$ by

$$\mathcal{L}f^{\mathcal{A}}(X) = \sup\{\langle \iota(X), Y \rangle - f^{\mathcal{B}}(Y) : \mathcal{B} \in \mathbb{W}, \iota : \mathcal{A} \rightarrow \mathcal{B} \text{ a tracial } W^*\text{-embedding}\}.$$

Example 3.16. Consider again $q_t^{\mathcal{A}}(X) = (1/2t)\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2$. A standard computation with norms and inner products shows that $\mathcal{L}q_t = q_{1/t}$.

Proposition 3.17. *Let f be a tracial W^* -function f .*

- (1) *The Legendre transform $\mathcal{L}f$ is an E -convex tracial W^* -function.*
- (2) *If $f \leq g$, then $\mathcal{L}f \geq \mathcal{L}g$.*
- (3) *We have $\mathcal{L}^2 f \leq f$ with equality if and only if f is E -convex.*
- (4) *$\mathcal{L}^2 f$ is the maximal E -convex function that is less than or equal to f .*

Proof. (1) If f is identically equal to $-\infty$ or $+\infty$, $\mathcal{L}f$ will be $+\infty$ or $-\infty$ respectively and there is nothing to prove. Hence, assume that f attains some finite value at some $Y \in L^2(\mathcal{B})_{\text{sa}}^m$ for some $\mathcal{B} \in \mathbb{W}$.

For any $\mathcal{A} \in \mathbb{W}$, the free product $\mathcal{A} * \mathcal{B}$ is isomorphic to some $\mathcal{C} \in \mathbb{W}$. Let $\iota_1 : \mathcal{A} \rightarrow \mathcal{C}$ and $\iota_2 : \mathcal{B} \rightarrow \mathcal{C}$ be the corresponding tracial W^* -embeddings. Then

$$\mathcal{L}f^{\mathcal{A}}(X) \geq \langle \iota_1(X), \iota_2(Y) \rangle - f^{\mathcal{C}}(\iota_2(Y)) > -\infty,$$

since $f^{\mathcal{C}}(\iota_2(Y)) = f^{\mathcal{B}}(Y) > -\infty$. Hence, $\mathcal{L}f$ is never equal to $-\infty$.

For each \mathcal{A} , the function $\mathcal{L}f^{\mathcal{A}}$ is a supremum of affine functions, and therefore it is convex and lower semi-continuous.

Let $\iota : \mathcal{A} \rightarrow \mathcal{B}$ be a tracial W^* -embedding and let $E : \mathcal{B} \rightarrow \mathcal{A}$ be the corresponding trace-preserving conditional expectation. Let $X \in L^2(\mathcal{B})_{\text{sa}}^m$. Let $\tilde{\iota} : \mathcal{A} \rightarrow \tilde{\mathcal{B}}$ be another inclusion. Let $\mathcal{M} \in \mathbb{W}$ be isomorphic to the amalgamated free product of \mathcal{B} and $\tilde{\mathcal{B}}$ over the subalgebra \mathcal{A} (or more precisely, over the images of $\iota(\mathcal{A}) \subseteq \mathcal{B}$ and $\tilde{\iota}(\mathcal{A}) \subseteq \tilde{\mathcal{B}}$) as in Proposition 2.22. Let $\rho : \mathcal{B} \rightarrow \mathcal{M}$ and $\tilde{\rho} : \tilde{\mathcal{B}} \rightarrow \mathcal{M}$ be the inclusions. Then for $Y \in L^2(\tilde{\mathcal{B}})$,

$$\begin{aligned} \mathcal{L}f^{\mathcal{B}}(X) &\geq \langle \rho(X), \tilde{\rho}(Y) \rangle_{L^2(\mathcal{M})_{\text{sa}}^m} - f^{\mathcal{M}}(\tilde{\rho}(Y)) \\ &= \langle \tilde{\iota} \circ E(X), Y \rangle_{L^2(\tilde{\mathcal{B}})_{\text{sa}}^m} - f^{\tilde{\mathcal{B}}}(Y), \end{aligned}$$

where we have used free independence with amalgamation to compute the inner product, and we have used the fact that f is a tracial W^* -function. Because $\tilde{\iota} : \mathcal{A} \rightarrow \tilde{\mathcal{B}}$ and Y were arbitrary, we have

$$\mathcal{L}f^{\mathcal{B}}(X) \geq \mathcal{L}f^{\mathcal{A}}(E(X)),$$

which establishes condition (2) in the definition of E -convexity.

It only remains to show that f is a tracial W^* -function. Suppose $\iota : \mathcal{A} \rightarrow \mathcal{B}$ is a tracial W^* -inclusion. If $\iota' : \mathcal{B} \rightarrow \mathcal{C}$ is a tracial W^* -inclusion, then so is $\iota' \circ \iota$, which implies that

$$\mathcal{L}f^{\mathcal{A}}(X) \geq \sup_{\substack{\iota' : \mathcal{B} \rightarrow \mathcal{C} \\ Y \in L^2(\mathcal{C})_{\text{sa}}^m}} \langle \iota' \circ \iota(X), Y \rangle_{L^2(\mathcal{C})_{\text{sa}}^m} - f^{\mathcal{C}}(Y) = \mathcal{L}f^{\mathcal{B}}(\iota(X)).$$

If $E : \mathcal{B} \rightarrow \mathcal{A}$ is the conditional expectation corresponding to ι , then by the preceding argument

$$\mathcal{L}f^{\mathcal{A}}(X) = \mathcal{L}f^{\mathcal{A}}(E \circ \iota(X)) \leq \mathcal{L}f^{\mathcal{B}}(\iota(X)).$$

Thus, $\mathcal{L}f^{\mathcal{A}} = \mathcal{L}f^{\mathcal{B}} \circ \iota$, so $\mathcal{L}f$ is a tracial W^* -function.

(2) This is immediate from the definition and the properties of suprema and infima.

(3) By definition of $\mathcal{L}f$, for every $\mathcal{A} \in \mathbb{W}$ and $X, Y \in L^2(\mathcal{A})_{\text{sa}}^m$, we have

$$\mathcal{L}f^{\mathcal{A}}(X) \geq \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - f^{\mathcal{A}}(Y),$$

hence

$$\mathcal{L}f^{\mathcal{A}}(X) + f^{\mathcal{A}}(Y) \geq \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}.$$

Hence, given an inclusion ι of \mathcal{A} into \mathcal{B} and $Y \in L^2(\mathcal{A})_{\text{sa}}^m$ and $X \in L^2(\mathcal{B})_{\text{sa}}^m$, we have

$$f^{\mathcal{A}}(Y) = f^{\mathcal{B}}(\iota(Y)) \geq \langle \iota(Y), X \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - \mathcal{L}f^{\mathcal{B}}(X).$$

Taking the supremum on the right-hand side, $f^{\mathcal{A}}(Y) \geq \mathcal{L}^2 f^{\mathcal{A}}(Y)$. Thus, $f \geq \mathcal{L}^2 f$.

Now suppose that f is E -convex, and we must show that $f = \mathcal{L}^2 f$. If f is identically $-\infty$ or $+\infty$, there is nothing to prove. Otherwise, fix \mathcal{A} . Because $f^{\mathcal{A}}$ is convex and lower semi-continuous, classical results about convex functions tell us that $f^{\mathcal{A}}$ can be expressed as the supremum of a family of affine functions $(g_\alpha)_{\alpha \in I}$, where

$$g_\alpha(X) = \langle X, Z_\alpha \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + c_\alpha$$

with $Z_\alpha \in L^2(\mathcal{A})_{\text{sa}}^m$ and $c_\alpha \in \mathbb{R}$. Let $\iota : \mathcal{A} \rightarrow \mathcal{B}$ be an inclusion and let $E : \mathcal{B} \rightarrow \mathcal{A}$ be the corresponding conditional expectation. If $Y \in L^2(\mathcal{B})_{\text{sa}}^m$, then by the E -convexity property

$$\langle \iota(Z_\alpha), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - f^{\mathcal{B}}(Y) \leq \langle \iota(Z_\alpha), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - f^{\mathcal{A}, \sigma}(E(Y)) \leq \langle Z_\alpha, E(Y) \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - \langle Z_\alpha, E(Y) \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - c_\alpha = -c_\alpha.$$

Therefore, $\mathcal{L}f^{\mathcal{A}}(Z_\alpha) \leq -c_\alpha$, which implies that

$$\mathcal{L}^2 f^{\mathcal{A}}(X) \geq \langle X, Z_\alpha \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - \mathcal{L}f^{\mathcal{A}}(Z_\alpha) \geq \langle X, Z_\alpha \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + c_\alpha = g_\alpha(X).$$

Therefore,

$$\mathcal{L}^2 f^{\mathcal{A}}(X) \geq \sup_{\alpha \in I} g_\alpha(X) = f^{\mathcal{A}}(X).$$

So $\mathcal{L}^2 f = f$ as desired. Conversely, if $f = \mathcal{L}^2 f$, then f is E -convex because it is the Legendre transform of some function.

(4) We already showed that $\mathcal{L}^2 f$ is E -convex and $\mathcal{L}^2 f \leq f$. Moreover, if g is E -convex and $g \leq f$, then $\mathcal{L}g \geq \mathcal{L}f$ and hence $g = \mathcal{L}^2 g \leq \mathcal{L}^2 f$. \square

Remark 3.18. It follows from E -convexity that for every $\iota : \mathcal{A} \rightarrow \mathcal{B}$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$ and $Y \in L^2(\mathcal{B})_{\text{sa}}^m$, we have

$$\langle \iota(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - f^{\mathcal{B}}(Y) \leq \langle X, E[Y] \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - f^{\mathcal{A}}(E[Y]),$$

where $E : \mathcal{B} \rightarrow \mathcal{A}$ is the conditional expectation corresponding to ι . Therefore,

$$\mathcal{L}f^{\mathcal{A}}(X) = \sup_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} (\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - f^{\mathcal{A}}(Y)).$$

Hence, if f is E -convex, there is no need to consider a larger W^* -algebra when computing the Legendre transform, and moreover $\mathcal{L}f^{\mathcal{A}}$ agrees with the classical Legendre transform of $f^{\mathcal{A}}$ as a function on the real Hilbert space $L^2(\mathcal{A})_{\text{sa}}^m$.

The next lemma states the relationship between Legendre transforms and subgradients, which is exactly analogous to the behavior of classical Legendre transforms. We will use this lemma many times.

Lemma 3.19. *Let f be an E -convex W^* -function, let $\mathcal{A} \in \mathbb{W}$ and $X, Y \in L^2(\mathcal{A})_{\text{sa}}^m$. Then $f^{\mathcal{A}}(X) + \mathcal{L}f^{\mathcal{A}}(Y) = \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$ if and only if $Y \in \partial f^{\mathcal{A}}(X)$.*

Proof. First, suppose that $f^{\mathcal{A}}(X) + \mathcal{L}f^{\mathcal{A}}(Y) = \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$. By definition of $\mathcal{L}f$, we have for all $X' \in L^2(\mathcal{A})_{\text{sa}}^m$ that

$$\langle X', Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - f^{\mathcal{A}}(X') \leq \mathcal{L}f^{\mathcal{A}}(Y) = \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - f^{\mathcal{A}}(X),$$

hence, $f^{\mathcal{A}}(X') \geq f^{\mathcal{A}}(X) + \langle X' - X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$, so $Y \in \partial f^{\mathcal{A}}(X)$.

Conversely, if $Y \in \partial f^{\mathcal{A}}(X)$ and if $\iota : \mathcal{A} \rightarrow \mathcal{B}$ is a W^* -inclusion, then by Lemma 3.13, $\iota(Y) \in \partial f^{\mathcal{B}}(\iota(X))$. Hence, for any $Z \in L^2(\mathcal{B})_{\text{sa}}^m$, we have

$$\langle Z, \iota(Y) \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - f^{\mathcal{B}}(Z) \leq \langle \iota(X), \iota(Y) \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - f^{\mathcal{B}}(\iota(X)) = \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - f^{\mathcal{A}}(X).$$

Since ι , \mathcal{B} , and Z were arbitrary, the supremum defining $\mathcal{L}f^{\mathcal{A}}(Y)$ is attained at the point X , so that $f^{\mathcal{A}}(X) + \mathcal{L}f^{\mathcal{A}}(Y) = \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$. \square

3.4. A non-commutative Monge-Kantorovich duality.

Definition 3.20. If f is a tracial W^* -function and $\mu \in \Sigma_m$, then we define $\mu(f) = f^{\mathcal{A}}(X)$, where $\mathcal{A} \in \mathbb{W}$ is (isomorphic to) the GNS representation of μ and X is the canonical generating m -tuple.

If f is a tracial W^* -function, for every \mathcal{A} and every $X \in \mathcal{A}_{\text{sa}}^m$ with $\lambda_X = \mu$, we have $\mu(f) = f^{\mathcal{A}}(X)$. This follows by the definition of tracial W^* -function and the fact that $W^*(X)$ is isomorphic to the GNS representation of μ .

Definition 3.21. Let us call a pair (f, g) of tracial W^* -functions *admissible* if they take values in $(-\infty, \infty]$ and for every $\mathcal{A} \in \mathbb{W}$,

$$f^{\mathcal{A}}(X) + g^{\mathcal{A}}(Y) \geq \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \text{ for all } X, Y \in L^2(\mathcal{A})_{\text{sa}}^m.$$

Proposition 3.22. *Let $\mu, \nu \in \Sigma_m$. The following quantities are equal:*

- (1) $C(\mu, \nu)$.
- (2) $\inf\{\mu(f) + \nu(g) : (f, g) \text{ admissible}\}$.
- (3) $\inf\{\mu(f) + \nu(\mathcal{L}f) : f \text{ a tracial } W^*\text{-function not identically } \infty\}$.
- (4) $\inf\{\mu(f) + \nu(g) : (f, g) \text{ admissible and } E\text{-convex}\}$.
- (5) $\inf\{\mu(f) + \nu(\mathcal{L}f) : f \text{ } E\text{-convex not identically } \infty\}$.

Here all the functions under consideration take values in $(-\infty, \infty]$.

Proof. (1) \leq (2) Let (\mathcal{A}, X, Y) be a coupling of μ and ν , and let (f, g) be an admissible pair. Then

$$\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \leq f^{\mathcal{A}}(X) + g^{\mathcal{A}}(Y) = \mu(f) + \nu(g).$$

Taking the supremum over couplings on the left-hand side and the infimum over admissible pairs (f, g) on the right-hand side, we have (1) \leq (2).

(2) \leq (3). It is clear from the definition of $\mathcal{L}f$ that $f^{\mathcal{A}}(X) + \mathcal{L}f^{\mathcal{A}}(Y) \geq \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$. Therefore, $(f, \mathcal{L}f)$ is always an admissible pair, and hence (3) is the infimum over a smaller set than (2).

(3) \leq (1). Define

$$f^{\mathcal{A}}(X) = \begin{cases} 0, & \text{if } X \in L^\infty(\mathcal{A}) \text{ and } \lambda_X = \mu, \\ +\infty, & \text{otherwise.} \end{cases}$$

Note that f is a tracial W^* -function. Let \mathcal{A} be the GNS-representation of ν with the canonical generators Y . Then $\mathcal{L}f^{\mathcal{A}}(Y)$ is the supremum of $\langle \iota(Y), X \rangle_{L^2(\mathcal{B})_{\text{sa}}^m}$ where $\iota : \mathcal{A} \rightarrow \mathcal{B}$ is an inclusion and $X \in L^\infty(\mathcal{B})_{\text{sa}}^m$ satisfies $\lambda_X = \mu$. In particular for a non-commutative law ν , letting (\mathcal{A}, Y) be the GNS realization of ν , we have $\nu(\mathcal{L}f) = \mathcal{L}f^{\mathcal{A}}(Y) = C(\mu, \nu)$. Moreover, $\mu(f) = 0$ and hence $C(\mu, \nu) = \mu(f) + \nu(\mathcal{L}f)$.

(2) \leq (4). This is immediate since (4) is the infimum over a smaller set.

(4) \leq (5). Suppose that f is E -convex. Then $(f, \mathcal{L}f)$ is admissible as noted above. Also, $\mathcal{L}f$ is always E -convex, so (5) is the infimum over a smaller set than (4).

(5) \leq (3). Let f be a tracial W^* -function. Then $\mathcal{L}^2 f \leq f$ and $(\mathcal{L}^2 f, \mathcal{L}f)$ is an E -convex admissible pair. Therefore,

$$\mu(f) + \nu(\mathcal{L}f) \geq \mu(\mathcal{L}^2 f) + \nu(\mathcal{L}f).$$

Of course, since $\mathcal{L}(\mathcal{L}^2 f) = \mathcal{L}^2(\mathcal{L}f) = \mathcal{L}f$, the term on the right-hand side participates in the infimum (5). Since the f on the left-hand side was chosen arbitrarily, (3) \geq (5). \square

Proposition 3.23. *Let (\mathcal{A}, X, Y) be a coupling of μ and $\nu \in \Sigma_m$. The following are equivalent:*

- (1) *The coupling is optimal.*
- (2) *There exists an admissible pair (f, g) such that $\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = f^{\mathcal{A}}(X) + g^{\mathcal{A}}(Y)$.*
- (3) *There exists a tracial W^* -function f such that $\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = f^{\mathcal{A}}(X) + \mathcal{L}f^{\mathcal{A}}(Y)$.*
- (4) *There exists an admissible, E -convex pair (f, g) such that $\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = f^{\mathcal{A}}(X) + g^{\mathcal{A}}(Y)$.*
- (5) *There exists an E -convex f such that $\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = f^{\mathcal{A}}(X) + \mathcal{L}f^{\mathcal{A}}(Y)$.*
- (6) *There exists an E -convex W^* -function f such that Y is a subgradient vector to $f^{\mathcal{A}}$ at the point X .*

Proof. It is immediate from the previous proposition that each of the conditions (2) – (5) implies (1).

For the converse implication, assume the coupling is optimal. Let

$$f^{\mathcal{B}}(Z) = \begin{cases} 0, & \text{if } \lambda_Z = \mu, \\ +\infty, & \text{otherwise.} \end{cases}$$

As in the proof of the previous proposition, we have $\mu(f) + \nu(\mathcal{L}f) = C(\mu, \nu)$, or equivalently $\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = f^{\mathcal{A}}(X) + \mathcal{L}f^{\mathcal{A}}(Y)$. We also have $C(\mu, \nu) = \mu(\mathcal{L}^2 f) + \nu(\mathcal{L}f) \leq \mu(f) + \nu(\mathcal{L}f) = C(\mu, \nu)$. Thus, the pair $(\mathcal{L}^2 f, \mathcal{L}f)$ fulfills all of the criteria of (2) – (5).

The equivalence of (5) and (6) follows from Lemma 3.19. \square

3.5. A decomposition result for optimal couplings. As an initial application of duality, we present the following result that expresses an optimal coupling (X, Y) in terms of another optimal coupling (X', Y') with $\mathcal{B} = W^*(X') = W^*(Y')$.

Theorem 3.24. *Let $\mu, \nu \in \Sigma_{m, \mathbb{R}}$, and let (\mathcal{A}, X, Y) be an optimal coupling of μ and ν . Then there exists a subalgebra $\mathcal{B} \subseteq \mathcal{A}$ with the following properties, letting $X' = E_{\mathcal{B}}[X]$ and $Y' = \mathbb{E}_{\mathcal{B}}[Y]$:*

- (1) $\mathcal{B} = W^*(X') = W^*(Y')$.
- (2) $X - X', X' - Y'$, and $Y' - Y$ are orthogonal.
- (3) (\mathcal{A}, X', Y') is an optimal coupling of $\lambda_{X'}$ and $\lambda_{Y'}$. Similarly, (\mathcal{A}, X, Y') and (\mathcal{A}, X', Y) are optimal couplings of the respective laws.

We may choose \mathcal{B} to be contained in $W^*(X)$ (or symmetrically, we may choose it to be contained in $W^*(Y)$).

Furthermore, there exists some optimal coupling (\mathcal{A}, X, Y) and a \mathcal{B} satisfying (1) – (3) with respect to this coupling such that $W^*(X, \mathcal{B})$ and $W^*(Y, \mathcal{B})$ are freely independent with amalgamation over \mathcal{B} .

Proof. Let

$$\mathcal{B} = \{\mathcal{B} \subseteq \mathcal{A} : \langle E_{\mathcal{B}}[X], E_{\mathcal{B}}[Y] \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}\},$$

which is partially ordered by inclusion. We claim that \mathcal{B} has a minimal element. By Zorn's lemma, it suffices to show that every chain in \mathcal{B} has a lower bound. Consider a chain $\mathcal{C} \subseteq \mathcal{B}$, and let $\mathcal{C} = \bigcap_{\mathcal{B} \in \mathcal{C}} \mathcal{B}$. We claim that $\lim_{\mathcal{B} \in \mathcal{C}} E_{\mathcal{B}}[X] = E_{\mathcal{C}}[X]$. Let

$$\delta = \inf_{\mathcal{B} \in \mathcal{C}} \|E_{\mathcal{B}}[X]\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2.$$

Given $\epsilon > 0$, there exists $\mathcal{B}_0 \in \mathcal{C}$ such that $\|E_{\mathcal{B}_0}[X]\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 < \delta^2 + \epsilon^2$. Then for all $\mathcal{B} \in \mathcal{C}$ with $\mathcal{B} \subseteq \mathcal{B}_0$, we have

$$\|E_{\mathcal{B}}[X] - E_{\mathcal{B}_0}[X]\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 = \|E_{\mathcal{B}_0}[X]\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - \|E_{\mathcal{B}}[X]\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 \leq \delta^2 + \epsilon - \delta^2 = \epsilon^2.$$

This implies that $Z = \lim_{\mathcal{B} \in \mathcal{C}} E_{\mathcal{B}}[X]$ exists in L^2 . But clearly $Z \in \bigcap_{\mathcal{B} \in \mathcal{C}} \mathcal{B} = \mathcal{C}$, and $\langle Z, W \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = \langle Z, W \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$ for $W \in L^2(\mathcal{C})_{\text{sa}}^m$. Thus, $\lim_{\mathcal{B} \in \mathcal{C}} E_{\mathcal{B}}[X] = E_{\mathcal{C}}[X]$. By the same token $\lim_{\mathcal{B} \in \mathcal{C}} E_{\mathcal{B}}[Y] = E_{\mathcal{C}}[Y]$. Therefore,

$$\langle E_{\mathcal{C}}[X], E_{\mathcal{C}}[Y] \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = \lim_{\mathcal{B} \in \mathcal{C}} \langle E_{\mathcal{B}}[X], E_{\mathcal{B}}[Y] \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}.$$

Therefore, $\mathcal{C} \in \mathcal{B}$ as desired.

So by Zorn's lemma, \mathcal{B} has some minimal element, which we will call \mathcal{B} . Let $X' = E_{\mathcal{B}}[X]$ and $Y' = E_{\mathcal{B}}[Y]$. Now $W^*(X') \subseteq \mathcal{B}$ and we have

$$\langle X', E_{W^*(X')}[Y'] \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = \langle X', Y' \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}.$$

By minimality of \mathcal{B} , we have $\mathcal{B} = W^*(X')$, and similarly, $\mathcal{B} = W^*(Y')$. Hence, (1) holds.

To show that \mathcal{B} can be chosen inside $W^*(X)$, note that

$$\langle E_{W^*(X)}[X], E_{W^*(X)}[Y] \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = \langle X, E_{W^*(X)}[Y] \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}.$$

Thus, we can apply the same argument with \mathcal{B} replaced by the elements of \mathcal{B} contained inside $W^*(X)$.

To prove (2), since $X' = E_{\mathcal{B}}[X]$ is orthogonal to \mathcal{B} , it is immediate that $X - X'$ and $X' - Y'$ are orthogonal. Similarly, $X - X'$ and $Y' - Y$ are orthogonal. Next, note that

$$\begin{aligned} \|X - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 &= \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - 2\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \|Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 \\ &= (\|X - E_{\mathcal{B}}[X]\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + \|E_{\mathcal{B}}[X]\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2) - 2\langle E_{\mathcal{B}}[X], E_{\mathcal{B}}[Y] \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \|Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 \\ &= \|X - E_{\mathcal{B}}[X]\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + \left(\|E_{\mathcal{B}}[X]\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - 2\langle E_{\mathcal{B}}[X], Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \|Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 \right) \\ &= \|X - X'\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + \|X' - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2. \end{aligned}$$

Thus, $X - X'$ and $X' - Y$ satisfy the Pythagorean identity, so they are orthogonal. In particular, $X - X'$ is orthogonal to $Y' - Y = (Y' - X') - (Y - X')$.

To prove (3), by Proposition 3.23, there exists an admissible pair of E -convex W^* -functions f and g such that $f^{\mathcal{A}}(X) + g^{\mathcal{A}}(Y) = \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$. By construction of \mathcal{B} and by E -convexity,

$$\begin{aligned} f^{\mathcal{A}}(X') + g^{\mathcal{A}}(Y') &\geq \langle X', Y' \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \\ &= \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \\ &= f^{\mathcal{A}}(X) + g^{\mathcal{A}}(Y) \\ &\geq f^{\mathcal{A}}(X') + g^{\mathcal{A}}(Y'). \end{aligned}$$

This implies that (X', Y') is an optimal coupling. By similar reasoning, since $\langle X, Y' \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = \langle X', Y' \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}$ and $f^{\mathcal{A}}(X') \leq f^{\mathcal{A}}(X)$, we see that (X', Y) is an optimal coupling, and symmetrically (X, Y') is an optimal coupling.

Let \mathcal{A}_1 be a copy of $W^*(X, \mathcal{B})$ and let \mathcal{A}_2 be a copy of $W^*(Y, \mathcal{B})$. Let $\tilde{\mathcal{A}} = \mathcal{A}_1 *_{\mathcal{B}} \mathcal{A}_2$ be the amalgamated free product (with its canonical trace $\tilde{\tau}$). Let $\tilde{X}, \tilde{X}', \tilde{Y}$, and \tilde{Y}' be the images of the original variables in $\tilde{\mathcal{A}}$. Then using free independence

$$\begin{aligned} \|\tilde{X} - \tilde{Y}\|_{L^2(\tilde{\mathcal{A}})_{\text{sa}}^m}^2 &= \|\tilde{X} - \tilde{X}'\|_{L^2(\tilde{\mathcal{A}})_{\text{sa}}^m}^2 + \|\tilde{X}' - \tilde{Y}'\|_{L^2(\tilde{\mathcal{A}})_{\text{sa}}^m}^2 + \|\tilde{Y}' - \tilde{Y}\|_{L^2(\tilde{\mathcal{A}})_{\text{sa}}^m}^2 \\ &= \|X - X'\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + \|X' - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + \|Y' - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 \\ &= \|X - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2. \end{aligned}$$

Therefore, (\tilde{X}, \tilde{Y}) is also an optimal coupling of μ and ν . The subalgebra $\mathcal{B} \subseteq \tilde{\mathcal{A}}$ also satisfies

$$\langle E_{\mathcal{B}}[\tilde{X}], E_{\mathcal{B}}[\tilde{Y}] \rangle_{\tilde{\tau}} = \langle \tilde{X}, \tilde{Y} \rangle_{\tilde{\tau}},$$

and satisfies (1). Thus, the same arguments as above show that \mathcal{B} in $\tilde{\mathcal{A}}$ satisfies (2) and (3). \square

4. THE DISPLACEMENT INTERPOLATION

If (\mathcal{A}, X, Y) is an L^2 -optimal coupling of $\mu, \nu \in \Sigma_m$, then one can consider the displacement interpolation $X_t = (1-t)X + tY$ for $t \in [0, 1]$. As shown in Proposition A.22 the corresponding family of laws defines a geodesic in $(\Sigma_m, d_W^{(2)})$. In this section, we study how the displacement interpolation interacts with non-commutative Monge-Kantorovich duality and use this to prove Theorem 1.5.

By Proposition 3.23, there exists an E -convex function f such that $\langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = f^{\mathcal{A}}(X) + \mathcal{L}f^{\mathcal{A}}(Y)$, or equivalently $Y \in \partial f^{\mathcal{A}}(X)$. Letting q_t be the W^* -function $q_t^{\mathcal{A}}(X) = (1/2t)\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2$, we observe that $X_t \in \partial f_t^{\mathcal{A}}(X)$ where $f_t = (1-t)q_1 + tf$. Hence, $X \in \partial(\mathcal{L}f_t)^{\mathcal{A}}(X_t)$. In order to show that $X \in L^2(W^*(X_t))_{\text{sa}}^m$, we want to understand the regularity properties of $\mathcal{L}f_t$.

It is well-known that for a convex function f on a Hilbert space H , the Legendre transform of $f(x) + (t/2)\|x\|^2$ is given by the inf-convolution $g_t = \inf_{y \in H} [f^*(y) + (1/2t)\|x - y\|^2]$, where f^* is the Legendre transform of f . Furthermore, g_t has a Lipschitz gradient for every $t > 0$, and it satisfies the Hamilton-Jacobi equation

$$\frac{d}{dt}g_t = -\frac{1}{2}\|\nabla g_t\|^2.$$

This can be checked by hand, or deduced for instance from [6, §2, Theorem 1]; also relevant to Hamilton-Jacobi equations on Hilbert space are [7, 8, 18, 19, 49, 46].

In this section, we adapt the theory of inf-convolutions to the setting tracial W^* -functions. In §4.1, we define inf-convolutions of W^* -functions and prove their basic properties. In §4.2, we describe how inf-convolutions interact with E -convexity and semi-concavity. In §4.3, we conclude the proof of Theorem 1.5.

We emphasize that the novelty in our work is not in the form of the Hamilton-Jacobi equation but rather in the fact that we study variables from infinite-dimensional non-commutative algebras and want the function to be defined consistently with respect to inclusions of these algebras (that is, to be a tracial W^* -function). This means for instance that if f and g are tracial W^* -functions and $f \square g$ is their inf-convolution as defined below, then $(f \square g)^{\mathcal{A}}$ need not agree with the inf-convolution of $f^{\mathcal{A}}$ and $g^{\mathcal{A}}$ as functions on the Hilbert space $L^2(\mathcal{A})_{\text{sa}}^m$ (Remark 4.5); however, they do agree if f and g are E -convex (Lemma 4.6). Hence, a notion of viscosity solutions compatible with our theory of inf-convolutions will thus have to take into account the inclusions of one tracial W^* -algebra into another.

4.1. Inf-convolutions. We begin with the definition and basic properties of the inf-convolution.

Definition 4.1. Let f, g be two W^* -functions with values in $[-\infty, \infty]$. We define the *inf-convolution* $f \square g$ by

$$(f \square g)^{\mathcal{A}}(X) = \inf \{ f^{\mathcal{B}}(\iota(X) - Y) + g^{\mathcal{B}}(Y) \mid \iota : \mathcal{A} \rightarrow \mathcal{B} \text{ embedding, } Y \in L^2(\mathcal{B})_{\text{sa}}^m \}.$$

Lemma 4.2. *The object $f \square g$ is a W^* -function.*

Proof. Let $\iota : \mathcal{A} \rightarrow \mathcal{B}$ be an inclusion, and we first show that

$$(4.1) \quad (f \square g)^{\mathcal{A}}(X) \leq (f \square g)^{\mathcal{B}}(\iota(X)).$$

If $\iota' : \mathcal{B} \rightarrow \mathcal{C}$ is another inclusion and $Y \in L^2(\mathcal{B})_{\text{sa}}^m$ as in the definition of $(f \square g)^{\mathcal{B}}$, then of course $\iota' \circ \iota$ is an inclusion and which can be used in the definition of $(f \square g)^{\mathcal{A}}$. This shows (4.1).

Conversely, suppose that $\iota' : \mathcal{A} \rightarrow \mathcal{C}$ is an inclusion and $Y \in L^2(\mathcal{C})_{\text{sa}}^m$ as in the definition of $(f \square g)^{\mathcal{A}}$. Then let $\tilde{\mathcal{C}}$ be the free product of \mathcal{B} and \mathcal{C} with amalgamation over the images of \mathcal{A} in the respective algebras. Then the image of Y in $\tilde{\mathcal{C}}$ participates in the infimum defining $(f \square g)^{\mathcal{B}}(\iota(X))$ and hence $(f \square g)^{\mathcal{B}}(\iota(X)) \leq (f \square g)^{\mathcal{A}}(X)$. \square

Lemma 4.3. *The inf-convolution is commutative and associative, that is, if f, g, h are W^* -functions, then $f \square g = g \square f$ and $(f \square g) \square h = f \square (g \square h)$.*

Proof. We have

$$(f \square g)^{\mathcal{A}}(X) = \inf_{\iota : \mathcal{A} \rightarrow \mathcal{B}} \inf_{Y \in L^2(\mathcal{B})_{\text{sa}}^m} [f^{\mathcal{B}}(\iota(X) - Y) + g^{\mathcal{B}}(Y)].$$

We substitute $Z = \iota(X) - Y$ and thus obtain

$$\inf_{\iota : \mathcal{A} \rightarrow \mathcal{B}} \inf_{Y \in L^2(\mathcal{B})_{\text{sa}}^m} [f^{\mathcal{B}}(Z) + g^{\mathcal{B}}(\iota(X) - Z)] = (g \square f)^{\mathcal{A}}(X).$$

For associativity,

$$(4.2) \quad \begin{aligned} ((f \square g) \square h)^{\mathcal{A}}(X) &= \inf_{\substack{\iota_1 : \mathcal{A} \rightarrow \mathcal{B} \\ Y \in L^2(\mathcal{B})_{\text{sa}}^m}} ((f \square g)^{\mathcal{B}}(\iota_1(X) - Y) + h^{\mathcal{B}}(Y)) \\ &= \inf_{\substack{\iota_1 : \mathcal{A} \rightarrow \mathcal{B} \\ Y \in L^2(\mathcal{B})_{\text{sa}}^m}} \inf_{\substack{\iota_2 : \mathcal{B} \rightarrow \mathcal{C} \\ Z \in L^2(\mathcal{C})_{\text{sa}}^m}} (f^{\mathcal{C}}(\iota_2(\iota_1(X) - Y) - Z) + g^{\mathcal{C}}(Z) + h^{\mathcal{C}}(\iota_2(Y))). \end{aligned}$$

We claim that is equal to

$$(4.3) \quad \inf_{\substack{\iota : \mathcal{A} \rightarrow \mathcal{B} \\ Y, Z \in L^2(\mathcal{B})_{\text{sa}}^m}} (f^{\mathcal{C}}(\iota(X) - Y - Z) + g^{\mathcal{B}}(Z) + h^{\mathcal{B}}(Y)),$$

or in other words, in our earlier expression we can without loss of generality impose the condition that $\mathcal{C} = \mathcal{B}$ and $\iota_2 = \text{id}$. The reason is that if we allowed Z to come only from the smaller algebra \mathcal{C} , then the infimum could only increase, hence by shrinking \mathcal{C} to \mathcal{B} , (4.3) \geq (4.2). On the other hand, if in (4.2), we allowed Y to come from the larger algebra \mathcal{C} instead of \mathcal{B} , then the infimum could only decrease, and hence by enlarging \mathcal{B} to \mathcal{C} , we see that (4.2) \leq (4.3). Now the expression (4.3) is symmetric in g and h , and hence

$$(f \square g) \square h = (f \square h) \square g.$$

This relation, together with commutativity, implies the associativity relation since

$$(f \square g) \square h = (g \square f) \square h = (g \square h) \square f = f \square (g \square h). \quad \square$$

The relationship between inf-convolution and Legendre transform is exactly what one would expect based on the classical case.

Lemma 4.4. *Let f and g be W^* -functions. Then*

$$\mathcal{L}(f \square g) = \mathcal{L}f + \mathcal{L}g.$$

Proof. Observe that

$$\begin{aligned} \mathcal{L}(f \square g)^{\mathcal{A}}(X) &= \sup_{\substack{\iota_1 : \mathcal{A} \rightarrow \mathcal{B} \\ Y \in L^2(\mathcal{B})_{\text{sa}}^m}} (\langle \iota_1(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - (f \square g)^{\mathcal{B}}(Y)) \\ &= \sup_{\substack{\iota_1 : \mathcal{A} \rightarrow \mathcal{B} \\ Y \in L^2(\mathcal{B})_{\text{sa}}^m}} \left(\langle \iota_1(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - \inf_{\substack{\iota_2 : \mathcal{B} \rightarrow \mathcal{C} \\ Z \in L^2(\mathcal{C})_{\text{sa}}^m}} (f^{\mathcal{C}}(\iota_2(Y) - Z) + g^{\mathcal{C}}(Z)) \right), \end{aligned}$$

where we take the supremum over \mathcal{B} and $\mathcal{C} \in \mathbb{W}$ and inclusions $\iota_1 : \mathcal{A} \rightarrow \mathcal{B}$ and $\iota_2 : \mathcal{B} \rightarrow \mathcal{C}$ and $Y \in L^2(\mathcal{B})_{\text{sa}}^m$ and $Z \in L^2(\mathcal{C})_{\text{sa}}^m$. This can be rewritten as

$$\sup_{\substack{\iota_1: \mathcal{A} \rightarrow \mathcal{B} \\ Y \in L^2(\mathcal{B})_{\text{sa}}^m}} \sup_{\substack{\iota_2: \mathcal{B} \rightarrow \mathcal{C} \\ Z \in L^2(\mathcal{C})_{\text{sa}}^m}} (\langle \iota_2(\iota_1(X)), \iota_2(Y) \rangle_{L^2(\mathcal{C})_{\text{sa}}^m} - f^{\mathcal{C}}(\iota_2(Y) - Z) - g^{\mathcal{C}}(Z)).$$

We can assume without loss generality that $\mathcal{B} = \mathcal{C}$ and $\iota_2 = \text{id}$. Indeed, allowing Y to range over the larger space $L^2(\mathcal{C})_{\text{sa}}^m$ rather than $L^2(\mathcal{B})_{\text{sa}}^m$ would only increase the supremum, but on the other hand, restricting Z to the smaller space $L^2(\mathcal{B})_{\text{sa}}^m$ instead of $L^2(\mathcal{C})_{\text{sa}}^m$ would only decrease the supremum. Thus, taking $\mathcal{B} = \mathcal{C}$ and renaming ι_1 to ι , we obtain

$$\sup_{\iota: \mathcal{A} \rightarrow \mathcal{B}} \sup_{Y, Z \in L^2(\mathcal{B})_{\text{sa}}^m} (\langle \iota(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - f^{\mathcal{B}}(Y - Z) - g^{\mathcal{B}}(Z)).$$

Substituting $Z' = Y - Z$, we have

$$(4.4) \quad \sup_{\iota: \mathcal{A} \rightarrow \mathcal{B}} \sup_{Z, Z' \in L^2(\mathcal{B})_{\text{sa}}^m} (\langle \iota(X), Z + Z' \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - f^{\mathcal{B}}(Z') - g^{\mathcal{B}}(Z)) \\ = \sup_{\iota: \mathcal{A} \rightarrow \mathcal{B}} \sup_{Z, Z' \in L^2(\mathcal{B})_{\text{sa}}^m} (\langle \iota(X), Z' \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - f^{\mathcal{B}}(Z') + \langle \iota(X), Z \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - g^{\mathcal{B}}(Z)).$$

We want to show that this is equal to

$$(4.5) \quad \mathcal{L}f^{\mathcal{A}}(X) + \mathcal{L}g^{\mathcal{A}}(X) = \sup_{\iota_1: \mathcal{A} \rightarrow \mathcal{B}_1} \sup_{Z' \in L^2(\mathcal{B}_1)_{\text{sa}}^m} (\langle \iota_1(X), Z' \rangle_{L^2(\mathcal{B}_1)_{\text{sa}}^m} - f^{\mathcal{B}_1}(Z')) \\ + \sup_{\iota_2: \mathcal{A} \rightarrow \mathcal{B}_2} \sup_{Z \in L^2(\mathcal{B}_2)_{\text{sa}}^m} (\langle \iota_2(X), Z \rangle_{L^2(\mathcal{B}_2)_{\text{sa}}^m} - g^{\mathcal{B}_2}(Z)).$$

The only difference between the two expressions is that the latter allows $\iota_1 : \mathcal{A} \rightarrow \mathcal{B}_1$ and $\iota_2 : \mathcal{A} \rightarrow \mathcal{B}_2$ to be different, but the former takes them to be the same, and thus a priori the (4.4) \leq (4.5). However, in (4.5), for any given $\mathcal{B}_1, \mathcal{B}_2, \iota_1$ and ι_2 , let \mathcal{B} be the free product of \mathcal{B}_1 and \mathcal{B}_2 with amalgamation over the subalgebras $\iota_1(\mathcal{A})$ in the first factor and $\iota_2(\mathcal{A})$ in the second factor. Allowing Z' and Z to range over $L^2(\mathcal{B})_{\text{sa}}^m$ rather than $L^2(\mathcal{B}_1)_{\text{sa}}^m$ and $L^2(\mathcal{B}_2)_{\text{sa}}^m$ respectively only increases the suprema over Z and Z' , and hence (4.5) remains unchanged when we restrict to the case $\iota_1 = \iota_2$, so it equals (4.4). \square

Remark 4.5. Suppose f and g are tracial W^* -functions. Let $f^{\mathcal{A}} \square g^{\mathcal{A}}$ denote the classical inf-convolution of $f^{\mathcal{A}}$ and $g^{\mathcal{A}}$ as functions on the Hilbert space $L^2(\mathcal{A})_{\text{sa}}^m$. Then $(f \square g)^{\mathcal{A}} \leq f^{\mathcal{A}} \square g^{\mathcal{A}}$. However, the following example shows that two functions do not necessarily agree. Take $m = 2$, and $f^{\mathcal{A}}(X_1, X_2) = (1/2) \|(X_1, X_2)\|_{L^2(\mathcal{A})^2}^2$ and $g^{\mathcal{A}}(X_1, X_2) = \tau_{\mathcal{A}}([X_1, X_2]^2)$. The formula for g is to be understood in the sense of affiliated operators (see §A); since $i[X_1, X_2]$ is a self-adjoint affiliated operator, $-[X_1, X_2]^2$ is positive and hence $\tau_{\mathcal{A}}([X_1, X_2]^2)$ is well-defined in $[-\infty, 0]$; see Theorem A.3 (4). Then $g^{\mathbb{C}} = 0$ because \mathbb{C} is commutative, and hence also $f^{\mathbb{C}} \square g^{\mathbb{C}} = 0$. On the other hand, let $\iota : \mathbb{C} \rightarrow M_2(\mathbb{C})$ be the canonical inclusion, and let

$$Y_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad Y_2 = \begin{pmatrix} 0 & i \\ -i & 0 \end{pmatrix}, \quad [Y_1, Y_2] = \begin{pmatrix} -2i & 0 \\ 0 & 2i \end{pmatrix}.$$

Then for $x_1, x_2, t \in \mathbb{R}$,

$$(f \square g)^{\mathbb{C}}(x_1, x_2) \leq \frac{1}{2} \|\iota(x_1) - tY_1\|_{L^2(M_2(\mathbb{C}))}^2 + \frac{1}{2} \|\iota(x_2) - tY_2\|_{L^2(M_2(\mathbb{C}))}^2 + t^4 \tau_{M_2(\mathbb{C})}([Y_1, Y_2]^2) \\ = \frac{1}{2} \|\iota(x_1) - tY_1\|_{L^2(M_2(\mathbb{C}))}^2 + \frac{1}{2} \|\iota(x_2) - tY_2\|_{L^2(M_2(\mathbb{C}))}^2 - 4t^4.$$

The first two terms are quadratic in t , and thus, taking the infimum over $t \in \mathbb{R}$, we see that $(f \square g)^{\mathbb{C}} = -\infty < f^{\mathbb{C}} \square g^{\mathbb{C}}$.

4.2. Inf-convolutions and regularity of E -convex functions.

Lemma 4.6. *If f and g are E -convex W^* -functions with $f < \infty$, then $f \square g$ is E -convex. Moreover, for any E -convex f and g , we have*

$$(4.6) \quad (f \square g)^{\mathcal{A}}(X) = \inf_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} (f^{\mathcal{A}}(X - Y) + g^{\mathcal{A}}(Y)).$$

Proof. We prove the second claim first. Clearly,

$$(f \square g)^{\mathcal{A}}(X) \leq \inf_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} (f^{\mathcal{A}}(X - Y) + g^{\mathcal{A}}(Y)).$$

For the opposite inequality, suppose that $\iota : \mathcal{A} \rightarrow \mathcal{B}$ is an embedding and $Y \in L^2(\mathcal{B})_{\text{sa}}^m$. Let $E : \mathcal{B} \rightarrow \mathcal{A}$ be the conditional expectation. Then by E -convexity of f and g ,

$$f^{\mathcal{A}}(X - E[Y]) + g^{\mathcal{A}}(E[Y]) \leq f^{\mathcal{B}}(\iota(X) - Y) + g^{\mathcal{B}}(Y),$$

and hence

$$\inf_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} (f^{\mathcal{A}}(X - Y) + g^{\mathcal{A}}(Y)) \leq (f \square g)^{\mathcal{A}}(X).$$

Now let us show that $f \square g$ is E -convex when $f < \infty$. If g is identically ∞ , then $f \square g$ is identically ∞ , so there is nothing to prove. Suppose $g^{\mathcal{B}}(Y)$ is finite for some \mathcal{A} and Y . Then $(f \square g)^{\mathcal{A}}(X) < \infty$ everywhere because, letting \mathcal{C} be the free product of \mathcal{A} and \mathcal{B} and letting $\iota_1 : \mathcal{A} \rightarrow \mathcal{C}$ and $\iota_2 : \mathcal{B} \rightarrow \mathcal{C}$ be the corresponding inclusions,

$$(f \square g)^{\mathcal{A}}(X) \leq f^{\mathcal{C}}(\iota_1(X) - \iota_2(Y)) + g^{\mathcal{C}}(\iota_2(Y)) < \infty.$$

To prove convexity of $(f \square g)^{\mathcal{A}}$, let $X_0, X_1 \in L^2(\mathcal{A})_{\text{sa}}^m$, and let $X_t = (1 - t)X_0 + tX_1$ for $t \in (0, 1)$. If $Y_0, Y_1 \in L^2(\mathcal{A})_{\text{sa}}^m$ and if $Y_t = (1 - t)Y_0 + tY_1$, then

$$\begin{aligned} (f \square g)^{\mathcal{A}}(X_t) &\leq f^{\mathcal{A}}(X_t - Y_t) + g^{\mathcal{A}}(Y_t) \\ &\leq (1 - t)f^{\mathcal{A}}(X_0 - Y_0) + tf^{\mathcal{A}}(X_1 - Y_1) + (1 - t)g^{\mathcal{A}}(Y_0) + tg^{\mathcal{A}}(Y_1). \end{aligned}$$

Since Y_0 and Y_1 were arbitrary, we can take the infimum over Y_0 and Y_1 and apply (4.6) to conclude that

$$(f \square g)^{\mathcal{A}}(X_t) \leq (1 - t)(f \square g)^{\mathcal{A}}(X_0) + t(f \square g)^{\mathcal{A}}(X_1).$$

This shows that $(f \square g)^{\mathcal{A}}$ is convex. Furthermore, since $f \square g < \infty$, this relation implies that if $(f \square g)^{\mathcal{A}}$ is $-\infty$ at one point in $L^2(\mathcal{A})_{\text{sa}}^m$, then it is $-\infty$ everywhere. Moreover, if $(f \square g)^{\mathcal{B}}$ is $-\infty$, then so $(f \square g)^{\mathcal{A}}$, as we can see by considering the free product of \mathcal{A} and \mathcal{B} .

It is automatic from these facts that $(f \square g)^{\mathcal{A}}$ is lower semi-continuous, since convexity automatically implies lower semi-continuity at points where $(f \square g)^{\mathcal{A}} < \infty$.

Finally, we must show the monotonicity of $(f \square g)$ under conditional expectation. Let $\iota : \mathcal{A} \rightarrow \mathcal{B}$ be an embedding and let $E : \mathcal{B} \rightarrow \mathcal{A}$ be the corresponding conditional expectation. If $X, Y \in L^2(\mathcal{A})_{\text{sa}}^m$, then

$$(f \square g)^{\mathcal{A}}(E[X]) \leq f^{\mathcal{A}}(E[X] - E[Y]) + g^{\mathcal{A}}(E[Y]) \leq f^{\mathcal{B}}(X - Y) + g^{\mathcal{B}}(Y).$$

Since Y on right-hand side was arbitrary, we conclude by (4.6) that $(f \square g)^{\mathcal{A}}(E[X]) \leq (f \square g)^{\mathcal{B}}(X)$ as desired. \square

Observation 4.7. *For $t \in (0, \infty)$, let $q_t^{\mathcal{A}}(X) = (1/2t)\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2$. For $s, t \in (0, \infty)$, because q_s and q_t are E -convex and take finite values, we have*

$$q_s \square q_t = \mathcal{L}^2(q_s \square q_t) = \mathcal{L}(\mathcal{L}q_s + \mathcal{L}q_t) = \mathcal{L}(q_{1/s} + q_{1/t}) = \mathcal{L}(q_{1/(s+t)}) = q_{s+t}.$$

Then by associativity of inf-convolution, for any W^ -function f , we have*

$$q_s \square (q_t \square f) = (q_s \square q_t) \square f = q_{s+t} \square f.$$

Thus, $(q_t \square (\cdot))_{t>0}$ defines a semigroup acting on W^ -functions. This is the W^* -analog of the Hopf-Lax semigroup.*

Definition 4.8. If f is a W^* function, we say that f is convex if $f^{\mathcal{A}}$ is convex for every $\mathcal{A} \in \mathbb{W}$. We say that f is semi-concave if $q_t - f$ is convex for some $t > 0$.

Lemma 4.9. *Suppose f and g are W^* -functions and $q_t - f$ is convex. Then $q_t - f \square g$ is convex.*

Proof. Note that

$$\begin{aligned} (q_t - f \square g)^{\mathcal{A}}(X) &= \sup_{\substack{\iota: \mathcal{A} \rightarrow \mathcal{B} \\ Y \in L^2(\mathcal{B})_{\text{sa}}^m}} (q_t^{\mathcal{B}}(\iota(X)) - f^{\mathcal{B}}(\iota(X) - Y) - g^{\mathcal{B}}(Y)) \\ &= \sup_{\substack{\iota: \mathcal{A} \rightarrow \mathcal{B} \\ Y \in L^2(\mathcal{B})_{\text{sa}}^m}} \left(q_t^{\mathcal{B}}(\iota(X) - Y) - f^{\mathcal{B}}(\iota(X) - Y) + \frac{1}{t} \langle \iota(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - q_t^{\mathcal{B}}(Y) - g^{\mathcal{B}}(Y) \right). \end{aligned}$$

The right-hand side is the supremum of a family convex functions of X and therefore is convex. \square

As a consequence of Lemmas 4.6 and 4.9, if f is E -convex, then $q_t \square f$ is an E -convex and semi-concave function. The next results give a characterization of such functions as well as some of their regularity properties. These results are quite close to the standard results about convex functions on a Hilbert space, so we do not claim any originality, but nonetheless we include the proofs for the sake of completeness.

Proposition 4.10. *Let f be an E -convex W^* -function that is not identically ∞ or $-\infty$. Then the following are equivalent:*

- (1) $f = q_t \square g$ for some E -convex function g .
- (2) $q_t - f$ is convex.
- (3) $q_t - f$ is E -convex.
- (4) $\mathcal{L}f - q_{1/t}$ is convex and lower semi-continuous.
- (5) $\mathcal{L}f - q_{1/t}$ is E -convex.

Moreover, in this case, $f < \infty$ everywhere.

Proof. (1) \implies (2) follow from Lemma 4.9.

(2) \implies (3). Because $q_t - f$ takes finite values everywhere, by Lemma 3.10, it suffices to show that for every $X \in L^2(\mathcal{A})$, there exists a some $Z \in \partial f^{\mathcal{A}}(X) \cap L^2(W^*(X))_{\text{sa}}^m$. Because $q_t - f$ is convex, it has a subgradient vector Z at X , so that

$$q_t^{\mathcal{A}}(X') - f^{\mathcal{A}}(X') - q_t^{\mathcal{A}}(X) + f^{\mathcal{A}}(X) \geq \langle X' - X, Z \rangle_{L^2(\mathcal{A})_{\text{sa}}^m},$$

which implies that

$$(4.7) \quad f^{\mathcal{A}}(X') - f^{\mathcal{A}}(X) \leq \langle X - X', Z \rangle + \frac{1}{2t} (\|X'\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2) = \langle X' - X, Z + (1/t)X \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \frac{1}{2t} \|X' - X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2.$$

Because f is E -convex, there exists some $Y \in \partial f^{\mathcal{A}}(X) \cap L^2(W^*(X))_{\text{sa}}^m$. Of course,

$$(4.8) \quad f^{\mathcal{A}}(X') - f^{\mathcal{A}}(X) \geq \langle X' - X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}.$$

This implies that

$$\langle X' - X, Z + (1/t)X - Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \geq -\frac{1}{2t} \|X' - X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2$$

for all X' . We $X' = X + tW$ for some $W \in L^2(\mathcal{A})_{\text{sa}}^m$, divide by t , and then send $t \rightarrow 0$ to obtain $\langle W, Z + (1/t)X - Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \geq 0$ for all W . Therefore, $Y = Z + (1/t)X$. This implies $Z \in L^2(W^*(X))_{\text{sa}}^m$.

(3) \implies (4). Note that

$$\begin{aligned} \mathcal{L}f^{\mathcal{A}}(X) - q_{1/t}^{\mathcal{A}}(X) &= \sup_{\substack{\iota: \mathcal{A} \rightarrow \mathcal{B} \\ Y \in L^2(\mathcal{B})_{\text{sa}}^m}} \left(\langle \iota(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - \frac{t}{2} \|\iota(X)\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 - f^{\mathcal{B}}(Y) \right) \\ &= \sup_{\substack{\iota: \mathcal{A} \rightarrow \mathcal{B} \\ Z \in L^2(\mathcal{B})_{\text{sa}}^m}} \left(\langle \iota(X), Z + t\iota(X) \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - \frac{t}{2} \|\iota(X)\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 - f^{\mathcal{B}}(Z + t\iota(X)) \right) \\ &= \sup_{\substack{\iota: \mathcal{A} \rightarrow \mathcal{B} \\ Z \in L^2(\mathcal{B})_{\text{sa}}^m}} \left(-\frac{1}{2t} \|Z\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 + \frac{1}{2t} \|Z + t\iota(X)\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 - f^{\mathcal{B}}(Z + t\iota(X)) \right). \end{aligned}$$

Because $q_t - f$ is convex and lower semi-continuous, the right-hand side is the supremum of convex lower semi-continuous functions of X , and therefore is convex and lower semi-continuous.

(4) \implies (5). Let $h = \mathcal{L}f$. Since f is not identically $-\infty$ or ∞ , the same is true of h . We assumed in (3) that $h - q_{1/t}$ is convex and lower semi-continuous. Moreover, if $E : \mathcal{B} \rightarrow \mathcal{A}$ is a conditional expectation, then $h^{\mathcal{B}}(X) < \infty$ implies $(h - q_{1/t})^{\mathcal{B}}(X) < \infty$ implies $(h - q_{1/t})^{\mathcal{A}}(E[X]) < \infty$ implies $h^{\mathcal{A}}(E[X]) < \infty$. Thus, it remains to show that $h^{\mathcal{A}}(E[X]) \leq h^{\mathcal{B}}(X)$ whenever $h^{\mathcal{A}}(E[X])$ is finite. As in Lemma 3.10, it suffices to show that for every \mathcal{A} and $X \in L^2(\mathcal{A})_{\text{sa}}^m$ with $h^{\mathcal{A}}(X) < \infty$, there exists some subgradient vector $Y \in L^2(\mathbb{W}^*(X), \tau|_{\mathbb{W}^*(X)})_{\text{sa}}^m$. By E -convexity of h , there exists some $Z \in \partial h^{\mathcal{A}}(X) \cap L^2(\mathbb{W}^*(X))_{\text{sa}}^m$. Then we claim that $Z - tX \in \partial h^{\mathcal{A}}(X)$. To prove this, observe that by convexity of $h - q_{1/t}$, for $s \in (0, 1)$, and $X' \in L^2(\mathcal{A})_{\text{sa}}^m$,

$$\begin{aligned} s(h - q_{1/t})^{\mathcal{A}}(X') &\geq (h - q_{1/t})^{\mathcal{A}}((1-s)X + sX') - (1-s)(h - q_{1/t})^{\mathcal{A}}(X) \\ &\geq h^{\mathcal{A}}(X) + \langle (1-s)X + sX' - X, Z \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - q_{1/t}^{\mathcal{A}}((1-s)X + sX') - h^{\mathcal{A}}(X) \\ &\quad + q_{1/t}^{\mathcal{A}}(X) + s(h - q_{1/t})^{\mathcal{A}}(X) \\ &= s(h - q_{1/t})^{\mathcal{A}}(X) + s\langle X' - X, Z \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} q_{1/t}^{\mathcal{A}}(X) - q_{1/t}^{\mathcal{A}}((1-s)X + sX') \\ &= s(h - q_{1/t})^{\mathcal{A}}(X) + s\langle X' - X, Z \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \frac{t}{2}\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - \frac{t}{2}\|X + s(X' - X)\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 \\ &= s(h - q_{1/t})^{\mathcal{A}}(X) + s\langle X' - X, Z - tX \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \frac{ts^2}{2}\|X' - X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2. \end{aligned}$$

Dividing by s and sending $s \rightarrow 0^+$, we obtain

$$(h - q_{1/t})^{\mathcal{A}}(X') \geq (h - q_{1/t})^{\mathcal{A}}(X) + \langle X' - X, Z - tX \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}.$$

Hence, $Z - tX \in \partial(h - q_{1/t})^{\mathcal{A}}(X)$. Since $Z - tX \in L^2(\mathbb{W}^*(X))_{\text{sa}}^m$, the proof is complete.

(5) \implies (1). Since $\mathcal{L}f - q_{1/t}$ is E -convex, we have $\mathcal{L}f - q_{1/t} = \mathcal{L}g$ for some E -convex function g by Proposition 3.17. Thus, since g and $q_{1/t}$ are both E -convex, we have

$$f = \mathcal{L}^2 f = \mathcal{L}(\mathcal{L}g + q_{1/t}) = \mathcal{L}\mathcal{L}(g \square q_t) = g \square q_t,$$

where the last line follows because $g \square q_t$ is E -convex by Lemma 4.6.

Finally, (1) implies that $f < \infty$ everywhere. Indeed, if $X \in L^2(\mathcal{A})_{\text{sa}}^m$, and if Y is some point where $g^{\mathcal{B}}(Y) < \infty$, then let \mathcal{C} be the free product of \mathcal{A} and \mathcal{B} and let $\iota_1 : \mathcal{A} \rightarrow \mathcal{C}$ and $\iota_2 : \mathcal{B} \rightarrow \mathcal{C}$ be the corresponding inclusions. Then $(g \square q_t)^{\mathcal{A}}(X) \leq \frac{1}{2t}\|\iota_1(X) - \iota_2(Y)\|_{L^2(\mathcal{C})_{\text{sa}}^m}^2 + g^{\mathcal{C}}(\iota_2(Y)) < \infty$. \square

Proposition 4.11. *Let f be an E -convex \mathbb{W}^* -function taking values in \mathbb{R} . Then the following are equivalent:*

- (1) $q_t - f$ is convex.
- (2) If $\mathcal{A} \in \mathbb{W}$ and $Y \in \partial f^{\mathcal{A}}(X)$ and $Y' \in \partial f^{\mathcal{A}}(X')$, then $\|Y - Y'\|_{L^2(\mathcal{A})_{\text{sa}}^m} \leq (1/t)\|X - X'\|_{L^2(\mathcal{A})_{\text{sa}}^m}$.
- (3) If $\mathcal{A} \in \mathbb{W}$, then $\partial f^{\mathcal{A}}(X)$ consists of a single point $\nabla f^{\mathcal{A}}(X) \in L^2(\mathbb{W}^*(X))_{\text{sa}}^m$, and $\nabla f^{\mathcal{A}}$ defines a $(1/t)$ -Lipschitz function $L^2(\mathcal{A})_{\text{sa}}^m \rightarrow L^2(\mathcal{A})_{\text{sa}}^m$.
- (4) For each \mathcal{A} and $X \in L^2(\mathcal{A})_{\text{sa}}^m$ and $Y \in \partial f^{\mathcal{A}}(X)$, we have

$$(4.9) \quad \langle X' - X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \leq f^{\mathcal{A}}(X') - f^{\mathcal{A}}(X) \leq \langle X' - X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \frac{1}{2t}\|X' - X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2$$

for all $X' \in L^2(\mathcal{A})_{\text{sa}}^m$.

Proof. (1) \implies (2). By the previous Proposition 4.10, $\mathcal{L}f - q_{1/t}$ is E -convex. Let $Y \in \partial f^{\mathcal{A}}(X)$ and $Y' \in \partial f^{\mathcal{A}}(X')$. Then by Lemma 3.19, we have $X \in \partial \mathcal{L}f^{\mathcal{A}}(Y)$ and $X' \in \partial \mathcal{L}f^{\mathcal{A}}(Y')$. By the same argument as (4) \implies (5) in the proof of Proposition 4.10, we have $Z = X - tY \in \partial(\mathcal{L}f - q_{1/t})$ and $Z' := X' - tY' \in \partial(\mathcal{L}f - q_{1/t})$. It follows that

$$\langle Z', Y - Y' \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \leq \mathcal{L}f^{\mathcal{A}}(Y) - \mathcal{L}f^{\mathcal{A}}(Y') \leq \langle Z, Y - Y' \rangle_{L^2(\mathcal{A})_{\text{sa}}^m},$$

hence

$$\begin{aligned}
0 &\leq \langle Z' - Z, Y' - Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \\
&= \langle X' - X - t(Y' - Y), Y' - Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \\
&= \langle X' - X, Y' - Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + t \|Y' - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 \\
&\leq \|X' - X\|_{L^2(\mathcal{A})_{\text{sa}}^m} \|Y' - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m} + t \|Y' - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2.
\end{aligned}$$

Therefore, $\|Y' - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m} \leq (1/t)\|X' - X\|_{L^2(\mathcal{A})_{\text{sa}}^m}$ as desired.

(2) \implies (3). By taking $X = X'$ in (2), we see that there is a unique $Y \in \bar{\partial}f^{\mathcal{A}}(X)$ and that $X \mapsto Y$ is a $(1/t)$ -Lipschitz function. By Lemma 3.10, we know that $\bar{\partial}f^{\mathcal{A}}(X)$ contains some point in $L^2(W^*(X))_{\text{sa}}^m$, and this point must equal Y .

(3) \implies (4). Let \mathcal{A} and X be given. By our assumption of (3), there is a unique point $Y = \nabla f^{\mathcal{A}}(X)$ in $\bar{\partial}f^{\mathcal{A}}(X)$. Let $X' \in L^2(\mathcal{A})_{\text{sa}}^m$. The lower bound $\langle X' - X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \leq f^{\mathcal{A}}(X') - f^{\mathcal{A}}(X)$ follows immediately from convexity. For the upper bound, let $X_t = (1-t)X' + tX$ and let $Y_t = \nabla f^{\mathcal{A}}(X_t)$.

For $n \in \mathbb{N}$, observe that

$$\begin{aligned}
f^{\mathcal{A}}(X') - f^{\mathcal{A}}(X) &= \sum_{j=1}^n (f^{\mathcal{A}}(X_{j/n}) - f^{\mathcal{A}}(X_{(j-1)/n})) \\
&\leq \sum_{j=1}^n \langle X_{j/n} - X_{(j-1)/n}, Y_{j/n} \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \\
&= \sum_{j=1}^n \langle X_{j/n} - X_{(j-1)/n}, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \sum_{j=1}^n \|X_{j/n} - X_{(j-1)/n}\|_{L^2(\mathcal{A})_{\text{sa}}^m} \|Y_{j/n} - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m} \\
&\leq \langle X' - X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \sum_{j=1}^n \frac{1}{n} \|X' - X\|_{L^2(\mathcal{A})_{\text{sa}}^m} \frac{1}{t} \|X_{j/n} - X\|_{L^2(\mathcal{A})_{\text{sa}}^m} \\
&\leq \langle X' - X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \frac{1}{t} \|X' - X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 \sum_{j=1}^n \frac{j}{n^2} \\
&= \langle X' - X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \frac{1}{t} \|X' - X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 \frac{n(n+1)}{2n^2}.
\end{aligned}$$

Taking $n \rightarrow \infty$ shows that $f^{\mathcal{A}}(X') - f^{\mathcal{A}}(X) \leq \langle X' - X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + (1/2t)\|X' - X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2$ as desired.

(4) \implies (1). Let $\mathcal{A} \in \mathbb{W}$. We show that $(q_t - f)^{\mathcal{A}}$ is convex by exhibiting a subgradient vector for every $X \in L^2(\mathcal{A})_{\text{sa}}^m$. Let $Y \in \bar{\partial}f^{\mathcal{A}}(X)$ and let $X' \in L^2(\mathcal{A})_{\text{sa}}^m$. By (4),

$$\begin{aligned}
(q_t - f)^{\mathcal{A}}(X') - (q_t - f)^{\mathcal{A}}(X) &\geq \frac{1}{2t} \|X'\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - \frac{1}{2t} \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - \langle X' - X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - \frac{1}{2t} \|X' - X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 \\
&= \langle X' - X, -Y + (1/t)X \rangle_{L^2(\mathcal{A})_{\text{sa}}^m}.
\end{aligned}$$

Hence, $-Y + (1/t)X$ is a subgradient vector for $q_t - f$ at X as desired. \square

4.3. Main results on the displacement interpolation. We start out by proving Theorem 1.5 which states that if (\mathcal{A}, X, Y) is an L^2 optimal coupling and $X_t = (1-t)X + tY$, then $W^*(X_t) = W^*(X, Y)$ for all $t \in (0, 1)$.

Proof of Theorem 1.5. By Proposition 3.23, there exists an E -convex function f such that $Y \in \partial f^{\mathcal{A}}(X)$. Let $f_t = (1-t)q_1 + tf$, where $q_1^{\mathcal{A}}(X) = (1/2)\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2$. Since $(1-t)X$ is a subgradient to $(1-t)q_1$ at X and tY is a subgradient to $f^{\mathcal{A}}$ at X , we have $X_t \in \bar{\partial}f_t^{\mathcal{A}}(X)$. By Lemma 3.19, we have $X \in \bar{\partial}\mathcal{L}f_t^{\mathcal{A}}(X_t)$. Since $f_t - q_{1/(1-t)} = f_t - (1-t)q_1 = tf$ is E -convex, $q_{1-t} - \mathcal{L}f_t$ is E -convex by Proposition 4.10. Hence, by Proposition 4.11, $\bar{\partial}\mathcal{L}f_t^{\mathcal{A}}(X_t)$ consists of a single point which is in $L^2(W^*(X_t))_{\text{sa}}^m$. But we already know that $X \in \bar{\partial}\mathcal{L}f_t^{\mathcal{A}}(X_t)$, and therefore $X \in L^2(W^*(X_t))_{\text{sa}}^m$.

A symmetrical argument shows that $Y \in L^2(W^*(X_t))_{\text{sa}}^m$. Therefore, $W^*(X, Y) \subseteq W^*(X_t)$. The reverse inclusion $W^*(X_t) \subseteq W^*(X, Y)$ is obvious since $X_t = (1-t)X + tY$. \square

It follows from the triangle inequality that (\mathcal{A}, X_s, X_t) is an optimal coupling of the laws of X_s and X_t (see Proposition A.22). Another way to show that is, given an E -convex function f such that $Y \in \partial f^{\mathcal{A}}(X)$, to derive E -convex functions $f_{t,s}$ for $s, t \in [0, 1]$ such that $X_t \in \partial f_{t,s}^{\mathcal{A}}(X_s)$. The next proposition gives an explicit construction of $f_{t,s}$ from f , and gives the properties of $f_{t,s}$. The specific cases relevant to the displacement interpolation are then summarized in Corollary 4.13. All of these results are completely analogous to the classical statements.

Proposition 4.12. *Let f be an E -convex function. For $s, t \in [0, 1]$, define $f_{t,s}$ as follows: For $s = 0$, set*

$$f_{t,0} = (1-t)q_1 + tf; \quad f_{0,t} = \mathcal{L}f_{t,0};$$

if $s > 0$ and $s \leq t$, set

$$f_{t,s}^{\mathcal{A}}(X) = \inf_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} \left(\frac{t}{2s} \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - \frac{t-s}{s} \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \frac{(t-s)(1-s)}{2s} \|Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + (t-s)f^{\mathcal{A}}(Y) \right);$$

if $s > 0$ and $s \geq t$, set

$$f_{t,s}^{\mathcal{A}}(X) = \sup_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} \left(\frac{t}{2s} \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - \frac{t-s}{s} \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \frac{(t-s)(1-s)}{2s} \|Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + (t-s)f^{\mathcal{A}}(Y) \right).$$

(In particular, $f_{t,t} = q_1$ for all $t \in [0, 1]$.) Then we have the following:

- (1) $f_{t,s}$ is E -convex and $f_{s,t} = \mathcal{L}f_{t,s}$.
- (2) If $s \leq t$, then $f_{t,s} - \frac{1-t}{1-s}q_1$ is E -convex for $s < 1$ and $\frac{t}{s}q_1 - f_{t,s}$ is E -convex for $s > 0$.
- (3) If $t \leq s$, then $f_{t,s} - \frac{t}{s}q_1$ is E -convex for $s > 0$ and $\frac{1-t}{1-s}q_1 - f_{t,s}$ is E -convex for $s < 1$.
- (4) In particular, if $s \in (0, 1)$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$, then $\partial f_{t,s}^{\mathcal{A}}(X)$ consists of a unique point $\nabla f_{t,s}^{\mathcal{A}}(X)$ and $\nabla f_{t,s}^{\mathcal{A}}$ is Lipschitz.
- (5) Suppose $0 \leq s < t \leq 1$. If $u \in (s, t)$, then

$$f_{u,s} = \frac{t-u}{t-s}q_1 + \frac{u-s}{t-s}f_{t,s}$$

and

$$f_{t,u} = \left(\frac{t-s}{u-s}q_1 \right) \square \left(\frac{t-u}{t-s}f_{t,s} \left(\frac{t-s}{t-u}(\cdot) \right) \right).$$

- (6) Suppose $0 \leq s < t \leq 1$ and $X, Y \in L^2(\mathcal{A})_{\text{sa}}^m$ with $Y \in \partial f_{t,s}(X)$. For $u \in [s, t]$, let

$$X_u = \frac{t-u}{t-s}X + \frac{u-s}{t-s}Y.$$

Then $X_u \in \partial f_{u,s}(X)$ and $Y \in \partial f_{t,u}(X_u)$.

- (7) For $s, t, u \in (0, 1)$, we have $\nabla f_{u,t} \circ \nabla f_{t,s} = \nabla f_{u,s}$.

The next corollary describes the most relevant cases of the proposition for optimal transport; the claims are special cases of (4) and (6) of the proposition.

Corollary 4.13. *Let (\mathcal{A}, X, Y) be an optimal coupling of $\mu, \nu \in \Sigma_m$. Let f be an E -convex function such that $Y \in \partial f(X)$. Let $f_{t,s}$ be as in Proposition 4.12. Let $X_t = (1-t)X + tY$ for $t \in [0, 1]$. Then $X_t \in \partial f_{t,s}(X_s)$ for all $s, t \in [0, 1]$. In particular, if $s \in (0, 1)$, then $f_{t,s}$ has a Lipschitz gradient and we have $X_t = \nabla f_{t,s}(X_s)$.*

In order to prove Proposition 4.12, we need the following scaling relation for Legendre transform.

Lemma 4.14. *Let f be a tracial W^* -function and let $c > 0$. Then $\mathcal{L}(cf)^{\mathcal{A}}(X) = c\mathcal{L}f^{\mathcal{A}}(c^{-1}X)$.*

Proof. Observe that

$$\begin{aligned} \mathcal{L}(cf)^{\mathcal{A}}(X) &= \sup_{\substack{\iota: \mathcal{A} \rightarrow \mathcal{B} \\ Y \in L^2(\mathcal{B})_{\text{sa}}^m}} (\langle \iota(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - cf^{\mathcal{B}}(Y)) \\ &= \sup_{\substack{\iota: \mathcal{A} \rightarrow \mathcal{B} \\ Y \in L^2(\mathcal{B})_{\text{sa}}^m}} (c\langle \iota(c^{-1}X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - cf^{\mathcal{B}}(Y)) \\ &= c\mathcal{L}f^{\mathcal{A}}(c^{-1}X). \end{aligned} \quad \square$$

The bulk of the proof of the proposition is the following lemma which explains how $f_{s,t}$ were obtained through addition of and inf-convolution with quadratic functions, using the same idea as in the proof of Theorem 1.5.

Lemma 4.15. *Consider the setup of Proposition 4.12. If $0 < s < t \leq 1$, then*

$$(4.10) \quad f_{t,s} = \frac{1-t}{1-s}q_1 + \frac{t-s}{1-s} \left[\left(\frac{1}{s}q_1 \right) \square \left((1-s)f \left(\frac{1}{1-s}(\cdot) \right) \right) \right]$$

$$(4.11) \quad \mathcal{L}f_{t,s} = \left(\frac{1-s}{1-t}q_1 \right) \square \left[\frac{t-s}{1-s} [sq_1 + (1-s)\mathcal{L}f] \left(\frac{1-s}{t-s}(\cdot) \right) \right]$$

and

$$(4.12) \quad f_{t,s} = \frac{t}{s}q_1 \square \left[\frac{t-s}{t} [(1-t)q_1 + tf] \left(\frac{t}{t-s}(\cdot) \right) \right]$$

$$(4.13) \quad \mathcal{L}f_{t,s} = \frac{s}{t}q_1 + \frac{t-s}{t} \left[\left(\frac{1}{1-t}q_1 \right) \square t\mathcal{L}f \left(\frac{1}{t}(\cdot) \right) \right].$$

Proof. Fix $\mathcal{A} \in \mathbb{W}$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$, and evaluate the right-hand side of (4.10) at X to obtain

$$\begin{aligned} & \frac{1-t}{1-s}q_1(X) + \frac{t-s}{1-s} \left[\left(\frac{1}{s}q_1 \right) \square \left((1-s)f \left(\frac{1}{1-s}(\cdot) \right) \right) \right] (X) \\ &= \frac{1-t}{2(1-s)} \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + \frac{t-s}{1-s} \inf_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} \left[\frac{1}{2s} \|X - Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + (1-s)f^{\mathcal{A}} \left(\frac{1}{1-s}Y \right) \right], \end{aligned}$$

where we have used the result from Lemma 4.6 that it suffices to take the infimum over $Y \in L^2(\mathcal{A})_{\text{sa}}^m$ rather than Y in $L^2(\mathcal{B})_{\text{sa}}^m$ for some larger tracial W^* -algebra \mathcal{B} . Next, we substitute $(1-s)Y$ instead of Y to obtain

$$\begin{aligned} & \frac{1-t}{2(1-s)} \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + \frac{t-s}{1-s} \inf_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} \left[\frac{1}{2s} \|X - (1-s)Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + (1-s)f^{\mathcal{A}}(Y) \right] \\ &= \inf_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} \left[\frac{1-t}{2(1-s)} \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + \frac{t-s}{2s(1-s)} \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + \frac{t-s}{s} \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} \right. \\ & \quad \left. + \frac{(t-s)(1-s)}{2s} \|Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + (t-s)f^{\mathcal{A}}(Y) \right]. \end{aligned}$$

Combining the two coefficients in front of $\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2$, we arrive at the formula for $f_{t,s}^{\mathcal{A}}(X)$.

The equation (4.11) is obtained from (4.10) by applying the Legendre transform, using the fact that $\mathcal{L}(cq_1) = c^{-1}q_1$ for $c > 0$, the relation between Legendre transform and inf-convolution in Lemma 4.4, and the scaling relation Lemma 4.14.

The proof of (4.12) is similar to the proof of (4.10), and then (4.13) is obtained by taking the Legendre transform. \square

Proof of Proposition 4.12. (1) It is immediate that $f_{t,0}$ and $f_{0,t}$ are E -convex and are Legendre transforms of each other. Also, in the case of $s = t$, we have $f_{t,s} = q_1$, so there is nothing to prove. For $0 < s < t \leq 1$, it follows from Lemma 4.15 that $f_{t,s}$ is E -convex for because it is expressed by applying scaling, addition of quadratics, and inf-convolution with quadratics to f .

Next, we show that for $0 < s < t$, we have $\mathcal{L}f_{t,s} = f_{s,t}$. We evaluate $\mathcal{L}f_{t,s}$ starting from (4.12) as

$$\mathcal{L}f_{t,s} = \frac{s}{t}q_1 + \mathcal{L} \left[\frac{t-s}{t} [(1-t)q_1 + tf] \left(\frac{t}{t-s}(\cdot) \right) \right]$$

Now we evaluate the second term at some $X \in L^2(\mathcal{A})_{\text{sa}}^m$, where $\mathcal{A} \in \mathbb{W}$. Using Remark 3.18, we may compute the Legendre transform of an E -convex function by taking the Hilbert-space Legendre transform for each \mathcal{A} (without considering a larger W^* -algebra \mathcal{B}). This yields

$$\sup_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} \left[\langle X, Y \rangle - \frac{t-s}{t} [(1-t)q_1 + tf]^{\mathcal{A}} \left(\frac{t}{t-s}Y \right) \right].$$

We substitute $\frac{t-s}{t}Y$ for Y to obtain

$$\sup_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} \left[\frac{t-s}{t} \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - \frac{(t-s)(1-t)}{2t} \|Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - (t-s)f(Y) \right].$$

Adding back the term $\frac{s}{t}q_1^A(X)$, we obtain

$$\sup_{Y \in L^2(\mathcal{A})_{\text{sa}}^m} \left[\frac{s}{2t} \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - \frac{s-t}{t} \langle X, Y \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} + \frac{(s-t)(1-t)}{2t} \|Y\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 + (s-t)f(Y) \right],$$

which is precisely $f_{s,t}$.

Therefore, for $s > t > 0$, we have $f_{t,s} = \mathcal{L}f_{s,t}$, hence $f_{t,s}$ is E -convex and $\mathcal{L}f_{t,s} = f_{s,t}$. This is the last remaining case.

(2) If $s = t$, then $f_{t,s} = q_1$, so there Let $s \leq t$. If $s \in (0, 1)$, then by (4.10), $f_{t,s}$ is $\frac{1-t}{1-s}q_1$ plus an E -convex function, hence $f_{t,s} - \frac{1-t}{1-s}q_1$ is E -convex. If $s \in (0, 1]$, then by (4.12), $f_{t,s}$ is the inf-convolution of $\frac{t}{s}q_1$ with an E -convex function and therefore $\frac{t}{s}q_1 - f_{t,s}$ is E -convex by Proposition 4.10.

(3) Let $s \geq t$. Then $f_{t,s} = \mathcal{L}f_{s,t}$. Thus, can we argue symmetrically to (2) using (4.11) and (4.13).

(4) This follows from (2) and (3) together with Proposition 4.11.

(5) Consider the first relation $f_{u,s} = \frac{t-u}{t-s}q_1 + \frac{u-s}{t-s}f_{t,s}$. If $s = 0$, this follows from direct computation from the definition of $f_{u,0}$ and $f_{t,0}$. In the case $s > 0$, we apply (4.10) to get

$$\begin{aligned} \frac{t-u}{t-s}q_1 + \frac{u-s}{t-s}f_{t,s} &= \frac{t-u}{t-s}q_1 + \frac{u-s}{t-s} \frac{1-t}{1-s}q_1 + \frac{u-s}{1-s} \left[\left(\frac{1}{s}q_1 \right) \square \left((1-s)f \left(\frac{1}{1-s}(\cdot) \right) \right) \right] \\ &= \frac{1-u}{1-s}q_1 + \frac{u-s}{1-s} \left[\left(\frac{1}{s}q_1 \right) \square \left((1-s)f \left(\frac{1}{1-s}(\cdot) \right) \right) \right] \\ &= f_{u,s}. \end{aligned}$$

Analogously, using (4.13), we obtain for $s \in (0, 1)$ that

$$\mathcal{L}f_{t,u} = \frac{u-s}{t-s}q_1 + \frac{t-u}{t-s}\mathcal{L}f_{t,s};$$

in fact, this relation also holds when $s = 0$ by evaluating $\mathcal{L}f_{t,u}$ on the left-hand side with (4.13) and evaluating $\mathcal{L}f_{t,0}$ on the right-hand side $\mathcal{L}[(1-t)q_1 + tf] = (\frac{1}{1-t}q_1) \square \mathcal{L}f(\frac{1}{t}(\cdot))$. Taking Legendre transforms of the previous equation implies that

$$f_{t,u} = \left(\frac{t-s}{u-s}q_1 \right) \square \left(\frac{t-u}{t-s}f_{t,s} \left(\frac{t-s}{t-u}(\cdot) \right) \right).$$

(6) Since $X \in \partial q_1^A(X)$ and $Y \in \partial f_{t,s}^A(X)$, we have

$$X_u = \frac{t-u}{t-s}X + \frac{u-s}{t-s}Y \in \partial \left[\frac{t-u}{t-s}q_1 + \frac{u-s}{t-s}f_{t,s} \right]^A(X) = \partial f_{u,s}^A(X).$$

Since $Y \in \partial f_{t,s}^A(X)$, we have $X \in \partial(\mathcal{L}f_{t,s})^A(Y)$. Hence, using the same relation as in the proof (5),

$$X_u = \frac{u-s}{t-s}Y + \frac{t-u}{t-s}X \in \partial \left[\frac{u-s}{t-s}q_1 + \frac{t-u}{t-s}\mathcal{L}f_{t,s} \right]^A(Y) = \partial(\mathcal{L}f_{t,u})^A(Y).$$

So $X_u \in \partial(\mathcal{L}f_{t,u})^A(Y)$, so that $Y \in \partial f_{t,u}(X_u)$.

(7) In light of (4), for $s, t \in (0, 1)$, the functions $f_{s,t}$ and $f_{t,s}$ have Lipschitz gradients. They are Legendre transforms of each other, which implies that $X \in \partial f_{s,t}(Y)$ if and only if $Y \in \partial f_{t,s}(X)$. Hence, $\nabla f_{s,t} = (\nabla f_{t,s})^{-1}$.

Suppose that $s < u < t$. Let $Y = \nabla f_{t,s}(X)$, and let $X_u = \frac{u-s}{t-s}X + \frac{t-u}{t-s}Y$. Then by (6), $X_u = \nabla f_{u,s}(X)$ and $Y = \nabla f_{t,u}(X_u)$, hence $\nabla f_{t,s}(X) = Y = \nabla f_{t,u}(X_u) = \nabla f_{t,u} \circ \nabla f_{u,s}(X)$.

So $\nabla f_{t,s} = \nabla f_{t,u} \circ \nabla f_{u,s}$. By applying $\nabla f_{s,u} = (\nabla f_{u,s})^{-1}$ on the right, we obtain $\nabla f_{t,s} \circ \nabla f_{s,u} = \nabla f_{t,u}$. By taking inverses, $\nabla f_{u,s} \circ \nabla f_{s,t} = \nabla f_{u,t}$. In fact, using composition and inverses in this way, we can achieve all permutations of u, s , and t . The only remaining case is when some of s, t, u are equal to each other, but this follows from the relations $\nabla f_{t,t} = \text{id}$ and $\nabla f_{s,t} = (\nabla f_{t,s})^{-1}$. \square

5. OPTIMAL COUPLINGS, QUANTUM INFORMATION THEORY, AND OPERATOR ALGEBRAS

In this section, we give several indications of why non-commutative optimal couplings are significantly more complicated than the commutative case by making connections to other results in operator algebras and quantum information theory. Specifically, using results from [52], we show that there exist $n \times n$ matrix tuples for which an optimal coupling requires a tracial W^* -algebra of arbitrarily large dimension. Next, based on [33] and [42], we conclude that there exist matrix tuples for which the optimal coupling requires a non-Connes-embeddable tracial W^* -algebra (that is, it cannot even be approximated by couplings in finite-dimensional algebras). Next, we show that the topology induced by the Wasserstein distance is strictly stronger than the weak- $*$ topology on $\Sigma_{m,R}$, and we characterize the points at which the two topologies agree. Finally, we show that $\Sigma_{m,R}$ with the Wasserstein distance is not separable based on [58, Theorem 1].

5.1. Completely positive and factorizable maps. We recall some standard definitions in operator algebras; see e.g. [60]. If \mathcal{A} is a tracial W^* -algebra, we denote by $M_n(\mathcal{A})$ the algebra $M_n(L^\infty(\mathcal{A})) \cong M_n(\mathbb{C}) \otimes L^\infty(\mathcal{A})$ equipped with the trace $\text{tr}_n \otimes \tau_{\mathcal{A}}$ and the weak- $*$ topology given by the entrywise weak- $*$ topology on $L^\infty(\mathcal{A})$; it is a standard fact that $M_n(\mathcal{A})$ is indeed a tracial W^* -algebra. If $\Phi : \mathcal{A} \rightarrow \mathcal{B}$ is a linear map between tracial W^* -algebras, then we define $\Phi^{(n)} : M_n(\mathcal{A}) \rightarrow M_n(\mathcal{B})$ as the map obtained from entrywise application of Φ . If \mathcal{A} is a tracial W^* -algebra and $a \in L^\infty(\mathcal{A})$, then we say that $a \geq 0$ if $a = x^*x$ for some $x \in L^\infty(\mathcal{A})$; this is equivalent to a defining a positive operator on $L^2(\mathcal{A})$ by left multiplication.³

Definition 5.1. We say that Φ is *completely positive* if for every $n \in \mathbb{N}$, if $a \in M_n(\mathcal{A})$ with $a \geq 0$, then $\Phi^{(n)}(a) \geq 0$. For tracial W^* -algebras \mathcal{A} and \mathcal{B} , we denote by $\text{CP}(\mathcal{A}, \mathcal{B})$ the space of completely positive maps $\mathcal{A} \rightarrow \mathcal{B}$. We denote by $\text{UCPT}(\mathcal{A}, \mathcal{B})$ the space of unital completely positive trace-preserving maps. These maps are known in quantum information theory as *quantum channels* from \mathcal{A} to \mathcal{B} .

Definition 5.2 (Anantharaman-Delaroche [1]). Let \mathcal{A} and \mathcal{B} be tracial W^* -algebras. A linear map $\Phi : \mathcal{A} \rightarrow \mathcal{B}$ is said to be *factorizable* if there exist tracial W^* -inclusions $\iota_1 : \mathcal{A} \rightarrow \mathcal{C}$ and $\iota_2 : \mathcal{B} \rightarrow \mathcal{C}$ such that $\Phi = \iota_2^* \circ \iota_1$, where $\iota_2^* : \mathcal{C} \rightarrow \mathcal{B}$ is the conditional expectation adjoint to ι_2 . We also say that Φ *factors through* \mathcal{C} if there exist ι_1 and ι_2 as above.

We denote the space of factorizable maps by $\text{FM}(\mathcal{A}, \mathcal{B})$. We denote by $\text{FM}_{\text{fin}}(\mathcal{A}, \mathcal{B})$ the set of maps that factorize through a finite-dimensional algebra \mathcal{C} .

Proposition 5.3. *Let \mathcal{A} , \mathcal{B} , and \mathcal{C} be tracial W^* -algebras.*

- (1) *We have $\text{FM}(\mathcal{A}, \mathcal{B}) \subseteq \text{UCPT}(\mathcal{A}, \mathcal{B})$.*
- (2) *$\text{UCPT}(\mathcal{A}, \mathcal{B})$, $\text{FM}(\mathcal{A}, \mathcal{B})$, and $\text{FM}_{\text{fin}}(\mathcal{A}, \mathcal{B})$ are convex sets.*
- (3) *$\text{UCPT}(\mathcal{A}, \mathcal{B})$ and $\text{FM}(\mathcal{A}, \mathcal{B})$ are closed in the pointwise weak- $*$ topology.*
- (4) *If $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$ and $\Psi \in \text{UCPT}(\mathcal{B}, \mathcal{C})$, then $\Psi \circ \Phi \in \text{UCPT}(\mathcal{A}, \mathcal{C})$. The same holds with UCPT replaced by FM.*

This proposition is well-known in operator algebras. For the sake of exposition, let us recall why $\text{FM}(\mathcal{A}, \mathcal{B}) \subseteq \text{UCPT}(\mathcal{A}, \mathcal{B})$. Let $\Phi \in \text{FM}(\mathcal{A}, \mathcal{B})$, and take a factorization $\Phi = \iota_2^* \iota_1$ where $\iota_1 : \mathcal{A} \rightarrow \mathcal{C}$ and $\iota_2 : \mathcal{B} \rightarrow \mathcal{C}$ are tracial W^* -inclusions. Since ι_1 and ι_2 are $*$ -homomorphisms, they are completely positive and unital. Now observe that $\langle \iota_2^*(c), b \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} = \langle c, \iota_2(b) \rangle_{L^2(\mathcal{C})_{\text{sa}}^m} \geq 0$ for $c \in M_n(\mathcal{C})_+$ and $b \in M_n(\mathcal{B})_+$; it follows that $\iota_2^*(c) \geq 0$ in $M_n(\mathcal{B})$ for every $c \in M_n(\mathcal{C})_+$. Since ι_2 is unital, ι_2^* is trace-preserving, and since ι_2 is trace-preserving, ι_2^* is unital. Finally, one verifies directly that UCPT is closed under composition, hence, $\iota_2^* \iota_1 \in \text{UCPT}(\mathcal{A}, \mathcal{B})$.

To show that factorizable maps are closed under composition in (4), one uses amalgamated free products. For convexity of $\text{FM}(\mathcal{A}, \mathcal{B})$, see e.g. [12, Lemma 2.3.6].

The next lemma summarizes some well-known facts about completely positive maps.

Lemma 5.4. *Let \mathcal{A} and \mathcal{B} be tracial W^* -algebras, and let $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$.*

- (1) *$\Phi(X^*) = \Phi(X)^*$ for all $X \in \mathcal{A}$.*
- (2) *Φ extends to a contractive map $L^2(\mathcal{A}) \rightarrow L^2(\mathcal{B})$.*

³Of course, definitions make sense more generally for C^* -algebras.

(3) *There exists a unique $\Phi^* \in \text{UCPT}(\mathcal{B}, \mathcal{A})$ such that $\langle X, \Phi^*(Y) \rangle_{L^2(\mathcal{A})} = \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})}$ for $X \in L^2(\mathcal{A})$ and $Y \in L^2(\mathcal{B})$.*

The connection between factorizable maps and non-commutative optimal couplings is as follows.

Observation 5.5. *Let \mathcal{A} and \mathcal{B} be tracial W^* -algebras and let $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$ and $Y \in L^\infty(\mathcal{B})_{\text{sa}}^m$. Then*

$$C(\lambda_X, \lambda_Y) = \sup_{\Phi \in \text{FM}(\mathcal{A}, \mathcal{B})} \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m}.$$

Proof. In fact, we will show that the two sets $\{\langle X', Y' \rangle_{L^2(\mathcal{C})_{\text{sa}}^m} : (\mathcal{C}, X', Y') \text{ a coupling}\}$ and $\{\langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} : \Phi \in \text{FM}(\mathcal{A}, \mathcal{B})\}$ are equal. Suppose that (\mathcal{C}, X', Y') is a coupling of λ_X and λ_Y . Since $\lambda_{X'} = \lambda_X$, there is a tracial W^* embedding $\iota_1 : W^*(X) \rightarrow \mathcal{C}$ sending X to X' . Similarly, there is a tracial W^* -embedding $\iota_2 : W^*(Y) \rightarrow \mathcal{C}$ sending Y to Y' . Let $\phi_1 : W^*(X) \rightarrow \mathcal{A}$ and $\phi_2 : W^*(Y) \rightarrow \mathcal{B}$ be the canonical inclusion maps, and let $\Phi = \phi_2 \iota_2^* \iota_1 \phi_1^* : \mathcal{A} \rightarrow \mathcal{B}$, which is factorizable by Proposition 5.3 (4) because it is a composition of factorizable maps. Moreover,

$$\langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} = \langle \iota_1 \phi_1^*(X), \iota_2 \phi_2^*(Y) \rangle_{L^2(\mathcal{C})_{\text{sa}}^m} = \langle \iota_1(X), \iota_2(Y) \rangle_{L^2(\mathcal{C})_{\text{sa}}^m}.$$

Conversely, given $\Phi \in \text{FM}(\mathcal{A}, \mathcal{B})$, we may factorize it as $\iota_2^* \iota_1$ for tracial W^* -embeddings $\iota_1 : \mathcal{A} \rightarrow \mathcal{C}$ and $\iota_2 : \mathcal{B} \rightarrow \mathcal{C}$, and let $X' = \iota_1(X)$ and $Y' = \iota_2(Y)$ to obtain a coupling (\mathcal{C}, X', Y') of λ_X and λ_Y . \square

5.2. Matrix tuples with optimal couplings of large dimension. This connection allows us to address a natural question: Suppose that μ and ν are non-commutative laws that can be realized by self-adjoint tuples X and Y in a finite-dimensional algebra; then is there a non-commutative optimal coupling (\mathcal{A}, X', Y') of μ and ν such that \mathcal{A} is finite-dimensional? And do we have some control over the dimension? The classical analog of this question certainly has a positive answer. Indeed, if μ and ν are finitely supported measures on \mathbb{R}^m , with supports S and T respectively, then a classical optimal coupling is given by a measure π on the product space $S \times T$. Hence, there exist random variables X and $Y \in L^2(\pi, \mathbb{R}^m)$ such that (\mathcal{A}, X, Y) is an optimal coupling of μ and ν , where \mathcal{A} is the finite-dimensional algebra $L^\infty(S \times T, \pi)$ equipped with the trace coming from π .

Our first negative result in the non-commutative setting shows that, even in situations when an optimal coupling can occur in a finite dimensional algebra, there can be no control over its dimension. This is a consequence of the following result of Musat and Rørdam [52]. Here $\text{FM}_{\text{fin}}(\mathcal{A}, \mathcal{B})$ denotes the space factorizable maps that factorize through a finite-dimensional tracial $*$ -algebra.

Theorem 5.6 (Musat-Rørdam [52, Theorem 4.1]). *If $n \geq 11$, then $\text{FM}_{\text{fin}}(M_n(\mathbb{C}), M_n(\mathbb{C}))$ is not closed, hence there exist factorizable maps $M_n(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ that do not factor through any finite-dimensional algebra.*

In order to translate this result into a statement about non-commutative optimal couplings, we use the following lemma, which is an application of the hyperplane separation theorem, vector space duality, and adjointness of tensor and hom functors.

Lemma 5.7. *Let $L_{\mathbb{R}}(M_n(\mathbb{C})_{\text{sa}}, M_m(\mathbb{C})_{\text{sa}})$ denote the space of real linear transformations $M_n(\mathbb{C}) \rightarrow M_m(\mathbb{C})$. Let $K \subseteq L_{\mathbb{R}}(M_n(\mathbb{C})_{\text{sa}}, M_m(\mathbb{C})_{\text{sa}})$ be a closed convex set, and let $\Phi \notin K$. Then there exists $k \leq \min(n^2, m^2)$ and $X \in M_n(\mathbb{C})_{\text{sa}}^k$ and $Y \in M_m(\mathbb{C})_{\text{sa}}^k$ such that*

$$\langle \Phi(X), Y \rangle_{L^2(M_m(\mathbb{C})_{\text{sa}}^k)} > \sup_{\Psi \in K} \langle \Psi(X), Y \rangle_{L^2(M_m(\mathbb{C})_{\text{sa}}^k)}.$$

Proof. Recall that there is a linear isomorphism

$$T : L_{\mathbb{R}}(M_n(\mathbb{C})_{\text{sa}}, M_m(\mathbb{C})_{\text{sa}}) \rightarrow L_{\mathbb{R}}(M_n(\mathbb{C})_{\text{sa}} \otimes M_m(\mathbb{C})_{\text{sa}}, \mathbb{R}) = (M_n(\mathbb{R})_{\text{sa}} \otimes M_m(\mathbb{R})_{\text{sa}})^*$$

that sends $\Psi \in L_{\mathbb{R}}(M_n(\mathbb{C})_{\text{sa}}, M_m(\mathbb{C})_{\text{sa}})$ to the map

$$\psi : M_n(\mathbb{C})_{\text{sa}} \otimes M_m(\mathbb{C})_{\text{sa}} \rightarrow \mathbb{R} : A \otimes B \mapsto \langle \Psi(A), B \rangle_{L^2(M_m(\mathbb{C})_{\text{sa}})}.$$

Of course, $M_n(\mathbb{C})_{\text{sa}} \otimes M_m(\mathbb{C})_{\text{sa}}$ is finite-dimensional, so the double dual is isomorphic to the original space. Applying the hyperplane separation theorem on the real inner-product space $M_n(\mathbb{C})_{\text{sa}} \otimes M_m(\mathbb{C})_{\text{sa}}$, we conclude that there exists some $v \in M_n(\mathbb{C})_{\text{sa}} \otimes M_m(\mathbb{C})_{\text{sa}}$ such that

$$T(\Phi)(v) > \sup_{\Psi \in K} T(\Psi)(v).$$

Let us decompose v into a sum of simple tensors $v = \sum_{j=1}^k X_j \otimes Y_j$. The smallest k for which this is possible is called the *tensor rank* of v . We claim that the tensor rank is at most $\min(n^2, m^2)$. The reason is that for real vector spaces V and W , we can identify $V \otimes W$ with $L_{\mathbb{R}}(V, W)$ and then apply the singular value decomposition of the matrix in $L_{\mathbb{R}}(V, W)$ corresponding to a given tensor v . Since a matrix in $L_{\mathbb{R}}(V, W)$ has rank at most $\min(\dim V, \dim W)$, it follows that the tensor rank of v is at most $\min(\dim V, \dim W)$. In particular, taking $V = M_n(\mathbb{C})_{\text{sa}}$ and $W = M_m(\mathbb{C})_{\text{sa}}$, we see that our vector $v \in M_n(\mathbb{C}) \otimes M_m(\mathbb{C})$ has tensor rank at most $\min(n^2, m^2)$.

Let $X = (X_1, \dots, X_k)$ and $Y = (Y_1, \dots, Y_k)$. Then for $\Psi \in L_{\mathbb{R}}(M_n(\mathbb{C})_{\text{sa}}, M_m(\mathbb{C})_{\text{sa}})$, we have

$$T(\Psi)(v) = \sum_{j=1}^k \langle \Psi(X_j), Y_j \rangle_{L^2(M_m(\mathbb{C})_{\text{sa}})} = \langle \Psi(X), Y \rangle_{L^2(M_m(\mathbb{C})_{\text{sa}})^k}$$

Thus, by our choice of v , the tuples X and Y satisfy the desired properties. \square

Corollary 5.8. *If $n \geq 11$ and $d \in \mathbb{N}$, then there exist $X, Y \in M_n(\mathbb{C})_{\text{sa}}^{n^2}$ such that for every optimal coupling (\mathcal{A}, X', Y') of λ_X and λ_Y , \mathcal{A} must have dimension at least d . In particular, if d is sufficiently large, then*

$$C(\lambda_X, \lambda_Y) > \sup_{U \in \mathcal{U}(M_n(\mathbb{C}))} \langle UXU^*, Y \rangle_{M_n(\mathbb{C})}.$$

Proof. Let $\text{FM}_d(M_n(\mathbb{C}), M_n(\mathbb{C}))$ denote the set of UCPT maps $M_n(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ that factorize through a tracial W^* -algebra $\mathcal{A} = (A, \tau)$ of dimension at most d . As a consequence of the Artin-Wedderburn theorem, every such $*$ -algebra A is a direct sum of at most d matrix algebras of size at most $d^{1/2}$; see e.g. [25]. Moreover, the trace $\tau_{\mathcal{A}}$ is a convex combination of the traces on each component. From these facts, it is not hard to see that $\text{FM}_d(M_n(\mathbb{C}), M_n(\mathbb{C}))$ is compact.

By Theorem 5.6, there exists $\Phi \in \overline{\text{FM}_{\text{fin}}(M_n(\mathbb{C}), M_n(\mathbb{C}))}$ that does not factor through a finite-dimensional algebra, and hence $\Phi \in \overline{\text{FM}_{\text{fin}}(M_n(\mathbb{C}), M_n(\mathbb{C}))} \setminus \text{FM}_d(M_n(\mathbb{C}), M_n(\mathbb{C}))$.

Also, we also remark that a completely positive map $\Phi : M_n(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ satisfies $\Phi(A^*) = \Phi(A)^*$, and therefore it restricts to a real-linear transformation $M_n(\mathbb{C})_{\text{sa}} \rightarrow M_n(\mathbb{C})_{\text{sa}}$, and Φ is uniquely determined by its restriction to self-adjoint elements. Thus, we can Lemma 5.7 to conclude that there exists $k \leq n^2$ and $X, Y \in M_n(\mathbb{C})_{\text{sa}}^k$ such that

$$\langle \Phi(X), Y \rangle_{L^2(M_n(\mathbb{C})_{\text{sa}})^k} > \sup_{\Psi \in \text{FM}_d(M_n(\mathbb{C}), M_n(\mathbb{C}))} \langle \Psi(X), Y \rangle_{L^2(M_n(\mathbb{C})_{\text{sa}})^k},$$

for some $k \leq n^2$. We can without loss of generality take $k = n^2$ because we can always add additional zero entries to our tuples without changing the value of the inner product of $\Psi(X)$ and Y . Hence, by the proof of Observation 5.5 any $\Psi \in \text{FM}_d(M_n(\mathbb{C}), M_n(\mathbb{C}))$ cannot produce an optimal coupling. \square

5.3. Optimal couplings and the Connes embedding problem. The situation is even more wild than this. Based on the work of [42] on Tsirelson's problem and the Connes embedding problem, as well as work of [33], we can conclude that for some n , there exist $X, Y \in M_n(\mathbb{C})_{\text{sa}}^{n^2}$, such that a non-commutative optimal coupling of λ_X and λ_Y cannot even be *approximated* by couplings in finite-dimensional tracial $*$ -algebras. We begin with some background on the Connes embedding problem, which includes first the definition of ultraproducts of tracial W^* -algebras, a tool to turn approximate embeddings into literal embeddings; [12, Appendix A], [13, §2], [2, §5.4].

Let $\beta\mathbb{N}$ denote the *Stone-Ćech compactification* of the natural numbers; it is a compact space containing \mathbb{N} as an open subset,⁴ and satisfies the universal property that every function from \mathbb{N} into a compact topological space K extends uniquely to a continuous function $\beta\mathbb{N} \rightarrow K$. In particular, if $(x_n)_{n \in \mathbb{N}}$ is a bounded sequence of real or complex numbers, then $\lim_{n \rightarrow \mathcal{U}} x_n$ exists for every $\mathcal{U} \in \beta\mathbb{N}$. The Stone-Ćech compactification $\beta\mathbb{N}$ is characterized up to a canonical homeomorphism by its universal property. One construction of $\beta\mathbb{N}$ is by way of ultrafilters, which is why we have used the letter \mathcal{U} for elements of $\beta\mathbb{N}$. In this framework, the elements of $\beta\mathbb{N} \setminus \mathbb{N}$ are known as *non-principal ultrafilters* and the limit $\lim_{n \rightarrow \mathcal{U}} x_n$ is also called an *ultralimit*.

⁴Here \mathbb{N} is equipped with the discrete topology. To simplify notation, we view \mathbb{N} as a subset of $\beta\mathbb{N}$ rather than considering an inclusion map $\eta : \mathbb{N} \rightarrow \beta\mathbb{N}$.

Ultraproducts of tracial von Neumann algebras are defined as follows. For $n \in \mathbb{N}$, let $\mathcal{A}_n = (A_n, \tau_n)$ be a sequence of tracial W^* -algebras. Let $\prod_{n \in \mathbb{N}} \mathcal{A}_n$ be the set of sequences $(a_n)_{n \in \mathbb{N}}$ such that $\sup_n \|a_n\|_{L^\infty(\mathcal{A}_n)} < \infty$, which is a $*$ -algebra. Let

$$I_{\mathcal{U}} = \left\{ (a_n)_{n \in \mathbb{N}} \in \prod_{n \in \mathbb{N}} \mathcal{A}_n : \lim_{n \rightarrow \mathcal{U}} \|a_n\|_{L^2(\mathcal{A}_n)} = 0 \right\}.$$

Using the non-commutative Hölder's inequality for L^2 and L^∞ , one can show that $I_{\mathcal{U}}$ is an ideal in $\prod_{n \in \mathbb{N}} \mathcal{A}_n$, and therefore, $\prod_{n \in \mathbb{N}} \mathcal{A}_n / I_{\mathcal{U}}$ is a $*$ -algebra. We denote by $[a_n]_{n \in \mathbb{N}}$ the equivalence class in $\prod_{n \in \mathbb{N}} \mathcal{A}_n / I_{\mathcal{U}}$ of a sequence $(a_n)_{n \in \mathbb{N}} \in \prod_{n \in \mathbb{N}} \mathcal{A}_n$. Furthermore, we define a trace on $\prod_{n \in \mathbb{N}} \mathcal{A}_n / I_{\mathcal{U}}$ as follows. If $(a_n)_{n \in \mathbb{N}} \in \prod_{n \in \mathbb{N}} \mathcal{A}_n$, then $(\tau_n(a_n))_{n \in \mathbb{N}}$ is a bounded sequence in \mathbb{C} and therefore $\lim_{n \rightarrow \mathcal{U}} \tau_n(a_n)$ exists. Since $|\tau_n(a_n)| \leq \|a_n\|_{L^2(\mathcal{A}_n)}$, we have $\lim_{n \rightarrow \mathcal{U}} \tau_n(a_n) = 0$ whenever $(a_n)_{n \in \mathbb{N}} \in I_{\mathcal{U}}$. Therefore, there is a well-defined map

$$\tau_{\mathcal{U}} : \prod_{n \in \mathbb{N}} \mathcal{A}_n / I_{\mathcal{U}} \rightarrow \mathbb{C}$$

given by $\tau_{\mathcal{U}}([a_n]_{n \in \mathbb{N}}) = \lim_{n \rightarrow \mathcal{U}} \tau_n(a_n)$. It turns out the pair $(\prod_{n \in \mathbb{N}} \mathcal{A}_n / I_{\mathcal{U}}, \tau_{\mathcal{U}})$ is already a tracial W^* -algebra; see [2, Proposition 5.4.1]. The proof is based on the fact that a tracial C^* -algebra is a W^* -algebra if and only if the operator-norm unit ball is complete in the L^2 norm [2, Proposition 2.6.4]. See also [12, Appendix A].

We call the tracial W^* -algebra $(\prod_{n \in \mathbb{N}} \mathcal{A}_n / I_{\mathcal{U}}, \tau_{\mathcal{U}})$ the *ultraproduct of $(\mathcal{A}_n)_{n \in \mathbb{N}}$ with respect to \mathcal{U}* and we denote it by

$$\prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n := \left(\prod_{n \in \mathbb{N}} \mathcal{A}_n / I_{\mathcal{U}}, \tau_{\mathcal{U}} \right).$$

The inspiration for this notation is that ultraproduct is defined using a combination of Cartesian product and ultralimits (of the L^2 -norm and the trace); in contrast to Cartesian products, the ultraproduct only cares about the asymptotic behavior of a sequence as $n \rightarrow \mathcal{U}$.

Definition 5.9. A tracial W^* -algebra \mathcal{A} is *Connes-embeddable* if there exist finite-dimensional tracial $*$ -algebras \mathcal{A}_n for $n \in \mathbb{N}$, an ultrafilter $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$, and a tracial W^* -embedding $\phi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$. The *Connes embedding problem* is the question of whether every tracial W^* -algebra with separable predual is Connes-embeddable.

Embeddings into ultraproducts are closely related to convergence of non-commutative laws in $\Sigma_{m,R}$.

Lemma 5.10. *Let $(\mathcal{A}_n)_{n \in \mathbb{N}}$ be a sequence of tracial W^* -algebras and let \mathcal{A} be another tracial W^* -algebra. Let $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$ with $\|X\|_{L^\infty(\mathcal{A})_{\text{sa}}^m} \leq R$ and suppose that X generates \mathcal{A} as a W^* -algebra. Let $X_n \in L^\infty(\mathcal{A}_n)_{\text{sa}}^m$ with $\|X_n\|_{L^\infty(\mathcal{A}_n)_{\text{sa}}^m} \leq R$. Then the following are equivalent:*

- (1) $\lim_{n \rightarrow \mathcal{U}} \lambda_{X_n} = \lambda_X$ with respect to the weak- $*$ topology on $\Sigma_{m,R}$.
- (2) There exists a tracial W^* -embedding $\phi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ such that $\phi(X) = [X_n]_{n \in \mathbb{N}}$.

Proof. (1) \implies (2). Let $Y = [X_n]_{n \in \mathbb{N}} \in L^\infty(\prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n)_{\text{sa}}^m$. Let $\tau_{\mathcal{U}}$ be the trace on the ultraproduct. Then for every $p \in \mathbb{C}\langle x_1, \dots, x_d \rangle$, we have

$$\lambda_Y(p) = \tau_{\mathcal{U}}(p(Y)) = \lim_{n \rightarrow \mathcal{U}} \tau_n(p(X_n)) = \lim_{n \rightarrow \mathcal{U}} \lambda_{X_n}(p) = \lambda_X(p).$$

Because $\lambda_Y = \lambda_X$, Lemma 2.33 implies that there is a W^* -embedding $\mathcal{A} = W^*(X) \rightarrow W^*(Y) \hookrightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$.

(2) \implies (1). Let $\phi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ as above be a tracial W^* -embedding with $\phi(X) = [X_n]_{n \in \mathbb{N}}$. Using the fact that ϕ preserves addition and multiplication as well as the definition of the trace $\tau_{\mathcal{U}}$ on the ultraproduct,

$$\lambda_X(p) = \tau(p(X)) = \tau_{\mathcal{U}}(\phi(p(X))) = \tau_{\mathcal{U}}(p(\phi(X))) = \lim_{n \rightarrow \mathcal{U}} \tau_n(p(X_n)) = \lim_{n \rightarrow \mathcal{U}} \lambda_{X_n}(p).$$

Therefore, $\lim_{n \rightarrow \mathcal{U}} \lambda_{X_n} = \lambda_X$ in the weak- $*$ topology, as desired. \square

Definition 5.11. Let $\Sigma_{m,R}^{\text{fin}}$ be the set of non-commutative laws μ in $\Sigma_{m,R}$ such that $\mu = \lambda_X$ for some $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$ where \mathcal{A} is a finite-dimensional tracial $*$ -algebra.

The following statement is almost a corollary of Lemma 5.10.

Lemma 5.12. *Let \mathcal{A} be a tracial W^* -algebra generated by $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$ with $\|X\|_{L^\infty(\mathcal{A})^m} \leq R$. Then \mathcal{A} is Connes-embeddable if and only if λ_X is in the weak- $*$ closure of $\Sigma_{m,R}^{\text{fin}}$ in $\Sigma_{m,R}$.*

Proof. If λ_X is in the closure of $\Sigma_{m,R}^{\text{fin}}$, then Lemma 5.10 implies that \mathcal{A} is Connes-embeddable. Conversely, suppose that \mathcal{A} is Connes-embeddable and $\iota : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ is an embedding into some ultraproduct of finite-dimensional tracial $*$ -algebras. Let $X_n = (X_n^{(1)}, \dots, X_n^{(m)}) \in L^\infty(\mathcal{A}_n)^m$ such that $[X_n]_{n \in \mathbb{N}} = \iota(X)$, and let us also write $X = (X^{(1)}, \dots, X^{(m)})$. By replacing $X_{n,j}$ with $(X_n^{(j)} + (X_n^{(j)})^*)/2$, we can assume without loss of generality that $X_n^{(j)}$ is self-adjoint. By assumption $M := \sup_{n \in \mathbb{N}} \|X_n\|_{L^\infty(\mathcal{A})_{\text{sa}}^m} < \infty$.

Although M may be larger than R , we can fix this issue through a standard argument with functional calculus. Let $f : [-M, M] \rightarrow [-R, R]$ be given by $f(t) = \text{sgn}(t) \min(|t|, R)$. By the Weierstrass approximation theorem, there exists a sequence of polynomials $(f_k)_{k \in \mathbb{N}}$ converging uniformly to f on $[-M, M]$. Note that $f_k(\iota(X^{(j)})) = [f_k(X_n^{(j)})]_{n \rightarrow \mathbb{N}}$. By the spectral mapping theorem, for each j, k , and n ,

$$\|f_k(X_n^{(j)}) - f(X_n^{(j)})\|_{L^\infty(\mathcal{A}_n)} \leq \sup_{t \in [-M, M]} |f_k(t) - f(t)|,$$

and the same estimate holds for $f_k(\iota(X^{(j)})) - f(\iota(X^{(j)}))$. Taking $k \rightarrow \infty$, we obtain $f(\iota(X^{(j)})) = [f(X_n^{(j)})]_{n \in \mathbb{N}}$ for each j . Since $\|X^{(j)}\|_{L^\infty(\mathcal{A})} \leq R$, we have $f(\iota(X^{(j)})) = \iota(X_j)$. Let $Y_n = (f(X_n^{(1)}), \dots, f(X_n^{(m)}))$. Then $\|Y\|_{L^\infty(\mathcal{A}_n)^m} \leq R$ and $\iota(X) = [Y_n]_{n \in \mathbb{N}}$, hence λ_X is in the closure of $\Sigma_{m,R}^{\text{fin}}$ by Lemma 5.10. \square

Decades of work found many equivalent problems in operator algebras and quantum information theory; for a survey, see e.g. [13, 59]. In particular, building on the established connections with quantum information theory, Haagerup and Musat showed the following result.

Theorem 5.13 (Haagerup-Musat [33, Theorem 3.6, 3.7]). *A factorizable map $\Phi : M_n(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ admits a factorization through a Connes-embeddable algebra if and only if it is in the closure of $\text{FM}_{\text{fin}}(M_n(\mathbb{C}), M_n(\mathbb{C}))$. Moreover, the Connes embedding problem has a positive answer if and only if*

$$\text{FM}(M_n(\mathbb{C}), M_n(\mathbb{C})) = \overline{\text{FM}_{\text{fin}}(M_n(\mathbb{C}), M_n(\mathbb{C}))} \text{ for all } n \in \mathbb{N}.$$

A negative answer to the Connes embedding problem was announced in [42]. This implies the following corollary.

Corollary 5.14. *There exist $n \in \mathbb{N}$ and $X, Y \in M_n(\mathbb{C})_{\text{sa}}^{n^2}$ such that*

$$C(\lambda_X, \lambda_Y) = \sup_{\Phi \in \text{FM}(M_n(\mathbb{C}), M_n(\mathbb{C}))} \langle \Phi(X), Y \rangle_{L^2(M_n(\mathbb{C}))_{\text{sa}}^m} > \sup_{\Phi \in \text{FM}_{\text{fin}}(M_n(\mathbb{C}), M_n(\mathbb{C}))} \langle \Phi(X), Y \rangle_{L^2(M_n(\mathbb{C}))_{\text{sa}}^m}.$$

Moreover, a non-commutative optimal coupling of λ_X and λ_Y does not exist in any Connes-embeddable tracial W^ -algebra.*

Proof. Let $K = \overline{\text{FM}_{\text{fin}}(M_n(\mathbb{C}), M_n(\mathbb{C}))}$, which is compact and convex. Because the Connes embedding problem has a negative answer [42], there exists $\Phi \in \text{FM}(M_n(\mathbb{C}), M_n(\mathbb{C})) \setminus K$. By Lemma 5.7, there exist $X, Y \in M_n(\mathbb{C})^{n^2}$ such that

$$\langle \Phi(X), Y \rangle_{L^2(M_n(\mathbb{C}))_{\text{sa}}^{n^2}} > \sup_{\Psi \in K} \langle \Psi(X), Y \rangle_{L^2(M_n(\mathbb{C}))_{\text{sa}}^{n^2}}.$$

Hence, by Theorem 5.13, if Ψ factors through a Connes-embeddable algebra, then $\langle \Psi(X), Y \rangle_{L^2(M_n(\mathbb{C}))_{\text{sa}}^{n^2}}$ cannot be optimal. Thus, by the proof of Observation 5.5, a coupling of λ_X and λ_Y in a Connes-embeddable algebra cannot be optimal. \square

Remark 5.15. Although Corollary 5.14 is much stronger than Corollary 5.8 as stated, they are based on different types of phenomena. Corollary 5.14 relies on the existence of factorizable maps $M_n(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ that cannot be approximated by elements of $\text{FM}_{\text{fin}}(M_n(\mathbb{C}), M_n(\mathbb{C}))$ (of which there are not yet explicit examples known). Meanwhile, Corollary 5.8 relies on the existence of factorizable maps that are approximated by elements of $\text{FM}_{\text{fin}}(M_n(\mathbb{C}), M_n(\mathbb{C}))$ but are not in $\text{FM}_{\text{fin}}(M_n(\mathbb{C}), M_n(\mathbb{C}))$ (of which [52]

gave explicit examples). Thus, the proof of Corollary 5.8 shows that for $n \geq 11$ and $d \in \mathbb{N}$, there exist tuples X and Y from $M_n(\mathbb{C})_{\text{sa}}^{n^2}$ such that

$$\sup_{\Phi \in \text{FM}_{\text{fin}}(M_n(\mathbb{C}), M_n(\mathbb{C}))} \langle \Phi(X), Y \rangle_{L^2(M_n(\mathbb{C}))_{\text{sa}}^{n^2}} > \sup_{\Psi \in \text{FM}_d(M_n(\mathbb{C}), M_n(\mathbb{C}))} \langle \Psi(X), Y \rangle_{L^2(M_n(\mathbb{C}))_{\text{sa}}^{n^2}}.$$

Hence, a coupling on an algebra \mathcal{A} of dimension at most d cannot even be optimal among couplings in Connes-embeddable algebras.

5.4. The Wasserstein and weak-* topologies. At the beginning, we equipped $\Sigma_{m,R}$ with the weak-* topology as a subset of the algebraic dual of $\mathbb{C}\langle x_1, \dots, x_m \rangle$. Meanwhile, because $d_W^{(2)}$ defines a metric on $\Sigma_{m,R}$, it induces another topology, which we will call the *Wasserstein topology*. We will show that the Wasserstein topology is strictly stronger than the weak-* topology. This is to be contrasted with classical probability theory where the weak-* topology on the space of probability measures on $[-R, R]^m$ is metrized by the L^2 -Wasserstein distance.

Our first step is to prove an ultraproduct characterization of Wasserstein convergence analogous to Lemma 5.10.

Lemma 5.16. *Let $(\mathcal{A}_n)_{n \in \mathbb{N}}$ be a sequence of tracial W^* -algebras and let \mathcal{A} be another tracial W^* -algebra. Let $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$ with $\|X\|_{L^\infty(\mathcal{A})_{\text{sa}}^m} \leq R$ and suppose that X generates \mathcal{A} . Let $X_n \in L^\infty(\mathcal{A}_n)_{\text{sa}}^m$ with $\|X_n\|_{L^\infty(\mathcal{A}_n)_{\text{sa}}^m} \leq R$. Then the following are equivalent:*

- (1) $\lim_{n \rightarrow \mathcal{U}} \lambda_{X_n} = \lambda_X$ with respect to Wasserstein distance.
- (2) There exists a tracial W^* -embedding $\phi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ and a factorizable map $\Phi_n \in \text{FM}(\mathcal{A}, \mathcal{A}_n)$ (for each $n \in \mathbb{N}$) such that

$$\phi(X) = [X_n]_{n \in \mathbb{N}}, \quad \phi(Z) = [\Phi_n(Z)]_{n \in \mathbb{N}} \text{ for all } Z \in L^\infty(\mathcal{A}).$$

Proof. (1) \implies (2). The limit $\lim_{n \rightarrow \mathcal{U}} \lambda_{X_n} = \lambda_X$ in Wasserstein distance means that there exists tracial W^* algebras \mathcal{B}_n and tracial W^* -embeddings $\pi_n : W^*(X_n) \rightarrow \mathcal{B}_n$ and $\rho_n : \mathcal{A} \rightarrow \mathcal{B}_n$ such that $\|\pi_n(X_n) - \rho_n(X)\|_{L^2(\mathcal{B}_n)_{\text{sa}}^m} \rightarrow 0$ as $n \rightarrow \mathcal{U}$. Let \mathcal{C}_n be the free product of \mathcal{A}_n and \mathcal{B}_n with amalgamation over $W^*(X_n)$, and let $\tilde{\pi}_n : \mathcal{A}_n \rightarrow \mathcal{C}_n$ and $\tilde{\rho}_n : \mathcal{A} \rightarrow \mathcal{C}_n$ be the corresponding tracial W^* -embeddings. It is straightforward to check that these induce tracial W^* -embeddings

$$\tilde{\pi} : \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{C}_n, \quad \tilde{\rho} : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{C}_n$$

such that $\tilde{\pi}(\phi(X)) = \tilde{\pi}([X_n]_{n \in \mathbb{N}}) = \tilde{\rho}(X)$. Since $\tilde{\pi} \circ \phi$ and $\tilde{\rho}$ are tracial W^* -embeddings, we have $\tilde{\pi}(\phi(Z)) = \tilde{\rho}(Z)$ for all $Z \in L^\infty(\mathcal{A})$ (because for instance every element of $L^\infty(\mathcal{A})$ can be approximated in $L^2(\mathcal{A})$ by non-commutative polynomials of X).

Let $\tilde{\pi}_n^*$ and $\tilde{\rho}_n^*$ be the trace-preserving conditional expectations adjoint to $\tilde{\pi}_n$ and $\tilde{\rho}_n$. We claim that for $Y = [Y_n]_{n \in \mathbb{N}} \in \prod_{n \rightarrow \mathcal{U}} \mathcal{C}_n$, we have

$$\tilde{\pi}^*(Y) = [\tilde{\pi}_n^*(Y_n)]_{n \in \mathbb{N}}.$$

Let $\tilde{\mathcal{A}} = \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ and $\tilde{\mathcal{C}} = \prod_{n \rightarrow \mathcal{U}} \mathcal{C}_n$. Note that $[\tilde{\pi}_n^*(Y_n)]_{n \in \mathbb{N}}$ is in the subalgebra $\tilde{\mathcal{A}} = \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$. Moreover, for every $Z = [Z_n] \in \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$, we have

$$\langle Y, \tilde{\pi}(Z) \rangle_{L^2(\tilde{\mathcal{C}})} = \lim_{n \rightarrow \mathcal{U}} \langle Y_n, \tilde{\pi}_n(Z_n) \rangle_{L^2(\mathcal{C}_n)} = \lim_{n \rightarrow \mathcal{U}} \langle \tilde{\pi}_n^*(Y_n), Z_n \rangle_{L^2(\mathcal{A}_n)} = \langle [\tilde{\pi}_n^*(Y_n)]_{n \in \mathbb{N}}, Z \rangle_{L^2(\tilde{\mathcal{A}})}.$$

Thus, $\tilde{\pi}^*(Y) = [\tilde{\pi}_n^*(Y_n)]_{n \in \mathbb{N}}$, as desired. As noted above, for every $Z \in \mathcal{A}$, we have $\tilde{\pi}(\phi(Z)) = \tilde{\rho}(Z)$ and hence $\phi(Z) = \tilde{\pi}^* \tilde{\rho}(Z) = \tilde{\pi}^* \tilde{\rho}(Z)$. This implies that

$$[\tilde{\pi}_n^* \tilde{\rho}(Z)]_{n \in \mathbb{N}} = \tilde{\pi}^* \tilde{\rho}(Z) = \phi(Z).$$

Therefore, $\Phi_n := \tilde{\pi}_n^* \tilde{\rho}_n$ is a factorizable map fulfilling condition (2).

(2) \implies (1). Let ϕ and Φ_n be as in (2). Then $[X_n]_{n \in \mathbb{N}} = \phi(X) = [\Phi_n(X)]_{n \in \mathbb{N}}$ belongs to $\prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$. Letting E_n be the trace-preserving conditional expectation $\mathcal{A}_n \rightarrow W^*(X_n)$, the map $E_n \circ \Phi_n : W^*(X) \rightarrow W^*(X_n)$ is factorizable by Proposition 5.3 (4), hence by Observation 5.5,

$$C(\lambda_{X_n}, X) \geq \langle E_n \circ \Phi_n(X), X_n \rangle_{L^2(W^*(X_n))_{\text{sa}}^m} = \langle \Phi_n(X), X_n \rangle_{L^2(\mathcal{A}_n)_{\text{sa}}^m}.$$

Therefore,

$$\begin{aligned}
\lim_{n \rightarrow \mathcal{U}} d_W^{(2)}(\lambda_{X_n}, \lambda_X)^2 &= \lim_{n \rightarrow \mathcal{U}} \left(\|X_n\|_{L^2(\mathcal{A}_n)_{\text{sa}}}^2 + \|X\|_{L^2(\mathcal{A})_{\text{sa}}}^2 - 2C(\lambda_{X_n}, \lambda_X) \right) \\
&\leq \lim_{n \rightarrow \mathcal{U}} \left(\|X\|_{L^2(\mathcal{A})_{\text{sa}}}^2 + \|X_n\|_{L^2(\mathcal{A}_n)_{\text{sa}}}^2 - 2\langle \Phi_n(X), X_n \rangle_{L^2(\mathcal{A}_n)_{\text{sa}}} \right) \\
&= \|\phi(X)\|_{L^2(\prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n)_{\text{sa}}}^2 + \|\phi(X)\|_{L^2(\prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n)_{\text{sa}}}^2 - 2\langle \phi(X), \phi(X) \rangle_{L^2(\prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n)_{\text{sa}}} \\
&= 0.
\end{aligned}$$

Hence, $\lim_{n \rightarrow \mathcal{U}} \lambda_{X_n} = \lambda_X$ in Wasserstein distance. \square

The next corollary was observed in [11, Proposition 1.4(b)], and can be proved in several ways (see for instance [35, Lemma 2.10, Corollary 2.11] for another method), but we will deduce it as a consequence of the ultraproduct characterizations for weak-* and Wasserstein convergence.

Corollary 5.17. *The Wasserstein topology on $\Sigma_{m,R}$ is refines the weak-* topology.*

Proof. Fix $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. Using the Urysohn subsequence principle, it suffices to show that if $\mu_n, \mu \in \Sigma_{m,R}$ and $\lim_{n \rightarrow \mathcal{U}} \mu_n = \mu$ in the Wasserstein distance, then $\lim_{n \rightarrow \mathcal{U}} \mu_n \rightarrow \mu$ in the weak-* topology. Letting (\mathcal{A}_n, X_n) and (\mathcal{A}, X) be the GNS realizations of μ_n and μ , Lemma 5.16 implies that there is a tracial W^* -embedding $\mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ with $\phi(X) = [X_n]_{n \in \mathbb{N}}$. By Lemma 5.10, this implies that $\lim_{n \rightarrow \mathcal{U}} \mu_n = \mu$ in the weak-* topology. \square

The next observation is closely related.

Lemma 5.18. *The metric $d_W^{(2)}$ is weak-* lower semi-continuous on $\Sigma_{m,R} \times \Sigma_{m,R}$.*

Proof. Fix $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. Again using the Urysohn subsequence principle, it suffices to show that for every pair of sequences $(\mu_n)_{n \in \mathbb{N}}$ and $(\nu_n)_{n \in \mathbb{N}}$ in $\Sigma_{m,R}$, letting $\mu = \lim_{n \rightarrow \mathcal{U}} \mu_n$ and $\nu = \lim_{n \rightarrow \mathcal{U}} \nu_n$, we have $d_W^{(2)}(\mu, \nu) \leq \lim_{n \rightarrow \mathcal{U}} d_W^{(2)}(\mu_n, \nu_n)$. Let $(\mathcal{A}_n, X_n, Y_n)$ be an optimal couplings of μ_n and ν_n . Let (\mathcal{B}, X) and (\mathcal{C}, Y) be the GNS realizations of μ and ν . By Lemma 5.10, there exist tracial W^* -embeddings $\phi : \mathcal{B} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ and $\psi : \mathcal{C} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ such that $\phi(X) = [X_n]_{n \in \mathbb{N}}$ and $\psi(Y) = [Y_n]_{n \in \mathbb{N}}$. Then

$$d_W^{(2)}(\mu, \nu) \leq \|\phi(X) - \psi(Y)\|_{L^2(\prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n)_{\text{sa}}} = \lim_{n \rightarrow \mathcal{U}} \|X_n - Y_n\|_{L^2(\mathcal{A}_n)} = \lim_{n \rightarrow \mathcal{U}} d_W^{(2)}(\mu_n, \nu_n). \quad \square$$

We will use Lemmas 5.10 and 5.12 to characterize when the Wasserstein and weak-* topologies agree at a point in $\Sigma_{m,R}$ in terms of a certain stability property. To fix terminology, if S is a set and \mathcal{T}_1 and \mathcal{T}_2 are two topologies on S , we say that \mathcal{T}_1 and \mathcal{T}_2 agree at $x \in S$ if every \mathcal{T}_1 -neighborhood of x is contained in a \mathcal{T}_2 -neighborhood of x and vice versa. If the topologies are metrizable, this is equivalent to saying that a sequence x_n converges to x with respect to \mathcal{T}_1 if and only if it converges to x with respect \mathcal{T}_2 . Furthermore, if \mathcal{U} is a given non-principal ultrafilter on \mathbb{N} , then agreement of the two topologies at x is equivalent to the claim that $\lim_{n \rightarrow \mathcal{U}} x_n = x$ with respect to \mathcal{T}_1 if and only if $\lim_{n \rightarrow \mathcal{U}} x_n = x$ with respect to \mathcal{T}_2 .

Definition 5.19 (FM-lifting). Let \mathcal{A} be a tracial W^* -algebra with separable predual, and let \mathcal{U} be a free ultrafilter on \mathbb{N} . If \mathcal{A}_n is a sequence of tracial W^* -algebras and $\phi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ is a tracial W^* -embedding, then an FM-lifting of ϕ is a sequence $(\Phi_n)_{n \in \mathbb{N}}$, where $\Phi_n \in \text{FM}(\mathcal{A}, \mathcal{A}_n)$, such that $\phi(Z) = [\Phi_n(Z)]_{n \in \mathbb{N}}$ for all $Z \in L^\infty(\mathcal{A})$.

Note that the sequence Φ_n in Lemma 5.16 (2) is an FM-lifting of ϕ . In other words, Lemma 5.16 describes convergence in Wasserstein distance in terms of ultraproduct embeddings that have FM-liftings.

Definition 5.20 (FM-stability). We say that \mathcal{A} is FM-stable if every tracial W^* -embedding $\phi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ into the ultraproduct of any sequence of tracial W^* -algebras \mathcal{A}_n has an FM-lifting.

Our notion of FM-stability is analogous and closely related to the notions of tracial stability and UCP-stability studied in [5, 34]. Analogously to [5, Remark 2.2], the definition of FM-stability can be restated as an approximation property without reference to ultraproducts. This implies in particular that the definition is independent of the choice of non-principal ultrafilter \mathcal{U} (hence it amounts to the same

thing whether require the lifting condition for a particular non-principal ultrafilter or for all non-principal ultrafilters).

Proposition 5.21. *Let $\mu \in \Sigma_{m,R}$ and let (\mathcal{A}, X) be the GNS realization of μ as in Proposition 2.31. Then the following are equivalent:*

- (1) *The weak-* and Wasserstein topologies on $\Sigma_{m,R}$ agree at μ .*
- (2) *\mathcal{A} is FM-stable.*

Proof. (1) \implies (2). Let $\mathcal{U} \in \beta\mathbb{N} \setminus \mathbb{N}$. Assume that the weak-* and Wasserstein topologies agree at μ . Let $(\mathcal{A}_n)_{n \in \mathbb{N}}$ be a sequence of tracial W^* -algebras, and let $\phi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ be a tracial W^* -embedding. Express $\phi(X)$ as $[X_n]_{n \in \mathbb{N}}$ where $X_n \in L^2(\mathcal{A}_n)_{\text{sa}}^m$ and $\sup_n \|X_n\|_{L^\infty(\mathcal{A}_n)_{\text{sa}}^m} < \infty$. Arguing with functional calculus as in Lemma 5.12, we can arrange that $\|X_n\|_{L^\infty(\mathcal{A}_n)_{\text{sa}}^m} \leq R$. By Lemma 5.10, we have $\lambda_{X_n} \rightarrow \lambda_X$ in the weak-* topology on $\Sigma_{m,R}$. Hence, by hypothesis $\lambda_{X_n} \rightarrow \lambda_X$ in the Wasserstein distance as $n \rightarrow \mathcal{U}$. By Lemma 5.16, this implies that ϕ has an FM-lifting.

(2) \implies (1). Conversely, suppose that \mathcal{A} is FM-stable. To show that the weak-* and Wasserstein topologies on $\Sigma_{m,R}$ agree at μ , using the Urysohn subsequence principle, it suffices to show that if $(\mu_n)_{n \in \mathbb{N}}$ is a sequence such that $\mu_n \rightarrow \mu$ weak-* as $n \rightarrow \mathcal{U}$, then $d_W^{(2)}(\mu_n, \mu) \rightarrow 0$ as $n \rightarrow \mathcal{U}$. Let (\mathcal{A}_n, X_n) be the GNS-realization of μ_n . By Lemma 5.10, the tuple $[X_n]_{n \in \mathbb{N}}$ in $\prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ has the same law as X , and therefore, there exists a tracial W^* -embedding $\phi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ with $\phi(X) = [X_n]_{n \in \mathbb{N}}$. By FM-stability of \mathcal{A} , there exist factorizable completely positive maps $\Phi_n : \mathcal{A} \rightarrow \mathcal{A}_n$ such that $\phi(Z) = [\Phi_n(Z)]_{n \in \mathbb{N}}$ for all $Z \in L^\infty(\mathcal{A})$. Hence, by Lemma 5.16, $\lim_{n \rightarrow \mathcal{U}} \mu_n = \mu$ in Wasserstein distance. \square

Next, we will show that using work of Connes [17] that if the weak-* and Wasserstein topologies agree at μ and the corresponding tracial W^* -algebra \mathcal{A} is Connes-embeddable, then in fact \mathcal{A} is approximately finite-dimensional. We recall the following theorem of Connes [17] that shows that approximate finite-dimensionality is equivalent to semi-discreteness for tracial W^* -algebras (and these are also equivalent, famously, to the two other conditions of injectivity and amenability); related proofs can also be found in [61], [67, §XIV], [12, §6.2, 6.3, 9.3], [2, §11].

Theorem 5.22 (Connes [17]). *Let $\mathcal{A} = (A, \tau)$ be a tracial W^* -algebra with separable predual. The following are equivalent:*

- (1) *\mathcal{A} is approximately finite-dimensional (AFD), that is, there exists a sequence $(A_k)_{k \in \mathbb{N}}$ of finite-dimensional subalgebras with $A_k \subseteq A_{k+1}$ such that $\bigcup_{k \in \mathbb{N}} A_k$ is dense in A with respect to $\|\cdot\|_{L^2(\mathcal{A})}$.*
- (2) *\mathcal{A} is semi-discrete, that is, there exists nets $(\Phi_\alpha)_{\alpha \in I}$ and $(\Psi_\alpha)_{\alpha \in I}$ of completely positive maps $\Phi_\alpha : \mathcal{A} \rightarrow M_{n(\alpha)}(\mathbb{C})$ and $\Psi_\alpha : M_{n(\alpha)}(\mathbb{C}) \rightarrow \mathcal{A}$ such that $\Psi_\alpha \circ \Phi_\alpha(Z) \rightarrow Z$ in the weak-* topology for every $Z \in L^\infty(\mathcal{A})$.*

We recall a few more results about AFD algebras, which are well-known in operator algebras. We recall that a II_1 -factor is an infinite-dimensional tracial von Neumann algebra with trivial center.

Lemma 5.23.

- (1) *Let \mathcal{A} be an AFD tracial W^* -algebra, let $(\mathcal{B}_n)_{n \in \mathbb{N}}$ be II_1 -factors, and let \mathcal{U} be a free ultrafilter on \mathbb{N} . If ϕ and ψ are two embeddings of \mathcal{A} into $\prod_{n \rightarrow \mathcal{U}} \mathcal{B}_n$, then there exists a unitary $U \in \prod_{n \rightarrow \mathcal{U}} \mathcal{B}_n$ such that $U\phi(Z)U^* = \psi(Z)$ for $Z \in L^\infty(\mathcal{A})$. See [44, 5].*
- (2) *If $(\mathcal{B}_n)_{n \in \mathbb{N}}$ are II_1 -factors and U is a unitary in $\prod_{n \rightarrow \mathcal{U}} \mathcal{B}_n$, then there exist unitaries $U_n \in L^\infty(\mathcal{B}_n)$ such that $U = [U_n]_{n \in \mathbb{N}}$.⁵*

Corollary 5.24. *Let \mathcal{A} be an AFD tracial W^* -algebra. Then \mathcal{A} is FM-stable.*

Proof. If $\mathcal{A} = \mathbb{C}$, then the conclusion is immediate, so assume that $\mathcal{A} \neq \mathbb{C}$. Let $\phi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ be a tracial W^* -embedding. Let \mathcal{B} be the tracial free product $\mathcal{A} * \mathcal{A}_n * L^\infty[0, 1]$ (where $L^\infty[0, 1]$ has the trace

⁵Every unitary u in a tracial W^* -algebra can be expressed as e^{ix} for some self-adjoint x using Borel functional calculus (Theorem A.3 (3)). Suppose $U = e^{iX}$ is unitary in $\prod_{n \rightarrow \mathcal{U}}$. Arguing as in the proof of Lemma 5.12, X can be expressed as $[X_n]_{n \in \mathbb{N}}$ where $X_n \in L^\infty(\mathcal{B}_n)_{\text{sa}}$, and we have $[e^{iX_n}]_{n \in \mathbb{N}} = U$.

coming from Lebesgue measure). Then \mathcal{B} is a II_1 factor by [68, Theorem 3.7] since $\mathcal{A} \neq \mathbb{C}$ and $L^\infty[0, 1]$ is diffuse. For each, n , there is a tracial W^* -embedding $\iota_n : \mathcal{A}_n \rightarrow \mathcal{B}_n$. Let ι be the induced map

$$\iota : \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{B}_n.$$

By construct, there also exists a tracial W^* -embedding $\psi_n : \mathcal{A} \rightarrow \mathcal{B}_n$. This sequence produces a tracial W^* -embedding $\psi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{B}_n$. By Lemma 5.23, there exists a unitary $U_n \in L^\infty(\mathcal{B}_n)$ such that, letting $U = [U_n]_{n \in \mathbb{N}}$, we have $U\iota \circ \phi(Z)U^* = \psi(Z)$ for $Z \in L^\infty(\mathcal{A})$.

Let $\Phi_n : \mathcal{A} \rightarrow \mathcal{A}_n$ be given by $\Phi_n(Z) = \iota_n^*[U_n^*\psi_n(Z)U_n]$. As observed in the proof of Proposition 5.21, ultraproducts respect conditional expectations and therefore for $Z \in \mathcal{A}$, we have

$$[\Phi_n(Z)]_{n \in \mathbb{N}} = [\iota_n^*[U_n^*\psi_n(Z)U_n]]_{n \in \mathbb{N}} = \iota^*[U_n^*\psi_n(Z)U_n]_{n \in \mathbb{N}} = \iota^*(U^*\psi(Z)U) = \iota^*\iota\phi(Z) = \phi(Z).$$

Thus, Φ_n is the desired lifting of ϕ to a sequence of factorizable maps. \square

Remark 5.25. In fact, [5, Theorem 2.6] implies the converse of Corollary 5.24: If \mathcal{A} is Connes-embeddable and FM-stable, then \mathcal{A} is AFD. The same statement is implied by the next proposition provided that \mathcal{A} is finitely generated.

Proposition 5.26. *Let μ be in the weak- $*$ closure of $\Sigma_{m,R}^{\text{fin}}$, and let (\mathcal{A}, X) be the GNS realization of μ . The following are equivalent:*

- (1) \mathcal{A} is approximately finite-dimensional.
- (2) \mathcal{A} is FM-stable.
- (3) The weak- $*$ and Wasserstein topologies agree at μ .
- (4) μ is in the Wasserstein closure of $\Sigma_{m,R}^{\text{fin}}$.

Proof. (1) \implies (2) by Corollary 5.24.

(2) \implies (3) by Proposition 5.21.

(3) \implies (4) Since two topologies agree at μ and μ is in the weak- $*$ closure of $\Sigma_{m,R}^{\text{fin}}$, it follows that μ is in the Wasserstein closure of $\Sigma_{m,R}^{\text{fin}}$.

(4) \implies (1). Assume that (4) holds and we will show that \mathcal{A} is semi-discrete, hence approximately finite-dimensional by Connes' theorem. Fix a free ultrafilter \mathcal{U} on \mathbb{N} . Let μ_n be a sequence in $\Sigma_{n,R}^{\text{fin}}$ such that $\lim_{n \rightarrow \mathcal{U}} d_W^{(2)}(\mu_n, \mu) = 0$. Let $(\mathcal{A}_n, X_n, Y_n)$ be an optimal coupling of μ and μ_n . Since $W^*(X_n) \cong W^*(X) = \mathcal{A}$, we can assume without loss of generality that $\mathcal{A} \subseteq \mathcal{A}_n$ and $X_n = X$. Let $\Phi_n : \mathcal{A} = W^*(X) \rightarrow W^*(Y_n)$ be the associated factorizable map. Since $W^*(Y_n)$ is finite-dimensional, if we can show that $\Phi_n^*\Phi_n(Z) \rightarrow Z$ in $L^2(\mathcal{A})_{\text{sa}}^m$ as $n \rightarrow \mathcal{U}$ for every $Z \in \mathcal{A}$, that will imply semi-discreteness of \mathcal{A} and finish the argument.

The convergence of $\Phi_n^*\Phi_n(Z)$ follows by a similar argument to Proposition 5.21. Let \mathcal{B}_n be the free product of two copies of \mathcal{A}_n with amalgamation over $W^*(Y_n)$ and let π_n and ρ_n be the two inclusions of \mathcal{A} into the first and second copies of \mathcal{A}_n . Then $\Phi_n^*\Phi_n = \pi_n^*\rho_n$. Now π_n and ρ_n induce maps

$$\pi, \rho : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{B}_n.$$

Moreover, $\|\pi_n(X) - \rho_n(X)\|_{L^2(\mathcal{B}_n)_{\text{sa}}^m} \leq 2\|X - Y_n\|_{L^2(\mathcal{A}_n)_{\text{sa}}^m} \rightarrow 0$, and therefore, $\pi(X) = \rho(X)$, so $\pi = \rho$ on all of $L^\infty(\mathcal{A})$. This implies that $\pi^*\rho(Z) = Z$ for $Z \in L^\infty(\mathcal{A})$, hence $\lim_{n \rightarrow \mathcal{U}} \|\pi_n^*\rho_n(Z) - Z\|_{L^2(\mathcal{A})} = 0$. \square

Corollary 5.27. *For $m > 1$ and $R > 0$, $\Sigma_{m,R}$ is not compact with respect to the Wasserstein topology.*

Proof. The identity map from $\Sigma_{m,R}$ with the Wasserstein topology to $\Sigma_{m,R}$ with the weak- $*$ topology is a continuous bijection. If the domain were compact, then it would be a homeomorphism. The previous proposition would then imply that every $\mu \in \Sigma_{m,R}$ that generates a Connes-embeddable tracial W^* -algebra would in fact generate an AFD tracial W^* -algebra. However, there are many finitely generated tracial W^* -algebras that are not AFD. \square

5.5. Non-separability of the Wasserstein space. We just showed that $\Sigma_{m,R}$ with the Wasserstein distance is not compact for $m > 1$, but in fact we will show that it is not separable using the results of Gromov [30], Olshanskii [55], and Ozawa [58]. We first recall some terminology about groups and their associated W^* -algebras.

Let Γ be a group and let $\ell^2(\Gamma)$ be the Hilbert space of square-summable functions on Γ . Let $\lambda : \Gamma \rightarrow B(\ell^2(\Gamma))$ be the *left regular representation* given by $u(g)\delta_h = \delta_{gh}$, where $\delta_g \in \ell^2(\Gamma)$ is the function which is 1 at g and zero elsewhere. The W^* -subalgebra of $B(\ell^2(\Gamma))$ generated by the unitary operators $u(g)$ for $g \in \Gamma$ is called the *group von Neumann algebra of Γ* . The map $\tau : L(\Gamma) \rightarrow \mathbb{C}$ given by $T \mapsto \langle \delta_e, T\delta_e \rangle$ is a faithful normal trace on $L(\Gamma)$, so that it is a tracial W^* -algebra.

Definition 5.28. A discrete group Γ is said to have *property (T)* if there exist generators g_1, \dots, g_m and an increasing function $f : [0, \infty) \rightarrow [0, \infty)$ with $\lim_{\epsilon \rightarrow 0^+} f(\epsilon) = 0$ with the following property: For every unitary representation π of Γ on a Hilbert space H and every unit vector $\xi \in H$, if $\max_{j \in [m]} \|\pi(g_j)\xi - \xi\| < \epsilon$, then there exists $\eta \in H$ such that $\pi(g)\eta = \eta$ for all $g \in \Gamma$ and $\|\eta - \xi\| < f(\epsilon)$.

Theorem 5.29 (Gromov [30], Olshanskii [55], and Ozawa [58, Theorem 1]). *There exists a group Γ with property (T) that admits uncountable family $\{\Gamma_\alpha\}_{\alpha \in I}$ of quotient groups that are simple and pairwise non-isomorphic. (In fact, such a family of quotient groups exists for every group Γ that is hyperbolic, torsion-free, and non-cyclic.)*

The next lemma will allow us to translate this result into a statement about the space of non-commutative laws. While the space of non-commutative laws is defined in terms of self-adjoint generators, it is natural in the group setting to consider unitary rather than self-adjoint generators of a tracial W^* -algebra. However, this issue is easily resolved by taking real and imaginary parts of operators. More precisely, if a is an operator in a tracial W^* -algebra \mathcal{A} , let $\operatorname{Re}(a) = (a + a^*)/2$ and $\operatorname{Im}(a) = (a - a^*)/2i$. Then $\operatorname{Re}(a)$ and $\operatorname{Im}(a)$ are self-adjoint and $a = \operatorname{Re}(a) + i\operatorname{Im}(a)$ and $\|a\|_{L^2(\mathcal{A})}^2 = \|\operatorname{Re}(a)\|_{L^2(\mathcal{A})}^2 + \|\operatorname{Im}(a)\|_{L^2(\mathcal{A})}^2$.

Lemma 5.30. *Let Γ be a group with property (T), and let $g_1, \dots, g_m \in \Gamma$ and $f : [0, \infty) \rightarrow [0, \infty)$ be as in Definition 5.28. Let $q_1 : \Gamma \rightarrow \Gamma_1$ and $q_2 : \Gamma \rightarrow \Gamma_2$ be quotient group homomorphisms. For $j = 1, 2$, let $\pi_j : \Gamma \rightarrow L(\Gamma_j)$ be the quotient map q_j composed with the left regular representation of Γ_j and let*

$$X_j = (\operatorname{Re}(\pi_j(g_1)), \operatorname{Im}(\pi_j(g_1)), \dots, \operatorname{Re}(\pi_j(g_m)), \operatorname{Im}(\pi_j(g_m))) \in L(\Gamma_j)_{\text{sa}}^{2m}.$$

If $f(d_W^{(2)}(\lambda_{X_1}, \lambda_{X_2})) < 1/2$, then $\ker(q_1) = \ker(q_2)$ and hence $\Gamma_1 = \Gamma_2$.

Proof. Let \mathcal{A} be a tracial W^* -algebra and let $\iota_j : L(\Gamma_j) \rightarrow \mathcal{A}$ be tracial W^* -embeddings such that $\|\iota_1(X_1) - \iota_2(X_2)\|_{L^2(\mathcal{A})_{\text{sa}}^{2m}} = d_W^{(2)}(\lambda_{X_1}, \lambda_{X_2})$. Note that for $j = 1, \dots, m$,

$$\begin{aligned} \|\iota_1(\pi_1(g_j)) - \iota_2(\pi_2(g_j))\|_{L^2(\mathcal{A})}^2 &= \|\iota_1(\operatorname{Re}(\pi_1(g_j))) - \iota_2(\operatorname{Re}(\pi_2(g_j)))\|_{L^2(\mathcal{A})}^2 + \|\iota_1(\operatorname{Im}(\pi_1(g_j))) - \iota_2(\operatorname{Im}(\pi_2(g_j)))\|_{L^2(\mathcal{A})}^2 \\ &\leq \|\iota_1(X_1) - \iota_2(X_2)\|_{L^2(\mathcal{A})_{\text{sa}}^{2m}}^2. \end{aligned}$$

Let $\pi : \Gamma \rightarrow B(L^2(\mathcal{A}))$ be the map given by $\pi(g)\xi = \iota_1(\pi_1(g))\xi\iota_2(\pi_2(g^{-1}))$ for $\xi \in L^2(\mathcal{A})$; note that this is a unitary representation. The vector $\widehat{1}$ in $L^2(\mathcal{A})$ satisfies

$$\begin{aligned} \left\| \pi(g_j)\widehat{1} - \widehat{1} \right\|_{L^2(\mathcal{A})} &= \left\| \iota_1(\pi_1(g_j))\widehat{1} - \widehat{1}\iota_2(\pi_2(g_j)) \right\|_{L^2(\mathcal{A})} \\ &= \|\iota_1(\pi_1(g_j)) - \iota_2(\pi_2(g_j))\|_{L^2(\mathcal{A})} \\ &\leq d_W^{(2)}(\lambda_{X_1}, \lambda_{X_2}). \end{aligned}$$

Hence, by property (T), there exists some $\eta \in L^2(\mathcal{A})$ such that $\|\widehat{1} - \eta\|_{L^2(\mathcal{A})} \leq f(d_W^{(2)}(\lambda_{X_1}, \lambda_{X_2}))$ and $\pi(g)\eta = \eta$ for all $g \in \Gamma$. The latter condition implies that $\iota_1(\pi_1(g))\eta = \eta\iota_2(\pi_2(g))$ for all $g \in \Gamma$. Therefore, using the triangle inequality and the fact that $\iota_j(\pi_j(g))$ is unitary,

$$\left\| \iota_1(\pi_1(g))\widehat{1} - \widehat{1}\iota_2(\pi_2(g)) \right\|_{L^2(\mathcal{A})} \leq 2 \left\| \widehat{1} - \eta \right\|_{L^2(\mathcal{A})} \leq 2f(d_W^{(2)}(\lambda_{X_1}, \lambda_{X_2})) < 1.$$

Hence,

$$|\tau_{\mathcal{A}}(\iota_1(\pi_1(g))) - \tau_{\mathcal{A}}(\iota_2(\pi_2(g)))| \leq \|\iota_1(\pi_1(g)) - \iota_2(\pi_2(g))\|_{L^2(\mathcal{A})} < 1.$$

Now observe that

$$\tau_{\mathcal{A}}(\iota_j(\pi_j(g))) = \tau_{L(\Gamma_j)}(\pi_j(g)) = \delta_{\pi_j(g)=1} = \delta_{g \in \ker(q_j)}.$$

Since $\delta_{g \in \ker(q_j)}$ is either zero or one and $|\delta_{g \in \ker(q_1)} - \delta_{g \in \ker(q_2)}| < 1$, we have $\ker(q_1) = \ker(q_2)$. \square

We can now prove Theorem 1.8 that shows that for sufficiently large m , $\Sigma_{m,1}$ is not separable with respect to $d_W^{(2)}$. The method is similar to [58, Proof of Theorem 2].

Proof of Theorem 1.8. First, we show that $\Sigma_{2m,1}$ is not separable for some m . Let Γ be a property (T) group with an uncountable family $(\Gamma_\alpha)_{\alpha \in I}$ of non-isomorphic quotients. Let $\pi_\alpha : \Gamma \rightarrow L(\Gamma_\alpha)$ be the quotient map composed with the left regular representation. Let g_1, \dots, g_m and $f : [0, \infty) \rightarrow [0, \infty)$ witness property (T). Let ϵ be sufficiently small that $f(\epsilon) < 1/2$. Let

$$X_\alpha = (\operatorname{Re}(\pi_\alpha(g_1)), \operatorname{Im}(\pi_\alpha(g_1)), \dots, \operatorname{Re}(\pi_\alpha(g_m)), \operatorname{Im}(\pi_\alpha(g_m))).$$

For $\alpha \neq \beta$ in I , since Γ_α and Γ_β are not isomorphic, the lemma implies that $f(d_W^{(2)}(\lambda_{X_\alpha}, \lambda_{X_\beta})) \geq 1/2$, and therefore $d_W^{(2)}(\lambda_{X_\alpha}, \lambda_{X_\beta}) \geq \epsilon$. Hence, $\{\lambda_{X_\alpha}\}_{\alpha \in I}$ is an uncountable ϵ -separated set in $\Sigma_{2m,1}$ with respect to the Wasserstein distance.

To prove that $\Sigma_{m,R}$ is not separable for general $m > 1$ and $R > 0$, we first observe that there is a bijection between $\Sigma_{m,R}$ and $\Sigma_{m,R'}$ given by rescaling the non-commutative random variables. Hence, for each m , if we prove non-separability for one value of R , then it holds for all values of R . Furthermore, we can define a map $\Sigma_{m,R} \rightarrow \Sigma_{m+1,R}$ sending the law of (X_1, \dots, X_m) to the law of $(X_1, \dots, X_m, 0)$. It is straightforward to show that this map is isometric with respect to the Wasserstein distance. Hence, if $\Sigma_{m,R}$ is not separable, then $\Sigma_{m',R}$ is not separable for $m' \geq m$. Therefore, to prove the theorem, it suffices to show that for some value of R , $\Sigma_{2,R}$ is not separable.

We already know that for some m , $\Sigma_{m,1}$ is not separable. Hence, for some $\epsilon > 0$, there is an uncountable family $(\mu_\alpha)_{\alpha \in I}$ of laws in $\Sigma_{m,1}$ that is ϵ -separated with respect to the Wasserstein distance. Let $(\mathcal{A}_\alpha, X_\alpha)$ be the GNS realization of μ_α , where $X_\alpha = (X_{\alpha,1}, \dots, X_{\alpha,m})$. Consider the tracial W^* -algebra $M_m(\mathcal{A}_\alpha)$ with the trace $\tau_\alpha \otimes \operatorname{tr}_m$, and let $Y_\alpha \in M_m(\mathcal{A}_\alpha)_{\text{sa}}$ be the diagonal matrix with entries $X_{\alpha,1} + 4, X_{\alpha,2} + 8, \dots, X_{\alpha,m} + 4m$. Let $U_\alpha \in M_m(\mathbb{C}) \subseteq M_m(\mathcal{A}_\alpha)$ be the matrix of an m -cycle permutation. By functional calculus, U_α can be expressed as e^{iZ_α} for some self-adjoint $Z_\alpha \in M_m(\mathbb{C}) \subseteq M_m(\mathcal{A}_\alpha)$ with $\|Z_\alpha\|_{L^\infty(M_m(\mathbb{C}))} \leq \pi/2$. Since U_α is the inclusion into $M_m(\mathcal{A}_\alpha)$ of an element of $M_m(\mathbb{C})$ that is independent of α , there is in fact a polynomial p such that $U_\alpha = p(Z_\alpha)$, and Z_α and p are independent of α . We claim that $d_W^{(2)}(\lambda_{Y_\alpha, Z_\alpha}, \lambda_{Y_\beta, Z_\beta}) \geq (1/K)d_W^{(2)}(\mu_\alpha, \mu_\beta)$ for some $K > 0$, which will imply that $\Sigma_{2,4m+1}$ is not separable and thus prove the theorem.

To accomplish this, we will express $X_{\alpha,j} \otimes I_m$ in $M_m(\mathcal{A}_\alpha)$ as a function of Y_α and Z_α (in an explicit way which allows us to estimate Wasserstein distances), using a well-known matrix amplification trick. We first recall a foundational result that the weak- $*$ topology of a W^* -algebra can be recovered from any faithful representation on a Hilbert space; see e.g. [63, Corollary 1.13.3, Proposition 1.16.2, Theorem 1.16.7]. In particular, \mathcal{A}_α can be faithfully represented on $H = L^2(\mathcal{A})$ and $M_m(\mathcal{A}_\alpha) = \mathcal{A}_\alpha \otimes M_m(\mathbb{C})$ can be faithfully represented on the Hilbert space $H \otimes \mathbb{C}^m = H^{\oplus m}$. Moreover, all the facts about spectral theory and functional calculus on $B(H)$ and $B(H^{\oplus m})$ can be applied to the operators from \mathcal{A}_α and $M_n(\mathcal{A}_\alpha)$. In particular,

$$\operatorname{Spec}(Y_\alpha) = \bigcup_{j=1}^m (\operatorname{Spec}(X_{\alpha,j}) + 4j) \subseteq \bigcup_{j=1}^m [4j - 1, 4j + 1].$$

Let γ_j be the rectangular contour in \mathbb{C} bounding the rectangle $[4j - 2, 4j + 2] \times [-1, 1]$, so that γ_j is separated from $\operatorname{Spec}(Y_\alpha)$ by a distance of 1. Using the Cauchy integral formula and functional calculus,

$$\int_{\gamma_j} (z - 4j)(z - X_{\alpha,k})^{-1} dz = \delta_{j,k} X_{\alpha,k}.$$

Hence,

$$\int_{\gamma_j} (z - 4j)(z - Y_\alpha)^{-1} dz = X_{\alpha,j} \otimes e_{j,j},$$

where $e_{j,j}$ is the j th diagonal matrix unit in $M_m(\mathbb{C})$. In particular, $X_{\alpha,j} \otimes e_{j,j} \in W^*(Y_\alpha)$ and thus $X_{\alpha,j} \otimes e_{k,\ell} = U_\alpha^{k-j}(X_{\alpha,j} \otimes e_{j,j})U_\alpha^{j-\ell} \in W^*(Y_\alpha, Z_\alpha)$ for every $k, \ell = 1, \dots, m$; this implies that Y_α and Z_α generate $M_m(\mathcal{A}_\alpha)$. Moreover,

$$(5.1) \quad X_{\alpha,j} \otimes I_m = \sum_{k=1}^m \int_{\gamma_j} U_\alpha^k(z - 4j)(z - Y_\alpha)^{-1} U_\alpha^{-k} dz = \sum_{k=1}^m \int_{\gamma_j} p(Z_\alpha)^k (z - 4j)(z - Y_\alpha)^{-1} \bar{p}(Z_\alpha)^k dz.$$

Let $\alpha \neq \beta$. Then an optimal coupling of $\lambda_{Y_\alpha, Z_\alpha}$ and $\lambda_{Y_\beta, Z_\beta}$ on the tracial W^* -algebra \mathcal{B} produces two tracial W^* -embeddings $\iota_\alpha : M_m(\mathcal{A}_\alpha) \rightarrow \mathcal{B}$ and $\iota_\beta : M_m(\mathcal{A}_\beta) \rightarrow \mathcal{B}$. Because the Cauchy integral representation (5.1) can be expressed as a Riemann integral, we have

$$\iota_\alpha(X_{\alpha,j} \otimes I_m) = \sum_{k=1}^m \int_{\gamma_j} p(\iota_\alpha(Z_\alpha))^k (z - 4j)(z - \iota_\alpha(Y_\alpha))^{-1} \bar{p}(\iota_\alpha(Z_\alpha))^k dz,$$

and the same holds for β . Using the resolvent identity and non-commutative Hölder's inequality,

$$\begin{aligned} \|(z - \iota_\alpha(Y_\alpha))^{-1} - (z - \iota_\beta(Y_\beta))^{-1}\|_{L^2(\mathcal{B})} &\leq \|(z - \iota_\alpha(Y_\alpha))^{-1}\|_{L^\infty(\mathcal{B})} \|Y_\alpha - Y_\beta\|_{L^2(\mathcal{B})} \|(z - \iota_\beta(Y_\beta))^{-1}\|_{L^\infty(\mathcal{B})} \\ &\leq \|Y_\alpha - Y_\beta\|_{L^2(\mathcal{B})}. \end{aligned}$$

Furthermore, one checks easily that $\|p(\iota_\alpha(Z_\alpha)) - p(\iota_\beta(Z_\beta))\|_{L^2(\mathcal{B})} \leq C_p \|\iota_\alpha(Z_\alpha) - \iota_\beta(Z_\beta)\|_{L^2(\mathcal{B})}$ for some constant C_p (since $\|Z_\alpha\|_{L^\infty(M_m(\mathcal{A}_\alpha))}$ is bounded by univocal constant). By estimating the difference between $p(\iota_\alpha(Z_\alpha))^k (z - 4j)(z - \iota_\alpha(Y_\alpha))^{-1} \bar{p}(\iota_\alpha(Z_\alpha))^k$ and $p(\iota_\beta(Z_\beta))^k (z - 4j)(z - \iota_\beta(Y_\beta))^{-1} \bar{p}(\iota_\beta(Z_\beta))^k$ and applying the triangle inequality for integrals, we obtain for some constant C'_p that

$$\|\iota_\alpha(X_{\alpha,j} \otimes I_m) - \iota_\beta(X_{\beta,j} \otimes I_m)\|_{L^2(\mathcal{B})} \leq C'_p \left(\|\iota_\alpha(Y_\alpha) - \iota_\beta(Y_\beta)\|_{L^2(\mathcal{B})}^2 + \|\iota_\alpha(Z_\alpha) - \iota_\beta(Z_\beta)\|_{L^2(\mathcal{B})}^2 \right)^{1/2}.$$

Since $(X_{\alpha,1} \otimes I_m, \dots, X_{\alpha,m} \otimes I_m)$ has the same non-commutative law as X_α , we obtain

$$\epsilon \leq d_W^{(2)}(\lambda_{X_\alpha}, \lambda_{X_\beta}) \leq m^{1/2} C'_p d_W^{(2)}(\lambda_{(Y_\alpha, Z_\alpha)}, \lambda_{(Y_\beta, Z_\beta)}).$$

Hence, $\{\lambda_{(Y_\alpha, Z_\alpha)}\}_{\alpha \in I}$ is $\epsilon/(m^{1/2} C'_p)$ -separated in $\Sigma_{2,4m+1}$, as desired. \square

We remark that a similar non-separability result in the context of model theory for operator algebras was shown in [4, Proposition 4.2.9]. In the model theoretic context, one often encounters triples (Ω, \mathcal{T}, d) where (Ω, \mathcal{T}) is a topological space and d is a metric on Ω that is lower semi-continuous with respect to \mathcal{T} and generates a topology that is at least as strong as \mathcal{T} ; such a triple (Ω, \mathcal{T}, d) is called a *topometric space* [9]. In particular, $\Sigma_{m,R}$ with the weak- $*$ topology and Wasserstein distance is a topometric space by Corollary 5.17 and Lemma 5.18. It was shown in [9, Proposition 3.20] that if (Ω, \mathcal{T}, d) is a topometric space and (Ω, \mathcal{T}) is second countable and locally compact, then the density character of (Ω, d) is either countable or equal to the continuum. Hence, as a corollary of Theorem 1.8, the density character of $(\Sigma_{m,R}, d_W^{(2)})$ is the continuum (of course since $\Sigma_{m,R}$ with the weak- $*$ topology is compact and metrizable, it is in particular second countable and locally compact).

6. FURTHER REMARKS

6.1. Non-commutative optimal couplings and random matrix theory. One of the motivations for our paper was the following question.

Question 6.1. *Suppose that $X^{(N)}, Y^{(N)}$ are random m -tuple of self-adjoint $N \times N$ matrices with probability distributions $\mu^{(N)}$ and $\nu^{(N)}$ respectively. Let $\mu, \nu \in \Sigma_{n,R}$. Suppose that almost surely*

$$\limsup_{N \rightarrow \infty} \|X^{(N)}\|_{L^\infty(M_N(\mathbb{C}))_{\text{sa}}}^m < R, \quad \limsup_{N \rightarrow \infty} \|Y^{(N)}\|_{L^\infty(M_N(\mathbb{C}))_{\text{sa}}}^m < R, \quad \lim_{N \rightarrow \infty} \lambda_{X^{(N)}} = \mu, \quad \lim_{N \rightarrow \infty} \lambda_{Y^{(N)}} = \nu.$$

Does the classical L^2 -Wasserstein distance of $\mu^{(N)}$ and $\nu^{(N)}$ (as probability measures on $M_N(\mathbb{C})_{\text{sa}}^m$) converge to the non-commutative L^2 -Wasserstein distance of μ and ν ?

The results of [31, 23] combined with [41] give a positive answer when $\mu^{(N)}$ is a random matrix model with density proportional to $e^{-N^2 V^{(N)}}$ where $V^{(N)} : M_N(\mathbb{C})_{\text{sa}}^m \rightarrow \mathbb{R}$ is a sufficiently regular convex function such as the trace of a non-commutative polynomial, and where $\nu^{(N)}$ has density proportional to $e^{-N^2 \|X\|_{L^2(M_N(\mathbb{C}))_{\text{sa}}}^2}$. The convexity of $V^{(N)}$ is crucial for all these arguments. By contrast, the present work shows that Question 6.1 can have a negative answer due to the obstruction of Connes-embeddability.

Proposition 6.2. *Let $X, Y \in M_n(\mathbb{C})_{\text{sa}}^m$ be matrix tuples such that an optimal coupling of λ_X and λ_Y requires a non-Connes-embeddable tracial W^* -algebra as in Corollary 5.14. Suppose $X^{(N)}$ and $Y^{(N)}$ are random (or even deterministic) elements of $M_N(\mathbb{C})_{\text{sa}}^m$ that converge in non-commutative law to X and Y . Then the classical Wasserstein distance of the probability distributions of $X^{(N)}$ and $Y^{(N)}$ on $M_N(\mathbb{C})_{\text{sa}}^m$ (with the L^2 norm associated to the normalized trace tr_N) does not converge to $d_W^{(2)}(\lambda_X, \lambda_Y)$.*

Before proving the proposition, we make some preliminary observations. Let $\Sigma_{m,R}^{\text{app}}$ denote the space of Connes-embeddable non-commutative laws in $\Sigma_{m,R}$. Let $d_{W,\text{app}}^{(2)}$ be the non-commutative Wasserstein distance on $\Sigma_{m,R}^{\text{app}}$ defined using only couplings in Connes-embeddable tracial W^* -algebras. Since $\Sigma_{m,R}^{\text{app}}$ is the weak- $*$ closure of $\Sigma_{m,R}^{\text{fin}}$, it is weak- $*$ compact, which implies the existence of optimal Connes-embeddable couplings. Moreover, the same reasoning as in Lemma 5.18 shows that $d_{W,\text{app}}^{(2)}$ is weak- $*$ lower semi-continuous. Of course, Corollary 5.14 shows that $d_{W,\text{app}}^{(2)}$ can be strictly greater than $d_W^{(2)}$ (however, we do not know whether these two metrics generate the same topology on $\Sigma_{m,R}^{\text{app}}$).

Proof of Proposition 6.2. Suppose that $X^{(N)}$ and $Y^{(N)}$ are random variables on the diffuse probability space (Ω, P) . Let $\mu^{(N)}$ and $\nu^{(N)}$ be the classical probability distributions of $X^{(N)}$ and $Y^{(N)}$ as random variables with values in the vector space $M_N(\mathbb{C})_{\text{sa}}^m$ equipped with inner product associated to tr_N . Let $\hat{\mu}^{(N)}$ and $\hat{\nu}^{(N)}$ be the non-commutative laws of $X^{(N)}$ and $Y^{(N)}$ as elements of the tracial W^* -algebra $L^\infty(\Omega, P; M_N(\mathbb{C}))$ with the trace given by $\mathbb{E} \circ \text{tr}_N$. A classical coupling of the probability distributions $\mu^{(N)}$ and $\nu^{(N)}$ on the probability space (Ω, P) can be interpreted as a non-commutative coupling on the tracial W^* -algebra $(L^\infty(\Omega, P; M_N(\mathbb{C})), \mathbb{E} \circ \text{tr}_N)$, which is Connes-embeddable. Therefore,

$$\liminf_{N \rightarrow \infty} d_W(\mu^{(N)}, \nu^{(N)}) \geq \liminf_{N \rightarrow \infty} d_{W,\text{app}}^{(2)}(\hat{\mu}^{(N)}, \hat{\nu}^{(N)}) \geq d_{W,\text{app}}^{(2)}(\lambda_X, \lambda_Y) > d_W^{(2)}(\lambda_X, \lambda_Y). \quad \square$$

This problem cannot be removed using free probabilistic regularity conditions (conditions such as finite free entropy, finite free Fisher information and so forth; see the introduction of [16] for context).

Proposition 6.3. *Again, let $X, Y \in M_n(\mathbb{C})_{\text{sa}}^m$ be as in Corollary 5.14. Let S be a free semicircular m -tuple freely independent of X and Y . Then $X + t^{1/2}S$ and $Y + t^{1/2}S$ have finite free microstate entropy (defined in [74]) and finite free Fisher information (defined in [76]). However, $d_{W,\text{app}}^{(2)}(\lambda_{X+t^{1/2}S}, \lambda_{Y+t^{1/2}S}) > d_W^{(2)}(\lambda_{X+t^{1/2}S}, \lambda_{Y+t^{1/2}S})$ for sufficiently small $t > 0$. Hence, as in Proposition 6.2, there do not exist random matrix approximations for $\lambda_{X+t^{1/2}S}$ and $\lambda_{Y+t^{1/2}S}$ whose classical Wasserstein distance converges to $d_W^{(2)}(\lambda_{X+t^{1/2}S}, \lambda_{Y+t^{1/2}S})$.*

Proof. By [77, Theorem 3.9], $X + t^{1/2}S$ and $Y + t^{1/2}S$ have finite free microstate entropy, and by [76, Corollary 6.14], they have finite free Fisher information. The free product of $M_N(\mathbb{C})$ and $W^*(S)$ is Connes-embeddable by [77, Proposition 3.3]. Hence,

$$d_W^{(2)}(\lambda_X, \lambda_{X+t^{1/2}S}) \leq d_{W,\text{app}}^{(2)}(\lambda_X, \lambda_{X+t^{1/2}S}) \leq (mt)^{1/2},$$

and the same holds with X replaced by Y . Thus, using the triangle inequality, $d_W^{(2)}(\lambda_{X+t^{1/2}S}, \lambda_{Y+t^{1/2}S}) < d_{W,\text{app}}^{(2)}(\lambda_{X+t^{1/2}S}, \lambda_{Y+t^{1/2}S})$ for sufficiently small $t > 0$, since this holds at $t = 0$. The same argument as in Proposition 6.2 rules out the possibility of the classical Wasserstein distance for random matrix models converging to $d_W^{(2)}(\lambda_{X+t^{1/2}S}, \lambda_{Y+t^{1/2}S})$. \square

Thus, at the very least, Question 6.1 needs to be reformulated using the Connes-embeddable version of the non-commutative Wasserstein distance. Even with such a modification, our results illustrate why this

question is so difficult.⁶ Indeed, in light of §5.4, random matrix models cannot converge in Wasserstein distance to the limiting non-commutative law unless that limiting law produces an approximately finite-dimensional tracial W^* -algebra. However, “good behavior” in random matrix theory and free probability often entails generating a tracial W^* -algebra that is “similar to” a free group von Neumann algebra, which is far from being approximately finite-dimensional (see e.g. [67, §XIV.3]). The random matrix question suggests a more general question in the framework of non-commutative optimal couplings.

Question 6.4. *Suppose that $\mu_n, \nu_n \in \Sigma_{m,R}$ and $\mu_n \rightarrow \mu$ and $\nu_n \rightarrow \nu$ weak-*. Under what conditions does $d_W^{(2)}(\mu_n, \nu_n) \rightarrow d_W^{(2)}(\mu, \nu)$?*

Monge-Kantorovich duality provides one avenue to attack this question. Indeed, suppose that (f_n, g_n) are admissible pairs of E -convex functions minimizing $\mu_n(f_n) + \nu_n(g_n)$. Suppose that (f, g) is an admissible pair minimizing $\mu(f) + \nu(g)$. To give a positive answer to Question 6.4, it suffices to show that $\mu_n(f_n) + \nu_n(g_n) \rightarrow \mu(f) + \nu(g)$. Suppose that we somehow show that $f_n \rightarrow f$ and $g_n \rightarrow g$ uniformly on each operator norm ball, so that $\mu_n(f_n) - \mu_n(f) \rightarrow 0$ and $\nu(g_n) - \nu(g) \rightarrow 0$.

Then it remains to show that $\mu_n(f) \rightarrow \mu(f)$ and $\nu_n(g) \rightarrow \nu(g)$. If f and g take finite values everywhere, then for each $\mathcal{A} \in \mathbb{W}$, $f^{\mathcal{A}}$ and $g^{\mathcal{A}}$ will define continuous functions on $L^2(\mathcal{A})_{sa}^m$, and in particular, $\lambda \mapsto \lambda(f)$ and $\lambda \mapsto \lambda(g)$ are continuous with respect to Wasserstein distance. However, we only assumed weak-* convergence of $\mu_n \rightarrow \mu$ and $\nu_n \rightarrow \nu$. Thus, in order to obtain the convergence of the Wasserstein distance, we would want the stronger condition that f and g are continuous with respect to convergence in law, that is, $\lambda \mapsto \lambda(f)$ and $\lambda \mapsto \lambda(g)$ are weak-* continuous on $\Sigma_{m,R}$ for each $R > 0$.

The examples of Monge-Kantorovich duality in [41, Lemma 9.10, Remark 9.11] use functions that are continuous with respect to the weak-* topology on $\Sigma_{m,R}$. However, we doubt that the optimizers (f, g) in the Monge-Kantorovich duality can always be chosen to be weak-* continuous. Nonetheless, it is worth investigating in future research how E -convex functions and Legendre transforms behave with respect to convergence in law.

6.2. Bimodule couplings and UCPT-convex functions. Another operator-algebraic analog of the idea of coupling arises from bimodules over von Neumann algebras, which have been very important in many areas of von Neumann algebras. For further background, see [12, Appendix F] and [2, §13].

Definition 6.5. If A and B are W^* -algebras, then a *Hilbert A - B -bimodule* is a Hilbert space H with an A - B -bimodule structure, such that the associated maps $A \rightarrow B(H)$ and $B \rightarrow B(H)$ are weak-* continuous. Given tracial W^* -algebras $\mathcal{A} = (A, \tau)$ and $\mathcal{B} = (B, \sigma)$ and a A - B -bimodule H , we say that a vector $\xi \in H$ is *bitracial* if $\langle \xi, a\xi \rangle = \tau(a)$ for $a \in A$ and $\langle \xi, \xi b \rangle = \sigma(b)$ for $b \in B$.

For example, suppose that there are tracial W^* -embeddings $\iota_1 : \mathcal{A} \rightarrow \mathcal{C}$ and $\iota_2 : \mathcal{B} \rightarrow \mathcal{C}$. Then $L^2(\mathcal{C})$ is a Hilbert $L^\infty(\mathcal{A})$ - $L^\infty(\mathcal{B})$ -bimodule and $\xi = \hat{1} \in L^2(\mathcal{C})$ is a bitracial vector. Thus, bimodules with bitracial vectors are a generalization of pair of tracial W^* -embeddings. In the case of a pair of embeddings ι_1 and ι_2 , there is an associated factorizable map $\iota_2^* \iota_1 : \mathcal{A} \rightarrow \mathcal{B}$. In a similar way, general $L^\infty(\mathcal{A})$ - $L^\infty(\mathcal{B})$ -bimodules with bitracial vectors correspond to general UCPT-maps.

Lemma 6.6 (See [2, §13.1.2]). *Let \mathcal{A}, \mathcal{B} be tracial W^* -algebras. If H is a Hilbert $L^\infty(\mathcal{A})$ - $L^\infty(\mathcal{B})$ -bimodule and $\xi \in H$ is a bitracial vectors, then there exists a unique $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$ such that $\langle \xi, a\xi b \rangle = \tau_{\mathcal{B}}(\Phi(a)b)$ for all $a \in L^\infty(\mathcal{A})$ and $b \in L^\infty(\mathcal{B})$. Conversely, $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$, there exists a Hilbert $L^\infty(\mathcal{A})$ - $L^\infty(\mathcal{B})$ -bimodule H and a bitracial vector ξ satisfying $\langle \xi, a\xi b \rangle = \tau_{\mathcal{B}}(\Phi(a)b)$. If we further demand that H is generated by ξ as a Hilbert $L^\infty(\mathcal{A})$ - $L^\infty(\mathcal{B})$ -bimodule, then the pair (H, ξ) is unique up to isomorphism.*

The bimodules and their associated UCPT-maps lead to an alternative notion of couplings for non-commutative random variables.

Definition 6.7. Let μ and $\nu \in \Sigma_m$ be non-commutative laws, and let (\mathcal{A}, X) and (\mathcal{B}, Y) be the GNS realizations of μ and ν respectively. A *bimodule coupling* of μ and ν is a Hilbert \mathcal{A} - \mathcal{B} -bimodule H together

⁶Questions of large- N convergence in mean field games are also extremely subtle and require regularity of the putative model for the large- N limit.

with a bitracial vector ξ . We define $C_{\text{bim}}(\mu, \nu)$ to be the supremum of $\sum_{j=1}^m \langle \xi, X_j \xi Y_j \rangle$ over all bimodule couplings of μ and ν , or equivalently,

$$C_{\text{bim}}(\mu, \nu) = \sup_{\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})} \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m}.$$

We then define $d_{\text{bim}}(\mu, \nu)$ by

$$d_{\text{bim}}(\mu, \nu)^2 = \sum_{j=1}^m \mu(x_j^2) + \sum_{j=1}^m \nu(x_j^2) - 2C_{\text{bim}}(\mu, \nu).$$

We remark that $d_{\text{bim}}(\mu, \nu)$ is the infimum of $\|X\xi - \xi Y\|$ in H^m over Hilbert $L^\infty(\mathcal{A})$ - $L^\infty(\mathcal{B})$ -bimodules with bitracial vectors (this follows from (6.1) below). Moreover, the existence of optimal bimodule couplings can be deduced from the compactness of $\text{UCPT}(\mathcal{A}, \mathcal{B})$ in the pointwise weak-* topology. The properties of C_{bim} and d_{bim} are quite similar to those of C and $d_W^{(2)}$ only with factorizable maps replaced by general UCPT maps, but we will see in Corollary 6.12 that they do not agree in general. But first, for completeness, we give proofs of some of the basic properties with the aid of the following lemma.

Lemma 6.8. *Let \mathcal{A} and \mathcal{B} be tracial W^* -functions and $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$. Let $X \in L^\infty(\mathcal{A})_{\text{sa}}^m$ and $Y \in L^\infty(\mathcal{B})_{\text{sa}}^m$ with $\|X\|_{L^\infty(\mathcal{A})_{\text{sa}}^m} \leq R$ and $\|Y\|_{L^\infty(\mathcal{B})_{\text{sa}}^m} \leq R$. Then for $i_1, \dots, i_\ell \in \{1, \dots, m\}$, we have*

$$\|\Phi(X_{i_1} \dots X_{i_\ell}) - Y_{i_1} \dots Y_{i_\ell}\|_{L^2(\mathcal{B})} \leq \ell R^{\ell-1} \left(\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - 2\langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} + \|Y\|_{L^2(\mathcal{B})_{\text{sa}}^m} \right)^{1/2}.$$

Proof. Let H be an $L^\infty(\mathcal{A})$ - $L^\infty(\mathcal{B})$ bimodule with a bitracial vector ξ such that $\langle \Phi(Z), W \rangle_{L^2(\mathcal{B})} = \langle \xi, Z\xi W \rangle_{L^2(\mathcal{B})}$ for all $Z \in L^\infty(\mathcal{A})$ and $W \in L^\infty(\mathcal{B})$. Direct computation shows that

$$\begin{aligned} (6.1) \quad \|\Phi(Z) - W\|_{L^2(\mathcal{B})}^2 &= \|\Phi(Z)\|_{L^2(\mathcal{B})}^2 - 2\text{Re}\langle \Phi(Z), W \rangle_{L^2(\mathcal{B})} + \|W\|_{L^2(\mathcal{B})}^2 \\ &\leq \|Z\|_{L^2(\mathcal{B})}^2 - 2\text{Re}\langle \Phi(Z), W \rangle_{L^2(\mathcal{B})} + \|W\|_{L^2(\mathcal{B})}^2 \\ &= \|Z\xi - \xi W\|^2. \end{aligned}$$

This implies that

$$\begin{aligned} \|\Phi(X_{i_1} \dots X_{i_\ell}) - Y_{i_1} \dots Y_{i_\ell}\|_{L^2(\mathcal{B})} &\leq \|X_{i_1} \dots X_{i_\ell} \xi - \xi Y_{i_1} \dots Y_{i_\ell}\| \\ &\leq \sum_{k=1}^{\ell} \|X_{i_1} \dots X_{i_k} \xi Y_{i_{k+1}} \dots Y_{i_\ell} - X_{i_1} \dots X_{i_{k-1}} \xi Y_{i_k} \dots Y_{i_\ell}\| \\ &\leq \sum_{k=1}^{\ell} \|X_{i_1} \dots X_{i_{k-1}}\|_{L^\infty(\mathcal{A})} \|X_{i_k} \xi - \xi Y_{i_k}\| \|Y_{i_{k+1}} \dots Y_{i_\ell}\|_{L^\infty(\mathcal{B})} \\ &\leq \ell R^{\ell-1} \|X\xi - \xi Y\| \\ &= \ell R^{\ell-1} \left(\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - 2\langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} + \|Y\|_{L^2(\mathcal{B})_{\text{sa}}^m} \right)^{1/2}. \quad \square \end{aligned}$$

Proposition 6.9. *$(\Sigma_{m,R}, d_{\text{bim}})$ is a complete metric space. If $\lambda, \mu \in \Sigma_{m,R}$, then*

$$(6.2) \quad |\lambda(x_{i_1} \dots x_{i_\ell}) - \mu(x_{i_1} \dots x_{i_\ell})| \leq \ell R^{\ell-1} d_{\text{bim}}(\lambda, \mu) \leq \ell R^{\ell-1} d_W^{(2)}(\lambda, \mu),$$

and in particular, the topology generated by d_{bim} refines the weak- topology, and the topology generated by $d_W^{(2)}$ refines the topology generated by d_{bim} . Moreover, d_{bim} is lower semi-continuous on $\Sigma_{m,R} \times \Sigma_{m,R}$ with respect to the weak-* topology.*

Proof. In the following, let λ, μ , and $\nu \in \Sigma_{m,R}$, and let (\mathcal{A}, X) , (\mathcal{B}, Y) , and (\mathcal{C}, Z) be their respective GNS realizations.

First, we prove (6.2). If $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$, then using (6.1),

$$\begin{aligned} |\lambda(x_{i_1} \dots x_{i_\ell}) - \mu(x_{i_1} \dots x_{i_\ell})| &= |\tau_{\mathcal{B}}(\Phi(X_{i_1} \dots X_{i_\ell}) - Y_{i_1} \dots Y_{i_\ell})| \\ &\leq \ell R^{\ell-1} \left(\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - 2\langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} + \|Y\|_{L^2(\mathcal{B})_{\text{sa}}^m} \right)^{1/2}. \end{aligned}$$

Taking the infimum over Φ , we obtain the first inequality of (6.2). The second inequality follows because $d_{\text{bim}}(\lambda, \mu) \leq d_W^{(2)}(\lambda, \mu)$ since $\text{FM}(\mathcal{A}, \mathcal{B}) \subseteq \text{UCPT}(\mathcal{A}, \mathcal{B})$.

Next, we show that d_{bim} is a metric on $\Sigma_{m,R}$ (postponing the proof of completeness to the end). Clearly, $d_{\text{bim}}(\lambda, \mu) \geq 0$. If $d_{\text{bim}}(\lambda, \mu) = 0$, then by (6.2), we have $\lambda = \mu$. Because every UCPT map has a UCPT adjoint, we have

$$C_{\text{bim}}(\lambda, \mu) = \sup_{\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})} \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} = \sup_{\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})} \langle X, \Phi^*(Y) \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} = C_{\text{bim}}(\mu, \lambda),$$

and hence $d_{\text{bim}}(\lambda, \mu) = d_{\text{bim}}(\mu, \lambda)$. To prove the triangle inequality, we use the fact that UCPT maps are closed under composition.⁷ Let $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$ and $\Psi \in \text{UCPT}(\mathcal{B}, \mathcal{C})$ be UCPT maps corresponding to optimal bimodule couplings between λ and μ and between μ and ν respectively, so that

$$\begin{aligned} d_{\text{bim}}(\lambda, \mu)^2 &= \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - 2\langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} + \|Y\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \\ &= \left(\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - \|\Phi(X)\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \right) + \|\Phi(X) - Y\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \\ &\geq \|\Phi(X) - Y\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \end{aligned}$$

and

$$\begin{aligned} d_{\text{bim}}(\mu, \nu)^2 &= \|Z\|_{L^2(\mathcal{C})_{\text{sa}}^m}^2 - 2\langle \Psi^*(Z), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} + \|Y\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \\ &= \left(\|Z\|_{L^2(\mathcal{C})_{\text{sa}}^m}^2 - \|\Psi^*(Z)\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \right) + \|\Psi^*(Z) - Y\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \\ &\geq \|\Psi^*(Z) - Y\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2. \end{aligned}$$

Then

$$\begin{aligned} d_{\text{bim}}(\lambda, \nu)^2 &\leq \|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - 2\langle \Psi \circ \Phi(X), Z \rangle_{L^2(\mathcal{C})_{\text{sa}}^m} + \|Z\|_{L^2(\mathcal{C})_{\text{sa}}^m}^2 \\ &= \left(\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - \|\Phi(X)\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \right) + \|\Phi(X) - \Psi^*(Z)\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 + \left(\|Z\|_{L^2(\mathcal{C})_{\text{sa}}^m}^2 - \|\Psi^*(Z)\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \right) \\ &\leq \left(\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - \|\Phi(X)\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \right) + \|\Phi(X) - Y\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \\ &\quad + 2\|\Phi(X) - Y\|_{L^2(\mathcal{B})_{\text{sa}}^m} \|\Psi^*(Z) - Y\|_{L^2(\mathcal{B})_{\text{sa}}^m} \\ &\quad + \|\Psi^*(Z) - Y\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 + \left(\|Z\|_{L^2(\mathcal{C})_{\text{sa}}^m}^2 - \|\Psi^*(Z)\|_{L^2(\mathcal{B})_{\text{sa}}^m}^2 \right) \\ &\leq d_{\text{bim}}(\lambda, \mu)^2 + 2d_{\text{bim}}(\lambda, \mu)d_{\text{bim}}(\mu, \nu) + d_{\text{bim}}(\mu, \nu)^2. \end{aligned}$$

It follows from (6.2) that the d_{bim} -topology refines the weak-* topology, and the Wasserstein topology refines the d_{bim} -topology.

Next, we show that d_{bim} is lower semi-continuous with respect to the weak-* topology. Fix a non-principal ultrafilter \mathcal{U} on \mathbb{N} , and suppose that $(\lambda_n)_{n \in \mathbb{N}}$ and $(\mu_n)_{n \in \mathbb{N}}$ are sequences in $\Sigma_{m,R}$ and (\mathcal{A}_n, X_n) and (\mathcal{B}_n, Y_n) are their respective GNS realizations. Let $\lambda = \lim_{n \rightarrow \mathcal{U}} \lambda_n$ and $\mu = \lim_{n \rightarrow \mathcal{U}} \mu_n$. Let $\mathcal{A} = \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ and $\mathcal{B} = \prod_{n \rightarrow \mathcal{U}} \mathcal{B}_n$. Let $X = [X_n]_{n \in \mathbb{N}} \in L^2(\mathcal{A})_{\text{sa}}^m$ and $[Y_n]_{n \in \mathbb{N}} \in L^2(\mathcal{B})_{\text{sa}}^m$. By Lemma 5.10, X and Y have non-commutative laws λ and μ respectively. Let $\Phi_n \in \text{UCPT}(\mathcal{A}_n, \mathcal{B}_n)$ such that $C_{\text{bim}}(\lambda_n, \mu_n) = \langle \Phi(X_n), Y_n \rangle_{L^2(\mathcal{B}_n)_{\text{sa}}^m}$. If $(Z_n)_{n \in \mathbb{N}}$ and $(Z'_n)_{n \in \mathbb{N}}$ are sequences in $\prod_{n \in \mathbb{N}} \mathcal{A}_n$ and if $\lim_{n \rightarrow \mathcal{U}} \|Z_n - Z'_n\|_{L^2(\mathcal{A}_n)} = 0$, then $\lim_{n \rightarrow \mathcal{U}} \|\Phi_n(Z_n) - \Phi_n(Z'_n)\|_{L^2(\mathcal{B}_n)} = 0$ because each Φ_n is a contraction with respect to the L^2 norms on \mathcal{A} and \mathcal{B} . Therefore, the equivalence class $[\Phi_n(Z_n)]_{n \in \mathbb{N}}$ in \mathcal{B} only depends on the equivalence class $[Z_n]_{n \in \mathbb{N}}$ in \mathcal{A} , so that the sequence Φ_n produces a well-defined map $\Phi : \mathcal{A} \rightarrow \mathcal{B}$. It is straightforward to check that $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$. Let $\Phi' : W^*(X) \rightarrow W^*(Y)$ be the composition of the inclusion $W^*(X) \rightarrow \mathcal{A}$, the map $\Phi : \mathcal{A} \rightarrow \mathcal{B}$, and the trace-preserving conditional expectation $\mathcal{B} \rightarrow W^*(Y)$. Then

$$C_{\text{bim}}(\lambda, \mu) \geq \langle \Phi'(X), Y \rangle_{L^2(W^*(Y))_{\text{sa}}^m} = \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} = \lim_{n \rightarrow \mathcal{U}} \langle \Phi_n(X_n), Y_n \rangle_{L^2(\mathcal{B}_n)} = \lim_{n \rightarrow \mathcal{U}} C_{\text{bim}}(\lambda_n, \mu_n).$$

This implies that $d_{\text{bim}}(\lambda, \mu) \leq \lim_{n \rightarrow \mathcal{U}} d_{\text{bim}}(\lambda_n, \mu_n)$, so d_{bim} is weak-* lower semi-continuous as desired.

⁷The corresponding notion for bimodules is the *Connes fusion*, and the proof of the triangle inequality is quite natural from this viewpoint; however, we will use a more elementary argument.

Finally, we show that $(\Sigma_{m,R}, d_{\text{bim}})$ is complete. Let $(\lambda_n)_{n \in \mathbb{N}}$ be a Cauchy sequence with respect to d_{bim} . Using (6.2), for each $i_1, \dots, i_\ell \in \{1, \dots, m\}$, the sequence $(\lambda_n(x_{i_1} \dots x_{i_\ell}))_{n \in \mathbb{N}}$ is Cauchy and hence converges in \mathbb{C} to some limit $\lambda(x_{i_1} \dots x_{i_\ell})$. Extend λ linearly to a map on $\mathbb{C}\langle x_1, \dots, x_m \rangle \rightarrow \mathbb{C}$, and then it is straightforward to check that $\lambda \in \Sigma_{m,R}$ using Definition 2.25. Then because d_{bim} is weak-* lower semi-continuous,

$$d_{\text{bim}}(\lambda_n, \lambda) \leq \liminf_{k \rightarrow \infty} d_{\text{bim}}(\lambda_n, \lambda_k) \leq \sup_{k \geq n} d_{\text{bim}}(\lambda_n, \lambda_k).$$

The right-hand side goes to zero as $n \rightarrow \infty$ because $(\lambda_n)_{n \in \mathbb{N}}$ was assumed to be Cauchy in d_{bim} . This shows that $\lambda_n \rightarrow \lambda$ in d_{bim} as desired. \square

We saw in the preceding argument that $C_{\text{bim}}(\mu, \nu) \geq C(\mu, \nu)$. In the commutative setting, we have equality by a similar argument as in [11, Theorem 1.5]. (For further discussion of bimodules over commutative tracial W^* -algebras, see [2, Example 13.1.2].)

Lemma 6.10. *Let μ and $\nu \in \Sigma_{m,R}$ be non-commutative laws that can be realized by elements of commutative tracial W^* -algebras. Then $C_{\text{bim}}(\mu, \nu) = C(\mu, \nu)$, and there exists an optimal coupling in a commutative tracial W^* -algebra.*

Proof. Let (\mathcal{A}, X) and (\mathcal{B}, Y) be the GNS realizations of μ and ν . Consider an optimal bimodule coupling given by a Hilbert \mathcal{A} - \mathcal{B} -bimodule H and a bitracial vector $\xi \in H$. Let $X'_j \in B(H)$ be the operator of left multiplication by X_j , and let $Y_j \in B(H)$ be the operator of right multiplication by Y_j . Let M be the W^* -subalgebra of $B(H)$ generated by $X' = (X'_1, \dots, X'_m)$ and $Y' = (Y'_1, \dots, Y'_m)$. Since X'_i and Y'_j commute and X'_i and X'_j commute and Y'_i and Y'_j commute, M is commutative. Let $\tau : M \rightarrow \mathbb{C}$ be the map $T \mapsto \langle \xi, T\xi \rangle$. Since M is commutative, τ is a trace (it is a state and satisfies $\tau(ab) = \tau(ba)$). We have not shown that it is normal or faithful, but nonetheless, the map $\gamma = \lambda_{(X', Y')} : \mathbb{C}\langle x_1, \dots, x_{2m} \rangle \rightarrow \mathbb{C}$ given by $p \mapsto \tau(p(X', Y'))$ is still an element of $\Sigma_{2m,R}$ according to Definition 2.25. Moreover, since ξ was a bitracial vector for \mathcal{A} and \mathcal{B} , we have $\tau(p(X')) = \tau_{\mathcal{A}}(p(X)) = \mu(p)$ and $\tau(p(Y')) = \tau_{\mathcal{B}}(p(Y)) = \nu(p)$. Therefore, γ has the marginals μ and ν . If $(\mathcal{C}, (\widehat{X}, \widehat{Y}))$ is the GNS realization of γ , then \mathcal{C} is commutative because for any non-commutative polynomials p and q in $2m$ variables,

$$\|(pq - qp)(\widehat{X}, \widehat{Y})\|_{L^2(\mathcal{C})}^2 = \gamma[(pq - qp)^*(pq - qp)] = \tau((pq - qp)^*(pq - qp)(X', Y')) = 0,$$

and non-commutative polynomials of X and Y are dense in $L^2(\mathcal{C})$ (by Lemma 2.34). Moreover,

$$\langle \widehat{X}, \widehat{Y} \rangle_{L^2(\mathcal{C})_{\text{sa}}^m} = \sum_{j=1}^m \gamma(x_j x_{m+j}) = \sum_{j=1}^m \tau(X'_j Y'_j) = \sum_{j=1}^m \langle \xi, X_j \xi Y_j \rangle.$$

Hence, $(\mathcal{C}, (\widehat{X}, \widehat{Y}))$ is a coupling in a commutative tracial W^* -algebra which is also an optimal bimodule coupling of μ and ν . \square

For general non-commutative laws, the inequality $C(\mu, \nu) \leq C_{\text{bim}}(\mu, \nu)$ can be strict, even for non-commutative laws of matrix tuples. We can deduce this from another result of Haagerup and Musat that $\text{FM}(M_n(\mathbb{C}), M_n(\mathbb{C}))$ is in general strictly smaller than $\text{UCPT}(M_n(\mathbb{C}), M_n(\mathbb{C}))$, and in particular there is an explicit non-factorizable UCPT map on $M_3(\mathbb{C})$.

Theorem 6.11 (Haagerup-Musat [32, Example 3.1], [33, Theorems 5.2 and 5.6]). *For $n > 1$, let $W_n^- : M_n(\mathbb{C}) \rightarrow M_n(\mathbb{C})$ be the Holevo-Werner channel $W_n^-(x) = \frac{1}{n-1}(\text{Tr}_n(x) - x^t)$. Then W_n^- is a UCPT map, and it is factorizable if and only if $n \neq 3$.*

Combining non-factorizability of W_3^- with Lemma 5.7 similarly to the proof of Corollary 5.14, we deduce the following corollary.

Corollary 6.12. *There exist $X, Y \in M_3(\mathbb{C})_{\text{sa}}^9$ such that $C_{\text{bim}}(\lambda_X, \lambda_Y) > C(\lambda_X, \lambda_Y)$.*

This shows that the metrics d_{bim} and $d_W^{(2)}$ are distinct. It is unclear to us whether d_{bim} and $d_W^{(2)}$ generate the same topology. However, the results of §5.4 about the Wasserstein distance adapt to the UCPT setting without much difficulty. For instance, we have the following analog of Lemma 5.16.

Lemma 6.13. *Let $(\mu_n)_{n \in \mathbb{N}}$ and μ be non-commutative laws. Let (\mathcal{A}, X) be the GNS realization of μ . Let \mathcal{A}_n be a tracial W^* -algebra and $X_n \in L^\infty(\mathcal{A}_n)_{\text{sa}}^m$ such that $\lambda_{X_n} = \mu_n$. Then the following are equivalent:*

- (1) $\lim_{n \rightarrow \mathcal{U}} d_{\text{bim}}(\mu_n, \mu) = 0$.
- (2) *There exists a tracial W^* -embedding $\phi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ and there exists $\Phi_n \in \text{UCPT}(\mathcal{A}, \mathcal{A}_n)$ such that*

$$\phi(X) = [X_n]_{n \in \mathbb{N}}, \quad \phi(Z) = [\Phi_n(Z)]_{n \in \mathbb{N}} \text{ for all } Z \in L^\infty(\mathcal{A}).$$

Proof. (1) \implies (2). By Lemma 5.10, there is a tracial W^* -embedding $\phi : \mathcal{A} \rightarrow \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ with $\phi(X) = [X_n]_{n \in \mathbb{N}}$. Let $\Phi_n \in \text{UCPT}(\mathcal{A}, \mathcal{A}_n)$ such that $\langle \Phi_n(X), X_n \rangle_{L^2(\mathcal{A}_n)_{\text{sa}}^m} = C_{\text{bim}}(\mu_n, \mu)$. As in the previous lemma, there exists $\Phi \in \text{UCPT}(\mathcal{A}, \prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n)$ such that

$$\Phi(Z) = [\Phi_n(Z)]_{n \in \mathbb{N}} \text{ for all } Z \in L^\infty(\mathcal{A}).$$

It remains to show that $\Phi = \phi$. Let $X_n = (X_n^{(1)}, \dots, X_n^{(m)})$ and $X = (X^{(1)}, \dots, X^{(m)})$. Using (6.1), for every $i_1, \dots, i_\ell \in \{1, \dots, m\}$, we have

$$\begin{aligned} \|\Phi_n(X^{(i_1)}, \dots, X^{(i_\ell)}) - X_n^{(i_1)}, \dots, X_n^{(i_\ell)}\|_{L^2(\mathcal{A}_n)} &\leq \ell R^{\ell-1} \left(\|X\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 - 2\langle \Phi_n(X), X_n \rangle + \|X_n\|_{L^2(\mathcal{A})_{\text{sa}}^m}^2 \right)^{1/2} \\ &= \ell R^{\ell-1} d_{\text{bim}}(\mu_n, \mu). \end{aligned}$$

Taking $n \rightarrow \mathcal{U}$, we obtain

$$\|\Phi(X^{(i_1)}, \dots, X^{(i_\ell)}) - \phi(X^{(i_1)}, \dots, X^{(i_\ell)})\|_{L^2(\prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n)} \leq \lim_{n \rightarrow \mathcal{U}} \ell R^{\ell-1} d_{\text{bim}}(\mu_n, \mu) = 0.$$

Hence, $\Phi(p(X)) = \phi(p(X))$ for every non-commutative polynomial p . Since non-commutative polynomials are in X are dense in $L^2(\mathcal{A})$ and Φ and ϕ are both contractions with respect to the L^2 norm, we have $\Phi = \phi$.

- (2) \implies (1). The proof is the same as in Lemma 5.16, so we leave the details to the reader. \square

In a completely analogous way to Proposition 5.21, one can deduce that the weak- $*$ and d_{bim} topologies agree at some point $\mu \in \Sigma_{m,R}$ if and only if the corresponding tracial W^* -algebra \mathcal{A} obtained from the GNS construction is UCPT-stable, meaning that every tracial W^* -algebra embedding from \mathcal{A} into some ultraproduct $\prod_{n \rightarrow \mathcal{U}} \mathcal{A}_n$ of tracial W^* -algebras lifts to a sequence $(\Phi_n)_{n \in \mathbb{N}}$ where $\Phi_n \in \text{UCPT}(\mathcal{A}, \mathcal{A}_n)$. Furthermore, if \mathcal{A} is Connes-embeddable, then these two conditions are also equivalent to \mathcal{A} being approximately finite-dimensional; the proof is essentially the same as that of [5, Theorem 2.6] or that of Proposition 5.26. However, it is unknown how FM-stability and UCPT-stability are related in the non-Connes-embeddable setting.

To circle back to Monge-Kantorovich duality, given the relationship of optimal couplings with factorizable maps on the one hand and E -convex functions on the other hand, one might wonder whether there is an alternative version of the theory of convex functions and Legendre transforms that is based on UCPT maps rather than factorizable maps. Indeed, this is possible, and we will sketch here some of the basic properties and the parts of the proof that are different from the E -convex case.

Definition 6.14. A W^* -function with values in $[-\infty, \infty]$ is UCPT-convex if either f is identically $-\infty$, or else for every \mathcal{A} , $f^{\mathcal{A}}$ is a convex and lower semi-continuous function with values in $(-\infty, \infty]$, and we have $f^{\mathcal{A}}(X) \leq f^{\mathcal{B}}(\Phi(X))$ for every $\mathcal{A}, \mathcal{B} \in \mathbb{W}$ and $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$.

Definition 6.15. The UCPT-Legendre transform of a tracial W^* -function f is the tracial W^* -function $\mathcal{K}f$ given by

$$(\mathcal{K}f)^{\mathcal{A}}(X) = \sup_{\substack{\mathcal{B} \in \mathbb{W} \\ \Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B}) \\ Y \in L^2(\mathcal{B})_{\text{sa}}^m}} \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - f^{\mathcal{B}}(Y).$$

We have the following analog of Proposition 3.17.

Proposition 6.16. *If f, g be a tracial W^* -functions.*

- (1) $\mathcal{K}f$ is UCPT-convex.
- (2) If $f \leq g$, then $\mathcal{K}f \geq \mathcal{K}g$.
- (3) We have $\mathcal{K}^2 f \leq f$ with equality if and only if f is UCPT-convex.

(4) $\mathcal{K}^2 f$ is the maximal UCPT-convex function that is less than or equal to f .

The proof is essentially the same as that of Proposition 3.17, modulo the necessary changes to work with UCPT maps rather than tracial W^* -embeddings and conditional expectations. For instance, to show monotonicity of $\mathcal{K}f$ under UCPT maps, suppose that $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$. If $\Psi \in \text{UCPT}(\mathcal{B}, \mathcal{C})$, then $\Psi \circ \Phi \in \text{UCPT}(\mathcal{A}, \mathcal{C})$. Therefore,

$$\mathcal{K}f^{\mathcal{A}}(X) \geq \sup_{\substack{\mathcal{C} \in \mathbb{W} \\ \Psi \in \text{UCPT}(\mathcal{B}, \mathcal{C}) \\ Y \in L^2(\mathcal{C})_{\text{sa}}^m}} (\langle \Psi \circ \Phi(X), Y \rangle - f^{\mathcal{C}}(Y)) = \mathcal{K}f^{\mathcal{B}}(\Phi(X)).$$

The relationship between the UCPT Legendre transform and the E -convex Legendre transform is as follows (compare the relationship between the E -convex Legendre transform and the Hilbert-space Legendre transform).

Corollary 6.17. *Let f be a tracial W^* -function.*

- (1) *If f is UCPT-convex, then f is E -convex.*
- (2) *$\mathcal{K}f \geq \mathcal{L}f$.*
- (3) *$\mathcal{K}^2 f \leq \mathcal{L}^2 f$.*
- (4) *If f is UCPT-convex, then $\mathcal{K}f = \mathcal{L}f$.*

Proof. (1) and (2) are immediate from the definitions of \mathcal{L} and \mathcal{K} since every tracial W^* -embedding is a UCPT map.

(3) Observe that $\mathcal{K}^2 f$ is E -convex by (1) and $\mathcal{K}^2 f \leq f$. Therefore, Proposition 3.17 (4) implies that $\mathcal{K}^2 f \leq \mathcal{L}^2 f$.

(4) We already know that $\mathcal{L}f \leq \mathcal{K}f$. For the reverse inequality, the idea is already contained in the proof of Proposition 6.16 (3). Note that for $\mathcal{A}, \mathcal{B} \in \mathbb{W}$ and $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$ and $Y \in L^2(\mathcal{B})_{\text{sa}}^m$, we have

$$\langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m} - f^{\mathcal{B}}(Y) \leq \langle X, \Phi^*(Y) \rangle_{L^2(\mathcal{A})_{\text{sa}}^m} - f^{\mathcal{A}}(\Phi^*(Y)) \leq \mathcal{L}f^{\mathcal{A}}(X).$$

Taking the supremum over \mathcal{B} , Φ , and Y , we obtain $\mathcal{K}f \leq \mathcal{L}f$. □

The UCPT-analog of Monge-Kantorovich duality is as follows.

Definition 6.18. A pair of tracial W^* -functions (f, g) with values in $(-\infty, \infty]$ is said to be UCPT-admissible if for every $\mathcal{A}, \mathcal{B} \in \mathbb{W}$ and $X \in L^2(\mathcal{A})_{\text{sa}}^m$ and $Y \in L^2(\mathcal{B})_{\text{sa}}^m$ and $\Phi \in \text{UCPT}(\mathcal{A}, \mathcal{B})$, we have

$$f^{\mathcal{A}}(X) + g^{\mathcal{B}}(Y) \geq \langle \Phi(X), Y \rangle_{L^2(\mathcal{B})_{\text{sa}}^m}.$$

Proposition 6.19. *$C_{\text{bim}}(\mu, \nu)$ is equal to the infimum of $\mu(f) + \nu(g)$ over all UCPT-admissible pairs of tracial W^* -functions, as well as the infimum of $\mu(f) + \nu(g)$ over all UCPT-admissible pairs of UCPT-convex functions.*

The proof is the same as that of Proposition 3.22; similarly, there is an UCPT analog of Proposition 3.23. However, although there is an analog of Monge-Kantorovich duality, there are many questions about bimodule couplings for which the answer is not immediately clear:

- Is there a bimodule analog of the displacement interpolation?
- Is there a bimodule analog of the L^p Wasserstein distance for $p \neq 2$?
- Is there a useful subgradient characterization of UCPT-convexity analogous to Lemma 3.10?
- Do d_{bim} and $d_W^{(2)}$ generate the same topology on $\Sigma_{m, R}$?

APPENDIX A. NON-COMMUTATIVE LAWS AND COUPLINGS FOR L^p VARIABLES

Although we have focused in this paper on the non-commutative L^2 -Wasserstein distance, Biane and Voiculescu [11] also defined L^p Wasserstein distance for $p \in [1, \infty)$. Although they only defined the Wasserstein distance for tuples of bounded operators, it is natural to extend the theory to non-commutative L^p spaces. In this section, after reviewing the properties of affiliated operators to a tracial W^* -algebra, we define laws, couplings, and L^p Wasserstein distance for m -tuples of self-adjoint operators in non-commutative L^p space, and show the existence of optimal couplings and Wasserstein geodesics.

A.1. Affiliated operators, L^p spaces. For background on unbounded operators, refer for instance to [62, §VIII]. We recall that if H is a Hilbert and $T : H \supseteq \text{dom}(T) \rightarrow H$ is a closed densely defined unbounded operator, then T has a polar decomposition as $U|T|$ where $|T|$ is a positive self-adjoint operator with $\text{dom}(|T|) = \text{dom}(T)$ and U is a partial isometry [62, §VIII.9].

We quote without proof the basic definitions and results about affiliated operators and non-commutative L^p spaces. Affiliated operators were first studied by Murray and von Neumann in [50, §XVI], and the non-commutative L^p spaces were studied in [24]. For a modern exposition of affiliated operators and L^p spaces in English, see [20] as well as [2, §7.2].

Let $\mathcal{A} = (A, \tau)$ be a tracial W^* -algebra. To avoid ambiguity, we use the notation $H_{\mathcal{A}}$ rather than $L^2(\mathcal{A})$ for the completion of A with respect to the inner product $(a, b) \mapsto \tau(a^*b)$. We still use the notation \widehat{a} for the element of $H_{\mathcal{A}}$ corresponding to $a \in A$.

Definition A.1. Let $\mathcal{A} = (A, \tau)$ be a tracial W^* -algebra, and let us view A as a subset of $B(H_{\mathcal{A}})$ as in Theorem 2.8. A closed densely defined operator $T : \text{dom}(T) \rightarrow H_{\mathcal{A}}$ with polar decomposition $U|T|$ is *affiliated to A* if $U \in A$ and the spectral projection $1_S(|T|) \in A$ for every Borel set $S \subseteq [0, \infty)$. We denote the set of affiliated operators by $\text{Aff}(\mathcal{A})$.

Example A.2. Let (Ω, P) be a probability space and let \mathcal{A} be $L^\infty(\Omega, P)$ equipped with the trace given by integration against P . Then $\text{Aff}(\mathcal{A})$ can be canonically identified with measurable functions on Ω that are finite almost everywhere, viewed as unbounded multiplication operators on $L^2(\Omega, P)$.

Theorem A.3. Let $\mathcal{A} = (A, \tau)$ be a tracial W^* -algebra.

- (1) $\text{Aff}(\mathcal{A})$ satisfies the following properties:
 - If $T \in \text{Aff}(\mathcal{A})$, then $T^* \in \text{Aff}(\mathcal{A})$.
 - If $T_1, T_2 \in \text{Aff}(\mathcal{A})$, then $T_1|_{\text{dom}(T_1) \cap \text{dom}(T_2)} + T_2|_{\text{dom}(T_1) + \text{dom}(T_2)}$ is closeable, and its closure is in $\text{Aff}(\mathcal{A})$.
 - If $T_1, T_2 \in \text{Aff}(\mathcal{A})$, then $T_1 T_2|_{T_2^{-1}(\text{dom}(T_1))}$ is closeable and its closure is in $\text{Aff}(\mathcal{A})$.
- In this way, $\text{Aff}(\mathcal{A})$ can be equipped with the structure of a $*$ -algebra.
- (2) The canonical inclusion $A \rightarrow \text{Aff}(\mathcal{A})$ is a $*$ -homomorphism. Moreover, if $T \in \text{Aff}(\mathcal{A})$ is a bounded operator, then $T \in A$.
- (3) If $T \in \text{Aff}(\mathcal{A})$ is a normal operator, and f is a Borel function on its spectrum, then $f(T) \in \text{Aff}(\mathcal{A})$. If f is bounded, then $f(T) \in A$. There is a unique probability measure μ_T on \mathbb{C} such that $\tau(f(T)) = \int f d\mu_T$ for all bounded Borel functions f . The spectrum of T is exactly the closed support of μ_T .
- (4) Let $\text{Aff}(\mathcal{A})_+$ be the set of positive operators affiliated to $\mathcal{A} = (A, \tau)$. Then τ extends to a map $\text{Aff}(\mathcal{A})_+ \rightarrow [0, \infty]$ satisfying

$$\tau(T) = \lim_{n \rightarrow \infty} \tau(f_n(T)),$$

whenever f_n is any sequence of nonnegative Borel functions increasing to the identity function on $[0, \infty)$.

Definition A.4. For a tracial W^* -algebra $\mathcal{A} = (A, \tau)$ and $p \in [1, \infty)$, we define

$$L^p(\mathcal{A}) = \{T \in \text{Aff}(\mathcal{A}) : \tau(|T|^p) < \infty\},$$

and we write

$$\|T\|_{L^p(\mathcal{A})} = \tau(|T|^p)^{1/p}.$$

As stated above, $L^\infty(\mathcal{A}) = A$ and $\|T\|_{L^\infty(\mathcal{A})}$ is the norm on A .

Theorem A.5. Let $\mathcal{A} = (A, \tau)$ as above.

- (1) For $p \in [1, \infty]$, $\|\cdot\|_p$ defines a norm on $L^p(\mathcal{A})$, and $L^p(\mathcal{A})$ is a complete with respect to this norm, hence it is a Banach space.
- (2) A is a dense subspace of $L^p(\mathcal{A})$ for $p \in [1, \infty)$.
- (3) Let $p, p_1, p_2 \in [1, \infty]$ with $1/p = 1/p_1 + 1/p_2$. If $T_1 \in L^{p_1}(\mathcal{A})$ and $T_2 \in L^{p_2}(\mathcal{A})$, then $T_1 T_2 \in L^p(\mathcal{A})$ and $\|T_1 T_2\|_p \leq \|T_1\|_{p_1} \|T_2\|_{p_2}$.
- (4) For $p \in [1, \infty]$, if $T \in L^p(\mathcal{A})$, then $T^* \in L^p(\mathcal{A})$ with $\|T\|_p = \|T^*\|_p$.

- (5) τ extends uniquely to a bounded map $L^1(\mathcal{A}) \rightarrow \mathbb{C}$ (still denoted by τ or $\tau_{\mathcal{A}}$) that satisfies $\tau(a^*) = \overline{\tau(a)}$
- (6) Let $p \in [1, \infty)$ and let $1/p + 1/q = 1$. Then $L^q(\mathcal{A})$ may be canonically identified with the dual of $L^p(\mathcal{A})$ through the pairing $(T_1, T_2) \mapsto \tau(T_1 T_2)$. In particular, this yields an identification between $L^1(\mathcal{A})$ and A_* .
- (7) If A_* is separable, then $L^p(\mathcal{A})$ is separable for $p \in [1, \infty)$.

Theorem A.6. Let $\mathcal{A} = (A, \tau)$, and $T \in \text{Aff}(\mathcal{A})$. Then $T \in L^2(\mathcal{A})$ if and only if $\widehat{1} \in \text{dom}(T)$. There is unitary isomorphism of Hilbert spaces $\phi : L^2(\mathcal{A}) \rightarrow H_{\mathcal{A}}$ given by $T \mapsto T\widehat{1}$. Furthermore, for $T \in L^2(\mathcal{A})$ and $a \in A$, we have $\phi(aT) = a\phi(T)$.

The next lemma can be deduced from well-known facts about C^* -algebras as well as the properties of affiliated operators and L^p spaces in the previous section.

Lemma A.7. Let $\iota : \mathcal{A} \rightarrow \mathcal{B}$ be a tracial W^* -embedding. Then $\|\iota(a)\| = \|a\|$ for $a \in A$. Moreover, for $a \in \mathcal{A}$ and $p \in [1, \infty)$, we have $\iota((a^*a)^{p/2}) = (\iota(a)^*\iota(a))^{p/2}$. Hence, $\|\iota(a)\|_{L^p(\mathcal{B})} = \|a\|_{L^p(\mathcal{A})}$ for $a \in \mathcal{A}$, and therefore, ι extends to an isometric linear map $L^p(\mathcal{A}) \rightarrow L^p(\mathcal{B})$ for every $p \in [1, \infty)$. In fact, ι extends to an injective $*$ -homomorphism $\text{Aff}(\mathcal{A}) \rightarrow \text{Aff}(\mathcal{B})$.

Notation A.8. If $\iota : \mathcal{A} \rightarrow \mathcal{B}$ is a tracial W^* -embedding, we will denote the extended map $\text{Aff}(\mathcal{A}) \rightarrow \text{Aff}(\mathcal{B})$ also by ι .

Proposition A.9. Let $\iota : \mathcal{A} = (A, \tau) \rightarrow \mathcal{B} = (B, \sigma)$ be a tracial W^* -embedding. Let $E : L^2(\mathcal{B}) \rightarrow L^2(\mathcal{A})$ be the adjoint of the map $\iota : L^2(\mathcal{A}) \rightarrow L^2(\mathcal{B})$.

- (1) For $p \in [1, \infty]$, E to a unique bounded linear map $L^p(\mathcal{B}) \rightarrow L^p(\mathcal{A})$, and (denoting the extended map still by E) we have $\|E(b)\|_{L^p(\mathcal{A})} \leq \|b\|_{L^p(\mathcal{B})}$.
- (2) For all $b \in L^1(\mathcal{B})$, we have $\tau(E(b)) = \sigma(b)$; in other words, E is trace-preserving.
- (3) For all $b \in L^1(\mathcal{B})$, we have $E(b^*) = E(b)^*$.
- (4) If $b \in L^p(\mathcal{B})$ and $a \in L^q(\mathcal{A})$ with $1/p + 1/q = 1$, then $E[\iota(a)b] = aE[b]$ and $E[b\iota(a)] = E[b]a$.

Sketch of proof. (1) Fix $p, q \in [1, \infty]$ with $1/p + 1/q = 1$. If $a \in L^\infty(\mathcal{A})$ and $b \in L^\infty(\mathcal{B})$, we have

$$|\langle a, E[b] \rangle_{\mathcal{A}}| = |\langle \iota(a), b \rangle_{L^2(\mathcal{B})_{\text{sa}}}^m| \leq \|\iota(a)\|_{L^q(\mathcal{B})} \|b\|_{L^p(\mathcal{B})} = \|a\|_{L^q(\mathcal{A})} \|b\|_{L^p(\mathcal{B})}.$$

By density of $L^\infty(\mathcal{A})$ in $L^q(\mathcal{A})$ and the duality of $L^q(\mathcal{A})$ and $L^p(\mathcal{A})$, we obtain $\|E[b]\|_{L^p(\mathcal{A})} \leq \|b\|_{L^p(\mathcal{B})}$, and the extension follows from this.

(2) Since $\iota(1) = 1$ and $E = \iota^*$, we obtain $\tau(E(b)) = \sigma(b)$ for $b \in L^2(\mathcal{B})$ and this extends to $L^1(\mathcal{B})$ by density.

(3) (4) The claims are first checked for $b \in L^\infty(\mathcal{B})$ and $a \in L^\infty(\mathcal{A})$ using the properties of the trace and the fact that E is the adjoint of ι , by similar reasoning as in Lemma 1.17. Then we use density of L^∞ in L^p to conclude. \square

Notation A.10. Let $\mathcal{A} = (A, \tau)$ be a tracial W^* -algebra. Let $X = (X_1, \dots, X_m) \in \text{Aff}(\mathcal{A})^m$. We denote by $W^*(X)$ the smallest W^* -subalgebra of A to which X_1, \dots, X_m are affiliated operators. Equivalently, letting $X_j = U_j|X_j|$ be a polar decomposition of X_j , $W^*(X)$ is the weak- $*$ closure of the $*$ -algebra generated by U_j and $f(|X_j|)$ for bounded Borel functions $f : \mathbb{R} \rightarrow \mathbb{C}$ and $j = 1, \dots, m$. We view $W^*(X)$ as a tracial W^* -algebra, where the trace is the simply the restriction of τ .

Notation A.11. Let $\mathcal{A} = (A, \tau)$ and $\mathcal{B} = (B, \sigma)$ be tracial W^* -algebras. For $p \in [1, \infty]$, we equip $L^p(\mathcal{A})^m$ with the norm

$$\|(X_1, \dots, X_m)\|_{L^p(\mathcal{A})^m} = \begin{cases} \left(\sum_{j=1}^m \tau((X_j^* X_j)^{p/2}) \right)^{1/p}, & p < \infty, \\ \max_{j=1, \dots, m} \|X_j\|_A, & p = \infty \end{cases}$$

Note that $L^p(\mathcal{A})_{\text{sa}}^m$ is a real subspace of $L^p(\mathcal{A})^m$. For $X = (X_1, \dots, X_m)$ and $Y = (Y_1, \dots, Y_m)$ in $L^2(\mathcal{A})^m$, we define

$$\langle X, Y \rangle_{L^2(\mathcal{A})^m} = \sum_{j=1}^m \tau(X_j^* Y_j) = \sum_{j=1}^m \langle X_j, Y_j \rangle_{\mathcal{A}}.$$

Given a tracial W^* -embedding $\iota : \mathcal{A} \rightarrow \mathcal{B}$ and the corresponding conditional expectation $E : \mathcal{B} \rightarrow \mathcal{A}$, if $X \in L^1(\mathcal{A})^m$ and $Y \in L^1(\mathcal{B})^m$, we write

$$\iota(X) := (\iota(X_1), \dots, \iota(X_m)), \quad E[Y] := (E[Y_1], \dots, E[Y_m]).$$

Note that ι and E both preserve the real subspaces of self-adjoint tuples.

A.2. Laws and the Wasserstein distance for L^p variables. To extend the notion of non-commutative laws and couplings to L^p variables, we use a fairly standard trick in operator algebras, namely transforming an unbounded operator into a bounded operator using functional calculus. If $(X_1, \dots, X_m) \in L^p(\mathcal{A})_{\text{sa}}^m$, then $\arctan(X) := (\arctan(X_1), \dots, \arctan(X_m))$, where $\arctan(X_j)$ is defined by functional calculus, is an m -tuple of bounded self-adjoint operators with has a non-commutative law $\lambda_{\arctan(X)} \in \Sigma_{m, \pi/2}$. Rather than defining λ_X directly, we will work with $\lambda_{\arctan(X)}$ instead. The analogous procedure in classical probability theory would be to study a probability distribution μ on \mathbb{R}^m through the compactly supported probability distribution $\arctan_* \mu$ obtained by pushing forward μ by the function $(x_1, \dots, x_d) \mapsto (\arctan(x_1), \dots, \arctan(x_d))$.

Given a law $\lambda \in \Sigma_{m, \pi/2}$, the following criterion decides whether $\lambda = \lambda_{\arctan(X)}$ for some m -tuple X in a non-commutative L^p -space:

Lemma A.12. *Let $\lambda \in \Sigma_{m, \pi/2}$. For each j there is a measure λ_j on $[-\pi/2, \pi/2]$ with $\int f(x) d\lambda_j(x) = \lambda(f(x_j))$ for all non-commutative polynomials f . The following are equivalent:*

- (1) $\tan \in L^p(\lambda_j)$ for every j .
- (2) There exists a tracial W^* -algebra \mathcal{A} and $X \in L^p(\mathcal{A})_{\text{sa}}^m$ such that $\lambda = \lambda_{\arctan(X)}$.

Proof. (1) \implies (2). Let (\mathcal{A}, Y) be the GNS realization of λ given by Proposition 2.31. Then λ_j is the spectral distribution of Y_j with respect to τ . Because $\tan \in L^p(\lambda_j)$ for each j , we know that λ_j has no mass at $\pm\pi/2$ and therefore $X_j = \tan(Y_j)$ is a well-defined self-adjoint operator affiliated to \mathcal{A} . Using Theorem A.3 (3), we have $\mu_{\tan(Y_j)} = \tan_* \mu_{Y_j} = \tan_* \lambda_j$. Hence, $\tau(|X_j|^p) = \int |t|^p d\mu_{X_j}(t) = \int |\tan t|^p d\lambda_j(t) < \infty$. Therefore, $X_j \in L^p(\mathcal{A})_{\text{sa}}$ and $\lambda_{\arctan(X_j)} = \lambda$.

(2) \implies (1). If X is as in (2), then let $Y_j = \arctan(X_j)$. The spectral distribution of Y_j with respect to τ is thus $\mu_{Y_j} = \lambda$. \square

Definition A.13. We define $\Sigma_m^{(p)}$ as the set of non-commutative laws λ in $\Sigma_{m, \pi/2}$ such that $\tan \in L^2(\lambda_j)$ for every j . We define the weak- $*$ topology on $\Sigma_m^{(p)}$ as the restriction of the weak- $*$ topology on $\Sigma_{m, \pi/2}$.

Next, we define couplings of laws in $\Sigma_m^{(p)}$. It will be useful to have two different points of view on couplings, one more measure-theoretic, and the other more probabilistic.

Definition A.14. Given $\mu, \nu \in \Sigma_m^{(p)}$, a *measure-theoretic coupling* of μ and ν is a law $\gamma \in \Sigma_{2m}^{(p)}$ such that $\gamma(f(x_1, \dots, x_m)) = \mu(f(x_1, \dots, x_m))$ and $\gamma(f(x_{m+1}, \dots, x_{2m})) = \nu(f(x_1, \dots, x_m))$ for all $f \in \mathbb{C}\langle x_1, \dots, x_m \rangle$. We denote by $\Gamma^{(p)}(\mu, \nu) \subseteq \Sigma_{2m}^{(p)}$ the space of measure-theoretic couplings.

Definition A.15. A *probabilistic coupling* of μ and ν is a tuple (\mathcal{A}, X, Y) , where \mathcal{A} is a tracial W^* -algebra and $X, Y \in L^p(\mathcal{A})_{\text{sa}}^m$ with $\lambda_{\arctan(X)} = \mu$ and $\lambda_{\arctan(Y)} = \nu$.

Of course, if (\mathcal{A}, μ, ν) is a probabilistic coupling, then $\gamma = \lambda_{\arctan(X), \arctan(Y)}$ is a measure-theoretic coupling. Conversely, if γ is a measure-theoretic coupling, then a probabilistic coupling can be obtained from the GNS construction of γ .

Definition A.16 (Wasserstein distance). For a given $\mu, \nu \in \Sigma_m^{(p)}$, we define $d_W^{(p)}(\mu, \nu)$ to be the infimum $\|X - Y\|_{L^p(\mathcal{A})_{\text{sa}}^m}$ over all probabilistic couplings (\mathcal{A}, X, Y) with $\mathcal{A} \in \mathbb{W}$.

Proposition A.17. *The Wasserstein distance $d_W^{(p)}$ defines a metric on the set $\Sigma_m^{(p)}$ which makes it into a complete metric space.*

The argument to show that $d_W^{(p)}$ is a metric on $\Sigma_m^{(p)}$ is exactly the same as in [11]. The hardest axiom to verify is the triangle inequality, but this follows because a coupling of μ_1 and μ_2 and a coupling of μ_2 and μ_3 can be joined by taking the amalgamated free product of the tracial W^* -algebras corresponding to the two couplings over the subalgebra generated by the variables corresponding to μ_2 . Hence, $\Sigma_m^{(p)}$ is a metric space.

A.3. Optimal couplings and Wasserstein geodesics.

Definition A.18. A *probabilistic coupling* (\mathcal{A}, X, Y) of μ and $\nu \in \Sigma_m^{(p)}$ is said to be *optimal* if $\|X - Y\|_{L^p(\mathcal{A})_{\text{sa}}^m} = d_W^{(p)}(\mu, \nu)$; in this case we also say that the corresponding *measure-theoretic coupling* is *optimal*. We denote the space of optimal measure-theoretic couplings by $\Gamma_{\text{opt}}^{(p)}(\mu, \nu)$.

To show the existence of optimal couplings, we use a certain type of continuity and compactness.

Lemma A.19. *Let $p \in [1, \infty)$. Then $\Gamma^{(p)}(\mu, \nu)$ is compact in the weak-* topology. For $\gamma \in \Gamma^{(p)}(\mu, \nu)$, the quantity*

$$N(\gamma, p) := \|X - Y\|_{L^p(\mathcal{A})_{\text{sa}}^m} \text{ where } \gamma = \lambda_{\arctan(X), \arctan(Y)},$$

where $(X, Y) \in L^2(\mathcal{A})_{\text{sa}}^m$ is a $2m$ -tuple with $\lambda_{\arctan(X), \arctan(X)}$, only depends on γ , and moreover $\gamma \mapsto N(\gamma, p)$ is continuous.

Proof. First, $\Gamma^{(p)}(\mu, \nu)$ is compact because it is a closed subset of $\Sigma_{2m, \pi/2}$, which is compact. To show well-definedness and continuity $\gamma \mapsto N(\gamma, p)$, first note that for any polynomial ϕ , the map

$$\lambda_{Z, Z'} \mapsto \|\phi(Z) - \phi(Z')\|_{L^p(\mathcal{A})_{\text{sa}}^m}$$

is well-defined and continuous on $\Sigma_{2m, \pi}$. Indeed, let M be an upper bound on $|\phi|$ and let g_m be a sequence of polynomials that converge uniformly on $[-2M, 2M]$ to the function $|\cdot|^{p/2}$. Then $g_m((\phi(Z) - \phi(Z'))^*(\phi(Z) - \phi(Z')))$ converges to $|\phi(Z) - \phi(Z')|^p$ and the rate of convergence is uniform for all \mathcal{A} and all (Z, Z') with $\|(Z, Z')\|_{\infty} \leq \pi$ because of the spectral mapping theorem. The continuity of

$$\lambda_{Z, Z'} \mapsto \sum_{j=1}^m g_m(\phi(Z) - \phi(Z')) = \lambda_{Z, Z'}(g_m(\phi(x_j) - \phi(x_{m+j})))$$

is immediate by definition of the weak-* topology. Thus, the continuity of $\lambda_{Z, Z'} \mapsto \|\phi(Z) - \phi(Z')\|_p$ follows from uniform convergence. Similarly, using uniform convergence, we can generalize ϕ from a polynomial to an arbitrary continuous real-valued function on $[-\pi/2, \pi/2]$.

Now let $\phi_m \in C([-\pi/2, \pi/2]; \mathbb{R})$ be a sequence such that $|\phi_m| \leq |\tan|$ and $\phi_m \rightarrow \tan$ pointwise. Suppose that $\gamma \in \Gamma^{(p)}(\mu, \nu)$ and $\gamma = \lambda_{(Z, Z')}$. Then $(\tan(Z), \tau(Z')) \in L^p(\mathcal{A})_{\text{sa}}^m$. Also,

$$\begin{aligned} \|\tan(Z) - \tan(Z')\|_p - \|\phi_m(Z) - \phi_m(Z')\| &\leq \sum_{j=1}^m \|\tan(Z_j) - \phi_m(Z_j)\|_{L^p(\mathcal{A})} + \sum_{j=1}^m \|\tan(Z'_j) - \phi_m(Z'_j)\|_{L^p(\mathcal{A})} \\ &= \sum_{j=1}^m \left(\int |\tan - \phi_m|^p d\mu_j \right)^{1/p} + \sum_{j=1}^m \left(\int |\tan - \phi_m|^p d\nu_j \right)^{1/p}, \end{aligned}$$

where μ_j and ν_j are the measures on $[-\pi/2, \pi/2]$ representing the j th marginals of μ and ν respectively. The bound on the right-hand side only depends on μ and ν and thus is a uniform bound for all $\gamma \in \Gamma^{(p)}(\mu, \nu)$. Furthermore, by the dominated convergence theorem $|\tan - \phi_m| \rightarrow 0$ in $L^p(\mu_j)$ and $L^p(\nu_j)$. Therefore, the map sending γ to $\|\tan(Z) - \tan(Z')\|_{L^p(\mathcal{A})_{\text{sa}}^m}$ is continuous as the uniform limit of continuous maps. \square

Corollary A.20. *For each $\mu, \nu \in \Sigma_m^{(p)}$, the space of optimal couplings $\Gamma_{\text{opt}}^{(p)}$ is nonempty and compact.*

Given the existence of L^2 optimal couplings, all the theorems from §3 and §4 can be generalized to $\Sigma_m^{(2)}$ with the appropriate changes to notation. Almost no change is needed for the proofs since E -convex functions were defined for L^2 non-commutative random variables to begin with.

Just as in the case of classical probability theory, the existence of optimal couplings and the ability to take convex combinations of non-commutative random variables immediately leads to the existence of geodesics between any two points in $\Sigma_m^{(p)}$.

Definition A.21. Let (Ω, d) be a metric space. A *geodesic* in (Ω, d) is a continuous map $g : I \rightarrow \Omega$, where $I \subseteq \mathbb{R}$ is an interval (of positive length), such that for all $t_1 < t_2 < t_3$ in I , we have

$$d(g(t_1), g(t_3)) = d(g(t_1), g(t_2)) + d(g(t_2), g(t_3)).$$

The geodesic is said to be *constant speed* if $d(g(t_1), g(t_2))/(t_2 - t_1)$ is constant for all $t_1 < t_2$.

Proposition A.22. *Let $p \in [1, \infty)$. Let $\mu, \nu \in \Sigma_d^{(p)}$, and let (\mathcal{A}, X, Y) be a probabilistic optimal coupling. For $t \in [0, 1]$, let $X_t = (1-t)X + tY$. Let $\mu_t = \lambda_{\arctan(X_t)}$. Then $t \mapsto \mu_t$ is a constant speed geodesic in $(\Sigma_m^{(p)}, d_W^{(p)})$. Moreover, for $s, t \in [0, 1]$, (\mathcal{A}, X_s, X_t) is an optimal coupling of μ_s and μ_t .*

Proof. Of course,

$$d_W^{(p)}(\mu_s, \mu_t) \leq \|X_s - X_t\|_{L^p(\mathcal{A})_{sa}^m} = |s - t| \|X - Y\|_{L^p(\mathcal{A})_{sa}^m}.$$

Thus, for $s < t$,

$$\begin{aligned} \|X - Y\|_{L^p(\mathcal{A})_{sa}^m} &= d_W^{(p)}(X, Y) \\ &\leq d_W^{(p)}(X, X_s) + d_W^{(p)}(X_s, X_t) + d_W^{(p)}(X_t, Y) \\ &\leq \|X - X_s\|_{L^p(\mathcal{A})_{sa}^m} + \|X_s - X_t\|_{L^p(\mathcal{A})_{sa}^m} + \|X_t - Y\|_{L^p(\mathcal{A})_{sa}^m} \\ &= \|X - Y\|_{L^p(\mathcal{A})_{sa}^m}. \end{aligned}$$

Thus, all the inequalities are forced to be equalities. Hence, $d_W^{(p)}(\mu_s, \mu_t) = (t - s)d_W^{(p)}(\mu, \nu)$. Therefore, $t \mapsto \mu_t$ is a constant speed geodesic. Also, $d_W^{(p)}(\mu_s, \mu_t) = \|X_s - Y_t\|_{L^p(\mathcal{A})_{sa}^m}$, so that (\mathcal{A}, X_s, X_t) is an optimal coupling. \square

REFERENCES

- [1] Claire Anantharaman-Delaroche. On ergodic theorems for free group actions on noncommutative spaces. *Probab. Theory Rel. Fields*, 135:520–546, 2006.
- [2] Claire Anantharaman-Delaroche and Sorin Popa. An introduction to II_1 factors. 2021.
- [3] Greg W. Anderson, Alice Guionnet, and Ofer Zeitouni. *An Introduction to Random Matrices*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2009.
- [4] Scott Atkinson, Isaac Goldbring, and Srivatsav Kunnawalkam Ellayavalli. Factorial relative commutants and the generalized Jung property for II_1 factors. arXiv:2004.02293.
- [5] Scott Atkinson and Srivatsav Kunnawalkam Ellayavalli. On ultraproduct embeddings and amenability for tracial von Neumann algebras. *International Mathematics Research Notices*, 2021(4):2882–2918, 2021.
- [6] V. Barbu and G. da Prato. Global existence for the Hamilton-Jacobi equations in Hilbert space. *Annali della Scuola Normale Superiore di Pisa Classe di Scienze 4e serie*, 8(2):257–284, 1981.
- [7] V. Barbu and G. da Prato. Hamilton-Jacobi equations in Hilbert spaces; variational and semigroup approach. *Annali di Matematica pura ed applicata*, 142:303–349, 1985.
- [8] V. Barbu and G. da Prato. A note on a Hamilton-Jacobi equation in Hilbert space. *Nonlinear Analysis: Theory, Methods & Applications*, 9(12):1337–1345, 1985.
- [9] Itai Ben Yaacov. Topometric spaces and perturbations of metric structures. *Log. Anal.*, 1:235–272, 2008.
- [10] Charles Bertucci, Mérouane Debbah, Jean-Michel Lasry, and Pierre-Louis Lions. A spectral dominance approach to random matrices. Preprint arXiv:2105.08983, 2021.
- [11] Philippe Biane and Dan-Virgil Voiculescu. A free probability analogue of the Wasserstein metric on the trace-state space. *Geometric and Functional Analysis*, 11:1125–1138, 2001.
- [12] Nathaniel P. Brown and Narutaka Ozawa. *C*-algebras and finite-dimensional approximations*, volume 88 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, 2008.
- [13] Valerio Capraro. A survey on Connes’ embedding conjecture. arXiv preprint at arXiv:1003.2076, 2010.
- [14] Pierre Cardaliaguet, François Delarue, Jean-Michel Lasry, and Pierre-Louis Lions. *The Master Equation and the Convergence Problem in Mean Field Games*, volume 2 of *Annals of Mathematics Studies*. Princeton University Press, 2019.
- [15] René Carmona, Mark Cerenzia, and Aaron Zeff Palmer. The Dyson and Coulomb games. *Ann. Henri Poincaré*, 21:2897–2949, 2020.
- [16] Ian Charlesworth and Brent Nelson. Free Stein irregularity and dimension. arXiv:1902.02379, 2019.
- [17] Alain Connes. Classification of injective factors. Cases II_1 , II_∞ , III_λ , $\lambda \neq 1$. *Ann. of Math. (2)*, 104(1):73–115, 1976.
- [18] Michael G. Crandall and Pierre-Louis Lions. Hamilton-Jacobi equations in infinite dimensions I. uniqueness of viscosity solutions. *J. Funct. Anal.*, 62:379–396, 1985.
- [19] Michael G. Crandall and Pierre-Louis Lions. Hamilton-Jacobi equations in infinite dimensions II. existence of viscosity solutions. *J. Funct. Anal.*, 65:368–405, 1986.
- [20] Ricardo Correa da Silva. Lecture notes on non-commutative L_p -spaces. arXiv:1803.02390, 2018.
- [21] Yoann Dabrowski. A Laplace principle for Hermitian Brownian motion and free entropy I: the convex functional case. arXiv:1604.06420, 2017.
- [22] Yoann Dabrowski. A non-commutative path space approach to stationary free stochastic differential equations. arxiv:1006.4351, 2010.
- [23] Yoann Dabrowski, Alice Guionnet, and Dimitri Shlyakhtenko. Free transport for convex potentials. arXiv:1701.00132, 2016.

- [24] Jacques Dixmier. Formes linéaires sur un anneau d'opérateurs. *Bulletin de la Société Mathématique de France*, 81:9–39, 1953.
- [25] Edward G. Effros. *Dimensions and C*-algebras*, volume 46 of *CBMS Regional Conference Series in Mathematics*. American Mathematical Society, Providence, 1981.
- [26] Wilfrid Gangbo, Sergio Mayorga, and Andrzej Świąch. Finite dimensional approximations of Hamilton-Jacobi-Bellman equations in spaces of probability measures. *SIAM J. Math. Anal.*, 53(2):1320–1356, 2021.
- [27] Wilfrid Gangbo, Alpár Mészáros, Chenchen Mou, and Jianfeng Zhang. Mean field games master equations with non-separable Hamiltonians and displacement monotonicity. arXiv preprint arXiv:2101.12362, 2021.
- [28] Wilfrid Gangbo and Alpár R Mészáros. Global well-posedness of master equations for deterministic displacement convex potential mean field games. arXiv preprint arXiv:2004.01660, 2021.
- [29] Wilfrid Gangbo and Adrian Tudorascu. On differentiability in the Wasserstein space and well-posedness for Hamilton-Jacobi equations. *J. Math Pures Appl.*, 125:119–174, 2018.
- [30] Mikhael Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, page 75–263. Springer, New York, 1987.
- [31] Alice Guionnet and Dimitri Shlyakhtenko. Free monotone transport. *Inventiones Mathematicae*, 197(3):613–661, 09 2014.
- [32] Uffe Haagerup and Magdalena Musat. Factorization and dilation problems for completely positive maps on von neumann algebras. *Comm. Math. Phys.*, 303(2):555–594, 2011.
- [33] Uffe Haagerup and Magdalena Musat. An asymptotic property of factorizable completely positive maps and the connes embedding problem. *Comm. Math. Phys.*, 338(2):721–752, 2015.
- [34] Don Hadwin and Tatiana Shulman. Tracial stability for C*-algebras. *Integral Equations Operator Theory*, 90(1), 2018.
- [35] Benjamin Hayes, David Jekel, Brent Nelson, and Thomas Sinclair. A random matrix approach to absorption in free products. *International Mathematics Research Notices*, 2021(3):1919–1979, 2021.
- [36] Fumio Hiai, Dénes Petz, and Yoshimichi Ueda. Free transportation cost inequalities via random matrix approximation. *Probab. Theory Related Fields*, 130(2):199–221, 2004.
- [37] Fumio Hiai and Yoshimichi Ueda. Free transportation cost inequalities for noncommutative multi-variables. *Infinite Dimensional Analysis, Quantum Probability, and Related Topics*, 9:391–412, 2006.
- [38] David Jekel. An elementary approach to free entropy theory for convex potentials. arXiv:1805.08814, 2018. To appear in *Analysis and PDE Journal*.
- [39] David Jekel. Conditional expectation, entropy, and transport for convex gibbs laws in free probability. *International Mathematics Research Notices IMRN*, 2020, 2020.
- [40] David Jekel. *Evolution equations in non-commutative probability*. PhD thesis, University of California, Los Angeles, 2020.
- [41] David Jekel, Wuchen Li, and Dimitri Shlyakhtenko. Tracial non-commutative smooth functions and the free Wasserstein manifold. arXiv:2101.06572, 2021.
- [42] Zhengfeng Ji, Anand Natarajan, Thomas Vidick, John Wright, and Henry Yuen. MIP*=RE. arXiv:2001.04383, 2020.
- [43] Richard Jordan, David Kinderlehrer, and Felix Otto. The variational formulation of the Fokker–Planck equation. *SIAM Journal on Mathematical Analysis*, 29(1):1–17, 1998.
- [44] Kenly Jung. Amenability, tubularity, and embeddings into \mathcal{R}^ω . *Math. Ann.*, 338(1):241–248, 2007.
- [45] John D. Lafferty. The density manifold and configuration space quantization. *Transactions of the American Mathematical Society*, 305(2):699–741, 1988.
- [46] J.M. Lasry and P.L. Lions. A remark on regularization in Hilbert spaces. *Israel J. Math.*, 55:257–266, 1986.
- [47] Wuchen Li. Transport information geometry: Riemannian calculus on probability simplex. arXiv:1803.06360 [math], 2018.
- [48] Dusa McDuff. Uncountably many II_1 factors. *Ann. of Math. (2)*, 90(2):372–377, 1969.
- [49] Pierre-Louis Lions Michael G. Crandall. Hamilton-Jacobi equations in infinite dimensions III. *J. Funct. Anal.*, 68:214–247, 1986.
- [50] F. J. Murray and J. von Neumann. On rings of operators. *Annals of Mathematics*, 37(1):116–229, 1936.
- [51] F.J. Murray and J. von Neumann. On rings of operators, II. *Trans. Amer. Math.Soc.*, 41:208–248, 1937.
- [52] Magdalena Musat and Mikael Rørdam. Non-closure of quantum correlation matrices and factorizable channels that require infinite dimensional ancilla (with an appendix by Narutaka Ozawa). *Comm. Math. Phys.*, 375:1761–1776, 2020.
- [53] Brent Nelson. Free monotone transport without a trace. *Communications in Mathematical Physics*, 334(3):1245–1298, 2015.
- [54] Brent Nelson. Free transport for finite depth subfactor planar algebras. *Journal of Functional Analysis*, 268(9):2586–2620, 2015.
- [55] A. Yu. Olshanskii. On residualizing homomorphisms and g-subgroups of hyperbolic groups. *Internat. J. Algebra Comput.*, 3(4):365–409, 1993.
- [56] Felix Otto. The geometry of dissipative evolution equations the porous medium equation. *Communications in Partial Differential Equations*, 26(1-2):101–174, 2001.
- [57] Felix Otto and Cédric Villani. Generalization of an inequality by Talagrand and links with the logarithmic Sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, 2000.
- [58] Narutaka Ozawa. There is no separable universal II_1 factor. *Proc. Amer. Math. Soc.*, 132:487–90, 2004.

- [59] Narutaka Ozawa. About the connes embedding conjecture: Algebraic approaches. *Japanese Journal of Mathematics*, 8:147–183, 2013.
- [60] Vern Paulsen. *Completely Bounded Maps and Operator Algebras*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2003.
- [61] Sorin Popa. A short proof of “injectivity implies hyperfiniteness” for finite von neumann algebras. *J. Operator Theory*, 16:261–272, 1986.
- [62] Michael Reed and Barry Simon. *Methods of Modern Mathematical Physics I: Functional Analysis*. Academic Press, New York, 1972.
- [63] Shōichirō Sakai. *C*-algebras and W*-algebras*, volume 60 of *Ergebnisse der Mathematik und ihrer Grenzgebiete*. Springer-Verlag, Berlin Heidelberg, 1971.
- [64] Dimitri Shlyakhtenko. Free fisher information for non-tracial states. *Pacific J. Math*, 211:375–390, 2003.
- [65] Masamichi Takesaki. *Theory of Operator Algebras I*, volume 124 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, Berlin Heidelberg, 2002.
- [66] Masamichi Takesaki. *Theory of Operator Algebras II*, volume 125 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, Berlin Heidelberg, 2003.
- [67] Masamichi Takesaki. *Theory of Operator Algebras III*, volume 127 of *Encyclopaedia of Mathematical Sciences*. Springer-Verlag, Berlin Heidelberg, 2003.
- [68] Yoshimichi Ueda. Factoriality, type classification and fullness for free product von Neumann algebras. *Adv. Math.*, 228(5):2647–2671, 2011.
- [69] Cédric Villani. *Optimal Transport: Old and New*, volume 338 of *Grundlehren Der Mathematischen Wissenschaften*. Springer, Berlin, 2009.
- [70] Dan-Virgil Voiculescu. Symmetries of some reduced free product C^* -algebras. In Huzihiro Araki, Calvin C. Moore, Şerban-Valentin Stratila, and Dan Voiculescu, editors, *Operator Algebras and their Connections with Topology and Ergodic Theory*, pages 556–588. Springer, Berlin, Heidelberg, 1985.
- [71] Dan-Virgil Voiculescu. Addition of certain non-commuting random variables. *Journal of Functional Analysis*, 66(3):323–346, 1986.
- [72] Dan-Virgil Voiculescu. Limit laws for random matrices and free products. *Inventiones mathematicae*, 104(1):201–220, Dec 1991.
- [73] Dan-Virgil Voiculescu. The analogues of entropy and Fisher’s information measure in free probability theory, I. *Communications in Mathematical Physics*, 155(1):71–92, 1993.
- [74] Dan-Virgil Voiculescu. The analogues of entropy and of Fisher’s information measure in free probability theory, II. *Inventiones Mathematicae*, 118:411–440, 1994.
- [75] Dan-Virgil Voiculescu. The analogues of entropy and of Fisher’s information measure in free probability theory, III: Absence of Cartan subalgebras. *Geometric and Functional Analysis*, 6:172–199, 1996.
- [76] Dan-Virgil Voiculescu. The analogues of entropy and of Fisher’s information measure in free probability theory V. *Inventiones Mathematicae*, 132:189–227, 1998.
- [77] Dan-Virgil Voiculescu. A strengthened asymptotic freeness result for random matrices with applications to free entropy. *International Mathematics Research Notices*, 1998(1):41–63, 1998.
- [78] Dan-Virgil Voiculescu, Kenneth J. Dykema, and Alexandru Nica. *Free Random Variables*, volume 1 of *CRM Monograph Series*. American Mathematical Society, Providence, RI, 1992.
- [79] Mark M. Wilde. *Quantum Information Theory*. Cambridge University Press, 2013.
- [80] Edward Witten. A mini-introduction to information theory. *La Rivista del Nuovo Cimento*, 43:187–227, 2020.

Email address: wgangbo@math.ucla.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES, LOS ANGELES, CA 90095

Email address: djekel@ucsd.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, SAN DIEGO, LA JOLLA, CA 92093

Email address: ksnam@g.ucla.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES, LOS ANGELES, CA 90095

Email address: shlyakht@math.ucla.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, LOS ANGELES, LOS ANGELES, CA 90095