

White Paper: “Non-commutative Optimal Transport” (IPAM Long Program, Spring 2025)

Authors (listed alphabetically):

Roy Araiza (University of Illinois at Urbana-Champaign)
Paul Ayers (McMaster University)
Emily Beatty (École Normale Supérieure de Lyon/Université Lyon 1)
Zachary Brannan (University of Ottawa)
Saroj Chhatoi (Laboratoire d'analyse et d'architecture des systèmes, LAAS-CNRS)
Valerii Chuiko (McMaster University)
Rolando de Santiago (California State University, Long Beach, CSU Long Beach)
Rie Fujii (Ochanomizu University)
Tryphon Georgiou (University of California, Irvine, UCI)
Augusto Gerolin (University of Ottawa)
Farnaz Heidar-Zadeh (Queen's University)
John Hopper (University of California, Los Angeles, UCLA)
Marius Junge (University of Illinois at Urbana-Champaign)
Therese Landry (University of California, Santa Barbara, UCSB)
Spencer Lee (Michigan State University)
Zhiyi Lin (University of Ottawa)
Milad Marvian (University of New Mexico)
Saniya Minase (Colorado School of Mines)
Nataliia Monina (University of Ottawa)
Trung Nghia Nguyen (Rutgers University New Brunswick/Piscataway)
Đorđe Nikolić (University of California, Santa Barbara, UCSB)
Pavlo Pelikh (University of Ottawa)
Simone Rademacher (Ludwig-Maximilians-Universität München)
Denys Ruban (University of Ottawa)
Kanon Sakurai (Ochanomizu University)
Wesley Sanchez (McMaster University)
Dimitri Shlyakhtenko (Institute for Pure and Applied Mathematics, UCLA)
Oliver Tse (Eindhoven University of Technology)
Luke Visser (Eindhoven University of Technology)
Jia Wang (University of Illinois at Urbana-Champaign)
Peixue Wu (Institute for Quantum Computing, University of Waterloo)
Yefei Zhang (University of Illinois at Urbana-Champaign)

doi:10.82930/B6WC7X

Table of Contents:

1. Executive Summary
2. Introduction
3. Bridging Existing NOT Concepts
4. NOT in Free Probability
5. Gradient Structures from Dilations
6. Quantum Thermodynamics
7. Quantum Optimal Control and Algorithms
8. Quantum Many-Body Problem
9. Computational Aspects for Quantum (ML, SDP)

Abbreviations:

BKM	Bogoliubov-Kubo-Mori inner product
DFT	Density-Functional Theory
DMFT, kDMFT	(k-body) Density Matrix Functional Theory
DM, kDM	(k-body) reduced Density Matrix
GHZ	Greenberger–Horne–Zeilinger
GNS	Gelfand Naimark Segal
KMS	Kubo-Martin-Schwinger condition
ML	Machine Learning
NOT	Non-commutative Optimal Transport
OT	Optimal Transport
PINN	Physics-Informed Neural Network
QMOT	Quantum Multi-marginal Optimal Transport
QMS	Quantum Markov Semigroup
QOC	Quantum Optimal Control
QOT	Quantum Optimal Transport
QSDE	Quantum Stochastic Differential Equation
SDE	Stochastic Differential Equation
SDP	Semidefinite Program
SOS	Sum of Squares
SYK	Sachdev-Ye-Kitaev algorithm
vNA	von Neumann Algebra

1. Executive Summary

This document summarizes the activities and outcomes of the Long Program “[Non-commutative Optimal Transport](#),” which was held at the Institute of Pure and Applied Mathematics (IPAM) from March 10 to June 13, 2025. It also briefly explores some of the current open questions and future directions in the field of electronic structure theory, as well as related fields that were discussed during the program.

This document surveys key topics identified by the program participants:

Bridging Existing NOT Concepts: Generalizing classical OT quantities to the quantum setting is a challenging task due to the barriers posed by entanglement and non-commutativity. This has led to diverging definitions of NOT concepts, which have little in common with each other. A key area of ongoing research, explored by two working groups, lies in the reunification of these concepts and the building of bridges between existing NOT formulations.

NOT in Free Probability: Analogizing from the classical setting, free optimal transport, free entropy, and free optimal control provide a framework for studying random matrix models, their large- N limits, and structural properties of vNAs.

Gradient Structures from Dilations: Recent work has shown the existence of a gradient structure for the Lindblad equation under certain symmetry conditions, which has been effectively used to prove entropy inequalities. We explore a way to derive gradient structures from dilations and microscopic dynamics by using large deviations.

Quantum Thermodynamics: The notion of work, recently adapted to the quantum setting, provides a control-theoretic framework for studying quantum state transportation along a path. Thanks to the entropic nature of the dissipative components of the work along such a path, new estimates can be derived.

Quantum Optimal Control and Algorithms: Both QOC and variational quantum circuit optimization are known to suffer from barren plateaus and many bad local minima as the number of qubits in a quantum computer increases. Using optimal transport techniques has proven effective in mitigating these issues for certain problems and circuit parameterizations. We explore how optimal transport techniques can be used to enhance the performance of QOC protocols and analyze the computational complexity of quantum algorithms.

Quantum Many-Body Problem: NOT provides a mathematical framework for reconstructing quantum states and their properties directly from their reduced density matrices. This leads to new mathematical and computational techniques that are especially effective for applications where the number of quantum particles becomes (infinitely) large.

Computational Aspects for Quantum: This section presents recent advances in computational strategies for QOT, highlighting the development of entropy-based regularization methods, SDP approaches, and ML techniques to overcome the challenges posed by QOT's non-commutative structure and high dimensionality.

2. Introduction

In recent decades, optimal transport (OT) has emerged as a vibrant and versatile field, offering deep insights and powerful tools across a broad spectrum of disciplines—from economics and geometry to meteorology, fluid mechanics, statistics, engineering, and design. More recently, driven by challenges in artificial intelligence, quantum physics, quantum information theory, and electronic structure theory, the scope of OT has expanded to encompass fundamentally different classes of state spaces. These include spaces of density matrices, operators, and operator algebras—structures that are inherently non-commutative.

To foster collaboration across the diverse mathematical communities engaging with this emerging field, the program brought together researchers working on non-commutative optimal transport (NOT). Key topics of focus included:

- Evolution equations for density operators
- Wasserstein distances for density matrices and operators
- Semidefinite relaxations of optimal transport problems
- Free Probability and its connections to OT
- Functional inequalities in matrix and operator settings
- Computational methods and machine learning methods for NOT

The overarching goal was to develop new analytical, geometric, computational, and statistical tools to address foundational questions in the field and extend the applicability of NOT to a wide range of problems in non-commutative probability, geometry, quantum theory, and beyond.

Non-commutative Optimal Transport: static and dynamical formulations

Efforts to construct a non-commutative analogue of the classical Monge-Kantorovich theory trace back roughly thirty years. In the context of non-commutative geometry and probability, early contributions were made by Alain Connes and Dan Voiculescu. These foundational studies explored static formulations based on duality principles—a theme revisited in [Workshop I](#) of the program.

Simultaneously, in the classical setting, a shift toward dynamic formulations was taking shape through the work of Benamou, Brenier, McCann, and Otto. The formulation of the Fokker–Planck equation as a gradient flow of the free energy in the Wasserstein metric, formalized by Jordan, Kinderlehrer, and Otto (JKO), marked a turning point. This variational viewpoint had profound implications for geometry, thermodynamics, and fluid mechanics.

Inspired by these developments, Carlen and Maas initiated a quantum counterpart of the JKO framework. They showed that certain quantum dynamical semigroups, specifically QMS, can be interpreted as gradient flows of the von Neumann entropy. This opened the door to new variational methods, the study of functional inequalities, and the development of non-commutative structures.

[Workshop II](#) explored dynamical formulations of NOT, aiming to unify various perspectives on quantum dynamics. While quantum mechanics admits multiple equivalent formulations—such as wavefunctions, matrix mechanics, path integrals, variational principles, and density matrices—the variational lens inspired by OT has emerged as particularly promising. In particular, it offers new avenues for discovering and interpreting principles of motion and equilibrium in quantum systems.

A central challenge lies in developing a robust theory of large deviations in the non-commutative setting. In classical OT, the Schrödinger Bridge problem provides a deep connection between large deviations and entropic transport. Extending such insights to quantum systems remains an open and intriguing problem.

[Workshop III](#) focused on algorithmic and statistical frameworks for NOT. It emphasized advances in convex optimization and semidefinite programming, along with potential synergies with ML. These computational tools are particularly crucial for addressing high-dimensional and many-body problems that become intractable using classical approaches.

Activities and focal points

Recent research has proposed a variety of NOT frameworks. Harmonization and differentiation of different concepts of NOT, aiming to answer the question: "What is NOT?" took center stage during the program.

One direction involves defining quantum analogues of Wasserstein distances between quantum states, such as qubits or channels. These distances often incorporate entropy-based regularization, facilitating numerical computation through generalized Sinkhorn-like algorithms. In such settings, classical transport plans are replaced by joint quantum states with specified marginal density matrices, and cost functions are modeled by Hamiltonians.

Duality approach. Alternative approaches adopt dual formulations involving transposed operators or quantum channels, leading to different geometric and operational insights. Still others connect to semiclassical limits of quantum mechanics, where transport-like problems emerge as approximations to full quantum evolution.

Unified framework. The talk by the student Fanch Coudreuse delivered a first general framework for static and dynamic concepts. This general framework can also be modified to include aspects of NOT in free probability.

NOT in quantum chemistry. The multimarginal perspective of NOT provides a new perspective on the traditional constrained-search approach to DM functional theory. The (regularized) NOT approach enables results (some previously established only heuristically) to be formally proved and facilitates the development of new computational approaches for many-electron systems in chemistry and physics.

NOT in free probability and random matrix theory. NOT techniques have been successfully applied in the context of free probability theory to obtain isomorphism and structure results for vNAs, dilation theory for semigroups, as well as in random matrix theory. A number of new promising directions using ideas such as gradient flows, entropy regularization, as well as techniques from mean field games, have been discussed during the program.

Long-Term Impact: The thematic program has created a new scientific community focused on developing mathematical, computational, and statistical NOT methods for a diverse class of problems within and beyond mathematics. As a result, both short- and long-term research agendas were established for advancing mathematical and computational techniques, as well as (potentially new) applications.

3. Bridging Existing NOT Concepts

Introduction: Over the last decade, much research has been focused on generalizing classical OT to quantum systems. Entanglement and non-commutativity make this both a difficult and fundamental task. Basic properties like the triangle inequality cannot be obtained with direct translations of the classical definition of OT, and many quantum geometric structures do not naturally induce cost functions. This leads to a range of quantities with different properties, underlying geometries, and applications, referred to as “non-commutative Wasserstein distances.” The aim of this section is to highlight progress in unifying these definitions. Free probability and multi-marginal OT require other approaches, and are discussed in Sections 4 and 8, respectively.

State of the art: Many existing works define QOT quantities by direct generalization of the classical definition of OT. This perspective often makes both the triangle inequality and faithfulness inaccessible, and those quantities that are faithful are restricted in the geometries they can represent. Others have quantized the Kantorovich dual formulation of quantum 1-Wasserstein distances by defining a notion of the Lipschitz constant of operators. Others have further quantized the Benamou-Brenier dynamical formulation of the 2-Wasserstein distance, though these three approaches, which coincide in the classical setting, diverge far in the non-commutative setting. Even within frameworks, the choice of cost parameters gives quantities that are difficult to relate to one another.

OT via group distance (static vs dynamic): For classical OT, the Kantorovich and Benamou-Brenier formulations are equivalent; however, in QOT, this is no longer true. The relationship between dynamical and static QOT for cost functions on groups allows for more explicit exploration. Leveraging on the work of [Peyre et al.](#) on unbalanced OT, we may formulate static QOT on finite compact groups and extract admissible solutions via the quantum Sinkhorn algorithm. Surprisingly, all couplings are separable, even though on commuting inputs, this reduces to the 2-Wasserstein distance. A common feature of these models is the choice of inclusion of observables, or equivalently, the interpretation of the restrictions of couplings to subsystems. This leads to barriers like the zero-infinity law for the static distance and violation of the triangle inequality. A qualitative difference between the cost of couplings can be derived from dynamical QOT and cost functions of measures on the group and their respective channels.

Harmonization within Lipschitz approaches: Building on the works of Connes, Lott, and Rieffel, the dual approach to Wasserstein distances through Lipschitz seminorms offers insights into quantum NOT analogues. This involves matricial Lipschitz

seminorms, combining Rieffel's axiomatic approach with Effros and Ruan's operator space axioms. The resulting L-Wasserstein distances have applications in contraction coefficients, mixing-time estimates, Hamiltonian simulation, and the implementation costs of quantum processes in vNAs. Notably, using Lipschitz semi-norms through commutators, [Araiza et al.](#) demonstrated lower bounds for circuit costs and recovered 1-Wasserstein complexity via the quantum 1-Wasserstein distance introduced by [De Palma et al.](#) This distance can also be seen as a dual Lipschitz seminorm, facilitating the application of non-commutative geometry and operator algebras, thus bridging various proposals for quantum Wasserstein distances.

Harmonization within static approaches: The classical Kantorovich problem can be directly translated in two ways: as an infimum of the expectation value of a cost operator over all quantum states on a tensor product of spaces with prescribed marginals, or as an infimum of a cost function over quantum transport plans. For the first to give a non-degenerate quantity with zero self-distance, the cost operator is restricted to an analogue of the discrete metric. This means some works have represented purely the discrete metric, and others have defined quantities with non-zero self-distance. The second method via transport plans cannot easily be approximated numerically. Key issues in this area include unifying the coupling and transport plan formalisms and establishing coherent cost matrix choices.

Work done at IPAM: The following topics were explored by the program participants.

1. Transportation cost and contraction coefficient for channels on von Neumann algebras ([Araiza, Junge, Wu \(2025\)](#)): As elaborated above, the dual approach to quantum Wasserstein distances via matricial Lipschitz seminorms leads to many interesting applications. During the long program, Araiza-Junge-Wu formalized these seminorms and their corresponding Wasserstein distances, which led to the introduction of the complete Lipschitz cost and complete Lipschitz constant of channels acting on arbitrary vNAs. This generalized approach leads to many key applications to expected group word length, Carnot-Caratheodory distance, and for non-symmetric channels which admit strictly contractive Lipschitz constants, one can obtain entropy contraction and mixing-time estimates.

2. Quantum Wasserstein distances for quantum permutation groups ([Anshu, Jekel, Landry \(2025\)](#)): The authors considered a notion of Lipschitz functions on the quantum permutation group S_n^+ and the associated notion of 1-Wasserstein distance on the state-space of the underlying C^* -algebra. This metric is a quantum analog of the classical Hamming distance, and its construction relies on ideas in free probability developed by Biane and Voiculescu.

3. Wasserstein distances on quantum structures: an overview (Beatty (2025)): This survey covers existing definitions of all quantum Wasserstein distances introduced before the start of the long program, highlighting gaps in the literature and open problems. Much of the work was completed while the author was at the long program.

4. $C^{1,\alpha}$ -Regularity for Shannon Entropy-regularized Optimal Transport (Baasandorj, Di Marino, Gerolin (2025)): This article establishes $C^{1,\alpha}$ -regularity for Schrödinger potentials arising in entropy-regularized OT for probability measures. That result holds under the assumption that the marginal distributions are supported on bounded, convex sets, and their densities are uniformly bounded. This can be seen as an entropic analogue of Caffarelli's regularity theory for the Monge–Ampère equation. This highlights the importance of Lipschitz regularity in both OT and NOT frameworks.

5. Highlight: Semiclassical couplings within the static framework: In existing works, couplings of quantum states are defined as bipartite states with prescribed partial traces, and transport plans are specified as weighted pairs of pure states. Significant progress in one of the working groups has shown that these can be reformulated as classical probability measures on a bipartite Hilbert space, and the cost of such a probability measure is an integral of some functional with respect to it. This unifies the existing coupling and transport plan frameworks within the static setting to be under one definition, differing only by a choice of cost functional on the product space.

6. Connection between static and dynamic formulation via group distance: During the long program, we devised a static formulation of QOT on groups and examined properties of solutions obtained via the quantum Sinkhorn algorithm. For specific group induced models, we found a way to distinguish densities with zero static distance via the dynamical distance on channels. In addition, we explored the possibility of extending this pattern to QOT on graphs.

Outlook: The harmonization of QOT quantities is an area that could allow quantum applications to take advantage of all key features of classical OT, instead of being tied to the properties offered by one specific definition. Key challenges in this area include the extension of the semiclassical formalism to identify practical cost functions and determine conditions on cost functions giving OT quantities with properties such as metricity, calculability, and links to physical systems, motivating choices of Lipschitz constant definitions within the commutator formalism, extending existing quantities to non-faithful quantum states, exploring Lagrangian costs as a link between the Lipschitz and dynamical formulations, and extending non-commutative QOT problems from the group setting to graphs via transfer methods on groupoids. Similarly, a new connection

to NOT can be derived from the theory of non-local games via interpreting strategies as couplings.

4. NOT in Free Probability

Introduction: In this version of NOT, the emphasis is on obtaining non-commutative analogs of the solution to the Monge problem, with optimality as a guiding principle for the construction of the Monge map (as in the approach by Brenier in the classical case). The key difference with other versions of NOT lies in a different definition of the notion of the “law” — in free probability, this notion is based on generalizations of duality between measures and non-commutative functions, rather than on considering a non-commutative notion of density.

Highlights: Free OT is an active area of research that has seen important applications to random matrix theory, free probability, and operator algebras. An introductory talk on free probability was given during [Tutorials](#), and several talks were given in [Workshop II](#). An active working group in Quantum Mean Field Games Free Probability Theory met throughout the program.

Free Probability as a Large-N Limit: Voiculescu’s free probability studies non-commutative probability spaces. These consist of a (non-commutative) unital $\ast$ -algebra N of “random variables” or “observables” endowed with a unital linear functional that assigns to each random variable its expected value. An example occurs in random matrix theory, where the algebra consists of matrices with random entries, and the expectation functional computes the expected value of the matrix trace.

Voiculescu discovered that his notion of *free independence* describes the asymptotic behavior of certain random matrix models. Associated with a potential function, one can define free Gibbs laws. Under suitable assumptions, these are limits of laws of random multi-matrix models obtained from Gibbs measures on spaces of matrices.

One can view free probability as a way to describe the “large- N ” limit of random matrix models. The limit can be described as mean-field, with the “algebra of observables” consisting of functions that are obtained by matrix operations as opposed to the larger algebra of functions of individual entries.

NOT for Random Matrix Models: Biane and Voiculescu introduced a notion of Wasserstein distance suitable for free probability theory in the early 2000s. In the early 2010s, Guionnet and Shlyakhtenko showed that in the case of free Gibbs states with potentials which are small perturbations of quadratic, non-commutative transport exists

and is given by “monotone” maps in the sense of Brenier. Furthermore, writing free Gibbs states as large- N limits of laws of random matrix models, the transport maps are themselves limits of the (classical probability) OT maps defined at each finite N . The finite- N OT maps involve not only matrix algebra (addition and multiplication of matrices) but also operations on individual entries. The Guionnet-Shlyakhtenko result shows that these operations become “more and more matricial” for large N .

Subsequently, Figalli and Guionnet were able to push this convergence result much further, showing that the finite- N OT maps admit an expansion in $1/N$. This was used to prove universality results for eigenvalues of certain random matrix ensembles.

On the free probability side, constructions of non-commutative transport maps were used to prove isomorphisms between certain vNAs. For example, it was shown by Dabrowski, Guionnet, and Shlyakhtenko (and, under different conditions, by Gangbo, Jekel, Li, Shlyakhtenko) that under convexity and growth assumptions on the potential, the vNAs associated to free Gibbs states are isomorphic. It was also shown that the so-called q -deformed free group factors with a finite number of generators are isomorphic to each other for small values of q .

Free Probability and the Mean Field Limit of Control Theory: Classical mean-field game theory uses a continuous object to describe the joint state of a game with an infinite number of infinitesimal players. This represents a mean-field limit of certain games where the number of players goes to infinity. Optimal control questions in the mean-field limit are then described by PDEs, which can be used to derive approximately optimal control strategies for the finite games.

In the mean-field approach to free control questions, Gangbo, Jekel, Li, and Palmer considered optimal control questions in the free probability limit, as well as related control questions for random matrix processes. There are new notions of noise: those that extend directly from the classical (e.g., common noise), generalizations to the free setting (e.g., individual noise and free semi-circulars), or ones unique to the non-commutative setting (e.g., Linbladans).

Defining the value function in the non-commutative setting required the use of iterated infima. Results related to Jekel’s work (see outlook) show that these complications are an inherent part of the free OT. If the infimums are attained, then they solve the Euler-Lagrange equations. This has the simplicity of producing a PDE but has the challenge of showing that the infimum is attained—which is difficult in the infinite-dimensional Hilbert space setting natural to free probability.

Another PDE arises from the value function itself, which avoids the existence of infima, as achieved by [Gangbo et al.](#) using the dynamic programming principle. We should point out the subtlety of this result, as the value function is often first defined pointwise, then integrated over the distribution of players. In contrast, in the non-commutative setting, there is no analog of points.

Approaches to Transport via Gradient Flow: A central question in the classification of vNAs is to determine whether distinct non-abelian free groups have non-isomorphic vNAs. In a related direction, one can ask whether there exists analytic criteria that one can impose on a vNA that would imply either containment of a free group factor or isomorphism to a free group factor. By solving an analog of the Monge-Ampere equation, Guionnet-Shlyakhtenko proved a variant of Brenier's theorem to construct invertible transport maps from the q -deformed free group factors to free group factors, thereby proving isomorphism results.

An approach undertaken during the program emulates how one might prove Brenier's theorem, which transports absolutely continuous measures on $\mathbb{R}^d$ to the standard Gaussian. This can be achieved by minimizing the entropy relative to the Gaussian of the push forward of the measure. By abstract methods, one can show that the optimizer would, in fact, give an OT map. Inspired by this, the group considered the case of vNAs generated by a tuple with finite free Fisher Information and attempted to show that one could construct a transport map from the generators to elements that are freely independent and whose laws are the semi-circular. This required finding equivalent formulations of relative entropy from the classical setting to construct a functional that would be suitable for operator algebras. When an appropriate modification was found, minimizing the equivalent functional proved to be challenging. However, significant progress was made by finding appropriate regularizations of the simplified functional that allow for better equations for the critical points.

Outlook: Future directions in free optimal transport include finding free probability analogs of entropy-regularized transport, as well as the existence of a comparison principle for a class of PDEs that arise in free non-commutative transport and control. Other directions include a fractal dimension quantity related to infinitesimal behavior of the Biane-Voiculescu-Wasserstein distance [[Shlyakhtenko's talk](#) in [Workshop II](#)]. There is also an emerging connection with continuous model theory from logic [see the work of David Jekel on this topic, as well as [his talk](#) during [Workshop II](#)].

5. Gradient Structures from Dilations

Introduction: Gradient structures provide a variational and geometric framework for describing dissipative dynamics, offering a bridge between thermodynamic principles and analytic tools. In the quantum domain, Carlen and Maas introduced a formulation of the Lindblad equation as a gradient flow of the quantum relative entropy with respect to a non-commutative Riemannian metric. This has led to substantial progress, including new quantum functional inequalities and improved understanding of entropy decay.

However, this structure was not directly derived from microscopic physical models. A key objective of this section is to investigate how such gradient structures may arise from dilations of Lindbladians—such as those involving quantum stochastic processes or system-environment interactions—motivated by large deviation theory. Additionally, we explore how transfer principles based on Lie group symmetries may relate classical and quantum gradient flows. Applications include refined convergence analysis for QMS and new perspectives on dynamical QOT, enriching the interface between geometry, analysis, and quantum thermodynamics.

Highlights: [Workshop II](#) addressed open problems and current research in the theory of QMS. Lectures conducted outside of workshop weeks introduced the Hudson–Parthasarathy quantum stochastic calculus, derivations of QSDEs via transfer principles, and the emergence of QSDEs from asymptotic limits of system-environment models.

Dilations of QMS: The *Stinespring dilation* theorem provides a powerful structural result: any completely positive trace-preserving map—describing the noisy evolution of a quantum system—can be modeled as a unitary evolution on an enlarged Hilbert space followed by a partial trace. This allows seemingly irreversible quantum evolutions to be viewed as projections of reversible dynamics in a higher-dimensional setting.

Markov dilations further refine this idea by embedding QMS into probabilistic frameworks. The foundational work of [Kümmerer and Maassen](#) introduced “essentially commutative” dilations, and [Junge and Mei](#) extended these definitions to the setting of tracial vNAs. Physically, such dilations offer a reversible microscopic description of systems interacting with thermal environments.

This line of inquiry is closely tied to developments in matrix concentration and non-commutative probability. In particular, [Huang and Tropp](#) established Poincaré-type inequalities for semiclassical matrix-valued processes using random matrix theory. Extending these results to fully quantum regimes via Markov dilations could pave the way for new transport inequalities and robust functional bounds relevant to quantum computing and simulation.

Quantum stochastic calculus: Quantum stochastic calculus generalizes classical stochastic analysis to non-commutative probability spaces, providing a rigorous language for describing quantum systems subject to noise. At the heart of this theory is the *Hudson–Parthasarathy calculus*, which formulates QSDEs on Fock space using quantum analogues of Itô integrals.

This framework naturally models *unitary dilations* of QMS and plays a central role in quantum control, optics, and measurement theory. A notable strength is its structural similarity to classical Itô calculus, enabling the transfer of classical stochastic methods into quantum settings via operator-valued processes and adapted filtrations. Quantum stochastic calculus captures the interplay between dynamics and continuous measurements, supporting both the theoretical analysis and practical design of dissipative quantum systems grounded in microscopic physical models.

Gradient structures for dilated QMS: Large deviation theory connects microscopic randomness to macroscopic behavior via rate functions, which quantify the likelihood of rare events. In many classical systems, these functionals can be reinterpreted as dissipation potentials, yielding gradient flow formulations for the corresponding equations. This approach has led to thermodynamically grounded interpretations of PDEs and facilitated connections between probability and geometry.

In the quantum setting, complications arise due to the lack of individual realizations of stochastic processes. One strategy is to operate at the level of classical path measures—where large deviation tools apply—and then contract the resulting gradient structures to the level of quantum states (density matrices). The Carlen-Maas gradient structure is based on GNS symmetry, and the dilation approach could generalize results to cases with less symmetry.

Work done at IPAM: Working groups worked on explicit constructions of Stinespring dilations and investigated gradient structures for toy examples. Tse gave a tutorial series on gradient structures and several talks on deriving gradient structures from microscopic dynamics via large deviations. Carlen gave a talk in [Workshop II](#) about

convergence to steady-state systems in quantum thermodynamic systems. Junge gave multiple talks connecting quantum mechanics to Lie algebras and the transference principle.

Explicit Stinespring Dilations. While the Stinespring theorem ensures the existence of dilations, finding explicit constructions remains challenging. We developed two examples. One involved pure decay—a two-level quantum system transitioning from an excited to a ground state. We identified a unitary dilation modeling this decay process. Interestingly, the same construction extends to a scenario with two decay channels, provided the environment is initialized in a specific state.

Toy Model: Gradient Structure from a Dilated Lindbladian. To better understand the full derivation process, we analyzed a toy model involving a two-level system with a jump operator given by the Pauli-X matrix and no Hamiltonian term. The Lindbladian was dilated via classical Brownian noise, producing an SDE that could be solved exactly. From this, we extracted a gradient structure at the level of path measures and attempted to contract it to the level of density matrices. Both the rate function and dissipation potential were successfully contracted, though the full characterization of the resulting gradient flow on quantum states remains ongoing.

Beyond GNS Symmetry. The Carlen–Maas framework utilizes the symmetry of the Lindblad operator with respect to the GNS inner product, establishing a differential structure for the involved non-commutative geometry. This symmetry leads to further thermodynamic symmetries, such as KMS and BKM. Subsequent research revealed that the Lindblad equation represents a gradient flow of the Umegaki relative entropy if it is symmetric with respect to the BKM inner product. We explored the construction of a non-commutative Riemannian metric in this broader context, considering recent advancements by Melchior Wirth, who reformulated such Lindbladians in Dirichlet form. This may facilitate the definition of suitable transport metrics beyond the GNS-symmetric case. The difference between these symmetries was discussed in a plenary presentation during the retreat.

Outlook: Having understood a basic example, the next step is to study more complex and physically relevant cases. For example, one can add a Hamiltonian term to investigate the interaction between conservative and dissipative terms, and investigate different types of dilations. Once such gradient structures are established, they can be used to derive functional inequalities (e.g., Poincaré, log-Sobolev) and inform the design of control strategies for high-dimensional open quantum systems.

6. Quantum Thermodynamics

Introduction: Quantum thermodynamics aims to explain and quantify at a fundamental level the entropic and energetic interactions between quantum processes. The key notions are free energy, work, and entropy, very much as in the classical regime. In the quantum regime, these fundamental notions require proper interpretation and definition. Quantum mechanics can be argued to be an L^2 -theory (vectors), in contrast to classical probability, which can be viewed as an L^1 -theory (distribution functions and densities). Basic operations are described by Hamiltonians and unitary operations. In contrast to classical thermodynamics, dissipation is derived from the Born–Markov approximation of the interaction with another quantum system—the heat bath. From this vantage point, quantum thermodynamics can be formulated in terms of density operators, quantum channels, and system Hamiltonians. Thence, work can be analyzed through the lens of quantum information theory as well as the theory of open quantum systems. Conceptualization and building of quantum heat engines require all aspects of quantum mechanical theory.

Highlight: A talk on “[Sub-Riemannian geometry of Optimal Transport](#)” was delivered by Tryphon Georgiou during [Workshop I](#). This was the starting point of our discussion, combining NOT, density paths, and quantum thermodynamics.

Energetic exchanges along curves in density space: A starting point of the analysis is the notion of free energy of a quantum system, namely, the difference between the internal energy and the entropy of the quantum state. The observable for the so-called work-rate is the derivative of the system Hamiltonian. A standard decomposition of the work in terms of change of free energy and dissipative terms ensues. The dissipative terms can be analyzed mathematically using Spohn’s entropy production. A specific choice of the Lindbladian derivative of the density path leads to paradigmatic estimates for work along the path.

Controllability of quantum trajectories and cycles: Controllability of the density path is obtained through the choice of a time-dependent Lindbladian and the underlying Hamiltonian path. A priori, there is no guarantee that periodic actuation leads to periodicity of the density paths. Even in simple model cases, generating quantum density cycles requires proper mathematical insight.

Work done at IPAM: The focus of the work at IPAM has been to clarify the thermodynamic aspects of the interaction between quantum systems. Specifically,

building on the blueprint of stochastic thermodynamics, we investigated quantitative metrics that explain energetic and entropic exchanges between quantum systems. After exploring proposals in the literature and drawing on analogies with classical thermodynamics, the focus has been on obtaining bounds on what can and cannot be achieved through such interactions as detailed in the following:

- (1) Derived upper and lower bounds on entropy production and energy exchange under suitable scenarios that expose the contribution of the speed of traversing a density path.
- (2) Obtained an operational interpretation of the interaction between Hamiltonian and dissipative parts for Lindblad generators by analyzing replacer channels for Lindbladian control. Concrete calculations reveal asymptotic periodicity for controllable paths.
- (3) Initiated a study of Lindbladian control for half densities with potential applications to the holonomy of density paths, building on discussions at IPAM.

Outlook: The work at IPAM shall be combined with previous results on reachability and Weyl calculus obtained by participants in preparing for the program. More general explicit solutions for density paths are expected to lead to further insights on energetic exchanges between quantum systems. A range of Lindbladians that are physically motivated needs to be investigated. Rates of convergence to periodic orbits are of great importance and need to be quantified. Such rates will help provide robustness guarantees for traversing quantum thermodynamic cycles.

7. Quantum Optimal Control and Algorithms

Introduction: In order to run any quantum algorithm, control pulses must be optimized via QOC to implement the basic gates used in the algorithm with a low gate error. Similarly, many problems can be formulated as variational problems, where the final state is generated by a circuit of parameterized gates. Variationally optimizing these gates can be a challenging problem due to the presence of limitations such as barren plateaus. To potentially address this challenge, we can optimize over a parametrized quantum circuit to reach the target state, using the quantum 1-Wasserstein distance as the objective function. This parametrization can be implemented as angles of rotation for gates or control pulse amplitudes.

Barren Plateaus: Solving variational problems on quantum computers is often hindered by the phenomenon known as *barren plateaus*. In the presence of a barren plateau, the cost function or its gradient becomes exponentially concentrated around the mean as

the number of qubits increases. This implies that, as the system size grows, the optimization landscape becomes increasingly flat and lacks informative features. As a result, it becomes difficult for optimization algorithms to identify meaningful descent directions, whether by evaluating differences in the cost function or by using gradient-based methods.

Recent approaches to circumvent such barren landscapes include using alternative objective functions, e.g., the quantum 1-Wasserstein distance (quantum Earth-mover's distance), as introduced by [De Palma et al.](#) Exploiting the unique properties of this 1-Wasserstein distance leads to more efficient and stable quantum learning, as shown by [Kiani et al.](#)

Landscape of variational quantum algorithms: Optimizing unstructured parametrized quantum circuits to generate arbitrary quantum states is not expected to be efficient, as their landscapes suffer from barren plateaus and also the existence of exponentially many bad local minima. Still, when restricted to preparing particular resource states for practical applications, variational optimization of the circuit can potentially be efficient. An example of such a state is the n -qubit GHZ state, a resource state for quantum metrology that can be used to achieve the Heisenberg limit in estimating parameters. As discussed before, it has been shown that in contrast to commonly used cost functions such as infidelity and trace distance, the quantum 1-Wasserstein cost function can provably avoid a barren plateau. The advantage of using optimal OT cost functions to enhance the optimization of quantum circuits for this toy model highlights the need for a systematic study to explore the landscape that various parametrizations of quantum circuits and cost functions can generate.

Landscape of pulse-based optimal control: Pulse-based QOC can be considered a special case of a variational quantum algorithm, in which the parameterization is not performed at the gate level, but rather at the level of the Hamiltonian used to control the quantum computer. A key question is whether the advantages of using the quantum 1-Wasserstein distance can also be applied to certain control pulse parameterizations in QOC. Analysis of pulse parameterizations is difficult due to the nonlinear dependence of the quantum state on pulse parameters; therefore, we conducted numerical experiments to analyze pulse parameterizations.

Limitations of quantum algorithms: QOT offers a versatile framework for deriving fundamental limits on quantum algorithms. By analyzing circuit depth and volume using quantum Wasserstein metrics, one can obtain general lower bounds on attainable performance. This viewpoint has yielded concrete limitations for several algorithms. Examples include the quantum approximate optimization algorithm (see [De Palma et](#)

[al.](#)), the dissipative approach to state preparation (see [Ding et al.](#)), and quantum SYK algorithms (see [Anschatz](#)).

Work done at IPAM: During this program, two working groups were formed to discuss QOT and NOT for quantum algorithms. These groups eventually merged to focus on mitigating barren plateaus in variational problems by using cost functions from OT. Numerical simulations of the variational circuit, described above, as well as an Ising spin-chain Hamiltonian system, were performed; heatmaps of the optimization landscape were made to graphically show the existence of barren plateaus (and non-existence when using the 1-Wasserstein distance in the variational circuit).

During the workshops, several presentations on the applications of NOT techniques for quantum algorithms and optimal control were delivered. In [Workshop I](#), Marvian [presented a talk](#) on the application of OT techniques to demonstrate the limitations of constant-depth quantum circuits in performing optimization and to highlight their application as the cost function of variational algorithms capable of learning quantum circuits. In [Workshop II](#), Jordan [presented a talk](#) on continuous quantum measurements and optimal control, as well as the potential application of OT methods. Moreover, Araiza presented a definition of transport cost for channels on general vNAs, sketching the potential application to quantum field theory. In another [presentation](#) by Wu, it was shown how circuit depth and volume can be mapped to OT bounds via Lipschitz geometry, yielding a general principle for quantifying fundamental limitations of quantum algorithms.

Outlook: We do not expect that using the 1-Wasserstein distance is a cure-all for the barren plateau issue. Rather, the cost function and circuit/pulse parameterization should be tailored to the goal. The 1-Wasserstein distance is an effective cost function for the goal of GHZ state preparation, as it captures local structure via the Hamming distance. For other classes of goals, e.g., compiling a family of quantum gates or states, other cost functions could be similarly effective. An important open question in designing optimal strategies is to understand the similarities and differences between optimizing variational circuits versus optimizing controlled Hamiltonians.

8. Quantum Many-Body Problem

Introduction: Although the state of a quantum system is completely specified by its density matrix, using the density matrix directly for quantum many-body problems is computationally intractable and conceptually challenging because its dimension grows exponentially with increasing system size. This motivates approaches based on

quantum marginals, also known as reduced density matrices, wherein most of the dimensions of the density matrix are traced out. Typically, one considers the 1- or 2-body reduced density matrix, as these objects are computationally tractable and conceptually facile, as they have the dimensionality of 1- and 2-particle systems, respectively.

The enormous simplification associated with reduced density matrices might seem extreme, but, for identical particles, information about the full density matrix can be reconstructed by (anti)symmetry. A key insight is that NOT can be used to reconstruct the full density matrix from its reduced density matrices or marginals. Specifically, given a cost (usually extracted from the physical Hamiltonian), minimization of the many-body cost subject to the constraint that the many-body density matrix has correct marginals allows one to not only define (and compute) the energy as a functional of the reduced density matrices, but also to reconstruct the full density matrix from its marginals. This insight recasts classic problems from electronic structure theory in the language of NOT.

For identical particles, all the marginals are the same, which is a very special (and unconventional) case of the general NOT problem. However, the same framework can be extended to treat non-identical particles (with potentially different marginals). The same problem statement can then be used to model quantum channels.

Highlights: Open problems and current research topics on quantum many-body theory were addressed in [Workshop I](#). Connections to reduced DMFT and quantum many-body physics were presented in a series of seminars and developed by several (in)formal working groups.

1-body Reduced DMFT: The most computationally efficient and conceptually straightforward approaches to the quantum many-body problem are based on the 1DM. This is an emerging research direction in molecular quantum chemistry, motivated by the failures of DFT for systems with strong electron correlation. In 1DMFT, one minimizes the energy as a functional of the 1DM, subject to the constraint that the 1DM's spectrum lies between zero and one (the Pauli exclusion principle) and the trace of the 1DM is the number of electrons. 1DMFT would be computationally feasible if good approximations for the interparticle repulsion energy as a functional of the 1DM were known.

As mentioned above, the *exact* functional, which was first formulated by Levy and Valone as a *constrained search problem* more than 40 years ago, can be viewed as a NOT problem. Using the NOT framework, and especially regularization, we can provide (nontrivial) benchmark data for the exact 1DM functional for the first time. Not only does

this provide a tool for 1DMFT that is analogous to existing tools for evaluating the exact density functional, but acquiring benchmark data for 1DMFT is computationally easier than its density functional theory counterpart.

At IPAM, benchmark data were generated for broad families of 1DMs and different costs/interparticle interactions. Alternative regularization techniques were explored. In addition, exact constraints and approximation strategies for the exact 1DM functional were developed, mostly by extending analogous results from DFT.

In DFT, classical OT theory can be used only to study a (relatively unimportant) limiting case: the strictly correlated limit. By contrast, NOT provides a full and rigorous theoretical framework for 1DMFT, which should be further pursued. Unfortunately, the absence of efficient computational methods means direct application of the NOT framework for *practical* 1DMFT computations is currently intractable.

***k*-body reduced density matrix functionals:** Reduced density matrices of higher order are useful because they allow the energy of many-particle systems (e.g., $k = 2$ for electrons and $k \geq 4$ for nucleons) to be computed exactly. As noted by Feliciangeli, Gerolin, and Portinale ($k = 1$) and by Ayers and Gerolin (general k), NOT provides an approach to the universal functional for k DMFT. This approach has advantages over traditional variational approaches for optimizing the k DM because the feasible set of k DMs is encapsulated in very complicated constraints by the k DMFT functional.

At IPAM, we developed the mathematical foundations of k DMFT using von Neumann-entropy regularization, proving the existence, duality, and related properties of the exact functional. Future work will include a proof of the existence of the maximizer for the dual formulation for arbitrary-order k DM, which will pave the way for the full characterization of the solutions and the development of theoretical guarantees for the convergence of practical computational algorithms.

OT & quantum many-body physics: One of the goals of mathematical analysis in quantum many-body physics is to rigorously derive effective equations in the limit of large particle numbers. At [Workshop I](#), recent advances in the effective description of bosonic systems were presented. These new results adopt a probabilistic framework, using tools such as Wasserstein metrics to quantify the convergence of many-body dynamics toward their effective descriptions. This innovative perspective not only reformulates existing results on the convergence of reduced particle densities but it also provides fresh mathematical and physical insights into this highly challenging problem. Going forward, we aim to further develop this approach and address complex questions, including systems with singular interactions and the thermodynamic limit.

In addition, as part of one of the working groups, the participants established a formal mathematical bridge between objects appearing in minimization problems on Fock spaces (particle physics and chemistry) and the objects encountered in QMOT (reduced density matrices, partial traces, marginal constraints, and optimality conditions).

Classical Limits: In the classical limit (where Planck’s constant approaches zero), the motion of electrons becomes strictly correlated, and DFT can be recast as a classical OT problem. Standard OT algorithms fail due to the problem’s extreme dimensionality (1 marginal per electron, with each electron’s position discretized by $\sim 10^4$ (or more) grid points). Participants approached this problem by

- Developing greedy algorithms that allowed each electron (pair) to be treated sequentially using a Monge ansatz.
- Representing the cost functional from entropically-regularized classical OT as graphs. The computational cost of evaluating each graph’s contribution is proportional to its tree width. Practical methods for including all graphs with tree width no greater than two were developed and shown to be asymptotically exact.

Quantum Multi-Marginal OT (QMOT) and QOT Barycenters: Comparing, summarizing, and reducing the dimensionality of empirical probability measures on a certain space are fundamental tasks in statistics and ML. Wasserstein barycenters of probability distributions possess many desirable properties and have had a significant impact both within mathematics and in broader applications. The goal of this research was to develop mathematical foundations for the quantum analog of Optimal Transport Barycenters.

We established the mathematical foundations of QOT Barycenters, including results on existence, duality, and related properties of a wide class of cost operators. We derived explicit analytical solutions in the case where the quantum states belong to a specific class of Gaussian states when the cost operator is defined via quadratures.

Our framework unifies existing approaches to QOT, encompassing both transport between density operators (as developed by [Golse et al.](#)) and quantum channels (following [De Palma and Trevisan](#)). Moreover, in cases where such analytical solutions are available, we demonstrate that the QOT barycenter coincides with the solution to a QMOT problem, formulated with a quantum analogue of the Gangbo–Święch cost. A general proof of this equivalence would establish a quantum counterpart of the foundational result by [Agueh and Carlier](#) (2010) in the classical setting of OT for probability measures.

Future directions include obtaining analytical solutions for more general Gaussian states and characterizing primal-dual solutions of these classes of variational problems.

Further, a suitably constructed continuous family of barycenters has been sought to provide a candidate notion of a surface in classical and non-commutative Wasserstein spaces. Interest in surfaces in a Wasserstein space arose in classical stochastic thermodynamics, where area integrals of work-density prescribe the work extracted by traversing a thermodynamic cycle (as explained in Movilla Miangolarra *et al.*). Introducing such notions of surfaces in classical and quantum Wasserstein spaces was a topic of discussion, as it may be of value in the development of quantum thermodynamics.

9. Computational Aspects for Quantum

Introduction: In OT, there are several well-established computational strategies, including linear programming, ML approaches, and convex regularization. In particular, convex regularization with the Shannon entropy yields the famous Sinkhorn algorithm, which has become a standard and relatively efficient computational approach. QOT, however, introduces fundamentally new challenges due to its non-commutative structure and formulation as a SDP. Even when regularized, QOT lacks a computationally practical analogue of the Sinkhorn algorithm. [Workshop III](#) presented state-of-the-art computational approaches for OT and highlighted the need to develop new computational strategies for both classical and quantum OT.

Inspired by the various definitions of QOT as discussed in Section 3, we discussed a range of computational frameworks. For some, we investigated direct approaches based on SDP and SOS hierarchies. For others, including those arising from reduced DMFT, we explored convex regularization techniques, machine learning-inspired methods, and strategies for generating synthetic data and constructing training samples. Our discussions aimed to bridge computational insights from classical OT and modern ML to develop new tools for tackling QOT and related problems.

Highlights: Participants developed bespoke neural network architectures (e.g., in the context of reduced density matrices) to compute NOT objects and generated appropriate data to train them. Participants investigated the properties of regularization beyond the von Neumann entropy and the properties of algorithms utilising them.

Regularization: Following the example of classical OT, we solve coupling-based QOT problems by adding convex regularization. The framework for such regularizations was

established by a paper by [Portinale, Caputo, Gerolin, and Monina](#), where it was shown that for finite-dimensional spaces the solution to the original QOT problem is recovered as the regularization parameter vanishes. Using convex duality principles, we recover a dual functional that is an unconstrained optimization over operators on Hilbert spaces, which, after discretization, can be recast as a matrix optimization problem. As a concave functional, the dual problem becomes amenable to gradient ascent and its refinements (e.g., quasi-Newton methods).

Additionally, thermal regularization is essential for computing two-electron reduced density matrices (2DMs) in systems with strong correlation or near-degeneracy. In these cases, the ground state spans a quasi-degenerate manifold, leading to many nearly equivalent solutions and numerical instability. By introducing an entropic penalty, thermal regularization promotes mixed-state ensembles, stabilizes the optimization, and yields physically meaningful results. It also reflects the statistical behavior of finite-temperature systems, making it a necessary tool—not just for numerical robustness, but for capturing ensemble-averaged physics beyond zero-temperature formalisms.

At IPAM, we developed algorithms and software for the exact 1DMFT functional using Tsallis entropy regularization (section 8). This approach eliminates the error associated with spectral truncation and allows the underlying (unregularized) universal functional to be estimated with machine precision. Future work will explore the optimal regularization strategy and the relevance of these computations for developing improved, approximate, 1DM functionals.

Machine learning strategies: Motivated by the recent work of [Amos et al.](#), which proposes an approach to solving classical OT problems using neural networks, we considered several ML approaches that could inspire new tools for NOT. Physics-Informed Neural Networks (PINNs) have gained attention as promising approaches for approximating solutions of PDEs while preserving physical consistency by incorporating physical laws into the loss function. These approaches can motivate new NOT loss functions, since some of the optimality conditions, particularly for regularized NOT, can be rewritten as Euler-Lagrange equations. The participants also discussed potential applications of PINNs to inverse problems, namely parameter estimation, and possible connections with objects in NOT.

Following discussions, several ideas were proposed to maximize the NOT functional via a neural network with marginal constraint measures as inputs, and maximizers of the NOT Dual problem as outputs. One possibility included the analogue of the mentioned work of Amos et al. via *unsupervised* learning. Another idea was based on *supervised*

learning, which would generate the training data of marginals and corresponding optimal potentials (variables of the Dual NOT)—these samples would subsequently be used to train the network. In both approaches, one could consider incorporating the PINNs-type constraints in the loss function.

SDP Implementation and SOS: In order to facilitate the optimization landscape experiments for variational quantum algorithms performed in Section 7, we wrote codes for computing 1-Wasserstein distances by solving the SDP formulation. In these experiments, because the dimension of the Hilbert space grows exponentially with the number of qubits, computing the 1-Wasserstein distance becomes impractical for even a moderate number of qubits. Consequently, it is crucial to develop computationally efficient methods for approximating the 1-Wasserstein distance in the context of quantum computing. In particular, the Lipschitz constant of the Hamiltonian in the dual formulation of the 1-Wasserstein distance can be replaced by an upper bound given by the operator norm of the decomposition of the Hamiltonian into a sum of Pauli strings, which results in an SDP that is cheaper to solve.

The Doherty–Parrilo–Spedalieri (DPS) hierarchy provides a sequence of semidefinite relaxations that outer-approximate the cone of separable bipartite quantum states. DPS hierarchy is a non-commutative analogue of the SOS/Lasserre hierarchy: degree $2k$ moment matrix in the SOS hierarchy becomes a k -symmetric extension of the bipartite state, and the SOS constraint is replaced by positivity under every partial transposition. The resulting cones are nested and converge asymptotically to the exact cone of separable states. This could be used to approximate the 2-Wasserstein distance, as proposed by [Beatty et al.](#) However, the method is still subject to SDP solvers, which are prone to the curse of dimensionality.

Outlook: NOT remains a computationally challenging area, and as described in Section 1, a wide variety of problems fall under its umbrella. Future directions in this area involve developing new solvers for quantities that are useful and increasing the robustness of current solvers—i.e., solving for larger electron systems in reduced DM problems. Furthermore, we need to characterize the properties of the algorithms developed during the long program, including convergence guarantees.