Technische Universität München
TUM School of Computation, Information and Technology

The curse of dimension in multi-marginal optimal transport

From a convex-geometric point of view to an efficient computational method

Daniela Vögler

Vollständiger Abdruck der von der TUM School of Computation, Information and Technology der Technischen Universität München zur Erlangung einer

Doktorin der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitz: Prof. Dr. Daniel Matthes

Prüfer*innen der Dissertation:

1. Prof. Dr. Gero Friesecke
2. Prof. Dr. Jean-David Benamou
3. Prof. Dr. Codina Cotar

Die Dissertation wurde am 04.10.2022 bei der Technischen Universität München eingereicht und durch die TUM School of Computation, Information and Technology am 27.04.2023 angenommen.

Zusammenfassung

Wir streben nach neuen Berechnungsverfahren für symmetrischen multi-marginalen Optimalen Transport (MMOT) mit Wenig-Körper Interaktionskosten. Bei wachsender Anzahl an Marginalen wird MMOT prohibitiv hochdimensional. Wir untersuchen die Suffizienz des klassischen dünnbesetzten Ansatzes. Außerdem stellen wir eine effiziente Methode vor, die Näherungslösungen im Fall von 2-Körper Interaktionen berechnet. Zuletzt zeigen wir, wie die Methode auf k-Körper Interaktionen für $k > 2$ angewendet werden kann.

Summary

We work towards new computational schemes for symmetric multi-marginal optimal transport (MMOT) with few-body interaction costs. MMOT becomes prohibitively high-dimensional for a growing number of marginals. We investigate the validity of the classic sparse ansatz (Monge's ansatz). Further, we introduce an efficient method that computes approximate solutions in the case of 2-body interactions. Finally, we provide an ingredient that allows us to apply the method to k-body interactions for $k > 2$.

Success is not final, failure is not fatal: it is the courage to continue that counts.

- *Winston S. Churchill*

Acknowledgments

It is my joy to have typed in the last corrections and finally, be able to thank my supporters.

First and foremost, I want to thank my advisor, Prof. Gero Friesecke, for his continuous support. He was open to me working at the intersection of analysis and discrete mathematics and even ventured with me into this promising terrain. For this and all the things I learned while working with him, I am truly thankful.

Secondly, I want to thank all the people behind the scenes of this dissertation: Daniel Matthes, Codina Cotar, and Jean-David Benamou for serving on my thesis committee; Agnieszka Baumgärtel for patiently answering my many questions on diverse topics; Dr. Carl Friedrich Kreiner for always having his door open and waiting behind it with advice and support; the elite program TopMath for having me; Dr. Kathrin Ruf for agreeing to be my mentor; and especially Lydia Weber and Katja Kröss for their never-ending supply of answers, reminders, and support.

Next, I want to thank TUM itself: M7 (especially Prof. Gero Friesecke, Prof. Marco Cicalese, and Prof. Simone Warzel) for providing an environment that allowed this dissertation to develop and mature; Frauke Bäcker for always having an open ear and keeping the place (M7) running; Rufat and Sören for always bringing light into my day; Dr. Hans-Peter Kruse for our heart-warming teamwork on various exercise classes; and finally, the students for keeping us motivated.

Over the last years, I have been lucky to meet and work with incredibly bright people. While I want to thank TUM itself and Prof. Gero Friesecke for making this possible, I would be remiss if I did not highlight a few of said bright people and thank them for their contribution: Prof. Andreas S. Schulz for his more than generous financial support and the time he invested in our collaborations; Prof. Guillaume Carlier for inviting me to a short research stay in Paris, our chalk-hand-hours on the blackboard during his summer term in Munich, and all the hours he invested in our collaborations.

Let me conclude the list of acknowledgements with the most important people in my life. Danke, Oma, für unzählige Apfelmaultaschen, Liwanzerl und dafür, dass Du mir bei „Mensch ärger Dich nicht“ gezeigt hast, was Ehrgeiz ist. Danke, Opa, für unzählige Kniffelrunden, Biathlonrennen und Deinen Glauben an mich. Es tut mir leid, dass Ihr das Einreichen meiner Dissertation nicht mehr miterleben konntet. Danke, Renate und Jürgen, für Eure Unterstützung, Nachsicht und Großzügigkeit. Danke, Thomas, für Deinen Humor, Deine Lockerheit und unsere Verbindung. Danke, Mama und Papa, für Eure Zuverlässigkeit, Geduld und Stärke. Danke, Regina, dafür, dass Du genau zum richtigen Zeitpunkt in mein Leben zurückgekommen bist, immer an mich glaubst und mir zeigst, was da draußen auf uns wartet. Danke, Andreas, dafür, dass Du mein Alltagsheld, mein Herz und mein Fels in der Brandung bist; diese Dissertation ist Dir gewidmet.

Finally, let me apologize to all the people I forgot to acknowledge. Finishing a dissertation is a stressful time that makes you forget seemingly unforgettable acts of kindness and support.

sgd. Daniela Vögler

List of contributed articles

Core publications

A.1 Daniela Vögler
Geometry of Kantorovich polytopes and support of optimizers for repulsive multi-marginal optimal transport on finite state spaces
J. Math. Anal. Appl., 502(1):125147, (2021)

A.2 Gero Friesecke, Andreas S. Schulz, and Daniela Vögler
Genetic Column Generation: Fast Computation of High-Dimensional Multi-marginal Optimal Transport Problems
SIAM J. Sci. Comput., 44(3):A1632-A1654, (2022)

Further publication

B.1 Guillaume Carlier, Gero Friesecke, and Daniela Vögler
Convex geometry of finite exchangeable laws and de Finetti style representation with universal correlated corrections
Probab. Theory Relat. Fields, (2022)

I, Daniela Vögler, am principle author of Core Publication A.1 and A.2.

Contents

1 Introduction **3**

2 The theory of optimal transport **7**
2.1 Two-marginal optimal transport 7
2.2 Multi-marginal optimal transport 15

3 Motivation **19**
3.1 Electronic structure in optimal transport 19
3.2 Discretization 25
3.3 The curse of dimension 27

4 Convex geometry of transport polytopes **29**
4.1 Notions of convexity 29
4.2 Insufficiency of Monge's ansatz in multi-marginal optimal transport 31
4.3 Quasi-Monge ansatz 36

5 Numerical approaches to multi-marginal optimal transport **41**
5.1 The standard in two marginals: Sinkhorn 42
5.2 Genetic Column Generation 44
5.3 How 'hard' are the problems? 48
5.4 Conclusion and outlook 49

6 Convex geometry of N-representable measures **51**
6.1 Application in the column-generation-context and definition of N-representability 51
6.2 Extremal N-representable measures 56
6.3 Application in the context of finitely extendible random vectors . . 59

Bibliography **65**

A Core publications **73**
A.1 Geometry of Kantorovich polytopes and support of optimizers for repulsive multi-marginal optimal transport on finite state spaces . . 73
A.2 Genetic Column Generation: Fast Computation of High-Dimensional Multimarginal Optimal Transport Problems 106

B Further publication **133**

B.1 Convex geometry of finite exchangeable laws and de Finetti style representation with universal correlated corrections 133

1 Introduction

Since its conception in 1781, the theory of optimal transport has gone from the problem of shovelling a pile of sand into a hole with the minimum effort possible to a well-developed and active field of research with connections to numerous branches in mathematics and a variety of applications. In more recent history, efforts have been made to understand an extension to multiple marginals. While the original problem aims at optimally coupling two marginals, the multi-marginal problem aims at optimally coupling multiple marginals. Such multi-marginal optimal transport (MMOT) problems arise for example in fluid dynamics, electronic structure, and data science.

The high-dimensional nature of MMOT problems affects the computation of their solutions. MMOT problems suffer from the curse of dimension: if we want to couple N marginals each of which lives on ℓ discretization points the admissible couplings consist of ℓ^N components. The ensuing exponential growth of the storage complexity renders standard methods infeasible for growing N.

In fluid dynamics, electronic structure, and data science the number of marginals N corresponds to the number of timesteps, electrons, and datasets, respectively. Consequently, these applications call for large N. With the standard being infeasible, new computational methods (possibly ones tailored to specific applications) are required. In this dissertation, we work towards such a new computational method for the problems arising in electronic structure. We use methods from convex analysis, optimization, discrete mathematics, and measure as well as probability theory.

Outline. In Chapter 2, we successively introduce the theories of two- and multi-marginal optimal transport.
In Chapter 3, we explain what motivates this dissertation. We shine a light on the role of MMOT in electronic structure: continuous symmetric MMOT problems arise as strong-interaction limit. Applying a structure-preserving discretization leads to discrete symmetric MMOT problems. We provide a measure-theoretic and linear-programming formulation of these discrete problems. Despite the dimension-reduction due to the symmetry, these problems still suffer from the curse of dimension.
We dedicated Chapter 4, Chapter 5, and 6 to Core Publication A.1, Core Publication A.2, and Further Publication B.1, respectively. In each of these chapters,

we highlight the main contributions of the article at hand and sketch the literary landscape surrounding them.
In Core Publication A.1, we numerically study the insufficiency of Monge's ansatz (the classic sparse ansatz in optimal transport) for symmetric MMOT on finite state spaces with a uniform marginal constraint. This study is limited to small numbers of marginals N and states ℓ. Despite its limitation, the study bore fruit: it inspired a necessary support-condition on optimizers that applies to general $N, \ell \geq 2$ and symmetric two-body interactions whose diagonal entries exceed a specified problem-dependent threshold. If N and ℓ display specified relationships, the necessary support-condition proved to be strong enough to guarantee the existence of a unique optimizer and not only imply its compliance with Monge's ansatz but also explicitly provide this optimizer. In Section 4.1, we recall relevant notions of convexity. In Section 4.2, we discuss the insufficiency of the classic sparse ansatz, Monge's ansatz. In Section 4.3, we relay a new sparse ansatz for symmetric MMOT on finite state spaces which is referred to as quasi-Monge. It was recently introduced in [44], where it was proven to be sufficient.
In Core Publication A.2, we introduce a simple and extremely efficient computational method, called Genetic Column Generation, that provides approximate solutions for symmetric MMOT problems on finite state spaces with two-body interaction costs. In Section 5.1, we describe the standard computational method in two marginals, Sinkhorn's algorithm, and explain, why its multi-marginal adaptation suffers from the curse of dimension. In Section 5.2, we present Genetic Column Generation in the context of the simplex method. In Section 5.3, we elaborate on recent complexity-theoretic work on MMOT. Finally, in Section 5.4, we draw a conclusion on Genetic Column Generation.
In Further Publication B.1, we shine a light on the convex geometry of the set of k-point marginals of certain high-dimensional objects. The key advancement consists in an explicit polynomial formula that expresses the extreme points of the set in terms of their one-point marginal. In Section 6.1, we relate the findings in Further Publication B.1 to Genetic Column Generation. In Section 6.2, we present the findings and preceding literature in their bare form. In Section 6.3, we translate our findings into probabilistic nomenclature. This leads to a finite form of de Finetti's theorem, which we compare and contrast with the original, infinite, form and an alternative finite form.
In the appendix, we included the contributed articles: Core Publication A.1, Core Publication A.2, and Further Publication B.1. Each article is preceded by a summary, a declaration of the present authors contribution, as well as the publisher's permission to include.

Notation. We tried to keep the notation as intuitive and simple as possible

and explained each piece of notation at first appearance with the exception of $\mathcal{P}(X)$ and $C_b(X)$ whose definitions we pre-empt. With $\mathcal{P}(X)$, we denote the set of (Borel) probability measures on X. For most parts, X will be a finite state space consisting of a finite number of pairwise distinct elements. In Chapter 6, for example, we go as far as allowing X to be a Polish space, a much more general notion of state space. Finally, $C_b(X)$ consists of all bounded continuous functions on X.

2 The theory of optimal transport

The goal of this chapter is to briefly introduce the theory of optimal transport (OT). We pay tribute to the origins of this very active field of research while focussing on the aspects of the theory relevant to this dissertation. The question stearing this introduction lies at the heart of the field and discusses if we can transport efficiently or even optimally without splitting mass. While this question has always been of theory-driven interest, the potential use of OT in high-dimensional computations has prompted a data-driven interest [40]. Transports that do not split mass move all the mass located at a source point x to a uniquely determined sink point $y = T(x)$. Storing 'no-mass-split transports' is much cheaper than storing 'mass-split transports': for a given source point x one only has to remember the unique sink point y instead of remembering multiple sink points $y_1, \dots, y_m$ and how much mass is transported to each one of them. The central task of the present dissertation is to use the data-reduction implied by sparse optimizers to design an efficient computational method for high-dimensional OT problems. This makes the question whether sparse optimizers do exist of central importance.

Our presentation of the topic leans on the excellent textbook by Santambrogio [77] titled *"Optimal Transport for Applied Mathematicians"*. We drew further inspiration from the lecture notes [40] on OT by Gero Friesecke which are perfectly suited to guide novices. As indicated by its name, the book by Santambrogio tells the OT-story from an application-driven point of view. For complementary reading, we refer the reader to the classic textbooks by Villani [82, 83] and the probability-flavoured accounts by Rachev and Rüschendorf [74, 75].

2.1 Two-marginal optimal transport

In 1781 Gaspard Monge [68] initialized the study of a transport problem:

> *Move a pile of sand from a grass-area into a sandbox in such a manner that no sand is left on the grass and the sandbox is filled up evenly. Minimize your effort.*

In modern mathematical language, densities

$$f_1, f_2 : \mathbb{R}^d \to \mathbb{R} \text{ with } f_1, f_2 \geq 0 \text{ and } \int f_1 = \int f_2 = 1$$

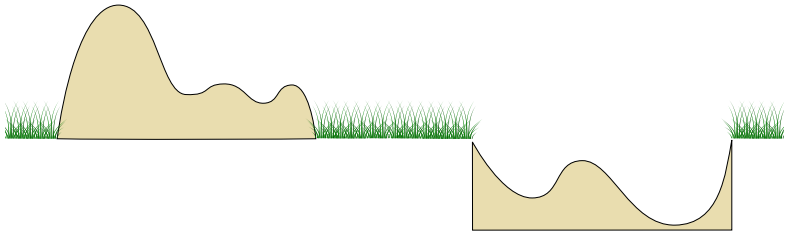

Figure 2.1: Visualization of the motivating problem inspired by a similar illustration in [82].

embodied the pile of sand on the grass and the hole in the sandbox. The natural choice for d is 2 or 3. Note that the mass of sand that we have to move equals the mass of sand that the sandbox requires. To represent the transport itself, Monge used measurable maps

$$T : \mathbb{R}^d \to \mathbb{R}^d, x_1 \mapsto T(x_1).$$

For any x_1, T transports all the mass located at x_1 to $T(x_1)$. The covered distance

$$|T(x_1) - x_1|$$

measures the effort it takes to move one unit of mass from x_1 to $T(x_1)$. Here $|\cdot|$ denotes the Euclidean norm. Monge wanted to minimize the effort of the transport

$$\int_{\mathbb{R}^d} |T(x_1) - x_1| f_1(x_1) \, \mathrm{d}x_1.$$

So far we have not subjected the maps T to any constraints. Therefore, $T(x) = x$ would be a valid choice with transport cost 0: without any constraints the identity would describe an optimal transport, independent of the choice of f_2 (the hole). This is hardly what Monge wanted! How did he incorporate the target density f_2 in his mathematical model of the problem? We delay answering this question until we have introduced a more general version of Monge's problem.

In our general problem version, X_1 and X_2 denote the source respectively sink space. Textbook accounts often present (at least part of) the subject in the general framework of Polish spaces (which are complete and separable metric spaces). Here, the reader can still think of the underlying spaces X_1 and X_2 to be compact subsets of the Euclidean space $\mathbb{R}^d$. For our intents and purposes - bringing the reader in touch with the roots of the later presented research - we believe the simplification to be appropriate. Given X_1 and X_2, probability measures $\mu_1 \in \mathcal{P}(X_1), \mu_2 \in \mathcal{P}(X_2)$ and a more general cost function

$c : X_1 \times X_2 \to \mathbb{R} \cup \{+\infty\}$ replace the densities f_1, f_2 and the 'covered distance' $|x_1 - x_2|$. The notion of transport - a map T transports μ_1 onto μ_2 - stays the same. Figure 2.2 shows what roles the problem-inputs X_1, X_2, μ_1, μ_2, and c play in the 'sandbox-example'.

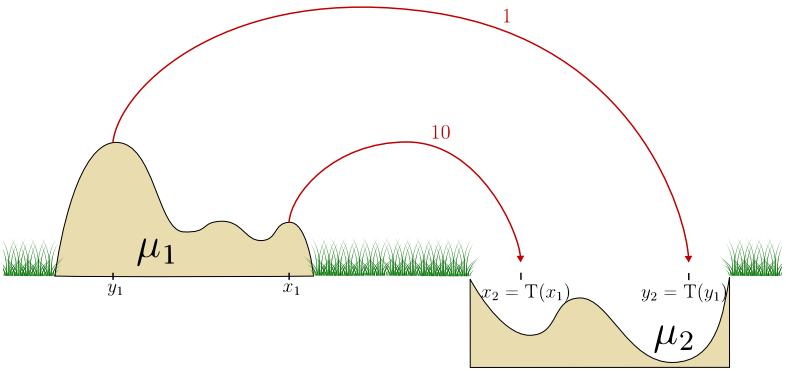

Figure 2.2: The present figure visualizes all components of Monge's OT. We take X_1 and X_2 to be the surface area beneath the pile of sand and above the hole in the sandbox, respectively; μ_1 and μ_2 then are probability measures on X_1, respectively, X_2 and form the pile of sand, respectively, the hole in the sandbox. The arrows represent a transport T: for any given source point x_1, the arrow originating at x_1 points to the uniquely determined sink point $T(x_1)$. The numbers above the arrows show how much it costs to transport one unit of mass from the source point of the arrow to its sink point. The given example highlights the fact that we deal with a general cost function c, which does not necessarily increase with the distance from source point to sink point.

We are now ready to formalize Monge's OT problem. For a given input-tuple $(X_1, X_2, \mu_1, \mu_2, c)$, it reads as follows. (Check below for a definition of the push-forward $T_\# \mu_1$ of μ_1 along a measurable map T.)

Monge's OT problem aims to

$$\text{Minimize} \quad I[T] := \int_{X_1} c(x_1, T(x_1)) \, \mathrm{d}\mu_1(x_1)$$

$$\text{subject to} \quad T \in \{T : X_1 \to X_2 : \quad T \text{ is measurable and fulfills } \underbrace{T_\# \mu_1 = \mu_2}_{\substack{\text{in words: } T \\ \text{transports } \mu_1 \\ \text{onto } \mu_2}}\}.$$

Now, how did Monge incorporate the target μ_2? Take a measurable map $T : X_1 \to X_2$ that is admissible regarding an instance $(X_1, X_2, \mu_1, \mu_2, c)$ of Monge OT - a so-called transport map. The notion of transport - T pushes μ_1 onto μ_2 - takes on the form '$T_{\#}\mu_1 = \mu_2$'. For a measurable map $T : X_1 \to X_2$ and a probability measure $\mu_1 \in \mathcal{P}(X_1)$, the measure $T_{\#}\mu_1 \in \mathcal{P}(X_2)$ is defined by

$$T_{\#}\mu_1(A) = \mu_1\left(T^{-1}(A)\right) \quad \text{for all measurable } A \subset X_2, \tag{2.1}$$

$$\text{or} \int_{X_2} \varphi \mathrm{d}\left(T_{\#}\mu_1\right) = \int_{X_1} \varphi \circ T \mathrm{d}\mu_1 \quad \text{for all measurable functions } \varphi \text{ on } X_2. \tag{2.2}$$

We refer to $T_{\#}\mu_1$ as the push-forward of μ_1 along T. In Figure 2.3 the reader can see that the push-forward constraint, $T_{\#}\mu_1 = \mu_2$, ensures that the mass of sand moved to a part of the hole $A \subset X_2$ - $\mu_1\left(T^{-1}(A)\right)$ - corresponds to the mass of sand required to fill the part A evenly - $\mu_2(A)$.

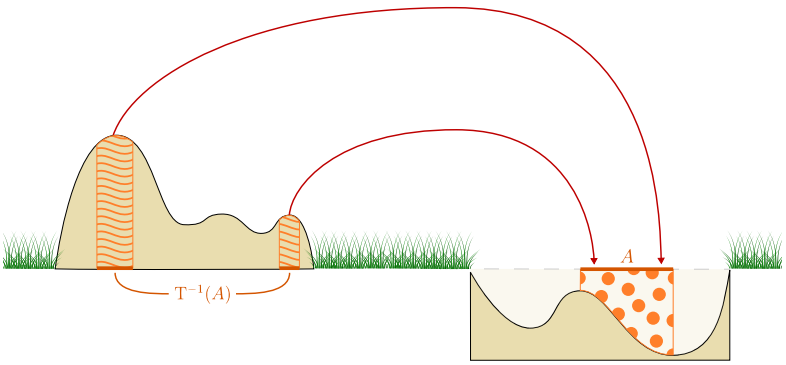

Figure 2.3: The present figure visualizes the meaning of the constraint $T_{\#}\mu_1 = \mu_2$. The probability measures μ_1 and μ_2 are fixed inputs for any given instance of Monge's OT. Any such instance then only admits a measurable map T as admissible transport if it fulfills $T_{\#}\mu_1 = \mu_2$, that is: for any subset A of X_2, the mass of sand that is moved to A - $\left(T_{\#}\mu_1\right)(A) = \mu_1\left(T^{-1}(A)\right)$ (mass with wavy lines) - has to equal the mass of sand required at A - $\mu_2(A)$ (mass with spots).

Many questions surrounding Monge's OT problem - as for instance the fundamental issue whether minimizers do exist - remained open for more than 150 years. In came Leonid Kantorovich [56] in 1942 with a change of perspective that led to a different OT problem formulation.

Kantorovich's OT problem aims to

$$\begin{aligned} &\text{Minimize} && C[\gamma] := \int_{X_1 \times X_2} c\,(x_1, x_2)\,\mathrm{d}\gamma\,(x_1, x_2) \\ &\text{subject to} && \gamma \in \{\gamma \in \mathcal{P}\,(X_1 \times X_2) : \underbrace{M_{X_1}\gamma = \mu_1 \text{ and } M_{X_2}\gamma = \mu_2}_{\text{in words: } \gamma \text{ couples } \mu_1 \text{ and } \mu_2}\}. \end{aligned}$$

The one-point marginal maps $M_{X_1} : \mathcal{P}\,(X_1 \times X_2) \to \mathcal{P}\,(X_1)$ and $M_{X_2} : \mathcal{P}\,(X_1 \times X_2) \to \mathcal{P}\,(X_2)$ are defined as follows.

$$\begin{aligned} M_{X_1}\gamma\,(A_1) &= \gamma\,(A_1 \times X_2) \quad \text{for all measurable } A_1 \subset X_1 \\ M_{X_2}\gamma\,(A_2) &= \gamma\,(X_1 \times A_2) \quad \text{for all measurable } A_2 \subset X_2. \end{aligned}$$

They extract the one-point marginals regarding the first (M_{X_1}) and the second (M_{X_2}) coordinate. Paralleling the term 'transport maps', transport plans refer to probability measures on $X_1 \times X_2$ that are admissible regarding an instance $(X_1, X_2, \mu_1, \mu_2, c)$ of Kantorovich OT.

Where is the transport in Kantorovich's OT problem formulation? Just as in Monge OT, any admissible 'point' represents a possible transport. Here, the admissible 'points' are probability measures on $X_1 \times X_2$ that fulfill the given marginal constraints. Any such probability measure γ captures a transport in a 'bookkeeping' manner: for $A_1 \subset X_1$ and $A_2 \subset X_2$ measurable, the amount of mass moving from A_1 to A_2 is given by $\gamma\,(A_1 \times A_2)$. The first marginal constraint

$$\gamma\,(A_1 \times X_2) = \mu_1\,(A_1) \quad \text{for all } A_1 \subset X_1 \text{ measurable}$$

ensures that it is exactly the pile of sand that we transport. The second marginal constraint

$$\gamma\,(X_1 \times A_2) = \mu_2\,(A_2) \quad \text{for all } A_2 \subset X_2 \text{ measurable}$$

ensures that we fill up the hole evenly. As $\gamma\,(X_1 \times A_2)$ equals all the mass that we transport to A_2, the second marginal constraint parallels the push-forward constraint. For an illustration of the differences between Kantorovich and Monge transports see Figure 2.4.

How do both problems relate to one another? Any transport map T gives rise to a transport plan, denoted by γ_T, with the same transport cost. Assume we face a transport problem and take both formulations - Kantorovich's as well as Monge's - into account. Any given transport map T fulfills - by definition -

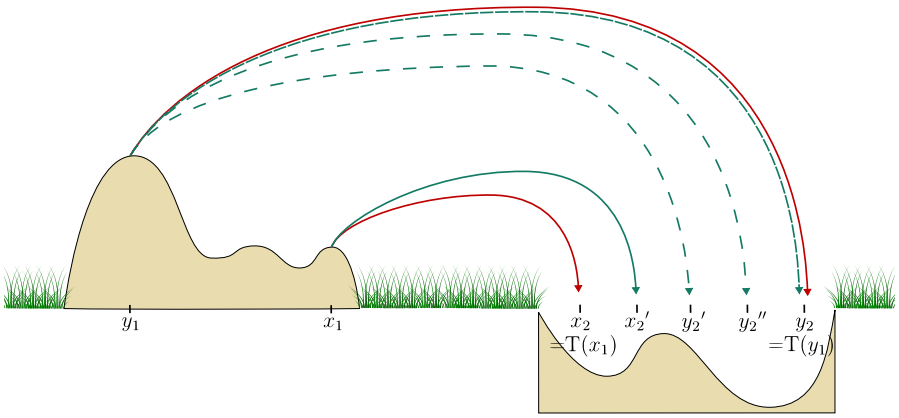

Figure 2.4: The present figure points out the difference between Kantorovich (green) and Monge (red) OT. The role of the spaces X_1 and X_2, the measures μ_1 and μ_2, and the cost function c is the same in both formulations. The difference consists in the manner in which they model the transport. Regarding Monge's OT, any given source point $x_1 \in X_1$ only chooses its sink point $T(x_1) = x_2 \in X_2$. This choice then allocates all the mass at x_1 to the sink $T(x_1)$. The appearance of each individual transport arrow carries meaning: the smaller the distance between two dashes, the more mass is moved. Accordingly, we show all the red arrows in their entirety. Regarding Kantorovich's OT, any given source point may choose multiple sink points and distribute its mass among them. The illustration shows that the source point y_1 partners with at least three sink points - y_2, y_2', and y_2'' - and moves most of its mass to y_2; the source point x_1, however, only partners with x_2' and moves all its mass there, no mass goes to x_2. While maps from X_1 to X_2 embody transports in Monge OT, probability measures on $X_1 \times X_2$ take on this role in Kantorovich OT. The reader can think of them as a plan laying out how much mass is moved from any source point $x_1 \in X_1$ to any sink point $x_2 \in X_2$.

$T_{\#}\mu_1 = \mu_2$. Let $(id, T) : X_1 \to X_1 \times X_2$ be defined by $x_1 \mapsto (x_1, T(x_1))$. Then the probability measure $\gamma_T = (id, T)_{\#}\mu_1$ on $X_1 \times X_2$ fulfills both marginal conditions and therefore is a transport plan. Plugging in γ_T in the objective function of Kantorovich's OT problem $C[\cdot]$ and using (2.2) shows that $C[\gamma_T] = I[T]$ does indeed hold. Kantorovich adapted Monge's OT problem by relaxing the idea of what transport is. In his textbook, Santambrogio presents a relaxation result, which - in his own words - *"means, roughly speaking, that (KP) is somehow the minimal extension of (MP) which has some chances to admit a minimizer"* [77, p. 3]. Here (KP) and (MP) refer to Kantorovich's respectively Monge's OT problem. The result also implies that the infima of both formulations coincide (under the

given - rather general - assumptions).
Further, if a transport plan γ is of the form $(id, T)_{\#}\mu_1$ for a measurable map $T : X_1 \to X_2$ then T is a transport map. Transport plans that split mass as shown in Figure 2.4 have no corresponding transport map. Accordingly, a transport map can push a source measure μ_1 that is given by δ_{a_1} for some $a_1 \in X_1$ only onto another Dirac measure $\mu_2 = \delta_{a_2}$ for some $a_2 \in X_2$ (with $a_2 = T(a_1)$).
Before coming to a conclusion on the relation between both formulations, we draw the attention of the reader again to Figure 2.4. It already indicates "Monge $\subset$ Kantorovich": transport maps induce exactly those transport plans that consist of complete arrows only. Readers with a background in discrete mathematics might see the similarity to integer points in a polytope - like permutation matrices in the set of biostochastic matrices. In Chapter 4, we will take a deeper look at this similarity.
Let us finally summarize the relationship between Kantorovich and Monge OT. Kantorovich introduced a new OT framework that elegantly enclosed Monge's notion of transport: Monge's transport maps directly yield admissible transport plans of the same cost and nicer (linear) constraints replace the disconcerting one governing the set of Monge transports.

Next, we present two of the most important results within the theory of two-marginal OT.

Kantorovich applies a linear objective function to probability measures that fulfill linear equality constraints. This description raises an array of questions: Does a dual problem exist? What does it look like? How do primal and dual relate? The following theorem answers all of these questions.

Theorem 2.1 (Kantorovich duality)**.** *For compact subsets X_1, X_2 of $\mathbb{R}^d$, we consider the probability measures $\mu_1 \in \mathcal{P}(X_1)$, $\mu_2 \in \mathcal{P}(X_2)$ and a continuous cost function $c : X_1 \times X_2 \to \mathbb{R}$. The following duality relation holds.*

$$\min_{\substack{\gamma \in \mathcal{P}(X_1 \times X_2) \\ M_{X_1}\gamma = \mu_1 \\ M_{X_2}\gamma = \mu_2}} \int_{X_1 \times X_2} c \mathrm{d}\gamma = \max_{\substack{\varphi_1 \in C_b(X_1), \varphi_2 \in C_b(X_2) \\ \text{subject to } \varphi_1 \oplus \varphi_2 \leq c}} \int_{X_1} \varphi_1 \mathrm{d}\mu_1 + \int_{X_2} \varphi_2 d\mu_2.$$

For $\varphi_1 \in C_b(X_1)$ and $\varphi_2 \in C_b(X_2)$, the function $\varphi_1 \oplus \varphi_2$ on $X_1 \times X_2$ is defined as $(\varphi_1 \oplus \varphi_2)(x_1, x_2) := \varphi_1(x_1) + \varphi_2(x_2)$.

A few remarks on this duality. Both optimal values are attained. The term 'optimal transport plan' refers to a minimizer of the primal problem, while the term 'Kantorovich potential' refers to a maximizer of the dual problem. The primal problem admits a solution under fairly general assumptions: we only require X_1 and X_2 to be Polish spaces and $c : X_1 \times X_2 \to [0, +\infty]$ to be lower

semi-continuous. By loosening the assumptions on the problem-specifying objects X_1, X_2, μ_1, μ_2, and c to that degree, we lose dual attainment.

Kantorovich's duality often plays an important role in proofs of the upcoming 'Monge result'. Before we present it, let us briefly and nonrigorously elaborate on how to conceive it. Kantorovich's duality, as stated above, not only provides us with a primal minimizer γ and a dual maximizing pair (φ_1, φ_2) but also relates these optimizers via complementary slackness:

$$\varphi_1(x_1) + \varphi_2(x_2) = c(x_1, x_2) \text{ on the support of } \gamma. \tag{2.3}$$

As $\varphi_1(x_1) + \varphi_2(x_2) \leq c(x_1, x_2)$ holds everywhere, (2.3) yields: for any supported pair (x_1, x_2), x_1 minimizes $x \mapsto c(x, x_2) - \varphi_1(x) - \varphi_2(x_2)$. Neglecting the boundary and assuming differentiability leads us to

$$\nabla\varphi_1(x_1) = \nabla_{x_1} c(x_1, x_2) \text{ on the support of } \gamma. \tag{2.4}$$

Now take the quadratic cost function $c(x_1, x_2) = \frac{1}{2}|x_1 - x_2|^2$, plug it into (2.4) and solve for x_2:

$$x_2 = x_1 - \nabla\varphi_1(x_1) \text{ on the support of } \gamma. \tag{2.5}$$

Exactly one x_2 for every x_1. The following result points out that the optimal transport plan for the quadratic cost is induced by a transport map ($T(x_1) = x_1 - \nabla\varphi_1(x_1)$) and unique. We present a version that we can immediately generalize to strictly convex costs.

Theorem 2.2 (Existence of optimal transport plans that are induced by transport maps)**.** *Let the compact set $X = X_1 = X_2 \subset \mathbb{R}^d$ be connected and with negligible boundary. As usual, μ_1 and μ_2 denote probability measures on X; let μ_1 be absolutely continuous. We consider the quadratic cost function $c : X \times X \to \mathbb{R}$ defined through $c(x_1, x_2) := \frac{1}{2}|x_1 - x_2|^2$ (where $|\cdot|$ denotes the Euclidean norm). Recall that the phrase 'γ is induced by T' means $\gamma = (id, T)_{\#}\mu_1$.*
Then the optimal transport plan γ is unique and induced by the transport map T given by

$$T(x_1) = x_1 - \nabla\varphi_1(x_1),$$

where φ_1 is a Kantorovich potential.

A few remarks on this sparsity result. If we replace the quadratic cost with any cost c of the form

$$c(x_1, x_2) = h(x_1 - x_2) \quad \text{for } h \text{ strictly convex,}$$

the statement of the theorem remains true. These cost functions, again, allow us to solve (2.4) for x_2 and maintain the status 'exactly one x_2 for every x_1' on the support of an optimal γ. A generalization of equation (2.5),

$$x_2 = x_1 - (\nabla h)^{-1} (\nabla \varphi_1 (x_1)),$$

embodies the dependence of x_2 on x_1. Accordingly, the expression for the transport map T needs to be updated.
Observe that it is crucial that we can solve equation (2.4) uniquely for x_2. The so-called twist condition on cost functions c requires the maps $x_2 \mapsto \nabla_{x_1} c(x_1, x_2)$ to be injective for every x_1: one x_2 for every possible function value. Chapters 9 and 10 in [83] present results of the same flavour as the above theorem for cost functions governed by a variety of assumptions including the twist condition. The presentation starts off with the quadratic cost and moves along in a very detailed manner. The bibliographical notes in these chapters provide a very nuanced view on the origins of these results. Here we want to keep it simple and refer the reader to Brenier's work on the quadratic cost [16, 17] and Gangbo's and McCann's work on strictly convex costs [45].

In this section we introduced the theory of two-marginal OT. The question - can we transport optimally when using maps only - was centre stage. But as Santambrogio already wrote: *"the theory of optimal transport cannot be reduced to the existence and the properties of optimal maps. The success of this theory can be associated to the many connections it has with many other branches of mathematics"* [77, p. xvii]. OT appears for example in research on partial differential equations, gradient flows, machine learning and density functional theory. Particularly the last mention leads generally to more than two marginals.

2.2 Multi-marginal optimal transport

When shifting our focus from two-marginal to multi-marginal OT, we move from coupling two marginals to coupling three or more marginals. How does this shift alter the transport models of the last section? For $N > 2$ - with $N \in \mathbb{N}$ denoting the number of marginals - the new problem-specifying objects are simple to deduce: instead of two, we consider N underlying spaces $X_1, X_2, \ldots, X_N$ (the reader can, again, think of them as compact subsets of the Euclidean space $\mathbb{R}^d$); instead of two, we consider N marginals $\mu_1 \in \mathcal{P}(X_1), \ldots, \mu_N \in \mathcal{P}(X_N)$; instead of cost functions on the two-fold product $X_1 \times X_2$, we consider cost functions $c : X_1 \times \ldots \times X_N \to \mathbb{R} \cup \{+\infty\}$ on the N-fold product. The adaptions of both OT problem formulations are equally straight-forward:

Monge's MMOT problem aims to

$$\begin{aligned} \text{Minimize} \quad & I[T] := \int_{X_1} c\left(x_1, T_2\left(x_1\right), \ldots, T_N\left(x_1\right)\right) \mathrm{d}\mu_1(x_1) \\ \text{subject to} \quad & T \in \{T : X_1 \to X_2 \times \ldots \times X_N : T \text{ is measurable} \\ & \qquad \text{and fulfills } \underbrace{T_{i\#}\mu_1 = \mu_i}_{\substack{\text{in words: } T_i \\ \text{transports } \mu_1 \\ \text{onto } \mu_i}} \text{ for } i = 2, \ldots, N\}. \end{aligned}$$

Kantorovich's MMOT problem aims to

$$\begin{aligned} \text{Minimize} \quad & C[\gamma] := \int_{X_1 \times \ldots \times X_N} c\left(x_1, \ldots, x_N\right) \mathrm{d}\gamma\left(x_1, \ldots, x_N\right) \\ \text{subject to} \quad & \gamma \in \{\gamma \in \mathcal{P}\left(X_1 \times \ldots \times X_N\right) : \quad \underbrace{M_{X_i}\gamma = \mu_i \text{ for } i = 1, \ldots, N}_{\text{in words: } \gamma \text{ couples } \mu_1, \mu_2, \ldots, \mu_N}\}. \end{aligned}$$

The shift from two to multiple marginals preserves key characteristics of Kantorovich's OT formulation: linear objective function, linear constraints. Again, dual problems arise. They take on the form

$$\begin{aligned} \text{Maximize} \quad & \sum_{i=1}^{N} \int_{X_i} \varphi_i\left(x_i\right) \mathrm{d}\mu_i\left(x_i\right) \\ \text{subject to} \quad & \varphi_1 \oplus \varphi_2 \oplus \ldots \oplus \varphi_N \leq c. \end{aligned}$$

For an account of strong duality (that is equality of primal and dual optimal value) as well as primal and dual attainment we refer the reader to Kellerer [58].

The shift from two to multiple marginals loosens our grip on Monge solutions (which refer to optimal transport plans that are induced by transport maps). What prompts us to lose our grip? Why is Monge's ansatz 'less sufficient' than in the two-marginal setting? We do not answer these questions fully but rather give a formal indication, a pointer if you will. Take a look back at our nonrigorous derivation of Theorem 2.2 and try to generalize it to the multi-marginal setting: formulate complementary slackness for MMOT and use the minimizer-property of x_1, specified above, to derive

$$\nabla_{x_1} c\left(x_1, x_2, \ldots, x_N\right) = \nabla \varphi_1\left(x_1\right).$$

For $x_1, \ldots, x_N \in \mathbb{R}^d$, we deal with a system of d individual equations that we want to solve for $(N-1)d$ variables. In the two-marginal case, we have to solve d

equations for d variables. It is much harder to express $(x_2, \dots, x_N)$ as a function of x_1 than it is to express x_2 as a function of x_1.
The literature pertaining the existence (and uniqueness) of Monge solutions in MMOT consists of individual puzzle pieces each of which discusses a specific cost function or a class of cost functions. In a review on MMOT [72], Pass provides *"a more unified view of what is known"* [72, p. 1772]. This includes a sufficient condition on cost functions that locally guarantees a d-dimensional support set of optimal transport plans as well as a sufficient condition on cost functions (for an absolutely continuous first marginal μ_1) that (globally) guarantees the existence and uniqueness of a Monge solution. Even though some patterns have emerged, we do not fully understand the whole picture yet. In Section 4.2, we shine a light on the insufficiency of Monge's ansatz in the multi-marginal setting, the focus being on (symmetric) MMOT on finite state spaces. We will explain why Core Publication A.1 works towards a better understanding of the insufficiency of Monge's ansatz in said setting.

MMOT problems appear, for example, in fluid dynamics, economics, electronic structure, and data science. In fluid dynamics, MMOT arises through time discretization from a minimal-action problem for incompressible fluids; see Brenier [18] and Benamou, Carlier, and Nenna [10], respectively, for a high- and low-level account. The marginal constraints enforce incompressibility after each time step $0, 1/(N-1), 2/(N-1), \dots, 1$. In economics, MMOT arises in the effort to configure teams; see Carlier and Ekeland [20]. Consider for example the housing market. For construction and sale of a house, a team consisting of a plumber, carpenter, painter, et cetera, and a buyer need to be matched. The marginal constraints incorporate the distributions of the individual team members; μ_1 would be the distribution of plumbers, μ_2 the distribution of carpenters, μ_3 the distribution of painters, $\dots$, and μ_N the distribution of buyers. In electronic structure, MMOT arises as the strong-interaction limit of the Hohenberg-Kohn functional; two working groups - Cotar, Friesecke, Klüppelberg [25] and Buttazzo, De Pascale, Gori-Giorgi [19] - independently initiated the use of MMOT in this context. Simply put, the goal is to find the spatial distribution of an N-electron-molecule that minimizes the energy in the system. The marginal constraints incorporate the spatial distributions of the individual electrons. In data science, MMOT arises in the effort to compute the barycenter of datasets; see Agueh and Carlier [1]. The marginal constraints incorporate the individual datasets in the database. A trend becomes apparent: throughout these MMOT applications, the concept of physical transport is set aside. The transport plan γ couples incompressibility constraints, distributions of individual team members, spatial distributions of individual electrons, and datasets in a database instead of source and sink distribution. Accordingly, we use

the term 'coupling' interchangeably with the term 'transport plan' for γ. In the next chapter, we go into the details of the motivating application in electronic structure.

3 Motivation

In this chapter, we introduce the reader to the curse of dimension tackled in this dissertation. After the brief optimal transport (OT) overview in the previous chapter and before going into the details of the contributed articles, we present the key application of multi-marginal optimal transport (MMOT): modelling the spatial distribution in many-particle systems. We derive the OT problems arising in this application, first in their continuous and then their discrete form. Finally, we explain the curse of dimension looming over the discrete problems.

In Section 3.1, we explain how OT arises in many-electron physics in an easy-to-digest manner: we rid the path from physics to OT of distractions such as the underlying function spaces and some explicit formulas for the separate energy functionals. The reader can find these details in the references [25, 72, 26] on which we also based our presentation. Before going into the details of this presentation, let us refer the reader to a recent excellent review on the topic titled "The strong-interaction limit of density functional theory" by Friesecke, Gerolin, and Gori-Giorgi [42]. It provides an extensive overview in comprehensive writing.

3.1 Electronic structure in optimal transport

A system of N electrons forms a so-called ground state if it minimizes the energy it contains. This physics-snippet might resonate with some of the readers that took advanced physics courses in their school days and it already points to the motivating question of this section: How can we find these ground states?

We start off by capturing the ground state of an N-electron-system in mathematical language. Wave functions ψ represent the possible states of the N-electron-system. They are functions of the positions $(x_1, x_2, \ldots, x_N)$ and spins $(s_1, s_2, \ldots, s_N)$ of the N electrons. They are complex-valued but - when taken to the absolute value squared - they have a probabilistic interpretation: for any wave function ψ, $|\psi(x_1, s_1, x_2, s_2, \ldots, x_N, s_N)|^2$ captures - in the role of a density function - the probability that the N electrons are located at the positions $(x_1, x_2, \ldots, x_N)$ with spins $(s_1, s_2, \ldots, s_N)$. The density $|\psi|^2$ integrates to 1. In this probabilistic sense the wave functions describe the state of the N-electron-system. Any such state implies a certain amount of energy. We capture the relationship

between wave function and energy in the energy functional E with

$$E[\psi] = T[\psi] + V_{\text{ne}}[\psi] + V_{\text{ee}}[\psi].$$

Here, T represents the kinetic energy, V_{ne} the nuclei-electron interaction energy, and V_{ee} the electron-electron interaction energy. We can realize our goal of determining the ground state by minimizing the energy $E[\psi]$ among all valid wave functions and extracting a minimizer. This minimizer corresponds to a ground state.

When attempting to find the ground state computationally, we face admissible objects that become increasingly intractable with an increasing number of electrons. Wave functions map N-tuples of space-spin coordinates onto the complex plane, i.e.,

$$(x_1, s_1, x_2, s_2, \ldots, x_N, s_N) \mapsto \psi\left(x_1, s_1, x_2, s_2, \ldots, x_N, s_N\right) \in \mathbb{C}.$$

Let us ignore, for a second, the spin and assume that a cube in $\mathbb{R}^3$ contains all possible positions of the electrons. In order to minimize the energy functional E computationally, we need to store wave functions. Let's say we discretize the cube by 1000 gridpoints (10 gridpoints for each space dimension). Then the system of the N electrons can form 1000^N configurations: each electron can choose from 1000 different positions. A wave function then assigns a complex value to each of the 1000^N configurations. Already for moderately sized N, storing these 1000^N complex numbers - and by that the wave function - becomes infeasible. We simply cannot handle wave functions as admissible objects because of their storage complexity.

Do we really need all the information contained in a wave function ψ to compute its energy $E[\psi]$? Take the energy component V_{ne}, for example. How does it depend on ψ? In order to answer this question, we introduce a 'projection' of wave functions to functions of one variable. A single particle density ρ arises from a wave function ψ in the following manner: take the absolute value squared of ψ, integrate out first the spins

$$\rho_N\left(x_1, x_2, \ldots, x_N\right) = \sum_{\substack{\text{all possible} \\ \text{spin combinations} \\ (s_1, s_2, \ldots, s_N)}} \left|\psi\left(x_1, s_1, x_2, s_2, \ldots, x_N, s_N\right)\right|^2$$

and then all but the position of electron 1

$$\rho\left(x_1\right) = N \int_{\mathbb{R}^{3(N-1)}} \rho_N\left(x_1, x_2, \ldots, x_N\right) \mathrm{d}x_2 \ldots \mathrm{d}x_N.$$

We inserted the normalization factor N to comply with a convention in physics: the single particle density should integrate to the number of electrons. In an intermediate step - by integrating out the spins - we produce the density ρ_N. It captures the probability that the N electrons are located at the positions $(x_1, x_2, \ldots, x_N)$ in its role as a density function. As an immediate consequence of an anti-symmetry constraint on wave functions, $|\psi|^2$ and therefore ρ_N are symmetric, that is invariant under argument-permutation:

$$\rho_N(x_1, x_2, \ldots, x_N) = \rho_N\left(x_{\sigma(1)}, x_{\sigma(2)}, \ldots, x_{\sigma(N)}\right) \text{ for all permutations } \sigma \text{ of } \{1, 2, \ldots, N\}.$$

The symmetry allows us to see the N electrons as indistinguishable (the probability of a configuration does only depend on the N chosen positions not on how the numbered electrons are assigned to these positions); the symmetry implies that the single particle density ρ will not change if we give the role of electron 1 to any of the other electrons. It does not matter which $(N-1)$ of the electron-positions are integrated out. Hence, $\int_A \rho$ for any $A \subset \mathbb{R}^3$ measures the probability that electron j is in A, independent of the choice of j. V_{ne} quantifies the interaction energy between nuclei and electrons; it does not depend on what specific configurations the N electrons form, it only depends on the position of the electron as an individual; so $V_{\text{ne}}[\psi]$ depends on ψ only through its single particle density ρ.

The reduced ψ-dependence of V_{ne} motivates a partition of the original energy-minimization-problem $\inf_\psi E[\psi]$ into a double minimization:

$$\inf_\psi E[\psi] = \inf_\rho \left\{V_{\text{ne}}[\rho] + F_{\text{HK}}[\rho]\right\} \tag{3.1}$$

with

$$F_{\text{HK}}[\rho] := \inf_{\substack{\psi \\ \text{subject to } \psi \mapsto \rho}} \left\{T[\psi] + V_{\text{ee}}[\psi]\right\}. \tag{3.2}$$

Instead of minimizing over all wave functions ψ (left-hand side of (3.1)), we minimize over all wave functions ψ that have a specified single particle density ρ and then over all single particle densities (right-hand side of (3.1)). This partition goes back to the work of Lieb(1983) [64], Levy(1979) [63], as well as Hohenberg and Kohn(1964) [52]. Simply put, one can attribute (3.1) to Hohenberg and Kohn, and the direct variational understanding (3.2) of F_{HK} to Levy and Lieb. With their direct variational understanding Levy and Lieb replaced an earlier, more indirect, and non-variational one by Hohenberg and Kohn.
A few remarks regarding the notation: $\inf_\psi E[\psi]$ denotes the original task to minimize $E[\psi]$ among all valid wave functions; we chose the inf-notation to cover the case when the minimal value is not attained. We write $\psi \mapsto \rho$ if ψ has single

particle density ρ. F_{HK} is called Hohenberg-Kohn functional.

So far the problem has remained intractable: on the right-hand side of (3.1) the ψ-dependence is hidden yet still present. The first instance of ψ-dependence was cleared by Lieb himself early on. So far we understand the admissible objects regarding the outer minimization only as single particle densities that arise from wave functions. This description still relies on the high-dimensional objects - the wave functions; the single particle densities are only implicitly understood as projection of them. In his contribution [64], Lieb provided an explicit description of admissible ρs that no longer relied on wave functions. The second instance of ψ-dependence creates a challenge that we face to this day. Even though at a glance everything on the right-hand side of (3.1) seems to be nicely low-dimensional now, wave functions are still present as admissible objects in the inner minimization making up the Hohenberg-Kohn functional (3.2). At this point we knock on the door of a very active field of research in physics and chemistry: density functional theory. At its heart lies the Hohenberg-Kohn functional and the goal to make the double minimization work. To that end researchers have traded in exactness for low-dimensionality: they are mostly concerned with approximating the Hohenberg-Kohn functional as good as possible while keeping the ψ-dependence low-dimensional. We refer readers that are interested in entering this field to [70].

How do we approximate a system consisting of N interacting electrons using lower-dimensional sub-systems? The most naive ansatz goes as follows. Let us assume, for a second, that the N electrons position themselves independently of each other. Then the probability that they form the configuration $(x_1, x_2, \ldots, x_N)$ disintegrates into a product of individual position probabilities, i.e.,

$$\rho_N(x_1, x_2, \ldots, x_N) = \frac{1}{N^N} \rho(x_1) \cdot \rho(x_2) \cdot \ldots \cdot \rho(x_N).$$

The factor $1/N^N$ corrects the normalization discrepancy: ρ_N integrates to 1 while ρ integrates to N. The independence assumption allows us to replace an N-variable-function by a product of (N) one-variable-functions. In the discretized example from above, we replace 1000^N input arguments by 1000 (note that the one-variable-functions are identical). We disrupted the exponential growth by treating the N-electron-system as if it consisted of N independent one-electron-systems. The independence ansatz is one of the bedrocks, approximation schemes in density functional theory are built upon. Further building blocks include corrections that account for interactions between the electrons.

There is an argument to be made for going down a different - if not the opposite - path, one that starts off with a strong interaction assumption [78, 80]. It is

common knowledge that electrons repel each other. The resulting energy influences the ground state through the component V_{ee} - the electron-electron interaction energy. It introduces a certain amount of correlation between the electrons. In order to understand these correlations we follow a common approach and put an emphasis on the interaction: we 'assume' strong interaction. The gained insights are relevant and interesting in their own right but also complement weak-interaction approaches like the independence ansatz.

To incorporate the strong interaction assumption mathematically, we introduce a parameter $\alpha > 0$ into the Hohenberg-Kohn functional

$$F^{\alpha}_{\text{HK}}[\rho] := \inf_{\substack{\psi \\ \text{subject to } \psi \mapsto \rho}} \{\alpha T[\psi] + V_{\text{ee}}[\psi]\}$$

and consider the limit $\alpha \to 0$. By decreasing α, we increase the influence of the electron-electron interaction energy and thereby the prominence of the interaction. In this strong-interaction limit, the minimization problem making up the parameter-dependent Hohenberg-Kohn functional F^{α}_{HK} takes on the form of an OT problem. (For the due references, see below.)

Theorem 3.1 (OT-form of strong-interaction limit of Hohenberg-Kohn functional)**.** *For any single particle density ρ*

$$\lim_{\alpha \to 0} F^{\alpha}_{HK}[\rho] = \inf_{\substack{\gamma \in \mathcal{P}_{sym}\left(\mathbb{R}^{3N}\right) \\ \gamma \mapsto \rho/N}} \int_{\mathbb{R}^{3N}} \sum_{1 \leq i < j \leq N} \frac{1}{|x_i - x_j|} \mathrm{d}\gamma\left(x_1, x_2 \ldots, x_N\right). \tag{3.3}$$

The attentive reader might notice some discrepancies with the general MMOT problem formulation from Section 2.2. $\mathcal{P}_{\text{sym}}\left(\mathbb{R}^{3N}\right)$ denotes the set of symmetric probability measures on $\mathbb{R}^{3N}$. Symmetry refers to an invariance under argument-permutation, i.e.,

$$\begin{aligned} \gamma \text{ is symmetric} \longleftrightarrow \gamma\left(A_1 \times A_2 \times \ldots \times A_N\right) &= \gamma\left(A_{\sigma(1)} \times A_{\sigma(2)} \times \ldots \times A_{\sigma(N)}\right) \\ &\text{for all measurable } A_1, A_2, \ldots, A_N \subset \mathbb{R}^3 \text{ and} \\ &\text{all permutations } \sigma \text{ of } \{1, 2, \ldots, N\}. \end{aligned}$$

Note that symmetric probability measures cannot have different one-point marginals: for any two coordinates $i, j \in \{1, 2, \ldots, N\}$ the respective one-point marginals $M_{X_i}\gamma$ and $M_{X_j}\gamma$ coincide. For $\gamma \in \mathcal{P}_{\text{sym}}\left(\mathbb{R}^{3N}\right)$ and $\mu \in \mathcal{P}\left(\mathbb{R}^3\right)$, we write $\gamma \mapsto \mu$ if one and thereby all one-point marginals of γ are equal to μ. Here we slightly abuse this notation. The single particle density ρ itself is not a probability measure on $\mathbb{R}^3$, however, it induces a probability measure μ_ρ on $\mathbb{R}^3$ in its role as a density function:

$$\mu_\rho(A) = \int_A \frac{\rho(x)}{N} \mathrm{d}x \quad \text{for } A \subset \mathbb{R}^3 \text{ measurable.}$$

The constraint $\gamma \mapsto \rho/N$ then means $\gamma \mapsto \mu_\rho$. This abuse of notation - writing ρ/N instead of μ_ρ - will follow us throughout the rest of this section.

How can we interpret the OT problem in (3.3)? The reader might feel like the cost function $c(x_1, x_2, \ldots, x_N) = \sum_{1\leq i<j\leq N} \frac{1}{|x_i - x_j|}$ dropped from the sky, even though it was implicitly present in the electron-electron interaction energy V_{ee} throughout this section:

$$V_{\text{ee}}[\psi] = \int_{\mathbb{R}^{3N}} \sum_{\substack{\text{all possible} \\ \text{spin combinations} \\ (s_1,s_2,\ldots,s_N)}} \sum_{1\leq i<j\leq N} \frac{1}{|x_i - x_j|} \\ |\psi(x_1, s_1, x_2, s_2, \ldots, x_N, s_N)|^2 \mathrm{d}x_1 \mathrm{d}x_2 \ldots \mathrm{d}x_N.$$

Compare the formula for V_{ee} and the OT problem in (3.3), and you will see that $\mathrm{d}\gamma(x_1, x_2, \ldots, x_N)$ takes on the role of $\sum |\psi(x_1, s_1, x_2, s_2, \ldots, x_N, s_N)|^2 \mathrm{d}x_1 \mathrm{d}x_2 \ldots \mathrm{d}x_N$. Recall the definition of ρ_N,

$$\rho_N(x_1, x_2, \ldots, x_N) := \sum |\psi(x_1, s_1, x_2, s_2, \ldots, x_N, s_N)|^2,$$

as well as its interpretation as position density, and you will understand that γ embodies the spatial distribution of the N-electron-system: $\gamma(A_1 \times A_2 \times \ldots \times A_N)$ (for $A_1, A_2, \ldots, A_N \subset \mathbb{R}^3$ measurable) gives the probability that the N-electron-system forms a configuration in $A_1 \times A_2 \times \ldots \times A_N$. The marginal condition ensures that each individual electron is distributed according to ρ/N. Symmetric probability measures $\gamma \in \mathcal{P}_{\text{sym}}(\mathbb{R}^{3N})$ can be the independent measure but also model intricate correlations. The connection between independent and symmetric measures is discussed in Chapter 6. Overall, we seek a spatial distribution γ of the N electrons that respects the single particle density ρ and minimizes the interaction energy between the electrons.

There is a line of references that worked towards a general statement on the OT-form of the Hohenberg-Kohn functional in the strong-interaction limit. In [25, 19], the two working groups (Cotar, Friesecke, Klüppelberg and Buttazzo, De Pascale, Gori-Giorgi) independently introduced the interpretation/formulation of the limit problem as OT problem. Theorem 3.1, in the presented general N form, is owed to Cotar, Friesecke, and Klüppelberg [26]. Their work succeeds previous results for $N = 2$ [25] and $N = 3$ [13]. Allowing physics literature, we may draw this line even further back to include Seidl [78] as well as Seidl, Gori-Giorgi, and Savin [79]. Therein the limit problem was already indicated as well as the lack of a rigorous proof. Together with [80] these authors also coined the term 'strictly correlated electrons' which describes a predicted state of the N-electron-system

in the strong-interaction limit. In this state the position of the first electron - x_1 - uniquely determines the positions - $x_2, x_3, \ldots, x_N$ - of the remaining $N-1$ electrons. As a result we can write the positions of the remaining electrons as functions of the position of the first electron

$$x_2 = T_2(x_1) \qquad x_3 = T_3(x_1) \qquad \ldots \qquad x_N = T_N(x_1).$$

Of course the numbering of the electrons here is arbitrary as they are indistinguishable. Such 'strictly correlated electron'-states yield a sparse structure of the transport plan γ; it would be of Monge-form with $T_2, T_3, \ldots, T_N$ being the corresponding transport maps:

$$\gamma = (id, T_2, T_3, \ldots, T_N)_{\#} \frac{\rho}{N}.$$

So far the advantage of the independence ansatz over the strong-interaction limit was a dramatic reduction of the storage complexity. In the strong-interaction limit, we still had to deal with admissible objects of exponential-in-N-size, the γs. Having a sparse structure of the γs guaranteed would level the playing field. Investigating whether Monge's ansatz (or other sparse approaches) are sufficient is a 'must' if we want to perform error-controlled strong-interaction computations.

3.2 Discretization

To perform computations regarding the continuous MMOT problems

$$\begin{aligned} \text{Minimize} \quad & \int_{\mathbb{R}^{3N}} \sum_{1 \le i < j \le N} \frac{1}{|x_i - x_j|} \mathrm{d}\gamma(x_1, x_2 \ldots, x_N) \\ \text{subject to} \quad & \gamma \in \{\gamma \in \mathcal{P}_{\text{sym}}\left(\mathbb{R}^{3N}\right) : \quad \gamma \mapsto \rho/N\}, \end{aligned}$$

which arise in the application of MMOT to electronic structure problems, we chose a discretization that preserves the OT-structure. We landed on the following discrete MMOT problems.

$$\text{Minimize} \quad \int_{X^N} c(x_1, \ldots, x_N)\, \mathrm{d}\gamma(x_1, \ldots, x_N) \tag{3.4}$$

$$\text{subject to} \quad \gamma \in \{\gamma \in \mathcal{P}_{\text{sym}}\left(X^N\right) : \quad \gamma \mapsto \lambda^*\}. \tag{3.5}$$

We discretized $\mathbb{R}^3$ or more generally $\mathbb{R}^d$ by a finite selection

$$X = \{a_1, a_2, \ldots, a_\ell\} \subset \mathbb{R}^d$$

of pairwise distinct points and replaced the Coulomb cost $\sum \frac{1}{|x_i - x_j|}$ by a general symmetric cost function $c : X^N \to \mathbb{R}$. A cost function $c : X^N \to \mathbb{R}$ is called symmetric if

$$c(x_1, x_2, \ldots, x_N) = c\left(x_{\sigma(1)}, x_{\sigma(2)}, \ldots, x_{\sigma(N)}\right) \quad \text{for all } x_1, x_2, \ldots, x_N \in X^N \text{ and all permutations } \sigma \text{ of } \{1, 2, \ldots, N\}.$$

We denote by $\mathcal{P}_{\text{sym}}\left(X^N\right)$ the set of symmetric probability measures on X^N where a measure $\gamma \in \mathcal{P}\left(X^N\right)$ is called symmetric if

$$\gamma(A_1 \times \ldots \times A_N) = \gamma\left(A_{\sigma(1)} \times \ldots \times A_{\sigma(N)}\right) \quad \text{for all } A_1, \ldots, A_N \subset X \text{ and all permutations } \sigma \text{ of } \{1, \ldots, N\}.$$

The symmetry of the cost functions allows us to hold on to the symmetry of the admissible objects: let us consider, for a second, the 'unsymmetrized' MMOT problem we create by replacing $\mathcal{P}_{\text{sym}}\left(X^N\right)$ with $\mathcal{P}\left(X^N\right)$; for these problems, symmetric cost functions c always give rise to a symmetric optimizer $\gamma \in \mathcal{P}_{\text{sym}}\left(X^N\right)$. For $\gamma \in \mathcal{P}_{\text{sym}}\left(X^N\right)$ and $\lambda^* \in \mathcal{P}(X)$, we write $\gamma \mapsto \lambda^*$ if one and thereby all one-point marginals of γ are equal to λ^*. Hereby, we compute the one-point marginal of γ with respect to the first coordinate as follows.

$$M_X \gamma(A_1) = \gamma\left(A_1 \times X^{N-1}\right) \quad \text{for all } A_1 \subset X.$$

Of course, we tailor the choice of discretization points X and their weights λ^* to approximate their continuous counterparts. For an example of how to do this, we refer the reader to Section 2 in Core Publication A.2. Therein, we also justify the chosen discretization by means of a convergence result.

We can express the discrete MMOT problem (3.4)-(3.5) in formalism adherent to discrete optimization. Identify an admissible object $\gamma \in \mathcal{P}_{\text{sym}}\left(X^N\right)$ and the marginal constraint $\lambda^* \in \mathcal{P}(X)$ with a tensor of order N, $\left(\gamma_{i_1 i_2 \cdots i_N}\right)_{i_1, i_2, \ldots, i_N = 1}^{\ell} \in \mathbb{R}^{\ell^N}$, and a vector, $\left(\lambda_i^*\right)_{i=1}^{\ell} \in \mathbb{R}^{\ell}$, via

$$\gamma_{i_1 i_2 \cdots i_N} := \gamma\left(\{(a_{i_1}, a_{i_2}, \ldots, a_{i_N})\}\right) \qquad \lambda_i^* := \lambda^*\left(\{a_i\}\right).$$

Then, we can write the discrete MMOT problem (3.4)-(3.5) as follows.

$$\text{Minimize} \quad \sum_{i_1,i_2,\dots,i_N=1}^{\ell} c\,(a_{i_1}, a_{i_2}, \dots, a_{i_N})\,\gamma_{i_1 i_2 \cdots i_N} \tag{3.6}$$

$$\text{subject to} \quad \gamma \in \{\gamma \in \mathbb{R}^{\ell^N} : \quad \gamma \text{ is symmetric}, \tag{3.7}$$

$$\gamma \geq 0, \tag{3.8}$$

$$\sum_{i_2,\dots,i_N=1}^{\ell} \gamma_{j i_2 \cdots i_N} = \lambda_j^* \text{ for } j \in \{1,\dots,\ell\}\}. \tag{3.9}$$

The nonnegativity constraint '$\gamma \geq 0$' means that every entry $\gamma_{i_1 i_2 \cdots i_N}$ is nonnegativ.

Monge's ansatz is best discussed in the context of a uniform marginal constraint

$$\overline{\lambda} := \sum_{i=1}^{\ell} \frac{1}{\ell} \delta_{a_i}.$$

In Section 2.1 we saw that maps can transport Dirac measures only onto other Dirac measures. The present problems enforce finite sums of Dirac measures as marginals and require every marginal to be the same. For transport maps (other than the identity) to exist, we need the weights that scale the individual Dirac measures to be quantized. This requirement makes the uniform measure $\overline{\lambda}$ the clear choice of marginal. The relevance of the uniform marginal goes beyond the theoretical aspect that it gives rise to transport maps. It approximates absolutely continuous probability measures through equi-mass discretization (see [22] for further reading).

3.3 The curse of dimension

The storage complexity of the discrete MMOT problems (3.4)- (3.5)/ (3.6)-(3.9) grows exponentially with the number of marginals N. The natural number N regulates the dimension of the problem: for $N = 2$ the admissible objects are $\ell \times \ell$-matrices (tensors of order 2), for $N = 3$ they are $\ell \times \ell \times \ell$-'cubes' (tensors of order 3), and for a general $N > 1$ they are tensors of order N. These tensors consist of ℓ^N individual entries. As any admissible object is symmetric, we can partition these ℓ^N entries into 'symmetry classes'. For any symmetric tensor γ, all the entries in the same class carry the same value; for example

$$\gamma_{12\cdots 2} = \gamma_{21\cdots 2} = \cdots = \gamma_{22\cdots 1}.$$

By only storing the value carried within a class once, we can reduce the number of unknowns from ℓ^N to $\binom{N+\ell-1}{N}$ [44]. How bad is $\binom{N+\ell-1}{N}$? We answer this question with the help of an example on electronic structure. We want to distribute N electrons over $\ell = 5 \cdot N$ sites. Then $\binom{N+\ell-1}{N}$ still grows exponentially with the number of marginals N.

$$\begin{aligned}\binom{N+\ell-1}{N} = \binom{N+5N-1}{N} &= \binom{6N-1}{N} \\ &= \frac{(6N-1)\cdot(6N-2)\cdot\ldots\cdot(5N+1)\cdot 5N}{N\cdot(N-1)\cdot\ldots\cdot 2\cdot 1} \\ &\geq 5^N.\end{aligned}$$

As long as $\ell/N > 1$ holds, this exponential growth persists - even if the number of sites ℓ is not a multiple of the number of marginals N. The reader can check this by generalizing the above computations. The curse of dimension consists in the exponential growth of the number of unknowns with the number of marginals.

To break the curse, we seek an ansatz for the discrete MMOT problems (3.4)-(3.5)/ (3.6)-(3.9) that

(i) is low-dimensional,

(ii) is sufficient in the sense that any instance of the discrete MMOT problems admits a solution that takes on the form prescribed by the ansatz, and

(iii) gives rise to an efficient computational method to solve the discrete MMOT problems.

4 Convex geometry of transport polytopes

A transport polytope encloses the set of probability measures that are admissible to an optimal transport (OT) problem. These OT problems can be either of Kantorovich- or Monge-type. While the word 'transport' describes what the admissible objects represent, the word 'polytope' describes what set the admissible objects form. In the first part of this chapter we recall basic notions of convexity - such as the term 'polytope' - and briefly explain how to test whether a given ansatz space is sufficient. In the second part we discuss the insufficiency of Monge's ansatz in multi-marginal optimal transport (MMOT) and draw a comparison with the two-marginal setting. In the third and final part we introduce the reader to a sufficient ansatz space for MMOT problems - the Quasi-Monge ansatz.

In Section 4.1 we focus on the notions of convexity relevant to this dissertation; for a textbook account of these notions and further reading material on convexity, we refer the reader to [53, 76].

4.1 Notions of convexity

A set $K \subset \mathbb{R}^d$ is convex if for two of its elements $x, y \in K$ the line segment $\{(1-t)x + ty : 0 \leq t \leq 1\}$ lies in K, i.e.,

$$x, y \in K \text{ and } t \in [0,1] \quad \rightarrow \quad (1-t)x + ty \in K.$$

The line segment connecting x and y consists of convex combinations of the endpoints. A convex combination of $x_1, x_2, \ldots, x_m \in \mathbb{R}^d$ is a linear combination

$$t_1 x_1 + t_2 x_2 + \ldots + t_m x_m$$

that satisfies the convexity constraint

$$t_1, t_2, \ldots, t_m \geq 0 \quad \text{and} \quad \sum_{i=1}^{m} t_i = 1.$$

As is easily seen using induction, a convex set $K \subset \mathbb{R}$ is closed under taking convex combinations of its elements:

$$x_1, x_2, \ldots, x_m \in K,\ t_1, t_2, \ldots, t_m \geq 0 \text{ and } \sum_{i=1}^{m} t_i = 1 \quad \rightarrow \quad \sum_{i=1}^{m} t_i x_i \in K.$$

For an arbitrary set $V \subset \mathbb{R}^d$, we refer to the smallest convex set enclosing V as the convex hull of V and denote it as $conv(V)$. To produce the convex hull of a set $V \subset \mathbb{R}^d$, take the set V itself and include all convex combinations of its elements. A set $V \subset \mathbb{R}^d$ is convex if and only if it equals its convex hull.

The line segment $\{(1-t)x+ty : 0 \leq t \leq 1\}$ connecting x with y is itself convex and therefore contains all convex combinations of its elements. As is apparent from its definition, it exactly consists of the convex combinations of its endpoints x and y. This description shaves off all convex combinations of non-endpoints. Does a similar minimal description exist for general convex sets? What is the role of x and y? An element e of a convex set K is an extreme point of K if

$$x, y \in K,\ t \in (0,1) \text{ and } (1-t)x + ty = e \quad \text{implies} \quad x = y = e.$$

By Minkowski's theorem, any compact convex set in $\mathbb{R}^d$ equals the convex hull of its extreme points.

The discrete MMOT problems (3.4)-(3.5)/(3.6)-(3.9) always admit extremal solutions (which are extreme points of the set of admissible objects that solve the problem). By Minkowski's theorem, the set of admissible objects equals the convex hull of its extreme points. Given an instance of the considered MMOT problems, take an arbitrary optimal transport plan (these always exist as we minimize a continuous objective function on a compact set in $\mathbb{R}^{\ell^N}$); write it as convex combination of extreme points; now any extreme point e_i with positive weight factor $t_i > 0$ yields the same objective value as the original solution and therefore is an extremal solution. We can use the existence of extremal solutions to give a simple criterion that indicates whether a given ansatz is sufficient: if all extreme points took on the prescribed form, the considered ansatz would be sufficient.

Conversely, a single incompatible extreme point yields the insufficiency of an ansatz. Let us fix the set of admissible objects. Then the choice of cost tensor $(c\,(a_{i_1}, a_{i_2}, \ldots, a_{i_N}))_{i_1,i_2,\ldots,i_N=1}^{\ell}$ completely determines an instance of the discrete MMOT problems (3.4)-(3.5)/(3.6)-(3.9). An ansatz is insufficient if a problem instance - which means a cost - exists that fails to admit a solution of the prescribed form. An element e of a convex set $K \subset \mathbb{R}^d$ is an exposed point of K if there exists a linear function $b \in \mathbb{R}^d$ that attains its minimum on K exactly at e, i.e.,

$$b^T x > b^T e \quad \text{for all } x \in K \setminus \{e\}.$$

A single incompatible extreme point only renders an ansatz insufficient if it is also an exposed point. In the present setting all extreme points are exposed points. Why?

Our set of admissible objects consists of all nonnegative tensors of order N that satisfy a symmetry and marginal constraint. We deal with a set in $\mathbb{R}^{\ell^N}$ that arises as the intersection of finitely many closed halfspaces. Rockafellar calls such sets polyhedral and shows that any bounded such set is the convex hull of finitely many points - a polytope. As our set of admissible objects is bounded, it forms a polytope. Hinting at the nature of the considered problems, we refer to these polytopes as transport polytopes.
By definition, every exposed point of these polytopes is an extreme point. By Straszewicz's theorem (see, for example, [76, p.167]) any extreme point of a closed convex set can be approximated arbitrarily well with exposed points. As the convex hull of finitely many points, the transport polytopes only have finitely many extreme points each of which - by Straszewicz's theorem - must be an exposed point.
Finally, let us summarize the sufficiency criterion: a chosen ansatz will be sufficient if and only if all extreme points take on the prescribed form.

4.2 Insufficiency of Monge's ansatz in multi-marginal optimal transport

Monge's ansatz goes back to the origin of OT. As the OT-equivalent of the physical term 'strictly correlated electrons', it has a foundation in physics. Last but not least it is sparse. For N electrons on ℓ discretization points, general transport plans assign probabilities to ℓ^N electron configurations. In Monge's ansatz the electrons are strictly correlated: the position of the first electron uniquely determines the positions of the remaining $(N-1)$ electrons. We can describe the strict correlation using maps $T_2, T_3, \dots, T_N$: if x_1 is the position of the first electron $T_2(x_1), T_3(x_1), \dots, T_N(x_1)$ are the position of the remaining electrons. Each map consists of ℓ values and we deal with $(N-1)$ maps. The result is a storage complexity of $\ell \cdot (N-1)$; it grows only linear in N. Monge's ansatz significantly disrupts the exponential growth of the storage complexity, but is the ansatz sufficient?

Let us concentrate on the MMOT problems that lie at the heart of this dissertation: symmetric MMOT problems on finite state spaces. See (3.4)-(3.5)/(3.6)-(3.9) for an explicit account of these problems. An admissible object γ is of Monge-type if there exist transport maps $T_2, T_3, \dots, T_N$ such that

$$\gamma = S_N \left((id, T_2, T_3, \dots, T_N)_{\#} \lambda^* \right). \tag{4.1}$$

Recall that $T : X \to X$ is a transport map if it pushes λ^* onto itself, i.e., $T_{\#}\lambda^* = \lambda^*$. In (4.1), we adapted the 'standard' Monge approach to the symmetric setting

using the symmetrization operator in N variables S_N. It maps $\mathcal{P}\left(X^N\right)$ onto $\mathcal{P}_{\text{sym}}\left(X^N\right)$: given any probability measure γ on X^N, its symmetrization $S_N\gamma$ is defined by

$$S_N\gamma\left(A_1 \times \ldots \times A_N\right) = \frac{1}{N!}\sum_{\sigma}\gamma\left(A_{\sigma(1)} \times \ldots \times A_{\sigma(N)}\right) \text{ for all } A_1, \ldots, A_N \subset X.$$

The sum runs over all permutations σ of $\{1, 2, \ldots, N\}$. Applying the symmetrization operator to the originally sparse state $(id, T_2, T_3, \ldots, T_N)_{\#}\lambda^*$, distributes its mass among more configurations in X^N. The state becomes 'less sparse'. Nevertheless, the storage complexity of the symmetrized Monge states (4.1) remains unchanged; the $(N-1)$ transport maps still contain all the necessary information. For finite state spaces, Monge's ansatz is best discussed in the context of a uniform marginal constraint

$$\overline{\lambda} := \sum_{i=1}^{\ell}\frac{1}{\ell}\delta_{a_i}.$$

Why? For any $\lambda^* \in \mathcal{P}(X)$, we can compute the push-forward of λ^* along $T : X \to X$ via

$$T_{\#}\lambda^* = \sum_{i=1}^{\ell}\lambda^*\left(\{a_i\}\right)\delta_{T(a_i)}.$$

A nonuniform marginal immediately opens up the question whether transport maps - other than the identity - exist. A uniform marginal yields a clearly defined set of transport maps: T transports $\overline{\lambda}$ onto itself if and only if it is a permutation of $X = \{a_1, a_2, \ldots, a_\ell\}$. The corresponding Monge states take on the form

$$S_N\left((id, T_2, T_3, \ldots, T_N)_{\#}\,\overline{\lambda}\right) \quad \text{with } T_2, T_3, \ldots, T_N \text{ being permutations of } X.$$

In the following we will focus our attention on the case of a uniform marginal.

For $N = 2$ marginals and any number of discretization points $\ell \geq 2$, Monge's ansatz is sufficient. In the general (unsymmetrized) case, this sufficiency is an immediate consequence of the well-known Birkhoff-von Neumann theorem [14, 84]. In OT-jargon, it states that any general transport plan is extremal if and only if it is of Monge-type. Applying the symmetrization operator S_2 to the set of general transport plans gives rise to the set of symmetric transport plans - our set of admissible objects. By linearity of S_2, any extreme point in the symmetric case is the symmetrization of an extreme point in the general case. Seeing as Monge states in the symmetric case correspond to symmetrized general Monge states, Monge's ansatz is also sufficient in the symmetric case.

Already for $N = 3$ marginals and $\ell = 3$ discretization points, Monge's ansatz becomes insufficient. The key to sufficiency in the two-marginal setting was the Birkhoff-von Neumann theorem. Can we lift it to a higher-dimensional setting? In its original form, the theorem states that a bistochastic matrix is extremal if and only if it is a permutation matrix. Bistochastic matrices have nonnegative entries and every one of their rows and columns sums to 1. A bistochastic matrix is a permutation matrix if every one of its rows and columns contains a single positive entry. Due to the summation constraints, this entry has to carry the value 1. Consequently, permutation matrices are exactly those bistochastic matrices that only contain integer values. The (in OT relevant) higher-dimensional analogue of bistochastic matrices are certain tensors of order N. Their entries are nonnegative and every one of their coordinate hyperplanes sums to 1. Such a 'N-stochastic tensor' is a 'permutation tensor' if every one of its coordinate hyperplanes contains a single positive entry. Due to the summation constraints, this entry has to carry the value 1. Consequently, 'permutation tensors' are exactly those 'N-stochastic tensors' that only contain integer values. In OT-jargon, bistochastic matrices and 'N-stochastic tensors' correspond to general transport plans, permutation matrices and 'permutation tensors' correspond to Monge states. While for $N = 2$ all extremal bistochastic matrices are integer, for $N \geq 3$ noninteger states appear among the extremal 'N-stochastic tensors'. So while for $N = 2$ Monge's ansatz is sufficient for unsymmetrized OT, for $N = 3$ it is not. The references [27, 62, 65] provide and study noninteger extreme points for the case of $N = 3$ marginals. Do all these extreme points vanish after symmetrization? Sadly, the answer is no. In [41], Friesecke focuses on the case of $N = 3$ marginals on $\ell = 3$ discretization points. He thoroughly investigates the polytope formed by the symmetric transport plans. His explicit list of all 22 extreme points shows that only 7 of them are of Monge-type. Pairwise costs

$$c\left(x_1, x_2, \ldots, x_N\right) = \sum_{1 \leq i < j \leq N} v\left(x_i, x_j\right)$$

expose only 8 of the original 22 extreme points, 3 out of these 8 are not of Monge-type. By putting the set of admissible objects centre stage, Friesecke is able to quantify the insufficiency of Monge's ansatz.

Core Publication A.1 seamlessly continues the efforts of [41] to gain insight into the insufficiency of Monge's ansatz through a proper convex-geometric analysis of the set of admissible objects. For small problem-parameters N and ℓ, we computed all extreme points of the transport polytope formed by our set of admissible objects (3.5) (with $\lambda^* = \overline{\lambda}$) and checked whether they take on the form prescribed by Monge's ansatz. Further, we answered the questions: Which of the extreme points can we expose with pairwise costs? What is the (Monge:non-Monge)-ratio

among them? Let us turn to explicit examples of extreme points. One of the '$N = \ell = 3$-examples' of a non-Monge extreme point given in [41] looks as follows.

$$\frac{1}{2}S_3\delta_{a_1,a_2,a_2} + \frac{1}{2}S_3\delta_{a_1,a_3,a_3}. \tag{4.2}$$

Here, δ_{a_1,a_2,a_2} and δ_{a_1,a_3,a_3} denote $\delta_{a_1} \otimes \delta_{a_2} \otimes \delta_{a_2}$ and $\delta_{a_1} \otimes \delta_{a_3} \otimes \delta_{a_3}$, respectively. A higher-dimensional version for $N = \ell > 3$ looks as follows.

$$\frac{1}{\ell-1}S_N\delta_{a_1,a_2,a_2,\dots,a_2} + \frac{1}{\ell-1}S_N\delta_{a_1,a_3,a_3,\dots,a_3} + \dots + \frac{1}{\ell-1}S_N\delta_{a_1,a_\ell,a_\ell,\dots,a_\ell}. \tag{4.3}$$

This symmetric transport plan is a non-Monge extreme point. Why? It is extreme due to the one-to-one relation between the sites $a_2, a_3, \dots, a_\ell$ and the $\ell - 1$ symmetrized Dirac measures. It is non-Monge due to a $1/\ell$- quantization of Monge states, which will be explained in the upcoming section. Being a Monge state is a global task, local non-Monge structures are enough to disrupt its fulfillment. In Core Publication A.1, we show how to create a non-Monge extreme point for fixed $N = 3$ and arbitrary $\ell > 3$ on the basis of (4.2): take (4.2) rescale it and add measures of the form $\frac{1}{\ell}S_3\delta_{a_i,a_i,a_i}$. We can take the same approach to generate non-Monge extreme points for fixed $N > 3$ and arbitrary $\ell' > \ell = N$ on the basis of (4.3). Overall, we face non-Monge extreme points for any $3 \leq N \leq \ell$.

For the sake of completeness, we include here one more example of a non-Monge extreme point that qualifies for any $\ell \geq 2$ and any $N \geq 3$:

$$\frac{N}{(N-1)\ell}S_N\delta_{a_1,a_1,\dots,a_1,a_2} + \frac{N-2}{(N-1)\ell}S_N\delta_{a_2,a_2,\dots,a_2} + \frac{1}{\ell}S_N\delta_{a_3,a_3,\dots,a_3} + \dots + \frac{1}{\ell}S_N\delta_{a_\ell,a_\ell,\dots,a_\ell}.$$

To see that it is indeed a non-Monge extreme point, lean on the arguments regarding (4.3). We will see in the upcoming section that being a Monge state is equivalent to being $1/\ell$-quantized. This characterization makes it clear why we had to exclude the two-marginal case: for $N = 2$, the example becomes a Monge extreme point.

For the discrete MMOT problems (3.4)-(3.5)/(3.6)-(3.9) with $N \geq 3$ marginals, $\ell \geq 2$ discretization points, and with $\lambda^* = \bar{\lambda}$ a uniform marginal constraint, Monge's ansatz is insufficient. We saw that the relevant transport polytopes exhibit non-Monge extreme points. We can expose these extreme points with general cost functions. As any such cost function behaves identically to its symmetrization, we can assume the exposing cost functions to be symmetric. So, for symmetric cost functions Monge's ansatz is insufficient, but what happens if we restrict the class of cost functions further? Can we find a class that only sees Monge extreme points?

It is a common trade in continuous OT research to pick a specific cost function or class of cost functions and study the behaviour of the solutions they admit. In [46], Gangbo and Święch picked a multi-marginal version of the quadratic cost

$$\sum_{1\leq i<j\leq N} |x_i - x_j|^2$$

and showed that Kantorovich's MMOT problem admits a unique solution and this solution is of Monge-type. This result was the first of its kind and lit the path for what was to come: a line of literature discussing the existence and uniqueness of Monge-type solutions for Kantorovich's MMOT problem. We briefly want to mention two line segments that concern the Coulomb cost

$$\sum_{1\leq i<j\leq N} \frac{1}{|x_i - x_j|}.$$

Colombo, De Pascale, and Di Marino showed in [23] that Kantorovich's MMOT problem admits a unique symmetric solution and this solution is of Monge-type. Their result only holds in the $(d = 1)$-case, i.e., for marginal probabilities on $\mathbb{R}$. In [71], Pass showed that for $d \geq 3$ Kantorovich's MMOT problems admit solutions that exceed the restriction on support-dimension set by Monge's ansatz.

In Core Publication A.1, we showed that - for certain constellations of the parameters N, ℓ and for a far-reaching class of repulsive pairwise costs - the discrete MMOT problems (3.4)-(3.5)/(3.6)-(3.9) (with $\lambda^* = \overline{\lambda}$) admit a unique solution and this solution is of Monge-type. In contrast to [23] and [46], we consider finite instead of continuous state spaces. A pairwise and symmetric cost $c(x_1, x_2, \ldots, x_N) = \sum v(x_i, x_j)$ belongs to the said class if and only if its diagonal entries $v(x, x)$ lie above a problem-dependent threshold. We designed this class as a capsule surrounding discretizations of the Coulomb cost. The Monge-result was an immediate consequence of a necessary support-condition on optimizers, which holds for any $N, \ell \geq 2$. We introduced this condition in Core Publication A.1; it states that an optimal transport plan assigns nonzero probabilities only to those configurations that distribute the N electrons as uniformly as possible among the ℓ discretization points. Figure 4.1 considers 5 electrons on 3 discretization points and depicts all attainable configurations on the left and some unattainable configurations on the right. For certain constellations of N and ℓ - for example $N = 5$ and $\ell = 3$ - there exists a single possibility to achieve a uniform marginal by overlaying the attainable configurations and the corresponding transport plan is of Monge-type.

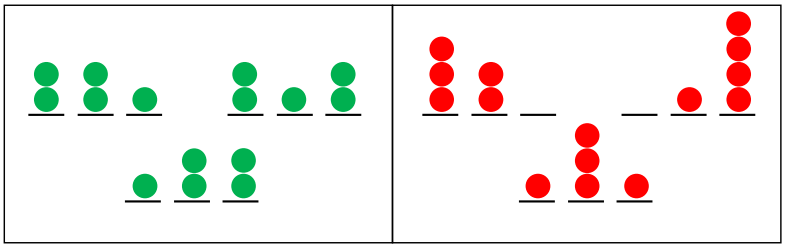

Figure 4.1: With the present figure, we illustrate the necessary support-condition on optimizers established in Core Publication A.1. We consider multi-marginal optimal transport problems on a finite state space X with a uniform marginal constraint and a restriction to symmetric transport plans only. Recall that we can interpret the mass a transport plan assigns to $(x_1, x_2, \dots, x_N) \in X^N$ as probability of the electron configuration that puts electron 1 on x_1, electron 2 on x_2, and so on. The necessary support-condition tells us that, for certain repulsive pairwise costs, any optimizer can assign positive probabilities only to those electron configurations that distribute the N electrons as uniformly as possible among the given sites in X. On the left the reader can see all attainable and on the right some unattainable configurations for the case of $N = 5$ electrons on $|X| = 3$ sites.

Even for repulsive pairwise cost functions, we cannot guarantee the existence of a Monge-type solution to our Kantorovich MMOT problem. Any computational method moving within the low-dimensional Monge class is in danger of never reaching an optimizer - no matter how many updates it performs. Further, Monge states have a quantized character. Optimizing over Monge states means optimizing over the integer elements of a certain polytope and integer optimization is truly not something to long for. In the upcoming section, we will explain why Monge means integer and how to cure the rigidity and insufficiency of Monge's ansatz.

4.3 Quasi-Monge ansatz

We cut to the chase and immediately define quasi-Monge states. For the moment, we assume a uniform marginal to ensure comparability with Monge's ansatz. A probability measure γ on X^N is a quasi-Monge state if there exist N maps $T_1, T_2, \dots, T_N : X \to X$ and a probability measure α on X that fulfill the constraint

$$\frac{1}{N} \sum_{k=1}^{N} T_{k\#}\alpha = \overline{\lambda}$$

and give rise to γ via the formula

$$\gamma = S_N\left((T_1, T_2, T_3, \ldots, T_N)_{\#}\,\alpha\right) = \sum_{r=1}^{\ell} \alpha\left(\{a_r\}\right) S_N \delta_{T_1(a_r), T_2(a_r), \ldots, T_N(a_r)}.$$

Recall that a probability measure γ on X^N is a Monge state if there exist $N-1$ maps $T_2, \ldots, T_N : X \to X$ that fulfill the constraints

$$T_{k\#}\overline{\lambda} = \overline{\lambda} \quad \text{for all } k = 2, \ldots, N$$

and give rise to γ via the formula

$$\gamma = S_N\left((id, T_2, T_3, \ldots, T_N)_{\#}\,\overline{\lambda}\right) = \sum_{r=1}^{\ell} \frac{1}{\ell} S_N \delta_{a_r, T_2(a_r), \ldots, T_N(a_r)}.$$

The quasi-Monge ansatz was introduced by myself and Friesecke in [44] where we also show that a probability measure on X^N is a Monge state if and only if it is a quasi-Monge state with $\alpha = \overline{\lambda}$.

Quasi-Monge states correspond to the sparse elements of a certain coefficient polytope. We can identify any symmetrized Dirac measure $S_N\delta_{a_{i_1}, a_{i_2}, \ldots, a_{i_N}}$ (for $1 \le i_1 \le i_2 \le \ldots \le i_N \le \ell$) by its one-point marginal

$$M_X S_N \delta_{a_{i_1}, a_{i_2}, \ldots, a_{i_N}} = \frac{1}{N} \sum_{k=1}^{N} \delta_{a_{i_k}},$$

as M_X projects the set of symmetrized Dirac measures bijectively onto the set of $1/N$-quantized probability measures

$$\mathcal{P}_{\frac{1}{N}}(X) := \left\{ \lambda \in \mathcal{P}(X) : \lambda\left(\{a_i\}\right) \in \left\{0, \frac{1}{N}, \frac{2}{N}, \ldots, \frac{N}{N}\right\} \text{ for all } i \in \{1, \ldots, \ell\} \right\},$$

see [44]. We capture the inverse relation with the notation γ_λ, i.e., for

$$\lambda = \frac{1}{N} \sum_{k=1}^{N} \delta_{a_{i_k}} \in \mathcal{P}_{\frac{1}{N}}(X),$$

γ_λ is defined as

$$\gamma_\lambda := S_N \delta_{a_{i_1}, a_{i_2}, \ldots, a_{i_N}}.$$

The γ_λs are exactly the extreme points of the set of symmetric probability measures $\mathcal{P}_{\mathrm{sym}}\left(X^N\right)$. By Minkowski's theorem, we can write any admissible object γ, that is any element of

$$\left\{\gamma \in \mathcal{P}_{\mathrm{sym}}\left(X^N\right) : \quad \gamma \mapsto \overline{\lambda}\right\}, \tag{4.4}$$

as convex combinations of said extreme points

$$\gamma = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda \gamma_\lambda \qquad \text{for some coefficients } \alpha = (\alpha_\lambda)_\lambda \text{ that fulfill}$$

$$\alpha_\lambda \geq 0 \text{ for all } \lambda \quad \text{and} \quad \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda = 1.$$

To impose the marginal constraint on the coefficients, proceed as follows: interpret the elements of $\mathcal{P}_{\frac{1}{N}}(X)$ as column vectors, glue these column vectors together to form a $\ell \times \binom{N+\ell-1}{N}$-matrix A ($\binom{N+\ell-1}{N}$ is the cardinality of $\mathcal{P}_{\frac{1}{N}}(X)$), match the order of the coefficients $(\alpha_\lambda)_\lambda$ to the order of the matrix A, and subject the coefficients to the constraint

$$A\alpha \left(= \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda \lambda \right) = \overline{\lambda}.$$

A symmetric probability measure γ on X^N then fulfills $\gamma \mapsto \overline{\lambda}$ if and only if its coefficients α fulfill $A\alpha = \overline{\lambda}$. As the columns of A sum to 1, we already enforce the normalization of the coefficients by imposing the marginal constraint which is why we drop it when formulating our coefficient polytope

$$\left\{ \alpha \in \mathbb{R}^{\binom{N+\ell-1}{N}} : \quad \alpha \geq 0 \text{ and } A\alpha = \overline{\lambda} \right\}. \tag{4.5}$$

In Core Publication A.1, we show that the transport polytope (4.4) is in a linear and bijective relation with the coefficient polytope (4.5). The pairing between coefficients and admissible objects is one-to-one: exactly one collection of coefficients for one admissible object. The transition from (4.4) to (4.5) makes it clear and tractable what it means to be sparse in the symmetric setting. Monge states exactly correspond to the $1/\ell$-quantized coefficients and quasi-Monge states exactly correspond to the coefficients that have at most ℓ nonzero entries.

Monge's ansatz turns the MMOT problems under consideration into integer optimization problems. Yes, coefficients of Monge-type are sparse but due to their $1/\ell$-quantization they are also predominantly of combinatorial nature. The quasi-Monge ansatz is just sparse - no quantization attached.

The quantization and the thereby implied rigidity of Monge's ansatz stems from its origin in unsymmetrized OT. There, Monge's ansatz consists of probability measures on X^N that distribute all their mass along the graph of a function of the first component:

$$(id, T_2, T_3, \ldots, T_N)_{\#} \overline{\lambda}. \tag{4.6}$$

The ansatz requires every component T_k to do exactly what is necessary for (4.6) to fulfill the marginal constraint: every component T_k has to preserve the uniform measure. In the symmetrized case, Monge's ansatz consists of the symmetrizations of measures of the form (4.6):

$$S_N\left((id, T_2, T_3, \ldots, T_N)_{\#}\,\overline{\lambda}\right).$$

Still, every component T_k has to preserve the uniform measure. Here, however it is unnecessary to constrain component wise. Take a probability measure α on X and N maps $T_1, T_2, \ldots, T_N : X \to X$ that fulfill

$$\frac{1}{N}\sum_{k=1}^{N} T_{k\#}\alpha = \overline{\lambda} \tag{4.7}$$

and with

$$S_N\left((T_1, T_2, \ldots, T_N)_{\#}\,\alpha\right) \tag{4.8}$$

you have a probability measure on X^N that (i) for $T_1 = id$ is the symmetrization of a measure that distributes all its mass along the graph of a function of the first component and (ii) fulfills the marginal constraint [44]. In the symmetrized case, the components have to preserve the uniform measure only on average for (4.8) to fulfill the marginal constraint: If you push α forward along the individual components T_k and overlay the push-forwards, the result has to be uniform. The quasi-Monge ansatz consists of probability measures on X^N that take on the form (4.8) subject to the constraint (4.7). This new ansatz takes advantage of the symmetry in two ways: it abolishes the special role of T_1 as identity and requires its components to preserve the uniform measure only on average.

As the quasi-Monge ansatz is more flexible than Monge's ansatz, we can apply it to an arbitrary marginal constraint $\lambda^* \in \mathcal{P}(X)$. The result is probability measures on X^N that take on the form

$$S_N\left((T_1, T_2, \ldots, T_N)_{\#}\,\alpha\right) \quad \text{subject to the constraint} \quad \frac{1}{N}\sum_{k=1}^{N} T_{k\#}\alpha = \lambda^*,$$

where, as before, α is a probability measure on X and $T_1, T_2, \ldots, T_N$ are maps from X to X. The quasi-Monge ansatz has a storage complexity of $\ell \cdot (N+1)$ - an increase when compared to Monge's ansatz by $\ell \cdot 2$. The ℓ values of the map T_1 and the probabilities α assigns to the ℓ singletons make up the additional unknowns. Note that both storage complexities, Monge's and quasi-Monge's, grow linearly with the number of marginals N. In [44], we show that the quasi-Monge

ansatz is - in contrast to Monge's ansatz - sufficient for the MMOT problems under consideration (3.4)-(3.5)/(3.6)-(3.9). Sufficient means that any instance of the said MMOT problems admits a solution that takes on the form prescribed by the ansatz. Overall, the quasi-Monge ansatz is flexible, sparse, and sufficient. Through these qualities, the ansatz gives rise to an efficient computational method to solve the discrete MMOT problems (3.4)-(3.5)/(3.6)-(3.9). For further reading, have a look at the next chapter.

5 Numerical approaches to multi-marginal optimal transport

The curse of dimension is the biggest challenge on the quest for an efficient numerical method for optimal transport (OT) problems with many marginals. The curse consists in the exponential growth of the storage complexity (see Section 3.3). OT problems with $N = 20$ marginals and $\ell = 1000$ discretization points - which correspond to 20 electrons distributed over 1000 sites in the electronic structure application of multi-marginal optimal transport (MMOT) - yield trail states that consist of 10^{60} real numbers. This exponential growth of the problem size hinders us to even store trial or optimal states for problems with a large number of marginals. How can we then solve these problems?

A good place to start looking for an answer is the two-marginal setting. What is it that works there? Can we adapt the method to multiple marginals? Is the adaptation still scalable? In Section 5.1, we introduce the reader to the standard computational method in two marginals, Sinkhorn's algorithm, and explain, why its multi-marginal adaptation suffers from the curse of dimension. Our presentation mostly follows the excellent textbook [73] by Peyré and Cuturi titled "Computational Optimal Transport". As the MMOT adaptation of Sinkhorn's algorithm did not break the curse, we next shine a light on the linear programming aspect of OT. Two- as well as multi-marginal OT problems are linear programs. To solve these well-understood optimization problems, a variety of methods is readily available. Here, we will focus on the simplex method. In contrast to Sinkhorn's algorithm, it operates on sparse states. In Section 5.2, we explain why - even though it operates on sparse states - the simplex method itself still suffers from the curse of dimension and how to adapt the method to break the curse for electronic structure MMOT. This breaking of the curse is the key achievement of Core Publication A.2. We call said adaptation of the simplex method Genetic Column Generation. It is a simple and extremely efficient heuristic that provides approximate solutions for the problems at hand. Here, 'approximate' simply refers to the fact that we cannot certify optimality in general; in the test cases where the optimal solution is known our method proved to be exact. Our letting go of exactness or predetermined approximation quality is justified in light of the recent complexity results by ourselves (Core Publication A.2) and by Altschuler and Boix-Adserà [4, 5]. In Section 5.3, we elaborate on the complexity-theoretic work of the latter authors and its implications. Simply put, their work helps us to identify for which

problems the curse can be broken with exact, approximate, or heuristical methods. Finally, in Section 5.4, we draw an overall conclusion.

Before going into details, let us refer the reader to a recent excellent review titled "Optimal transportation, modelling and numerical simulation" by Benamou [7]. It provides a bird's eye view on the topic which not only reveals the relevant concepts but also connections between them. For further reading on numerical approaches to the MMOT problems arising in electronic structure, we refer the reader to [42].

5.1 The standard in two marginals: Sinkhorn

Why make a fuss about the case of two marginals? The curse of dimension manifests with a growing number of marginals. For two marginals, the number of unknowns is ℓ^2 (with ℓ being the number of discretization points). So, why not simply apply the simplex method and be done with it? The answer lies in the applications of two-marginal OT, as for example statistical machine learning. We can use OT problems to measure the difference of two probability densities or histograms. Consider the two-marginal Kantorvich OT problem, let μ and ν denote the marginals and assume the cost function c to be a distance. Then, the optimal transport cost, that is the optimal value of the problem, defines a distance between the marginals μ and ν [83]. A key task in statistical machine learning consists in finding the right measuring tape - that is notion of distance - to compare probability densities or histograms. While the OT notion of distance makes for an intuitive and powerful candidate from a theoretical perspective, its computational cost based on standard linear programming methods prohibits its use when μ and ν have a high number of support points. With fine discretization and high resolution histograms dominating the field of machine learning, the relevance of OT distances entirely depends on whether we can find creative ways to efficiently solve or approximate the underlying problems.

In 2013, Marco Cuturi [28] proposed to regularize the classic OT distances by adding an entropic term to the original objective function. The resulting regularized problems provide approximate solutions for the original problem and can be solved by a powerful alternating optimization process. This process goes by the names of Sinkhorn's algorithm [81], iterative proportional fitting procedure [33], and Bregman iterative projections [15]. Simplex-like methods move from one sparse state to the next until they have reached the optimum. In contrast, Sinkhorn's algorithm has the subgoal to achieve maximal entropy and as a consequence mollifies or regularizes the previously sparse states. Computationally

speaking, the regularization has a stabilizing effect.

Let us now complement our words with formulas. The above-mentioned entropically regularized OT problems take on the form

$$\begin{aligned} \text{Minimize} \quad & \sum_{i,j} c_{ij}\gamma_{ij} + \varepsilon \sum_{i,j} \gamma_{ij} \left(\log\left(\gamma_{ij}\right) - 1\right) \\ \text{subject to} \quad & \gamma \text{ has marginals } \mu \text{ and } \nu. \end{aligned}$$

The new objective function is strictly convex. Consequently, the regularized problem admits a unique solution γ_ε^*. While solutions to the original, 'unregularized', problem sit on the boundary of the set of admissible trail states, the solution γ_ε^* of the regularized problems sits in the 'interior'. The regularization moves the solutions away from sparse towards more densely occupied states. For decreasing regularization ($\varepsilon \to 0$), the solutions $(\gamma_\varepsilon^*)_\varepsilon$ converge to the solution of the original, 'unregularized', problem that has maximal entropy.

To understand the alternating strategy of Sinkhorn's algorithm, we shine a light on the regularized solutions γ_ε^* from a different vantage point. Let us firstly note that γ_ε^* also solves the problem

$$\begin{aligned} \text{Minimize} \quad & KL\left(\gamma|\kappa\right) \\ \text{subject to} \quad & \gamma \text{ has marginals } \mu \text{ and } \nu. \end{aligned}$$

Here, KL denotes the Kullback-Leibler divergence, which is defined as

$$KL\left(\gamma|\hat{\gamma}\right) := \sum_{i,j} \gamma_{ij} \left(\log\left(\frac{\gamma_{ij}}{\hat{\gamma}_{ij}}\right) - 1\right) + \hat{\gamma}_{ij}$$

for componentwise positive $\hat{\gamma}$, and κ denotes the Gibbs kernel associated to the cost c, which is defined as $\kappa_{ij} := e^{-\frac{c_{ij}}{\varepsilon}}$. Simply put, γ_ε^* being a solution of the above divergence minimization means that γ_ε^* is a projection of the Gibbs kernel κ onto the given transport polytope. We can connect the Gibbs kernel κ and the regularized solution γ_ε^* also in algebraic terms:

$$(\gamma_\varepsilon^*)_{ij} = u_i \kappa_{ij} v_j$$

for two (unknown) vectors u and v. While this algebraic connection does not have the visual appeal of its geometric predecessor, it does yield important computational advantages. It allows us to consider $\#i + \#j$ (u and v) instead of $\#i \cdot \#j$ (γ_ε) unknowns. Now, Sinkhorn's algorithm indeed works with these $\#i + \#j$ unknowns in form of the vectors u and v: it consists of alternating updates

$$u^{(n+1)} := \frac{\mu}{\kappa v^{(n)}} \qquad \text{and} \qquad v^{(n+1)} := \frac{\nu}{\kappa^T u^{(n+1)}}, \tag{5.1}$$

where we apply the division operator componentwise. The reader can think of the process as follows: after starting at a random positive vector $v^{(0)}$, Sinkhorn's algorithm alternatingly updates u and v in such a way that the measures $\gamma_\varepsilon^{(m)}$ defined by

$$\left(\gamma_\varepsilon^{2n+1}\right)_{ij} := u_i^{(n+1)} \kappa_{ij} v_j^n \qquad \left(\gamma_\varepsilon^{2n+2}\right)_{ij} := u_i^{(n+1)} \kappa_{ij} v_j^{(n+1)}$$

fulfill the first marginal constraints for odd $m = 2n + 1$ and the second for even $m = 2n + 2$. Sinkhorn's algorithm is particularly successful at computing OT distances at large scale, that is pairwise distances of a large number of histograms. While already the matrix-vector-products in (5.1) invite parallelism, the matrix-matrix-products that appear in large scale computations do even more so.

Now to the case of $N \geq 3$ marginals. For further reading, we refer to our primary sources [8, 9, 69]. The appealing updating steps (5.1) become unwieldy for multiple marginals. The once innocent matrix-vector-product becomes a multilinear map with $N - 1$ input vectors, whose evaluation requires the summation of ℓ^{N-1} terms (with ℓ being the number of discretization points). As it is the problem size that yields the scale of the summation, it is once again the curse of dimension that throws a wrench into our plans by hindering us to successfully lift the standard in two to multiple marginals. The new updating steps only allow for a fairly low number of discretization points and marginals. Benamou, Carlier, and Nenna [9] proposed a heuristic mesh refinement strategy that moved the little number of available discretization points to the 'right' areas. They determine these 'right' areas based on the expectation that the support of regularized solutions is a mollified version of the sparse support of the original solutions. Even in the case of entropically regularized OT the sparsity of optimizers is used (at least heuristically) to circumvent the curse.

5.2 Genetic Column Generation

In this section, we present the scientific contribution of Core Publication A.2. As in the article itself, the focus lies on introducing a new computational method called Genetic Column Generation. While in principle the method can be adapted to all kinds of MMOT problems, the article targets the discrete MMOT problems that arise in the electronic structure application of OT (see Section 3.1). These problems aim to find the minimum of a linear objective function on a high-dimensional symmetric transport polytope. In contrast to Sinkhorn's algorithm, our method targets the original, 'unregularized', problem. In contrast to Sinkhorn, we embrace sparsity.

The problems at hand. In Core Publication A.2, we apply the Genetic Column Generation method to the following problems.

$$\begin{aligned} \text{Minimize} \quad & \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} c_\lambda \alpha_\lambda \\ \text{subject to} \quad & \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \lambda \alpha_\lambda = \lambda^* \\ & \alpha \geq 0 \end{aligned} \tag{5.2}$$

Here, we face a so-called extremal formulation of the symmetric MMOT problems (3.4)-(3.5)/(3.6)-(3.9). To reach this formulation, write the previously admissible objects γ as convex combination of the extreme points $\{\gamma_\lambda\}$ of the set of symmetric probability measures $\mathcal{P}_{\text{sym}}\left(X^N\right)$ with X denoting the finite state space $\{a_1, a_2, \ldots, a_\ell\}$. So, just as Dantzig and Wolfe did in their pioneering articles [29, 30] on the concept of column generation, we express the admissible objects as convex combinations of the extreme points of a suitably chosen superset. The coefficient of γ_λ is denoted by α_λ. The corresponding cost coefficient c_λ accounts for the transport cost of γ_λ, i.e.,

$$c_\lambda := \int_{X^N} c \mathrm{d}\gamma_\lambda. \tag{5.3}$$

For answers to the questions - what do the extreme points γ_λ look like, why can we index the extreme points $\{\gamma_\lambda\}_\lambda$ via their one-point marginal $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$, and why does (5.2) embody the marginal and normalization constraint - we refer the reader to Section 4.3. While the extremal formulation reduces the number of variables to the actual number of unknowns - $\binom{N+\ell-1}{N}$, this reduction does not suffice to break the curse of dimension (see Section 3.3).

How does Genetic Column Generation relate to the simplex method? As the set of admissible objects of the problems at hand is bounded, they admit an optimal solution. A fundamental truth in linear programming theory tells us that we can choose this solution to be an extreme point. The simplex method is built on this fundamental truth. It moves from one extreme point (= vertex) to the next along cost reducing edges. It looks for these edges through exhaustive search. See [12] for an excellent introduction to the simplex method. Our method also moves from one extreme point to the next along cost reducing edges. We, however, look for these edges through genetic sampling, not exhaustive search and by that circumvent the curse of dimension. Generally speaking, column generation is an adaptation of the simplex method towards flat linear programs which are

linear programs with a high number of variables and comparatively low number of equality constraints. If the number of variables become too high, the linear program in its entirety becomes intractable or even unstorable. But, the simplex method operates on extreme points which for these flat linear programs are sparse, i.e., have few nonzero entries (see, for example, [11]). Column generation makes use of this fact and only stores the sparse data. To move from one sparse data point to a better one, we need to find a cost reducing edge. More concretely, we have to find a new variable that if we switch it on, i.e., allow it to be nonzero, improves our cost. (Switching on a variable means activating a column in the constraint matrix: hence, the name column generation.) The simplex method uses exhaustive search for this task. As there are as many places to look as there are variables, exhaustive search is out of the question for the said flat linear programs. This is why in column generation one casts this search as an optimization problem that looks for the variable that would improve our current cost the most. This problem is referred to as pricing problem. Whether we can apply column generation successfully to a given flat linear program, depends entirely on whether we can treat the pricing problem efficiently. In the classic example of the cutting stock problem the pricing problem turns out to be a knapsack problem that can in practice often be solved efficiently via dynamic programming (see [48] for the original reference and [66] for an excellent review on the concept of column generation).

Our pricing problem takes on the form

$$\begin{aligned} \text{Maximize} \quad & \lambda^T y^* - \frac{N}{2}\lambda^T C\lambda + \frac{N}{2}\operatorname{diag}(C)^T \lambda && (5.4)\\ \text{subject to} \quad & \sum_{i=1}^{|X|} \lambda_i = N && (5.5)\\ & \lambda \in \mathbb{N}_0^{|X|} && (5.6) \end{aligned}$$

or equivalently

$$\begin{aligned} \text{Maximize} \quad & \lambda^T y^* - \frac{N^2}{2}\lambda^T C\lambda + \frac{N}{2}\operatorname{diag}(C)^T \lambda && (5.7)\\ \text{subject to} \quad & \lambda \in \mathcal{P}_{\frac{1}{N}}(X). && (5.8) \end{aligned}$$

Here, we slightly abuse the notion and identify a probability measure $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$ with its vector counterpart $\lambda \in \mathbb{R}^{|X|}$ defined by $\lambda_i := \lambda(\{a_i\})$. The vector y^* denotes the dual to our current sparse data point (more concretely: our sparse data point gives rise to a small primal linear program, y^* is a solution to its dual). The matrix $C = (C_{ij})_{ij} \in \mathbb{R}^{|X|\times|X|}$ arises from the transport cost $c(x_1, x_2, \ldots, x_N) = \sum_{1\le i<j\le N} v(x_i, x_j)$ via

$$C_{ij} = v(a_i, a_j)$$

and diag(C) denotes the vector $\begin{pmatrix} C_{11} & C_{22} & \dots & C_{|X||X|} \end{pmatrix}^T$ consisting of the diagonal entries of C. Note that we consider pairwise interactions only. It is this restriction that allows us to compute the high-dimensional integrals (5.3) with the quadratic form $\frac{N^2}{2}\lambda^T C\lambda - \frac{N}{2}\text{diag}\,(C)^T\,\lambda$. See Section 6.1 for the simplification. As in general the objective of the pricing problem takes on the form $\lambda^T y^* - c_\lambda$, the restriction to pairwise interactions simplifies our pricing objective to a quadratic function. Even though we now 'only' face a quadratic objective, the overall integer optimization problem with a quadratic objective and linear equality constraints is NP-complete, as we prove in Core Publication A.2. We postpone the question on what exactly NP-completeness means to the next section. Now, let us just say, we face a problem that becomes too intractable to be solved with a growing number of marginals and discretization points. Solving the problem means finding the best edge to move forward. But, an improving edge suffices. Accordingly, we are content with finding a column λ with positive objective value. Note: best edge means best λ, improving edge means λ with positive objective value. To find such an improving λ, we employ a genetic sampling method: we admit only 'successful' λs as parents, randomly select one of the parents, then randomly select one of the nearest neighbours of said parent as child, and finally test the child for a positive objective value. Now, how exactly does it work? Our variables α embody the positional distribution of the N electrons in a molecule. This distribution is a stochastic superposition of electron configurations. Our N electrons form configuration j with probability α_j. At any iteration in the Genetic Column Generation optimization process, we sit at a sparse data point. This means our current distribution is a superposition of few configurations. Any admissible λ in our pricing problem represents such a configuration: each quantized mass unit in a component λ_j means that an electron is sitting on the corresponding discretization point a_j. Now, when faced with the task to guess good λs, we choose to learn from our current sparse data points. We randomly choose one of the few configurations that appears with nonzero probability. Then, we randomly choose one of the N electrons. Let's say in the chosen configuration λ it sits on discretization point a_j. And, finally we randomly choose one of the neighbours of a_j and move the electron there. In this manner, we have produced a genetic sample $\hat{\lambda}$ which we now can test for acceptance using our low-complexity acceptance criterion (that is positive objective value regarding the pricing problem).

The sparse states we assume in each iteration belong to the class of quasi-Monge states introduced in Section 4.3. With our Genetic Column Generation method we move from one quasi-Monge state to the next through genetic sampling and subsequent low-complexity testing. So, the quasi-Monge ansatz is not only low-dimensional and sufficient (as discussed in Section 4.3) but also gives rise to an

efficient computational method for the discrete MMOT problems (3.4)-(3.5)/(3.6)-(3.9) for pairwise interacting cost functions.

5.3 How 'hard' are the problems?

In this section, we will shine a light on recent complexity results in MMOT theory by Altschuler and Boix-Adserà. Before we start off with a reminder on what "NP-complete" means, let us note that in general hardness results concern so-called decision problems while we here deal with optimization problems. As the hardness of the decision version of a problem implies the hardness of the corresponding optimization problem, we will not go into the details of decision problems and restrict our attention to optimization problems accepting a slight misuse of terms.

What does it mean to be an NP-complete problem? Let's take our pricing problem (5.4)-(5.8). We have shown that it is NP-complete which means that it belongs to the class NP and it is at least as hard to solve as any other problem in the class NP (see, for example, [57, 24]). It is the latter of the two conditions that concerns the difficulty of solving the problem. Due to the central unanswered question in complexity theory - does P equal NP - it is of comparative instead of explicit nature. Garey and Johnson give the following, more concrete, description in their seminal book [47] on the theory of NP-completeness: If a problem is NP-complete then finding a polynomial time algorithm that solves it - if at all possible - corresponds to a major breakthrough in theoretical computer science.

Hardness of MMOT. In Section 5.1 and 5.2, we took a closer look at Sinkhorn's algorithm and the simplex method. With an increase of the problem size, the curse of dimension tightens the noose around their bottleneck more and more until it becomes unpassable. But, there is still hope. Both, the simplex method and Sinkhorn's algorithm, were discussed in a general, off-the-shelf form. If we focussed on specific classes of cost functions, we might be able to use the implied structure to simplify the problem and derive a computational method for this simplification. In [21], for example, they focussed on matching for teams problems (in particular the special case of Wasserstein barycenter) and presented two efficient computational methods. But, for which classes of cost functions do there exist efficient or to be more concrete polynomial time methods? In their complexity-theoretic work [4, 5], Altschuler and Boix-Adserà show us how to find answers: if there exists a polynomial time method that exactly / approximately solves a class of MMOT problems then the Ellipsoid / Multiplicative Weights Update (MWU) algorithm runs in polynomial time. Here, the time is polynomial in the number of marginals N and the number of discretization points ℓ for exactness and in the number of

marginals N, the number of discretization points ℓ, and $C_{\max}/\varepsilon$ (with $C_{\max}$ denoting the maximal entry in the cost tensor and ε denoting the approximation error) for approximation. The Ellipsoid and the MWU algorithm then run in polynomial time if and only if their bottleneck is solvable in polynomial time, exactly for the Ellipsoid and approximately for MWU. Interestingly, solving their bottlenecks means solving the pricing problem of the column generation routine exactly (Ellipsoid) or approximately (MWU). Overall, Altschuler and Boix-Adserà provide us with the tools we need to show polynomial run time in the tractable cases or NP-hardness in the intractable cases. They use these tools themselves on an array of problems. On the one hand, they establish polynomial time solvability for costs of graphical structure, set-optimization structure, and low-rank plus sparse structure, as they call it. On the other hand, they establish the NP-hardness of low-rank costs, pairwise-interaction costs, and certain repulsive costs. In particular, they establish the NP-hardness of MMOT for the Coulomb-Buckingham potential regarding exact as well as approximate solvability. For the classic Coulomb potential, they conjecture as much (for exactness).

Even if we face an NP-hard problem, there are steps we can take. First of all, NP-hardness concerns the worst case. So, we advise checking whether you can slim down your problem class. If this is not possible, as described in [47], you have a choice to make: hold on to exactness or let go. If you hold on, you can work on cleverly arranging your exhaustive search using for example branch-and-bound or dynamic programming techniques. While these may lead to speed-ups, they will (in general) not allow you to break the curse. If you let go, you can work on finding a clever heuristic that produces not optimal but 'good' solutions 'fast'. While this allows you to break the curse, you lack (in general) a certificate of your solution being 'good' or optimal. With the Genetic Column Generation method, we let go of exactness and designed a heuristic that produces 'good' solutions 'fast'. In the test cases where the optimal solution is known our method even proved to be exact.

5.4 Conclusion and outlook

Genetic Column Generation is a simple and extremely efficient method that provides approximate solutions for the discrete MMOT problems arising in electronic structure physics. Here, the term 'approximate' simply refers to the fact that we cannot certify optimality or we choose not to in light of recent hardness results (see Section 5.2 and 5.3). In the test cases where the optimal solution is known the method proved to be exact. These test cases climaxed in a problem consisting of $2.58 \cdot 10^{31}$ unknowns ($N = 30$ elections on $\ell = 120$ discretization points in one

dimension). The overall number of genetic samples needed in the optimization process seems to grow only polynomially with the number of marginals - which constitutes computational evidence of us breaking the curse and makes the seemingly undoable problem size doable.

In our opinion, Genetic Column Generation is a very promising computational method worthy of further study. As next steps, we will develop more sophisticated sampling routines and investigate the possibility of performance guarantees.

As always, there is also something to learn from alternative methods, not only what to improve, but also what to cherish. Sinkhorn's algorithm profits from the stabilizing effect of the entropic regularization. In the context of Column Generation, various stabilization approaches exist (see [66, 34] for further reading), including a sophisticated dynamic boxstep method [51, 67, 37, 6] that proved to accelerate computations significantly. Recently, two further computational methods for the MMOT problems that arise in electronic structure physics appeared on the literary landscape. In [61, 60], the authors build upon the polynomial formula for extremal N-representable 2-point measures (6.7) to compute approximations using a semi-definite relaxation approach. In [3, 2], the authors chose a discretization different to ours: they kept the state space continuous and 'just' discretized the marginal constraint by a finite number of moment constraints. The corresponding Moment Constrained Optimal Transport problems admit sparse optimizers. Restricting the optimization to admit sparse states only yields a lower-dimensional yet nonlinearly constrained optimization problem. They used a stochastic particle method to solve these problems. The scaling of accuracy and computational cost with system size (in particular the number of moment constraints, limited to 52 for 100 electrons in what we believe to be the key simulation example in [2]) poses an important open question to investigate. We believe a key advantage of the Genetic Column Generation method to be its simplicity: we alternatingly solve small linear programs and perform effective sampling with a low-complexity acceptance criterion.

6 Convex geometry of N-representable measures

In the present chapter, we embed Further Publication B.1 in this dissertation. The key result is of convex-geometric nature. It gives us an explicit expression for the extreme points of a certain polytope that arises through projection of high-dimensional measures. In more detail, the key result gives us an explicit polynomial formula that expresses extremal N-representable k-point probability measures (for $N \geq k$) in terms of their one-point marginal. We state this result, in its bare form, in Section 6.2. In Section 6.1 and 6.3, we investigate the result and possible applications from a computational and probabilistic vantage point. We first bridge the gap to previous chapters and finally embark on an excursion into de Finetti-style representations for finitely extendible random vectors.

As in the previous chapters, we abbreviate optimal transport and multi-marginal optimal transport with OT and MMOT, respectively.

6.1 Application in the column-generation-context and definition of N-representability

The formula

$$M_2 S_N \delta_{a_{i_1}, a_{i_2}, \ldots, a_{i_N}} = \frac{N}{N-1} \lambda \otimes \lambda - \frac{1}{N-1} (id, id)_{\#} \lambda \quad \text{with } \lambda := \frac{1}{N} \sum_{k=1}^{N} \delta_{a_{i_k}} \tag{6.1}$$

goes back to [44] and enabled us to swiftly decide whether to reject or accept a genetic sample in our column generation process. It expresses the two-point marginal of a symmetrized Dirac measure (= an extreme point of the set of symmetric probability measures on X^N) in terms of a one-point probability measure λ on X. Here, the finite state space X consists of ℓ distinct points:

$$X := \{a_1, a_2, \ldots, a_\ell\}.$$

The two-point marginal map M_2 or, more generally, the k-point marginal map $M_k : \mathcal{P}\left(X^N\right) \to \mathcal{P}\left(X^k\right)$ for $1 \leq k \leq N-1$ is defined by

$$M_k \gamma(A) = \gamma(A \times X^{N-k}) \quad \text{for all } A \subset X^k \text{ and all } \gamma \in \mathcal{P}\left(X^N\right),$$

M_N shall equal the identity. In the genetic-column-generation-process, we randomly pick a candidate column from a genetic sampling pool and then test whether this candidate has the potential to improve our current approximate solution. Let λ denote our current candidate. Recall that any column and therefore also our candidate λ is a $\frac{1}{N}$-quantized probability measure, i.e.,

$$\lambda \in \mathcal{P}_{\frac{1}{N}}(X) := \{\lambda \in \mathcal{P}(X) : \quad \lambda(\{a_i\}) \in \{0, 1/N, 2/N, \ldots, N/N\} \text{ for all } i\}.$$

Consequently, we can write λ as follows.

$$\lambda = \frac{1}{N}\sum_{k=1}^{N} \delta_{a_{i_k}} \text{ for some } a_{i_1}, a_{i_2}, \ldots, a_{i_N} \in X.$$

The aforementioned test requires us to compute the transport cost of the extremal symmetric probability measure $S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}$:

$$\int_{X^N} c(x_1, x_2, \ldots, x_N)\, \mathrm{d}S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}(x_1, x_2, \ldots, x_N). \tag{6.2}$$

With a growing problem size, this computation becomes an infeasible task. The integration domain consists of ℓ^N elements and is - already for moderately sized N and ℓ - too big to allow an efficient evaluation of the integral expression (6.2). How can we avoid this infeasibility? For symmetric cost functions with pairwise structure $c(x_1, x_2, \ldots, x_N) = \sum_{1\le i<j\le N} v(x_i, x_j)$, (6.2) becomes:

$$\begin{aligned}
&\int_{X^N} \sum_{1\le i<j\le N} v(x_i, x_j)\, \mathrm{d}S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}(x_1, x_2, \ldots, x_N) \\
=&\binom{N}{2}\int_{X^2} v(x_1, x_2)\, \mathrm{d}M_2S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}(x_1, x_2) \\
=&\binom{N}{2}\int_{X^2} v(x_1, x_2)\, \mathrm{d}\left(\frac{N}{N-1}\lambda\otimes\lambda - \frac{1}{N-1}(id, id)_{\#}\lambda\right)(x_1, x_2).
\end{aligned} \tag{6.3}$$

The first equality is easy to understand and goes back to [43]; the second equality is due to formula (6.1) and goes back to [44]. In the first step, we reduce the size of the integration domain from ℓ^N to ℓ^2; in the second step, we replace the implicit description of the measure as two-point marginal of a high-dimensional object - $M_2S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}$ - with an explicit one that no longer relies on high-dimensional objects - $N/(N-1)\lambda\otimes\lambda - 1/(N-1)(id, id)_{\#}\lambda$. Overall, (6.3) allows us to compute the transport cost of the high-dimensional object $S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}$ through no more than $\mathcal{O}(\ell^2)$ arithmetic operations.

In Further Publication B.1, we generalized the success factor

$$M_2 S_N \delta_{a_{i_1}, a_{i_2}, \ldots, a_{i_N}} = \frac{N}{N-1} \lambda \otimes \lambda - \frac{1}{N-1} (id, id)_{\#} \lambda \quad \text{with } \lambda := \frac{1}{N} \sum_{k=1}^{N} \delta_{a_{i_k}}$$

to k-point marginals for $k = 2, 3, \ldots, N$. This means, we explicitly expressed the k-point marginal of a symmetrized Dirac measure - $M_k S_N \delta_{a_{i_1}, a_{i_2}, \ldots, a_{i_N}}$ - in terms of a one-point probability measure λ:

$$M_k S_N \delta_{a_{i_1}, a_{i_2}, \ldots, a_{i_N}} = F_{N,k}(\lambda) \quad \text{with } \lambda := \frac{1}{N} \sum_{k=1}^{N} \delta_{a_{i_k}}. \tag{6.4}$$

The expression $F_{N,k}(\lambda)$ mimics the key traits of $N/(N-1)\lambda \otimes \lambda - 1/(N-1)\,(id, id)_{\#} \lambda$: it is a polynomial of degree k with leading term $N^k / \prod_{j=0}^{k-1}(N-j) \lambda \otimes \lambda \otimes \ldots \otimes \lambda$ (prefactor times k-fold tensor product of λ with itself) and a series of $k-1$ corrections, with correction j being of order $1/N^j$. For $k = 2$, $F_{N,k}(\lambda)$ equals its predecessor $N/(N-1)\lambda \otimes \lambda - 1/(N-1)\,(id, id)_{\#} \lambda$. For further examples, we refer the reader to Section 6.2.

Just as we used the original formula (6.1) in the genetic column generation method for two-point interactions, we can use the new formula (6.4) for k-point interactions

$$c(x_1, x_2, \ldots, x_N) = \sum_{1 \le i_1 < i_2 < \ldots < i_k \le N} v(x_{i_1}, x_{i_2}, \ldots, x_{i_k}).$$

These range from two-point ($k = 2$) to N-point ($k = N$) interactions. For $k = 2$, k-point interactions revert to two-point interactions and the new formula reverts to the original one. An increase in k allows us to cover more complicated interactions. This comes at a price: the size of the integration domain increases exponentially with k. For ks close to N, we cover almost all cost functions, yet, are again unable to handle the integrals. For small ks, we still can handle the integrals and the new formula (6.4) can continue the success story of its predecessor.

Can we use the k-point interaction structure of cost functions to reduce the number of variables in our discrete MMOT problems? Recall the extremal formulation introduced in Section 5.2:

$$\begin{aligned} \text{Minimize} \quad & c^T \alpha \\ \text{subject to} \quad & \alpha \in \{\alpha \in \mathbb{R}^{\left|\mathcal{P}_{\frac{1}{N}}(X)\right|} : \quad A\alpha = \lambda^* \\ & \qquad\qquad\qquad\qquad \alpha \ge 0\}. \end{aligned}$$

To form the matrix A, proceed as follows. Interpret the elements of $\mathcal{P}_{\frac{1}{N}}(X)$ as column vectors by setting the i-th column entry λ_i to the probability of the i-th singleton $\lambda(\{a_i\})$. Then, put these column vectors next to each other, one after the other. The result is a $|X| \times \left|\mathcal{P}_{\frac{1}{N}}(X)\right|$-matrix - our matrix A. Any column $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$ has a corresponding variable α_λ. Accepting a new column λ means allowing the variable α_λ to be nonzero. So far, we have explained, how to streamline the evaluation process of a candidate column using the formulas (6.1) and (6.4). So, we have only discussed, how to efficiently decide whether to turn a variable α_λ on. But, can we use the k-point interaction structure of cost functions to reduce the number of variables, the α_λs, altogether?

To answer this question, we have to understand, how the α_λs arise. The extreme points of the set of symmetric probability measures on X^N, $\mathcal{P}_{\text{sym}}\left(X^N\right)$, are exactly the symmetrized Dirac measures

$$S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}} \quad \text{for } 1 \leq i_1 \leq i_2 \leq \ldots \leq i_N \leq \ell = |X|.$$

See, for example, [44]. The transport plans of our discrete MMOT problems (3.4)-(3.5)/ (3.6)-(3.9) were exactly the elements γ of $\mathcal{P}_{\text{sym}}\left(X^N\right)$ that fulfilled the marginal constraint $\gamma \mapsto \lambda^*$. The variables α_λ arose as coefficients when writing a transport plan γ as convex combination of the symmetrized Dirac measures. The one-point marginal map M_1 connects a coefficient with its symmetrized Dirac measure: α_λ is the coefficient of $S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}$ if and only if $S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}$ has one-point marginal λ , i.e., $M_1 S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}} = \lambda$. We showed in [44] that for any $1/N$-quantized probability measure λ on X exists exactly one symmetrized Dirac measure on X^N that has one-point marginal λ. We denote the symmetrized Dirac measure with one-point marginal λ by γ_λ. To summarize: the variable α_λ is the coefficient of the extremal symmetric probability measure γ_λ.

Now, we answer the question, whether we can use the k-point interaction structure of cost functions to reduce the number of the variables α_λ. For k-point interactions, a transport plan γ can influence its transport cost only through its k-point marginal $M_k\gamma$:

$$\begin{aligned}&\int_{X^N} \sum_{1\leq i_1<i_2<\ldots<i_k\leq N} v\left(x_{i_1}, x_{i_2}, \ldots, x_{i_k}\right) \mathrm{d}\gamma\left(x_1, x_2, \ldots, x_N\right) = \\ =&\binom{N}{k} \int_{X^k} v\left(x_1, x_2, \ldots, x_k\right) \mathrm{d}M_k\gamma\left(x_1, x_2, \ldots, x_k\right).\end{aligned}$$

Accordingly, we can replace the original set of transport plans with their k-point marginals and optimize over these instead. The new set of admissible objects then

consists of the elements μ_k of $M_k\mathcal{P}_{\text{sym}}\left(X^N\right)$ that have one-point marginal λ^*, i.e., $\mu_k \mapsto \lambda^*$. By linearity of the k-point marginal map M_k, the following holds: if $\alpha = (\alpha_\lambda)_\lambda$ are the coefficients of a transport plan γ,

$$\gamma = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda \gamma_\lambda,$$

then, α are also the coefficients of $M_k\gamma$,

$$M_k\gamma = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda M_k\gamma_\lambda.$$

We have written the lower-dimensional admissible objects, the $M_k\gamma$s, as convex combinations of the k-point marginals of extremal high-dimensional probability measures, the $M_k\gamma_\lambda$s. But, do we really need all these k-point marginals? In other words: Are all measures of the form

$$M_k\gamma_\lambda \quad \lambda \in \mathcal{P}_{\frac{1}{N}}(X) \tag{6.5}$$

extremal with respect to $M_k\mathcal{P}_{\text{sym}}\left(X^N\right)$? If this was not the case, we could shave off any coefficient α_λ that is associated with a non-extremal $M_k\gamma_\lambda$. Unfortunately, the answer is yes. In [44] ($k = 2$) and Further Publication B.1 ($2 < k < N$), we showed that all measures of the form (6.5) are pairwise distinct and extremal with respect to $M_k\mathcal{P}_{\text{sym}}\left(X^N\right)$. Consequently, we need all the $M_k\gamma_\lambda$s and cannot shave off any coefficients. The number of variables (=coefficients) of our discrete MMOT problems remains unchanged.

In Further Publication B.1, we refer to the elements of $M_k\mathcal{P}_{\text{sym}}\left(X^N\right)$ (k-point marginals of symmetric probability measures on X^N) as N-representable probability measures on X^k. In doing so, we follow the example set by Friesecke, Mendl, Pass, Cotar, and Klüppelberg in [43]; they introduced the concept of N-representability. A proper definition goes as follows.

Definition 6.1 (N-representability)**.** *A k-point probability measure $\mu_k \in \mathcal{P}\left(X^k\right)$ is called N-representable, if there exists a symmetric N-point probability measure $\gamma \in \mathcal{P}_{sym}\left(X^N\right)$ that has μ_k as its k-point marginal, i.e., $M_k\gamma = \mu_k$. The N-point measure γ is then called a representing measure of μ_k. We will denote the set of N-representable k-point probability measures on X^k as $\mathcal{P}_{N\text{-}rep}\left(X^k\right)$, i.e.,*

$$\mathcal{P}_{N\text{-}rep}\left(X^k\right) = M_k\mathcal{P}_{sym}\left(X^N\right).$$

6.2 Extremal N-representable measures

The set $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ of N-representable probability measures on X^k consists of all probability measures μ_k on X^k that are the k-point marginal of some symmetric probability measure on X^N (see definition at the end of Section 6.1). Here, ℓ distinct points make up the finite state space X:

$$X := \{a_1, a_2, \ldots, a_\ell\}.$$

The definition of N-representability says that we can 'extend' an N-representable probability measure on X^k to a symmetric probability measure on X^N. The latter form a convex subset of $\mathcal{P}\left(X^N\right)$. To be more concise, the set of symmetric probability measures on X^N, $\mathcal{P}_{\text{sym}}\left(X^N\right)$, arises as the intersection of finitely many closed halfspaces. As $\mathcal{P}_{\text{sym}}\left(X^N\right)$ is also bounded we can generate it by forming the convex hull of finitely many points (see, e.g., [76]). By definition, $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ is the image of $\mathcal{P}_{\text{sym}}\left(X^N\right)$ under the linear marginal map M_k. As such, $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ is itself the convex hull of finitely many points. Any selection of finitely many points that generates $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ contains its extreme points with the minimal selection consisting exactly of said extreme points. In what follows, we will give an abstract and explicit description of these extreme points.

Let us start with the abstract description of the extreme points of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$. We reiterate: $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ is the image of $\mathcal{P}_{\text{sym}}\left(X^N\right)$ under the linear map M_k. Consequently, if $\mathcal{P}_{\text{sym}}\left(X^N\right)$ equals the convex hull of a finite selection of points then $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ equals the convex hull of the k-point marginals of this finite selection of points. The extreme points of $\mathcal{P}_{\text{sym}}\left(X^N\right)$ are the symmetrized Dirac measures

$$S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}} \quad \text{for some } 1 \leq i_1 \leq i_2 \leq \ldots \leq i_N \leq \ell.$$

See, for example, [44]. By Minkowski's theorem (see Section 4.1), $\mathcal{P}_{\text{sym}}\left(X^N\right)$ equals the convex hull of its extreme points. Then, $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ equals the convex hull of points of the form

$$M_k S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}} \quad \text{for some } 1 \leq i_1 \leq i_2 \leq \ldots \leq i_N \leq \ell. \tag{6.6}$$

Now, any extreme point of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ must be of the form (6.6). But, are the measures of the form (6.6) pairwise distinct? And, which of them are extreme points? As for different ordered indices $(i_1, i_2, \ldots, i_N) \neq (j_1, j_2, \ldots, j_N)$ the extremal measures $S_N\delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}$ and $S_N\delta_{a_{j_1},a_{j_2},\ldots,a_{j_N}}$ have different one-point marginals [44], the measures (6.6) are pairwise distinct. Further, all of them are extreme points. For any $2 \leq k \leq N-1$. This was shown in [44] ($k = 2$) and Further Publication B.1 ($3 \leq k \leq N-1$). For the lift to ks greater than 2, we used the $k = 2$-result

as a starting point. Let us summarize the abstract description. For fixed $N \geq 2$ and $|X| = \ell \geq 2$, all sets $\mathcal{P}_{N\text{-rep}}(X^2)$, $\mathcal{P}_{N\text{-rep}}(X^3)$, ..., and $\mathcal{P}_{N\text{-rep}}\left(X^N\right)$ have the same number of extreme points. The marginal maps M_k connect these extreme points: for any $2 \leq k \leq N$, M_k maps the set of extreme points of $\mathcal{P}_{\text{sym}}\left(X^N\right)$ bijectively onto the set of extreme points of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$. So, even tough the sets $\mathcal{P}_{N\text{-rep}}(X^2)$, $\mathcal{P}_{N\text{-rep}}(X^3)$, ..., and $\mathcal{P}_{N\text{-rep}}\left(X^N\right)$ live in completely different spaces, dimension-wise, they all share part of their convex structure.

Let us move on to the explicit description of the extreme points of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$.

Before we will provide formulas that explicitly express extreme points of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$, i.e., measures of the form

$$M_k S_N \delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}} \quad \text{for some } 1 \leq i_1 \leq i_2 \leq \ldots \leq i_N \leq \ell,$$

in terms of their one-point marginal, we will take a closer look at these one-point marginals. The one-point marginal of a symmetrized Dirac measure $S_N \delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}$ is given by

$$M_1 S_N \delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}} = \frac{1}{N} \sum_{k=1}^{N} \delta_{a_{i_k}}$$

(see [44]). Different ordered indices $i_1 \leq i_2 \leq \ldots \leq i_N$ yield different one-point marginals. Overall, M_1 maps the set of symmetrized Dirac measures bijectively onto its range, the set $\mathcal{P}_{\frac{1}{N}}(X)$ of $1/N$-quantized probability measures on X [44]. From a convex geometric point of view, $\mathcal{P}_{\frac{1}{N}}(X) = M_1\left\{S_N \delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}\right\}$ plays a less fundamental role in $M_1 \mathcal{P}_{\text{sym}}\left(X^N\right) = \mathcal{P}_{N\text{-rep}}(X)$ than, say, $M_2\left\{S_N \delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}\right\}$ in $M_2 \mathcal{P}_{\text{sym}}\left(X^N\right) = \mathcal{P}_{N\text{-rep}}(X^2)$. For $k \geq 2$, $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ is a true subset of $\mathcal{P}\left(X^k\right)$ and the set $M_k\left\{S_N \delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}\right\}$ exactly consists of the extreme points of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$. For $k = 1$, $\mathcal{P}_{N\text{-rep}}(X)$ equals $\mathcal{P}(X)$ and $M_1\left\{S_N \delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}\right\}$ contains far more elements than the extreme points of $\mathcal{P}_{N\text{-rep}}(X)$. The latter exactly correspond to the Dirac measures δ_a for $a \in X$. The set $\mathcal{P}_{\frac{1}{N}}(X) = M_1\left\{S_N \delta_{a_{i_1},a_{i_2},\ldots,a_{i_N}}\right\}$ plays an important role when it comes to the convex geometry of its extensions. It will serve as a parameter space for the upcoming explicit description of extremal N-representable measures on X^k.

We will provide the explicit formulas in order of increasing k. For $k = 2$, a probability measure μ_2 on X^2 is an extreme point of $\mathcal{P}_{N\text{-rep}}(X^2)$ if and only if

$$\mu_2 = \frac{N}{N-1}\left[\lambda \otimes \lambda - \frac{1}{N}(id, id)_{\#}\lambda\right] \quad \text{for a } \lambda \in \mathcal{P}_{\frac{1}{N}}(X). \tag{6.7}$$

This characterization is due to [44]; some preliminary results were proven in [39, 43]. For $k = 3$, a probability measure μ_3 on X^3 is an extreme point of $\mathcal{P}_{N\text{-rep}}(X^3)$ if and only if

$$\mu_3 = \frac{N^2}{(N-1)(N-2)} \left[\lambda \otimes \lambda \otimes \lambda - \frac{3}{N} S_3 \left((id, id)_{\#} \lambda \otimes \lambda \right) + \frac{2}{N^2} (id, id, id)_{\#} \lambda \right] \quad \text{for a } \lambda \in \mathcal{P}_{\frac{1}{N}}(X). \tag{6.8}$$

This characterization is due to [61] and Further Publication B.1. In [61],the polynomial expression (6.8) appeared for the first time; in Further Publication B.1, we show that measures of the form (6.8) are exactly the extreme points of $\mathcal{P}_{N\text{-rep}}(X^3)$. For $k = 4$, a probability measure μ_4 on X^4 is an extreme point of $\mathcal{P}_{N\text{-rep}}(X^4)$ if and only if

$$\begin{aligned} \mu_4 = \frac{N^3}{(N-1)(N-2)(N-3)} \Bigg[& \lambda \otimes \lambda \otimes \lambda \otimes \lambda - \frac{6}{N} S_4 \left((id, id)_{\#} \lambda \otimes \lambda \otimes \lambda \right) \\ & + \frac{8 S_4 \left((id, id, id)_{\#} \lambda \otimes \lambda \right) + 3 S_4 \left((id, id)_{\#} \lambda \otimes (id, id)_{\#} \lambda \right)}{N^2} \\ & - \frac{6}{N^3} (id, id, id, id)_{\#} \lambda \Bigg] \qquad \text{for a } \lambda \in \mathcal{P}_{\frac{1}{N}}(X). \end{aligned} \tag{6.9}$$

This characterization can now be attributed fully to Further Publication B.1. Therein, we also provide an analogous characterization for $k \geq 5$: A probability measure μ_k on X^k is an extreme point of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ if and only if

$$\mu_k = F_{N,k}(\lambda) \quad \text{for a } \lambda \in \mathcal{P}_{\frac{1}{N}}(X).$$

Here, $F_{N,k}$ is a polynomial of degree k with leading term $N^{k-1} / \prod_{j=1}^{k-1} (N-j) \lambda^{\otimes k}$ (prefactor times k-fold tensor product of λ with itself) and a series of $(k-1)$ corrections, with correction j being of order $1/N^j$. For $k = 2, 3$, and 4, $F_{N,k}(\lambda)$ equals the polynomial expression in (6.7), (6.8), and (6.9), respectively. For the explicit formulas for general k, we refer the reader to Further Publication B.1.

The polynomials $F_{N,k}$ also characterize the extreme points of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ for more general state spaces X. Let X be a Polish space, i.e., a complete and separable metric space. Note, that Polish spaces include the continuous state spaces $\mathbb{R}^d$. Then, we show in Further Publication B.1 that for any $N \geq k \geq 2$, a probability measure μ_k on X^k is an extreme point of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ if and only if

$$\mu_k = F_{N,k}(\lambda) \quad \text{for a } \lambda \in \mathcal{P}_{\frac{1}{N}}(X).$$

We further show that each of these extreme points is also an exposed point. For finite state spaces, it is needless to show that we can expose the extreme points due to their finite number (see the end of Section 4.1). Overall, Further Publication B.1 fully (abstractly and explicitly) describes the extreme and exposed points of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ for finite as well as continuous state spaces.

6.3 Application in the context of finitely extendible random vectors

The polynomial formula $\mu_k = F_{N,k}(\lambda)$, which expresses extremal N-representable measures in terms of their one-point marginal, led us to a finite form of de Finetti's theorem. Here, we present this finite form and highlight its place in the literary landscape of de Finetti-style representations. We start off with a transition to probabilistic nomenclature, where we explain, for example, the terms 'random vector' and 'exchangeability'. By means of an illustrative example, we clarify that the notion of exchangeability covers more than independent and identically distributed random variables (iids). Next, we state a version of de Finetti's theorem whose generality is due to Hewitt and Savage [50]. In the theorem, de Finetti considers random vectors that can be extended to an 'infinitely long exchangeable random vector' (=infinitely extendible random vectors) and expresses them as convex mixture of iids. What happens if a random vector can only be extended by a finite amount? We illustrate - using a well-known example - that finitely extendible random vectors need not be a convex mixture of iids. However, they can be represented as signed mixtures of iids [32, 55, 59, 54] or as convex mixtures of certain correlated random variables, which are iids at the highest order but show lower order correlations (Further Publication B.1). We state both representations as well as de Finetti's theorem itself for random vectors of length 4. This restriction makes it easier to compare the three statements. Let us stress that it is unnecessary for their validity. In Further Publication B.1, we present further insights that go beyond what we discuss in the present chapter. Among others, we recover the de Finetti theorem in the presented Hewitt-Savage version, extend the Diaconis-Freedman error bounds, and present a connection to the Ewens sampling formula from genetics. For the details, see Further Publication B.1.

Transition to probabilistic nomenclature. Given k random variables Z_1, $Z_2, \ldots, Z_k$, we consider the vector $Z = (Z_1, Z_2, \ldots, Z_k)$ whose i-th component is given by the random variable Z_i. We call such a vector Z random vector of length k. A random vector $Z = (Z_1, Z_2, \ldots, Z_k)$ is called exchangeable if its law remains unchanged after its components have been permuted, i.e., for any permutation σ of $\{1, 2, \ldots, k\}$ the random vector $Z_\sigma = \big(Z_{\sigma(1)}, Z_{\sigma(2)}, \ldots, Z_{\sigma(k)}\big)$ has the same law as Z itself. Using the terms predating this section: A random vector is exchangeable if its law is symmetric. Now, we come to the issue of extendibility. A random vector $Z = (Z_1, Z_2, \ldots, Z_k)$ is called N-extendible (for $N \geq k$) if it equals (in law) the first k components of some exchangeable random vector of length N. Analogously, the random vector $Z = (Z_1, Z_2, \ldots, Z_k)$ is called infinitely-extendible if it equals (in law) the first k components of some exchangeable sequence $(Z_i)_{i\in\mathbb{N}}$ of random

variables. Here, a sequence $(Z_i)_{i\in\mathbb{N}}$ of random variables is called exchangeable if its law remains unchanged after a finite number of its components have been permuted, i.e., for any permutation σ of $\mathbb{N}$ that leaves all but finitely many natural numbers unchanged the sequence $\left(Z_{\sigma(i)}\right)_{i\in\mathbb{N}}$ has the same law as $(Z_i)_{i\in\mathbb{N}}$ itself. Using the terms predating this section: A random vector is N-extendible if its law is N-representable.

Example [49]**.** What follows is an example of a random vector that is exchangeable but whose components are not iids. Imagine yourself in front of an urn. It is filled with N balls; each ball has a certain colour; a single colour may appear multiple times. Now, you randomly draw a ball and place it in front of you instead of putting it back in the urn. You performed a draw without replacement. Keep going until you have k (for some $2 \leq k \leq N$) balls in front you. Let Z_i, for $1 \leq i \leq k$, be the colour of the i-th drawn ball. Then, $Z_1, Z_2, \ldots, Z_k$ are no iids, but they form an exchangeable random vector. Why are the colours of the drawn balls not independent? Suppose one of the colours - say, purple - appears only once in the urn, and we get the corresponding ball with our first draw. Then, the colour purple cannot appear in any of the following draws. So, $Z_1, Z_2, \ldots, Z_k$ cannot be independent. Why is the random vector $Z = (Z_1, Z_2, \ldots, Z_k)$ exchangeable? In the present experiment the probability of a given colour combination only depends on the overall number of balls N, the number of draws k, the colour-make up of the urn, and the colour-make of the draw; it does not depend on the order of the drawn colours. (We advise readers who want more insight into the dependencies to work out the details of a small example.) Consequently, any fixed colour combination (for example, (red, red, green, blue)) appears with the same probability as any of its permuted versions (for example, (blue, red, green, red)). This proves the random vector $Z = (Z_1, Z_2, \ldots, Z_k)$ to be exchangeable.

de Finetti-Hewitt-Savage theorem. Even though the terms 'exchangeable random vector' and 'iids' are not equivalent, they are of similar nature. Both concern joint distributions of identically distributed random variables. The present theorem puts the notions of exchangeability and extendibility in relation with iids: it identifies infinitely-extendible random vectors with convex mixtures of iids. We can trace back the origin of this representation to de Finetti [31]. In its original form it concerns random variables taking values in $\{0, 1\}$. A probability measure on $\{0, 1\}$ only depends on the probability $p \in [0, 1]$ of the singleton $\{1\}$. This is why the de Finetti representation for 01-random variables consists of an integral over the interval $[0, 1]$. After several intermediate steps, the generalization process of de Finetti's theorem climaxed in the celebrated manuscript by Hewitt and Savage [50]. Their work gave rise to the following version of de Finetti's theorem.

Theorem 6.2 (de Finetti-Hewitt-Savage theorem)**.** *Let Z_1, Z_2, Z_3, and Z_4 be random variables that take values in some Polish space X. Then, the random vector $Z = (Z_1, Z_2, Z_3, Z_4)$ is infinitely-extendible if and only if there exists a probability measure $\alpha \in \mathcal{P}(\mathcal{P}(X))$ such that*

$$\mathbb{P}(Z \in \cdot\,) = \int_{\mathcal{P}(X)} \lambda \otimes \lambda \otimes \lambda \otimes \lambda \quad \mathrm{d}\alpha(\lambda).$$

The analogous statement holds for any $k \geq 2$, not only $k = 4$. For an exchangeable sequence $(Z_i)_{i\in\mathbb{N}}$, there exists a unique choice of mixing measure α. Note that Dubins and Freedman [38] showed that for the representation to hold not any measurable space will do, i.e., the random variables cannot take values in an arbitrary space.

Example [49, 35, 36]**.** What follows is a textbook example for an exchangeable random vector that cannot be written as convex mixture of iids. Let us consider two random variables Z_1 and Z_2 taking values in $\{0, 1\}$. The random vector $Z = (Z_1, Z_2)$ takes the values $(0, 1)$ and $(1, 0)$ each with probability $1/2$, and the values $(0, 0)$ and $(1, 1)$ each with probability 0; i.e.,

$$\begin{aligned}\mathbb{P}(Z = (1,0)) = \mathbb{P}(Z = (0,1)) &= \frac{1}{2},\\ \mathbb{P}(Z = (0,0)) = \mathbb{P}(Z = (1,1)) &= 0.\end{aligned}$$

Clearly, Z is exchangeable. Now, assume we can write Z as a convex mixture of iids. Then,

$$\mathbb{P}(Z = (a,b)) = \int_{\mathcal{P}(\{0,1\})} \lambda(\{a\}) \cdot \lambda(\{b\})\, \mathrm{d}\alpha(\lambda).$$

Recall that we may identify any $\lambda \in \mathcal{P}(\{0, 1\})$ with the probability p it assigns the singleton $\{1\}$. This allows us to simplify the integration domain as well as the integrand:

$$\mathbb{P}(Z = (a,b)) = \int_0^1 p^{a+b} \cdot (1-p)^{2-(a+b)}\, \mathrm{d}\alpha(p).$$

Using $\mathbb{P}(Z = (0, 0)) = \mathbb{P}(Z = (1, 1)) = 0$, leads us to a contradiction:

$$\begin{aligned}0 = \mathbb{P}(Z = (0,0)) &= \int_0^1 (1-p)^2\, \mathrm{d}\alpha(p) && \to \alpha = \delta_1\\ 0 = \mathbb{P}(Z = (1,1)) &= \int_0^1 p^2\, \mathrm{d}\alpha(p) && \to \alpha = \delta_0.\end{aligned}$$

The example illustrates that finite extendibility (here 2-extendibility) is insufficient to guarantee a convex-mixture-of-iids-representation. Note that the random vector

Z is an exchangeable pair of random variables that we cannot extend. (To show the latter, assume we can extend Z to a random vector of length 3 and deduce a contradiction.) Consequently, Z and the random vectors that are admissible in the de Finetti-Hewitt-Savage theorem are on opposite ends of the extendibility scale. Our intuition tells us that the further a random vector can be extended the closer it should be to a convex mixture of iids. While this insight is important and its pursuit proved fruitful for Diaconis and Freedman, it still leaves us short of an explicit representation. In other words: An N-extendible random vector may very well be close to a convex mixture of iids but how do we write it out? In the following, we will present the reader with two finite forms of de Finetti's theorem that write out N-extendible random vectors explicitly.

First finite form of de Finetti's theorem: signed mixtures of iids. The beauty and significance of de Finetti's theorem lies (partly) in its simplicity. It expresses infinitely-extendible random vectors as convex mixtures of iids; iids are fundamental objects in probability theory and statistics, well-understood, and useful tools. The first finite form [32, 55, 59, 54] of de Finetti's theorem holds on to the integrand - iids - and its simplicity. In return, they have to make concessions regarding 'everything around the integrand': they relax the notion of mixture by considering signed instead of convex mixing measures; (as a consequence) their representation result only provides a necessary condition on N-extendibility. As demonstrated in [54], the signed mixing measure for exchangeable random vectors need not be unique. In contrast to de Finetti's theorem, we can apply the upcoming finite form to arbitrary measurable spaces.

Theorem 6.3 (First finite form of de Finetti's theorem: signed mixtures of iids)**.** *Let Z_1, Z_2, Z_3, and Z_4 be random variables that take values in an arbitrary space X. If the random vector $Z = (Z_1, Z_2, Z_3, Z_4)$ is N-extendible for some $N \geq 4$ then there exists a signed measure α on $\mathcal{P}(X)$ such that*

$$\mathbb{P}(Z \in \cdot\,) = \int_{\mathcal{P}(X)} \lambda \otimes \lambda \otimes \lambda \otimes \lambda \quad \mathrm{d}\alpha(\lambda).$$

The analogous statement holds for any $(N \geq) k \geq 2$, not only $k = 4$. For an exchangeable random vector $Z = (Z_1, Z_2, \ldots, Z_N)$, the choice of mixing measure is not necessarily unique. Note that even if $\mathbb{P}(Z \in \cdot\,) = \int \lambda^{\otimes 4} \mathrm{d}\alpha(\lambda)$, $\int \lambda^{\otimes 5} \mathrm{d}\alpha(\lambda)$ need not be a nonnegative measure.

Second finite form of de Finetti's theorem: convex mixtures of urn distributions. A look behind the scenes of the de Finetti-Hewitt-Savage theorem reveals a convex geometric side of the result. Hewitt and Savage actually show

that the extreme points of the set of symmetric probability measures on infinitely but countably many coordinates are exactly the measures of the form

$$\lambda \otimes \lambda \otimes \ldots =: \lambda^{\otimes\infty} \quad \text{for } \lambda \in \mathcal{P}(X).$$

The integral representation in the de Finetti-Hewitt-Savage theorem then represents a 'continuous convex combination' of these extreme points.
Naturally, some researchers working on finite exchangeability and extendibility took note and pursued a convex geometric approach. The key terms popping up in this context are 'hypergeometric distributions' and 'urn distributions'. They refer to the distribution of ball-colours when drawing k balls out of an urn containing N balls. We considered these distributions already in the first example of the present section, but will reiterate the important details to ensure readability. You stand in front of an urn. This urn contains N balls. Each ball has a certain colour. A single colour may appear multiple times. You successively draw k balls without replacement. The key terms 'hypergeometric distribution' and 'urn distribution' refer to the distribution of the k colours you drew. These distributions appear for example in a simple proof [49] of de Finetti's theorem for 01-random variables, as an intermediate step in the derivation of the signed measure representation in [59], and in the celebrated work by Diaconis and Freedman [35, 36] where they estimated the error between the law of a finitely extendible random vector and the closest convex mixture of iids. The lack of an explicit general closed form for the 'urn distributions' is present throughout these references. The 'urn distributions' are understood abstractly. Explicit probabilities are limited to special cases. Except in [36]. Therein, the probability of k-colour-combinations is given after the urn, the number of draws, and other variables were fixed as experiment parameters. With our polynomials $F_{N,k}(\lambda)$, we give an explicit general closed form for the 'hypergeometric' and 'urn distributions'. Hereby, N is the number of balls in the urn, k the number of draws and the $1/N$-quantized one-point probability measure λ on the colour space X (a Polish space) represents the composition of the urn. The closed form allowed us to write N-extendible random vectors as convex mixtures of $F_{N,k}(\lambda)$s in an explicit(!) manner just as Hewitt and Savage wrote infinitely-extendible random vectors as convex mixtures of iids.

Theorem 6.4 (Second finite form of de Finetti's theorem: convex mixtures of 'urn distributions')**.** *Let Z_1, Z_2, Z_3, and Z_4 be random variables that take values in some Polish space X. Then, the random vector $Z = (Z_1, Z_2, Z_3, Z_4)$ is N-extendible for some $N \geq 4$ if and only if there exists a probability measure $\alpha \in \mathcal{P}(\mathcal{P}(X))$ such*

that $\alpha\left(\mathcal{P}_{\frac{1}{N}}(X)\right) = 1$ *and*

$$\mathbb{P}(Z \in \cdot) = \int_{\mathcal{P}_{\frac{1}{N}}(X)} \frac{N^3}{(N-1)(N-2)(N-3)} \Big[\lambda \otimes \lambda \otimes \lambda \otimes \lambda - \frac{6}{N} S_4\left((id,id)_{\#}\lambda \otimes \lambda \otimes \lambda\right) + \frac{8S_4\left((id,id,id)_{\#}\lambda \otimes \lambda\right) + 3S_4\left((id,id)_{\#}\lambda \otimes (id,id)_{\#}\lambda\right)}{N^2} - \frac{6}{N^3}(id,id,id,id)_{\#}\lambda\Big] \quad \mathrm{d}\alpha(\lambda).$$

In Core Publication B.1, we provided an analogous statement for any $(N \geq)k \geq 2$, not only $k = 4$. Then, the integrand, which corresponds to the polynomial $F_{N,4}(\lambda)$, needs to be replaced with $F_{N,k}(\lambda)$. For an exchangeable random vector $Z = (Z_1, Z_2, \ldots, Z_N)$, there exists a unique choice of mixing measure α.

To put it in a nutshell: an N-extendible random vector may not be representable as a convex mixture of iids; the signed-mixture-of-iids-approach holds on to the iids and pays for it by allowing signed mixing measures; the convex-mixture-of-urn-distributions-approach replaces iids with their finite analogue - urn distributions - and can so keep the mixing convex.

Bibliography

[1] M. Agueh and G. Carlier. Barycenters in the Wasserstein Space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.

[2] A. Alfonsi, R. Coyaud, and V. Ehrlacher. Constrained overdamped Langevin dynamics for symmetric multimarginal optimal transportation. *Mathematical Models and Methods in Applied Sciences*, 32(03):403–455, 2022.

[3] A. Alfonsi, R. Coyaud, V. Ehrlacher, and D. Lombardi. Approximation of Optimal Transport problems with marginal moments constraints. *Mathematics of Computation*, 2020.

[4] J. M. Altschuler and E. Boix-Adserà. Hardness results for Multimarginal Optimal Transport problems. *Discrete Optimization*, 42:100669, 2021.

[5] J. M. Altschuler and E. Boix-Adserà. Polynomial-time algorithms for Multimarginal Optimal Transport problems with structure. *arXiv - preprint*, 2021.

[6] H. Ben Amor. Stabilisation de l'algorithme de génération de colonnes. École Polytechnique de Montréal, Montréal, Québec, Canada, Ph.D. Thesis, 2002.

[7] J.-D. Benamou. Optimal transportation, modelling and numerical simulation. *Acta Numerica*, 30:249–325, 2021.

[8] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna, and G. Peyré. Iterative Bregman Projections for Regularized Transportation Problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.

[9] J.-D. Benamou, G. Carlier, and L. Nenna. A Numerical Method to Solve Multi-Marginal Optimal Transport Problems with Coulomb Cost. In *Splitting Methods in Communication, Imaging, Science, and Engineering*, pages 577–601. Springer International Publishing, 2016.

[10] J.-D. Benamou, G. Carlier, and L. Nenna. Generalized incompressible flows, multi-marginal transport and Sinkhorn algorithm. *Numerische Mathematik*, 142(1):33–54, 2019.

[11] D. P. Bertsekas. *Convex Optimization Theory*. Athena Scientific, 06 2009.

[12] D. Bertsimas and J. N. Tsitsiklis. *Introduction to linear optimization*. 6. Athena Scientific, 1997.

[13] U. Bindini and L. De Pascale. Optimal transport with Coulomb cost and the semiclassical limit of density functional theory. *Journal de l'École polytechnique – Mathématique*, 4:909–934, 2017.

[14] G. Birkhoff. Tres observaciones sobre el algebra lineal. *Rev. Univ. Nac. Tucuman (A)*, 5(3):147–151, 1946.

[15] L. M. Bregman. The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics*, 7(3):200–217, 1967.

[16] Y. Brenier. Décomposition polaire et réarrangement monotone des champs de vecteurs. *C. R. Acad. Sci. Paris Sér. I Math.*, 305(19):805–808, 1987.

[17] Y. Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Communications on Pure and Applied Mathematics*, 44(4):375–417, 1991.

[18] Y. Brenier. The dual Least Action Problem for an ideal, incompressible fluid. *Archive for Rational Mechanics and Analysis*, 122(4):323–351, 1993.

[19] G. Buttazzo, L. De Pascale, and P. Gori-Giorgi. Optimal-transport formulation of electronic density-functional theory. *Physical Review A*, 85:062502, 6 2012.

[20] G. Carlier and I. Ekeland. Matching for teams. *Economic Theory*, 42(2):397–418, 2010.

[21] G. Carlier, A. Oberman, and E. Oudet. Numerical methods for matching for teams and Wasserstein barycenters. *ESAIM: M2AN*, 49(6):1621–1642, 2015.

[22] H. Chen, G. Friesecke, and C. B. Mendl. Numerical Methods for a Kohn-Sham Density Functional Model Based on Optimal Transport. *Journal of Chemical Theory and Computation*, 10(10):4360–4368, 10 2014.

[23] M. Colombo, L. De Pascale, and S. Di Marino. Multimarginal Optimal Transport Maps for One-dimensional Repulsive Costs. *Canadian Journal of Mathematics*, 67(2):350—368, 2015.

[24] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, Cambridge, Massachusetts, 3 edition, 2009.

[25] C. Cotar, G. Friesecke, and C. Klüppelberg. Density Functional Theory and Optimal Transportation with Coulomb Cost. *Communications on Pure and Applied Mathematics*, 66(4):548–599, 2013.

[26] C. Cotar, G. Friesecke, and C. Klüppelberg. Smoothing of Transport Plans with Fixed Marginals and Rigorous Semiclassical Limit of the Hohenberg–Kohn Functional. *Archive for Rational Mechanics and Analysis*, 228(3):891–922, 2018.

[27] J. Csima. Multidimensional Stochastic Matrices and Patterns. *Journal of Algebra*, 14(2):194–202, 2 1970.

[28] M. Cuturi. Sinkhorn Distances: Lightspeed Computation of Optimal Transport. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.

[29] G. B. Dantzig and P. Wolfe. Decomposition Principle for Linear Programs. *Operations Research*, 8(1):101–111, 1960.

[30] G. B. Dantzig and P. Wolfe. The Decomposition Algorithm for Linear Programs. *Econometrica*, 29(4):767–778, 1961.

[31] B. de Finetti. Funzione caratteristica di un fenomeno aleatorio. *Atti della R. Accademia Nazionale dei Lincei, Ser. 6. Memorie, Classe di Scienze Fisiche, Matematiche e Naturali*, 4:251–299, 1931.

[32] C. Dellacherie and P.-A. Meyer. *Probabilities and Potential B: Theory of Martingales*, volume 72 of *North-Holland Mathematics Studies*. North-Holland, 1982.

[33] W. E. Deming and F. F. Stephan. On a Least Squares Adjustment of a Sampled Frequency Table When the Expected Marginal Totals are Known. *The Annals of Mathematical Statistics*, 11(4):427 – 444, 1940.

[34] J. Desrosiers and M. E. Lübbecke. A Primer in Column Generation. In *Column Generation*, pages 1–32. Springer US, Boston, MA, 2005.

[35] P. Diaconis. Finite forms of de Finetti's theorem on exchangeability. *Synthese*, 36(2):271–281, 1977.

[36] P. Diaconis and D. Freedman. Finite Exchangeable Sequences. *The Annals of Probability*, 8(4):745 – 764, 1980.

[37] O. du Merle, D. Villeneuve, J. Desrosiers, and P. Hansen. Stabilized column generation. *Discrete Mathematics*, 194:229–237, 1999.

[38] L. E. Dubins and D. A. Freedman. Exchangeable processes need not be mixtures of independent, identically distributed random variables. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 48(2):115–132, 1979.

[39] M. Fichtl. Multi-Marginal Optimal Transportation and Many-Electron Systems. Technische Universität München, Department of Mathematics. Bachelor's Thesis. 9 2014.

[40] G. Friesecke. Lectures on Optimal Transport. Lecture Notes. TUM - Faculty of Mathematics, to appear in SIAM, 2019.

[41] G. Friesecke. A Simple Counterexample to the Monge Ansatz in Multimarginal Optimal Transport, Convex Geometry of the Set of Kantorovich Plans, and the Frenkel–Kontorova Model. *SIAM Journal on Mathematical Analysis*, 51(6):4332–4355, 2019.

[42] G. Friesecke, A. Gerolin, and P. Gori-Giorgi. The strong-interaction limit of density functional theory. *arXiv - preprint*, 2022.

[43] G. Friesecke, C. B. Mendl, B. Pass, C. Cotar, and C. Klüppelberg. N-density representability and the optimal transport limit of the Hohenberg-Kohn functional. *The Journal of Chemical Physics*, 139:164109, 2013.

[44] G. Friesecke and D. Vögler. Breaking the Curse of Dimension in Multi-Marginal Kantorovich Optimal Transport on Finite State Spaces. *SIAM Journal on Mathematical Analysis*, 50(4):3996–4019, 2018.

[45] W. Gangbo and R. J. McCann. The geometry of optimal transportation. *Acta Math*, 177:113–161, 1996.

[46] W. Gangbo and A. Święch. Optimal Maps for the Multidimensional Monge-Kantorovich Problem. *Communications on Pure and Applied Mathematics*, 51(1):23–45, 1998.

[47] M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman & Co., USA, 1979.

[48] P. C. Gilmore and R. E. Gomory. A Linear Programming Approach to the Cutting-Stock Problem. *Operations Research*, 9(6):849–859, 1961.

[49] D. Heath and W. Sudderth. De Finetti's Theorem on Exchangeable Variables. *The American Statistician*, 30(4):188–189, 1976.

[50] E. S. Hewitt and L. J. Savage. Symmetric measures on Cartesian products. *Transactions of the American Mathematical Society*, 80:470–501, 1955.

[51] W. W. Hogan, R. E. Marsten, and J. W. Blankenship. The BOXSTEP Method for Large Scale Optimization. *Operations Research*, 23, 06 1975.

[52] P. Hohenberg and W. Kohn. Inhomogeneous Electron Gas. *Phys. Rev.*, 136:B864–B871, 1964.

[53] L. Hörmander. *Notions of Convexity*. Birkhäuser Boston, 1 edition, 1994.

[54] S. Janson, T. Konstantopoulos, and L. Yuan. On a representation theorem for finitely exchangeable random vectors. *Journal of Mathematical Analysis and Applications*, 442(2):703–714, 2016.

[55] E. T. Jaynes. Some Applications and Extensions of the de Finetti Representation Theorem. *Bayesian Inference and Decision Techniques with Applications: Essays in Honor of Bruno de Finetti*, pages 31–50, 5 1986.

[56] L. V. Kantorovich. On the translocation of masses. *Dokl. Akad. Nauk. USSR*, 37:199–201, 1942. English translation in J. Math. Sci. 133, 4 (2006), 1381–1382.

[57] H. Kellerer, U. Pferschy, and D. Pisinger. *Knapsack Problems*. Springer-Verlag, Berlin Heidelberg, 1 edition, 2004.

[58] H. G. Kellerer. Duality Theorems for Marginal Problems. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 67(4):399–432, 1984.

[59] G. J. Kerns and G. J. Székely. Definetti's Theorem for Abstract Finite Exchangeable Sequences. *Journal of Theoretical Probability*, 19(3):589–608, 2006.

[60] Y. Khoo, L. Lin, M. Lindsey, and L. Ying. Semidefinite Relaxation of Multimarginal Optimal Transport for Strictly Correlated Electrons in Second Quantization. *SIAM Journal on Scientific Computing*, 42(6):B1462–B1489, 2020.

[61] Y. Khoo and L. Ying. Convex Relaxation Approaches for Strictly Correlated Density Functional Theory. *SIAM Journal on Scientific Computing*, 41:B773–B795, 01 2019.

[62] V. M. Kravtsov. Combinatorial properties of noninteger vertices of a polytope in a three-index axial assignment problem. *Cybernetics and Systems Analysis*, 1:25–33, 2007.

[63] M. Levy. Universal variational functionals of electron densities, first-order density matrices, and natural spin-orbitals and solution of the v-representability problem. *Proceedings of the National Academy of Sciences*, 76(12):6062–6065, 1979.

[64] E. H. Lieb. Density Functionals for Coulomb Systems. *International Journal of Quantum Chemistry*, 24(3):243–277, 1983.

[65] N. Linial and Z. Luria. On the Vertices of the d-Dimensional Birkhoff Polytope. *Discrete & Computational Geometry*, 51(1):161–170, 2014.

[66] M. E. Lübbecke and J. Desrosiers. Selected Topics in Column Generation. *Operations Research*, 53:1007–1023, 12 2005.

[67] R. E. Marsten. The Use of the Box Step Method in Discrete Optimization. *NBER Working Paper No. w0086*, 05 1975.

[68] G. Monge. Mémoire sur la théorie des déblais et des remblais. *Histoire de l'Académie Royale des Sciences de Paris*, page 666–704, 1781.

[69] L. Nenna. Numerical Methods for Multi-Marginal Optimal Transportation. Université Paris sciences et lettres, PhD Thesis, 2 2017.

[70] R. G. Parr and W. Yang. *Density-Functional Theory of Atoms and Molecules*. Oxford University Press, 1995.

[71] B. Pass. Remarks on the semi-classical Hohenberg-Kohn functional. *Nonlinearity*, 26(9):2731–2744, 2013.

[72] B. Pass. Multi-marginal optimal transport: Theory and applications. *ESAIM: M2AN*, 49(6):1771–1790, 2015.

[73] G. Peyré and M. Cuturi. *Computational Optimal Transport*, volume 11. 2019.

[74] S. T. Rachev and L. Rüschendorf. *Mass Transportation Problems - Volume I: Theory*, volume 1 of *Probability and its Applications*. Springer International Publishing, 1998.

[75] S. T. Rachev and L. Rüschendorf. *Mass Transportation Problems - Volume II: Applications*, volume 2 of *Probability and its Applications*. Springer International Publishing, 1998.

[76] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, reprint edition, 1997.

[77] F. Santambrogio. *Optimal Transport for Applied Mathematicians*, volume 87 of *Progress in Nonlinear Differential Equations and their Applications*. Springer International Publishing, 2015.

[78] M. Seidl. Strong-interaction limit of density-functional theory. *Physical Review A*, 60:4387–4395, 12 1999.

[79] M. Seidl, P. Gori-Giorgi, and A. Savin. Strictly correlated electrons in density-functional theory: A general formulation with applications to spherical densities. *Physical Review A*, 75:042511, 4 2007.

[80] M. Seidl, J. P. Perdew, and M. Levy. Strictly correlated electrons in density-functional theory. *Physical Review A*, 59:51–54, 1 1999.

[81] R. Sinkhorn. A Relationship Between Arbitrary Positive Matrices and Doubly Stochastic Matrices. *The Annals of Mathematical Statistics*, 35(2):876 – 879, 1964.

[82] C. Villani. *Topics in Optimal Transportation*, volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, 2003.

[83] C. Villani. *Optimal Transport: Old and New*, volume 338 of *Grundlehren der mathematischen Wissenschaften*. Springer, 2008.

[84] J. von Neumann. A Certain Zero-sum Two-person Game Equivalent to the Optimal Assignment Problem. *Contributions to the Theory of Games (AM-28)*, 2:5–12, 1953.

A Core publications

A.1 Geometry of Kantorovich polytopes and support of optimizers for repulsive multi-marginal optimal transport on finite state spaces

Author: Daniela Vögler

Summary. In the present article, we shine a light on the convex geometry of the polytopes

$$\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right) := \left\{\gamma \in \mathcal{P}_{\mathrm{sym}}\left(X^N\right) : \quad \gamma \mapsto \overline{\lambda}\right\}$$

and

$$\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right) := \left\{\mu \in M_2\mathcal{P}_{\mathrm{sym}}\left(X^N\right) : \quad \mu \mapsto \overline{\lambda}\right\}.$$

Here, X is a finite state space consisting of ℓ pairwise distinct elements $a_1, a_2, \ldots, a_\ell$; $\overline{\lambda}$ denotes the uniform probability measure on X; $\mathcal{P}_{\mathrm{sym}}\left(X^N\right)$ consists of all symmetric probability measures on the product space X^N; $M_2 : \mathcal{P}_{\mathrm{sym}}\left(X^N\right) \to \mathcal{P}_{\mathrm{sym}}\left(X^2\right)$ denotes the two-point marginal map defined by

$$M_2\gamma\left(A\right) = \gamma\left(A \times X^{N-2}\right) \quad \text{for } A \subset X^2.$$

Finally, $\cdot \mapsto \overline{\lambda}$ means that $\cdot$ has equal one-point marginals $\overline{\lambda}$. These sets arise as optimization domain for multi-marginal optimal transport problems on the finite state space X that enforce a uniform marginal in every component: $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ arises for the restriction to symmetric cost functions, $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ arises for the restriction to symmetric two-body interactions. Our work is a natural extension of the previous consideration of the $N = \ell = 3$-case by Friesecke in [41].

Overall, we work towards a better understanding of the insufficiency exhibited by Monge's ansatz, the classic sparse ansatz in optimal transport. First, we focus on small problem-parameters N and ℓ, compute all extreme points of $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ and $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$, and determine how many of the computed extreme points are of the form prescribed by Monge's ansatz. We complement this computational study with bite-sized theoretical results. Next, we only admit $\ell = 3$ discretization points as elements of the finite state space X. This makes $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ three-dimensional and allows us to visualize the role Monge's ansatz takes on. Despite its severe restrictions, this model problem bore fruit: it inspired a necessary support-condition

on optimizers that applies to general $N, \ell \geq 2$ and symmetric two-body interactions whose diagonal entries exceed a specified threshold. This threshold varies from problem to problem, from cost function to cost function. For

$$N = k \cdot \ell \text{ for some } k \in \mathbb{N} \text{ or } \left(N - \left\lfloor \frac{N}{\ell} \right\rfloor \cdot \ell \right) \in \{\ell - 1, 1\}\,,$$

the necessary support-condition proved to be strong enough to guarantee the existence of a unique optimizer and not only imply its compliance with Monge's ansatz but also explicitly provide this optimizer.

Own contribution. I, Daniela Vögler, am sole and, consequently, principal author of the present article. Under the guidance of my advisor, Prof. Gero Friesecke, I conducted the research, put pen to paper, and got the article published. While Prof. Gero Friesecke oversaw the process, the responsibility rested with me.

The publisher of the *Journal of Mathematical Analysis and Applications*, Elsevier, permits the inclusion of this publication:

> *"Author rights*
> *(...) Use and share their works for scholarly purposes (with full acknowledgement of the original article):*
> *(...) Include in a thesis or dissertation (provided this is not published commercially)"*
>
> [Elsevier. Copyright. https://www.elsevier.com/about/policies/copyright, 2022. [Online; accessed 16.05.2022].]

J. Math. Anal. Appl. 502 (2021) 125147

Contents lists available at ScienceDirect

Journal of Mathematical Analysis and Applications

www.elsevier.com/locate/jmaa

Geometry of Kantorovich polytopes and support of optimizers for repulsive multi-marginal optimal transport on finite state spaces

Daniela Vögler

Zentrum Mathematik, Technische Universität München, Boltzmannstr. 3, 85748 Garching bei München, Germany

ARTICLE INFO

Article history:
Received 19 September 2020
Available online 18 March 2021
Submitted by G. Carlier

Keywords:
Optimal transport
Monge's ansatz
N-representability
Birkhoff-von Neumann theorem
Density functional theory
Support-condition for optimizers

ABSTRACT

We consider symmetric multi-marginal Kantorovich optimal transport problems on finite state spaces with uniform-marginal constraint. Hereby the symmetry of the problem refers to an assumption on the cost function as well as a corresponding restriction of the set of admissible trial states where the former enables the latter. Note that the symmetry of this setting forces us to pick for each of the considered marginals one and the same probability measure. The said problems consist of minimizing a linear objective function over a high-dimensional polytope, here referred to as Kantorovich polytope. The presented results are of split nature, computational and theoretical. Within the computational part only small numbers of marginals N and marginal sites ℓ are considered. This restriction allows us to computationally determine all extreme points of the Kantorovich polytope and investigate how many of them are in compliance with the in optimal transport typical Monge ansatz. Singling out the results for $\ell = 3$ discretization points and pairwise symmetric cost functions enables us to visually compare Kantorovich's to Monge's ansatz space for a varying number of marginals. Finally we present a necessary support-condition for optimizers which is inspired by the insights the said model problem on three sites provided. This result is not limited to the case of $\ell = 3$ sites and applies to symmetric pair-costs whose diagonal entries lie above a cost-specific threshold. In case N and ℓ display certain relationships the discussed condition provides an optimizer in Monge-form and implies its uniqueness as a solution of the considered Kantorovich optimal transport problem.

Acknowledgments: The author thanks Gero Friesecke and Sören Behr for helpful discussions.

1. Introduction

In general multi-marginal Kantorovich optimal transport (OT) problems aim at coupling N probability measures $\lambda^{(1)}, \ldots, \lambda^{(N)}$ optimally with respect to a given cost function c (see (1.1) for a discrete symmetric

E-mail address: voegler@ma.tum.de.

https://doi.org/10.1016/j.jmaa.2021.125147
0022-247X/

OT problem). These problems arise in various fields of research, ranging from economics [8,11] through mathematical finance [2,21] and image processing [1,36] to electronic structure [14,6].

Here we consider a symmetric multi-marginal Kantorovich OT problem on finite state spaces given by

$$\text{Minimize} \int_{X^N} c(x_1,\dots,x_N)d\gamma(x_1,\dots,x_N) \text{ over } \gamma \in \mathcal{P}_{\text{sym}}\left(X^N\right) \text{ subject to } \gamma \mapsto \overline{\lambda}. \tag{1.1}$$

X denotes a finite state space as defined in (1.2), $c : X^N \to \mathbb{R}\cup\{+\infty\}$ an arbitrary symmetric cost function, $\overline{\lambda}$ the uniform marginal as defined in (1.3) and $\mathcal{P}_{\text{sym}}\left(X^N\right)$ the set of symmetric probability measures on X^N, where a probability measure γ on X^N is symmetric if

$$\gamma\left(A_1 \times \cdots \times A_N\right) = \gamma\left(A_{\sigma(1)} \times \cdots \times A_{\sigma(N)}\right) \text{ for all subsets } A_1,\dots,A_N \text{ of } X \text{ and all permutations } \sigma.$$

Any $\gamma \in \mathcal{P}_{\text{sym}}\left(X^N\right)$ fulfills $\gamma \mapsto \overline{\lambda}$ if and only if γ has equal one-point marginals $\overline{\lambda}$, i.e.,

$$\gamma\left(X^{k-1} \times A_k \times X^{N-k}\right) = \overline{\lambda}(A_k) \text{ for all subsets } A_k \text{ of } X \text{ and all } k \in \{1,\dots,N\}.$$

Multi-marginal OT problems of form (1.1) were already considered in [20] and [18]. While [18] discusses the validity of Monge's approach in the setting of 3 marginals and 3 sites, [20] introduces a sufficient ansatz space for problem (1.1) (see Section 2 as well as Remark 3.5 for information about the content of these papers). The present paper accompanies these previous considerations. In particular, some of the used nomenclature and notation is already introduced there.

Readers that are interested in computational methods regarding multi-marginal optimal transport are advised to have a look at [28,27]: Therein semidefinite relaxation approaches that are based on results in [20] are discussed.

For finite state spaces

$$X = \{a_1,\dots,a_\ell\} \tag{1.2}$$

consisting of ℓ distinct points $a_1,\dots,a_\ell$, the uniform probability measure

$$\overline{\lambda} = \sum_{i=1}^{\ell} \frac{1}{\ell}\delta_{a_i} \tag{1.3}$$

on X is the prototypical marginal. The corresponding multi-marginal Kantorovich OT problems, i.e., problems of form (1.1) with $\mathcal{P}\left(X^N\right)$ replacing $\mathcal{P}_{\text{sym}}\left(X^N\right)$, appear directly as assignment problems (see [39,5] for reviews) and arise from continuous problems via equi-mass discretization [10].

Note that the restriction to symmetric probability measures in problem (1.1) is motivated by a physical application. Modeling the electronic structure of a molecule with N electrons in a discretized setting, is a prototypical application of multi-marginal OT on finite state spaces. In this context, X corresponds to a set of ℓ discretization points in $\mathbb{R}^3$ and any coupling γ of the N marginals $\overline{\lambda},\dots,\overline{\lambda}$ describes a joint probability distribution regarding the electron positions in an N-electron molecule. Then the marginal condition ensures that each discretization point is occupied equally often and the cost function $c : X^N \to \mathbb{R}\cup\{+\infty\}$ embodies the electron interaction energy. As electrons are indistinguishable the considered cost functions are usually symmetric, i.e., invariant under argument permutation. These symmetric cost functions are 'dual' to the set of symmetric probability measures on the product space X^N in the following sense: There always exists an optimal coupling of $\overline{\lambda},\dots,\overline{\lambda}$ that is symmetric.

The interaction energy between electrons displays a pairwise structure, i.e., $c(x_1,\dots,x_N) = \sum_{1\le i<j\le N} v(x_i,x_j)$, with the Coulomb cost $\sum_{1\le i<j\le N} \frac{1}{|x_i-x_j|}$ being the prototypical example. Here $|\cdot|$ is

the Euclidean norm in $\mathbb{R}^d$. As discussed in Section 3, this pairwise structure allows us to reformulate the multi-marginal OT problem (1.1) as

$$\text{Minimize} \int_{X^2} v(x,y)d\mu(x,y) \text{ over } \mu \in \mathcal{P}_{N\text{-rep}}\left(X^2\right) \text{ subject to } \mu \mapsto \overline{\lambda}. \tag{1.4}$$

This reformulated problem was initially introduced in [19]. In (1.4), $\mathcal{P}_{N\text{-rep}}\left(X^2\right) \subseteq \mathcal{P}\left(X^2\right)$ can be interpreted as a 'reduced version' of $\mathcal{P}_{\text{sym}}\left(X^N\right)$.

The goal of this paper is to present new insights into problem (1.1) which arose out of a careful consideration of the polytope formed by the admissible trial states. Within this consideration the role Monge states take on in the said polytope is investigated.

One of the central questions in the theory of optimal transportation is: Under which assumptions exists an optimal coupling that is supported on a graph (over the first variable)? Such optimizers are then called Monge-solutions (see (1.6) - (1.8)). In the case of two marginals this question is well understood; the existence of Monge-solutions is always respectively under very general conditions guaranteed (see the renowned Birkhoff-von Neumann theorem [4,41] regarding finite state spaces respectively, e.g., [40] regarding continuous state spaces). For multiple marginals the understanding of this question does not reach the same generality. However, there are isolated examples for Monge- and non-Monge-solutions. For the former, see [24,7,32,14,6,12] as well as the fundamental paper by Gangbo and Święch [22] for an interesting selection. For the latter, we refer the interested reader to [9,33,19,34,15,17,13,38,31,23,18] regarding continuous state spaces as well as to [16,29,30,18] regarding finite state spaces.

In order to understand Monge's approach in the present setting, we first take a quick glance at the 'unsymmetrized' OT problem, i.e., problems of form (1.1) with $\mathcal{P}\left(X^N\right)$ replacing $\mathcal{P}_{\text{sym}}\left(X^N\right)$. Then, an optimal coupling $\gamma \in \mathcal{P}\left(X^N\right)$ of the N marginals $\overline{\lambda}, \dots, \overline{\lambda}$ is a Monge-solution if

$$\gamma = \sum_{\nu=1}^{\ell} \frac{1}{\ell} \delta_{T_1(a_\nu)} \otimes \cdots \otimes \delta_{T_N(a_\nu)} \text{ for } N \text{ permutations } T_1, \dots, T_N : X \to X, \tag{1.5}$$

where $T : X \to X$ is a permutation if there exists a permutation of indices $\tau : \{1, \dots, \ell\} \to \{1, \dots, \ell\}$ such that $T(a_\nu) = a_{\tau(\nu)}$ for all $\nu \in \{1, \dots, \ell\}$. Demanding that the T_ks are permutations ensures that γ is indeed a coupling of $\overline{\lambda}, \dots, \overline{\lambda}$: $T : X \to X$ is a permutation if and only if it pushes the uniform measure forward to itself, i.e., $T_{\#}\overline{\lambda} = \overline{\lambda}$. Here for any probability measure $\lambda = \sum_{\nu=1}^{\ell} \lambda_\nu \delta_{a_\nu}$ on X and any map $T : X \to X$ the push-forward $T_{\#}\lambda$ of λ under T is defined by $T_{\#}\lambda = \sum_{\nu=1}^{\ell} \lambda_\nu \delta_{T(a_\nu)}$. One may choose $T_1 = \text{id}$, i.e., $T_1(a) = a$ for all $a \in X$, by re-ordering the sum in (1.5).

Regarding (1.1) an admissible trial state $\hat{\gamma}$ is referred to as a (symmetrized) Monge state if it is the symmetrization of a probability measure γ of form (1.5), i.e.,

$$\hat{\gamma} = \sum_{\nu=1}^{\ell} \frac{1}{\ell} S\delta_{T_1(a_\nu)} \otimes \cdots \otimes \delta_{T_N(a_\nu)} \tag{1.6}$$

such that

$$T_k\#\overline{\lambda} = \overline{\lambda} \text{ for all } k \in \{1, \dots, N\}, \tag{1.7}$$

or equivalently,

$$T_1, \dots, T_N : X \to X \text{ are permutations.} \tag{1.8}$$

Here S denotes the linear symmetrization operator in N variables as defined in (2.4). Probability measures on X^N of form (1.6)-(1.8) are also said to be of Monge-type or in Monge-form and restricting the minimization problem (1.1) to such measures yields the corresponding Monge problem.

In Section 2 the 'Kantorovich-coupling-polytope', $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$, as defined in (2.2), will be identified with the 'coefficient-polytope' P_{coef}, as introduced in (2.9). Monge states then correspond to integer elements of the latter of the two polytopes. Both, the reformulation as well as the identification of Monge states with integer coefficients are based on results in [20].

This different view on the set of admissible trial states of problem (1.1) makes Monge's approach more accessible, in the sense of, it is easier to decide whether a given coefficient vector in P_{coef} embodies a Monge state or not. This grasp of the Monge concept allows us to numerically quantify the insufficiency of Monge's ansatz which in the present setting is established in [18]. In more detail, for small problem-parameters we determine all extremal coefficients and partition them into a Monge and a non-Monge class. The mere results of this classification are interpreted and complemented by 'small' theoretical results building upon the numerical ones.

In Section 3 the same numerical analysis is performed under the additional assumption of the cost function c in (1.1) displaying pairwise structure.

In Section 4, we will consider a model problem of optimally coupling the N marginals $\overline{\lambda},\dots,\overline{\lambda}$ with respect to a cost function of pairwise symmetric structure, where the finite state space X consists only of three states, i.e., $X=\{a_1,a_2,a_3\}$. In particular $\overline{\lambda}$ is the uniform probability measure on $\{a_1,a_2,a_3\}$ and the domain of the cost function is given by $\{a_1,a_2,a_3\}^N$. The present setting allows us to draw a visual comparison between Kantorovich's and Monge's ansatz as depicted in Fig. 1. We further compare both OT approaches by volume of (the convex hull of) the respective set of admissible trial states and establish a computationally simple upper bound on the optimal value in (1.4).

Taking a look at Fig. 1 the reader may notice that in each of the illustrations one of the extreme points creates a 'peak' in the front of the polytope. As indicated by the coloring each of these extreme points is of Monge-type. In Section 4, more specifically Theorem 4.1, these 'peaks' are identified as the unique solutions of OT problems with respect to very rudimentary repulsive cost functions based on the discrete metric. The results of Section 5 originated from the idea 'peaks solve repulsive OT problems'. For any number of marginals $N\geq 2$ and any number of states $\ell\geq 2$ we consider a large class of repulsive pair-costs with the underlying assumption on these cost functions being that their diagonal entries are constant and 'big enough' when compared to their off-diagonal entries. We show that in the given setting any optimizer of problem (1.1) only gives mass to tupels $(x_1,\dots,x_N)\in X^N$ for which the 'appearance-frequency' of elements of the finite state space X is as uniform as possible (see Theorem 5.2). So for $N=5$ and $\ell=3$ (a_1,a_2,a_1,a_3,a_2) is a valid support-tupel whereas (a_1,a_2,a_1,a_3,a_1) is not. This support-condition is more explicit than the notion of c-cyclical monotonicity, which appears frequently in the context of OT, and for certain parameter constellations it paints a very thorough picture of what optimizers look like. In case

$$N=k\cdot\ell \text{ for some } k\in\mathbb{N} \text{ or } \left(N-\left\lfloor\frac{N}{\ell}\right\rfloor\cdot\ell\right)\in\{\ell-1,1\} \tag{1.9}$$

the discussed condition provides an optimizer in Monge-form and implies its uniqueness. This optimizer represents the depicted 'peak' for $\ell=3$ and a higher-dimensional analogue for $\ell>3$. Note that for $\ell=3$ states (1.9) is fulfilled for any $N\geq 2$. So to put it in a nutshell, for any pair of parameters $N,\ell\geq 2$ fulfilling (1.9) we provide a large class of repulsive pair-costs which yield the idea 'peaks solve repulsive OT problems' to be true. As for why this identification of 'peaks' as optimizers is limited to parameter choices N,ℓ fulfilling (1.9) note the following. When keeping the number of states $\ell>3$ fixed and letting N tend to $+\infty$ starting at $N=\ell$ the geometric behavior can be described as follows: While the 'peak' remains intact for $N=\ell$ as well as $N=\ell+1$, for $N=\ell+2$ its representing measure (see Definition 3.1) on the product

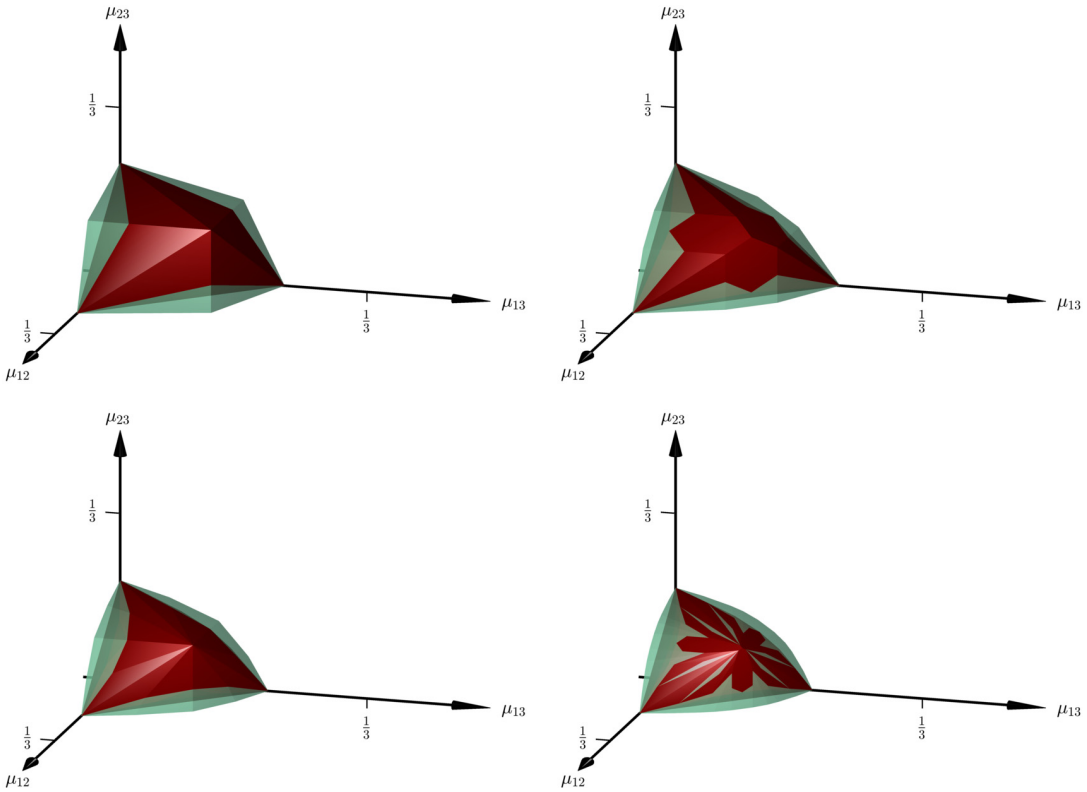

Fig. 1. The reduced Kantorovich (see (3.4)) respectively Monge (see (3.8)) polytope for N marginals and 3 states is visualized for $N = 3$ (top-left), 4 (top-right), 6 (bottom-left) and 10 (bottom-right) in green respectively red. The elements $(\mu_{ij})_{i,j=1}^{3}$ of the polytopes are parametrized by their off-diagonal entries μ_{12}, μ_{13} and μ_{23}. The present figure recovers the illustration of the reduced polytopes in the case of three marginals in [18]. If a two-dimensional face of the reduced Monge polytope belongs to the boundary of the reduced Kantorovich polytope the occupied area is depicted in red. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

space X^N blossoms into multiple extreme points; for $N = 2\ell - 1$ the blossom retracts into a closed state recreating the 'peak'.

In order to prove Theorem 5.2, we establish a lower bound on the nonzero entries of extreme points of the 'coefficient-polytope' P_{coef}, which - we believe - is in itself an interesting result.

2. Classification of the extreme points of a Kantorovich polytope

Throughout the paper we will consider the finite state space X given by (1.2). We will denote the set of probability measures on X as $\mathcal{P}(X)$. Each such probability measure $\lambda \in \mathcal{P}(X)$ can be canonically identified with a vector in $\mathbb{R}^\ell$ via $\lambda_i := \lambda(a_i)$. The vector $(\lambda_1, \dots, \lambda_\ell)$ fulfills $\sum_{i=1}^{\ell} \lambda_i = 1$ and $\lambda_i \geq 0$ for $i \in \{1, \dots, \ell\}$ and is therefore an element of the unit simplex. The probability measure λ can then be written as $\lambda = \sum_{i=1}^{\ell} \lambda_i \delta_i$, where here and below we use the shorthand notations

$$\delta_i := \delta_{a_i}, \qquad \delta_{i_1 \dots i_N} := \delta_{a_{i_1}} \otimes \dots \otimes \delta_{a_{i_N}} \tag{2.1}$$

for $a_i \in X$ being a single point in the finite state space and $(a_{i_1}, \dots, a_{i_N})$ being an element of the product space X^N.

As announced in the Introduction, we will now take a closer look at the set of admissible trial states of problem (1.1), i.e., the polytope

$$\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right) := \left\{\gamma \in \mathcal{P}_{\text{sym}}\left(X^N\right) : \gamma \mapsto \overline{\lambda}\right\}. \tag{2.2}$$

From this point on, we will refer to the elements of this set as *symmetric Kantorovich couplings*. The set $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ itself will be called *(symmetric) Kantorovich polytope for N marginals and ℓ states*. As within this paper we focus our attention on the symmetric case, the term symmetric will be dropped from time to time. It is easy to see that, as a result of the linearity of the marginal constraint and the finiteness of the state space X, $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ is a compact and convex set in $\mathbb{R}^{\ell^N}$ and therefore by Minkowski's theorem (see, e.g., [25]) the convex hull of its extreme points.

Recall the following basic definitions and notions of convexity (see, e.g., [25,37]). For $y_1,\dots,y_n\in\mathbb{R}^m$ and $\lambda_1,\dots,\lambda_n\geq 0$ such that $\sum_{i=1}^n\lambda_i=1$

$$\lambda_1 y_1+\dots+\lambda_n y_n=\sum_{i=1}^n\lambda_i y_i$$

is called a *convex combination of the points* $y_1,\dots,y_n$. A subset $K\subseteq\mathbb{R}^m$ is called *convex* if for each finite selection of points in K each possible convex combination of these points is again contained in K. For a subset $V\subseteq\mathbb{R}^m$ the *convex hull of* V, denoted as $\mathrm{conv}(V)$, corresponds to the set of all possible convex combinations of a finite selection of points in V. Obviously a set $K\subseteq\mathbb{R}^m$ is convex if and only if it is equal to its convex hull. Finally an element k of the convex set $K\subseteq\mathbb{R}^m$ is called an *extreme point* if $k=\lambda_1 y_1+\lambda_2 y_2$ for some $y_1,y_2\in K$ and $\lambda_1,\lambda_2>0$ such that $\lambda_1+\lambda_2=1$ implies $y_1=k=y_2$. For a considered convex set K the set of extreme points will from now on be denoted as $\mathrm{ext}(K)$.

As $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ is equal to the convex hull of its extreme points, we can use the extreme points to describe the convex structure of the set of symmetric Kantorovich couplings. Now it follows by a simple contradiction argument that for any given linear objective function there is always an optimizer that is an extreme point. Moreover, in our setting of finite states spaces, for any extreme point γ^* there is a function $c:X^N\to\mathbb{R}$ such that

$$\int_{X^N}c(x_1,\dots,x_N)d\gamma(x_1,\dots,x_N)>\int_{X^N}c(x_1,\dots,x_N)d\gamma^*(x_1,\dots,x_N)\text{ for any }\gamma\in\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right),\tag{2.3}$$

i.e., there is a cost function such that γ^* is the unique optimizer of the corresponding OT problem. This is a result of the fact that $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ is a bounded polyhedron, i.e., a polytope, of finite dimension and therefore only possesses finitely many extreme points each of whom is itself an exposed point (see, e.g., [37]), i.e., a point in the set $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ that fulfills (2.3) for some cost function $c:X^N\to\mathbb{R}$. As for any cost function $c:X^N\to\mathbb{R}$ there is always an optimizer that is an extreme point of $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ and vice versa for any extreme point γ^* of $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ there is a cost function $c:X^N\to\mathbb{R}$ such that γ^* is the unique optimizer, analyzing, how many of the extreme points of $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ are of Monge-type, is a good approach to investigate the validity of Monge's ansatz. Recall that in the given setting a probability measure $\gamma\in\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ is said to be of Monge-type or in Monge-form if there are N permutations $\tau_1,\dots,\tau_N:\{1,\dots,\ell\}\to\{1,\dots,\ell\}$ such that

$$\gamma=\sum_{i=1}^{\ell}\frac{1}{\ell}S\delta_{\tau_1(i)\tau_2(i)\dots\tau_N(i)},$$

where the *symmetrization operator* $S:\mathcal{P}\left(X^N\right)\to\mathcal{P}_{\mathrm{sym}}\left(X^N\right)$ is defined by

$$(S\gamma)\left(A_1\times\dots\times A_N\right)=\frac{1}{N!}\sum_{\sigma\in S_N}\gamma\left(A_{\sigma(1)}\times\dots\times A_{\sigma(N)}\right)\text{ for all }A_1,\dots,A_N\subseteq X\tag{2.4}$$

with S_N being the group of all permutations on the set $\{1,\dots,N\}$.

As in the given setting it is rather inconvenient and long-winded to check whether a given symmetric Kantorovich coupling is of Monge-type or not, we will derive in the following an alternative LP-formulation of problem (1.1), where Monge-states will correspond exactly to (rescaled) integer points in the corresponding polytope of admissible trial states. We start by taking a closer look at the convex geometry of the set of symmetric probability measures on the product space X^N.

Note that a probability measure $\gamma \in \mathcal{P}_{\text{sym}}\left(X^N\right)$ is an extreme point of $\mathcal{P}_{\text{sym}}\left(X^N\right)$ if and only if it is of the form

$$S\delta_{i_1\dots i_N} \text{ for some } 1 \leq i_1 \leq \dots \leq i_N \leq \ell \tag{2.5}$$

(see [20]). Therefore symmetric Kantorovich couplings, which are of Monge-type, are an average of ℓ not necessarily distinct extremal symmetric probability measures on the product space X^N with respect to the uniform measure. Below we will elaborate further on this characterization of couplings in Monge-form, which will be the basis for identifying Monge-states with the (rescaled) integer points in a certain polytope.

From now on, we will denote the set of extremal symmetric probability measures, i.e., measures of the form (2.5), as E^N_{sym}. It was shown in [20] that E^N_{sym} contains $\binom{N+\ell-1}{N}$ elements.

As for each pair of these extreme points their support is disjoint, one can immediately deduce the following result.

Proposition 2.1. *$\mathcal{P}_{sym}\left(X^N\right)$ is a simplex, i.e., the extremal symmetric probability measures on X^N are affinely independent.*

Hence, for every $\gamma \in \mathcal{P}_{\text{sym}}\left(X^N\right)$ there is a unique way to represent γ as a convex combination of extremal symmetric probability measures on X^N, i.e., there is a unique non-negative coefficient vector $\alpha \in \mathbb{R}^{|E^N_{\text{sym}}|}$ fulfilling $\sum \alpha_{i_1\dots i_N} = 1$ such that

$$\gamma = \sum_{1\leq i_1\leq\dots\leq i_N\leq\ell} \alpha_{i_1\dots i_N} S\delta_{i_1\dots i_N}. \tag{2.6}$$

As the extreme points of $\mathcal{P}_{\text{sym}}\left(X^N\right)$ can be parametrized using their one-point marginal, α can be interpreted as a probability measure on the set of these one-point marginals.

Given the k-point marginal map $M_k : \mathcal{P}\left(X^N\right) \to \mathcal{P}\left(X^k\right)$ for $1 \leq k \leq N-1$ with

$$\left(M_k\gamma\right)(A) := \gamma\left(A \times X^{N-k}\right) \text{ for all } A \subseteq X^k \tag{2.7}$$

for $\gamma \in \mathcal{P}\left(X^N\right)$, with the convention $M_N = id$, note that M_1 is a bijection from the set of extremal symmetric probability measures on X^N, i.e., measures of the form (2.5), to the set of $\frac{1}{N}$-quantized probability measures

$$\mathcal{P}_{\frac{1}{N}}(X) := \left\{\lambda \in \mathcal{P}(X) : \lambda(\{i\}) \in \left\{0, \frac{1}{N}, \dots, 1\right\}\right\} \tag{2.8}$$

(see [20]). Hereby the one-point marginal of a measure of form (2.5) is an empirical measure of the indices $(i_1, \dots, i_N)$, it holds

$$M_1 S\delta_{i_1\dots i_N} = \frac{1}{N}\sum_{j=1}^{N} \delta_{i_j}.$$

In the following $\psi_N : \mathcal{P}_{\frac{1}{N}}(X) \to E^N_{\text{sym}}$ will denote the corresponding inverse function.

This parametrization gives rise to the *coefficients-to-coupling map* $R : \mathcal{P}\left(\mathcal{P}_{\frac{1}{N}}(X)\right) \to \mathcal{P}_{\mathrm{sym}}\left(X^N\right)$. It maps an arbitrary probability measure α on $\mathcal{P}_{\frac{1}{N}}(X)$, which, via the underlying parametrization, corresponds to the coefficients in the representation (2.6), to the corresponding coupling γ, i.e., in pedestrian notation

$$R\alpha = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda \psi_N(\lambda),$$

or more elegantly

$$R\alpha = \int_{\mathcal{P}_{\frac{1}{N}}(X)} \psi_N(\lambda) d\alpha(\lambda).$$

As $\mathcal{P}_{\mathrm{sym}}\left(X^N\right)$ is a simplex, R is bijective. This enables us to establish the following isomorphic relationship between two alternative formulations of the set of symmetric Kantorovich couplings.

Lemma 2.2 *(isomorphic relationship between couplings and coefficients). The coefficients-to-coupling map R maps the polytope*

$$P_{coef} := \left\{\alpha \in \mathbb{R}^{|E^N_{sym}|} : A\alpha = \overline{\lambda}, \alpha \geq 0\right\} \tag{2.9}$$

linearly and bijectively to the set of symmetric Kantorovich couplings, i.e., $\mathcal{P}_{sym,\overline{\lambda}}\left(X^N\right)$ defined in (2.2). Here E^N_{sym} is the set of extremal symmetric probability measures on X^N and A is the matrix in $\mathbb{R}^{\ell \times |E^N_{sym}|}$, whose columns are given by the elements of $\mathcal{P}_{\frac{1}{N}}(X)$, i.e.,

$$A := \begin{pmatrix} \lambda_1^{(1)} & \lambda_1^{(2)} & \dots & \lambda_1^{(|E^N_{sym}|)} \\ \vdots & \vdots & & \vdots \\ \lambda_\ell^{(1)} & \lambda_\ell^{(2)} & \dots & \lambda_\ell^{(|E^N_{sym}|)} \end{pmatrix}. \tag{2.10}$$

The corresponding inverse map is also linear.

Proof. Linearity and injectivity of R as a map from P_{coef} to $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ are an immediate consequence of the linearity and injective of $R : \mathcal{P}\left(\mathcal{P}_{\frac{1}{N}}(X)\right) \to \mathcal{P}_{\mathrm{sym}}\left(X^N\right)$ as introduced above. We further know that any $\gamma \in \mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ is an element of $\mathcal{P}_{\mathrm{sym}}\left(X^N\right)$. Hence, applying the parametrization of extremal symmetric probability measures on X^N via their one-point marginals, there exist coefficients $\alpha \in \mathcal{P}\left(\mathcal{P}_{\frac{1}{N}}(X)\right)$, which are non-negative and whose entries sum to 1 such that

$$\gamma = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda \psi_N(\lambda) \tag{2.11}$$

and therefore $\gamma = R\alpha$ holds. Applying the linear marginal map M_1 to (2.11) yields the following.

$$\overline{\lambda} = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda \lambda$$

Therefore α corresponds to an element of P_{coef}. This implies surjectivity of the considered map R. Linearity of the corresponding inverse map is an immediate consequence of the fact that the extremal symmetric probability measures on X^N of the form (2.5) interpreted as vectors are linearly independent. □

Now it is easy to see that the extreme points of P_{coef} correspond exactly to the extremal symmetric Kantorovich couplings, in the sense that R maps the corresponding sets of extreme points bijectively to each other. By standard arguments of polyhedral optimization the extreme points of P_{coef} have a sparse structure, i.e., any extreme point of P_{coef} can have at most ℓ, that is the number of states in the finite state space X, non-zero entries (see, e.g., [3]). In [20] it was shown that this implies that any extremal Kantorovich coupling is a so called *Quasi-Monge state*, i.e., of the form $\sum_{\nu=1}^{\ell} \alpha_\nu S\delta_{T_1(a_\nu)} \otimes \cdots \otimes \delta_{T_N(a_\nu)}$ for N maps $T_1, \dots, T_N : X \to X$ such that $\frac{1}{N}\sum_{k=1}^{N} T_k \# \alpha = \overline{\lambda}$. Here we renounce from using the shorthand notations (2.1) in order to make it easier to draw a comparison with Monge's approach (1.6)-(1.8). The ansatz space of Quasi-Monge states increases the number of unknowns only by $2 \cdot \ell$ compared to the class of symmetrized Monge states and as every extremal Kantorovich coupling is a Quasi-Monge state, this ansatz space always contains an optimal coupling, in contrast to Monge's approach. Note further that obviously every symmetrized Monge state is a Quasi-Monge state. For further reading on this sufficient low-dimensional enlargement of the class of symmetrized Monge states we refer the interested reader to [20]. There also a characterization of Monge states in the given setting was established. A probability measure on the product space X^N is a symmetrized Monge state if and only if it is a Quasi-Monge state all of whose site weights $\alpha_1, \dots, \alpha_\ell$ are equal to $\frac{1}{\ell}$. In summary, we get the following corollary.

Corollary 2.3. *Extremal symmetric Kantorovich couplings correspond exactly, via the coefficients-to-coupling map R, to the extreme points of P_{coef}. Any of these extreme points of P_{coef} is the coefficient vector of a coupling in Monge-form if and only if it is an integer vector scaled by the factor $\frac{1}{\ell}$.*

This corollary gives us a numerically-convenient way to compute the set of extremal Kantorovich couplings and check whether they are of Monge-form or not. In addition we also want to consider Monge's approach by itself. For this purpose we introduce the sets

$$\mathcal{P}_{\text{sym,Monge}}\left(X^N\right) := \left\{\gamma \in \mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right) : \gamma \text{ is of Monge-form (1.6)} - \text{(1.8)}\right\} \tag{2.12}$$

and

$$\mathcal{P}^{\text{conv}}_{\text{sym,Monge}}\left(X^N\right) := \text{conv}\left(\mathcal{P}_{\text{sym,Monge}}\left(X^N\right)\right). \tag{2.13}$$

$\mathcal{P}_{\text{sym,Monge}}\left(X^N\right)$ is the set of all symmetrized Monge states. In the following we will refer to $\mathcal{P}^{\text{conv}}_{\text{sym,Monge}}\left(X^N\right)$ as the *(symmetric) Monge polytope for N marginals and ℓ states*. For simplicity we will once again drop the term symmetric from time to time. Note that if there exists an optimizer of problem (1.1) which is an element of $\mathcal{P}^{\text{conv}}_{\text{sym,Monge}}\left(X^N\right)$ then there exists a Monge-type minimizer.

Having the explanations leading up to Corollary 2.3 in mind, it is easy to see that $\mathcal{P}_{\text{sym,Monge}}\left(X^N\right)$ corresponds to the (scaled by $\frac{1}{\ell}$) integer elements of P_{coef}. These can be for example determined by a simple enumeration of all the ordered choices of $N-1$ permutations interpreted as coefficient vectors in P_{coef}. Checking which of these scaled integer coefficient vectors are extremal with respect to the convex hull of them as a whole, gives us the extremal elements of $\mathcal{P}^{\text{conv}}_{\text{sym,Monge}}\left(X^N\right)$.

The data in Fig. 2 was computed using MATLAB [26] and polymake [35].

It was already mentioned above, that the extreme points of the polytope P_{coef} have a sparse structure. In more detail, a coefficient vector $\alpha \in P_{\text{coef}}$ is extremal with respect to the polytope P_{coef} if and only if its nonzero entries correspond to a selection of columns of A which are linearly independent (see, e.g., [3]). That is why, the complexity of computing the extreme points of P_{coef}, and their number, increases faster with the number of states than with the number of marginals. Suppose you are looking at a setting where the number of marginals is equal to the number of states. Then, on the one hand, increasing the number of marginals by 1 yields $\binom{2N}{N+1} - \binom{2N-1}{N}$ more columns in A. On the other hand, an increase in the number of

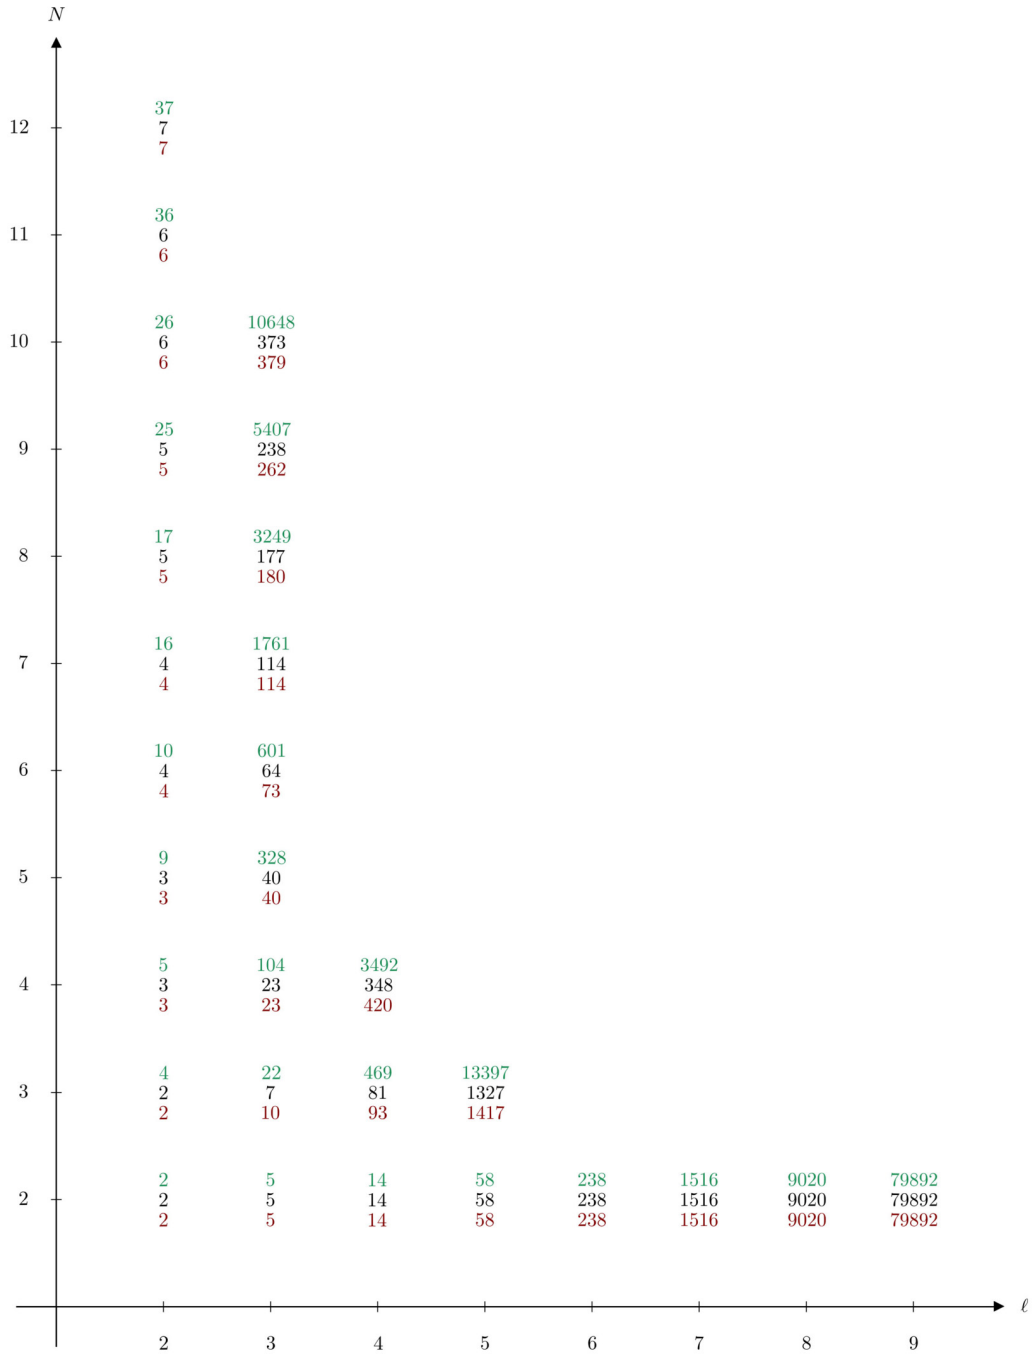

Fig. 2. The number of extreme points of the symmetric Kantorovich polytope $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}(X^N)$ respectively of the symmetric Monge polytope $\mathcal{P}^{\mathrm{conv}}_{\mathrm{sym,Monge}}(X^N)$ is given in green respectively red. The number of extreme points of $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}(X^N)$ that are of Monge-type (see (1.6)-(1.8)) is depicted in black. Here, as usual, N denotes the number of marginals and ℓ the number of states.

states by 1 enlarges the number of columns of A by $\binom{2N}{N} - \binom{2N-1}{N}$. Elementary computations show that in the second case A has $\frac{1}{N+1}\binom{2N}{N}$ more columns than in the first case. Moreover, in contrast to an increase in the number of marginals, an increase in the number of states also increases the number of rows of A by 1. Therefore then up to $\ell + 1$ columns of A can be linearly independent. Hence, an increase in the number of states leads to a faster increasing (compared to an increase in the number of marginals) number of subsets

of linearly independent columns of the constraint matrix A by yielding a steeper increase in the number of columns of A as well as by enlarging the dimension of the column space. Each of these subsets corresponds to an extreme point of P_{coef}.

Remark 2.4. This remark lists interpretations and observations regarding Fig. 2.

1) In the case $N = 2$, Fig. 2 shows that in the considered cases every extremal symmetric Kantorovich coupling is of Monge-type. In the given setting this means that every extreme point of $\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)$ is a symmetrized permutation matrix. It is easy to see, using the celebrated Birkhoff-von Neumann theorem [4,41] as well as the linearity of the symmetrization operator S (2.4), that this holds true for an arbitrary number ℓ of states. Note, however, that not every symmetrized permutation matrix is an extreme point of $\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)$, but only those symmetrized Monge-states whose corresponding coefficient vectors select linearly independent columns of A.
2) In the case $\ell = 2$, the number of extremal symmetric Kantorovich couplings which are of Monge-type increases by 1 each time the marginal number is even. It is easy to prove that this pattern will continue. Firstly, note that, in the case $\ell = 2$, every symmetric Kantorovich coupling of Monge-type is an extreme point of $\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)$. This follows by a support-argument regarding the corresponding coefficient vectors. Secondly, we take a look at the symmetrized Monge-states in this setting. We assume the marginal vectors $\lambda^{(1)}, \dots, \lambda^{(N+1)}$ are sorted in the columns of A by the first component in decreasing order, i.e.,

$$A = \begin{pmatrix} 1 & \frac{N-1}{N} & \dots & \frac{1}{N} & 0 \\ 0 & \frac{1}{N} & \dots & \frac{N-1}{N} & 1 \end{pmatrix}.$$

Then the symmetric Kantorovich couplings of Monge-type are exactly those couplings with coefficient vectors

$$\alpha^{(j)} = \frac{1}{2}e^{(j)} + \frac{1}{2}e^{(N+1-j+1)}$$

for $j = 1, 2, \dots, \left\lceil \frac{N+1}{2} \right\rceil$, where $e^{(i)}$ is the i-th unit vector.
3) The setting of 3 marginals and 3 sites, i.e., $N = \ell = 3$ is the main focus in [18]. There interested readers can find the 22 extreme points of the symmetric Kantorovich polytope explicitly listed including the information which extremal elements are of Monge-type and which are not. This list also shows which pairs of permutations (identifying T_1 with the identity) correspond to an extremal symmetric Kantorovich coupling. [18] also visualizes these 22 extremal states as molecular packings, where one can identify irreducible packings with extreme points.
4) Note that for each grid-point in Fig. 2 dividing the number depicted in black by the number depicted in green, i.e., '$\frac{\text{black}}{\text{green}}$', gives the ratio of extreme points of the symmetric Kantorovich polytope that are of Monge-type. For a fixed three element state space, i.e., $\ell = 3$, this ratio consistently decreases with growing N from 1 for 2 marginals to 0.035 for 10 marginals. Reversing the roles of N and ℓ, i.e., fixing the number of marginals N to three and letting the number of marginal states grow, also yields a consistently decreasing behavior of the considered ratio; starting from 0.5 for $\ell = 2$ and ending at 0.099 for $\ell = 5$.
The considered ratio has the following interesting probabilistic interpretation. Given a non-degenerate cost function, i.e., a cost function that yields a unique optimizer of problem (1.1), the probability of the corresponding optimizer being of Monge-type is given by the considered ratio. Here we obviously draw uniformly from the set of extremal symmetric Kantorovich couplings. Specific cost functions might always yield Monge-type optimizers, see Section 4.

5) For each grid-point in Fig. 2 the ratio '$\frac{\text{black}}{\text{red}}$' is an indicator of how much unnecessary information is contained in Monge's ansatz. For the considered cases this ratio is always above 0.8 except for one outlier at $N = \ell = 3$ where the ratio is given by 0.7. Hence, even though the Monge ansatz does not contain the 'entire information of the Kantorovich polytope', see 4), at least it does not entail 'a lot' of unnecessary information.
6) Finally, we want to give computationally determined examples of non-Monge extreme points in the case of $N = 3$ marginals.

$$\ell = 3: \ \frac{1}{2}\psi_N \begin{pmatrix} 2/3 \\ 1/3 \\ 0 \end{pmatrix} + \frac{1}{2}\psi_N \begin{pmatrix} 0 \\ 1/3 \\ 2/3 \end{pmatrix} \qquad \ell = 4: \ \frac{3}{8}\psi_N \begin{pmatrix} 2/3 \\ 1/3 \\ 0 \\ 0 \end{pmatrix} + \frac{3}{8}\psi_N \begin{pmatrix} 0 \\ 1/3 \\ 2/3 \\ 0 \end{pmatrix} + \frac{1}{4}\psi_N \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}$$

$$\ell = 5: \ \frac{3}{10}\psi_N \begin{pmatrix} 2/3 \\ 1/3 \\ 0 \\ 0 \\ 0 \end{pmatrix} + \frac{3}{10}\psi_N \begin{pmatrix} 0 \\ 1/3 \\ 2/3 \\ 0 \\ 0 \end{pmatrix} + \frac{1}{5}\psi_N \begin{pmatrix} 0 \\ 0 \\ 0 \\ 2/3 \\ 1/3 \end{pmatrix} + \frac{1}{5}\psi_N \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1/3 \\ 2/3 \end{pmatrix}$$

The $(\ell = 3)$-example was already given in [18]. The two remaining extreme points both consist of two components. One that is compatible with Monge's approach; it is given by the last $(\ell = 4)$ respectively the two last $(\ell = 5)$ terms of the corresponding sum. The remaining component arises from the $(\ell = 3)$-example and yields the non-Monge property of the considered extreme points. These considerations indicate how to construct non-Monge extreme points for growing ℓ: Assume the number of states ℓ to be no less than 6. Firstly, choose an increasing triple $\{i_1, i_2, i_3\}$ of pairwise distinct indices from the set $\{1, 2, \dots, \ell\}$. These indices mark the elements of the finite state space X that will be 'covered' by the $(\ell = 3)$-example of a non-Monge extreme point from above. In order to simplify notation we assume these 'covered' states to be a_1, a_2 and a_3, i.e., $\{i_1, i_2, i_3\} = \{1, 2, 3\}$. Moreover, let $\lambda^{(1)}, \dots, \lambda^{(\ell-3)} \in \mathcal{P}_{\frac{1}{3}}(\{a_4, \dots, a_\ell\})$ be $\frac{1}{3}$-quantized probability measures on $X \backslash \{a_1, a_2, a_3\}$ which form an extreme point of the symmetric Kantorovich polytope for 3 marginals and $\ell - 3$ states in Monge-form. Then

$$\frac{3}{2 \cdot \ell}\psi_N \begin{pmatrix} 2/3 \\ 1/3 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + \frac{3}{2 \cdot \ell}\psi_N \begin{pmatrix} 0 \\ 1/3 \\ 2/3 \\ 0 \\ \vdots \\ 0 \end{pmatrix} + \sum_{i=1}^{\ell-3} \frac{1}{\ell}\psi_N \begin{pmatrix} 0 \\ 0 \\ 0 \\ | \\ \lambda^{(i)} \\ | \end{pmatrix}$$

is a non-Monge extreme point for 3 marginals and ℓ states. The described construction reveals that from the single non-Monge extreme point for $\ell = 3$ states arise

$$\binom{\ell}{3} \cdot \text{Monge}_{\ell-3}$$

non-Monge extreme points for ℓ states. Hereby $\text{Monge}_{\ell-3}$ is the number of extreme points of the symmetric Kantorovich polytope for 3 marginals and $\ell - 3$ states in Monge-form.

The computational restriction regarding our extreme point investigation is visualized in Fig. 3 and 4: Fig. 3 depicts a super-exponential growth of the number of extremal symmetric Kantorovich couplings in the case of $N = 2$ marginals. Moreover, Fig. 4 indicates that the portion of permutation matrices which correspond to extremal symmetric Kantorovich couplings (in the case of $N = 2$ marginals) tends to a value close to 0.2.

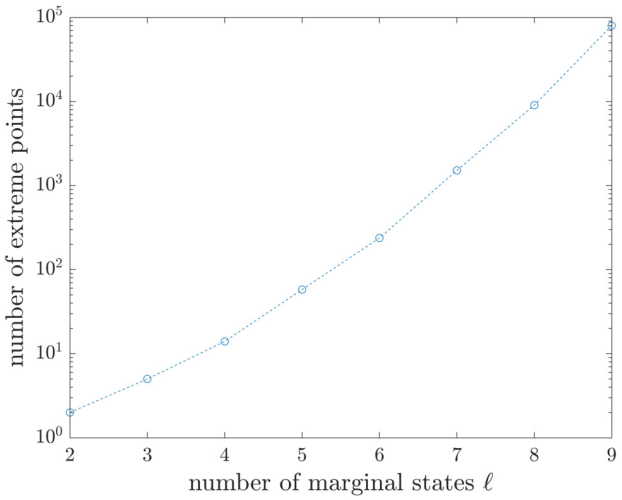

Fig. 3. The number of extremal symmetric Kantorovich couplings for $N = 2$ marginals is depicted in dependency of the number of marginal states ℓ.

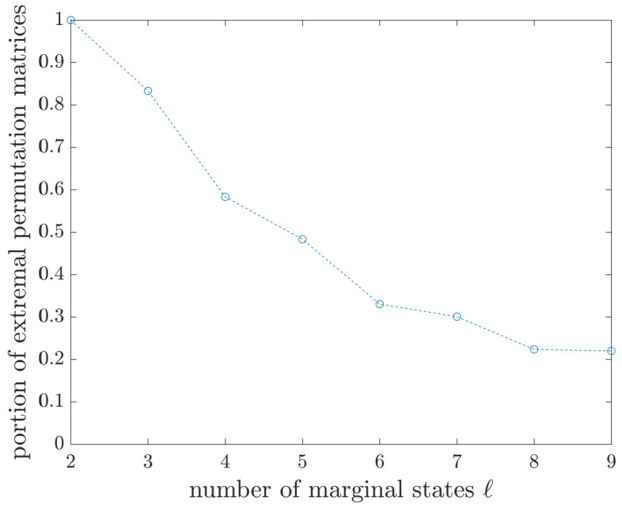

Fig. 4. The ratio of the number of extremal symmetrized permutation matrices to the number of permutation matrices is depicted in dependency of the number of marginal states ℓ.

3. Classification of the extreme points of a reduced Kantorovich polytope

In Section 2, we achieved a better understanding of the OT problem (1.1) by numerically analyzing the convex geometry of the set of admissible trial states, i.e., the set of symmetric Kantorovich couplings $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$. Motivated by applications in physics we assume from this point on that the given cost function has pairwise symmetric structure. Then the set of admissible trial states can be reduced by the linear map M_2 to a lower-dimensional polytope thereby decreasing the number of extremal states.

In more detail, we consider the OT problem (1.1) with $c : X^N \to \mathbb{R}$ being a cost function with pairwise symmetric structure, i.e.,

$$c(x_1, \ldots, x_N) = \sum_{1 \leq i < j \leq N} v(x_i, x_j) \text{ for all } (x_1, \ldots, x_N) \in X^N, \tag{3.1}$$

where $v : X^2 \to \mathbb{R}$ is a symmetric pair-potential, i.e., $v(x,y) = v(y,x)$ for all $(x,y) \in X^2$. Then the objective function of (1.1) can be rewritten as

$$\int_{X^N} c(x_1,\dots,x_N)d\gamma(x_1,\dots,x_N) = \binom{N}{2}\int_{X^2} v(x,y)d\left(M_2\gamma\right)(x,y), \tag{3.2}$$

where $\gamma \in \mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$ is an arbitrary symmetric Kantorovich coupling. This elementary reformulation was established in [19]. There also the concept of N-representability (see Definition 3.1) was introduced which we will use in the following to write the reduced set of admissible trial states in a compact manner.

Definition 3.1 *(N-representability).* A probability measure $\mu \in \mathcal{P}\left(X^k\right)$ is called N-representable if there exists a symmetric probability measure γ on the product space X^N, i.e., $\gamma \in \mathcal{P}_{\mathrm{sym}}\left(X^N\right)$, such that μ is its k-point marginal, i.e.,

$$\mu = M_k\gamma. \tag{3.3}$$

Any such symmetric probability measure on X^N that fulfills (3.3) is then called a representing measure of μ. In the following the set of N-representable k-plans will be denoted by $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$.

As we consider pairwise interactions, we will focus our attention on the set of N-representable 2-point measures, i.e., $\mathcal{P}_{N\text{-rep}}\left(X^2\right)$. Note, however, that cost functions c embodying k-particle interactions would give rise to a problem reformulation reducing the set of admissible trial states to a subset of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$. In the case $k = N$, c would be a symmetric cost which are, as mentioned in the introduction, 'dual' to the set of symmetric probability measures on the product space X^N. In the same manner, cost functions with symmetric pairwise structure have a dual relationship with the set of N-representable 2-plans.

By definition, the set of N-representable 2-point measures is the image of the set of symmetric probability measures on X^N under the map M_2, defined in (2.7), i.e., $M_2\left(\mathcal{P}_{\mathrm{sym}}\left(X^N\right)\right) = \mathcal{P}_{N\text{-rep}}\left(X^2\right)$. Combining this equality with (3.2) yields that (1.4) is an equivalent reformulation of the multi-marginal OT problem (1.1) for a cost function with pairwise symmetric structure (3.1). Here (1.4) can also be written as

$$\min_{\mu\in\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)} \int_{X^2} v(x,y)d\mu(x,y),$$

where $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ is the set of N-representable 2-plans having uniform marginal, i.e.,

$$\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right) := \left\{\mu \in \mathcal{P}_{N\text{-rep}}\left(X^2\right) : M_1(\mu) = \overline{\lambda}\right\}. \tag{3.4}$$

We will refer to the set $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ as *reduced Kantorovich polytope for N marginals and ℓ states*. The convex geometry of this set will be numerically analyzed in the following. Thereby the validity of Monge's approach in the given setting will be tested.

We have seen above that under the assumption of pairwise symmetric cost functions the OT problem (1.1), where the set of admissible trial states is given by the high-dimensional set $\mathcal{P}_{\mathrm{sym},\overline{\lambda}}\left(X^N\right)$, can be equivalently formulated as a minimization problem over the lower-dimensional set $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ (see (1.4)). The pairwise symmetric structure implies that any symmetric Kantorovich coupling influences the value of the objective function of problem (1.1) only through their respective two-point marginal (see (3.2)). The nature of this reformulation, applying the two-point marginal map M_2 on the set of symmetric Kantorovich couplings, however, entails that the new set of admissible trial states, i.e., the reduced Kantorovich polytope is only implicitly known. Only in the two-marginal (N=2) case the reduced Kantorovich polytope can be

understood in a straightforward manner: It corresponds to the set of symmetric bistochastic matrices scaled by the factor $\frac{1}{\ell}$ (see Remark 3.5 1) below for further consideration of the two-marginal case). Hence, in the case $N = 2$, $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right) = \mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)$ holds. For a better understanding of the multi-marginal ($N > 2$) case, we will in the following, as motivated, view the reduced Kantorovich polytope as the image of the set of symmetric Kantorovich couplings on X^N under the two-point marginal map, i.e.,

$$M_2(\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)) = \mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right). \tag{3.5}$$

As described in Section 2, $\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)$ corresponds to the convex hull of its extreme points. Combining this fact with (3.5) and the linearity of M_2 yields that the reduced Kantorovich polytope $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ is equal to the convex hull of the two-point marginals of extremal symmetric Kantorovich couplings, i.e.,

$$\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right) = \operatorname{conv}\left(\left\{M_2\gamma : \gamma \text{ is an extreme point of } \mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)\right\}\right). \tag{3.6}$$

The following proposition is an immediate consequence.

Proposition 3.2. *Any extreme point of the reduced Kantorovich polytope for N marginals and ℓ states is the two-point marginal of an extremal symmetric Kantorovich coupling.*

Now the question is whether or not M_2 represents a bijective relationship between the sets of extreme points of $\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)$ and $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$. The following remark sheds light on this issue applying the bijective relationship between $\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)$ and the polytope P_{coef} established in Lemma 2.2 and Corollary 2.3.

Remark 3.3. In Section 2 the extreme points of $\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)$, i.e., the set of admissible trial states of problem (1.1), are determined using the set's bijective relationship, captured in the coefficients-to-coupling map R introduced in Section 2, to the polytope P_{coef}. As explained above in more detail, the map R identifies any symmetric probability measure on X^N γ with a coefficient vector α, such that γ can be written as the respective convex combination of the extreme points of $\mathcal{P}_{\text{sym}}\left(X^N\right)$, i.e., (2.6) holds. These coefficients are unique due to the disjoint support of the extremal symmetric probability measures on X^N. It was proven in [20] that the two-point marginal map M_2 is a bijection between the sets of extreme points of $\mathcal{P}_{\text{sym}}\left(X^N\right)$ and $\mathcal{P}_{N\text{-rep}}\left(X^2\right)$, respectively. Due to the linearity of M_2, given a coefficient vector α and a symmetric probability measure γ on X^N, such that $\gamma = R\alpha$, i.e., (2.6) holds true, then

$$M_2\gamma = \sum_{1\le i_1\le\cdots\le i_N\le\ell} \alpha_{i_1,\ldots,i_N} M_2 S\delta_{i_1,\ldots,i_N}.$$

Only now, these coefficients α representing $M_2\gamma$ as a convex combination of the extreme points of the set of N-representable two-point measures may not be unique, rendering us unable to identify the extreme points of the reduced Kantorovich polytope with those of the coefficient-polytope P_{coef}.

The remark above illuminates why the extreme points of the set of symmetric Kantorovich couplings $\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)$ can not be identified with the extremal elements of the reduced Kantorovich polytope $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ via M_2. The two-point marginal map may for example map multiple extreme points of the set $\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)$ on a single point of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$; this point may lie on a face or in the interior of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ (see [18] for an well-illustrated example).

Nevertheless, it was established in Proposition 3.2 that every extremal element of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ has a representing measure that is itself extremal with respect to $\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^N\right)$. The extreme points of this set of symmetric Kantorovich couplings were in Corollary 2.3 identified with the extreme points of P_{coef}. In combination with the in Remark 3.3 established connection between P_{coef} and $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ this leads us to the following approach to determine the extremal elements of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$:

1. We start by determining the extremal elements of P_{coef}. This was already done within the considerations of Section 2.
2. Every such extreme point is multiplied by the matrix $T \in \mathbb{R}^{\ell^2 \times |E^N_{\text{sym}}|}$ which is constructed as follows. The matrix A as defined in (2.10) lists all the elements of $\mathcal{P}_{\frac{1}{N}}(X)$ as columns. It was proven in [20] that for any element λ of $\mathcal{P}_{\frac{1}{N}}(X)$ the following holds:

$$M_2\psi_N(\lambda) = \frac{N}{N-1}\lambda \otimes \lambda - \frac{1}{N-1}(\text{id},\text{id})\,\#\lambda, \tag{3.7}$$

where the map ψ_N was introduced in Section 2. Note that it was further established in [20] that measures of form (3.7) for $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$ are exactly the extreme points of $\mathcal{P}_{N\text{-rep}}(X^2)$. Now we construct T by replacing any column λ of A with $M_2\psi_N(\lambda)$ as given in (3.7) where we canonically identify matrices with vectors by gluing columns together.
3. Finally we check which points of the form

$$T\alpha \text{ for } \alpha \in \text{ext}(P_{\text{coef}})$$

are extremal with respect to $\text{conv}(\{T\alpha : \alpha \in \text{ext}(P_{\text{coef}})\})$ and therefore by (3.6) with respect to $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$.

Note that it is computationally more complex to determine the extremal elements of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ than those of $\mathcal{P}_{\text{sym},\overline{\lambda}}(X^N)$.

Now, we will incorporate Monge's approach in the reduced setting.

Definition 3.4. An element of the reduced Kantorovich polytope for N marginals and ℓ states is said so be *of Monge-type* or *in Monge-form* if it has a representing measure that is of Monge-form (see (1.6)-(1.8)).

This definition is consistent with our goal to check the validity of Monge's approach as any optimizer in Monge-form for problem (1.4) guarantees the existence of an optimizer in Monge-form for problem (1.1). The set of all elements of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ which are in Monge-form will be denoted as $\mathcal{P}_{N\text{-rep,Monge}}(X^2)$, i.e.,

$$\mathcal{P}_{N\text{-rep,Monge}}(X^2) := \{M_2\gamma : \gamma \text{ is of Monge-type (1.6)} - \text{(1.8)}\}.$$

Analogously to (2.13) we introduce the *reduced Monge polytope for N marginals and ℓ states* $\mathcal{P}^{\text{conv}}_{N\text{-rep,Monge}}(X^2)$ as follows.

$$\mathcal{P}^{\text{conv}}_{N\text{-rep,Monge}}(X^2) := \text{conv}\left(\mathcal{P}_{N\text{-rep,Monge}}(X^2)\right) \tag{3.8}$$

The extremal elements of the reduced Monge polytope can be determined in the same manner as those of the reduced Kantorovich polytope (see the description of the procedure above). Starting point are now the extremal elements of the Monge polytope $\mathcal{P}_{\text{sym,Monge}}(X^N)$ interpreted as coefficient vectors.

Checking which of the extreme points of the reduced Kantorovich polytope $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ correspond to an extremal element of the reduced Monge polytope $\mathcal{P}^{\text{conv}}_{N\text{-rep,Monge}}(X^2)$ tells us which of the extreme points of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ are of Monge-type.

The data in Fig. 5 was computed using MATLAB [26] and polymake [35].

Remark 3.5. What follows are interpretations and observations regarding Fig. 5.

1) Combining the convention $M_N = \text{id}$ with (3.5), it is obvious that the symmetric Kantorovich polytope for 2 marginals and ℓ states $\mathcal{P}_{\text{sym},\overline{\lambda}}(X^2)$ coincides with the reduced Kantorovich polytope for 2 marginals

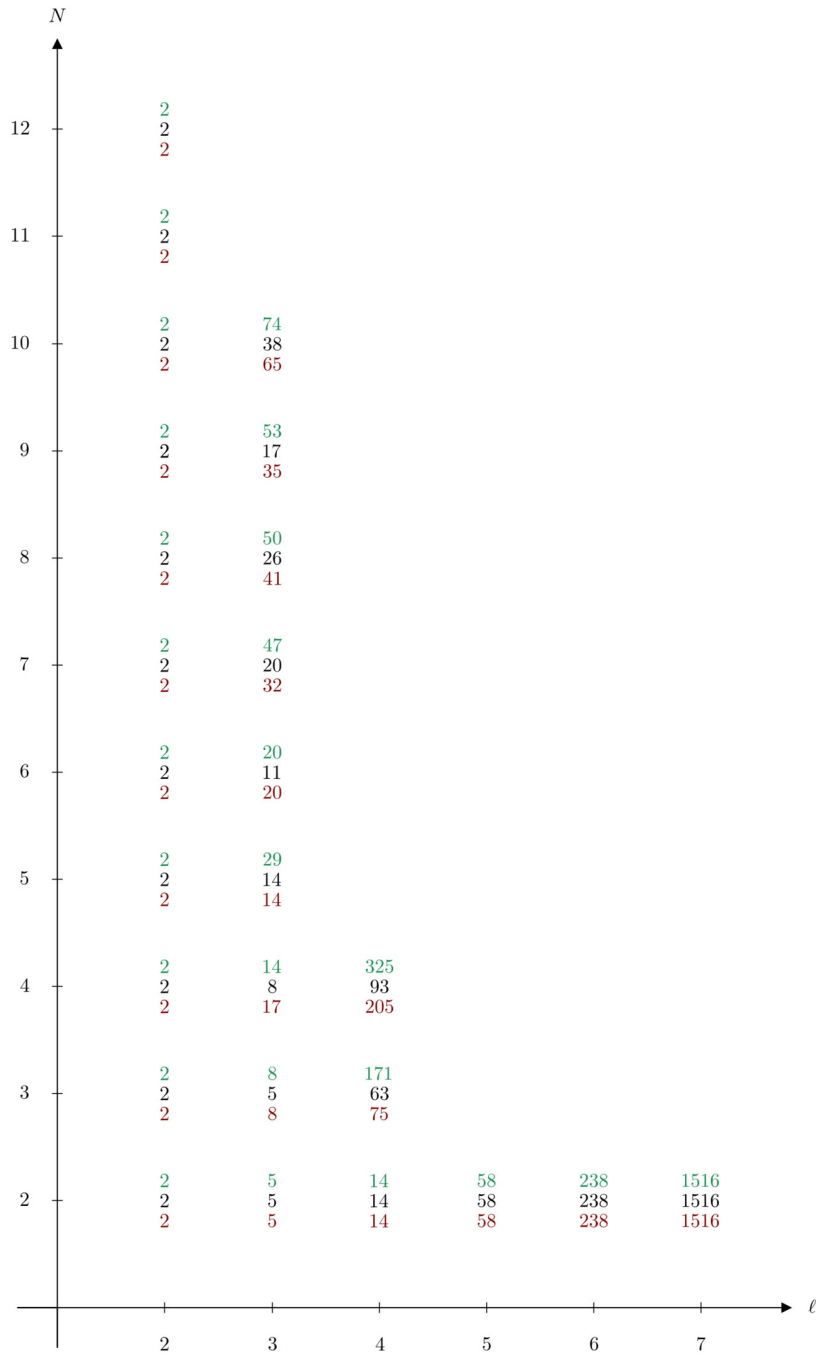

Fig. 5. The number of extreme points of the reduced Kantorovich polytope $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ respectively of the reduced Monge polytope $\mathcal{P}^{\text{conv}}_{N\text{-rep,Monge}}\left(X^2\right)$ is given in green respectively red. The number of extreme points of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ that are of Monge-type (see Definition 3.4) is depicted in black. Here, as usual, N denotes the number of marginals and ℓ the number of states.

and ℓ states $\mathcal{P}_{2\text{-rep},\overline{\lambda}}\left(X^2\right)$. This fact was already mentioned above. It was established in Remark 2.4 that every extreme point of $\mathcal{P}_{\text{sym},\overline{\lambda}}\left(X^2\right)$ is a symmetrized permutation matrix, i.e., the image of a permutation matrix under the symmetrization operator (2.4). In the setting of 2 marginals, symmetrized permutation matrices exactly correspond to symmetrized Monge states. See Remark 2.4 for further considerations of the case $N = 2$.

2) In the case $\ell = 2$, Fig. 5 depicts that in the considered cases, the reduced Kantorovich polytope $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ has two extreme points both of which are in Monge-form. Hence, in any considered case the line segment $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ coincides with the respective reduced Monge polytope $\mathcal{P}^{\text{conv}}_{N\text{-rep,Monge}}(X^2)$. One can prove by elementary arguments that this holds true for an arbitrary number of marginals N in the case of 2 sites. In a little more detail, considering the dimension of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ in the given case and parametrising the elements of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ by their off-diagonal element allows us to deduce that the two extreme points of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ are given by

$$\mu^{(1)} = M_2\left(\frac{1}{2}\psi_N(\delta_1) + \frac{1}{2}\psi_N(\delta_2)\right) \tag{3.9}$$

$$\mu^{(2)} = \begin{cases} M_2\left(\psi_N\left(\frac{1}{2}\delta_1 + \frac{1}{2}\delta_2\right)\right) & \text{if } N \text{ is even} \\ M_2\left(\frac{1}{2}\psi_N\left(\frac{N-1}{2N}\delta_1 + \frac{N+1}{2N}\delta_2\right) + \frac{1}{2}\psi_N\left(\frac{N+1}{2N}\delta_1 + \frac{N-1}{2N}\delta_2\right)\right) & \text{if } N \text{ is odd} \end{cases} \tag{3.10}$$

or in pedestrian notation,

$$\mu^{(1)} = \begin{pmatrix} \frac{1}{2} & 0 \\ 0 & \frac{1}{2} \end{pmatrix}$$

$$\mu^{(2)} = \begin{cases} \frac{1}{4(N-1)}\begin{pmatrix} N-2 & N \\ N & N-2 \end{pmatrix} & \text{if } N \text{ is even} \\ \frac{1}{4N}\begin{pmatrix} N-1 & N+1 \\ N+1 & N-1 \end{pmatrix} & \text{if } N \text{ is odd.} \end{cases}$$

As the coefficients in (3.9) and (3.10) are integer multiples of $\frac{1}{\ell}$ both extreme points are 2-point marginals of symmetric Kantorovich couplings in Monge-form and therefore they are themselves elements of the reduced Kantorovich polytope $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ which are of Monge-type (see Definition 3.4). Note that $\mu^{(2)}$, which is of Monge-type and therefore describes a correlated or in other words deterministic state, converges for $N \to \infty$ to the independent measure $\lambda \otimes \lambda$ for $\lambda = \left(\frac{1}{2}\delta_1 + \frac{1}{2}\delta_2\right)$.
These findings coincide with the results in [19], where a model problem of N particles on 2 sites was considered. There also the set of N-representable 2-plans $\mathcal{P}_{N\text{-rep}}(X^2)$ for X consisting of 2 distinct elements is illustrated. Imposing the here given marginal condition on these sets leads to the respective line segment $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$.

3) The case of 3 marginals and 3 sites, i.e., $N = \ell = 3$, is a minimal example of a point in the grid, both with respect to the sum of both parameters $N + \ell$ and with respect to the minimum of both parameters $\min\{N, \ell\}$, such that not every extremal element of the reduced Kantorovich polytope $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ is of Monge-type. In the considered case $\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)$ has 8 extreme points 5 of which are in Monge-form. By extension 3 of them are not. They are given by

$$\frac{1}{2}M_2(S\delta_{112}) + \frac{1}{2}M_2(S\delta_{233}) \tag{3.11}$$

and the two states one generates by imposing the role of the second site on the first and third site respectively. (3.11) is the unique optimizer of an OT problem stated in [18]. This problem corresponds to a molecular packing problem. See [18] for further reading.

4. A model problem: optimal couplings of N marginals on 3 sites for pairwise costs

In the following, we focus our attention on symmetric multi-marginal OT problems (1.1) on 3 sites, i.e., $X = \{a_1, a_2, a_3\}$. As in Section 3, we only consider pairwise symmetric costs and therefore are able to reformulate (1.1) as the lower-dimensional problem (1.4). In particular, the reduced Kantorovich polytope

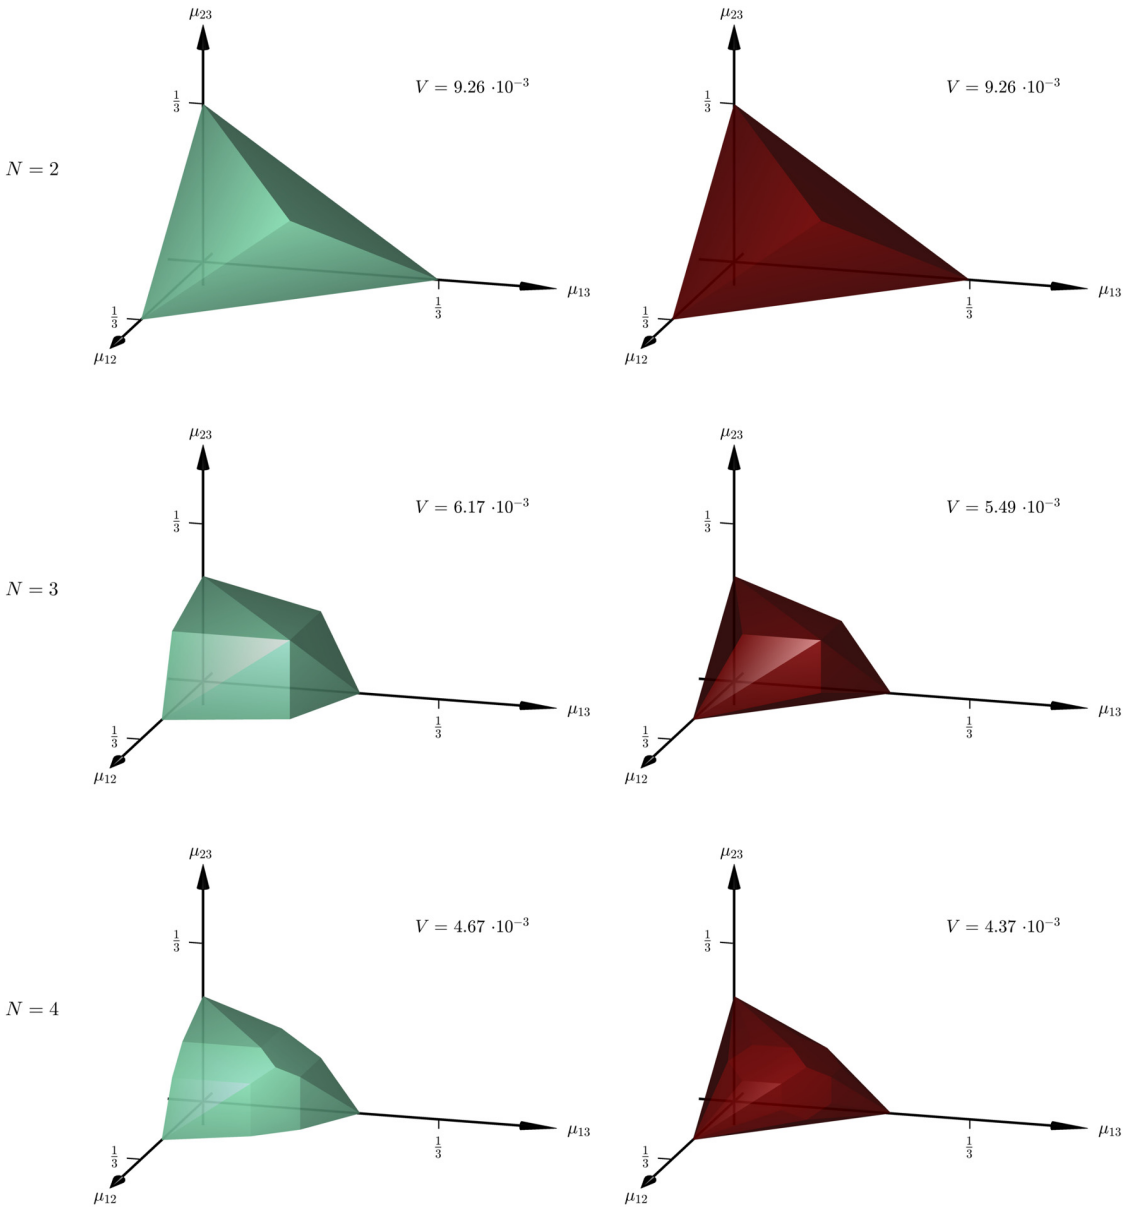

Fig. 6. The reduced Kantorovich respectively Monge polytope for N marginals and 3 states is visualized for $N = 2, 3, 4, 6$ and 10 in green respectively red. The plots are arranged in ascending order with respect to N. The elements $(\mu_{ij})_{i,j=1}^{3}$ of the polytopes are parametrized by their off-diagonal entries μ_{12}, μ_{13} and μ_{23}. In the case of three marginals the reduced polytopes were initially depicted in [18]. V indicates the volume of the corresponding polytope. The volumetric ratio, reduced Monge polytope to reduced Kantorovich polytope, is depicted in Fig. 8.

for N marginals and 3 states corresponds to the respective set of admissible trial states. It is easy to see that in the given setting these polytopes are three-dimensional. As by extension the reduced Monge polytope is at most three-dimensional, we are able to visually compare both approaches.

The visualizations in Fig. 6 were generated by extending the above explained calculations and routines in MATLAB [26] and polymake [35].

Note that by definition the reduced Monge polytope is always contained in the reduced Kantorovich polytope independently of the number of marginals N and the number of sites ℓ. Some of the extreme points of the reduced Monge polytope are also extremal with respect to the reduced Kantorovich polytope and some lie on faces or in the interior of the latter (see Fig. 5 for specific numbers).

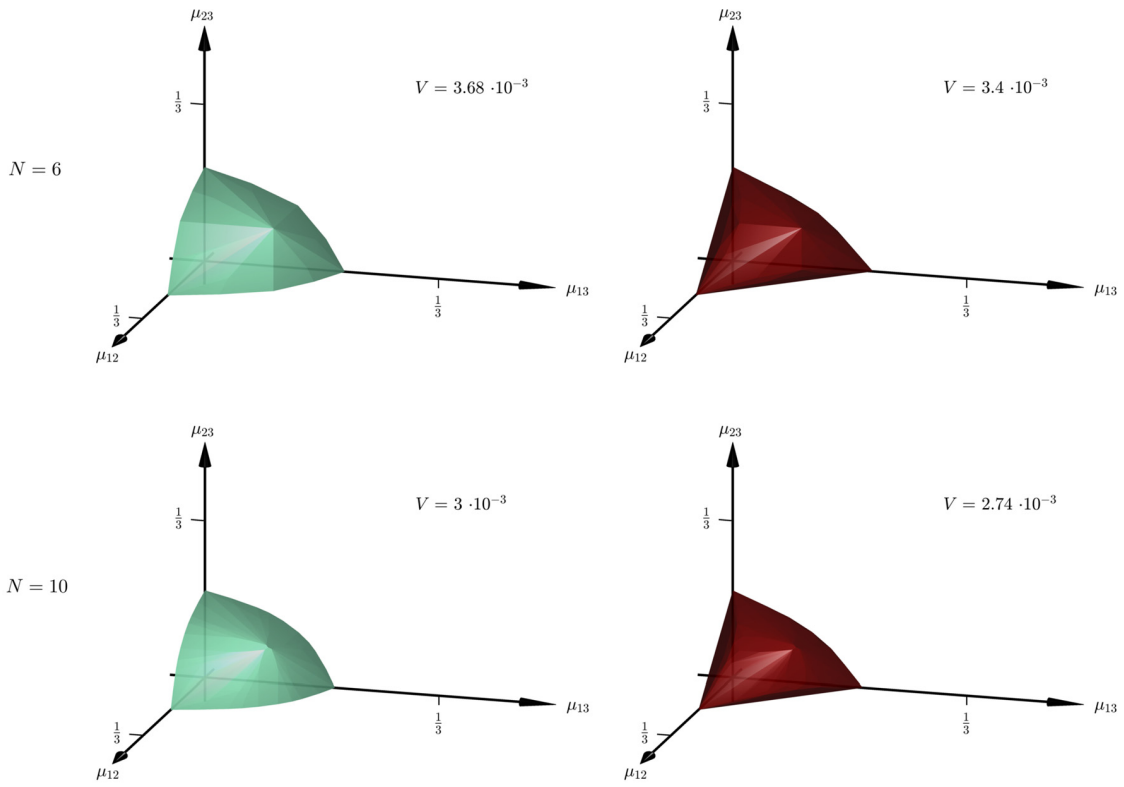

Fig. 6. (*continued*)

In the setting of $\ell = 3$ sites, the numerical analysis of the reduced setting discussed in Section 3 yields that in the case of $N = 2, 3, \dots, 9$ and 10 marginals there are always 5 prominent extreme points of the reduced Kantorovich polytope that are of Monge-type. This is also indicated by Fig. 1 as well as Fig. 6. In the illustrations they can be identified with the four extreme points on the co-ordinate axes, including the origin, as well as the 'peak' in the front of the polytopes. In formulas these extreme points can be written as depicted in Table 1. In Fig. 6, $EA^{(N)}$ corresponds to the origin, $ER^{(N)}$ to the 'peak' and $E12^{(N)}$ to the non-origin extreme point on the μ_{12}-axis. $E12^{(N)}$ assumes an exemplary role in Table 1. The corresponding extreme points on the μ_{13}- respectively μ_{23}-axis can be expressed analogously in abstract as well as matrix notation and will be denoted by $E13^{(N)}$ respectively $E23^{(N)}$.

So far we know by numerical analysis that $EA^{(N)}$, $ER^{(N)}$, $E12^{(N)}$, $E13^{(N)}$ and $E23^{(N)}$ are extreme points of the reduced Kantorovich polytope $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ in the cases of $N = 2, 3, \dots, 9$ and 10 marginals. One can prove that this holds true for a general number $N \geq 2$ of marginals. For $E12^{(N)}$, $E13^{(N)}$ and $E23^{(N)}$ one can show this by following the same approach that was taken in Remark 3.5 2). In case of $EA^{(N)}$ and $ER^{(N)}$ it is an immediate consequence of Theorem 4.1. In the following $d : X \times X \to \mathbb{R}$ will denote the discrete metric defined by

$$d(x,y) := \begin{cases} 1 & \text{if } x \neq y \\ 0 & \text{if } x = y. \end{cases}$$

Theorem 4.1. *We consider the reduced multi-marginal OT problem (1.4) for $N \geq 2$ marginals and $\ell = 3$ sites.*

a) For the attractive cost function $d : X \times X \to \mathbb{R}$ the unique minimizer is given by $EA^{(N)}$.

Table 1
The prominent extreme points $EA^{(N)}$, $ER^{(N)}$ and $E12^{(N)}$ of the reduced Kantorovich polytope are depicted in abstract and matrix notation. Hereby N corresponds to the number of marginals and $m \in \mathbb{N}_0$ is a non-negative integer, allowing us to distinguish between the various cases regarding N. As all the coefficients in the 'Abstract Notation'-column are integer multiples of $\frac{1}{3}$, these extreme points are of Monge-type.

Nomenclature		Abstract Notation	Matrix Notation
$EA^{(N)}$	$EA^{(N)}$	$\frac{1}{3}M_2\psi_N(\delta_1)$ $+\frac{1}{3}M_2\psi_N(\delta_2)$ $+\frac{1}{3}M_2\psi_N(\delta_3)$	$\begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}$
$ER^{(N)}$	$ER^{(3m)}$	$M_2\psi_N\left(\frac{1}{3}\delta_1 + \frac{1}{3}\delta_2 + \frac{1}{3}\delta_3\right)$	$\begin{pmatrix} \frac{N-3}{9(N-1)} & \frac{N}{9(N-1)} & \frac{N}{9(N-1)} \\ \frac{N}{9(N-1)} & \frac{N-3}{9(N-1)} & \frac{N}{9(N-1)} \\ \frac{N}{9(N-1)} & \frac{N}{9(N-1)} & \frac{N-3}{9(N-1)} \end{pmatrix}$
	$ER^{(3m+1)}$	$\frac{1}{3}M_2\psi_N\left(\frac{m}{N}\delta_1 + \frac{m}{N}\delta_2 + \frac{m+1}{N}\delta_3\right)$ $+\frac{1}{3}M_2\psi_N\left(\frac{m}{N}\delta_1 + \frac{m+1}{N}\delta_2 + \frac{m}{N}\delta_3\right)$ $+\frac{1}{3}M_2\psi_N\left(\frac{m+1}{N}\delta_1 + \frac{m}{N}\delta_2 + \frac{m}{N}\delta_3\right)$	$\begin{pmatrix} \frac{N-2}{9N} & \frac{N+1}{9N} & \frac{N+1}{9N} \\ \frac{N+1}{9N} & \frac{N-2}{9N} & \frac{N+1}{9N} \\ \frac{N+1}{9N} & \frac{N+1}{9N} & \frac{N-2}{9N} \end{pmatrix}$
	$ER^{(3m+2)}$	$\frac{1}{3}M_2\psi_N\left(\frac{m}{N}\delta_1 + \frac{m+1}{N}\delta_2 + \frac{m+1}{N}\delta_3\right)$ $+\frac{1}{3}M_2\psi_N\left(\frac{m+1}{N}\delta_1 + \frac{m}{N}\delta_2 + \frac{m+1}{N}\delta_3\right)$ $+\frac{1}{3}M_2\psi_N\left(\frac{m+1}{N}\delta_1 + \frac{m+1}{N}\delta_2 + \frac{m}{N}\delta_3\right)$	$\begin{pmatrix} \frac{N-2}{9N} & \frac{N+1}{9N} & \frac{N+1}{9N} \\ \frac{N+1}{9N} & \frac{N-2}{9N} & \frac{N+1}{9N} \\ \frac{N+1}{9N} & \frac{N+1}{9N} & \frac{N-2}{9N} \end{pmatrix}$
$E12^{(N)}$	$E12^{(2m)}$	$\frac{2}{3}M_2\psi_N\left(\frac{1}{2}\delta_1 + \frac{1}{2}\delta_2\right) + \frac{1}{3}M_2\psi_N\delta_3$	$\begin{pmatrix} \frac{N-2}{6(N-1)} & \frac{N}{6(N-1)} & 0 \\ \frac{N}{6(N-1)} & \frac{N-2}{6(N-1)} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}$
	$E12^{(2m+1)}$	$\frac{1}{3}M_2\psi_N\left(\frac{N-1}{2N}\delta_1 + \frac{N+1}{2N}\delta_2\right)$ $+\frac{1}{3}M_2\psi_N\left(\frac{N+1}{2N}\delta_1 + \frac{N-1}{2N}\delta_2\right)$ $+\frac{1}{3}M_2\psi_N(\delta_3)$	$\begin{pmatrix} \frac{N-1}{6N} & \frac{N+1}{6N} & 0 \\ \frac{N+1}{6N} & \frac{N-1}{6N} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}$

b) For the repulsive cost function $c_R : X \times X \to \mathbb{R}$ given by

$$c_R(x,y) := \begin{cases} \frac{1}{d(x,y)} & \text{if } x \neq y \\ B & \text{if } x = y \end{cases} \tag{4.1}$$

for some constant $B > 1$, the unique minimizer is given by $ER^{(N)}$.

Proof. In the following the elements of the reduced Kantorovich polytope will always be interpreted as matrices. Along those lines D respectively CR corresponds to the matrix notation of d respectively c_R, i.e., $D_{ij} := d(a_i, a_j)$ respectively $CR_{ij} := c_R(a_i, a_j)$ for $i, j \in \{1, 2, \dots, \ell\}$, and $\langle \cdot, \cdot \rangle$ denotes the standard matrix scalar product.

a) Note that by non-negativity of the cost function d the objective value of an arbitrary admissible state $\mu \in \mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ is non-negative, i.e., $\langle D, \mu \rangle \geq 0$. As $EA^{(N)}$ is admissible and yields an objective value of 0, i.e., $\langle D, EA^{(N)} \rangle = 0$, it is an optimizer of the corresponding problem (1.4). Positivity of D in its off-diagonal entries and the marginal constraint ensure that $EA^{(N)}$ is the unique optimizer.

b) To prove the second assertion, we drop the marginal constraint in a reformulated version of the considered problem (1.4) and calculate the extremal elements of $\mathcal{P}_{N\text{-rep}}\left(X^2\right)$, which solve the new optimization problem. There will be a unique convex combination of the optimal extreme points of $\mathcal{P}_{N\text{-rep}}\left(X^2\right)$, namely $ER^{(N)}$, that lies in $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$. This state then corresponds to the unique minimizer of problem (1.4).
We consider the problem

$$\min_{\mu \in \mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)} \langle CR, \mu \rangle. \tag{4.2}$$

Subsequently changing the objective function to $\langle CR - \mathbb{1}, \cdot \rangle$, where all the entries of $\mathbb{1} \in \mathbb{R}^{3\times 3}$ are given by 1, and plugging in the marginal constraint allows us to reformulate (4.2) as follows

$$\max_{\mu \in \mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)} \mu_{12} + \mu_{13} + \mu_{23}. \tag{4.3}$$

By dropping the marginal constraint, further restricting the admissible set to the extreme points (3.7) of $\mathcal{P}_{N\text{-rep}}\left(X^2\right)$ and rescaling the new admissible set by N^2 leads to the new optimization problem

$$\max_{\lambda \in N\mathcal{P}_{\frac{1}{N}}(X)} \lambda_1\lambda_2 + \lambda_1\lambda_3 + \lambda_2\lambda_3 = \lambda_1\lambda_2 + (N - \lambda_3)\lambda_3. \tag{4.4}$$

Assuming $(\lambda_1^*, \lambda_2^*, \lambda_3^*)$ is an optimizer of problem (4.4), then elementary calculations show that $(\lambda_1^*, \lambda_2^*)$ fulfills

$$(\lambda_1^*, \lambda_2^*) \in \begin{cases} \left\{\left(\frac{r}{2}, \frac{r}{2}\right)\right\} & \text{if } r \text{ is even} \\ \left\{\left(\frac{r-1}{2}, \frac{r+1}{2}\right), \left(\frac{r+1}{2}, \frac{r-1}{2}\right)\right\} & \text{if } r \text{ is odd,} \end{cases} \tag{4.5}$$

where $r := N - \lambda_3^*$. Otherwise $(\lambda_1^*, \lambda_2^*, \lambda_3^*)$ would not be optimal. Note that (4.4) always admits a maximizer as $\mathcal{P}_{\frac{1}{N}}(X)$ is finite.
This allows us to identify problem (4.4) with the one-parameter optimization problem

$$\max\left\{\max_{r\in 2\mathbb{N}_0, r\leq N} -\frac{3}{4}r^2 + Nr, \max_{r\in 2\mathbb{N}_0+1, r\leq N} -\frac{3}{4}r^2 + Nr - \frac{1}{4}\right\}. \tag{4.6}$$

Elementary calculations reveal that $r \in \{0, 1, \ldots, N\}$ is optimal regarding (4.6) if and only if

$$r \in \begin{cases} \{2m\} & \text{if } N = 3m \text{ for } m \in \mathbb{N}_0 \\ \{2m, 2m+1\} & \text{if } N = 3m+1 \text{ for } m \in \mathbb{N}_0 \\ \{2m+1, 2m+2\} & \text{if } N = 3m+2 \text{ for } m \in \mathbb{N}_0. \end{cases}$$

It immediately follows that $\lambda \in N\mathcal{P}_{\frac{1}{N}}(X)$ is optimal with respect to problem (4.4) if and only if

$$\lambda \in \begin{cases} N\left\{\left(\frac{m}{N}, \frac{m}{N}, \frac{m}{N}\right)\right\} & \text{if } N = 3m \text{ for } m \in \mathbb{N}_0 \\ N\left\{\left(\frac{m}{N}, \frac{m}{N}, \frac{m+1}{N}\right), \left(\frac{m}{N}, \frac{m+1}{N}, \frac{m}{N}\right), \left(\frac{m+1}{N}, \frac{m}{N}, \frac{m}{N}\right)\right\} & \text{if } N = 3m+1 \text{ for } m \in \mathbb{N}_0 \\ N\left\{\left(\frac{m}{N}, \frac{m+1}{N}, \frac{m+1}{N}\right), \left(\frac{m+1}{N}, \frac{m}{N}, \frac{m+1}{N}\right), \left(\frac{m+1}{N}, \frac{m+1}{N}, \frac{m}{N}\right)\right\} & \text{if } N = 3m+2 \text{ for } m \in \mathbb{N}_0. \end{cases} \tag{4.7}$$

Recall that (3.7) allows us (after dropping the factor N in (4.7)) to identify the given maximizers with exactly those extremal elements of $\mathcal{P}_{N\text{-rep}}\left(X^2\right)$ that maximize the sum of their off-diagonal entries. As by Minkowski's theorem any element of $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ can be written as convex combination of the extreme points of $\mathcal{P}_{N\text{-rep}}\left(X^2\right)$ and in any of the considered cases in (4.7) there are unique coefficients given by

$$\begin{aligned} \alpha_1 &= 1 \quad \text{if } N = 3m \text{ for } m \in \mathbb{N}_0 \\ \alpha_1 = \alpha_2 = \alpha_3 &= \frac{1}{3} \quad \text{else} \end{aligned}$$

that allow us to write $\overline{\lambda}$ as a convex combination of the respective optimizers in (4.7), it is easy to see that $ER^{(N)}$ as defined in Table 1 is the unique optimizer of (4.3) and thereby (4.2) for any $N \geq 2$. □

Remark 4.2. It is an immediate consequence of Theorem 4.1 that

$$\gamma_{GS} := \frac{1}{3} S\delta_{11\dots1} + \frac{1}{3} S\delta_{22\dots2} + \frac{1}{3} S\delta_{33\dots3}$$

respectively

$$\gamma_C := \frac{1}{3} S\delta_{1\tau(1)\dots\tau^{(N-1)}(1)} + \frac{1}{3} S\delta_{2\tau(2)\dots\tau^{(N-1)}(2)} + \frac{1}{3} S\delta_{3\tau(3)\dots\tau^{(N-1)}(3)},$$

where $\tau : \{1,2,3\} \to \{1,2,3\}$ is the cyclic permutation defined by $\tau(1) = 2$, $\tau(2) = 3$, $\tau(3) = 1$ and $\tau^{(i)}$ denotes the i-th composition of τ with itself, is a solution to the OT problem (1.1) for the Gangbo-Święch cost function $c_{GS} : X^N \to \mathbb{R}$ defined by

$$c_{GS}(x_1, \dots, x_N) := \sum_{1 \le i < j \le N} d(x_i, x_j) \tag{4.8}$$

respectively the Coulomb cost function $c_C : X^N \to \mathbb{R}$ defined by

$$c_C(x_1, \dots, x_N) := \sum_{1 \le i < j \le N} c_R(x_i, x_j).$$

Here (4.8) is a discretization of the pair-cost considered in [22]. Note that one could replace $d(\cdot,\cdot)$ in (4.8) with $d(\cdot,\cdot)^p$ for any $p > 1$ without changing c_{GS}.

Next, we examine the behavior of the sequences $\left(EA^{(N)}\right)_{N\ge2}$, $\left(ER^{(N)}\right)_{N\ge2}$, $\left(E12^{(N)}\right)_{N\ge2}$, $\left(E13^{(N)}\right)_{N\ge2}$ and $\left(E23^{(N)}\right)_{N\ge2}$ for N tending to ∞. Taking a look at the right column in Table 1, it is easy to see that the following holds true.

$$EA^{(N)} \xrightarrow{N\to\infty} \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix} =: EA^{(\infty)}, ER^{(N)} \xrightarrow{N\to\infty} \begin{pmatrix} \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \\ \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \\ \frac{1}{9} & \frac{1}{9} & \frac{1}{9} \end{pmatrix} =: ER^{(\infty)},$$

$$E12^{(N)} \xrightarrow{N\to\infty} \begin{pmatrix} \frac{1}{6} & \frac{1}{6} & 0 \\ \frac{1}{6} & \frac{1}{6} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix} =: E12^{(\infty)}$$

Here $E12^{(\infty)}$ assumes again an exemplary role and $E13^{(\infty)}$ as well as $E23^{(\infty)}$ are defined in an analogous manner. One can express these 'limit extreme points' in a more probabilistic manner $EA^{(\infty)} = \frac{1}{3}\delta_{11} + \frac{1}{3}\delta_{22} + \frac{1}{3}\delta_{33}$, $ER^{(\infty)} = \left(\frac{1}{3}\delta_1 + \frac{1}{3}\delta_2 + \frac{1}{3}\delta_3\right) \otimes \left(\frac{1}{3}\delta_1 + \frac{1}{3}\delta_2 + \frac{1}{3}\delta_3\right)$ as well as $E12^{(\infty)} = \frac{2}{3}\left(\frac{1}{2}\delta_1 + \frac{1}{2}\delta_2\right) \otimes \left(\frac{1}{2}\delta_1 + \frac{1}{2}\delta_2\right) + \frac{1}{3}\delta_3$ corresponding to the 'Abstract Notation'-column in Table 1 via (3.7).

In the following, $\mathcal{D}_{\infty\text{-rep},\overline{\lambda}}$ will denote the convex hull of these 'limit extreme points', i.e.,

$$\mathcal{D}_{\infty\text{-rep},\overline{\lambda}} = \operatorname{conv}\left(\left\{EA^{(\infty)}, ER^{(\infty)}, E12^{(\infty)}, E13^{(\infty)}, E23^{(\infty)}\right\}\right). \tag{4.9}$$

For an illustration of $\mathcal{D}_{\infty\text{-rep},\overline{\lambda}}$ see Fig. 7.

It was proven in [19] that N-representability becomes an increasingly stringent condition as N grows, in more detail, $\mathcal{P}_{N\text{-rep}}\left(X^2\right) \subseteq \mathcal{P}_{\hat{N}\text{-rep}}\left(X^2\right)$ for any $N \ge \hat{N} \ge 2$. It follows immediately that the reduced Kantorovich polytope for N marginals and 3 states $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ is contained in the reduced Kantorovich polytope $\mathcal{P}_{\hat{N}\text{-rep},\overline{\lambda}}\left(X^2\right)$ for $\hat{N}$ marginals and 3 states. As $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ is closed and convex, $\mathcal{D}_{\infty\text{-rep},\overline{\lambda}}$ is a subset of the reduced Kantorovich polytope $\mathcal{P}_{N\text{-rep},\overline{\lambda}}\left(X^2\right)$ for any number $N \ge 2$ of marginals and 3 sites. In summary we get the following chain of inequalities

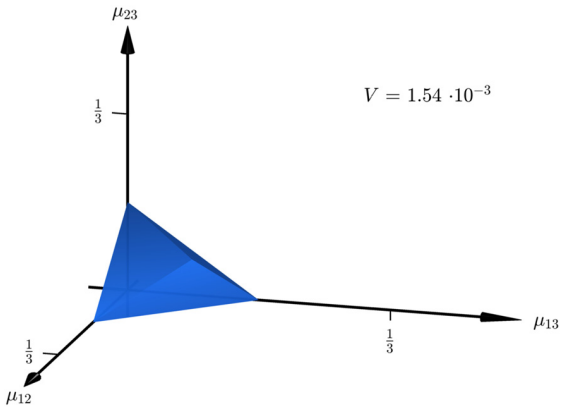

Fig. 7. The diamond-shaped polytope $\mathcal{D}_{\infty\text{-rep},\overline{\lambda}}$, as defined in (4.9), is depicted in blue. The elements $(\mu_{ij})_{i,j=1}^{3}$ of the polytope are parametrized by their off-diagonal entries μ_{12}, μ_{13} and μ_{23}. The volume of $\mathcal{D}_{\infty\text{-rep},\overline{\lambda}}$ is indicated in the upper-right corner.

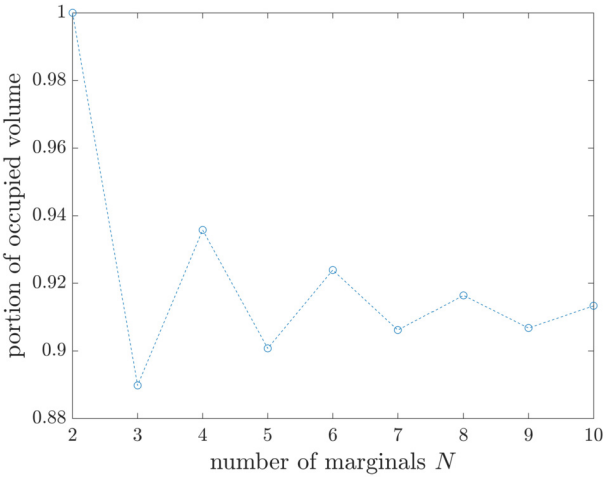

Fig. 8. The volumetric ratio, reduced Monge polytope to reduced Kantorovich polytope for N marginals and 3 states, is depicted in dependency of the number of marginals N.

$$\min_{\mu\in\mathcal{P}_{N\text{-rep},\overline{\lambda}}(X^2)} V[\mu] \leq \min_{\mu\in\mathcal{P}_{(N+1)\text{-rep},\overline{\lambda}}(X^2)} V[\mu] \leq \cdots \leq \min_{\mu\in\mathcal{D}_{\infty\text{-rep},\overline{\lambda}}} V[\mu] \leq V[ER^{(\infty)}], \tag{4.10}$$

where $V[\mu] := \int_{X^2} v(x,y)d\mu(x,y)$. The inequalities (4.10) show that for any number of marginals $N \geq 2$ we can find an upper bound of the optimal value in (1.4) by computing the objective value of the 'attractive limit extreme point' $EA^{(\infty)}$, the 'repulsive limit extreme point' $ER^{(\infty)}$ and the 'axis limit extreme points' $E12^{(\infty)}$, $E13^{(\infty)}$ as well as $E23^{(\infty)}$ and choosing the smallest one. Note that this improves the mean field approximation $V[ER^{(\infty)}]$, which is commonly used in physics. Here one usually considers repulsive pair-costs $v : X \times X \to \mathbb{R}$.

Finally, we note that the volume portion of the reduced Kantorovich polytope that is occupied by the reduced Monge polytope exhibits oscillatory behavior with decreasing amplitude when interpreted as a function of the number of marginals N, see Fig. 8. The considered volumetric ratio oscillates around a value above 0.9 where even marginals when directly compared to the odd marginals produce a higher ratio. In the sense that an optimizer in the occupied volume yields the existence of a Monge-solution, the Monge ansatz seems to be 'better' for an even number of marginals.

5. Lower bound on extremal coefficients

The results of this section were achieved in the pursuit of a generalization of Theorem 4.1. When replacing c_R in (4.1) with a general repulsive interaction, the optimization process no longer boils down to a maximization of the sum of off-diagonal entries. Even subtle differences in the off-diagonal cost coefficients could influence the optimization. Still, our intuition tells us, that - if the repulsion is strong enough - it is still best to distribute the $\frac{1}{N}$-quantized entries as uniformly as possible among the given ℓ sites. But what is the reason behind the non-existence of cases for which it is best to attain an unevenly distributed configuration with a very small probability? The answer is given in the following theorem.

Theorem 5.1 *(lower bound on extremal coefficients). Assume α to be an extreme point of the polytope P_{coef}, as defined in (2.9). Then each nonzero entry of α is bigger than or equal to $\frac{1}{\ell N^{\ell-1}}$, i.e., for all $\nu \in \left\{1, \dots, \binom{N+\ell-1}{N}\right\}$*

$$\alpha_\nu \neq 0 \to \alpha_\nu \geq \frac{1}{\ell N^{\ell-1}}.$$

Proof. Let α be an arbitrary extreme point of the polytope P_{coef}. Then - as discussed in Section 2 - the nonzero entries of α indicate a selection of columns of A that are linearly independent.

In case the cardinality of this selection is strictly less than ℓ, we add a suitable choice of elements of $\mathcal{P}_{\frac{1}{N}}(X)$ in order to form a basis B of $\mathbb{R}^\ell$. Otherwise the present selection already constitutes such a basis B. In the following A_B will denoted the $\ell \times \ell$-submatrix of A that consists of exactly those columns that are contained in B, accordingly α_B is the subvector of α that is reduced to those entries that correspond to elements of B. By construction of B,

$$A_B \alpha_B = \overline{\lambda} \tag{5.1}$$

holds. As A_B is invertible, (5.1) is equivalent to

$$\alpha_B = A_B^{-1}\overline{\lambda} = \frac{1}{\det(A_B)} C^T \overline{\lambda} \tag{5.2}$$

with C denoting the cofactor matrix of A_B. Each entry of C is the product of a sign factor and the determinant of an $(\ell-1)\times(\ell-1)$-submatrix of A_B. As the columns of A_B represent $\frac{1}{N}$-quantized probability measures, the entries of A_B are integer multiples of $\frac{1}{N}$ and therefore the entries of C are integer multiples of $\frac{1}{N^{\ell-1}}$. Utilizing Hadamard's inequality, one easily sees that $|\det(A_B)| \leq 1$ holds. Finally we recall that each of the entries of $\overline{\lambda}$ is given by $\frac{1}{\ell}$. Consequently, for the i-th entry $(\alpha_B)_i$ of α_B (with $i \in \{1, \dots, \ell\}$) the following holds

$$(\alpha_B)_i = |(\alpha_B)_i| = \frac{1}{|\det(A_B)|} |(C^T\overline{\lambda})_i| \geq |(C^T\overline{\lambda})_i| = \frac{k_i}{\ell N^{\ell-1}}$$

with k_i being some non-negative integer. As k_i is zero if and only if $(\alpha_B)_i$ is zero the proof of Theorem 5.1 is complete. □

We do not expect the lower bound established in Theorem 5.1 to be sharp. The key to unlocking an improvement and even potentially a quantization of the extremal coefficients lies in a deeper analysis of (5.2), particularly the entries of the cofactor matrix C and how they relate to one another as well as the determinant of A_B itself. This analysis, however, lies beyond the scope of this paper and we consider it to be subject to future research.

We consider Theorem 5.1 to be the bedrock the upcoming theorem is built upon.

Theorem 5.2 *(support of optimal couplings regarding 'repulsive OT problems'). Let $N \geq 2$ be the number of marginals and $\ell \geq 2$ the number of states. Consider the OT problem*

$$\text{Minimize} \int_{X^N} \sum_{1 \leq i < j \leq N} v(x_i, x_j) \mathrm{d}\gamma(x_1, \dots, x_N) \text{ over } \gamma \in \mathcal{P}_{sym}\left(X^N\right) \text{ subject to } \gamma \mapsto \overline{\lambda}, \tag{5.3}$$

with $v : X^2 \to \mathbb{R}$ being a symmetric pair-potential that fulfills $v(x,x) = B$ for all $x \in X$ and some constant

$$B > \left(\max_{x_i \neq x_j} v(x_i, x_j)\right) + N^{\ell+1}\ell \left(\left(\max_{x_i \neq x_j} v(x_i, x_j)\right) - \left(\min_{x_i \neq x_j} v(x_i, x_j)\right)\right).$$

Let γ be an optimizer of the considered problem. If γ gives mass to a point $(x_1, \dots, x_N) \in X^N$ then each state $a \in X$ appears either $\left\lfloor \frac{N}{\ell} \right\rfloor$ or $\left\lceil \frac{N}{\ell} \right\rceil$ times in the given tupel, i.e., $|\{i : x_i = a\}| \in \left\{ \left\lfloor \frac{N}{\ell} \right\rfloor, \left\lceil \frac{N}{\ell} \right\rceil \right\}$ for all $a \in X = \{a_1, \dots, a_\ell\}$.

(Note that in case $\frac{N}{\ell} = k$ for some $k \in \mathbb{N}$ each state appears exactly k times.)

Proof. Based on the proof of Theorem 4.1 assertion b) we start off by rewriting the objective function, i.e., the function that is to be minimized. Successively using the pairwise symmetric structure of the cost function, identifying the function $v : X^2 \to \mathbb{R}$ as well as the measure $M_2\gamma$ on X^2 with their respective matrix counterparts $(v_{ij})_{i,j=1}^{\ell}, ((M_2\gamma)_{ij})_{i,j=1}^{\ell}$ with $v_{ij} := v(a_i, a_j), (M_2\gamma)_{ij} := M_2\gamma(\{(a_i, a_j)\})$ and finally utilizing the marginal constraint allows us to write the objective value of any admissible γ independent of its diagonal-entries.

$$\begin{aligned}
\int_{X^N} \sum_{1 \leq i < j \leq N} v(x_i, x_j) \mathrm{d}\gamma(x_1, \dots, x_N) &= \binom{N}{2} \int_{X^2} v(x,y) \mathrm{d}(M_2\gamma)(x,y) \\
&= \binom{N}{2} \sum_{i,j=1}^{\ell} v_{ij} (M_2\gamma)_{ij} \\
&= \binom{N}{2} \left(B + \sum_{\substack{i,j=1 \\ i \neq j}}^{\ell} (v_{ij} - B)(M_2\gamma)_{ij} \right) \\
&= \binom{N}{2} B + \binom{N}{2} \sum_{\substack{i,j=1 \\ i \neq j}}^{\ell} (v_{ij} - B)(M_2\gamma)_{ij}
\end{aligned}$$

Now one easily sees that

$$\text{Maximize } C_v[\gamma] := \sum_{\substack{i,j=1 \\ i \neq j}}^{\ell} (B - v_{ij})(M_2\gamma)_{ij} \text{ over } \gamma \in \mathcal{P}_{\text{sym}}\left(X^N\right) \text{ subject to } \gamma \mapsto \overline{\lambda}, \tag{5.4}$$

is an equivalent problem formulation, in the sense that any admissible γ is optimal with respect to (5.3) if and only if it is optimal with respect to the problem at hand.

Now let γ be such an optimizer solving the problem stated in the considered theorem (5.3) as well as the problem given by (5.4). Firstly, we assume γ to be an extreme point of the symmetric Kantorovich polytope for N marginals and ℓ states, i.e., the set of admissible trial states. Then by Corollary 2.3 the to γ corresponding coefficient vector $\alpha = R^{-1}\gamma$ is itself an extreme point of the polytope P_{coef}. α then fulfills

$$\gamma = R\alpha = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda \psi_N(\lambda). \tag{5.5}$$

Recall that $\psi_N(\lambda)$ denotes the (uniquely determined) symmetrized Dirac-measure (2.5) with one-point marginal λ. Readers feeling lost regarding the present notation are advised to take a look back at Section 2, particularly pages 7-8. By linearity of $C_v[\cdot]$ it holds

$$C_v[\gamma] = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda C_v[\psi_N(\lambda)].$$

As already stated in (3.7), $M_2\psi_N(\lambda) = \frac{N}{N-1}\lambda \otimes \lambda - \frac{1}{N-1}(\mathrm{id}, \mathrm{id})\#\lambda$ holds yielding

$$C_v[\psi_N(\lambda)] = \frac{N}{N-1} \sum_{\substack{i,j=1 \\ i \neq j}}^{\ell} (B - v_{ij})\lambda_i\lambda_j.$$

In the following we will take a closer look at the objective value $C_v[\cdot]$ of the extremal symmetric probability measures $\psi_N(\lambda)$ by investigating the behavior of the function $f : \mathcal{P}_{\frac{1}{N}}(X) \to \mathbb{R}$ given by

$$f(\lambda) = C_v[\psi_N(\lambda)].$$

We split f into a dominant and a submissive term, denoted by d and s, respectively.

$$f(\lambda) = d(\lambda) + s(\lambda)$$

$$\text{with } d(\lambda) := \frac{N}{N-1} \sum_{\substack{i,j=1 \\ i \neq j}}^{\ell} (B - v_{i^*j^*})\lambda_i\lambda_j \text{ and } s(\lambda) := \frac{N}{N-1} \sum_{\substack{i,j=1 \\ i \neq j}}^{\ell} (v_{i^*j^*} - v_{ij})\lambda_i\lambda_j,$$

where i^*, j^* are indices fulfilling $v_{i^*j^*} = \max\limits_{\substack{i,j \\ i \neq j}} v_{ij}$.

Hereby, a more compact manner to write d is given by

$$d(\lambda) = \lambda^T D\lambda \quad \text{for} \quad D = (D_{ij}) \in \mathbb{R}^{\ell \times \ell} \text{ defined by } D_{ij} = \begin{cases} 0 & \text{for } i = j \\ B' & \text{for } i \neq j \end{cases}$$

$$\text{with } B' := \frac{N}{N-1}(B - v_{i^*j^*}). \tag{5.6}$$

By Taylor-expanding d at the uniform probability measure $\overline{\lambda}$ and utilizing the geometry of the matrix D embodied by its eigenspaces one easily sees that

$$d(\lambda) = d(\overline{\lambda}) - B'|\lambda - \overline{\lambda}|^2$$

with $|\cdot|$ denoting the Euclidean norm in $\mathbb{R}^\ell$. Now elementary arguments and calculations reveal the following. $\hat{\lambda}$ maximizes $d(\cdot)$ among the $\frac{1}{N}$-quantized probability measures $\mathcal{P}_{\frac{1}{N}}(X)$ if and only if r entries of $\hat{\lambda}$ are given by $\frac{m+1}{N}$ and the remaining $(\ell - r)$ entries correspond to $\frac{m}{N}$ with $m := \left\lfloor \frac{N}{\ell} \right\rfloor$ and $r := N - m\ell$. Any deviating $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$, i.e., any λ not obeying these restrictions regarding its entries, decreases the value of $d(\cdot)$ by at least $\frac{B'}{N^2}$. That is, for a rule-abiding $\hat{\lambda}$ and a deviating λ it holds, $d(\hat{\lambda}) - d(\lambda) \geq \frac{B'}{N^2}$.

Now we return to the consideration of γ which is assumed to be a solution of the problems (5.3) and (5.4) as well as (for now) an extreme point of both sets of admissible trial states. α denotes the coefficient vector underlying the representation of γ as a convex combination of extremal symmetric probability measures as given in (5.5). As already mentioned above, α is itself an extreme point of P_{coef}.

The next step is to derive a contradiction starting from the assumption $\alpha_\lambda > 0$ for a $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$ that deviates from the 'entry laws' described above. In the following this deviating $\frac{1}{N}$-quantized probability measure will be denoted by $\tilde{\lambda}$. Let further $\hat{\gamma} = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \hat{\alpha}_\lambda \psi_N(\lambda)$ be an admissible trial state whose coefficient vector $\hat{\alpha}$ only gives mass to law-abiding λs $\in \mathcal{P}_{\frac{1}{N}}(X)$. It is easy to see that such a state always exists. Then - with i_*, j_* denoting indices that fulfill $v_{i_* j_*} = \min_{\substack{i,j \\ i \neq j}} v_{ij}$ and $\hat{\lambda}$ denoting an arbitrary law-abiding $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$ - it holds

$$\begin{aligned}
C_v[\hat{\gamma}] - C_v[\gamma] &= \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \hat{\alpha}_\lambda (d(\lambda) + s(\lambda)) - \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda (d(\lambda) + s(\lambda)) \\
&= \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} (\hat{\alpha}_\lambda - \alpha_\lambda) s(\lambda) + \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \hat{\alpha}_\lambda d(\lambda) - \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda d(\lambda) \\
&\geq -\frac{N}{N-1}(v_{i^* j^*} - v_{i_* j_*}) + d(\hat{\lambda}) - \frac{1}{\ell N^{\ell-1}} d(\tilde{\lambda}) - \left(1 - \frac{1}{\ell N^{\ell-1}}\right) d(\hat{\lambda}) \\
&= -\frac{N}{N-1}(v_{i^* j^*} - v_{i_* j_*}) + \frac{1}{\ell N^{\ell-1}}(d(\hat{\lambda}) - d(\tilde{\lambda})) \\
&\geq -\frac{N}{N-1}(v_{i^* j^*} - v_{i_* j_*}) + \frac{B'}{\ell N^{\ell+1}}.
\end{aligned}$$

Hereby, elementary estimates gave us a lower bound on how much $\hat{\gamma}$ might loose to γ regarding the submissive function s. We further took advantage of the fact that all law-abiding λs produce the same dominant value $d(\hat{\lambda})$. The key step, however, is to establish an upper-bound on the portion of the objective value of γ which corresponds to the dominant function d. This bound is based on Theorem 5.1 as well as the priorly established fact that the law-abiding elements of $\mathcal{P}_{\frac{1}{N}}(X)$ are exactly those that maximize $d(\cdot)$ and $\tilde{\lambda}$ falls short by at least $\frac{B'}{N^2}$ in comparison.

Combining the definition of B' in (5.6) with the assumption on B now yields $C[\hat{\gamma}] > C[\gamma]$. This finalizes the contradiction. Hence, positivity of a coefficient α_λ implies that λ fulfills the 'entry laws'. Recalling the representation (5.5) of γ as well as the definition of ψ_N now reveals that the statement of Theorem 5.2 is true for the extremal γ.

As any non-extremal optimizer may be written as a convex combination of extremal ones the proof of Theorem 5.2 is complete. □

Remark 5.3. A question that naturally pops up in the given context is: What happens in the *non-symmetric case*? Hereby the *non-symmetric case* refers to optimal transport problems of the form

$$\text{Minimize} \int_{X^N} c(x_1, \dots, x_N) d\gamma(x_1, \dots, x_N)$$

over all probability measures γ on X^N whose i-th marginal is given by $\lambda^{(i)}$. Note that the N marginals $\lambda^{(1)}, \dots, \lambda^{(N)}$ are now allowed to differ from one another. In order to gain some understanding of this setting, we retrace some of the steps taken in this paper.

When formulating the corresponding coefficient polytope, one has to bare in mind, that now one works with Dirac measures instead of symmetrized Dirac measures. This results in an increase of entries in a coefficient vector α as well as an increase of the number of constraints.

We believe that in this unsymmetrized setting it is still possible to establish a lower bound on extremal coefficients - at least if the given marginals are quantized. Hereby a choice of N marginals is understood as quantized if there exists a positive real number q such that each one of the $N \cdot \ell$ marginal probabilities is an integer multiple of q. The lower bound then not only depends on the number of marginals N and the number of states ℓ but also on the quantization of the marginals q.

The proof of Theorem 5.2 heavily relies on the symmetry of the considered setting for example in terms of the polynomial formula regarding extremal N-representable two-point measures. As we do not see a straightforward manner to generalize the corresponding result we classify this task as subject to future research; a good question to initialize this research would be: What kind of cost functions does one want to consider?

The following consequence of Theorem 5.2 is a generalization of Theorem 4.1 assertion b).

Corollary 5.4. *We consider the reduced multi-marginal OT problem (1.4) for $N \geq 2$ marginals and $\ell = 3$ sites.*

For any symmetric cost function $v : X \times X \to \mathbb{R}$ that fulfills $v(x,x) = B$ for all $x \in X$ and some constant

$$B > \left(\max_{x_i \neq x_j} v(x_i, x_j) \right) + N^{\ell+1} \ell \left(\left(\max_{x_i \neq x_j} v(x_i, x_j) \right) - \left(\min_{x_i \neq x_j} v(x_i, x_j) \right) \right)$$

the unique minimizer is given by $ER^{(N)}$.

Proof. We start off with a change of venue and consider the 'unreduced' problem version (5.3). With the number of states ℓ being equal to three, Theorem 5.2 reduces the points an optimizer might give mass to already to such an extent that the optimizer's uniqueness follows. The two-point marginal of said optimizer is given by $ER^{(N)}$ which inherits the status of a unique optimizer from its representing measure. □

Recall that $ER^{(N)}$ is of Monge-type. Consequently, Corollary 5.4 provides a class of repulsive costs yielding a unique Monge optimizer. All of these examples, however, are set in a finite state space X consisting only of three elements.

The following discussion concerns a lift of Corollary 5.4 to a given $\ell' > 3$. So now the focus lies on the question whether or not the suitably adapted statement of Corollary 5.4 holds for any number of marginals $N \geq 2$ when paired with ℓ'. Note that in the introduction specific 'pairable' N's for any $\ell' > 3$ are given.

Already when increasing the number of states to $\ell = 4$ and keeping the number of marginals at $N = 2$ the representing measure of the 'peak' that is $ER^{(N)}$ blossoms into multiple extreme points of the symmetric Kantorovich polytope for 2 marginals and 4 states. That is, there exist multiple extreme points of the symmetric Kantorovich polytope for 2 marginals and 4 states that are in line with the support-restrictions provided by Theorem 5.2. Hence, the proof of Corollary 5.4 can not be lifted to the case of $\ell = 4$ states. However, we suspect that each one of the 'in line'-extreme points of the symmetric Kantorovich polytope is of Monge-type for $\ell = 4$ as well as $\ell = 5$ states and any number of marginals $N \geq 2$. So even though one is not able to identify a specific extreme point as unique optimizer we believe that the support-restriction suffices to at least classify the optimizer(s) as Monge. When increasing the number of states to $\ell = 6$ this door closes as

$$\frac{1}{4} S\delta_{123} + \frac{1}{4} S\delta_{145} + \frac{1}{4} S\delta_{246} + \frac{1}{4} S\delta_{356} \tag{5.7}$$

is an 'in-line' extreme point of the symmetric Kantorovich polytope for $N = 3$ marginals and $\ell = 6$ states that is not of Monge-type.

An interesting question to pose is now whether or not non-Monge states of form (5.7) remain extremal when moving from N-point to two-point costs. A more general question - we believe - this section accumulates to is the following. Given a symmetric pair-cost v that fulfills the condition on its diagonal stated in Theorem 5.2, what geometric attributes of the off-diagonal part of v decide whether the optimizer is of Monge-type or not - provided it exists a unique optimizer.

References

[1] M. Agueh, G. Carlier, Barycenters in the Wasserstein space, SIAM J. Math. Anal. 43 (2) (2011) 904–924.
[2] M. Beiglböck, P. Henry-Labordère, F. Penkner, Model-independent bounds for option prices - a mass transport approach, Finance Stoch. 17 (3) (7 2013) 477–501.
[3] D.P. Bertsekas, Convex Optimization Theory, 1 edition, Athena Scientific, Belmont, MA, 2009.
[4] G. Birkhoff, Tres observaciones sobre el algebra lineal, Rev. Univ. Nac. Tucumán Ser. A 5 (3) (1946) 147–151.
[5] R. Burkard, M. Dell'Amico, S. Martello, Assignment Problems. Revised Reprint, Society for Industrial and Applied Mathematics, 2012.
[6] G. Buttazzo, L. De Pascale, P. Gori-Giorgi, Optimal-transport formulation of electronic density-functional theory, Phys. Rev. A 85 (6) (2012) 062502.
[7] G. Carlier, On a class of multidimensional optimal transportation problems, J. Convex Anal. 10 (2) (02 2003) 517–529.
[8] G. Carlier, I. Ekeland, Matching for teams, Econ. Theory 42 (2) (2 2010) 397–418.
[9] G. Carlier, B. Nazaret, Optimal transportation for the determinant, ESAIM Control Optim. Calc. Var. 14 (4) (2008) 678–698.
[10] H. Chen, G. Friesecke, C.B. Mendl, Numerical methods for a Kohn-Sham density functional model based on optimal transport, J. Chem. Theory Comput. 10 (10) (2014) 4360–4368, PMID: 26588133.
[11] P.-A. Chiappori, R.J. McCann, L.P. Nesheim, Hedonic price equilibria, stable matching, and optimal transport: equivalence, topology, and uniqueness, Econom. Theory 42 (2) (Feb 2010) 317–354.
[12] M. Colombo, L. De Pascale, S. Di Marino, Multimarginal optimal transport maps for one-dimensional repulsive costs, Can. J. Math. 67 (2013) 350–368.
[13] M. Colombo, F. Stra, Counterexamples in multimarginal optimal transport with Coulomb cost and spherically symmetric data, Math. Models Methods Appl. Sci. 26 (06) (2016) 1025–1049.
[14] C. Cotar, G. Friesecke, C. Klüppelberg, Density functional theory and optimal transportation with Coulomb cost, Commun. Pure Appl. Math. 66 (4) (2013) 548–599.
[15] C. Cotar, G. Friesecke, B. Pass, Infinite-body optimal transport with Coulomb cost, Calc. Var. Partial Differ. Equ. 54 (1) (9 2015) 717–742.
[16] J. Csima, Multidimensional stochastic matrices and patterns, J. Algebra 14 (2) (1970) 194–202.
[17] S. Di Marino, A. Gerolin, L. Nenna, Optimal Transportation Theory with Repulsive Costs, ArXiv e-prints, 12 2015.
[18] G. Friesecke, A simple counterexample to the Monge ansatz in multi-marginal optimal transport, convex geometry of the set of Kantorovich plans, and the Frenkel-Kontorova model, ArXiv e-prints, 8 2018.
[19] G. Friesecke, C.B. Mendl, B. Pass, C. Cotar, C. Klüppelberg, N-density representability and the optimal transport limit of the Hohenberg-Kohn functional, J. Chem. Phys. 139 (16) (2013) 164109.
[20] G. Friesecke, D. Vögler, Breaking the curse of dimension in multi-marginal Kantorovich optimal transport on finite state spaces, SIAM J. Math. Anal. 50 (4) (2018) 3996–4019.
[21] A. Galichon, P. Henry-Labordère, N. Touzi, A stochastic control approach to no-arbitrage bounds given marginals, with an application to lookback options, Ann. Appl. Probab. 24 (1) (2014) 312–336.
[22] W. Gangbo, A. Święch, Optimal maps for the multidimensional Monge-Kantorovich problem, Commun. Pure Appl. Math. 51 (1) (1998) 23–45.
[23] A. Gerolin, A. Kausamo, T. Rajala, Non-existence of optimal transport maps for the multi-marginal repulsive harmonic cost, ArXiv e-prints, 5 2018.
[24] H. Heinich, Probléme de Monge pour probabilités, C. R. Math. Acad. Sci. Paris 334 (9) (12 2002) 793–795.
[25] L. Hoermander, Notions of Convexity, Birkhäuser, Boston, Massachusetts, 1994.
[26] The MathWorks Inc., MATLAB and Statistics Toolbox R2018b (MATLAB 9.5), Natick, Massachusetts, United States, 2018.
[27] Y. Khoo, L. Lin, M. Lindsey, L. Ying, Semidefinite relaxation of multimarginal optimal transport for strictly correlated electrons in second quantization, SIAM J. Sci. Comput. 42 (6) (2020) B1462–B1489.
[28] Y. Khoo, L. Ying, Convex relaxation approaches for strictly correlated density functional theory, SIAM J. Sci. Comput. 41 (4) (2019) B773–B795.
[29] V.M. Kravtsov, Combinatorial properties of noninteger vertices of a polytope in a three-index axial assignment problem, Cybern. Syst. Anal. 43 (1) (Jan 2007) 25–33.
[30] N. Linial, Z. Luria, On the vertices of the d-dimensional Birkhoff polytope, Discrete Comput. Geom. 51 (1) (Jan 2014) 161–170.
[31] A. Moameni, B. Pass, Solutions to multi-marginal optimal transport problems concentrated on several graphs, ESAIM Control Optim. Calc. Var. 23 (2) (2017) 551–567.
[32] B. Pass, Uniqueness and Monge solutions in the multimarginal optimal transportation problem, SIAM J. Math. Anal. 43 (6) (2011) 2758–2775.
[33] B. Pass, On the local structure of optimal measures in the multi-marginal optimal transportation problem, Calc. Var. Partial Differ. Equ. 43 (3) (Mar 2012) 529–536.

[34] B. Pass, Remarks on the semi-classical Hohenberg-Kohn functional, Nonlinearity 26 (9) (2013) 2731.
[35] The polymake team, polymake 3.2, Discrete Mathematics/Geometry Institut für Mathematik der Technischen Universität Berlin, Straße des 17. Juni 136, Berlin, Deutschland, 2018.
[36] J. Rabin, G. Peyré, J. Delon, M. Bernot, Wasserstein barycenter and its application to texture mixing, in: Scale Space and Variational Methods in Computer Vision, 2012, pp. 435–446.
[37] R.T. Rockafellar, Convex Analysis, Princeton University Press, New Jersey, 1997.
[38] M. Seidl, S. Di Marino, A. Gerolin, L. Nenna, K.J.H. Giesbertz, P. Gori-Giorgi, The strictly-correlated electron functional for spherically symmetric systems revisited, ArXiv e-prints, 2017.
[39] F.C.R. Spieksma, Multi Index Assignment Problems: Complexity, Approximation, Applications, Springer US, Boston, MA, 2000, pp. 1–12.
[40] C. Villani, Optimal Transport: Old and New, Springer Verlag, Berlin, Heidelberg, 2009.
[41] J. von Neumann, A certain zero-sum two person game equivalent to the optimal assignment problem, Contrib. Theory Games 11 (1953) 5–12.

A.2 Genetic Column Generation: Fast Computation of High-Dimensional Multimarginal Optimal Transport Problems

Authors: Gero Friesecke, Andreas S. Schulz, and Daniela Vögler

Summary. In the present article, we introduce a simple and efficient computational method that provides approximate solutions for the discrete symmetric multi-marginal optimal transport (MMOT) problems of the following form.

$$\begin{aligned} \text{Minimize} \quad & \int_{X^N} \sum_{1\le i<j\le N} v\left(x_i, x_j\right) \mathrm{d}\gamma\left(x_1, \ldots, x_N\right) \\ \text{subject to} \quad & \gamma \in \{\gamma \in \mathcal{P}_{\mathrm{sym}}\left(X^N\right): \quad \gamma \mapsto \lambda^*\}. \end{aligned}$$

Here, X denotes a finite state space consisting of ℓ pairwise distinct elements; $v : X \times X \to \mathbb{R}$ denotes a symmetric pair-cost; $\mathcal{P}_{\mathrm{sym}}\left(X^N\right)$ denotes the set of symmetric probability measures on the product space X^N; for $\lambda^* \in \mathcal{P}\left(X\right)$, $\gamma \mapsto \lambda^*$ means that γ has equal one-point marginals λ^*. We call the presented method Genetic Column Generation.

The difficulty of solving these problems lies not in the type of optimization problem but in the size of the problem. We face linear programs whose storage complexity grows exponentially with the number of marginals N. With our motivating application being electronic structure, we are interested in the large N regime. In this regime, storing the problems at hand is infeasible. Genetic Column Generation operates within a sparse ansatz space referred to as quasi-Monge. This sparse ansatz was recently introduced by Friesecke and I in [44]. Column generation itself is a pragmatic approach towards linear programs with an extremely high number of variables and few equality constraints. It is rooted in discrete optimization and allows us to practically implement the intended evolution within the sparse ansatz space. Moving from one sparse state to the best possible successor requires solving an intermediate (often combinatorial) optimization problem. In our setting, this intermediate problem turns out to be NP-complete (see the present article for the proof) which is why we opted for moving from one sparse state to a better, but not necessarily the best possible, successor. We do this by producing a 'good', but not necessarily optimal, solution for the intermediate problem through a genetic sampling process.

Overall, Genetic Column Generation alternatingly solves small linear programs

and performs a genetic sampling step. This simple construction allowed us to solve the above-mentioned discrete symmetric MMOT problems with up to $2.58 \cdot 10^{31}$ variables. In the test cases where the optimal solution is known (which included the problem with $2.58 \cdot 10^{31}$ variables), our method proved to be exact.

Own contribution. I, Daniela Vögler, am principal author of the present article. Being principal author entails (i) taking the lead in writing the article and (ii) contributing the most to the research conducted and insights gained. In more detail, I initiated to project; I acted as liaison between operations research personified by Prof. Andreas S. Schulz and analysis personified by Prof. Gero Friesecke; I applied the method column generation from discrete optimization to the symmetric multi-marginal optimal transport problems at hand and was significantly involved in implementing the algorithm; I also worked out the proof establishing the NP-completeness of our pricing problem. I would be remiss to leave out the significant impact on the project my co-authors had. Prof. Andreas S. Schulz, for example, initially came up with the idea to apply column generation and proposed to use the clique problem to establish the NP-completeness of our pricing problem. Prof. Gero Friesecke, for example, came up with the genetic sampling method including the interpretation of children as nearest-neighbours in Wasserstein space and supplied the justification of the used discretization. This article would not have come into existence without their openness to venture out of their respective fields of expertise. I want to take this opportunity to thank them!

The publisher of the *Journal of Scientific Computing*, SIAM, permits the inclusion of this publication[1]:

> *"Author's Rights*
> *(...) The Author may reproduce and distribute the Work (including derivative works) in connection with the Author's teaching, technical collaborations, conference presentations, lectures, or other scholarly works[.]"*
>
> [Society for Industrial and Applied Mathematics (SIAM). Consent to Publish. https://www.siam.org/Portals/0/Publications/Journals/Open_Access/SIAM_Consent_to_Publish.pdf, 2022. [Online; accessed 16.05.2022].]

[1]Core Publication A.2 was first published in the Journal of Scientific Computing in 44(3), 2022, published by the Society for Industrial and Applied Mathematics (SIAM).

SIAM J. SCI. COMPUT.
Vol. 44, No. 3, pp. A1632–A1654

GENETIC COLUMN GENERATION: FAST COMPUTATION OF HIGH-DIMENSIONAL MULTIMARGINAL OPTIMAL TRANSPORT PROBLEMS*

GERO FRIESECKE†, ANDREAS S. SCHULZ‡, AND DANIELA VÖGLER†

Abstract. We introduce a simple, accurate, and extremely efficient method for numerically solving multimarginal optimal transport (MMOT) problems arising in density functional theory. The method relies on (i) the sparsity of optimal plans (for N marginals discretized by ℓ gridpoints each, general Kantorovich plans require ℓ^N gridpoints, but the support of optimizers is of size $O(\ell \cdot N)$ [G. Friesecke and D. Vögler, *SIAM J. Math. Anal.,* 50 (2018), pp. 3996–4019], (ii) the method of column generation (CG) from discrete optimization which is novel in the optimal transport context, and (iii) ideas from machine learning. The well-known bottleneck in CG consists in generating new candidate columns efficiently; we prove that in our context, finding the best new column is an NP-complete problem. To overcome this bottleneck we use a genetic learning method tailor-made for MMOT in which the dual state within CG plays the role of an "adversary" in loose similarity to Wasserstein generative adversarial networks (GANs). On a sequence of benchmark problems with up to 120 gridpoints and up to 30 marginals, our method always finds the exact optimizers. Moreover, empirically the number of computational steps needed to find them appears to scale only polynomially when both N and ℓ are simultaneously increased (while keeping their ratio fixed to mimic a thermodynamic limit of the particle system).

Key words. optimal transport, column generation, high dimension, genetic algorithm

AMS subject classifications. 65K, 90C06, 68Q17

DOI. 10.1137/21M140732X

1. Introduction.

Multimarginal optimal transport (MMOT) suffers from the curse of dimension. If the marginals are discretized by ℓ gridpoints, optimal (or candidate) Kantorovich plans for the N-marginal problem require ℓ^N gridpoint values. While powerful and successful computational schemes centered around the Sinkhorn algorithm have been developed for two-marginal problems ($N = 2$) [Cut13, Sch16, Sch19, PC19, BS20], with recent extensions to a small number of marginals [BCN16, Nen17, BCN19], the high-dimensionality of multimarginal plans forbids the use of these schemes in practice beyond a handful of marginals.

On the other hand, in recent applications of MMOT to many-electron physics [CFK13, BDPGG12], data science [AC11], and fluid dynamics [Bre89, Nen17], N corresponds, respectively, to the number of electrons in a molecule, datasets in a database, or timesteps. This makes it highly desirable to develop computational schemes for MMOT with large N. Some recent advances were made in the context of the MMOT problem arising in many-electron physics [CFK13, BDPGG12], which is our key motivating application and the focus of this paper. In [FV18] two of the current authors obtained the following rigorous sparsity result (whose ancestor is the celebrated Brenier's theorem [Bre91]): after discretization, for any marginals and costs there exist optimizers which are superpositions of at most ℓ symmetrized Dirac mea-

*Submitted to the journal's Methods and Algorithms for Scientific Computing section March 24, 2021; accepted for publication (in revised form) January 14, 2022; published electronically June 23, 2022.

https://doi.org/10.1137/21M140732X

†Department of Mathematics, Technical University of Munich, D 85748 Garching, Germany (gf@ma.tum.de, voegler@ma.tum.de).

‡Operations Research, Technical University of Munich, D 80333 Munich, Germany (andreas.s.schulz@tum.de).

sures. Moreover, the structure of optimizers was shown to be closely related to the Monge ansatz of optimal transport theory, and a two-marginal formulation of the N-marginal problem was given. In [KY19], Khoo and Ying introduced and studied a semidefinite relaxation of the two-marginal formulation and presented an algorithm for the relaxed problem. In [ACEL21, ACE21], Alfonsi et al. established existence of sparse optimizers even in the situation when the state space is kept continuous and only the marginal constraints are discretized; moreover, they proposed a constrained Lagrangian particle method for the ensuing problem. Also, let us mention a recent advance not related to MMOT, namely that smooth *two*-marginal problems in high dimension are soluble with dimension-free computational rates, with potentially exponentially dimension-dependent constants [VMR+21].

Here we present a simple and extremely efficient algorithm for MMOT which combines MMOT sparsity, methods from high-dimensional discrete optimization, and recent advances in machine learning. Numerical results show that it allows the accurate computation of optimal plans with, say, $N = 30$ marginals and $\ell = 100$ gridpoints or basis functions per dimension (i.e., $\ell^N = 10^{60}$) with MATLAB on a laptop. In benchmark examples of this size where the exact solution is known, the algorithm always finds the exact optimizers (see section 7). Moreover, empirically (see Figure 5) the number of computational steps needed to find them scales only polynomially instead of exponentially in the thermodynamic limit when both N and ℓ get large, with their ratio N/ℓ remaining constant, although we cannot offer a rigorous proof of this fact. Instead, in section 6 we show that the pricing problem which our genetic learning method addresses is NP-complete. For a related result recently posted on arXiv see [ABA20a].

Our algorithm, which we call *Genetic Column Generation* (`GenCol`), is presented in this paper in detail in the context of the multimarginal optimal transport problems arising in many-electron physics. It is based on three ideas:

- the existence of extremely sparse optimizers as first pointed out and investigated in the present context by two of the authors in [FV18]. This breaks the curse of dimension with respect to storage complexity (but at the time we could not offer any algorithm).
- the method of column generation (CG), which is well established in discrete optimization but novel in MMOT. CG is a pragmatic approach to tackling certain extremely high-dimensional problems which originated in integer programming. The latter arises when looking for Monge plans for $\ell = 2$ and N large, in which case the unknown is a pair of vectors in $\{0,1\}^N$. We note that this is exactly the opposite regime to $N = 2$, ℓ large where the Sinkhorn algorithm works most successfully.
- a simple genetic method tailor-made for MMOT to overcome the well-known bottleneck in CG of generating new candidate columns efficiently. In our context new columns represent intricate spatial many-body correlation patterns of the system which are not known a priori; these are learned with the help of an "adversary" represented by the dual state within CG, in loose similarity to Wasserstein generative adversarial networks (GANs) [ACB17].

The underlying theory is described in sections 2–5. The algorithm, which in the end is rather simple, is presented in section 5.4. Numerical results for test problems up to sizes of $\ell^N \approx 10^{60}$ are given in section 7. Applications to more complex electronic structure problems will be given elsewhere.

2. MMOT, motivation, and discretization.

Multimarginal optimal transport. Many different problems in mathematics, science, and engineering can be cast in the form of the following general MMOT problem:

Minimize a cost functional

$$C[\gamma] = \int_{X_1 \times \cdots \times X_N} c(x_1, \ldots, x_N)\, d\gamma(x_1, \ldots, x_N) \tag{2.1}$$

over N-point probability measures

$$\gamma \in \mathcal{P}(X_1 \times \cdots \times X_N) \tag{2.2}$$

subject to the marginal constraints

$$M_{X_i}\gamma = \mu_i \quad (i = 1, \ldots, N). \tag{2.3}$$

Here the X_i are metric spaces (in practice, subsets of $\mathbb{R}^d$ for continuous problems and finite sets for discrete problems), the μ_i are given Borel probability measures on X_i, $\mathcal{P}(X_1 \times \cdots \times X_N)$ denotes the set of Borel probability measures on $X_1 \times \cdots \times X_N$, $c : X_1 \times \cdots \times X_N \to \mathbb{R} \cup \{+\infty\}$ is a cost function, and the marginal of γ with respect to the ith space X_i is the probability measure on X_i defined by

$$M_{X_i}\gamma(A) = \gamma(X_1 \times \cdots \times X_{i-1} \times A \times X_{i+1} \times \cdots \times X_N) \text{ for all measurable } A \subseteq X_i.$$

Optimizers are known as *optimal plans* or *Kantorovich plans*. Both the analysis and the numerical treatment of optimal transport problems have been the subject of intensive and fruitful research, with the focus overwhelmingly on two-marginal problems ($N = 2$); see [Vil09, San15, PC19] for wide-ranging surveys.

Multimarginal problems ($N > 2$), about which much less is known, have been considered for quite some time in operations research, probability theory, analysis, and mathematical economics [Pie68, Poo94, RR98, GS98, Spi00, CMN10, BDM12]. Recently, important examples of multimarginal problems with large N have emerged independently in many-electron physics [CFK13, BDPGG12], fluid dynamics [Bre89, Nen17], and data science [AC11]. The number N of marginals corresponds, respectively, to the number of particles, timesteps, or datasets in a database, motivating the interest in large N.

Physical motivation. A central example which we want to attack in this paper is MMOT with Coulomb cost, which arises as the strongly correlated limit of density functional theory (DFT). DFT is the most widely used method for numerical electronic structure computations in physics, chemistry, and materials science; see [Bec14] for a review. The strongly correlated limit was introduced by Seidl [Sei99]. As first noted and exploited in [CFK13, BDPGG12], the limit problem is an optimal transport problem, with

$$X_i = \bar{X} \subseteq \mathbb{R}^d \,\forall i, \;\; \mu_i = \mu \,\forall i, \;\; c(x_1, \ldots, x_N) = \sum_{1 \le i < j \le N} \frac{1}{|x_i - x_j|}, \tag{2.4}$$

where $\mu : \mathbb{R}^d \to \mathbb{R}$ is the single-particle density of the system, normalized so that it integrates to 1. See [CFK18] for a rigorous derivation from the underlying quantum

many-body system. In physics one is only interested in Kantorovich plans which are symmetric with respect to the x_i (as these model N-point position densities of electrons, which are symmetric by the laws of quantum theory). This means that for all permutations σ,

$$\gamma(A_1 \times \cdots \times A_N) = \gamma(A_{\sigma(1)} \times \cdots \times A_{\sigma(N)}) \text{ for any Borel subsets } A_1, \ldots, A_N \text{ of } \bar{X}.$$

Mathematically, this restriction does not alter the optimal cost because for equal marginals and a symmetric cost c (as in (2.4)), each nonsymmetric plan gives rise to a symmetric one with the same cost, by symmetrization. Also, for a symmetric plan, any one marginal condition implies the others. Thus in the situation (2.4), denoting the set of symmetric probability measures on $\bar{X}^N$ by $\mathcal{P}_{sym}(\bar{X}^N)$ and abbreviating $M_{X_1}\gamma = M_1\gamma$, the MMOT problem (2.1)–(2.3) reduces to

$$\text{Minimize } C[\gamma] = \int_{\bar{X}^N} c(x_1, \ldots, x_N)\, d\gamma(x_1, \ldots, x_N) \tag{2.5}$$

$$\text{over } \gamma \in \mathcal{P}_{sym}(\bar{X}^N) \tag{2.6}$$

$$\text{subject to } M_1\gamma = \mu \tag{2.7}$$

(symmetric MMOT). Here c can be any symmetric function on $\bar{X}^N$.

Corrections from the strongly correlated (MMOT) limit have been demonstrated to improve the accuracy of electronic structure simulations based on DFT [FGGSDS16]; but as yet no numerical method is available which can handle this limit reliably for anything other than small test systems with a few electrons.

Discretization. A simple, in the $N = 2$ case standard, structure-preserving discretization of (2.1)–(2.3) which preserves the favorable sparsity and duality properties of optimal transport is as follows. Suppose the X_i are compact subsets of $\mathbb{R}^d$, and $c : X_1 \times \cdots \times X_N \to \mathbb{R}$ is continuous. Let

$$\mu_i^{(\nu)} = \sum_{\alpha=1}^{\ell_i(\nu)} m_{i,\alpha}^{(\nu)} \delta_{a_{i,\alpha}^{(\nu)}},\ m_{i,\alpha}^{(\nu)} \geq 0,\ a_{i,\alpha}^{(\nu)} \in X_i, \tag{2.8}$$

be any sequence of finite sums of Dirac measures converging weak* in $\mathcal{M}(X_i) = (C(X_i))^*$ to μ_i. (Such approximations always exist. For instance, if X_i is the closure of an open bounded set with smooth boundary, one may partition X_i into distinct small cells $V_{i,\alpha}^{(\nu)} = X_i \cap Q_{i,\alpha}^{(\nu)}$ where the $Q_{i,\alpha}^{(\nu)}$ are disjoint cubes in $\mathbb{R}^d$ of sidelength $1/\nu$. One now picks any representative point $a_{i,\alpha}^{(\nu)}$ in $V_{i,\alpha}^{(\nu)}$ and places all the mass from $V_{i,\alpha}^{(\nu)}$ there, i.e., one sets $m_{i,\alpha}^{(\nu)} = \mu_i(V_{i,\alpha}^{(\nu)})$.) Then any plan $\gamma \in \mathcal{P}(X_1 \times \cdots \times X_N)$ satisfying the marginal conditions (2.3) must be of the following form, where we omit the superscript ν:

$$\gamma = \sum_{i_1=1}^{\ell_1} \cdots \sum_{i_N=1}^{\ell_N} \gamma_{i_1 \ldots i_N} \delta_{a_{1,i_1}} \otimes \cdots \otimes \delta_{a_{N,i_N}}, \tag{2.9}$$

so such a plan can be viewed as a tensor $(\gamma_{i_1 \ldots i_N})$ of order N, and (2.1)–(2.3) reduces

to the discrete problem

$$\text{Minimize } C[\gamma] = \sum_{i_1=1}^{\ell_1} \cdots \sum_{i_N=1}^{\ell_N} \gamma_{i_1 \ldots i_N} c(a_{1,i_1}, \ldots, a_{N,i_N}) \tag{2.10}$$

$$\text{subject to } \sum_{i_j \,:\, j \neq k} \gamma_{i_1 \ldots i_k \ldots i_N} = m_{k,i_k} \; \forall i_k \in \{1, \ldots, \ell_k\} \tag{2.11}$$

$$\text{subject to } \gamma \geq 0 \tag{2.12}$$

(with the last inequality understood componentwise). For symmetric MMOT, we may assume that the ℓ_i, $a_{i,\alpha}$, and $m_{i,\alpha}$ are independent of i, and the discrete problem reads as follows: Given a set of ℓ distinct discretization points,

$$X = \{a_1, \ldots, a_\ell\} \subset \mathbb{R}^d, \tag{2.13}$$

and a marginal $\lambda^* \in \mathcal{P}(X)$ which we may view as a vector in $\mathbb{R}^\ell$ whose ith component is given by $\lambda^*(\{a_i\})$,

$$\text{Minimize } C[\gamma] = \sum_{i_1, \ldots, i_N = 1}^{\ell} \gamma_{i_1 \ldots i_N} c(a_{i_1}, \ldots, a_{i_N}) \text{ over } \gamma \in \mathcal{P}_{sym}(X^N) \tag{2.14}$$

$$\text{subject to } \sum_{i_2, \ldots, i_N = 1}^{\ell} \gamma_{i_1 i_2 \ldots i_N} = \lambda^*_{i_1} \forall\; i_1 \in \{1, \ldots, \ell\} \tag{2.15}$$

$$\text{subject to } \gamma \geq 0. \tag{2.16}$$

The associated dual problem is

$$\text{Maximize } \sum_{i=1}^{\ell} y_i \lambda^*_i \text{ over } y \in \mathbb{R}^\ell \tag{2.17}$$

$$\text{subject to } \tfrac{1}{N}\big(y_{i_1} + \cdots + y_{i_N}\big) \leq c(a_{i_1}, \ldots, a_{i_N}) \; \forall i_1, \ldots, i_N \in \{1, \ldots, \ell\}; \tag{2.18}$$

it discretizes the continuous dual problem [BDPGG12] to maximize $\int_{\bar{X}} y \, d\mu$ over measurable functions $y : \bar{X} \to \mathbb{R}$ satisfying $\frac{1}{N}(y(x_1) + \cdots + y(x_N)) \leq c(x_1, \ldots, x_N)$ $\forall x_1, \ldots, x_N \in \bar{X}$, whose solutions are called *Kantorovich potentials*. By linear programming duality, the value of (2.14)–(2.16) equals that of (2.17)–(2.18).

Application of a well-known stability result in optimal transport theory (see Theorems 1.50 and 1.51 of [San15] in the context of two-marginal problems; the extension to N marginals is straightforward) immediately yields the following convergence result as $\nu \to \infty$.

THEOREM 2.1 (justification of discretization). *For any compact sets $X_1, \ldots, X_N$ in $\mathbb{R}^d$, any continuous cost $c : X_1 \times \cdots \times X_N \to \mathbb{R}$, and any discretization* (2.8) *of the marginals which converges weak* to these, the optimal cost of the discretized problem* (2.9)–(2.11) *converges to that of the continuous problem* (2.1)–(2.3). *Moreover, any sequence of optimizers $\gamma^{(\nu)}$ of the discretized problem converges—after passing to a subsequence—weak* to a minimizer of the continuous problem.*

More sophisticated discretizations can be considered. For instance, one can represent integrable marginals μ_k by piecewise linear finite elements and use effective cost coefficients obtained by integrating the continuous cost function against the tensor

products of these elements, as in [CFM14] where the Coulomb problem was simulated for the dihydrogen molecule. For smooth marginals and costs, this is expected to improve the discretization error from $O(1/\nu)$ to $O(1/\nu^2)$. Moreover, to alleviate the computational cost the elements could be chosen adaptively so that each element carries approximately the same marginal mass [CFM14]. In this paper we do not investigate such refinements and instead confine ourselves to the basic qualitative justification of the discretization (2.9)–(2.11) given in Theorem 2.1.

The discrete problems (2.10)–(2.12) and (2.14)–(2.16) are high-dimensional linear programs. For some costs with very special interaction structure (such as the Wasserstein barycenter problem) a transformation to low-dimensional linear programs is possible [COO15] (see also [ABA20b]), making standard methods from linear programming applicable. For general costs, including (2.4), such schemes become unfeasible beyond a handful of marginals, due to the curse of dimension.

3. Extremal formulation of symmetric MMOT. The starting point of the algorithm presented here is the following equivalent formulation of symmetric MMOT introduced in [FV18], in which (candidate and optimal) Kantorovich plans are expressed as convex combinations of extreme points of $\mathcal{P}_{sym}(X^N)$. This eliminates any redundancy in the parametrization of plans and thus reduces the problem dimension, while at the same time it keeps the problem in the form used in two of the pioneering articles on column generation [DW60, DW61].

It is not difficult to show (see [FV18]) that when X is a finite state space, (2.13), the extreme points of $\mathcal{P}_{sym}(X^N)$ can be uniquely recovered from their marginals, which are given by the $\frac{1}{N}$-quantized probability measures on X,

$$\mathcal{P}_{\frac{1}{N}}(X) := \left\{ \lambda \in \mathcal{P}(X) \,\middle|\, \lambda(\{a_i\}) \in \left\{ 0, \frac{1}{N}, \frac{2}{N}, \dots, \frac{N}{N} \right\} \forall\, i \in \{1, \dots, \ell\} \right\}. \tag{3.1}$$

To recover the corresponding extreme point, write an element λ from the above set in the form $\sum_{k=1}^{N} \frac{1}{N} \delta_{a_{i_k}}$ for some (not necessarily distinct) points $a_{i_1}, \dots, a_{i_N} \in X$, and set

$$\gamma_\lambda = S_N \delta_{a_{i_1}} \otimes \cdots \otimes \delta_{a_{i_N}}. \tag{3.2}$$

Here S_N is the symmetrizer defined by $(S_N\gamma)(A_1 \times \cdots \times A_N) = \frac{1}{N!} \sum_\sigma \gamma(A_{\sigma(1)} \times \cdots \times A_{\sigma(N)})$, with the sum running over all permutations of $\{1, \dots, N\}$. Of course any element of the set $\mathcal{P}_{sym}(X^N)$ is a convex combination of extreme points, but here something better is true.

LEMMA 3.1 ([FV18]). *Any element $\gamma \in \mathcal{P}_{sym}(X^N)$ can be represented uniquely as a convex combination of the above extreme points, that is,*

$$\gamma = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda \gamma_\lambda, \quad \alpha_\lambda \geq 0 \ \forall \lambda \in \mathcal{P}_{\frac{1}{N}}(X), \quad \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda = 1. \tag{3.3}$$

Here the uniqueness is obvious from the fact that the γ_λ have mutually disjoint support.

Since γ_λ has marginal λ, the marginal condition becomes

$$\lambda^* = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda \lambda. \tag{3.4}$$

Thus the MMOT problem (2.9)–(2.12) can be written as the following optimization problem over the coefficient vectors α. Here and below we identify probability measures $\lambda \in \mathcal{P}_{1/N}(X)$ with vectors in $\mathbb{R}^\ell$ whose ith component is given by $\lambda(\{a_i\})$.

$$\text{Minimize } c^T\alpha = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} c_\lambda \alpha_\lambda \tag{3.5}$$

$$\text{subject to } A\alpha = \lambda^* \tag{3.6}$$

$$\alpha \geq 0, \tag{3.7}$$

with cost coefficients

$$c_\lambda = \sum_{i_1,\dots,i_N=1}^{\ell} (\gamma_\lambda)_{i_1,\dots,i_N} c(a_{i_1},\dots,a_{i_N}) \tag{3.8}$$

and A being the $\ell \times \binom{N+\ell-1}{N}$ matrix defined by

$$A\alpha = \sum_{\lambda \in \mathcal{P}_{\frac{1}{N}}(X)} \alpha_\lambda \lambda, \tag{3.9}$$

that is, the columns of A are given—say, in alphabetical order—by the vectors in $\mathcal{P}_{1/N}(X)$. For instance, for $\ell = 5$ and $N = 3$,

$$A = \frac{1}{3}\begin{pmatrix} 30000 & 22221111000000000000 & 1111110000 \\ 03000 & 10002000222111000000 & 1110001110 \\ 00300 & 01000200100200221100 & 1001101101 \\ 00030 & 00100020010020102021 & 0101011011 \\ 00003 & 00010002001002010212 & 0010110111 \end{pmatrix}.$$

Note that the normalization condition that the α_λ must sum to 1 is automatically enforced by the marginal constraints $A\alpha = \lambda^*$.

In what follows, we refer to the linear program (3.5)–(3.7) as the *master problem* (MP). This is the problem we seek to tackle in this paper. Note that the curse of dimension is still present, as the number of unknowns still grows combinatorially in N; just that by exploiting symmetry we have reduced it from ℓ^N in (2.9)–(2.11) to $\binom{N+\ell-1}{N}$. For instance, for 25 particles and 100 gridpoints for discretizing the marginal, this reduces the number of unknowns from 10^{50} to about 10^{26}—still out of reach of conventional methods.

4. Sparsity of optimizers; sparse manifolds. A fundamental feature of the above MP which our algorithm exploits is the extreme sparsity of optimizers. As is well known in polyhedral optimization, the number of nonzero entries of extremal optimizers is governed by the number of equality constraints. In the context of MMOT, this number is much smaller than the number of unknowns, and the ensuing exact sparse ansatz was first introduced and investigated by two of the authors in [FV18], where the following result was proved.

THEOREM 4.1 ([FV18]). *For any ℓ and N, any symmetric cost function $c : X^N \to \mathbb{R}$, and any marginal $\lambda_* \in \mathcal{P}(X)$, there exists an optimizer α_* of* (3.5)–(3.9) *belonging to the manifold*

$$\mathcal{M}_\ell := \{\alpha \in \mathbb{R}^{\binom{N+\ell-1}{N}} \,|\, \alpha_\lambda \geq 0 \,\forall \lambda,\ A\alpha = \lambda^*,\ \alpha \text{ has at most } \ell \text{ nonzero entries}\}.$$

Thus, in our case the number of required nonzero entries is just ℓ, independently of N. (Strictly speaking, $\mathcal{M}_\ell$ is not a manifold but only an algebraic variety.)

In [FV18] we proposed the name *quasi-Monge states* for the elements in this sparse manifold because of a close connection with the Monge ansatz in optimal transport. More precisely, one can show [FV18] that each plan γ corresponding to a coefficient vector in $\mathcal{M}_\ell$ can be written in the form

$$\gamma = S_N \sum_{\nu=1}^{\ell} \mu_\nu \delta_{T_1(a_\nu)} \otimes \cdots \otimes \delta_{T_N(a_\nu)}$$

or, in optimal transport notation (with $(\ )_\sharp$ denoting the push-forward of a measure)

$$\gamma = S_N (T_1, \dots, T_N)_\sharp \mu$$

for N maps $T_1, \dots, T_N : X \to X$ and ℓ coefficients $\mu_\nu \geq 0$ which sum to 1. Restricting μ to be equal to the prescribed marginal λ^* is the classical Monge ansatz from optimal transport theory. But the latter is too restrictive for the validity of Theorem 4.1 when $N \geq 3$, even in the case of the uniform marginal $\lambda^* = \frac{1}{\ell} \sum_{\nu=1}^{\ell} \delta_{a_\nu}$ (see [Fri19] for simple counterexamples and [Vög19] for a systematic numerical study).

From a computational perspective it will be useful to work on a slightly larger ansatz manifold,

(4.1)
$$\mathcal{M}_{\ell'} := \{\alpha \in \mathbb{R}^{\binom{N+\ell-1}{N}} \,\big|\, \alpha_\lambda \geq 0 \ \forall \lambda,\ A\alpha = \lambda^*,\ \alpha \text{ has at most } \ell' \text{ nonzero entries}\},$$

where

$$\ell < \ell' = \ell + O(\ell). \tag{4.2}$$

In practice we will use

$$\ell' = \beta \ell, \tag{4.3}$$

where β is a hyperparameter in the `GenCol` algorithm (chosen to be 5 in all our simulations). The intuition behind the enlargement of $\mathcal{M}_\ell$ to $\mathcal{M}_{\ell'}$ is that it keeps the sparsity at an extremely low level but makes the problem less nonlinear. (In the—in practice unfeasible—limit $\ell' = \binom{N+\ell-1}{N}$, one would obtain back the original linear program.)

5. Genetic column generation.

5.1. Column generation. In light of Theorem 4.1 it is—in principle—possible to solve the MP *exactly* via an algorithm that runs only on the data-sparse manifolds $\mathcal{M}_\ell$ or $\mathcal{M}_{\ell+O(\ell)}$ without ever touching the MP in its entirety. But what should one do in practice? CG is a pragmatic approach from discrete optimization, of primal-dual type, in which the primal state evolves precisely on such a sparse manifold. Its development originated in integer programming, but it has been especially useful in 0/1-integer programming, where the unknown is a vector in $\{0,1\}^N$ (corresponding to the domain of Kantorovich plans for $\ell = 2$ and N marginals) and where CG can be used in association with branch-and-bound techniques. Successful applications include traveling salesman problems, airline scheduling, and vehicle routing (see, e.g., [LD05]). Concurrent to our work, in [ABA20b] the authors use CG in connection with optimal transport to numerically support interesting theoretical results based on the ellipsoid method.

Consider any linear program of the form of our MP (3.5)–(3.7), and suppose we are in the general situation (satisfied in our case) that the matrix A has far fewer rows than columns and the number of columns is far too large to use standard linear programming solvers (such as Gurobi [GO19]). In CG one starts off by reducing the MP to a problem with far fewer variables by admitting only a small sized subset of the columns of A as a new constraint matrix. As only those admissible coefficient vectors α of the MP that are supported on the chosen columns are admissible for the new problem, the MP can be viewed as a "relaxation" of this new problem. Now the idea is to suitably *generate* additional candidate *columns* for the reduced problem and use a duality based criterion to accept or reject them in order to decrease its optimal value and thereby the gap to the optimal value of the MP.

Let us now explain the method in detail. The first step in CG consists of choosing a small sized subset $I \subset \bar{I} = \left\{1, 2, \dots, \binom{N+\ell-1}{N}\right\}$ of the columns of the constraint matrix A of the MP (3.5)–(3.7). For any such I, A_I and c_I denote the submatrix of A (respectively, the subvector of c) that contains exactly the corresponding columns (respectively, entries). Replacing the original constraint matrix A of the MP (3.5)–(3.7) by A_I and the cost vector c by c_I yields the problem

$$\text{Minimize } c_I^T \alpha \tag{5.1}$$

$$\text{subject to } A_I \alpha = \lambda^* \tag{5.2}$$

$$\alpha \geq 0. \tag{5.3}$$

Problem (5.1)–(5.3) will be referred to as the *restricted master problem* (RMP). As long as $\sharp I \leq \ell'$, candidate or optimal primal states of the RMP, extended by zero to $\bar{I}$, stay in the sparse manifold $\mathcal{M}_{\ell'}$ defined in (4.1).

Given an RMP (5.1)–(5.3) induced by a reduced constraint matrix A_I, one would like to add "better" columns to A_I, i.e., columns that improve the optimal value of the RMP. These "better" columns are best understood from a dual point of view. The *dual of the RMP* (DRMP) is given by

$$\text{Maximize } y^T \lambda^* \tag{5.4}$$

$$\text{subject to } A_I^T y \leq c_I. \tag{5.5}$$

Replacing A_I by A and c_I by c in this problem yields the *dual of the master problem* (DMP). The DRMP differs from the DMP by imposing far fewer constraints ($\sharp I$ instead of $\sharp \bar{I}$) on the dual variables $y \in \mathbb{R}^\ell$.

Theoretical discussions of CG now continue with the following—for high-dimensional problems unfeasible—step, in which the dual problem is used to find the "best" additional column: Given a dual optimal solution y^* of the DRMP, solve the so-called *pricing problem* (PP)

$$\text{Maximize } \lambda^T y^* - c_\lambda \tag{5.6}$$

$$\text{subject to } \lambda \in \mathcal{P}_{\frac{1}{N}}(X). \tag{5.7}$$

This problem looks for the constraint of the DMP that is violated the most by the given optimal solution y^* of the DRMP.

But the PP suffers from the fundamental problem we seek to circumvent, namely the curse of dimension. In fact, we will show in section 6 that even for pairwise costs, in which case the evaluation of c_λ is simple (see section 5.2), this problem is NP-complete.

In practice, for high-dimensional problems one needs to replace (5.6)–(5.7) by the following:

(5.8) Efficiently find a new column λ such that $\lambda^T y^* - c_\lambda > 0$.

Any such column can be added to the restricted constraint matrix A_I. The new column represents a constraint of the full dual (DMP) which the solution y^* to the current DRMP violates. Adding this column to the matrix A_I "cuts off" y^* from the optimization domain of the DRMP, yielding a new dual optimal solution $\tilde{y}^*$. Except in degenerate cases, this also leads to a new primal solution and a decrease in cost. For the convenience of the reader we include the well-known theoretical justification of the acceptance criterion in (5.8).

LEMMA 5.1 (justification of acceptance criterion). *If $\lambda^T y^* - c_\lambda \le 0$ for all columns λ of the full constraint matrix A, then the current dual solution y^* of the DRMP solves the full dual problem DMP, and the current primal solution α_I of the RMP, extended by zeros, solves the full primal problem MP.*

Proof. Denote the current primal solution extended by zeros by $\overline{\alpha_I}$. By assumption, $A^T y^* \le c$, that is, y^* is admissible for the full dual problem. Using, in order of appearance, the definition of $\overline{\alpha_I}$, duality for the RMP, admissibility of y^*, and duality for the full MP gives

$$c^T \overline{\alpha_I} = c_I^T \alpha_I = \lambda^* y^* \le \max_{y\,:\,A^T y \le c} \lambda^* y = \min_{\substack{\alpha\,:\,A\alpha = \lambda^* \\ \alpha \ge 0}} c^T \alpha.$$

Since $\overline{\alpha_I}$ is admissible for the full primal problem, the assertion follows. ∎

5.2. Fast cost evaluation for candidate columns. A possible bottleneck in CG in addition to the large number of columns can be the cost evaluation of a new column, required by the acceptance criterion in (5.8). In the case of MMOT with large N, a priori this requires evaluation of a high-dimensional sum; see (2.14). We now show that, due to the special structure of the extreme points γ_λ of $\mathcal{P}_{sym}(X^N)$ and the fact that the costs of interest are of pairwise form, this cost evaluation can, in fact, be done extremely fast, requiring only an *N-independent* number of arithmetic operations.

First, it is elementary that whenever $c : X^N \to \mathbb{R}$ is of pairwise and symmetric form,

$$c(x_1, \dots, x_N) = \sum_{1 \le i < j \le N} w(x_i, x_j) \text{ for some symmetric } w : X \times X \to \mathbb{R}, \tag{5.9}$$

then for any $\gamma \in \mathcal{P}_{sym}(X^N)$,

$$C[\gamma] = \binom{N}{2} \sum_{i,j=1}^{\ell} (M_2\gamma)_{ij} w(a_i, a_j),$$

where $M_2\gamma$ is the two-point marginal of γ, defined by $M_2\gamma(A) = \gamma(A \times X^{N-2})$ for all subsets A of X^2. For further discussion of this representation and its usefulness in electronic structure, see [FMP+13]. In the case of the extreme points γ_λ, the following explicit formula for the two-point marginal in terms of the one-point marginal λ was derived in [FV18]; it shows that on these points the highly noninvertible projection map from $M_2\gamma$ to $M_1\gamma$ can be inverted.

LEMMA 5.2 ([FV18]). *For any* $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$,

$$M_2\gamma_\lambda = \frac{N}{N-1}\lambda\otimes\lambda - \frac{1}{N-1}\sum_{i=1}^{\ell}\lambda(\{a_i\})\delta_{a_i}\otimes\delta_{a_i}.$$

Moreover, γ_λ *is the unique element of* $\mathcal{P}_{sym}(X^N)$ *with this two-point marginal.*

(In optimal transport notation, the second term on the right-hand side equals $-\frac{1}{N-1}(id,id)_{\#}\lambda$.) This immediately yields the following simple expression for the cost c_λ of a column $\lambda \in \mathcal{P}_{1/N}(X)$. Any such λ is now again regarded as a vector in $\mathbb{R}^\ell$.

COROLLARY 5.3. *If* c *has the pairwise form* (5.9)*, and* λ *is any element of* $\mathcal{P}_{1/N}(X)$*, then*

$$c_\lambda = \frac{N^2}{2}\lambda^T C\lambda - \frac{N}{2}\mathrm{diag}(C)^T\lambda, \tag{5.10}$$

with $C = (C_{ij})_{i,j=1}^{\ell} \in \mathbb{R}^{\ell\times\ell}$ *defined by*

$$C_{ij} = w(a_i,a_j). \tag{5.11}$$

This reduces cost evaluation to just matrix-vector multiplication with a precomputed matrix of N-independent size ℓ, and shows that the acceptance criterion in the PP (5.8) is extremely cheap computationally.

5.3. Genetic method for generating new columns. To tackle (5.8), let us recall the physical meaning of columns in MMOT in the key example of electronic structure. Transport plans γ with N marginals correspond to the joint probability density of N electron positions in a continuous d-dimensional domain $\Omega \subseteq \mathbb{R}^d$, or on the ℓ discretization points of an ℓ-point discretization $X = \{a_1,\dots,a_\ell\} \subset \Omega$ (see section 2). The columns $\lambda \in \mathcal{P}_{1/N}(X)$ describe all the possible "pure" N-particle configurations, obtained by dropping the N electrons on the ℓ discretization points (while allowing to multiply occupy sites). The MP (3.5)–(3.9) seeks to determine a stochastic superposition of these electron configurations that minimizes the interaction energy—prototypically, the mutual Coulomb repulsion—while fulfilling the marginal constraint. The latter describes the single-electron density, that is, the total occupancy of each site. Finding promising new columns corresponds to guessing good new N-particle configurations for the given density and interaction.

We take the view that guessing such—intricately correlated—configurations from the vast number of possiblities must be *learned*. The best available information given a current RMP matrix A_I and a solution α_I to the RMP is *which columns are successful*, i.e., which ones correspond to a nonzero component of the vector α_I. But this is already very valuable many-body information. For instance, in the case of Coulomb repulsion, successful many-particle configurations will already keep the electrons spatially apart from each other and avoid unfavorable clustering. This suggests a genetic approach which performs random small mutations of currently successful many-body configurations. More precisely, we propose the following.

Genetic search rule. *Given an instance of a reduced constraint matrix* A_I *and a corresponding RMP solution* α_I*,*

1. *allow only columns* λ *of* A_I *with* $(\alpha_I)_\lambda > 0$ *to be parents,*
2. *pick a parent column at random,*
3. *create a child by moving one randomly chosen particle in the parent configuration from its location* $a \in X$ *to a randomly chosen neighboring site* $a' \in X$*.*

The last step is crucially based on the physical/geometric meaning of columns as N-particle configurations in a region of d-dimensional Euclidean space.

Our rule for creation of children has an interesting *metric* meaning in column space which has nothing to do with viewing columns as vectors in $\mathbb{R}^\ell$ and using neighbors with respect to standard distances on $\mathbb{R}^\ell$. Instead, children are obtained from parents by *moving a minimum amount of mass by a minimum nonzero Euclidean distance*. To formalize this, let $d : X \times X \to \mathbb{R}$ be the Euclidean metric $d(x, y) = |x - y|$ on $X = \{a_1, \dots, a_\ell\} \subseteq \mathbb{R}^d$ inherited from the ambient $\mathbb{R}^d$. Columns are probability measures on X, and for any two columns λ, $\lambda' \in \mathcal{P}_{1/N}(X)$ let us introduce their *Wasserstein-1 distance* (also called the earth mover's distance) inherited from the ground metric d,

$$W_1(\lambda, \lambda') = \min \left\{ \int_{X \times X} d(x, y) d\gamma(x, y) \,|\, \gamma \in \mathcal{P}(X \times X),\ \gamma \text{ has marginals } \lambda \text{ and } \lambda' \right\}.$$

Then rule 3 can be reformulated as follows:

3′. *Pick a random nearest neighbor of the parent in the column space $\mathcal{P}_{1/N}(X)$ with respect to the Wasserstein-1 distance induced by the Euclidean metric on $X \subseteq \mathbb{R}^d$.*

We remark that, due to the mass quantization in $\mathcal{P}_{1/N}(X)$, any of the Wasserstein-p distances with $p \in [1, \infty)$ could be used here instead.

We emphasize that this abstract description of our genetic search rule does not mean that in practice there would be any need to compute Wasserstein distances. In our numerical examples the discretization points are chosen as the intersection of some region $\Omega \subset \mathbb{R}^d$ with a uniform lattice $h\mathbb{Z}^d$ of mesh size $h > 0$. One then just needs to pick a random occupied lattice point and update a random component by $\pm h$. In more sophisticated discretizations, such as the adaptive one in 3D in [CFM14], one simply needs to keep a nearest-neighbor list for each discretization point and make a random choice from this list.

A less stochastic, but slower, variant of step 3 would be to generate the best child (in terms of (5.6)) among all neighboring sites a' of the location a, or among all children (Wasserstein-1-neighbors) of the parent configuration.

5.4. The GenCol algorithm. Based on the results and considerations in the previous sections, we propose the following simple algorithm. By an active column we mean a column λ for which $(\alpha_I)_\lambda > 0$.

The inner while loop generates new columns according to the genetic rule described in section 5.3 until the acceptance criterion from the PP (5.8) is satisfied.

The outer loop is a standard CG iteration in which new columns are added to the current matrix A_I of the RMP and the primal and dual solutions are updated.

To prevent the size of A_I from growing too large, the "oldest" inactive columns are cleared whenever a maximum allowed size has been reached. The maximum size is defined with the help of the hyperparameter β; we do not allow it to exceed the minimum size for exactness of the method (namely ℓ; see Theorem 4.1) by more than a factor β. The meaning of "oldest" is oldest with respect to having been found; the empirical rationale here is that older columns were found with the help of a less accurate dual solution.

Algorithm 1 Genetic column generation (`GenCol`).

Input: N (the no. of marginals), ℓ (the no. of sites), β (hyperparameter, chosen to be 5 in our simulations), w (pair potential), Euclidean coordinates of sites in $\mathbb{R}^d$, marginal λ^*

Output: Solution to the MMOT problem (3.5)–(3.7), (5.10), (3.9)

initialize A_I, compute c_I, samples $= 0$, iter $= 0$, gain $= -1$ **while** *iter* $\leq$ *maxiter* **do**

  α_I = solution to RMP (5.1)–(5.3) y^* = solution to dual problem DRMP (5.4)–(5.5) **while** *gain* ≤ 0 *and samples* $\leq$ *maxsamples* **do**

    parent = random active column of A_I child = new column obtained from parent by randomly moving one particle to a neighboring site compute c_{child} (cost of child) using (5.10) gain = child$^T y_* - c_{\text{child}}$ samples = samples + 1

  end

  A_I = [A_I,child], c_I = [c_I,c_{child}] **if** *number of columns of* $A_I \geq \beta \cdot \ell$ **then**

    clear oldest ℓ inactive columns

  end

  iter = iter + 1

end

6. NP-completeness of the pricing problem. The formula for fast cost evaluation derived in Corollary 5.3 means that the PP (5.6)–(5.7) for MMOT is a linearly constrained integer-optimization problem with quadratic objective:

$$\text{Maximize } \lambda^T y^* - \frac{N}{2}\lambda^T C\lambda + \frac{N}{2}\text{diag}(C)^T\lambda \tag{6.1}$$

$$\text{subject to } \sum_{i=1}^{\ell}\lambda_i = N, \tag{6.2}$$

$$\lambda \in \mathbb{N}_0^\ell. \tag{6.3}$$

To derive this form, we have rescaled the objective function and the computational domain by a factor N.

As the theory of NP-completeness evolves around decision problems, we start by formulating a "decision version" of the PP (6.1)–(6.3).

The objective function of the PP consists of a quadratic term and a linear term. The linear term depends on the cost matrix for the quadratic objective. Nevertheless, we formulate the *Pricing Decision Problem* (PDP) by treating both terms as independent.

> *PDP: Given natural numbers $N, \ell \in \mathbb{N}$, a cost matrix $V \in \mathbb{R}^{\ell\times\ell}$, a cost vector $a \in \mathbb{R}^\ell$, and a threshold $K \in \mathbb{R}$, does there exist a vector $\lambda \in \{0, 1, \ldots, N\}^\ell$ such that*
>
> $$\sum_{i=1}^{\ell}\lambda_i = N$$
>
> *and*
>
> $$\lambda^T V\lambda + a^T\lambda \geq K?$$

Even if we restricted our attention to the choices of input parameters covered by

our PP (6.1)–(6.3), we still would be able to establish the NP-completeness of the PDP; see Remark 6.3 below.

The remainder of this section is devoted to proving NP-completeness of the PDP. This result strongly calls into question the possibility of a polynomial time algorithm for the PDP. In fact, it also calls into question the possibility of a polynomial time algorithm that solves the PP (6.1)–(6.3), by the following argument. Suppose that such an algorithm exists. Given an instance of the PDP, one is now able to compute the optimal value of the corresponding PP and simply compare it to the threshold of the given instance. Thus, the PDP also would be solvable in polynomial time.

To prove NP-completeness of the PDP we will use the following elementary lemma, whose proof is included for completeness.

LEMMA 6.1. *Given a natural number q, let $E^q \in \mathbb{R}^{q\times q}$ be the matrix whose diagonal entries are equal to zero, whereas all off-diagonal entries are equal to one. Then*

$$\tilde{\lambda} \text{ with } \tilde{\lambda}_i = 1 \quad \forall\, i \in \{1,\dots,q\} \tag{6.4}$$

is the unique maximizer of the problem

$$\begin{aligned} &\text{Maximize } \lambda^T E^q \lambda \\ &\text{subject to } \sum_{i=1}^{q} \lambda_i = q, \\ &\lambda \in \mathbb{R}^q, \end{aligned}$$

whose optimal value is therefore given by $q(q-1)$.

Proof. The matrix E^q has the two eigenvalues $q-1$ and -1. Corresponding eigenvectors are given by $v_1 = \tilde{\lambda}$ as well as $v_2 = e_2 - e_1, v_3 = e_3 - e_1, \dots, v_q = e_q - e_1$. Here e_i denotes the ith standard unit vector in $\mathbb{R}^d$. Geometrically, the eigenvector v_1 takes us onto the hyperplane our optimization problem is "living" on. Any further movement corresponds to an addition of a linear combination of the eigenvectors $v_2,\dots,v_q$ with the negative eigenvalue -1, and thereby a decrease of the objective value. In formulas, let us write any admissible trial state $\lambda \in \mathbb{R}^q$ as a linear combination of eigenvectors, $\lambda = \sum_{i=1}^{q} \alpha_i v_i$ with $\alpha_1,\dots,\alpha_q \in \mathbb{R}$. Multiplication with $(1,\dots,1)$ immediately shows that, by admissibility of λ, we have $\alpha_1 = 1$, and therefore

$$\lambda = v_1 + \sum_{i=2}^{q} \alpha_i v_i.$$

As v_1 is perpendicular to $v_2,\dots,v_q$,

$$\lambda^T E^q \lambda = v_1^T E^q v_1 - \left|\sum_{i=2}^{q} \alpha_i v_i\right|^2 \begin{cases} = q(q-1) & \text{if } \alpha_2 = \dots = \alpha_q = 0, \\ < q(q-1) & \text{else.} \end{cases}$$

This establishes the optimality of $v_1 = \lambda$. ∎

The main result of this section is the following.

THEOREM 6.2. *The PDP is NP-complete.*

Before we come to the proof, let us recall what it is that needs to be proven. As discussed, for example, in [KPP04, CLRS09], a decision problem Q is classified as

NP-complete if it is (i) contained in the class NP and (ii) "at least as hard" as any other problem in NP. This class consists of those decision problems for which "yes" instances can be verified in polynomial time. Regarding (ii), we will use the concept of polynomial reducibility. A decision problem S is said to reduce (or transform) to another decision problem T in polynomial time if there exists a polynomial time function f that maps any instance I of S to an instances $f(I)$ of T in such a manner that I is a "yes" instance of S if and only if $f(I)$ is a "yes" instance of T. Then as polynomial solvability of T implies polynomial solvability of S, T is considered "harder" as S. Consequently, in order to prove that Q is NP-complete one needs to show that any problem in NP can be polynomially reduced to Q. As a result of the transitivity of polynomial reducibility, one can show this by picking a known NP-complete problem and proving that it reduces to Q. In the following we will establish the NP-completeness of the PDP, i.e., prove Theorem 6.2, using this common approach. We will show that the following *Clique Decision Problem* (CDP) reduces to the PDP in polynomial time:

> *CDP: Let $G := (V, E)$ be an undirected graph, and let $K' \in \mathbb{N}$ be a natural number that fulfills $K' \leq |V|$. Does there exist a clique of size at least K'?*

Recall that given an undirected graph $G := (V, E)$, a clique C corresponds to a subset of the vertices of G, i.e., $C \subseteq V$, such that every distinct pair of vertices $c_1, c_2 \in C$ is connected by an edge, i.e., $\{c_1, c_2\} \in E$. We refer the interested reader to [CLRS09] for a more detailed discussion of the CDP including a proof of its NP-completeness.

Proof. First, it is easy to see that the PDP is contained in the class NP. Assume that we are given a capacity N, a size parameter ℓ, a threshold K, a cost matrix V, and a cost vector a that yield the answer "yes" if used as input arguments of the PDP. Now let λ be one of those vectors about whose existence the PDP asks. Then λ corresponds to a certificate, whose size is polynomial in the size of the input, for which the capacity as well as the threshold constraint can be checked in polynomial time. Thus, overall the "yes" instance can be verified in polynomial time.

Next we prove that the CDP can be reduced to the PDP in polynomial time. A given instance $I = (G, K')$ of the CDP is hereby mapped to an instance $f(I) = (N, \ell, K, a, V)$ of the PDP where N is set to K', ℓ corresponds to the cardinality $|V|$ of the vertex set of G, the threshold K is given by $K'(K' - 1)$, the cost vector a equals the zero vector, and finally, the cost matrix V is set to be the adjacency matrix $A_G = (a_{ij})_{i,j=1}^{\ell}$ of the graph G which fulfills

$$a_{ij} = \begin{cases} 1 & \text{if } \{i, j\} \in E, \\ 0 & \text{else.} \end{cases}$$

Thereby the adjacency matrix is, as usual, based on a given order of the vertices. It is easy to see that for a given instance I, the matching instance $f(I)$ can be computed in polynomial time. What remains to be shown is that for the described mapping, I is a "yes" instance regarding the CDP if and only if $f(I)$ is a "yes" instance with respect to the PDP.

Let $I = (G, K')$ be such a "yes" instance regarding the CDP. Then $G = (V, E)$ contains a clique $C \subseteq V$ of size K'. Assume further that λ is the vector indicating

which vertices are contained in C, i.e., $\lambda \in \{0,1\}^\ell$ with

$$\lambda_i = \begin{cases} 1 & \text{if } i \in C, \\ 0 & \text{else.} \end{cases}$$

Then it is elementary to check that the entries of λ sum to K' and that $\lambda^T A_G \lambda = K'(K'-1)$. Consequently, the given λ triggers a "yes" answer of the PDP with respect to the instance $f(I)$.

Now assume $f(I)$ is a "yes" instance with respect to the PDP. Then there exists a vector $\lambda \in \{0,1,\dots,K'\}^\ell$ such that its entries sum to K' and it satisfies $\lambda^T A_G \lambda \geq K'(K'-1)$. We will show that this vector λ only consists of zero- and one-entries, where the latter indicate a set of vertices that forms a clique in the graph G of size K'. This immediately determines that I is a "yes" instance regarding the CDP.

In the following, $I \subseteq \{1,2,\dots,\ell\}$ will denote a set of indices that satisfies

$$\lambda_i > 0 \text{ implies } i \in I$$

as well as

$$|I| = K'.$$

One can generate I by filling up the set of "support-indices" of λ with an arbitrary choice of the remaining indices in $\{1,2,\dots,\ell\}$. This is possible because by assumption $K' \leq \ell$. The K'-dimensional vector consisting only of those entries of λ that correspond to indices in I will, in the following, be denoted by λ_I. Accordingly, $A_G^I = (a_{ij})_{i,j \in I}$ denotes the matrix that is built from only those rows and columns of A_G with indices in I.

By the assumption on λ and the construction of I we have

$$K'(K'-1) \leq \lambda^T A_G \lambda = \lambda_I^T A_G^I \lambda_I. \tag{6.5}$$

Since replacing A_G^I with $E^{K'}$ as defined in Lemma 6.1 does not decrease the value of the quadratic function for nonnegative input arguments, we have

$$\lambda_I^T A_G^I \lambda_I \leq \lambda_I^T E^{K'} \lambda_I. \tag{6.6}$$

Combining (6.5) and (6.6) and applying Lemma 6.1 for $q = K'$ yields

$$K'(K'-1) \leq \lambda^T A_G \lambda = \lambda_I^T A_G^I \lambda_I \leq \lambda_I^T E^{K'} \lambda_I \leq K'(K'-1). \tag{6.7}$$

Every inequality in (6.7) is now actually an equality; consequently λ_I equals the vector $\tilde{\lambda}$ from (6.4) with $q = K'$, and therefore we also have $A_G^I = E^{K'}$. Since A_G^I corresponds to the adjacency matrix of the subgraph of G induced by I, G indeed contains a clique of size K', namely I. This concludes the proof of Theorem 6.2. ∎

Remark 6.3. We claim that the PDP is still NP-complete if the cost matrix V and the cost vector a are restricted to be of the special form in (6.1)–(6.3). In order for the PDP to remain NP-complete, there need to exist reduced problems—RMP, DRMP—that underlie the choices of A_G and the zero vector as cost matrix V and cost vector a. Define I in such a manner that A_I corresponds to the $\ell \times \ell$ identity matrix, and let $C = -\frac{2}{N} A_G$. Then the objective of the PP (6.1)–(6.3) takes on the required form. Consequently Theorem 6.2 remains valid if one restricts the choices of cost matrix and cost vector to the ones arising in (6.1).

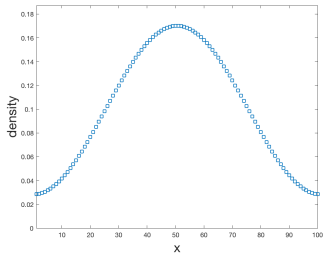

FIG. 1. *Prescribed single-particle density.*

7. Numerical results. All our tests were performed for our key motivating application, the Coulomb problem (2.1)–(2.4). For simplicity we used the regularized Coulomb interaction $w(x,y) = 1/\sqrt{\varepsilon^2 + |x-y|^2}$ with $\varepsilon = 0.1$. The tests were run with MATLAB on a MacBook Pro with a 2.4 GHz Intel Core i7 processor.

7.1. Ten electrons in one dimension with inhomogeneous density. As a first test we ran the `GenCol` algorithm on problem (2.1)–(2.4) with 10 electrons in a 1D interval discretized by $\ell = 100$ uniformly spaced gridpoints for the marginal density shown in Figure 1. We normalized the spacing to 1 and took the density as a function of the gridpoints $a_i = i \in \{1, \dots, \ell\} \subset \mathbb{R}$ to be $\lambda^*(\{a_i\}) = c_0 \cdot (0.2 + \sin^2(\frac{i}{\ell+1}))$, where c_0 is a normalization constant so that $\sum_i \lambda^*(\{a_i\}) = 1$.

We initialized the matrix A_I with the ℓ columns of the $\ell \times \ell$ identity matrix (to ensure that the optimization in the RMP (5.1)–(5.3) is feasible) as well as $(\beta - 1) \times \ell$ random columns, each of them obtained by dropping N particles randomly with respect to the uniform measure onto the grid. The results are given in Figure 2. After less than 7000 iterations the algorithm found what we believe to be the exact solution (within machine precision). Due to the problem size of 4.26×10^{13} possible columns, rigorous certification of the solution is out of the question, but we tested it both by a long (and, as it turned out, futile) biased search for better columns and by re-running the simulation many times, always ending up with the same state. Also, this behavior did not depend on any special choice of the hyperparameter β (the increase factor of the size of the column space beyond what is needed theoretically to capture optimizers). Increasing our choice $\beta = 5$ to 7 just led to a reduction of the number of iterations by a few percent, whereas decreasing it to 3 led to an increase by about 30 percent.

The multimarginal Kantorovich plan (or N-point density), visualized via its two-point marginal (or pair density), is seen to concentrate on the graphs of $N - 1 = 9$ maps, thereby accurately reproducing the known behavior of the continuous problem as predicted by Seidl [Sei99] and rigorously proved in [CDPDM15]. Overall only about 33000 columns out of the 4.26×10^{13} possible columns were sampled in order to find the ground state solution. The cost decreased steadily at an exponential rate (see Figure 3).

From an unsupervised learning perspective, the Kantorovich potential plays the same role in the `GenCol` algorithm for MMOT as it does in the Wasserstein-GAN algorithm [ACB17] for learning unknown distributions from data, namely that of an "adversary." In the initial stages the adversary is not much help (it looks to be close to a random potential) and the primal state has difficulty learning anything other

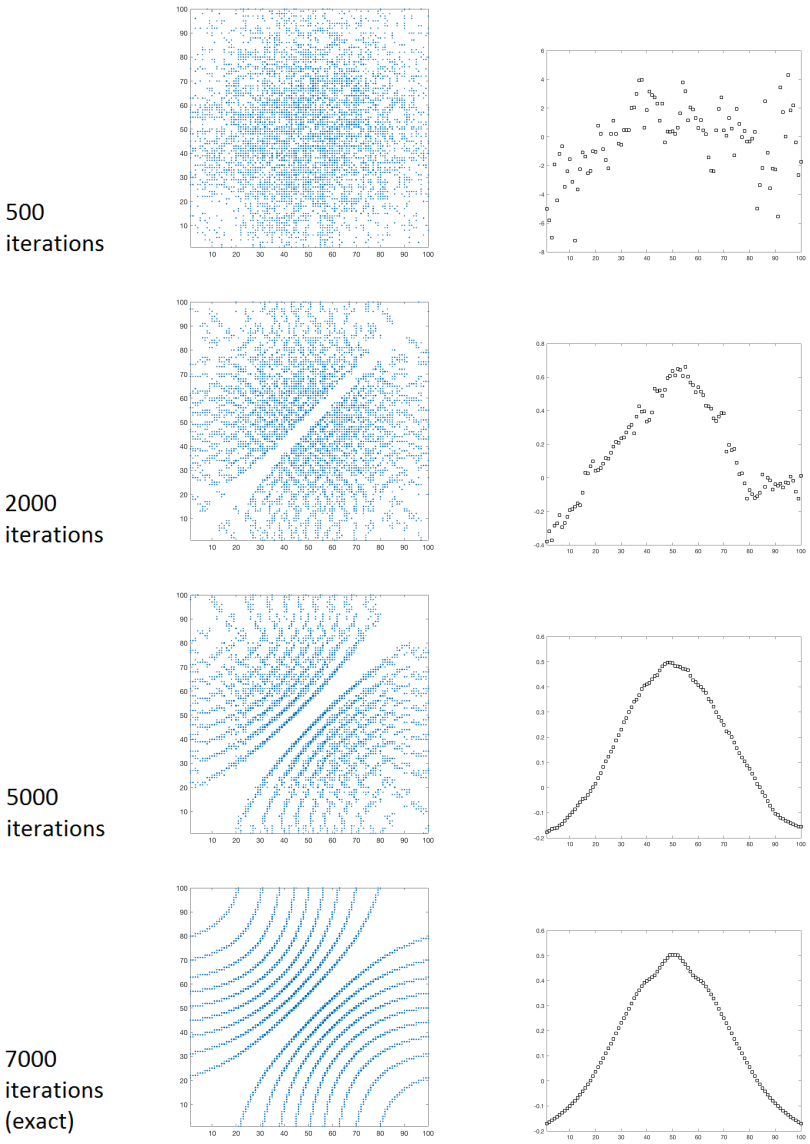

FIG. 2. *Solution to MMOT with Coulomb cost for* 10 *electrons in one dimension with the* `GenCol` *algorithm. The prescribed one-point marginal (single-electron density) is depicted in Figure* 1*, and was discretized by* 100 *gridpoints, resulting in* 4.26×10^{13} *unknowns (or "columns") in the full linear program. Left column: Evolution of the multimarginal plan, visualized via its two-point marginal (pair density). All gridpoints with nonzero values are shown, with larger markers indicating higher values. Right column: Evolution of the dual solution (Kantorovich potential). The final plan—believed to be the exact ground state within machine precision—was found using* 6789 *iterations (accepted columns) and* 33,283 *samples (genetically generated columns).*

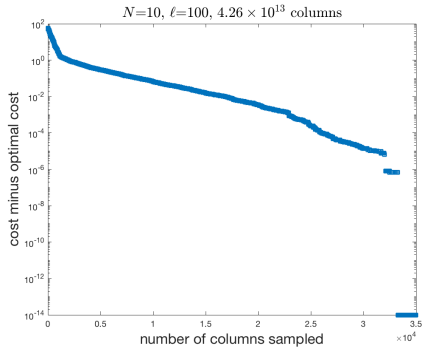

FIG. 3. *Evolution of the cost of the RMP solution for the* 10*-electron simulation from Figure* 2.

than the—physically obvious—fact that two electrons being extremely close is costly. As the number of iterations increases, the primal and dual state steadily acquire finer and finer characteristics until reaching optimality. We attribute the success of the `GenCol` algorithm in overcoming the vastness of the space of possible Kantorovich plans to the ability of the primal and dual state to "learn from each other."

7.2. Large N-electron systems in one dimension; cost scaling. We now empirically investigate the important issue of how the computational cost of the `GenCol` algorithm scales with system size. As a suite of test systems we choose MMOT with Coulomb cost in one dimension and homogeneous marginal λ^*, with an increasing number N of electrons and an increasing number ℓ of gridpoints. In fact, it is physically natural to *increase both parameters simultaneously* and consider a sequence of systems with

$$\text{increasing } N, \quad \text{increasing } \ell, \quad \frac{N}{\ell} \equiv const. \tag{7.1}$$

In the limit $N \to \infty$, $\ell \to \infty$, $\frac{N}{\ell} \equiv const$ (the so-called thermodynamic limit) the system approaches the 1D homogeneous electron gas. At fixed mesh size (normalized to 1 in our simulations), the condition $\frac{N}{\ell} \equiv const$ means physically that we increase the available volume proportionally to the number of particles, thereby allowing typical interparticle distances to stay unaltered, as happens in large molecules and solids in nature.

The above family of systems has the advantage that for integer values of $\frac{N}{\ell}$ the exact solution to (2.14)–(2.16) is known even after discretization (or, more precisely, it can be deduced via the same methods with which the exact solution for the continuous theory was derived in [CDPDM15]). It consists of the symmetrized Monge state

$$\gamma_{i_1,\dots,i_N} = S_N \sum_{i_1=1}^{\ell} \lambda^*_{i_1} \prod_{k=2}^{N} \delta_{i_k, i_1+(k-1)\frac{\ell}{N}}, \tag{7.2}$$

which represents a superposition of uniformly spaced N-particle configurations. Here $\delta_{i,j}$ denotes the Kronecker delta function.

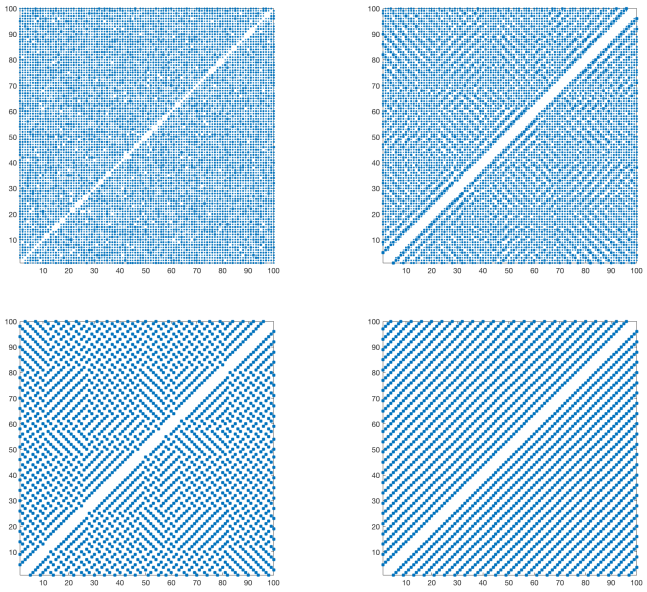

FIG. 4. *Solution to MMOT with Coulomb cost for* 25 *electrons in one dimension with the* `GenCol` *algorithm, with prescribed homogeneous marginal (single-electron density). The marginal was discretized by* 100 *gridpoints, resulting in* 1.0404×10^{26} *unknowns (or "columns") in the full linear program. Multimarginal plans are visualized via their two-point marginal (pair density) as described in Figure* 2. *Top row and bottom left:* 1000, 8000, *and* 9000 *iterations. Bottom right: exact solution, reached after* 9322 *iterations (accepted columns) and* 38,860 *samples (genetically generated columns).*

We ran the `GenCol` algorithm on the sequence of systems

$$\begin{cases} N = 5 \\ \ell = 20 \end{cases}, \quad \begin{cases} N = 10 \\ \ell = 40 \end{cases}, \quad \begin{cases} N = 15 \\ \ell = 60 \end{cases}, \quad \begin{cases} N = 20 \\ \ell = 80 \end{cases}, \quad \begin{cases} N = 25 \\ \ell = 100 \end{cases}, \quad \begin{cases} N = 30 \\ \ell = 120 \end{cases}, \tag{7.3}$$

with five different runs for each system. We initialized the matrix A_I with the ℓ columns of the identity matrix (for feasibility), augmented by $N \cdot \ell$ random columns. In every single case `GenCol` found the exact solution. See Figure 4 for the evolution of the Kantorovich plan for $N = 25$, $\ell = 100$. The number of iterations and genetic samples needed to find the exact solution are given in Table 1. The average runtime in MATLAB on a MacBook Pro was 48 seconds for 15 electrons on 60 gridpoints (i.e., 1.8×10^{15} unknowns) and 13 minutes for 30 electrons on 120 gridpoints (i.e., 2.6×10^{31} unknowns). All but a negligible amount of this CPU time decomposed as the number of iterations (= calls of the linear programming solver used, in our case the MATLAB function `linprog`) times the CPU time for solving an instance of the linear program.

Since each iteration only involves solving a linear program for at most $\beta \cdot \ell$ unknowns and ℓ constraints (where $\beta = 5$ in our case), and we limited the number of iterations in the linear programming solver used to $O(\ell^2)$, from a theoretical point of view the key limiting factor is the number of genetic samples needed. Figure 5 shows a log-log-plot of the average number of genetic samples needed for each system. While

TABLE 1

Number of iterations (accepted columns) and samples (genetically generated columns) needed by `GenCol` *to find the exact ground state for MMOT with Coulomb cost and homogeneous marginal in one dimension. The number ℓ of gridpoints was increased proportionally to the number N of electrons in line with* (7.3) *(see first column), and five different runs were performed for each system.*

System	Total number of columns	Accepted columns	Sampled columns	Sampled columns (average)
$N = 5$, $\ell = 20$	4.2504×10^4	101, 121, 116, 146, 117	467, 592, 485, 559, 455	511.6
$N = 10$, $\ell = 40$	8.2178×10^9	913, 757, 735, 915, 664	3853, 2768, 2872, 3912, 2762	3233.4
$N = 15$, $\ell = 60$	1.8240×10^{15}	2575, 2401, 2342, 2540, 2658	9901, 9301, 9141, 9967, 11812	10024.4
$N = 20$, $\ell = 80$	4.2879×10^{20}	5649, 5633, 4839, 5557, 5256	24856, 24227, 20272, 22872, 22265	22898.4
$N = 25$, $\ell = 100$	1.0404×10^{26}	10611, 9436, 8334, 10186, 9322	48188, 40939, 31371, 40724, 38860	40017.4
$N = 30$, $\ell = 120$	2.5759×10^{31}	15539, 14262, 15484, 15190, 14714	65566, 58283, 75729, 63004, 62759	65068.2

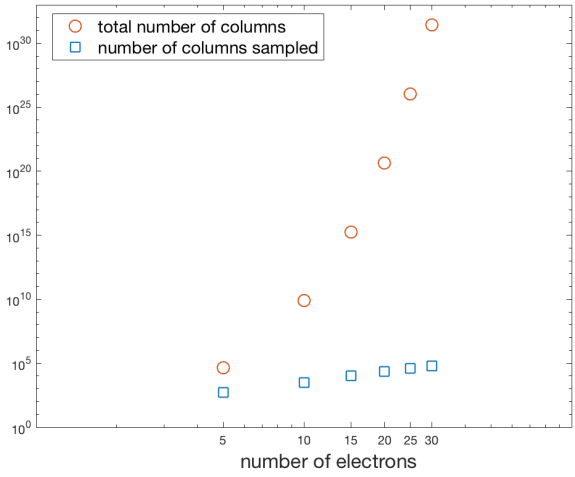

FIG. 5. *Double logarithmic plot of the total number of genetic samples needed by* `GenCol` *to find the exact ground state versus the number of electrons, for the systems described in Table* 1. *The plot gives the average number of samples over five runs (last column of the table), and the number of gridpoints was increased proportionally to the number of electrons (see the first column of the table).*

the system size (i.e., the number of unknowns) grows exponentially, the number of genetic samples needed for finding the exact solution appears to lie on a straight line, suggesting polynomial growth only. This is particularly remarkable in light of our result in section 6 that the PP—which our genetic sampling method addresses—is NP-complete.

8. Discussion and conclusions. The main advantage of our algorithm making it much faster than previous methods appears to be its simplicity: one just needs to solve low-dimensional linear programs. Moreover, after discretization no further approximations are made and the marginal constraints are automatically maintained, making the solution very accurate. Finally, we note that the method also gives the Kantorovich potential, which is needed in applications to electronic structure.

REFERENCES

[ABA20a] J. Altschuler and E. Boix-Adserà, *Hardness Results for Multimarginal Optimal Transport Problems*, preprint, https://arxiv.org/abs/2012.05398, 2020.

[ABA20b] J. Altschuler and E. Boix-Adserà, *Polynomial-Time Algorithms for Multimarginal Optimal Transport Problems with Decomposable Structure*, preprint, https://arxiv.org/abs/2008.03006v1, 2020.

[AC11] M. Agueh and G. Carlier, *Barycenters in the Wasserstein space*, SIAM J. Math. Anal., 43 (2011), pp. 904–924, https://doi.org/10.1137/100805741.

[ACB17] M. Arjovsky, S. Chintala, and L. Bottou, *Wasserstein GAN*, preprint, https://arxiv.org/abs/1701.07875, 2017.

[ACE21] A. Alfonsi, R. Coyaud, and V. Ehrlacher, *Constrained Overdamped Langevin Dynamics for Symmetric Multimarginal Optimal Transportation*, preprint, https://arxiv.org/abs/2102.03091, 2021.

[ACEL21] A. Alfonsi, R. Coyaud, V. Ehrlacher, and D. Lombardi, *Approximation of optimal transport problems with marginal moments constraints*, Math. Comp., 90 (2021), pp. 689–737.

[BCN16] J-D. Benamou, G. Carlier, and L. Nenna, *A numerical method to solve multimarginal optimal transport problems with Coulomb cost*, in Splitting Methods in Communication, Imaging, Science, and Engineering, Springer, 2016, pp. 577–601.

[BCN19] J.-D. Benamou, G. Carlier, and L. Nenna, *Generalized incompressible flows, multi-marginal transport and Sinkhorn algorithm*, Numer. Math., 142 (2019), pp. 33–54.

[BDM12] R. Burkard, M. Dell'Amico, and S. Martello, *Assignment Problems, revised reprint*, SIAM, 2012, https://doi.org/10.1137/1.9781611972238.

[BDPGG12] G. Buttazzo, L. De Pascale, and P. Gori-Giorgi, *Optimal-transport formulation of electronic density-functional theory*, Phys. Rev. A, 85 (2012), 062502.

[Bec14] A. Becke, *Perspective: Fifty years of density-functional theory in chemical physics*, J. Chem. Phys., 140 (2014), 18A301.

[Bre89] Y. Brenier, *The least action principle and the related concept of generalized flows for incompressible perfect fluids*, J. Amer. Math. Soc., 2 (1989), pp. 225–255.

[Bre91] Y. Brenier, *Polar factorization and monotone rearrangement of vector-valued functions*, Comm. Pure Appl. Math., 44 (1991), pp. 375–417.

[BS20] M. Bonafini and B. Schmitzer, *Domain Decomposition for Entropy Regularized Optimal Transport*, preprint, https://arxiv.org/abs/2001.10986, 2020.

[CDPDM15] M. Colombo, L. De Pascale, and S. Di Marino, *Multimarginal optimal transport maps for one-dimensional repulsive costs*, Canad. J. Math., 67 (2015), pp. 350–368.

[CFK13] C. Cotar, G. Friesecke, and C. Klüppelberg, *Density functional theory and optimal transportation with Coulomb cost*, Comm. Pure Appl. Math., 66 (2013), pp. 548–599.

[CFK18] C. Cotar, G. Friesecke, and C. Klüppelberg, *Smoothing of transport plans with fixed marginals and rigorous semiclassical limit of the Hohenberg–Kohn functional*, Arch. Ration. Mech. Anal., 228 (2018), pp. 891–922.

[CFM14] H. Chen, G. Friesecke, and C. Mendl, *Numerical methods for a Kohn–Sham density functional model based on optimal transport*, J. Chem. Theory Comput., 10 (2014), pp. 4360–4368.

[CLRS09] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein, *Introduction to Algorithms*, 3rd ed., The MIT Press, 2009.

[CMN10] P.-A. Chiappori, R. J. McCann, and L. P. Nesheim, *Hedonic price equilibria, stable matching, and optimal transport: Equivalence, topology, and uniqueness*, Econom. Theory, 42 (2010), pp. 317–354.

[COO15] G. Carlier, A. Oberman, and E. Oudet, *Numerical methods for matching for teams and Wasserstein barycenters*, ESAIM Math. Model. Numer. Anal., 49 (2015), pp. 1621–1642.

[Cut13] M. Cuturi, *Sinkhorn distances: Lightspeed computation of optimal transport*, in Advances in Neural Information Processing Systems 26, Curran Associates, Inc., 2013, pp. 2292–2300.

[DW60] G. B. Dantzig and P. Wolfe, *Decomposition principle for linear programs*, Oper. Res., 8 (1960), pp. 101–111.

[DW61] G. B. Dantzig and P. Wolfe, *The decomposition algorithm for linear programs*, Econometrica, 29 (1961), pp. 767–778.

[FGGSDS16] E. Fabiano, P. Gori-Giorgi, M. Seidl, and F. Della Sala, *Interaction-strength interpolation method for main-group chemistry: Benchmarking, limitations, and perspectives*, J. Chem. Theory Comput., 12 (2016), pp. 4885–4896.

[FMP+13] G. Friesecke, C. B. Mendl, B. Pass, C. Cotar, and C. Klüppelberg, *N-density representability and the optimal transport limit of the Hohenberg-Kohn functional*, J. Chem. Phys., 139 (2013), 164109.

[Fri19] G. Friesecke, *A simple counterexample to the Monge ansatz in multimarginal optimal transport, convex geometry of the set of Kantorovich plans, and the Frenkel–Kontorova model*, SIAM J. Math. Anal., 51 (2019), pp. 4332–4355, https://doi.org/10.1137/18M1207326.

[FV18] G. Friesecke and D. Vögler, *Breaking the curse of dimension in multi-marginal Kantorovich optimal transport on finite state spaces*, SIAM J. Math. Anal., 50 (2018), pp. 3996–4019, https://doi.org/10.1137/17M1150025.

[GO19] LLC Gurobi Optimization, *Gurobi Optimizer Reference Manual*, 2019.

[GS98] W. Gangbo and A. Święch, *Optimal maps for the multidimensional Monge-Kantorovich problem*, Comm. Pure Appl. Math., 51 (1998), pp. 23–45.

[KPP04] H. Kellerer, U. Pferschy, and D. Pisinger, *Knapsack Problems*, Springer-Verlag, 2004.

[KY19] Y. Khoo and L. Ying, *Convex relaxation approaches for strictly correlated density functional theory*, SIAM J. Sci. Comput., 41 (2019), pp. B773–B795, https://doi.org/10.1137/18M1207478.

[LD05] M. E. Lübbecke and J. Desrosiers, *Selected topics in column generation*, Oper. Res., 53 (2005), pp. 1007–1023.

[Nen17] L. Nenna, *Numerical Methods for Multi-marginal Optimal Transportation*, Ph.D. thesis, https://tel.archives-ouvertes.fr/tel-01471589/, 2017.

[PC19] G. Peyré and M. Cuturi, *Computational Optimal Transport*, preprint, https://arxiv.org/abs/1803.00567, 2019.

[Pie68] W. P. Pierskalla, *The multidimensional assignment problem*, Oper. Res., 16 (1968), pp. 422–431.

[Poo94] A. B. Poore, *Multidimensional assignment formulation of data association problems arising from multitarget and multisensor tracking*, Comput. Optim. Appl., 3 (1994), pp. 27–57.

[RR98] S. T. Rachev and L. R. Rüschendorf, *Mass Transportation Problems, Volume* I: *Theory; Volume* II: *Applications*, Springer, 1998.

[San15] F. Santambrogio, *Optimal Transport for Applied Mathematicians: Calculus of Variations, PDEs, and Modeling*, Birkhäuser Basel, 2015.

[Sch16] B. Schmitzer, *A sparse multiscale algorithm for dense optimal transport*, J. Math. Imaging Vis., 56 (2016), pp. 238–259.

[Sch19] B. Schmitzer, *Stabilized sparse scaling algorithms for entropy regularized transport problems*, SIAM J. Sci. Comput., 41 (2019), pp. A1443–A1481, https://doi.org/10.1137/16M1106018.

[Sei99] M. Seidl, *Strong-interaction limit of density-functional theory*, Phys. Rev. A, 60 (1999), pp. 4387–4395.

[Spi00] F. C. R. Spieksma, *Multi index assignment problems: Complexity, approximation, applications*, in Nonlinear Assignment Problems, Comb. Optim. 7, Springer, 2000, pp. 1–12.

[Vil09] C. Villani, *Optimal Transport: Old and New*, Springer-Verlag, 2009.

[VMR+21] A. Vacher, B. Muzellec, A. Rudi, F. Bach, and F.-X. Vialard, *A Dimension-free Computational Upper-Bound for Smooth Optimal Transport Estimation*, preprint, https://arxiv.org/abs/2101.05380, 2021.

[Vög19] D. Vögler, *Kantorovich vs. Monge: A Numerical Classification of Extremal Multi-marginal Mass Transports on Finite State Spaces*, preprint, https://arxiv.org/abs/1901.04568, 2019.

B Further publication

B.1 Convex geometry of finite exchangeable laws and de Finetti style representation with universal correlated corrections

Authors: Guillaume Carlier, Gero Friesecke, and Daniela Vögler

Summary. We shine a light on the convex geometry of the set $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ of N-representable k-point probability measures regarding a given state space X. Here, N and k denote natural numbers that fulfil $N \geq k$. The set $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ arises as the image of the set $\mathcal{P}_{\text{sym}}\left(X^N\right)$ of symmetric N-point probability measures under the linear k-point marginal map M_k, i.e.,

$$\mathcal{P}_{N\text{-rep}}\left(X^k\right) := \left\{ M_k\gamma \quad : \quad \gamma \in \mathcal{P}_{\text{sym}}\left(X^N\right) \right\}.$$

In symmetric multi-marginal optimal transport, a k-body interaction structure of cost functions allows us to reduce the optimization domain from a subset of $\mathcal{P}_{\text{sym}}\left(X^N\right)$ to a subset of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$. Here, we speak of a reduction with regards to the dimension of the product space: from X^N to X^k. Even though the elements of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ 'live' on a lower-dimensional product space, they were initially only understood implicitly as marginals of higher-dimensional measures.

In the present article, we explicitly provide universal polynomials $F_{N,k}$ that allow us to express any extreme point μ_k of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ in terms of its one-point marginal:

$$\mu_k = F_{N,k}\left(\lambda\right) \quad \text{for } \lambda = M_1\mu_k.$$

Non-extremal elements of $\mathcal{P}_{N\text{-rep}}\left(X^k\right)$ can be written as 'combinations' of these, now explicitly known, extreme points $\left(F_{N,k}\left(\lambda\right)\right)_\lambda$. Here, $F_{N,k}$ is a polynomial of degree k with leading term $N^{k-1}/\prod_{j=1}^{k-1}(N-j)\lambda^{\otimes k}$ (prefactor times k-fold tensor product of λ with itself) and a series of $(k-1)$ corrections, with correction j being of order $1/N^j$. The one-point marginal map M_1 maps the set of extremal N-representable k-point probability measures bijectively onto the set of $1/N$-quantized probability measures $\mathcal{P}_{\frac{1}{N}}\left(X\right)$, $F_{N,k}$ actually acts as a polynomial inverse in this relation. We first establish this full description of extreme points in the case of a finite state space X, before moving on to the case of a general state

space X, that is a Polish space. For $k = 2$ and 3, we recover previous results for finite state spaces from [44] and [61], respectively.

Own contribution. I, Daniela Vögler, do NOT claim to be principal author of the present article. Nevertheless, I was significantly involved in each step on the way from forming the ideas through conducting the research to putting pen to paper. In my opinion, each of the three authors contributed equally. Without each one of us bringing their specific set of skills to the table, providing such a beautiful result - that is relevant to multi-marginal optimal transport as well as finitely extendible random vectors - would not have been possible.

The publisher of *Probability Theory and Related Fields*, Springer Nature, permits the inclusion of this publication:

> *"Author reuse*
> *(...) Authors have the right to reuse their article's Version of Record, in whole or in part, in their own thesis. Additionally, they may reproduce and make available their thesis, including Springer Nature content, as required by their awarding academic institution."*
>
> [Springer Nature. Rights and Permissions. https://www.springer.com/gp/rights-permissions/obtaining-permissions/882, 2022. [Online; accessed 16.05.2022].]

Probability Theory and Related Fields
https://doi.org/10.1007/s00440-022-01115-2

Convex geometry of finite exchangeable laws and de Finetti style representation with universal correlated corrections

Guillaume Carlier[1] · Gero Friesecke[2] · Daniela Vögler[2]

Received: 7 July 2021 / Revised: 23 January 2022 / Accepted: 26 January 2022

Abstract

We present a novel analogue for finite exchangeable sequences of the de Finetti, Hewitt and Savage theorem and investigate its implications for multi-marginal optimal transport (MMOT) and Bayesian statistics. If $(Z_1, \ldots, Z_N)$ is a finitely exchangeable sequence of N random variables taking values in some Polish space X, we show that the law μ_k of the first k components has a representation of the form

$$\mu_k = \int_{\mathcal{P}_{\frac{1}{N}}(X)} F_{N,k}(\lambda)\, \mathrm{d}\alpha(\lambda)$$

for some probability measure α on the set of $\frac{1}{N}$-quantized probability measures on X and certain universal polynomials $F_{N,k}$. The latter consist of a leading term $N^{k-1}/\prod_{j=1}^{k-1}(N-j)\lambda^{\otimes k}$ and a finite, exponentially decaying series of correlated corrections of order N^{-j} $(j = 1, \ldots, k)$. The $F_{N,k}(\lambda)$ are precisely the extremal such laws, expressed via an explicit polynomial formula in terms of their one-point marginals λ. Applications include novel approximations of MMOT via polynomial convexification and the identification of the remainder which is estimated in the celebrated error bound of Diaconis and Freedman (Ann Probab 8(4):745–764, 1980) between finite and infinite exchangeable laws.

Guillaume Carlier
carlier@ceremade.dauphine.fr

Gero Friesecke
gf@ma.tum.de

Daniela Vögler
voegler@ma.tum.de

[1] CEREMADE, Université Paris Dauphine, PSL and INRIA-Paris, MOKAPLAN, Paris, France

[2] Faculty of Mathematics, Technische Universität München, Munich, Germany

Published online: 16 February 2022

Keywords N-representability · Finite exchangeability · De Finetti · Hewitt–Savage theorem · Multi-marginal optimal transport · Bayesian statistics · Extremal measures · Choquet theory

Mathematics Subject Classification 60G09 · 52A05 · 49Q22

1 Introduction

Multi-marginal optimal transport (MMOT) has attracted a great deal of attention in recent years. The relevance of MMOT to tackle challenging problems arising from electronic density functional theory was established in [6, 8]. In this context, one has to find the joint density of N electrons with fixed one-point marginal so as to minimize a total repulsive Coulombian cost. Even though the problem is difficult for large N, it is symmetric (invariant under permutations of the electrons) and only depends on the two-point marginal of the joint law of the N electrons (2-body interaction). Whether symmetries and few-body interactions are helpful to analyze such MMOT problems is a natural question. An interesting result from [9] relying on the fact that the Coulomb potential has a positive Fourier transform and the de Finetti, Hewitt and Savage theorem is that when one lets N go to $+\infty$, the optimal plan is the independent (infinite product) measure. This is in striking contrast with the more standard two-marginal optimal transport where, for typical costs including the Coulomb cost, optimal plans are sparse and concentrate on low-dimensional subsets of the product space [5, 8, 18]. The present paper is motivated by MMOT for a possibly large but finite number of marginals N and symmetric k-body (with $k \leq N$) interaction cost. We present a novel explicit analogue of the de Finetti, Hewitt and Savage theorem and investigate its implications for such problems. We also briefly indicate implications for Bayesian statistics.

The main technical novelty in our work is the construction of an explicit polynomial inverse of the marginal map from extremal N-representable k-point probability measures (see below for terminology) to 1-point probability measures. This extends previous results for 2-point [17] and 3-point [29] measures on finite state spaces to arbitrary k and general Polish spaces.

Our ensuing finite version of de Finetti yields a complete, finite, exponentially decaying series of correlated corrections which need to be added to the independent measure in the case of finite N. This explicitly identifies the remainder estimated in the celebrated error bound of Diaconis and Freedman [12] between finite and infinite exchangeable laws.

In the remainder of this introduction we first recall the celebrated de Finetti–Hewitt–Savage theorem, then describe in more detail what changes in the finite exchangeable case.

De Finetti–Hewitt–Savage Recall that a sequence $(Z_i)_{i\in\mathbb{N}}$ of random variables taking values in a Polish space X is called exchangeable if the law of $(Z_1, Z_2, \ldots)$ equals that of $(Z_{\sigma(1)}, Z_{\sigma(2)}, \ldots)$ for each finite permutation σ of $\mathbb{N}$, that is each permutation which leaves all but finitely many elements unchanged. The de Finetti–Hewitt–Savage theorem says that any such sequence is a convex mixture of i.i.d. sequences. In other

words, the law of $(Z_1, Z_2, \ldots)$ is a convex combination of independent measures,

$$\mu = \int_{\mathcal{P}(X)} \lambda^{\otimes\infty} d\alpha(\lambda) \tag{1.1}$$

for some probability measure (or prior) α on the set $\mathcal{P}(X)$ of probability measures on X. In Bayesian language, this says that the general infinite exchangeable sequence (Z_i) is obtained by first picking some distribution λ on X at random from some prior, then taking (Z_i) to be i.i.d. with distribution λ. For comprehensive reviews of exchangeability we refer the reader to Aldous [2] and Kallenberg [26].

Finite exchangeability; finite extendibility A sequence $(Z_1, \ldots, Z_N)$ is called finitely exchangeable if its law equals that of $(Z_{\sigma(1)}, \ldots, Z_{\sigma(N)})$ for any permutation σ of $\{1, \ldots, N\}$. For $k \leq N$, a sequence $(Z_1, \ldots, Z_k)$ is called finitely extendible if its law equals that of the first k elements of some finitely exchangeable sequence $(\tilde{Z}_1, \ldots, \tilde{Z}_N)$.[1]

For finite exchangeable sequences $(Z_1, \ldots, Z_N)$ it is well-known that the analogous representation to (1.1) with $\lambda^{\otimes\infty}$ replaced by $\lambda^{\otimes N}$ does not hold, see Diaconis [13],Diaconis and Freedman [12], Jaynes [25]; the error is known to be of order $\frac{1}{N}$ in total variation [4, 12].

The main approach for describing finite exchangeable sequences which has been introduced in the probability literature is to write such a sequence as a superposition of i.i.d. sequences but drop the requirement that the superposition of the laws be convex, i.e. allows signed measures α in (1.1), see Dellacherie and Meyer [11], Jaynes [25], Kerns and Székely [28], Janson, Konstantopoulos and Yuan [24]. For applications, this approach has limited appeal, for two reasons. First, the signed measure representation is not unique. Second, the superposition does not yield a probability measure for an arbitrary signed α, but remaining within probability measures is mandatory for recovering an exchangeable law by sampling (see below) and for our application to MMOT.

A very interesting second picture of finite exchangeability which appears not to have received the attention it deserves can be found in Kerns and Székely [28] and Kallenberg [26]. Namely, finite exchangeable sequences are convex mixtures of "urn sequences", or equivalently, finitely exchangeable laws are convex superpositions of symmetrized Dirac measures, the latter being the laws of urn sequences (as described further below). See Kerns and Székely ([28], top of p. 600), where such a representation appears as an intermediate step in the proof of the signed-measure representation. The laws of urn sequences are known to be the extreme points of the convex set of finitely exchangeable laws, see [28] for an elementary proof for finite state spaces and Kallenberg ([26] Proposition 1.8) for a general proof using advanced probability methods. Although not given in [26], the Kerns–Székely representation could be deduced from the statement of Proposition 1.8 via disintegration of measures.

The present work builds upon this picture, which turns out to be very useful for the applications we have in mind. Thus we view finite exchangeable laws as convex mixtures of urns. But we re-instate the idea from original de Finetti, kept in

[1] Analogously, $(Z_1, \ldots, Z_k)$ is called infinitely extendible if its law equals that of the first k elements of an infinite exchangeable sequence $(\tilde{Z}_1, \tilde{Z}_2, \ldots)$.

the signed-measure approach, that the parameter space of the superposition should consist of probability measures on the *original* Polish space X, not its N-fold product. In principle, this is possible by parametrizing urn laws by their one-point marginals, which are easily seen to be in 1–1 correspondence with these laws. In practice, to arrive at an explicit representation one needs an explicit formula for the inverse of this marginal map. By deriving such a formula, we obtain a unique representation of finitely exchangeable laws which sheds some light on their universal correlation structure and is useful for applications.

Main results In terms of laws, N-extendibility turns into what has been called N-representability [16]: for $k \leq N$, a k-point probability measure μ_k on X^k, or k-plan for short, is called N-representable if it is the k-point marginal of a symmetric N-point probability measure μ_N on X^N (see Definition 2.1).

As a first main result, we explicitly determine the extremal N-representable k-plans, that is, those that cannot be written as strict convex combinations of any other N-representable k-plans, and give a polynomial parametrization in terms of their one-point marginals. Focusing in this introduction for simplicity on the case $k = 4$, these are the probability measures

$$F_{N,4}(\lambda) = \frac{N^3}{(N-1)(N-2)(N-3)} \left[\lambda^{\otimes 4} - \frac{6}{N} S_4 \, \mathrm{id}_\#^{\otimes 2}\lambda \otimes \lambda^{\otimes 2} + \frac{8\, S_4 \, \mathrm{id}_\#^{\otimes 3}\lambda \otimes \lambda + 3\, S_4 \, \mathrm{id}_\#^{\otimes 2}\lambda \otimes \mathrm{id}_\#^{\otimes 2}\lambda}{N^2} - \frac{6}{N^3} \, \mathrm{id}_\#^{\otimes 4}\lambda \right]$$

where λ is a $\frac{1}{N}$-quantized probability measure on X, i.e. an empirical measure of the form $\frac{1}{N}\sum_{i=1}^N \delta_{x_i}$ for some—not necessarily distinct—points $x_i \in X$. It is not obvious, but part of our result, that these measures are nonnegative, and different for different λ.

The above expression can be viewed as a degree-4 symmetric polynomial in λ. Besides an overall positive prefactor, the polynomial has leading term $\lambda^{\otimes 4}$ which is homogeneous of degree 4 and uncorrelated, and alternating corrections of order $\frac{1}{N^j}$ which are homogeneous of degree $4-j$ and more and more strongly correlated. As N tends to infinity $F_{N,4}(\lambda)$ approaches the independent measure $\lambda^{\otimes 4}$, recovering the basis in the de Finetti representation for infinitely representable 4-plans implied by (1.1). The correlated corrections are of significant size even when N is quite large; see Fig. 1. All these findings persist for general k; see Theorem 4.5 for the general expression $F_{N,k}(\lambda)$ for extremal N-representable k-plans. Qualitatively, the corrections to independence form a finite exponentially decaying series; quantitatively the (rational) coefficients which appear can be related to the analytic continuation of the Ewens function from genetics, which we introduce for this purpose.

Our second contribution is to cast the abstract insight [26, 28] that finite exchangeables are convex mixtures of urn sequences into a quantitative polynomial formula. We show that any N-representable k-plan is a convex mixture of the $F_{N,k}(\lambda)$. More precisely, a k-plan is N-representable if and only if it is of the form

$$\mu_k = \int_{\mathcal{P}_{\frac{1}{N}}(X)} F_{N,k}(\lambda)\, \mathrm{d}\alpha(\lambda) \tag{1.2}$$

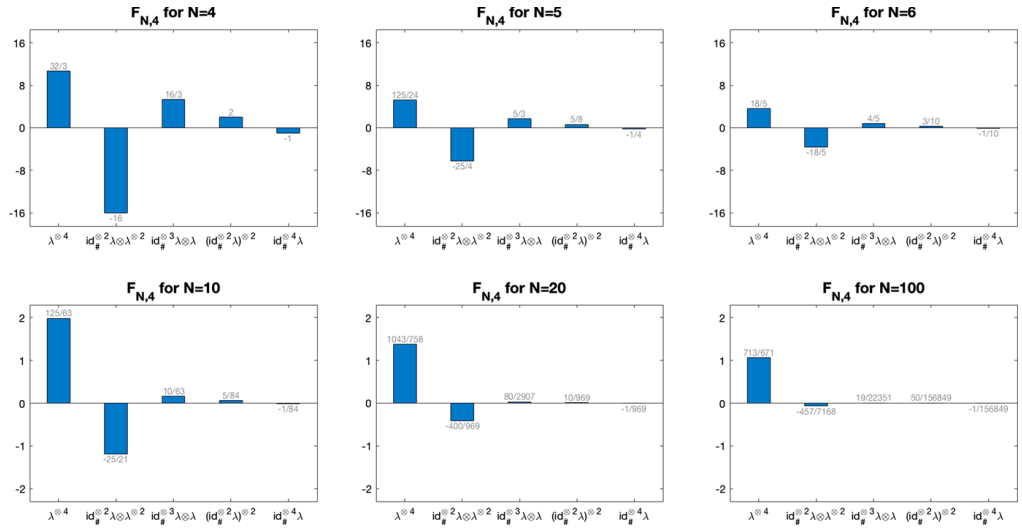

Fig. 1 Coefficients of the universal polynomial $F_{N,4}$ for different N. For $N = 5$ and 6, the second (correlated) term is bigger respectively equal in absolute value to the first (independent) term; for $N = 20$ its size is about 30% that of the first term. For large N, $F_{N,4}$ converges to the independent measure $\lambda^{\otimes 4}$, but even for $N = 100$ the deviation from the latter is still visible

for some probability measure α on the set $\mathcal{P}_{\frac{1}{N}}(X)$ of $\frac{1}{N}$-quantized probability measures on X. Moreover if $N = k$ the measure α is unique, giving a one-to-one parametrization of the laws of finitely exchangeable sequences. By contrast, the signed measure representation of such laws is not unique [24], and not all signed measures give rise to such a law.

Formula (1.2) generalizes de Finetti–Hewitt–Savage, (1.1), from infinitely to finitely representable measures, or from infinite to finite exchangeable sequences of random variables. Formally, in the limit $N \to \infty$ the domain of integration in (1.2) tends to all of $\mathcal{P}(X)$, and the integrand tends to the independent measure $\lambda^{\otimes k}$, recovering de Finetti (see Sect. 6 for a rigorous account). In Bayesian language, formula (1.2) says that the general N-extendible sequence $(Z_1, \ldots, Z_k)$ of X-valued random variables is obtained by first picking some $\frac{1}{N}$-quantized distribution λ on X at random from some prior, then sampling $(Z_1, \ldots, Z_k)$ from the correlated distribution $F_{N,k}(\lambda)$. In particular, by setting $k = N$ we conclude that *the general finite exchangeable sequence* $(Z_1, \ldots, Z_N)$ *is obtained by picking* λ *at random from its – in the case* $k = N$ *unique – prior* α *in* (1.2)*, then sampling* $(Z_1, \ldots, Z_N)$ *from* $F_{N,N}(\lambda)$.

Sampling from $F_{N,k}(\lambda)$ has a transparent probabilistic meaning which we now explain in the language of urns. Write a given $\frac{1}{N}$-quantized probability measure λ as $\frac{1}{N}\sum_{i=1}^{N}\delta_{x_i}$ for N not necessarily distinct points $x_1, \ldots, x_N \in X$. Now pick, in turn, k of these points at random *without replacement*, and denote the so-obtained sequence by $(Z_1, \ldots, Z_k)$. By construction, the law of this sequence is the k-point marginal μ_k of the symmetrization of the Dirac measure $\delta_{(x_1,\ldots,x_N)}$ on X^N. But the polynomial $F_{N,k}$ is precisely constructed as the inverse of the marginal map $\mu_k \mapsto \lambda$ (see eq. (4.6), eq. (4.8), and Theorem 4.5 below). Hence $F_{N,k}(\lambda) = \mu_k$, and so $(Z_1, \ldots, Z_k)$ is the sought-after finite N-extendible sequence. We find it quite remarkable that the extremal N-representable k-plans $F_{N,k}(\lambda)$—which emerge purely from

convex geometric considerations—have such a simple probabilistic meaning, being the laws of classical examples [2] of finite exchangeable sequences which are not infinitely extendible.

Recovering the prior from sampling A nice aspect of our representation of the law of a general finitely exchangeable sequence $(Z_1, \dots, Z_N)$ (eq. (1.2) with $k = N$) is that the prior α, which is unique when $k = N$, can be determined by sampling, as follows. Let

$$\left(Z_1^{(\nu)}, \dots, Z_N^{(\nu)}\right)_{\nu=1}^n$$

be a sequence of n independent samples in X^N. Form the $\mathcal{P}_{\frac{1}{N}}(X)$-valued sequence

$$\lambda^{(\nu)} = \frac{1}{N} \sum_{i=1}^N \delta_{Z_i^{(\nu)}}.$$

Then the empirical measure

$$\frac{1}{n} \sum_{\nu=1}^n \delta_{\lambda^{(\nu)}}$$

converges almost surely to α. See Corollary 5.2.

The paper is organized as follows. After introducing some notations and preliminaries in Sects. 2 and 3 deals with $\frac{1}{N}$-quantized measures. In Sect. 4, we focus on the finite case; we first recall the results of [17] and then identify the universal correlated polynomials $F_{N,k}(\lambda)$. Section 5 extends these findings to the case of a Polish state space X, and show in addition that the $F_{N,k}(\lambda)$ are in fact exposed N-representable k-plans. Section 6 discusses connections with the Hewitt and Savage theorem and the Diaconis–Freedman error bounds. Section 7 gives an unexpected connection with the Ewens sampling formula from genetics. Finally, Sect. 8 is devoted to applications to MMOT emphasizing the connection with convexification of polynomials.

2 Preliminaries and notations

In the sequel X will denote a Polish (i.e., complete and separable metric) space. The principal example we have in mind is $X = \mathbb{R}^d$, in which case all of our results are already new and interesting. In this case the metric is the usual Euclidean metric $d(x, y) = |x - y| = (\sum_{i=1}^d (x_i - y_i)^2))^{1/2}$. We denote by $\mathcal{P}(X)$ the set of Borel probability measures on X. Probability measures on X^k will be called k-plans. From now on, we fix two integers k and N with $1 \le k \le N$. Given $\gamma \in \mathcal{P}(X^N)$ we denote by $M_k\gamma$ the k-point-marginal of γ, i.e.,

$$(M_k\gamma)(A) := \gamma(A \times X^{N-k}) \text{ for every Borel subset } A \text{ of } X^k \tag{2.1}$$

(with the convention $M_N\gamma = \gamma$).

We denote by $C_b(X^N)$ the space of bounded and continuous functions on X^N, and by $\mathcal{S}_N$ the group of permutations of $\{1,\dots,N\}$. For $\gamma \in \mathcal{P}(X^N)$ and $\sigma \in \mathcal{S}_N$, the measure $\gamma^\sigma \in \mathcal{P}(X^N)$ is defined by

$$\int_{X^N} \varphi \mathrm{d}\gamma^\sigma = \int_{X^N} \varphi(x_{\sigma(1)},\dots,x_{\sigma(N)})\mathrm{d}\gamma(x_1,\dots,x_N)$$

for every test-function $\varphi \in C_b(X^N)$. A measure $\gamma \in \mathcal{P}(X^N)$ is called symmetric if $\gamma = \gamma^\sigma$ for every $\sigma \in \mathcal{S}_N$. If $\gamma \in \mathcal{P}(X^N)$ is arbitrary, its symmetrization $S_N\gamma$ is given by

$$S_N\gamma := \frac{1}{N!}\sum_{\sigma\in S_N} \gamma^\sigma. \tag{2.2}$$

The symmetrization operator $S_N : \gamma \mapsto S_N\gamma$ is a linear projection operator on $\mathcal{P}(X^N)$, i.e. it maps $\mathcal{P}(X^N)$ linearly into itself and satisfies $(S_N)^2 = S_N$; and γ is symmetric if and only if $S_N\gamma = \gamma$. The set of symmetric N-plans is denoted by $\mathcal{P}_{\mathrm{sym}}(X^N)$:

$$\mathcal{P}_{\mathrm{sym}}(X^N) := \{\gamma \in \mathcal{P}(X^N) \ : \ \gamma = S_N\gamma\}. \tag{2.3}$$

We shall use the notation # to denote the push-forward measure, that is, given two Polish spaces Y and Z, a Borel map T from Y to Z and a Borel probability measure μ on Y, then $T_\#\mu$ is the probability measure on Z defined by $T_\#\mu(B) = \mu(T^{-1}(B))$ for every Borel subset B of Z; equivalently, for every real-valued, continuous and bounded function φ on Z:

$$\int_Z \varphi \ \mathrm{d}T_\#\mu = \int_Y \varphi\circ T \ \ \mathrm{d}\mu.$$

We recall the following definition from [16].

Definition 2.1 For $N \in \mathbb{N}$ and $k \in \{1,\dots,N\}$, a k-plan $\mu_k \in \mathcal{P}(X^k)$ is said to be N-representable if it is the k-point marginal of a symmetric N-plan, that is to say if there exists $\gamma \in \mathcal{P}_{\mathrm{sym}}(X^N)$ such that $\mu_k = M_k\gamma$. We denote by $\mathcal{P}_{N-\mathrm{rep}}(X^k)$ the set of N-representable k-plans, i.e.:

$$\mathcal{P}_{N-\mathrm{rep}}(X^k) = \{M_kS_N\tilde{\gamma} \ : \ \tilde{\gamma} \in \mathcal{P}(X^N)\} = \{M_k\gamma \ : \ \gamma \in \mathcal{P}_{\mathrm{sym}}(X^N)\}.$$

In probabilistic terms, a symmetric N-plan $\gamma \in \mathcal{P}_{\mathrm{sym}}(X^N)$ is the law of a finite exchangeable random sequence $(Z_1,\dots,Z_N)$ with values in X^N, whereas $\mu_k = M_k\gamma \in \mathcal{P}_{N-\mathrm{rep}}(X^k)$ is the law of its first k-components $(Z_1,\dots,Z_k)$.

We will work with the following standard notion of convergence in $\mathcal{P}(X)$ (as well as $\mathcal{P}(X^k)$, $\mathcal{P}(X^N)$, …). Recall that $C_b(X)$ denotes the space of bounded continuous functions on X.

Definition 2.2 A sequence $(\mu_\nu)_{\nu\in\mathbb{N}}$ of probability measures in $\mathcal{P}(X)$ is said to converge narrowly to $\mu \in \mathcal{P}(X)$ if

$$\lim_{\nu\to\infty}\int_X \varphi\mathrm{d}\mu_\nu = \int_X \varphi\mathrm{d}\mu \ \text{ for all } \varphi \in C_b(X).$$

Thus, in applications to statistical physics where the probability measures live on a space of particle configurations, narrow convergence corresponds to convergence of bounded continuous observables.

We note the following basic topological property of the set of N-representable k-plans.

Lemma 2.3 *The set $\mathcal{P}_{N-\mathrm{rep}}(X^k)$ is closed under narrow convergence.*

Proof Let $\{\mu_\nu\}_{\nu\in\mathbb{N}}$ be a narrowly convergent sequence in $\mathcal{P}_{N-\mathrm{rep}}(X^k)$. Write μ_ν as $M_k\gamma_\nu$ for $\gamma_\nu \in \mathcal{P}_{\mathrm{sym}}(X^N)$. Since $\mu_\nu = M_k\gamma_\nu$ is narrowly convergent, it is tight and hence so is γ_ν. By Prokhorov's theorem, γ_ν has a narrowly convergent subsequence which converges to some $\gamma \in \mathcal{P}_{\mathrm{sym}}(X^N)$. Since M_k is narrowly continuous, μ_ν converges to $M_k\gamma \in \mathcal{P}_{N-\mathrm{rep}}(X^k)$.
We recall that narrow convergence on $\mathcal{P}(X)$ is metrizable. For instance one may start from the metric d on X, truncate it to the (topologically equivalent) bounded metric $\tilde{d}(x, y) := \min\{d(x, y), 1\}$, and use the associated 1-Wasserstein metric on $\mathcal{P}(X)$

$$W_1(\lambda_1, \lambda_2) := \inf_{\theta\in\Pi(\lambda_1,\lambda_2)} \left\{ \int_{X\times X} \tilde{d}(x, y)\, \mathrm{d}\theta(x, y) \right\} \text{ for all } (\lambda_1, \lambda_2) \in \mathcal{P}(X)^2, \quad (2.4)$$

where $\Pi(\lambda_1, \lambda_2)$ is the set of transport plans between λ_1 and λ_2, i.e., the set of Borel probability measures on $X \times X$ having λ_1 and λ_2 as marginals. Then the (bounded) metric W_1 metrizes narrow convergence on $\mathcal{P}(X)$ (that is, μ_ν converges narrowly to μ if and only if $W_1(\mu_\nu, \mu)$ tends to zero) and $(\mathcal{P}(X), W_1)$ is itself a Polish space.

Also we recall the definition of the total variation distance between two signed measures μ and ν on X:

$$\|\mu - \nu\|_{\mathrm{TV}} := \sup \Big\{ |\mu(A) - \nu(A)| \ : \ A \text{ Borel subset of } X \Big\}. \quad (2.5)$$

□

3 $\frac{1}{N}$-quantized probability measures

An important role will be played by the set of $\frac{1}{N}$-quantized probability measures on the Polish space X,

$$\mathcal{P}_{\frac{1}{N}}(X) := \left\{ \frac{1}{N} \sum_{i=1}^{N} \delta_{x_i} \ : \ x_1, \ldots, x_N \in X^N \text{ (not necessarily distinct)} \right\}. \quad (3.1)$$

It is easy to see that this set can also be written as

$$\mathcal{P}_{\frac{1}{N}}(X) = \left\{ \lambda \in \mathcal{P}(X) \ : \ \lambda(A) \in \left\{ 0, \frac{1}{N}, \ldots, 1 \right\} \text{ for every Borel subset } A \text{ of } X \right\}. \quad (3.2)$$

In the special case of finite state spaces X, this set was introduced—and utilized to parametrize extremal N-representable measures—in [17]. Let us collect two basic properties of this set which hold for general state spaces.

Lemma 3.1 (Quadratic constraint characterization and closedness of $\frac{1}{N}$-quantized probability measures)

(a) $\lambda \in \mathcal{P}(X)$ *belongs to* $\mathcal{P}_{\frac{1}{N}}(X)$ *if and only if*

$$\lambda \otimes \lambda - \frac{2k+1}{N} \mathrm{id}_{\#}^{\otimes 2} \lambda \geq -\frac{k(k+1)}{N^2}, \text{ for } k = 0, \ldots, N-1. \tag{3.3}$$

(b) $\mathcal{P}_{\frac{1}{N}}(X)$ *is closed under narrow convergence.*

In formula (3.3) and in the sequel, $\lambda^{\otimes \ell}$ denotes the ℓ-fold tensor product of λ with itself and $\mathrm{id}_{\#}^{\otimes \ell} \lambda$ is defined by

$$\int_{X^\ell} \varphi \, \mathrm{d}\, \mathrm{id}_{\#}^{\otimes \ell} \lambda := \int_X \varphi(x, \ldots, x) \mathrm{d}\lambda(x), \quad \text{for all } \varphi \in C_b(X^\ell).$$

Proof

(a) The nontrivial implication is that (3.3) implies that $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$. Let A be any Borel subset of X. Applying the measure on the left hand side of (3.3) to $A \times A$ gives $f(\lambda(A)) \geq 0$, where f is the scalar function $f(t) = t^2 - \frac{2k+1}{N} t + \frac{k(k+1)}{N^2} = (t - \frac{k}{N})(t - \frac{k+1}{N})$. But f is negative precisely in the open interval $(\frac{k}{N}, \frac{k+1}{N})$, whence $\lambda(A)$ does not lie in this interval. Since this holds for all $k = 0, \ldots, N-1$, it follows that λ belongs to the set (3.2).

(b) The maps $\lambda \in \mathcal{P}(X) \mapsto \lambda^{\otimes 2} \in \mathcal{P}(X^2)$ and $\lambda \in \mathcal{P}(X) \mapsto \mathrm{id}_{\#}^{\otimes 2} \lambda \in \mathcal{P}(X^2)$ are continuous with respect to narrow convergence, hence so is the left hand side of (3.3). The assertion now follows from a). □

4 Extremal *N*-representable *k*-plans on finite state spaces

Throughout this section, we assume $N \geq 2$ and restrict our attention to a finite state space X consisting of ℓ distinct points,

$$X = \{a_1, \ldots, a_\ell\}. \tag{4.1}$$

4.1 Extreme points

Our goal is to describe the geometry of the convex set of N-representable k-plans on the finite state space X, i.e., $\mathcal{P}_{N-\mathrm{rep}}(X^k)$. This set is a compact polyhedron in a finite-dimensional vector space and therefore coincides, by Minkowski's theorem (see e.g. [23]), with the convex hull of its extreme points. Therefore, classifying the extreme points is one way to characterize the geometry of the object. We recall that a point x

in a convex set K is an extreme point if, whenever $x = \alpha x_1 + (1-\alpha)x_2$ for some x_1, $x_2 \in K$ and some $\alpha \in (0,1)$, we have that $x_1 = x_2 = x$. The set of extreme points of K will be denoted ext K.

In [17] the extremal N-representable k-plans are determined in the case $k = 2$ and $k = N$. They correspond exactly to the symmetrized Dirac measures respectively their two-point marginals:

Theorem 4.1 [17]

(a) *A measure μ on X^N is an extreme point of $\mathcal{P}_{\mathrm{sym}}(X^N)$ if and only if it is of the form*

$$S_N \delta_{a_{i_1},\ldots,a_{i_N}} \qquad \text{for some } 1 \le i_1 \le \ldots \le i_N \le \ell. \tag{4.2}$$

Moreover different index vectors $(i_1,\ldots,i_N)$ with $1 \le i_1 \le \ldots \le i_N \le \ell$ yield different extreme points.

(b) A measure μ on X^2 is an extreme point of $\mathcal{P}_{N-\mathrm{rep}}(X^2)$ if and only if it is of the form

$$M_2 S_N \delta_{a_{i_1},\ldots,a_{i_N}} \qquad \text{for some } 1 \le i_1 \le \ldots \le i_N \le \ell. \tag{4.3}$$

(c) Moreover the marginal maps $M_2 : \mathrm{ext}\,(\mathcal{P}_{\mathrm{sym}}(X^N)) \to \mathrm{ext}\,(\mathcal{P}_{N-\mathrm{rep}}(X^2))$ and $M_1 : \mathrm{ext}\,(\mathcal{P}_{\mathrm{sym}}(X^N)) \to \mathcal{P}_{\frac{1}{N}}(X)$ are bijections.

Here (a) and the fact that the set of measures in (4.3) contains the set ext $(\mathcal{P}_{N-\mathrm{rep}}(X^2))$ of extremal N-representable two-plans is easy to see, but the reverse inclusion and the bijectivity of M_2 between extremal symmetric N-plans and extremal N-representable two-plans is nontrivial; geometrically it says that none of the corners of the high-dimensional polytope $\mathcal{P}_{\mathrm{sym}}(X^N)$ is mapped into the interior (or face interior or edge interior) of the low-dimensional polytope $\mathcal{P}_{N-\mathrm{rep}}(X^2)$ by the highly non-injective marginal map M_2. Using this nontrivial fact it is easy to extend Theorem 4.1 to an arbitrary choice of $k \in \{2,\ldots,N\}$.

Theorem 4.2 *A measure μ on X^k is an extreme point of $\mathcal{P}_{N-\mathrm{rep}}(X^k)$ if and only if it is of the form*

$$M_k S_N \delta_{a_{i_1},\ldots,a_{i_N}} \qquad \text{for some } 1 \le i_1 \le \ldots \le i_N \le \ell. \tag{4.4}$$

Moreover the marginal map $M_k : \mathrm{ext}\,(\mathcal{P}_{\mathrm{sym}}(X^N)) \to \mathrm{ext}\,(\mathcal{P}_{N-\mathrm{rep}}(X^k))$ is a bijection.

Proof We will abbreviate $\mathcal{P}_{\mathrm{sym}}(X^N) = \mathcal{P}_{\mathrm{sym}}$, $(i_1,\ldots,i_N) = i$, $\{(i_1,\ldots,i_N) : 1 \le i_1 \le \ldots \le i_N \le \ell\} = \mathcal{I}$. By the definition of N-representability, $\mathcal{P}_{N-\mathrm{rep}}(X^k) = M_k \mathcal{P}_{\mathrm{sym}}$. Using, in order of appearance, this fact, the linearity of M_k, and Theorem 4.1 (a), we have

$$\mathrm{ext}\,(\mathcal{P}_{N-\mathrm{rep}}(X^k)) = \mathrm{ext}\,(M_k \mathcal{P}_{\mathrm{sym}}) \subseteq M_k \mathrm{ext}\,(\mathcal{P}_{\mathrm{sym}}) = \{M_k S_N \delta_{a_{i_1},\ldots,a_{i_N}} : i \in \mathcal{I}\}. \tag{4.5}$$

To establish the reverse inclusion it suffices to show that the number of elements of the set on the left is bigger or equal that of the set on the right. By the fact that $M_2 = M_2 M_k$

and the linearity of M_2,

$$\text{ext}\,(M_2\mathcal{P}_{\text{sym}}) = \text{ext}\,(M_2M_k\mathcal{P}_{\text{sym}}) \subseteq M_2\text{ext}\,(M_k\mathcal{P}_{\text{sym}}) = M_2\text{ext}\,(\mathcal{P}_{N-\text{rep}}(X^k))$$

and consequently $|\text{ext}\,(M_2\mathcal{P}_{\text{sym}})| \leq |\text{ext}\,(\mathcal{P}_{N-\text{rep}}(X^k))|$, where $|\cdot|$ denotes the number of elements of a set. Combining this inequality with the bijectivity property of M_2 in Theorem 4.1 (b) and Theorem 4.1 (a) yields

$$\begin{aligned}|\text{ext}\,(\mathcal{P}_{N-\text{rep}}(X^k))| &\geq |\text{ext}\,(M_2\mathcal{P}_{\text{sym}})| = |\text{ext}\,(\mathcal{P}_{\text{sym}})| = |\{S_N\delta_{a_{i_1},\ldots,a_{i_N}}\colon\, i \in \mathcal{I}\}| \\ &\geq |\{M_kS_N\delta_{a_{i_1},\ldots,a_{i_N}}\colon\, i \in \mathcal{I}\}|.\end{aligned}$$

□

Note that if $k = 1$ (which is the case where N-representability does not impose geometric restrictions), the statement in Theorem 4.2 is false: $\mathcal{P}_{N-\text{rep}}(X^1) = \mathcal{P}(X)$ whose extreme points are the Dirac masses. In [17] it was established that for the set (4.1) consisting of ℓ distinct points, the cardinality of $\mathcal{P}_{\frac{1}{N}}(X)$—and hence, by Theorem 4.1 (b), the number of extreme points of $\mathcal{P}_{\text{sym}}(X^N)$—equals $\binom{N+\ell-1}{\ell-1}$. Now the following corollary is an immediate consequence of Theorem 4.2.

Corollary 4.3 *For any $k \in \{2,\ldots,N\}$, $\mathcal{P}_{N-\text{rep}}(X^k)$ has $\binom{N+\ell-1}{\ell-1}$ extreme points.*

Combining the isomorphisms M_k and M_1 from Theorems 4.2 respectively 4.1 shows that the extreme points of $\mathcal{P}_{N-\text{rep}}(X^k)$, i.e. the k-plans of form (4.4), can be uniquely recovered from their one-point marginals $\frac{1}{N}(\delta_{a_{i_1}}+\cdots+\delta_{a_{i_N}}) \in \mathcal{P}_{\frac{1}{N}}(X)$. But the above abstract reasoning does not provide a convenient formula for the recovery map. This issue is dealt with in the next section.

4.2 A universal polynomial formula for extreme points in terms of their one-point marginals

Our aim now is to derive an explicit polynomial formula for the extremal measures (4.4) in terms of their one-point marginals. In order to do so we consider any extremal N-representable k-plan

$$\mu_k := M_kS_N\delta_{x_1,\ldots,x_N} \tag{4.6}$$

for $(x_1,\ldots,x_N) \in X^N$.

In [17] it was shown that, in the case of $k = 2$, μ_2 can be expressed explicitly as

$$\mu_2 = \frac{N}{N-1}\lambda^{\otimes 2} - \frac{1}{N-1}\,\text{id}_{\#}^{\otimes 2}\,\lambda, \tag{4.7}$$

where

$$\lambda := M_1\mu_k = \frac{1}{N}\sum_{i=1}^{N}\delta_{x_i} \in \mathcal{P}_{\frac{1}{N}}(X) \tag{4.8}$$

is the one-point marginal of μ_k. (Recall the notation $\lambda^{\otimes \ell}$ and $\mathrm{id}_{\#}^{\otimes \ell}\,\lambda$ for the ℓ-fold tensor product of λ with itself respectively the push-forward of λ under the ℓ-fold cartesian product of the identity; see the end of Sect. 2.) A similar computation for $k = 3$ gives

$$\mu_3 = \frac{N^2}{(N-1)(N-2)}\left[\lambda^{\otimes 3} - \frac{3}{N} S_3\Big((\mathrm{id}_{\#}^{\otimes 2}\,\lambda)\otimes\lambda\Big) + \frac{2}{N^2}\,\mathrm{id}_{\#}^{\otimes 3}\,\lambda\right]. \tag{4.9}$$

(For a justification of (4.9) using our general results see the examples below Theorem 4.5.) In view of (4.7) and (4.9), it is natural to look for a similar *polynomial* of degree k in λ expression of μ_k, consisting of a mean field term $\lambda^{\otimes k}$ and corrections of order $\frac{1}{N^j}$ for $j = 1, \ldots, k-1$. As turns out, the j^{th} order correction is related to the partitions of the number j.

Definition 4.4 Let $\mathbb{N} = \{1, 2, 3, \ldots\}$ denote the set of positive integers. A partition of $j \in \mathbb{N}$ of length $n \in \mathbb{N}$ is a vector $\mathbf{p} = (p_1, \ldots, p_n) \in \mathbb{N}^n$ such that $\sum_{i=1}^n p_i = j$, $p_1 \geq \cdots \geq p_n$. For any partition $\mathbf{p}$ we denote its length by $n(\mathbf{p})$.

For example, the partitions of 4 are

$$\begin{aligned} &1+1+1+1\\ &2+1+1\\ &2+2\\ &3+1\\ &4. \end{aligned}$$

This corresponds in the above notation to $\mathbf{p} = (1,1,1,1) \in \mathbb{N}^4$, $\mathbf{p} = (2,1,1) \in \mathbb{N}^3$, $\mathbf{p} = (2,2) \in \mathbb{N}^2$, $\mathbf{p} = (3,1) \in \mathbb{N}^2$, and $\mathbf{p} = 4 \in \mathbb{N}$.

Theorem 4.5 *Let $N \geq 2$, $k \in \{2, \ldots, N\}$. Any extremal N-representable k-plan μ_k (see (4.6)) can be written in terms of its one-point marginal λ (see (4.8)) as*

$$\mu_k = \frac{N^{k-1}}{\prod_{i=1}^{k-1}(N-i)}\left[\lambda^{\otimes k} + \sum_{j=1}^{k-1}\frac{(-1)^j}{N^j} S_k P_j^{(k)}(\lambda)\right] =: F_{N,k}(\lambda) \tag{4.10}$$

where for $j = 1, \ldots, k-1$

$$P_j^{(k)}(\lambda) = \sum_{\substack{\mathbf{p}=(p_1,\ldots p_{n(\mathbf{p})})\ \text{partition}\\ \text{of } j \text{ with } j+n(\mathbf{p})\leq k}} d_{\mathbf{p}}^{(k)}\ \mathrm{id}_{\#}^{\otimes(p_1+1)}\lambda\ \otimes\ldots\otimes\ \mathrm{id}_{\#}^{\otimes(p_{n(\mathbf{p})}+1)}\lambda\ \otimes\lambda^{\otimes(k-j-n(\mathbf{p}))} \tag{4.11}$$

with positive coefficients $d_{\mathbf{p}}^{(k)}$ given by

$$d_{\mathbf{p}}^{(k)} = \frac{k!}{(k-j-n(\mathbf{p}))!}\prod_{i=1}^{n(\mathbf{p})}\frac{1}{p_i+1}\prod_{q\in Ran\,\mathbf{p}}\frac{1}{(|\mathbf{p}^{-1}(q)|)!}\quad . \tag{4.12}$$

Moreover the coefficients satisfy the sum rule

$$\sum_{\substack{\mathbf{p}\ \text{partition of } j \\ \text{with } j+n(\mathbf{p})\le k}} d_{\mathbf{p}}^{(k)} = \sum_{1\le i_1<\ldots<i_j\le k-1} i_1\cdot\ldots\cdot i_j =: c_j^{(k)} \quad (j=1,\ldots,k-1). \tag{4.13}$$

In particular, $||P_j^{(k)}(\lambda)||_{TV} = \int_{X^k} dP_j^{(k)}(\lambda) = c_j^{(k)}$.

In the last term in Eq. (4.12), a partition $\mathbf{p}$ is viewed as a map from the set of its component indices to $\mathbb{N}$; the range $\operatorname{Ran}\mathbf{p}$ of this map is the set of values taken by the components, and $|\mathbf{p}^{-1}(q)|$ denotes the number of components with value q. For example, for $\mathbf{p}=(3,1,1)$ and $q=1$, $|\mathbf{p}^{-1}(q)|=2$. The factor $(|\mathbf{p}^{-1}(q)|)!$ in the denominator says that a partition with many repeat components contributes much less than a partition with few repeat components.

Some remarks are in order.

(1) The first term in expression (4.10) for the extreme points is a mean field term and the remaining terms are correlation corrections. We emphasize that the $P_j^{(k)}$ are independent of N and hence the correlation corrections form a *finite series* in inverse powers of N.

(2) The $P_j^{(k)}$ are polynomials of degree $k-j$ in λ.

(3) Due to the presence of the signs $(-1)^j$, it is far from trivial that $F_{N,k}(\lambda)$ is a nonnegative measure; but it must be, e.g. because the left hand side of (4.10) equals (4.6). Nonnegativity relies on a subtle interplay between the explicit coefficients in Theorem 4.5 and the quantization condition $\lambda\in\mathcal{P}_{\frac{1}{N}}(X)$, and does not hold for arbitrary $\lambda\in\mathcal{P}(X)$.

(4) The coefficients $c_j^{(k)}=\sum i_1\cdot\ldots\cdot i_j$ introduced in (4.13) which measure the total mass of the jth-order correction to independence are related to the well-known Stirling numbers, particularly the (absolute) Stirling numbers of the first kind. For given natural numbers $q,r\in\mathbb{N}\cup\{0\}$ with $r\le q$ the corresponding absolute Stirling number of the first kind $s(q,r)$ gives the number of permutations of $\{1,\ldots,q\}$ that decompose into r cycles, with the convention that $s(q,r)$ is zero when exactly one of q and r is zero and that $s(0,0)=1$. From well-known expressions for Stirling numbers one can see that the following holds

$$c_j^{(k)} = s(k,k-j),$$

that is to say the present coefficients $c_j^{(k)}$ equal the number of permutations of $\{1,\ldots,k\}$ that decompose into $k-j$ cycles. For more information about Stirling numbers we refer the interested reader to [7].

(5) By combining Theorems 4.5 and 4.2, and the isomorphism property of M_1 from Theorem 4.1, we immediately obtain:

Corollary 4.6 *A measure μ_k on X^k is an extreme point of $\mathcal{P}_{N-\mathrm{rep}}(X^k)$ if and only if it is of the form $\mu_k = F_{N,k}(\lambda)$ for some $\lambda\in\mathcal{P}_{\frac{1}{N}}(X)$, with $F_{N,k}$ given by*

(4.10)–(4.12). *Moreover* $F_{N,k}(\lambda)$ *has one-point marginal* λ*, and* j*-point marginal* $F_{N,j}(\lambda)$ *for* $2 \leq j \leq k-1$.

We now write out the universal polynomials $F_{N,k}$ explicitly for small k.

Example: $k = 2$ The finite sum over j in (4.10) reduces to a single term for $j = 1$, $-\frac{1}{N} S_2 P_1^{(2)}(\lambda)$, and $P_1^{(2)}(\lambda)$ consists of a single term associated with the only partition $\mathbf{p} = 1$ of 1, $id_{\#}^{\otimes 2}\lambda$. Consequently

$$F_{N,2}(\lambda) = \frac{N}{N-1}\left[\lambda^{\otimes 2} - \frac{1}{N}\,\mathrm{id}_{\#}^{\otimes 2}\,\lambda\right].$$

This expression agrees with (4.7), and so Theorem 4.5 recovers [17] Theorem 2.1.

Example: $k = 3$ The finite series in (4.10) runs from $j = 1$ to $j = 2$, and for these two values of j, the partitions contributing to the sum in (4.11) are the partitions $\mathbf{p}$ of j satisfying $j + n(\mathbf{p}) \leq 3$. These partitions and the associated coefficient $d_{\mathbf{p}}$ given by (4.12) are

j	Partitions of j with $j + n(\mathbf{p}) \leq 3$	Our notation: $\mathbf{p} =$	Coefficient $d_{\mathbf{p}} =$
1	1	1	3
2	2	2	2

and consequently

$$F_{N,3}(\lambda) = \frac{N^2}{(N-1)(N-2)}\left[\lambda^{\otimes 3} - \frac{3}{N} S_3\Big((\mathrm{id}_{\#}^{\otimes 2}\,\lambda)\otimes\lambda\Big) + \frac{2}{N^2}\,\mathrm{id}_{\#}^{\otimes 3}\,\lambda\right].$$

Example: $k = 4$ By formulae (4.10)–(4.12), the partitions $\mathbf{p}$ of j contributing to the j^{th} order correction are:

j	Partitions of j with $j + n(\mathbf{p}) \leq 4$	Our notation: $\mathbf{p} =$	Coefficient $d_{\mathbf{p}} =$
1	1	1	6
2	2	2	8
	1+1	(1,1)	3
3	3	3	6

and consequently

$$F_{N,4}(\lambda) = \frac{N^3}{(N-1)(N-2)(N-3)}\left[\lambda^{\otimes 4} - \frac{6}{N} S_4\,\mathrm{id}_{\#}^{\otimes 2}\lambda\otimes\lambda^{\otimes 2} + \frac{8\,S_4\,\mathrm{id}_{\#}^{\otimes 3}\lambda\otimes\lambda + 3\,S_4\,\mathrm{id}_{\#}^{\otimes 2}\lambda\otimes\mathrm{id}_{\#}^{\otimes 2}\lambda}{N^2} - \frac{6}{N^3}\,\mathrm{id}_{\#}^{\otimes 4}\lambda\right].$$

Example: k = 5 The contributing partitions are

j	Partitions of j with $j + n(\mathbf{p}) \le 5$	Our notation: $\mathbf{p} =$	Coefficient $d_{\mathbf{p}} =$
1	1	1	10
2	2	2	20
	1+1	(1,1)	15
3	3	3	30
	2+1	(2,1)	20
4	4	4	24

and so

$$F_{N,5}(\lambda) = \frac{N^4}{(N-1)(N-2)(N-3)(N-4)} \Bigg[\lambda^{\otimes 5} - \frac{10}{N} S_5 \, \mathrm{id}_{\#}^{\otimes 2}\lambda \otimes \lambda^{\otimes 3} + \frac{20\, S_5 \, \mathrm{id}_{\#}^{\otimes 3}\lambda \otimes \lambda^{\otimes 2} + 15\, S_5 \, \mathrm{id}_{\#}^{\otimes 2}\lambda \otimes \mathrm{id}_{\#}^{\otimes 2}\lambda \otimes \lambda}{N^2} - \frac{30\, S_5 \, \mathrm{id}_{\#}^{\otimes 4}\lambda \otimes \lambda + 20\, S_5 \, \mathrm{id}_{\#}^{\otimes 3}\lambda \otimes \mathrm{id}_{\#}^{\otimes 2}\lambda}{N^3} + \frac{24}{N^4} \, \mathrm{id}_{\#}^{\otimes 5}\lambda \Bigg].$$

Example: j = 1 and j = k −1 In these cases only one partition of j satisfies $j+n(\mathbf{p}) \le k$ and formula (4.12) for the coefficient $d_{\mathbf{p}}$ becomes particularly simple:

j	Partitions of j with $j + n(\mathbf{p}) \le k$	Our notation: $\mathbf{p} =$	Coefficient $d_{\mathbf{p}} =$
1	1	1	$\frac{k(k-1)}{2}$
k-1	k-1	k-1	(k-1)!

It follows that the polynomials describing the first-order respectively order-(k − 1) contribution to $F_{N,k}(\lambda)$ are

$$P_1^{(k)}(\lambda) = \frac{k(k-1)}{2} \, \mathrm{id}_{\#}^{\otimes 2}\lambda \otimes \lambda^{\otimes(k-2)}, \qquad P_{k-1}^{(k)}(\lambda) = (k-1)! \, \mathrm{id}_{\#}^{\otimes k}\lambda.$$

We now discuss the error when truncating the finite series in (4.10). Retaining only the mean-field term gives

$$F_{N,k}(\lambda) = \frac{N^{k-1}}{\prod_{j=1}^{k-1}(N-j)} \left(\lambda^{\otimes k} + \varepsilon_{N,k}(\lambda) \right) \text{ with } \|\varepsilon_{N,k}(\lambda)\|_{\mathrm{TV}} \le \frac{C_k}{N} \qquad (4.14)$$

and keeping the first p correction terms ($p \in \{1, \dots, k-2\}$) we have

$$F_{N,k}(\lambda) = \frac{N^{k-1}}{\prod_{j=1}^{k-1}(N-j)} \left[\lambda^{\otimes k} + \sum_{j=1}^{p} \frac{(-1)^j}{N^j} S_k \, P_j^{(k)}(\lambda) + \varepsilon_{N,k,p}(\lambda) \right]$$
$$\text{with } \|\varepsilon_{N,k,p}(\lambda)\|_{\mathrm{TV}} \leq \frac{C_k}{N^{p+1}}, \tag{4.15}$$

with constants C_k independent of N and p. For example, to give explicit values,

$$C_k = \sum_{j=1}^{k-1} c_j^{(k)} \tag{4.16}$$

will do. Moreover the coefficients $d_{\mathbf{p}}^{(k)}$ and hence the C_k are independent of the size ℓ of the finite state space. Thus, for k fixed and any $N \geq 2$, retaining only the first p correlation terms captures the extreme points up to an error which decreases *exponentially* in p, the rate being *uniform* in the size of the finite state space and *improving* logarithmically with N.

Before proving Theorem 4.5 let us give a quick heuristic derivation of the formulae for the coefficients $c_j^{(k)} = \int dP_j^{(k)}(\lambda)$ which give the total mass of the jth order correction to independence for extremal N-representable k-plans. Expressions (4.7) and (4.9) suggest to try the ansatz

$$\mu_k = \frac{N^{k-1}}{\prod_{j=1}^{k-1}(N-j)} \left[\lambda^{\otimes k} + \sum_{j=1}^{k-1} \frac{(-1)^j}{N^j} c_j^{(k)} \nu_j \right] \tag{4.17}$$

with normalized measures ν_j (i.e. $\int d\nu_j = 1$) and a priori unknown but N-independent coefficients c_j. Consider for example $k = 4$. Integrating over X^k, using that μ_k and the ν_j are normalized, and multiplying both sides by the product $\prod_{j=1}^{k-1}(N-j)$ gives

$$N^3 \left(1 - \frac{c_1}{N} + \frac{c_2}{N^2} - \frac{c_3}{N^3}\right) = (N-1)(N-2)(N-3). \tag{4.18}$$

Expanding the right hand side into powers of N gives

$$(N-1)(N-2)(N-3) = N^3 - (1+2+3)N^2 + (1\cdot 2 + 1\cdot 3 + 2\cdot 3)N - (1\cdot 2\cdot 3)$$

so equating coefficients yields $c_1 = \sum_{1\leq i\leq 3} i \,(=6)$, $c_2 = \sum_{1\leq i<j\leq 3} ij \,(=11)$, $c_3 = \sum_{1\leq i<j<k\leq 3} ijk \,(=6)$, i.e. the asserted formulae for the c_j. Extending this heuristic argument to general k is straightforward.

Of course this argument is not a proof because it rests on the (as yet unjustified) ansatz (4.17) with N-independent coefficients. This ansatz is a corollary of the more detailed result (4.10)–(4.13) to whose proof we now turn.

We begin by eliminating the high-dimensional space $\mathcal{P}_{\mathrm{sym}}(X^N)$ which appears in (4.6).

Lemma 4.7 *Any extremal N-representable k-plan μ_k given by (4.6) can be written as*

$$\mu_k = \frac{(N-k)!}{N!} \sum_{\substack{(m_1,\dots,m_k)\in\{1,\dots,N\}^k,\\ \textit{pairwise distinct}}} \delta_{x_{m_1}\dots x_{m_k}}. \tag{4.19}$$

Proof Proceeding as in [17], more specifically the proof of Lemma 2.1, we rewrite μ_k by plugging in the definition of the symmetrization operator S_N and conditioning the sum over all permutations $\sigma : \{1,\dots,N\} \to \{1,\dots,N\}$ on the values on the first k integers:

$$\begin{aligned} N!\mu_k &= M_k \sum_{\sigma\in S_N} \delta_{x_{\sigma(1)}\dots x_{\sigma(N)}} \\ &= M_k \sum_{\substack{(m_1,\dots,m_k)\in\{1,\dots,N\}^k\\ \text{pairwise distinct}}} \sum_{\substack{\sigma\in S_N\\ \sigma(1)=m_1,\dots,\sigma(k)=m_k}} \delta_{x_{m_1}\dots x_{m_k} x_{\sigma(k+1)}\dots x_{\sigma(N)}} \\ &= (N-k)! \sum_{\substack{(m_1,\dots,m_k)\in\{1,\dots,N\}^k\\ \text{pairwise distinct}}} \delta_{x_{m_1}\dots x_{m_k}}. \end{aligned}$$

□

To establish Theorem 4.5 we will proceed by induction over k. The next lemma gives a deceptively simple recursion formula for extremal N-representable k-plans. Just like (4.19), it hides the inverse power series structure (4.10) and the combinatorial complexity of the coefficients (4.12) by expanding the symmetric plan μ_k in a non-symmetric basis of delta functions, leading to many terms with identical symmetrization.

Lemma 4.8 *Let $N \geq 2$, and consider the k-plans μ_k $(k = 2,\dots,N)$ defined by (4.6) for fixed $(x_1,\dots,x_N) \in X^N$. Then for $k = 1,\dots,N-1$ and $\lambda = M_1\mu_k$,*

$$\mu_{k+1} = \frac{N}{N-k}\mu_k \otimes \lambda - \frac{1}{N-k}\sum_{j=1}^{k} R_{j\#}\mu_k \tag{4.20}$$

where $R_j : X^k \to X^{k+1}$ is given by $R_j(z_1,\dots,z_k) := (z_1,\dots,z_k,z_j)$.

Proof We observe that $(m_1,\dots,m_k,m_{k+1}) \in \{1,\dots,N\}^{k+1}$ has pairwise distinct components if and only if $(m_1,\dots,m_k)$ has pairwise distinct components and $m_{k+1} \in \{1,\dots,N\}\setminus\{m_1,\dots,m_k\}$. So by inclusion-exclusion we get, using $(N-(k+1))! = (N-k)!/(N-k)$,

$$\mu_{k+1} = \frac{(N-(k+1))!}{N!} \sum_{\substack{(m_1,\dots,m_k)\in\{1,\dots,N\}^k \\ \text{pairwise distinct}}} \sum_{m_{k+1}\in\{1,\dots,N\}\setminus\{m_1,\dots,m_k\}} \delta_{x_{m_1}\dots x_{m_{k+1}}}$$

$$= \frac{(N-k)!}{N!\,(N-k)}\Big(\sum_{i=1}^{N} \sum_{\substack{(m_1,\dots,m_k)\in\{1,\dots,N\}^k \\ \text{pairwise distinct}}} \delta_{x_{m_1}\dots x_{m_k} x_i}$$

$$- \sum_{\substack{(m_1,\dots,m_k)\in\{1,\dots,N\}^k \\ \text{pairwise distinct}}} \sum_{j=1}^{k} \delta_{x_{m_1}\dots x_{m_k} x_{m_j}}\Big)$$

$$= \frac{N}{N-k}\mu_k \otimes \lambda - \frac{1}{N-k}\sum_{j=1}^{k} R_{j\#}\mu_k.$$

□

Next we derive a non-recursive formula in terms of *set partitions* of $\{1,\dots,k\}$. To state it we need to introduce some notation. Recall that a set partition of $\{1,\dots,k\}$ is a set $\mathcal{P}$ of pairwise disjoint nonempty subsets of $\{1,\dots,k\}$ (called blocks) whose union equals $\{1,\dots,k\}$. The set of all such partitions will be denoted Part_k. For a partition $\mathcal{P}\in\mathrm{Part}_k$, we denote by $n(\mathcal{P})$ the cardinality of $\mathcal{P}$, so that $\mathcal{P}=\{P_1,\dots,P_{n(\mathcal{P})}\}$ for some $P_i\subseteq\{1,\dots,k\}$, and introduce the combinatorial factor

$$\beta_{\mathcal{P}} := \prod_{i=1}^{n(\mathcal{P})} (|P_i|-1)! = \prod_{P\in\mathcal{P}} (|P|-1)!\,. \tag{4.21}$$

Next, each partition $\mathcal{P}$ induces a certain natural mapping $G_{\mathcal{P}} : \mathcal{P}(X)\to\mathcal{P}(X^k)$. Informally, this mapping pushes, for each block P of the partition $\mathcal{P}$, a tensor factor $\lambda\in\mathcal{P}(X)$ forward onto the diagonal of those cartesian factors $X_{i_1},\dots,X_{i_{|P|}}$ of the product space X^k whose indices belong to P. More precisely, if $\mathcal{P}=\{P_1,\dots,P_n\}$, define $G_{\mathcal{P}}(\lambda)$ by

$$G_{\mathcal{P}}(\lambda)(A_1\times\cdots\times A_k) = \prod_{P\in\mathcal{P}} \Big(\mathrm{id}_{\#}^{\otimes|P|}\,\lambda\Big)\Big(\prod_{i\in P} A_i\Big) \quad \text{for any } A_1,\dots,A_k\subseteq X. \tag{4.22}$$

For instance if $k=4$ and $\mathcal{P}=\{\{1,2\},\{3,4\}\}$, $G_{\mathcal{P}}(\lambda)=(\mathrm{id}_{\#}^{\otimes 2}\,\lambda)\otimes(\mathrm{id}_{\#}^{\otimes 2}\,\lambda)$, whereas if $k=5$ and $\mathcal{P}=\{\{1,3\},\{2,4,5\}\}$ then $G_{\mathcal{P}}(\lambda)$ is defined by

$$\int_{X^5}\varphi\ \mathrm{d}G_{\mathcal{P}}(\lambda) = \int_{X^2}\varphi(x,y,x,y,y)\,\mathrm{d}\lambda(x)\,\mathrm{d}\lambda(y), \quad \text{for all } \varphi\in C_b(X^5).$$

We then have the following representation formula.

Proposition 4.9 *Let μ_k be defined by (4.6), and let λ be its one-point marginal (4.8). Then*

$$\mu_k = \frac{(N-k)!}{N!} \sum_{\mathcal{P}\in \mathrm{Part}_k} (-1)^{k-n(\mathcal{P})} N^{n(\mathcal{P})} \beta_{\mathcal{P}} G_{\mathcal{P}}(\lambda). \tag{4.23}$$

Proof For $k = 1$ the assertion is obvious, and for $k = 2$ it easily follows from (4.7). Let us assume that (4.23) holds for $k \le N-1$. By (4.20)

$$\mu_{k+1} = \tfrac{(N-(k+1))!}{N!}(A+B), \quad \text{with } A := \sum_{\mathcal{P}\in \mathrm{Part}_k} (-1)^{k-n(\mathcal{P})} N^{n(\mathcal{P})+1} \beta_{\mathcal{P}} G_{\mathcal{P}}(\lambda)\otimes\lambda$$

$$\text{and } \ B := \sum_{j=1}^{k} \sum_{\mathcal{P}\in \mathrm{Part}_k} (-1)^{k+1-n(\mathcal{P})} N^{n(\mathcal{P})} \beta_{\mathcal{P}} R_{j\#} G_{\mathcal{P}}(\lambda).$$

Now let us partition Part_{k+1} into the two subsets Part^a_{k+1} and its complement Part^b_{k+1} where Part^a_{k+1} consists of all partitions $\mathcal{P}'$ of $\{1,\dots,k+1\}$ for which the singleton $\{k+1\}$ belongs to $\mathcal{P}'$. Thus $\mathcal{P}' \in \mathrm{Part}^a_{k+1}$ if and only if it can be written as $\mathcal{P}\cup\{\{k+1\}\}$ with $\mathcal{P} \in \mathrm{Part}_k$. Note then that $n(\mathcal{P}') = n(\mathcal{P}) + 1$, $\beta_{\mathcal{P}} = \beta_{\mathcal{P}'}$ and $G_{\mathcal{P}'}(\lambda) = G_{\mathcal{P}}(\lambda)\otimes\lambda$. Therefore we have

$$\begin{aligned}
&\sum_{\mathcal{P}'\in \mathrm{Part}^a_{k+1}} (-1)^{k+1-n(\mathcal{P}')} N^{n(\mathcal{P}')} \beta_{\mathcal{P}'} G_{\mathcal{P}'}(\lambda) \\
&\quad = \sum_{\mathcal{P}\in \mathrm{Part}_k} (-1)^{k-n(\mathcal{P})} N^{n(\mathcal{P})+1} \beta_{\mathcal{P}} G_{\mathcal{P}}(\lambda)\otimes\lambda \\
&\quad = A.
\end{aligned}$$

Partitions $\mathcal{P}'$ in Part^b_{k+1} are those for which $k+1$ does not form a singleton in $\mathcal{P}'$. This is the same as saying that the following map a^+_{k+1} from $\{(\mathcal{P}, P) : \mathcal{P} \in \mathrm{Part}_k,\ P \in \mathcal{P}\}$ to Part^b_{k+1} is a bijection: $a^+_{k+1}(\{P_1,\dots,P_n\}, P_i) := \mathcal{P}' := \{P_1,\dots,P_i\cup\{k+1\},\dots,P_n\}$ $(i = 1,\dots,n)$, or—in label-free notation

$$a^+_{k+1}(\mathcal{P}, P) := \{P\cup\{k+1\}\}\cup\{Q\in\mathcal{P} : Q\neq P\}.$$

We chose the notation a^+_{k+1} to emphasize the analogy with creation operators in quantum theory: the map a^+_{k+1} "creates" an extra entry $k+1$ in some block of the partition. Note that if $\mathcal{P}' = a^+_{k+1}(\mathcal{P}, P)$, then $n(\mathcal{P}) = n(\mathcal{P}')$, $\beta_{\mathcal{P}'} = \beta_{\mathcal{P}}\cdot|P|$, and $G_{\mathcal{P}'}(\lambda) = R_{j\#}G_{\mathcal{P}}(\lambda)$ for every $j \in P$, so that

$$\beta_{\mathcal{P}'} G_{\mathcal{P}'}(\lambda) = \beta_{\mathcal{P}} \sum_{j\in P} R_{j\#} G_{\mathcal{P}}(\lambda).$$

We thus have

$$\begin{aligned}
&\sum_{\mathcal{P}'\in \mathrm{Part}_{k+1}^{b}} (-1)^{k+1-n(\mathcal{P}')} N^{n(\mathcal{P}')} \beta_{\mathcal{P}'} G_{\mathcal{P}'}(\lambda) \\
&= \sum_{\mathcal{P}\in \mathrm{Part}_k} \sum_{P\in\mathcal{P}} (-1)^{k+1-n(\mathcal{P})} N^{n(\mathcal{P})} \beta_{\mathcal{P}} \sum_{j\in P} R_{j\#} G_{\mathcal{P}}(\lambda) \\
&= \sum_{j=1}^{k} \sum_{\mathcal{P}\in \mathrm{Part}_k} (-1)^{k+1-n(\mathcal{P})} N^{n(\mathcal{P})} \beta_{\mathcal{P}} \Big(\sum_{P\in\mathcal{P}\,:\, j\in P} 1 \Big) R_{j\#} G_{\mathcal{P}}(\lambda) \\
&= \sum_{j=1}^{k} \sum_{\mathcal{P}\in \mathrm{Part}_k} (-1)^{k+1-n(\mathcal{P})} N^{n(\mathcal{P})} \beta_{\mathcal{P}} R_{j\#} G_{\mathcal{P}}(\lambda) \\
&= B,
\end{aligned}$$

which gives the desired expression for $\mu_{k+1} = \frac{(N-(k+1))!}{N!}(A+B)$. □

It remains to match the unwieldy-to-evaluate expression (4.23) with the more explicit expansion (4.10), (4.11), (4.12).

We begin by dealing with the fact that expression (4.23) contains many terms with identical symmetrization. Since μ_k is symmetric, applying the symmetrization operator S_k to both sides gives

$$\mu_k = \sum_{\mathcal{P}\in \mathrm{Part}_k} (-1)^{k-n(\mathcal{P})} \frac{N^{n(\mathcal{P})}}{N\cdot(N-1)\cdot\ldots\cdot(N-k+1)} \beta_{\mathcal{P}}\, S_k\, G_{\mathcal{P}}(\lambda). \tag{4.24}$$

If $\mathcal{P}\in \mathrm{Part}_k$, then according to (4.22) $G_{\mathcal{P}}(\lambda)$ pushes $n(\mathcal{P})$ factors λ onto the k cartesian factors of the product space X^k. The different $S_k G_{\mathcal{P}}(\lambda)$'s which can arise from such a set partition are in bijective correspondence to the partitions $\mathbf{p}'$ of the number k, via

$$\mathbf{p}' = (p'_1, \ldots, p'_m) \longmapsto S_k\, \mathrm{id}_{\#}^{\otimes p'_1}\lambda \otimes \cdots \otimes \mathrm{id}_{\#}^{\otimes p'_m}\lambda. \tag{4.25}$$

The set partitions $\mathcal{P}\in \mathrm{Part}_k$ satisfying $S_k G_{\mathcal{P}}(\lambda) =$ r.h.s. of (4.25) for some given partition $\mathbf{p}' = (p'_1, \ldots, p'_m)$ of k are precisely those consisting of m sets $P_1, \ldots, P_m$ with cardinalities $|P_i| = p'_i$ for all i. Let us denote their totality by $\mathrm{Part}_k(\mathbf{p}')$. A canonical set partition in $\mathrm{Part}_k(\mathbf{p}')$ is

$$\begin{aligned}
&\mathcal{P}(\mathbf{p}') = \Big\{ P_1, \ldots, P_m \Big\} \ \text{ with} \\
&P_1 = \{1, \ldots, p'_1\},\ P_2 = \{p'_1 + 1, \ldots, p'_1 + p'_2\}, \ldots, \\
&P_m = \{p'_1 + \cdots + p'_{m-1}, \ldots, \underbrace{p'_1 + \cdots + p'_m}_{=k}\}.
\end{aligned}$$

For this set partition, as well as any other $\mathcal{P} \in \mathrm{Part}_k(\mathbf{p}')$,

$$n(\mathcal{P}) = m, \quad \beta_{\mathcal{P}} = (p'_1 - 1)! \cdot \ldots \cdot (p'_m - 1)!. \tag{4.26}$$

To reduce (4.23) to a sum over partitions $\mathbf{p}'$ of the number k, it remains to determine the number $\mathcal{N}(\mathbf{p}')$ of set partitions belonging to $\mathrm{Part}_k(\mathbf{p}')$. Let us fix any partition $\mathbf{p}' = (p'_1, \ldots, p'_m)$ of k. First of all we note that

$$\mathcal{N}(\mathbf{p}') = \frac{1}{\prod\limits_{q \in \mathrm{Ran}\,\mathbf{p}'} (|\mathbf{p}'^{-1}(q)|)!}\,\mathcal{N}' \tag{4.27}$$

where $\mathcal{N}'$ is the number of set partitions corresponding to $\mathbf{p}'$ endowed with an ordering of the blocks such that larger blocks come before smaller ones, $\mathcal{N}' = |\{(P'_1, \ldots, P'_m) : |P'_i| = p'_i \text{ for all } i\}|$. Here the combinatorial factor in the denominator accounts for the fact that any group of b equal-sized blocks in a set partition admits $b!$ orderings. But the number $\mathcal{N}'$ is straightforward to compute: choosing P'_1 means choosing p'_1 numbers out of k, so there are $\binom{k}{p'_1}$ choices; given P'_1, choosing P'_2 means choosing p'_2 numbers out of the remaining $k - p'_1$ numbers, yielding $\binom{k-p'_1}{p'_2}$ choices; and so on. It follows that

$$\mathcal{N}' = \binom{k}{p'_1} \cdot \binom{k - p'_1}{p'_2} \cdot \ldots \cdot \binom{k - (p'_1 + \cdots + p'_{m-1})}{p'_m} = \frac{k!}{p'_1! \cdot \ldots \cdot p'_m!}. \tag{4.28}$$

Combining (4.24), the bijectivity of the map (4.25), (4.26), (4.27), and (4.28) yields:

Proposition 4.10 *Let μ_k be any extremal N-representable k-plan (see (4.6)), and let λ be its one-point marginal (4.8). Then*

$$\begin{aligned}
\mu_k &= \sum_{p' \text{ partition of } k} c_{\mathbf{p}'} S_k \,\mathrm{id}_{\#}^{\otimes p'_1}\lambda \otimes \ldots \otimes \mathrm{id}_{\#}^{\otimes p'_{n(\mathbf{p}')}}\lambda \text{ with} \\
c_{\mathbf{p}'} &= (-1)^{k-n(\mathbf{p}')} \frac{k!}{N(N-1)\ldots(N-k+1)} N^{n(\mathbf{p}')} \frac{1}{p'_1 \cdot \ldots \cdot p'_{n(\mathbf{p}')}} \\
&\quad \cdot \frac{1}{\prod\limits_{q \in Ran\,\mathbf{p}'} (|\mathbf{p}'^{-1}(q)|)!}.
\end{aligned} \tag{4.29}$$

The expression for μ_k in (4.29) may be taken as an alternative definition of the polynomial $F_{N,k}(\lambda)$ introduced in Theorem 4.5.

4.3 End of the proof of Theorem 4.5

The last step in the proof of Theorem 4.5 is to match the above expression with the series given in the theorem. We would like to decompose (4.29) into terms of order

$\frac{1}{N^j}$ (times the overall order 1 prefactor in (4.10)). To this end, we re-write the N-dependent factors in (4.29) as

$$\frac{1}{N\cdot(N-1)\cdot\ldots\cdot(N-k+1)}N^{n(\mathbf{p}')}=\frac{N^{k-1}}{(N-1)\cdot\ldots\cdot(N-k+1)}\,\frac{1}{N^{k-n(\mathbf{p}')}}.$$

This together with the fact that the only partition $\mathbf{p}'$ of k with $k-n(\mathbf{p}')=0$ is $\mathbf{p}'=(1,\ldots,1)$, in which case the r.h.s. of (4.25) is $\lambda^{\otimes k}$ and $\frac{1}{p_1'\cdot\ldots\cdot p_m'}=1$, shows that

$$\mu_k=\frac{N^{k-1}}{(N-1)\ldots(N-k+1)}\Big(\lambda^{\otimes k}+\sum_{j=1}^{k-1}\frac{(-1)^j}{N^j}\sum_{\substack{\mathbf{p}'\text{ partition of }k\\ \text{with }n(\mathbf{p}')=k-j}}\frac{k!}{p_1'\cdot\ldots\cdot p_{n(\mathbf{p}')}'}$$
$$\cdot\frac{1}{\prod\limits_{q\in\operatorname{Ran}\mathbf{p}'}(|\mathbf{p}'^{-1}(q)|)!}S_k\,\mathrm{id}_{\#}^{\otimes p_1'}\lambda\otimes\ldots\otimes\mathrm{id}_{\#}^{\otimes p_{n(\mathbf{p}')}'}\lambda\Big).\tag{4.30}$$

The partitions $\mathbf{p}'$ of k of length $n(\mathbf{p}')=k-j$ are in bijective correspondence to the partitions $\mathbf{p}$ of the number j with length $n(\mathbf{p})\le k-j$, via

$$\begin{aligned}&\mathbf{p}=(p_1,\ldots,p_n)\text{ partition of }j\text{ of length }n\le k-j\\ &\longmapsto\mathbf{p}'=(p_1',\ldots,p_{k-j}')=\begin{cases}(p_1+1,\ldots,p_n+1) & \text{if } n=k-j\\ (p_1+1,\ldots,p_n+1,\underbrace{1,\ldots,1}_{k-j-n\text{ times}}) & \text{if } n<k-j.\end{cases}\end{aligned}\tag{4.31}$$

Moreover, for any two partitions $\mathbf{p}$, $\mathbf{p}'$ related by (4.31), we have

$$S_k\,\mathrm{id}_{\#}^{\otimes p_1'}\lambda\otimes\cdots\otimes\mathrm{id}_{\#}^{\otimes p_m'}\lambda=S_k\,\mathrm{id}_{\#}^{\otimes p_1+1}\lambda\otimes\cdots\otimes\mathrm{id}_{\#}^{\otimes p_n+1}\lambda\otimes\lambda^{\otimes k-(j+n)}\tag{4.32}$$

and

$$p_1'\cdot\ldots\cdot p_{n(\mathbf{p}')}'=(p_1+1)\cdot\ldots\cdot(p_n+1),\tag{4.33}$$

$$\prod_{q\in\operatorname{Ran}\mathbf{p}'}(|\mathbf{p}'^{-1}(q)|)!=\prod_{q\in\operatorname{Ran}\mathbf{p}}(|\mathbf{p}^{-1}(q)|)!\cdot(k-(j+n))!,\tag{4.34}$$

with the last factor above accounting for the $k-(j+n)$ components with value 1 occurring in $\mathbf{p}'$. Combining eq. (4.30), the bijectivity of (4.31), and identities (4.32)–(4.34) yields the desired expression for μ_k in (4.10)-(4.11)-(4.12). Let us finally prove formula (4.13). Observe that the detailed representation (4.10) implies (4.17), where each ν_j is a probability measure and the coefficients $c_j^{(k)}$ do not depend on N. Taking

the total mass of each side of (4.17) we get that for every $N \geq k$, one has

$$\prod_{j=1}^{k-1}(N-j) = N^{k-1} + \sum_{j=1}^{k-1}(-1)^j c_j^{(k)} N^{k-j-1}$$

so that $(-1)^j c_j^{(k)}$ is the coefficient of X^{k-j-1} in the polynomial $\prod_{j=1}^{k-1}(X-j)$ which proves (4.13) thanks to Vieta's formulas.

5 Extreme points and integral representation of *N*-representable *k*-plans on continuous state spaces

Now we return to the case of a general state space, i.e. we just assume that X is a Polish space. Importantly, the results in this section cover the prototypical continuous state space $X = \mathbb{R}^d$.

Recall the set $\mathcal{P}_{\frac{1}{N}}(X)$ of $\frac{1}{N}$-quantized probability measures on X introduced in (3.1), (3.2).

Given $\lambda \in \mathcal{P}(X)$ we define $F_{N,k}(\lambda) \in \mathcal{P}(X^k)$ by formulae (4.10), (4.11), (4.12); we note that the expressions in these formulae make sense for general Polish spaces X and general (not necessarily $\frac{1}{N}$-quantized) $\lambda \in \mathcal{P}(X)$.

5.1 De Finetti style representation

We now state a de Finetti style representation result for N-representable k-plans.

Theorem 5.1 *Let $N \geq k \geq 2$. A measure $\mu_k \in \mathcal{P}(X^k)$ is N-representable if and only if there exists $\alpha \in \mathcal{P}(\mathcal{P}(X))$ such that $\alpha(\mathcal{P}_{\frac{1}{N}}(X)) = 1$ and*

$$\mu_k = \int_{\mathcal{P}_{\frac{1}{N}}(X)} F_{N,k}(\lambda)\, d\alpha(\lambda) \tag{5.1}$$

where $F_{N,k}$ is defined by (4.10)–(4.12). Moreover, if $k = N$, the measure α in (5.1) is unique.

The meaning of (5.1) is that for every test function $\varphi \in C_b(X^k)$,

$$\int_{X^k} \varphi\, d\mu_k = \int_{\mathcal{P}_{\frac{1}{N}}(X)} \Phi(\lambda) d\alpha(\lambda),$$

where

$$\Phi(\lambda) := \int_{X^k} \varphi(x_1, \ldots, x_k) d\Big(F_{N,k}(\lambda)\Big)(x_1, \ldots, x_k).$$

Note that $\Phi : \lambda \mapsto \Phi(\lambda)$ is a continuous function on $\mathcal{P}(X)$ endowed with the narrow topology, by the narrow-to-narrow continuity of $F_{N,k}$. Moreover the sup norm of Φ is bounded by the sup norm of φ times the TV norm of $F_{N,k}(\lambda)$, the finiteness of the latter being clear from the definition. It follows that $\Phi \in C_b(\mathcal{P}(X))$, whence $\int_{\mathcal{P}_{\frac{1}{N}}(X)} \Phi(\lambda) d\alpha(\lambda)$ is well defined for any $\alpha \in \mathcal{P}(\mathcal{P}(X))$.

Theorem 5.1 generalizes the celebrated de Finetti–Hewitt–Savage theorem from infinitely representable to finitely representable measures, or - in probabilistic language - from infinitely to finitely extendible sequences of random variables. Formally, in the limit $N \to \infty$, the domain of integration $\mathcal{P}_{\frac{1}{N}}(X)$ in (5.1) approaches all of $\mathcal{P}(X)$ and the integrand tends to the independent measure $\lambda^{\otimes k}$, recovering de Finetti (see Sect. 6 for a rigorous account). Thus in the finitely representable case, the role of the independent measures $\lambda^{\otimes k}$ in de Finetti is taken by the universally correlated measures $F_{N,k}(\lambda)$, which contain corrections of order N^{-j} for $j = 1, \dots, k$.

The additional assumption $k = N$ for uniqueness cannot be omitted, see the next section for a simple counterexample.

Proof Recall that in a Polish space, finitely supported probability measures are dense with respect to narrow convergence. We know from Theorem 4.5 that when $\lambda := N^{-1} \sum_{i=1}^N \delta_{x_i} \in \mathcal{P}_{\frac{1}{N}}(X)$ then $F_{N,k}(\lambda) = M_k S_N \delta_{x_1 \dots x_N}$ and so $F_{N,k}(\lambda) \in \mathcal{P}_{N-\mathrm{rep}}(X^k)$. By convexity of $\mathcal{P}_{N-\mathrm{rep}}(X^k)$, any measure of the form (5.1) with a finitely supported probability measure α on $\mathcal{P}_{\frac{1}{N}}(X)$ also belongs to $\mathcal{P}_{N-\mathrm{rep}}(X^k)$. Now if $\alpha \in \mathcal{P}(\mathcal{P}_{\frac{1}{N}}(X))$ is arbitrary and μ_k is given by (5.1), we approximate α by a sequence of finitely supported measures α^n. We now use that the map $\lambda \mapsto F_{N,k}(\lambda)$ is continuous under narrow convergence. This follows immediately from the definition (4.10)–(4.12) and the continuity of the maps $\lambda \mapsto \lambda^{\otimes j}$ and $\lambda \mapsto \mathrm{id}_\#^{\otimes j} \lambda$. Hence $\mu_k^n := \int_{\mathcal{P}_{\frac{1}{N}}(X)} F_{N,k}(\lambda) \mathrm{d}\alpha^n(\lambda)$ converges narrowly to μ_k which therefore belongs to $\mathcal{P}_{N-\mathrm{rep}}(X^k)$, since the latter is closed under narrow convergence (see Lemma 2.3).

Conversely, given $\gamma \in \mathcal{P}(X^N)$ and $\mu_k := M_k S_N \gamma \in \mathcal{P}_{N-\mathrm{rep}}(X^k)$, let γ^n be a sequence of finitely supported probability measures which narrowly converges to γ, so that $\mu_k^n := M_k S_N \gamma^n$ converges to μ_k. Let us write $\gamma^n := \sum_{j \in J_n} \gamma_j^n \delta_{\mathbf{x}_j^n}$ where J_n is finite and $\mathbf{x}_j^n = (x_{1,j}^n, \dots x_{N,j}^n) \in X^N$. Using Theorem 4.5 again, we know that μ_k^n can be written as

$$\mu_k^n := \int_{\mathcal{P}_{\frac{1}{N}}(X)} F_{N,k}(\lambda) \, \mathrm{d}\alpha^n(\lambda) \tag{5.2}$$

with

$$\alpha^n := \sum_{j \in J_n} \gamma_j^n \delta_{\Lambda(\mathbf{x}_j^n)} = \Lambda_\# \gamma^n \tag{5.3}$$

where the map $\Lambda : X^N \mapsto \mathcal{P}_{\frac{1}{N}}(X)$ is defined for all $\mathbf{x} := (x_1, \ldots, x_N) \in X^N$ by

$$\Lambda(\mathbf{x}) := \frac{1}{N} \sum_{i=1}^{N} \delta_{x_i}.$$

Note that Λ is Lipschitz with respect to W_1; in particular it is continuous from X^N to $\mathcal{P}(X)$ endowed with the narrow topology. Since γ^n converges narrowly to γ, α^n converges narrowly to $\alpha := \Lambda_{\#}\gamma$. Since $\mathcal{P}_{\frac{1}{N}}(X)$ is closed (see Lemma 3.1), it follows from the Portmanteau theorem (see, e.g., [3]) and the fact that $\alpha^n(\mathcal{P}_{\frac{1}{N}}(X)) = 1$ for every n that $\alpha(\mathcal{P}_{\frac{1}{N}}(X)) = 1$. (Here we use the following part of the Portmanteau theorem: if $C \subseteq \mathcal{P}(X)$ is closed and $\alpha^n \in \mathcal{P}(\mathcal{P}(X))$ converges narrowly to α, then $\alpha(C) \geq \limsup_{n\to\infty} \alpha^n(C)$.) Finally, thanks to the narrow continuity of $F_{N,k}$ we deduce the desired representation by integrating (5.2) against a test function and passing to the limit $n \to \infty$:

$$\mu_k = \int_{\mathcal{P}_{\frac{1}{N}}(X)} F_{N,k}(\lambda) \mathrm{d}\alpha(\lambda).$$

Finally, assume $k = N$ and that $\mu \in \mathcal{P}_{\mathrm{sym}}(X^N)$ can be written as

$$\mu = \int_{\mathcal{P}_{\frac{1}{N}}(X)} F_{N,N}(\lambda) \mathrm{d}\alpha(\lambda). \tag{5.4}$$

Let $\psi : \mathcal{P}(X) \to \mathbb{R}$ be bounded and continuous for the narrow topology so that $\psi \circ \Lambda \in C_b(X^N)$. Let us now observe that if $\lambda = \Lambda(\mathbf{x}) \in \mathcal{P}_{\frac{1}{N}}(X)$ for some $\mathbf{x} = (x_1, \ldots, x_N) \in X^N$, then $S_N \delta_{\mathbf{x}} = F_{N,N}(\lambda)$ and since $\Lambda(\mathbf{x}) = \Lambda(x_{\sigma(1)}, \ldots, x_{\sigma(N)}) = \lambda$, for every $\sigma \in \mathcal{S}_N$, we have

$$\int_{X^N} \psi \circ \Lambda \, \mathrm{d}F_{N,N}(\lambda) = \psi(\lambda).$$

Taking $\psi \circ \Lambda$ as a test function in (5.4) (recall that Λ and hence $\psi \circ \Lambda$ is continuous) yields

$$\int_{X^N} \psi \circ \Lambda \, \mathrm{d}\mu = \int_{\mathcal{P}_{\frac{1}{N}}(X)} \psi(\lambda) \, \mathrm{d}\alpha(\lambda)$$

i.e. $\alpha = \Lambda_{\#}\mu$, showing in particular the uniqueness of α. □

From Theorem 5.1 and the law of large numbers, we easily deduce how to recover the prior α from sampling, as emphasized in our introduction.

Corollary 5.2 *Let $Z = (Z_1, \dots, Z_N)$ be a finitely exchangeable sequence of random variables with values in X, let $\mu \in \mathcal{P}_{\mathrm{sym}}(X^N)$ be the law of Z and let $\alpha \in \mathcal{P}(\mathcal{P}(X))$ be such that $\alpha(\mathcal{P}_{\frac{1}{N}}(X)) = 1$ and*

$$\mu = \int_{\mathcal{P}_{\frac{1}{N}}(X)} F_{N,N}(\lambda)\, d\alpha(\lambda). \tag{5.5}$$

Let $(Z^{(\nu)})_{\nu\in\mathbb{N}}$ be i.i.d drawn according to μ, and consider the $\mathcal{P}(X)$-valued sequence

$$\Lambda(Z^{(\nu)}) := \frac{1}{N}\sum_{i=1}^{N} \delta_{Z_i^{(\nu)}}.$$

Then, almost surely, the empirical measure $\frac{1}{n}\sum_{\nu=1}^{n} \delta_{\Lambda(Z^{(\nu)})}$ converges narrowly to α as $n \to \infty$.

Proof We have seen in the proof of Theorem 5.1 that $\alpha = \Lambda_{\#}\mu$ so that $\Lambda(Z)$ has law α. Hence $(\Lambda(Z^{(\nu)}))_\nu$ are i.i.d $\mathcal{P}(X)$-valued drawn according to α. Since $\mathcal{P}(X)$ endowed with the topology of narrow convergence is a Polish space, the claim follows from the strong law of large numbers for empirical measures on Polish spaces. □

5.2 Independent sequences as a convex mixture of extremal exchangeable sequences

Let us compare, in a simple example, the traditional "basis" for representing exchangeable sequences, independent sequences, with the new one advocated in this paper, extremal exchangeable sequences. This example illustrates that extremal exchangeables may not be a convex mixture of independents, but that independents are always a unique convex mixture of extremal exchangeables (Theorem 5.1).

Example Let X be the finite state space consisting of the three colors red, green, and blue, that is to say

$$X = \{\mathtt{r},\ \mathtt{g},\ \mathtt{b}\},$$

and let $N = k = 3$. Consider the probability measure

$$\lambda_* = \tfrac{2}{3}\delta_{\mathtt{r}} + \tfrac{1}{3}\delta_{\mathtt{g}},$$

which corresponds to an `rrg` urn (i.e. an urn containing two red balls and one green ball). Now consider a sequence of three independent random variables with law λ_*, that is to say

$$(Z_1, Z_2, Z_3) \sim \lambda_*^{\otimes 3}.$$

This sequence corresponds to three independent draws with replacement from an `rrg` *urn.* Our goal is to find the unique representation of this joint law as a convex mixture of extreme exchangeables.

Denote the above λ_* by `rrg`, and similiarly $\delta_{\mathtt{r}}$ by `rrr` etc. By Theorem 5.1 and the fact that λ_* contains only red and green balls, the joint law $\lambda_*^{\otimes 3}$ must be representable as a convex combination of those $F_{N,3}(\lambda)$ where λ contains only red and green balls, that is, λ=`rrr`, `rrg`, `rgg`, `ggg`. To determine the convex combination we need to compute the $F_{N,3}(\lambda)$, which is a straightforward task given our explicit formula from Sect. 4.2,

$$F_{N,3}(\lambda) = \frac{N^2}{(N-1)(N-2)}\left[\lambda^{\otimes 3} - \frac{3}{N} S_3\Big((\mathrm{id}_{\#}^{\otimes 2}\,\lambda)\otimes\lambda\Big) + \frac{2}{N^2}\,\mathrm{id}_{\#}^{\otimes 3}\,\lambda\right].$$

The result is given in the following table, where we identify probability measures $\lambda = \lambda_1\delta_{\mathtt{r}} + \lambda_2\delta_{\mathtt{g}} + \lambda_3\delta_{\mathtt{b}}$ in $\mathcal{P}(X)$ with their coefficient vectors in $\mathbb{R}^3$, and—analogously—probability measures in $\mathcal{P}(X^3)$ with their coefficient tensors in $\mathbb{R}^{3\times 3\times 3}$.

	$(F_{N,3}(\lambda))_{\cdot\cdot 1}$	$(F_{N,3}(\lambda))_{\cdot\cdot 2}$	$(F_{N,3}(\lambda))_{\cdot\cdot 3}$
$\lambda =$ `rrr`	$\begin{pmatrix} 1 & & \\ & 0 & \\ & & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & & \\ & 0 & \\ & & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & & \\ & 0 & \\ & & 0 \end{pmatrix}$
$\lambda =$ `rrg`	$\begin{pmatrix} 0 & \frac{1}{3} & \\ \frac{1}{3} & 0 & \\ & & 0 \end{pmatrix}$	$\begin{pmatrix} \frac{1}{3} & & \\ & 0 & \\ & & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & & \\ & 0 & \\ & & 0 \end{pmatrix}$
$\lambda =$ `rgg`	$\begin{pmatrix} 0 & & \\ & \frac{1}{3} & \\ & & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & \frac{1}{3} & \\ \frac{1}{3} & 0 & \\ & & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & & \\ & 0 & \\ & & 0 \end{pmatrix}$
$\lambda =$ `ggg`	$\begin{pmatrix} 0 & & \\ & 0 & \\ & & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & & \\ & 1 & \\ & & 0 \end{pmatrix}$	$\begin{pmatrix} 0 & & \\ & 0 & \\ & & 0 \end{pmatrix}$

On the other hand, the joint measure we seek to represent is

$$(\lambda_*^{\otimes 3})_{\cdot\cdot 1} = \begin{pmatrix} \frac{8}{27} & \frac{4}{27} & \\ \frac{4}{27} & \frac{2}{27} & \\ & & 0 \end{pmatrix}, \quad (\lambda_*^{\otimes 3})_{\cdot\cdot 2} = \begin{pmatrix} \frac{4}{27} & \frac{2}{27} & \\ \frac{2}{27} & \frac{1}{27} & \\ & & 0 \end{pmatrix}, \quad (\lambda_*^{\otimes 3})_{\cdot\cdot 3} = \begin{pmatrix} 0 & & \\ & 0 & \\ & & 0 \end{pmatrix}.$$

From the above explicit expressions one sees that

$$\lambda_*^{\otimes 3} = \tfrac{8}{27} F_{N,3}(\mathtt{rrr}) + \tfrac{4}{9} F_{N,3}(\mathtt{rrg}) + \tfrac{2}{9} F_{N,3}(\mathtt{rgg}) + \tfrac{1}{27} F_{N,3}(\mathtt{ggg}). \tag{5.6}$$

Thus the probability measure α in $\mathcal{P}(\mathcal{P}(X))$ in (5.1) is in our case given by $\alpha = \frac{8}{27}\delta_{\mathtt{rrr}} + \frac{4}{9}\delta_{\mathtt{rrg}} + \frac{2}{9}\delta_{\mathtt{rgg}} + \frac{1}{27}\delta_{\mathtt{ggg}}$.

The representation (5.6) has the following probabilistic meaning. *One can simulate drawing three balls with replacement from an* `rrg` *urn by*

- *first picking one of the four urns* `rrr`, `rrg`, `rgg`, `ggg` *with probabilities* $\frac{8}{27}$, $\frac{4}{9}$, $\frac{2}{9}$, $\frac{1}{27}$ *(i.e. probability ratios* $8 : 12 : 6 : 1$*),*
- *then drawing three balls without replacement from the picked urn.*

Moreover thanks to the uniqueness result in Theorem 5.1, the above choice of urns and probabilities provides the unique way of simulating the given draws with replacement by draws *without* replacement.

We emphasize that the converse (simulating draws without replacement by draws *with* replacement) is not possible, due to the well known fact that finite exchangeable laws may not be representable as convex mixtures of independents. For instance, $F_{N,3}(\mathtt{rrg})$ (the joint law for three draws without replacement from an `rrg` urn), being extremal thanks to Theorem 4.5, is not a convex combination of any other joint laws, and in particular not of any independent joint laws.

Let us also provide a simple example which shows that uniqueness of the measure α in (5.1) can fail when $k < N$.

Example Let X, N, and λ_* be as above, but $k = 2$. The measure $\lambda_*^{\otimes 2}$ (corresponding to two independent draws from a `rrg` urn) cannot just be represented by the right hand side of (5.6) with $F_{N,3}$ replaced by $F_{N,2}$, but alternatively by $\frac{2}{9}F_{N,2}(\mathtt{rrr}) + \frac{2}{3}F_{N,2}(\mathtt{rrg}) + \frac{1}{9}F_{N,2}(\mathtt{ggg})$, as the reader can easily check.

5.3 Extremal *N*-representable *k*-plans

The integral representation given by (5.1) in Theorem 5.1 will allow us to identify the set of extreme points of $\mathcal{P}_{N-\mathrm{rep}}(X^k)$ as

$$\begin{aligned} E_{N,k} &:= \{F_{N,k}(\lambda) \,:\, \lambda \in \mathcal{P}_{\frac{1}{N}}(X)\} \\ &= \{M_k S_N \delta_{x_1,\dots,x_N} \,:\, (x_1,\dots,x_N) \in X^N\} \end{aligned} \tag{5.7}$$

and to show in addition that all these points are exposed. Recall that if C is a convex subset of $\mathcal{P}(X^k)$ then $\mu \in C$ is an exposed point of C if there exists $\varphi \in C_b(X^k)$ such that $\int_{X^k} \varphi \,\mathrm{d}\mu < \int_{X^k} \varphi \,\mathrm{d}\nu$ for every $\nu \in C \setminus \{\mu\}$. It is obvious that exposed points are extreme points but the converse need not be true in general. For the set of N-representable k-plans, we have the following

Theorem 5.3 *Let* $N \geq k \geq 2$.

(a) *The set of extreme points of* $\mathcal{P}_{N-\mathrm{rep}}(X^k)$ *is given by the set* $E_{N,k}$ *defined in* (5.7).
(b) *Every such extreme point is also exposed.*

Proof

(a) Let μ be an extreme point of $\mathcal{P}_{N-\mathrm{rep}}(X^k)$. Let us write μ as in (5.1) and prove that α is a Dirac mass. If this was not the case, we could find $\varphi \in C_b(X)$ and $t \in \mathbb{R}$ such that $A_1 := \{\lambda \in \mathcal{P}(X) \,:\, \int_X \varphi \,\mathrm{d}\lambda > t\}$ and $A_2 := \{\lambda \in \mathcal{P}(X) \,:\, \int_X \varphi \,\mathrm{d}\lambda \leq t\}$

satisfy $\alpha(A_1) > 0$ and $\alpha(A_2) > 0$. Then decomposing α as $\alpha(A_1)\alpha_1 + \alpha(A_2)\alpha_2$ with α_i a probability measure giving full mass to $\mathcal{P}_{\frac{1}{N}}(X) \cap A_i$, the extremality of μ would imply

$$\int_{\mathcal{P}_{\frac{1}{N}}(X)} F_{N,k}(\lambda)\, \mathrm{d}\alpha_1(\lambda) = \int_{\mathcal{P}_{\frac{1}{N}}(X)} F_{N,k}(\lambda)\, \mathrm{d}\alpha_2(\lambda).$$

But recalling that for $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$ we have $M_1 F_{N,k}(\lambda) = \lambda$, this would also give

$$\int_{\mathcal{P}_{\frac{1}{N}}(X)} \lambda\, \mathrm{d}\alpha_1(\lambda) = \int_{\mathcal{P}_{\frac{1}{N}}(X)} \lambda\, \mathrm{d}\alpha_2(\lambda),$$

and integrating φ against this identity would lead to a contradiction. Extreme points of $\mathcal{P}_{N-\mathrm{rep}}(X^k)$ therefore belong to $E_{N,k}$

The reverse inclusion follows from (b). Alternatively, it follows from Theorem 4.2, for if $\mu = F_{N,k}(\lambda)$ for some $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$, and μ is a strict convex combination of two measures $\mu^1 \neq \mu^2$ in $\mathcal{P}_{N-\mathrm{rep}}(X^k)$, then the support of the μ^i must be contained the support of μ and thus—a fortiori—in $(\mathrm{supp}\,\check{})^{\mathrm{k}}$, contradicting the extremality of $F_{N,k}(\lambda)$ among N-representable k-point measures supported on the finite state space $(\mathrm{supp}\,\lambda)^k$.

(b) We need to show that if $\lambda := \frac{1}{N}\sum_{i=1}^N \delta_{x_i} \in \mathcal{P}_{\frac{1}{N}}(X)$, then $\mu := F_{N,k}(\lambda)$ is an exposed point of $\mathcal{P}_{N-\mathrm{rep}}(X^k)$. Again thanks to the integral representation (5.1), this amounts to finding $\varphi \in C_b(X^k)$ such that $\int_{X^k} \varphi \mathrm{d}\mu < \int_{X^k} \varphi \mathrm{d}\nu$ for every $\nu \in E_{N,k} \setminus \{\mu\}$.

First, let us rewrite $\lambda := \sum_{j=1}^{\ell} \lambda_j \delta_{y_j}$ with $\lambda_j > 0$ and $y_1, \ldots, y_\ell$ pairwise distinct. Next, we define

$$\varphi_0(z_1, \ldots, z_k) := \sum_{i=1}^{k} \min_{j=1,\ldots,\ell} d(z_i, y_j),\ \forall (z_1, \ldots, z_k) \in X^k.$$

Let u be a function in $C_b(X, \mathbb{R}^\ell)$ such that $u(y_1), \ldots, u(y_\ell)$ are linearly independent. For every $\theta \in \mathcal{P}(X)$, let us now define

$$\begin{aligned} J(\theta) &:= \frac{1}{2}\int_{X\times X} |u(x) - u(y)|^2 \mathrm{d}\theta(x)\mathrm{d}\theta(y) \\ &= \int_X |u(x)|^2 \mathrm{d}\theta(x) - \Big|\int_X u(x)\mathrm{d}\theta(x)\Big|^2 \end{aligned}$$

Obviously J is a concave quadratic functional and more precisely, defining

$$f_\lambda(x) := |u(x)|^2 - 2\Big(\int_X u(y)\mathrm{d}\lambda(y)\Big)\cdot u(x),\ \forall x \in X, \tag{5.8}$$

for every $\theta \in \mathcal{P}(X)$ we have

$$J(\theta) = J(\lambda) + \int_X f_\lambda(x)\mathrm{d}(\theta-\lambda)(x) - \Big|\int_X u(x)\mathrm{d}(\theta-\lambda)(x)\Big|^2.$$

In particular, this implies that

$$J(\theta) \le J(\lambda) + \int_X f_\lambda(x)\mathrm{d}(\theta-\lambda)(x) \tag{5.9}$$

and this inequality is strict unless

$$\int_X u\mathrm{d}\theta = \int_X u\mathrm{d}\lambda = \sum_{j=1}^{l} \lambda_j u(y_j). \tag{5.10}$$

Now if $\theta \in \mathcal{P}_{\frac{1}{N}}(X)$ and $\nu := F_{N,k}(\theta)$, we know (see (4.7)) that

$$M_2\nu = \frac{N}{N-1}\theta\otimes\theta - \frac{1}{N-1}\,\mathrm{id}_\#^{\otimes 2}\,\theta,\ \ M_1\nu = \theta$$

since $(x, y) \in X^2 \mapsto |u(x)-u(y)|^2$ vanishes on the diagonal and thanks to (5.9), we thus get

$$\begin{aligned}\frac{N}{N-1}J(\theta) &= \frac{1}{2}\int_{X^k} |u(z_1)-u(z_2)|^2\mathrm{d}\nu(z_1,\dots,z_k)\\ &\le \frac{1}{2}\int_{X^k} |u(z_1)-u(z_2)|^2\mathrm{d}\mu(z_1,\dots,z_k)\\ &\quad + \frac{N}{N-1}\int_{X^k} f_\lambda(z_1)\mathrm{d}(\nu-\mu)(z_1,\dots,z_k)\end{aligned}$$

and the last inequality is strict unless (5.10) holds. Defining

$$\varphi_1(z_1,\dots,z_k) := -\frac{1}{2}|u(z_1)-u(z_2)|^2 + \frac{N}{N-1}f_\lambda(z_1),\ \ \forall (z_1,\dots,z_k)\in X^k$$

we deduce that $\int_{X^k}\varphi_1\mathrm{d}\mu \le \int_{X^k}\varphi_1\mathrm{d}\nu,\ \ \forall \nu \in E_{N,k}$. Since $\int_{X^k}\varphi_0\mathrm{d}\mu = 0$ and $\varphi_0 \ge 0$, setting $\varphi := \varphi_0 + \varphi_1$ we also have

$$\int_{X^k}\varphi\mathrm{d}\mu \le \int_{X^k}\varphi\mathrm{d}\nu,\ \ \forall \nu \in E_{N,k}. \tag{5.11}$$

It remains to show that if $\theta \in \mathcal{P}_{\frac{1}{N}}(X)$ is such that (5.11) is an equality at $\nu = F_{N,k}(\theta)$ then necessarily $\theta = \lambda$. But in this equality case we must have $\int_{X^k}\varphi_0\mathrm{d}\nu = 0$ so that $\theta = M_1\nu$ is supported by the finite set $\{y_1,\dots y_\ell\}$ (i.e. $\theta = \sum_{j=1}^{\ell}\theta_j\delta_{y_j}$), but we must also have an equality in (5.9). The latter implies that $\sum_{j=1}^{\ell}\theta_j u(y_j) =$

$\sum_{j=1}^{\ell} \lambda_j u(y_j)$, i.e. $\lambda_j = \theta_j$ for $j = 1, \ldots, \ell$, since we have chosen $u(y_1), \ldots, u(y_\ell)$ linearly independent. □

Let us point out the following consequence of Theorem 5.3 (a).

Corollary 5.4 *The set of extreme points of $\mathcal{P}_{N-\mathrm{rep}}(X^k)$ is closed under narrow convergence.*

Proof We have to show that $E_{N,k}$ is closed under narrow convergence. Take a sequence $\mu^n = F_{N,k}(\lambda^n) \in E_{N,k}$ converging narrowly to some $\mu \in \mathcal{P}(X^k)$. Since the one-point marginals $M_1\mu^n = \lambda^n$ belong to $\mathcal{P}_{\frac{1}{N}}(X)$ and the marginal map M_1 is continuous under narrow convergence, $\lambda^n \to \lambda$ for some $\lambda \in \mathcal{P}(X)$. By the closedness of $\mathcal{P}_{\frac{1}{N}}(X)$ (see Lemma 3.1) we have $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$ and the continuity of $F_{N,k}$ yields $F_{N,k}(\lambda^n) \to F_{N,k}(\lambda)$, so $\mu = F_{N,k}(\lambda) \in E_{N,k}$. □

In some applications, one is interested in the following subset of $\mathcal{P}_{N-rep}(X^k)$ consisting of measures with no mass on the diagonal:

$$\mathcal{P}^{\mathrm{offdiag}}_{N-rep}(X^k) := \{\mu \in \mathcal{P}_{N-rep}(X^k) \,:\, \mu(diag_k(X)) = 0\} \tag{5.12}$$

where

$$diag_k(X) := \{(x, \ldots, x) \in X^k \,:\, x \in X\}. \tag{5.13}$$

The motivation for considering this set comes from singular k-body interactions. These interactions lead to costs of the form

$$c(x_1, \ldots, x_N) = \sum_{1 \le i_1 < \cdots < i_k \le N} \Phi(x_{i_1}, \ldots, x_{i_k})$$

for some measurable nonnegative Φ with $\Phi = +\infty$ on $diag_k(X)$, then the total cost $C[\gamma] = \int_{X^N} c\, d\gamma$ is infinite if $M_k\gamma \notin \mathcal{P}^{\mathrm{offdiag}}_{N-rep}(X^k)$. A prototypical example is the Coulomb cost, see Sect. 8, which led Khoo and Ying [29] to introduce the set (5.12)–(5.13) in the case $k = 2$ and X a finite state space. Recall the quantization constraint for measures $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$ that $\lambda(\{x\}) \in \{0, 1/N, \ldots, 1\}$.

Theorem 5.5

(a) *The set of extreme points of $\mathcal{P}^{\mathrm{offdiag}}_{N-rep}(X^k)$ is given by*

$$E^{\mathrm{offdiag}}_{N,k} := \left\{ F_{N,k}(\lambda) \,:\, \lambda \in \mathcal{P}_{\frac{1}{N}}(X) \,:\, \lambda(\{x\}) \in \{0, \tfrac{1}{N}, \ldots, \tfrac{k-1}{N}\} \right\}. \tag{5.14}$$

(b) *Every such extreme point is also exposed.*

Geometrically this means that intersecting the set of N-representable k-plans with the subspace of measures which vanish on the diagonal does not generate new extreme points.

Theorem 5.5 generalizes a recent result of [29] from $k = 2$ and finite state spaces X to general k and X.

Proof First we show that any extreme point μ is contained in the set (5.14). By the representation (5.1), $\mu = \int_A F_{N,k}(\lambda)\, d\alpha(\lambda)$ where $A = \{\lambda \in \mathcal{P}_{\frac{1}{N}}(X) : F_{N,k}(\lambda)(diag_k(X)) = 0\}$. The key point is that the polynomial $F_{N,k}(\lambda)$ factorizes on the diagonal:

$$\begin{aligned} F_{N,k}(\lambda)(\{(x,\dots,x)\}) &= \frac{N^{k-1}}{\prod_{j=1}^{k-1}(N-j)}\Big[\lambda(\{x\})^k + \sum_{j=1}^{k-1}\frac{(-1)^j}{N^j}\lambda(\{x\})^{k-j}c_j^{(k)}\Big] \\ &= \frac{N^{k-1}}{\prod_{j=1}^{k-1}(N-j)}\lambda(\{x\})\Big(\lambda(\{x\})-\tfrac{1}{N}\Big)\cdots\Big(\lambda(\{x\})-\tfrac{k-1}{N}\Big). \end{aligned} \tag{5.15}$$

The second equality follows because the coefficients $c_j^{(k)}$ are N-independent and satisfy

$$N^k\Big(1-\frac{c_1^{(k)}}{N}+\frac{c_2^{(k)}}{N^2}-+\cdots+(-1)^{k-1}\frac{c_{k-1}^{(k)}}{N^{k-1}}\Big)=\prod_{j=0}^{k-1}(N-j)$$

for all nonnegative integers N, see Eq. (4.18). Since λ is $\frac{1}{N}$-quantized, $\lambda(x)N$ is a nonnegative integer and can be substituted for N, yielding the asserted expression. The factorization identity (5.15) shows that $A = \{\lambda \in \mathcal{P}_{\frac{1}{N}}(X) : \lambda(x) \in \{0, \frac{1}{N}, \dots, \frac{k-1}{N}\}\}$. The rest of the proof of (a) is analogous to that of Theorem 5.3 (a) and the exposedness of the elements in $E_{N,k}^{\text{offdiag}}$ is immediate from Theorem 5.3 (b). □

5.4 Exponentially convergent approximation of extreme points

The error estimates (4.14), (4.15) when approximating $F_{N,k}(\lambda)$ by a truncated series can be extended in a straightforward manner to Polish spaces.

Theorem 5.6 *For any extreme point μ_k of $\mathcal{P}_{N-\text{rep}}(X^k)$, denoting its one-point marginal $M_1\mu_k \in \mathcal{P}_{\frac{1}{N}}(X)$ by λ we have*

$$\mu_k = \frac{N^{k-1}}{\prod_{j=1}^{k-1}(N-j)}\Big[\lambda^{\otimes k} + \varepsilon_{N,k}(\lambda)\Big] \text{ with } \|\varepsilon_{N,k}(\lambda)\|_{\text{TV}} \le \frac{C_k}{N} \tag{5.16}$$

and

$$\begin{aligned} \mu_k = \frac{N^{k-1}}{\prod_{j=1}^{k-1}(N-j)}&\left[\lambda^{\otimes k} + \sum_{j=1}^{p}\frac{(-1)^j}{N^j}S_k\, P_j^{(k)}(\lambda) + \varepsilon_{N,k,p}(\lambda)\right] \\ &\text{with } \|\varepsilon_{N,k,p}(\lambda)\|_{\text{TV}} \le \frac{C_k}{N^{p+1}}, \end{aligned} \tag{5.17}$$

with constants C_k independent of N and p. The explicit constants (4.16) will do. Conversely, for every $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$ there exists an extreme point μ_k of $\mathcal{P}_{N-\mathrm{rep}}(X^k)$ such that (5.16), (5.17) hold.

Proof By Theorem 5.3, an extreme point μ_k equals $F_{N,k}(\lambda)$, and in particular is supported on the finite set $(\operatorname{supp}\lambda)^k$. Estimates (5.16), (5.17) now follow from (4.14)–(4.15) and the fact that

$$||\varepsilon_{N,k}(\lambda)||_{TV(X^k)} = ||\varepsilon_{N,k}(\lambda)||_{TV((\operatorname{supp}\lambda)^k)}$$

(and the analogous identity for $\varepsilon_{N,k,p}$). □

Of course, via the integral representation (5.1) this translates into an analogous exponentially convergent approximation result for arbitrary (non-extremal) N-representable measures, see the next section.

6 Recovering the de Finetti–Hewitt–Savage theorem

The celebrated Hewitt–Savage theorem [22] (going back to de Finetti in the special case of the state space $X = \{0, 1\}$) says the following.

Theorem 6.1 (Hewitt–Savage [22]) *If $(Z_n)_n$ is an infinite exchangeable sequence of random variables with values in the Polish space X, then there exists $\alpha \in \mathcal{P}(\mathcal{P}(X))$ such that for every k, the law μ_k of $(Z_1, \ldots, Z_k)$ is given by*

$$\mu_k = \int_{\mathcal{P}(X)} \lambda^{\otimes k} \, d\alpha(\lambda). \tag{6.1}$$

Let us show how this theorem follows from Theorem 5.1 which therefore may be viewed as a non-asymptotic form of the Hewitt–Savage theorem. If μ_k is the law of the first k components of an infinite exchangeable sequence, it is N-representable for every N, so that there exists $\alpha^{N,k} \in \mathcal{P}(\mathcal{P}_{\frac{1}{N}}(X))$ such that

$$\mu_k = \int_{\mathcal{P}_{\frac{1}{N}}(X)} F_{N,k}(\lambda) \mathrm{d}\alpha^{N,k}(\lambda).$$

In particular, for the law μ_1 of Z_1 we have

$$\mu_1 := \int_{\mathcal{P}(X)} \lambda \, \mathrm{d}\alpha^{N,k}(\lambda).$$

Since X is Polish, μ_1 is tight so there exists $\Psi : X \to [0, +\infty]$ which is lower semicontinuous and coercive (i.e. with compact sublevel sets) such that

$$+\infty > \int_X \Psi(x) \mathrm{d}\mu_1(x) = \int_{\mathcal{P}(X)} \Big(\int_X \Psi(x) \mathrm{d}\lambda(x) \Big) \mathrm{d}\alpha^{N,k}(\lambda).$$

But since Ψ is coercive, $\lambda \in \mathcal{P}(X) \mapsto \int_X \Psi(x)\mathrm{d}\lambda(x)$ is coercive on $\mathcal{P}(X)$. It thus follows from Prokhorov's theorem that $\alpha^{N,k}$ is tight. Hence, taking a subsequence if necessary, $\alpha^{N,k}$ converges narrowly to some α^k as $N \to \infty$, and since $F_{N,k}(\lambda)$ converges uniformly to $\lambda^{\otimes k}$, thanks to (4.14) we get letting $N \to \infty$:

$$\mu_k = \int_{\mathcal{P}(X)} \lambda^{\otimes k}\, \mathrm{d}\alpha^k(\lambda).$$

It remains to show that α^k may be chosen independently of k. For this we simply observe that for $\ell \geq k$, $\mu_k = M_k(\mu_\ell)$. Since $M_k(\lambda^{\otimes \ell}) = \lambda^{\otimes k}$, thanks to the linearity and narrow continuity of M_k we thus get

$$\mu_k = \int_{\mathcal{P}(X)} \lambda^{\otimes k}\, \mathrm{d}\alpha^\ell(\lambda), \quad \text{for all } \ell \geq k.$$

We then use the same tightness argument as above to see that for some sequence $\ell_n \to \infty$, α^{ℓ_n} converges narrowly to some α as $n \to +\infty$ and the Hewitt–Savage representation (6.1) holds for this measure α.

Diaconis and Freedman (in [12]) obtained TV bounds between elements of $\mathcal{P}_{N-\mathrm{rep}}(X^k)$ and mixtures of independent measures $\lambda^{\otimes k}$. More precisely, Diaconis and Freedman showed that if $\mu_k \in \mathcal{P}_{N-\mathrm{rep}}(X^k)$ and $k \leq N$, there exists $\alpha \in \mathcal{P}(\mathcal{P}(X))$ such that, for some constant B_k (for results on their k-dependence see [12])

$$\|\mu_k - \int_{\mathcal{P}(X)} \lambda^{\otimes k}\mathrm{d}\alpha\|_{\mathrm{TV}} \leq \frac{B_k}{N}. \tag{6.2}$$

These $O(\frac{1}{N})$ bounds are recovered by using the integral representation of $\mu_k \in \mathcal{P}_{N-\mathrm{rep}}(X^k)$ of Theorem 5.1 and applying (5.16).

A better approximation of μ_k is obtained by adding the universally correlated correction terms from (5.17) to the independent measures $\lambda^{\otimes k}$ in (6.2) and applying (5.17). This yields:

Corollary 6.2 *For any measure μ_k in $\mathcal{P}_{N-\mathrm{rep}}(X^k)$, and with α as in Theorem 5.1, we have for $p = 1, \ldots, k-2$ and some constant B_k*

$$\|\mu_k - \int_{\mathcal{P}(X)} \frac{N^{k-1}}{\prod_{j=1}^{k-1}(N-j)}\Big[\lambda^{\otimes k} + \sum_{j=1}^{p} \frac{(-1)^j}{N^j} S_k P_j^{(k)}(\lambda)\Big]\, d\alpha\|_{\mathrm{TV}} \leq \frac{B_k}{N^{p+1}}\,. \tag{6.3}$$

7 Connection with the Ewens sampling formula

When looking at the coefficients arising in the classification (4.10)–(4.12) of extremal N-representable measures, we noticed some resemblance to those in the celebrated *Ewens distribution* or *Ewens sampling formula* from genetics, which—mathematically speaking—is a probability distribution on integer partitions. On the other hand, the sign

factor $(-1)^j$ in (4.10), or $(-1)^{k-n(\mathbf{p}')}$ in the alternative expression (4.29), means that some coefficients are negative, so the coefficients cannot be a probability distribution on integer partitions. *What is happening here?*

There is a precise connection which seems to us quite remarkable and which we now describe.

The Ewens distribution is usually stated in terms of the following parametrization of partitions. If $\mathbf{p}' = (p'_1, \ldots, p'_m)$ is a partition of the integer k, that is to say $p'_1 + \cdots + p'_m = k$ and $p'_1 \geq \ldots \geq p'_m \geq 1$, one denotes by m_q the number of components equal to q (i.e., the number of times the term q appears in the sum), $m_q := |\mathbf{p}'^{-1}(q)|$, and introduces the vector $(m_1, \ldots, m_k)$ (called allelic partition). For instance, the partition $\mathbf{p}' = (2, 1, 1, 1)$ of 5 corresponds to the allelic partition $(3, 1, 0, 0, 0)$, because 1 appears 3 times, 2 appears once, and 3, 4, 5 appear zero times. The Ewens distribution on the allelic partitions of the integer $k \in \mathbb{N}$ is now given by

$$\mathbb{P}(m_1, \ldots, m_k; \theta) = \frac{k!}{\theta(\theta+1)\ldots(\theta+k-1)} \prod_{q=1}^{k} \frac{\theta^{m_q}}{q^{m_q} m_q!}, \tag{7.1}$$

where $\theta > 0$ is a parameter. In genetics, this distribution models the genetic diversity of a population (at statistical equilibrium and under neutral selection), and θ is the *mutation rate*. More precisely, for a random sample of k genes taken from a population at a particular locus, Eq. (7.1) gives the probability that m_1 alleles (variant forms of the gene) appear exactly once, m_2 alleles appear exactly twice, and so on. The Ewens distribution has found widespread use in biology, statistics, probability, and even number theory [14], see [15, 27] for original papers and [10] for a recent review.

In terms of the standard representation of the partitions of k, expression (7.1) becomes

$$\frac{k!}{\theta(\theta+1)\ldots(\theta+k-1)} \theta^{n(\mathbf{p}')} \prod_{i=1}^{n(\mathbf{p}')} \frac{1}{p'_i} \frac{1}{\prod_{q \in \operatorname{Ran} \mathbf{p}'} (|\mathbf{p}'^{-1}(q)|)!} =: \mathtt{ew}(\mathbf{p}'; \theta). \tag{7.2}$$

Definition 7.1 Given any partition $\mathbf{p}'$ of the integer k, we define the complex Ewens function $\theta \mapsto \mathtt{ew}(\mathbf{p}'; \theta)$ by (7.2), where $\theta \in \mathbb{C}$ is a complex parameter and the arithmetic operations are the usual ones in $\mathbb{C}$.

For a picture when $k = 3$ see Fig. 2.

The complex Ewens function has the following properties:

1. $\mathtt{ew}(\mathbf{p}'; \cdot)$ has poles exactly at the negative integers $\theta = -1, \ldots, -(k-1)$, and is holomorphic on $\mathbb{C} \backslash \{-1, \ldots, -(k-1)\}$. In particular, $\mathtt{ew}(\mathbf{p}'; \cdot)$ is meromorphic on $\mathbb{C}$.
2. $\mathtt{ew}(\mathbf{p}'; \cdot)$ is nonzero for the trivial partition $\mathbf{p}' = k$, and zero only at $\theta = 0$ for all other partitions of k.
3. $\sum_{\mathbf{p}' \text{ partition of } k} \mathtt{ew}(\mathbf{p}'; \cdot)$ is identically equal to 1.

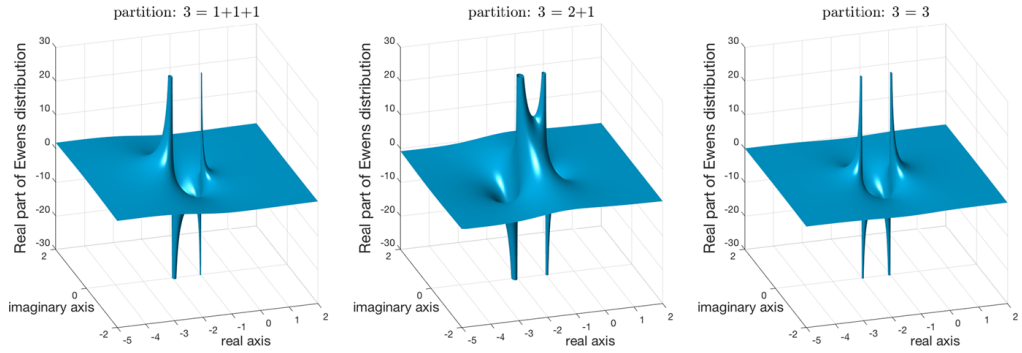

Fig. 2 The complex Ewens distribution $(\mathbf{p}', \theta) \mapsto \mathrm{ew}(\mathbf{p}'; \theta)$ for $k = 3$. Since there are three partitions $\mathbf{p}'$ of 3, it is a collection of three meromorphic functions of the complex variable θ. The poles of each function are located at $\theta = -1$ and $\theta = -2$. Restricting θ to the positive real axis yields the classical Ewens distribution, which is a probability measure for each θ, and has long been known to be relevant in genetics. Its analytic continuation to the negative real axis, which to our knowledge has not been introduced before, is a signed measure for each θ, and turns out to describe extremal N-representable measures

4. $\{\mathrm{ew}(\mathbf{p}'; \theta)\}_{\mathbf{p}'}$ partition of k is a probability measure (i.e., in addition real and nonnegative) when θ belongs to the nonnegative real axis, and a signed measure (i.e., in addition real but not nonnegative) when θ belongs to the negative real axis excluding the poles.

The first two properties are immediate from the definition. The third follows from the fact that the sum is identically equal to 1 for positive real θ and the identity theorem from complex analysis. Property 4 is an obvious consequence of Property 3 and the definition.

The connection to extremal N-representable measures is the following.

Theorem 7.2 *The coefficients $c_{\mathbf{p}'}$ in the expansion (4.29) of extremal N-representable k-point measures are given by $\mathrm{ew}(\mathbf{p}'; \theta)|_{\theta=-N}$; that is, they are given by the Ewens distribution evaluated at a probabilistically meaningless but by analytic continuation admissible negative parameter value.*

Properties 3 and 4 of the complex Ewens distribution thus recover the phenomenon discussed at the beginning of the section that the coefficients $c_{\mathbf{p}'}$ in (4.29) are only a signed measure of total mass 1 but not a probability measure. (As already noted below Theorem 4.5, due to this phenomenon it is highly nontrivial that the induced measures (4.23), $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$, are nevertheless probability measures.)

Some sort of correspondence between N-representable k-point measures and the Ewens distribution at $\theta = -N$ also exists at the poles. By Property 1, the complex Ewens distribution stops making mathematical sense at the points $-N$ when $N \in \mathbb{N}$, $N < k$; but these are precisely the N's for which the notion of N-representability (see Definition (2.1)) stops making sense as a k-plan cannot be N-representable when $N < k$. Thus the two poles seen in Fig. 2 at -1 and -2 reflect the fact that it makes no sense for a 3-point probability measure to be 1- or 2-representable.

8 Applications to optimal transport

Our interest in the structure of the set $\mathcal{P}_{N-\mathrm{rep}}(X^k)$ is motivated by symmetric multi-marginal optimal transport problems (MMOT). We briefly recall below the general form of MMOT problems and then consider a special class which leads to linear minimization problems over $\mathcal{P}_{N-\mathrm{rep}}(X^k)$.

MMOT Given a cost function $c \in C_b(X^N)$ and $(\rho_1, \dots, \rho_N) \in \mathcal{P}(X)^N$, the multi-marginal optimal transport with cost c and marginals $(\rho_1, \dots, \rho_N)$ consists in finding a transport plan between $(\rho_1, \dots, \rho_N)$, i.e. a joint probability $\gamma \in \mathcal{P}(X^N)$ having marginals $(\rho_1, \dots, \rho_N)$ and minimizing the total cost

$$\int_{X^N} c(x_1, \dots, x_N) \mathrm{d}\gamma((x_1, \dots, x_N).$$

This generalization of the more standard two-marginals Monge-Kantorovich optimal transport problem arises in different settings such as polar factorization of vector fields [20], Wasserstein barycenters [1, 19] and many-electron physics [6–8], just to name a few. Note that these problems in general dramatically suffer from the curse of dimensionality.

Symmetric MMOT with k-body interactions We now consider the special class of MMOT problems for which: (i) all marginals are equal, (ii) the cost c is symmetric (so that one can look for symmetric optimal plans) and (iii) the cost involves only k-body interactions with $2 \le k \le N$. More precisely, given integers $2 \le k \le N$ and $\Phi \in C_b(X^k)$ *symmetric* (i.e., $\Phi(x_{\sigma(1)}, \dots, x_{\sigma(k)}) = \Phi(x_1, \dots, x_k)$ for every $(x_1, \dots, x_k) \in X^k$ and every permutation $\sigma \in \mathcal{S}_k$), we consider the symmetric cost c_Φ defined on X^N by

$$c_\Phi(x_1, \cdots, x_N) := \frac{1}{\binom{N}{k}} \sum_{1 \le i_1 < i_2 \dots < i_k \le N} \Phi(x_{i_1}, \dots, x_{i_k}). \tag{8.1}$$

In this setting, N is the total number of particles in the system, X^N is the space of all N-particle configurations, and Φ is a k-body interaction potential. In practice N is much larger than k. Given $\rho \in \mathcal{P}(X)$ we are interested in the multi-marginal optimal transport problem

$$\begin{aligned} &C_{N,k}(\rho) \\ &:= \inf \Big\{ \int_{X^N} c_\Phi(x_1, \dots, x_N) \mathrm{d}\gamma(x_1, \dots, x_N) \ : \ \gamma \in \mathcal{P}_{\mathrm{sym}}(X^N), \ M_1\gamma = \rho \Big\}. \end{aligned} \tag{8.2}$$

An important example which has received much recent interest is $k = 2$, $X = \mathbb{R}^3$, $\Phi(x_1, x_2) = |x_1 - x_2|^{-1}$ (optimal transport with Coulomb cost [6, 8]), in which case (8.2) arises as the strictly correlated limit of Hohenberg-Kohn density functional theory [8]. Thanks to (8.1), the minimization can be reformulated in terms of the k-marginal

$\mu_k = M_k(\gamma) \in \mathcal{P}_{N-\mathrm{rep}}(X^k)$:

$$C_{N,k}(\rho) := \inf \Big\{ \int_{X^k} \Phi \mathrm{d}\mu_k \ : \ \mu_k \in \mathcal{P}_{N-\mathrm{rep}}(X^k), \ M_1(\mu_k) = \rho \Big\}. \tag{8.3}$$

8.1 A $\frac{1}{N}$-quantized polynomial convexification problem

The de Finetti style representation from Theorem 5.1, together with the fact that $M_1(F_{N,k}(\lambda)) = \lambda$ for every $\lambda \in \mathcal{P}_{\frac{1}{N}}(X)$, enables us to write $C_{N,k}(\rho)$ as follows.

Theorem 8.1 (Reformulation of symmetric multi-marginal optimal transport) *The functional $C_{N,k}$ can be rewritten as*

$$C_{N,k}(\rho) = \inf \Big\{ \int_{\mathcal{P}_{\frac{1}{N}}(X)} \Big(\int_{X^k} \Phi \ dF_{N,k}(\lambda) \Big) d\alpha(\lambda) \ : \ \alpha \in \mathcal{P}(\mathcal{P}_{\frac{1}{N}}(X)) \Big\}, \tag{8.4}$$

subject to the prescribed marginal constraint

$$\int_{\mathcal{P}_{\frac{1}{N}}(X)} \lambda \ d\alpha(\lambda) = \rho. \tag{8.5}$$

In view of formulae (4.10)–(4.12), one can observe that $\lambda \in \mathcal{P}_{\frac{1}{N}}(X) \mapsto \int_{X^k} \Phi \ \mathrm{d}F_{N,k}(\lambda)$ is a polynomial of degree k expression in the weights of the discrete measure λ, for instance

$$\int_{X^2} \Phi \ \mathrm{d}F_{N,2}(\lambda) = \frac{N}{N-1} \int_{X^2} \Phi(x, y) \mathrm{d}\lambda(x) \mathrm{d}\lambda(y) - \frac{1}{N-1} \int_X \Phi(x, x) \mathrm{d}\lambda(x)$$

and

$$\begin{aligned} \int_{X^3} \Phi \ \mathrm{d}F_{N,3}(\lambda) = & \frac{N^2}{(N-1)(N-2)} \int_{X^3} \Phi(x, y, z) \mathrm{d}\lambda(x) \mathrm{d}\lambda(y) \mathrm{d}\lambda(z) \\ & - \frac{3N}{(N-1)(N-2)} \int_{X^2} \Phi(x, x, y) \mathrm{d}\lambda(x) \mathrm{d}\lambda(y) \\ & + \frac{2}{(N-1)(N-2)} \int_X \Phi(x, x, x) \mathrm{d}\lambda(x). \end{aligned}$$

Defining the polynomial $P_{N,k}$ for every single marginal (not necessarily $\frac{1}{N}$-quantized) $\lambda \in \mathcal{P}(X)$ by

$$P_{N,k}(\lambda) := \int_{X^k} \Phi \ \mathrm{d}F_{N,k}(\lambda), \quad \text{for all } \lambda \in \mathcal{P}(X), \tag{8.6}$$

we see that minimization problem (8.4)–(8.5) is a $\frac{1}{N}$-quantized constrained version of the convexification of the polynomial $P_{N,k}$. More precisely, the convexification (or

convex envelope) of $P_{N,k}$, denoted $P^{**}_{N,k}$, is the largest convex function bounding $P_{N,k}$ from below on the whole of $\mathcal{P}(X)$, it is given by the formula:

$$P^{**}_{N,k}(\rho) := \inf_{\alpha \in \mathcal{P}(\mathcal{P}(X))} \left\{ \int_{\mathcal{P}(X)} P_{N,k}(\lambda) \mathrm{d}\alpha(\lambda) \ : \ \int_{\mathcal{P}(X)} \lambda \mathrm{d}\alpha(\lambda) = \rho \right\}. \tag{8.7}$$

We thus deduce from (8.4)–(8.5), the lower bound

$$C_{N,k}(\rho) \geq P^{**}_{N,k}(\rho), \ \forall \rho \in \mathcal{P}(X). \tag{8.8}$$

To understand the advantage of the formulation (8.4)–(8.5) compared to the initial multi-marginal problem (8.2), it is worth considering in detail the finite case where $X = X_\ell$ is finite with ℓ elements. In this case both (8.2) and (8.4)–(8.5) are linear programs which have $\binom{N+\ell-1}{\ell-1}$ variables and ℓ constraints. But the special structure of (8.4)–(8.5) makes it much more appealing from a computational viewpoint. Indeed, the computation of the value function $C_{N,k}$ given by (8.4)–(8.5) amounts to finding the convex envelope of the restriction of the polynomial $P_{N,k}$ to the finite set $\mathcal{P}_{\frac{1}{N}}(X)$. In the simplest case where $\ell = 2$, this amounts to computing the convex hull of $N+1$ points in the plane located in the graph of $P_{N,k}$. This convex envelope can be computed exactly in linear in N time thanks to the Graham scan algorithm [21] for instance.

8.2 Convergence as $N \to \infty$

Now we consider the asymptotic behavior of $C_{N,k}$ as $N \to \infty$. As first emphasized in [9] in the case of the Coulomb cost (or more general pairwise interactions with a potential having a positive Fourier transform), the Hewitt–Savage theorem enables one to drastically simplify the analysis of problems like (8.2) as $N \to \infty$. Since $\rho^{\otimes k}$ is admissible in the optimal transport problem (8.3) we have

$$C_{N,k}(\rho) \leq P_k(\rho), \ \forall \rho \in \mathcal{P}(X)$$

where

$$P_k(\rho) := \int_{X^k} \Phi \mathrm{d}\rho^{\otimes k}, \ \forall \rho \in \mathcal{P}(X) \tag{8.9}$$

but since $C_{N,k}$ is obviously convex, denoting by P^{**}_k the convexification of P_k defined in (8.9), this also gives

$$C_{N,k}(\rho) \leq P^{**}_k(\rho), \quad \text{for all } \rho \in \mathcal{P}(X). \tag{8.10}$$

Recalling (5.16), we observe that for some constant C_k we have for every $\lambda \in \mathcal{P}(X)$

$$P_{N,k}(\lambda) \geq \frac{(N-k)!N^k}{N!} \Big(P_k(\lambda) - \frac{C_k \|\Phi\|_\infty}{N}. \Big)$$

Taking convex envelopes and using (8.8) we thus get, for every $\rho \in \mathcal{P}(X)$:

$$\frac{(N-k)!N^k}{N!}\left(P_k^{**}(\rho) - \frac{C_k\|\Phi\|_\infty}{N}\right) \leq C_{N,k}(\rho) \leq P_k^{**}(\rho). \tag{8.11}$$

So $C_{N,k}$ converges uniformly on $\mathcal{P}(X)$ to P_k^{**} as $N \to +\infty$. We also have a Γ-convergence result:

Theorem 8.2 *As $N \to \infty$, the sequence of functionals $C_{N,k}$ defined in (8.2) Γ-converges with respect to the narrow topology of $\mathcal{P}(X)$ to P_k^{**}.*

Proof The Γ-limsup inequality obviously follows from (8.10). As for the Γ-liminf inequality, it follows directly from (8.11) and the lower semicontinuity of P_k^{**} with respect to the narrow topology. □

To conclude, we see that the asymptotic problem obtained by letting $N \to \infty$ in (8.3) amounts to computing the convex envelope of the polynomial of degree k interaction functional P_k. This might be a challenging computational task in general but we believe that the theory developed by Lasserre for polynomial optimization (see in particular [30]) might be useful in certain situations.

Acknowledgements G.C. gratefully acknowledges the hospitality of TU Munich, where this work was begun while he held a John Von Neumann visiting professorship and acknowledges the support of the Lagrange Mathematics and Computation Research Center.

References

1. Agueh, Martial, Carlier, Guillaume: Barycenters in the Wasserstein space. SIAM J. Math. Anal. **43**(2), 904–924 (2011)
2. Aldous, D.J.: Exchangeability and related topics. In: École d'été de probabilités de Saint-Flour, XIII—1983, volume 1117 of Lecture Notes in Math., pp. 1–198. Springer, Berlin (1985)
3. Billingsley, Patrick: Convergence of Probability Measures. Wiley, Hoboken (1999)
4. Bobkov, Sergey G.: Generalized symmetric polynomials and an approximate de Finetti representation. J. Theor. Probab. **18**(2), 399–412 (2005)
5. Brenier, Yann: Polar factorization and monotone rearrangement of vector-valued functions. Commun. Pure Appl. Math. **44**(4), 375–417 (1991)
6. Buttazzo, G., De Pascale, L., Gori-Giorgi, P.: Optimal-transport formulation of electronic density-functional theory. Phys. Rev. A **85**, 062502 (2012)
7. Comtet, Louis: Advanced Combinatorics: The Art of Finite and Infinite Expansions. Springer, Dordrecht (1974)
8. Cotar, Codina, Friesecke, Gero, Klüppelberg, Claudia: Density functional theory and optimal transportation with Coulomb cost. Commun. Pure Appl. Math. **66**(4), 548–599 (2013)
9. Cotar, Codina, Friesecke, Gero, Pass, Brendan: Infinite-body optimal transport with Coulomb cost. Calc. Var. Partial Differ. Equ. **54**(1), 717–742 (2015)
10. Crane, H.: The ubiquitous Ewens sampling formula. Stat. Sci. **31**(1), 1–19 (2016)
11. Dellacherie, C., Meyer, P.-A.: Probabilities and potential. B, volume 72 of North-Holland Mathematics Studies. North-Holland Publishing Co., Amsterdam,: Theory of martingales. Translated from the French by J. P, Wilson (1982)
12. Diaconis, P., Freedman, D.: Finite exchangeable sequences. Ann. Probab. **8**(4), 745–764 (1980)
13. Diaconis, Persi: Finite forms of de Finetti's theorem on exchangeability. Synthese **36**(2), 271–281 (1977). (**Foundations of probability and statistics, II**)

14. Donnelly, P., Grimmett, G.: On the asymptotic distribution of large prime factors. J. Lond. Math. Soc. **s2–47**(3), 395–404 (1993)
15. Ewens, W.J.: The sampling theory of selectively neutral alleles. Theor. Popul. Biol. **3**(1), 87–112 (1972)
16. Friesecke, G., Mendl, C.B., Pass, B., Cotar, C., Klüppelberg, C.: N-density representability and the optimal transport limit of the Hohenberg–Kohn functional. J. Chem. Phys. **139**(16), 164109 (2013)
17. Friesecke, Gero, Vögler, Daniela: Breaking the curse of dimension in multi-marginal Kantorovich optimal transport on finite state spaces. SIAM J. Math. Anal. **50**(4), 3996–4019 (2018)
18. Gangbo, Wilfrid, McCann, Robert J.: The geometry of optimal transportation. Acta Math. **177**(2), 113–161 (1996)
19. Gangbo, Wilfrid, Święch, Andrzej: Optimal maps for the multidimensional Monge-Kantorovich problem. Commun. Pure Appl. Math. **51**(1), 23–45 (1998)
20. Ghoussoub, Nassif, Moameni, Abbas: Symmetric Monge-Kantorovich problems and polar decompositions of vector fields. Geom. Funct. Anal. **24**(4), 1129–1166 (2014)
21. Graham, Ronald: An efficient algorithm for determining the convex hull of a finite planar set. Inf. Process. Lett. **1**, 132–133 (1972)
22. Hewitt, Edwin, Savage, Leonard J.: Symmetric measures on Cartesian products. Trans. Am. Math. Soc. **80**, 470–501 (1955)
23. Hörmander, Lars: Notions of Convexity. Birkhäuser, Boston (1994)
24. Janson, Svante, Konstantopoulos, Takis, Yuan, Linglong: On a representation theorem for finitely exchangeable random vectors. J. Math. Anal. Appl. **442**(2), 703–714 (2016)
25. Jaynes, E.: Some applications and extensions of the de Finetti representation theorem. In: Bayesian Inverence and Decision Techniques: Essays in Honor of Bruno de Finetti, pp. 31–42 (1986)
26. Kallenberg, Olav: Probabilistic Symmetries and Invariance Principles. Probability and Its Applications (New York). Springer, New York (2005)
27. Karlin, S., McGregor, J.: Addendum to a paper of W. Ewens. Theor. Popul. Biol. **3**(1), 113–116 (1972)
28. Kerns, G.J., Székely, G.J.: De Finetti's theorem for abstract finite exchangeable sequences. J. Theor. Probab. **19**(3), 589–608 (2006)
29. Khoo, Yuehaw, Ying, Lexing: Convex relaxation approaches for strictly correlated density functional theory. SIAM J. Sci. Comput. **41**(4), B773–B795 (2019)
30. Laraki, R., Lasserre, J.B.: Computing uniform convex approximations for convex envelopes and convex hulls. J. Convex Anal. **15**(3), 635–654 (2008)