Annales de la Faculté des Sciences de Toulouse
MATHÉMATIQUES

AYMERIC BARADAT AND ELIAS VENTRE

Convergence of the Sinkhorn algorithm when the Schrödinger problem has no solution

Tome XXXIII, nº 5 (2024), p. 1297–1371.

https://doi.org/10.5802/afst.1800

Publication membre du centre Mersenne pour l'édition scientifique ouverte
http://www.centre-mersenne.org/
e-ISSN : 2258-7519

Annales de la faculté des sciences de Toulouse Volume XXXIII, n° 5, 2024
pp. 1297-1371

Convergence of the Sinkhorn algorithm when the Schrödinger problem has no solution (*)

Aymeric Baradat (1) and Elias Ventre (2)

Abstract. — The Sinkhorn algorithm is one of the most popular methods for solving the Schrödinger problem: it is known to converge as soon as the latter has a solution, and with a linear rate when the solution has the same support as the reference coupling. Motivated by recent applications of the Schrödinger problem where structured stochastic processes lead to degenerate situations with possibly no solution, we show that the Sinkhorn algorithm still gives rise in this case to exactly two limit points, that can be used to compute the solution of a relaxed version of the Schrödinger problem, which appears as the Γ-limit of a problem where the marginal constraints are replaced by marginal penalizations. These results also allow to develop a theoretical procedure for characterizing the support of the solution – both in the original and in the relaxed problem – for any reference coupling and marginal constraints. We showcase promising numerical applications related to a model used in cell biology.

Résumé. — L'algorithme de Sinkhorn est une méthode largement utilisée pour calculer les solutions du problème de Schrödinger car il converge dès que ce dernier admet une solution, et à un taux linéaire dès que cette solution a le même support que la matrice de référence. Récemment, il a été proposé d'appliquer cette théorie à des situations biologiques modélisées par des processus stochastiques structurés donnant lieu à des problèmes de Schrödinger dégénérés pouvant ne pas avoir de solution. Dans cet article, nous démontrons qu'en l'absence de solution, l'algorithme de Sinkhorn admet deux valeurs d'adhérence qui permettent de calculer la solution d'un problème relaxé où les contraintes marginales du problème de Schrödinger sont remplacées par une pénalisation, dans la limite ou celle-ci tend vers l'infini. Notre analyse nous permet également de caractériser le support de la solution, à la fois dans le problème original et dans sa version relaxée. Enfin, nous présentons des applications numériques prometteuses pour l'étude d'un problème de biologie cellulaire.

(*) Reçu le 10 janvier 2023, accepté le 29 septembre 2023.

Keywords: Schrödinger problem, the Sinkhorn algorithm, matrix scaling.

2020 *Mathematics Subject Classification:* 65F35, 65B99, 92-08.

(1) Université Claude Bernard Lyon 1, CNRS UMR 5208, Institut Camille Jordan, Villeurbanne, France — aymeric.baradat@cnrs.fr

This work was supported by funding from French agency ANR (SingleStatOmics; ANR-18-CE45-0023-03).

(2) University of British Columbia, Vancouver, BC, Canada — eventre@math.ubc.ca

Article proposé par Radu Ignat.

1. Introduction

The Schrödinger problem has been introduced by Schrödinger himself in the 30's [32, 33] in the context of statistical mechanics. It is one of these problems in mathematics for which there is periodically a resurgence of interest, as witnessed by the numerous works which it was the object of for almost 100 years, among which [7, 10, 11, 28, 40]. The last of these resurgences, over the past twenty years, occurred because of its close links with Optimal Transport. On the one hand, the Schrödinger problem, that comes with a temperature parameter in its classical formulation, converges towards Optimal Transport when the temperature goes to zero [19, 21, 23]. On the other hand, it is often much easier to compute the solutions of the Schrödinger problem than the ones of the Optimal Transport problem [8, 26], thanks to the so-called Sinkhorn algorithm [34]. This algorithm converges exponentially fast (i.e. at a linear rate, following the usual terminology in the field), at least when it is applied to a reference matrix whose entries are all below bounded by a positive number.

It is known in the theory of matrix scaling that when the reference matrix has nonnegative but possibly cancelling entries, the data in the Schrödinger problem may be chosen in such a way that the latter admits no solution. This is the so-called *non-scalable* case. Also, when the data are located at the boundary of those for which there is a solution, the so-called *approximately scalable* case, the Schrödinger problem has a solution but the convergence of the Sinkhorn algorithm is not linear anymore. This slow-down is known to be related to the fact that in this approximately scalable case, the solution of the Schrödinger problem has a bigger set of zero entries than the reference matrix, or otherwise stated, that its support is smaller [1].

In this paper, we want to study the Sinkhorn algorithm in the degenerate case where the Schrödinger problem has no solution. Our main finding is that for such problems, the Sinkhorn algorithm leads to exactly two limit points, each of them being the solution of a Schrödinger problem with modified data, that we characterize themselves as solutions of auxiliary optimization problems. Also, we show that these limit points are related to a problem where the marginal constraints of the original problem are replaced by marginal penalizations. Moreover, the Schrödinger problem related to the modified data is seen to belong to the approximately scalable case in general. We therefore provide a new outlook on the question of the support of the solution in this case, allowing to design an approximate method for improving the Sinkhorn algorithm's convergence both in the approximately scalable and non-scalable cases.

For simplicity and because it fits with the context of our numerical explorations and needs, we decided to work in finite spaces, even though some of the results might be generalizable.

The Schrödinger problem in finite spaces

Let $\mathcal{D} = \{x_1, \dots, x_N\}$ and $\mathcal{F} = \{y_1, \dots, y_M\}$ be two nonempty finite spaces and $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$ be a nonnegative measure on $\mathcal{D} \times \mathcal{F}$. Of course, we can identify R with a matrix $R = (R_{ij}) \in \mathbb{R}_+^{N \times M}$ by setting $R_{ij} := R(\{(x_i, y_j)\})$. Assuming that R models the coupling between the initial and final positions of the particles of a large system, we interpret R_{ij} as the sum of the masses of all the particles being in x_i at the initial time, and in y_j at the final time.

Let us choose $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$. Once again, we see $\mu = (\mu_i)$ and $\nu = (\nu_j)$ as vectors of $\mathbb{R}_+^N$ and $\mathbb{R}_+^M$ respectively.

We call $\Pi(\mu, \nu)$ the subset of $\mathcal{M}_+(\mathcal{D} \times \mathcal{F})$ consisting of all those matrices $\overline{R}$ whose row and column sums give μ and ν respectively, that is, such that

$$\forall\, i = 1, \dots, N, \quad \sum_j \overline{R}_{ij} = \mu_i, \qquad \text{and} \qquad \forall\, j = 1, \dots, M, \quad \sum_i \overline{R}_{ij} = \nu_j.$$

In our interpretation, it means that for the system described by $\overline{R}$, the sum of the masses of all the particles being in $x_i \in \mathcal{D}$ at the initial time is μ_i, and the sum of the masses of all the particles being in $y_j \in \mathcal{F}$ at the final time is ν_j. In particular, for $\Pi(\mu, \nu)$ to be nonempty, μ and ν need to share their total mass.

Remark 1.1. — Let us point out to the readers acquainted with the notations used in the Optimal Transport literature that calling $X : \mathcal{D} \times \mathcal{F} \to \mathcal{D}$ and $Y : \mathcal{D} \times \mathcal{F} \to \mathcal{F}$ the canonical projections and denoting by $_\#$ the push forward operation on measures, the measure $\overline{R}$ belongs to $\Pi(\mu, \nu)$ provided $X_\#\overline{R} = \mu$ and $Y_\#\overline{R} = \nu$. Actually, we will not use these notations, and prefer to define $\mu^{\overline{R}} := X_\#\overline{R}$ and $\nu^{\overline{R}} := Y_\#\overline{R}$, see formula (2.1).

We call the Schrödinger problem w.r.t. R between μ and ν the convex optimization problem consisting in minimizing among $\Pi(\mu, \nu)$ the relative entropy w.r.t. R:

$$\mathrm{Sch}(R; \mu, \nu) := \min\big\{ H(\overline{R} \,|\, R) \;\big|\; \overline{R} \in \Pi(\mu, \nu) \big\}, \tag{1.1}$$

where for all $\overline{R} \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, the relative entropy of P w.r.t. R is defined by

$$H(\overline{R} \,|\, R) := \sum_{ij} \left\{ \overline{R}_{ij} \log \frac{\overline{R}_{ij}}{R_{ij}} + R_{ij} - \overline{R}_{ij} \right\},$$

taking the conventions $a \log \frac{a}{0} = +\infty$ if $a > 0$, and $0 \log 0 = 0 \log \frac{0}{0} = 0$. Notice that if $H(\overline{R} \,|\, R) < +\infty$, then for all i, j such that $R_{ij} = 0$, we also have $\overline{R}_{ij} = 0$, i.e., $\overline{R} \ll R$ in the sense of measures.

Remark 1.2. — Of course, as before, for a solution R^* to exist, μ and ν need to have the same total mass, and R^* will then have the same total mass as μ and ν.

By strict convexity of the relative entropy as a function of $\overline{R}$, when there is a solution, the latter is unique. Also, the relative entropy being lower semi-continuous w.r.t. $\overline{R}$ and $\Pi(\mu, \nu)$ being compact, the existence of a solution R^* for $\mathrm{Sch}(R; \mu, \nu)$ is equivalent to the existence of a $\overline{R} \in \Pi(\mu, \nu)$ satisfying $H(\overline{R} \,|\, R) < +\infty$. In what follows, such an $\overline{R}$ is called a competitor for $\mathrm{Sch}(R; \mu, \nu)$.

Heuristically, we seek for the measure R^* that is the closest possible to R in the entropic sense while imposing its first and second marginals.

In virtue of the Sanov theorem [29], this problem has an interpretation in terms of large deviations. It is also known to be connected to Optimal Transport problems, see [5, 12, 19, 21]: if for all i, j, c_{ij} models the cost to transport a unit of mass from x_i to y_j, and $R_{ij} \propto \exp(-c_{ij}/\varepsilon)$ for some small $\varepsilon > 0$, then the solution of $\mathrm{Sch}(R; \mu, \nu)$ is a good approximation of a solution of the Optimal Transport problem between μ and ν, of cost (c_{ij}).

The Sinkhorn algorithm

When the solution of $\mathrm{Sch}(R; \mu, \nu)$ exists, it is well known for a very long time that this solution turns out to be the limit of the sequences $(P^n)_{n \in \mathbb{N}^*}$ and $(Q^n)_{n \in \mathbb{N}^*}$ appearing in the following so-called Sinkhorn algorithm, also called IPFP for *iterative proportional fitting procedure* [7, 16, 24, 34, 35]:

$$\begin{cases} Q^0 := R, \\ \forall\, n \geqslant 0, \quad P^{n+1} := \arg\min\big\{H(P \,|\, Q^n), \qquad \mu^P = \mu\big\}, \\ \forall\, n \geqslant 0, \quad Q^{n+1} := \arg\min\big\{H(Q \,|\, P^{n+1}), \quad \nu^Q = \nu\big\}. \end{cases} \tag{1.2}$$

This formulation is implicit as it involves minimization problems. In fact, easy results concerning these problems, detailed in Corollary 2.9 below, give access to an explicit and easily computable version, which takes the following form, when expressed in terms of the so called *dual variables* or *potentials*

$(a^n)_{n\in\mathbb{N}^*} \in (\mathbb{R}_+^N)^{\mathbb{N}}$ and $(b^n)_{n\in\mathbb{N}} \in (\mathbb{R}_+^M)^{\mathbb{N}}$:

$$\begin{cases} \forall\, j, \qquad b_j^0 := 1, \\ \forall\, n \geqslant 0, \quad \forall\, i,j, \quad a_i^{n+1} := \dfrac{\mu_i}{\sum_{j'} b_{j'}^n R_{ij'}}, \qquad P_{ij}^{n+1} := a_i^{n+1} b_j^n R_{ij}, \\ \forall\, n \geqslant 0, \quad \forall\, i,j, \quad b_j^{n+1} := \dfrac{\nu_j}{\sum_{i'} a_{i'}^{n+1} R_{i'j}}, \quad Q_{ij}^{n+1} := a_i^{n+1} b_j^{n+1} R_{ij}. \end{cases} \tag{1.3}$$

A reason for the popularity of this algorithm is that in a lot of contexts, the sequences of potentials (a^n) and (b^n), and hence the sequence of couplings (P^n) and (Q^n) converge at a linear rate, and the limit of (P^n) and (Q^n) coincide with the unique solution of $\mathrm{Sch}(R;\mu,\nu)$. For this reason, the Sinkhorn algorithm is nowadays one of the most efficient ways to compute approximate solutions of Optimal Transport problems [2, 8, 26].

Observe that a priori, the existence of a solution for the Schrödinger problem is not necessary to give a meaning to the Sinkhorn algorithm. Actually, we will see that there are lots of situations where the Schrödinger problem has no solution, and yet the Sinkhorn algorithm is perfectly well defined. These are the cases that we want to study in this text.

A degenerate case

As we just said, our aim is to study the Sinkhorn algorithm in the cases where the existence of a solution of the Schrödinger problem is either false, or at least nontrivial. This includes the case where μ and ν do not have the same total mass, see Remark 1.2. However, this is not the main new situation that we want to encompass, since the Sinkhorn algorithm behaves trivially under normalization. More interestingly, we will give a detailed study of the case where some entries of R cancel, or in Optimal Transport terms, when the cost function takes the value $+\infty$.

In that situation, it can be hard to exhibit a competitor, since the natural candidate that is usually chosen, namely, the product measure of μ and ν, is not absolutely continuous w.r.t. R in general. In fact, there are cases where it is easy to see that no competitor exists. We give in Appendix A an explicit and simple example of such a case. To illustrate our findings, we also describe the behaviour of the Sinkhorn algorithm applied to this example.

Note that beyond the theoretical interest, there are practical motivations for studying cases where the problem has no solution. Indeed, the Schrödinger problem can be used as follows. Suppose that μ and ν are some

observed densities of a random phenomenon at two different timepoints, obtained for instance by building the empirical distributions associated to some collected data. Suppose also that we have at our disposal a good model for this phenomenon, that is, a reference stochastic process chosen based on our knowledge of the system prior to the observations of μ and ν. Let us call R the coupling of this process between the two studied timepoints. If we believe enough in our model and in our data, but still the marginals of R are not μ and ν, then it is reasonable to try to improve our model by looking for the coupling that is the closest to R (for instance in the entropic sense), but which is compatible with the data: this means solving $\mathrm{Sch}(R; \mu, \nu)$.

Now imagine that the support of the coupling R is a strict subset of $\mathcal{D}\times\mathcal{F}$, or in terms of matrices that R has zero entries. This can be perfectly justified in some applications (think for instance of a nondecreasing process, like the size of some randomly growing phenomenon). Then, small measurement errors due to imprecision of the devices or even to too restricted samplings may result in the non-existence of any coupling with marginals μ and ν being absolutely continuous with respect to R: the Schrödinger problem would thus have no solution. In that case, we would like to be able to find a coupling which explains the best the data while being entropically close to R. Some methods are available for doing so, like for instance algorithms solving the so-called unbalanced problem [6], but at the cost of introducing a new parameter quantifying the balance between the proximity to the data and to the reference coupling, whose value will often be arbitrarily chosen. We show in this article that interestingly, the Sinkhorn algorithm allows to overcome this choice in the specific situation where the data are more trustworthy than the model.

In particular, we were motivated by an application of the Sinkhorn algorithm related to systems biology, and more specifically to the treatment of single-cell data. The quick progresses of acquisition methods for such data raises the hope of a better understanding of the cell-differentiation process, which would in turn pave the way for major medical breakthroughs. In the seminal papers [18, 31], Schiebinger and his coauthors suggest to analyse the collected data through an approach based on Optimal Transport and more specifically on the Schrödinger problem.

In this field, the unknown is the law of the evolution of the quantity of mRNA molecules in the cells through time: this evolution cannot be followed, as our techniques of measurement destroy the cells. Hence, to study it between two timepoints, the approach consists in:

(i) choosing a reference theoretical model R, where for all i, j, R_{ij} is the expected quantity of cells whose mRNA levels are given by the vector x_i at the initial time, and by y_j at the final time;
(ii) measuring the mRNA levels of samples of cells at the initial and final times to get approximate distributions μ and ν of these levels among the population of cells under study;
(iii) solving the Schrödinger problem $\mathrm{Sch}(R; \mu, \nu)$ to get a law R^* that is close to our model R, but which explains the data.

In the case of Schiebinger, R is the coupling produced by a Brownian motion between two time points, and therefore has a full support. In a separated work [38], the second author argues that a more realistic model would be obtained by replacing the Brownian motion by a *piecewise deterministic Markov process* as described in [15]. For such models, dynamical constraints involving mRNAs half-life times lead to a degenerate R and the corresponding Schrödinger problem could thus have no solution, not because of a lack in the model, but because of inaccuracies in the measurements. Our results show that the Sinkhorn algorithm can still be used in this situation, without any pre-treatment of the data. We refer once again to Appendix A for a further discussion on this topic.

Contributions

Our results are divided into two parts. In a first part, we will assert that the Sinkhorn algorithm can always be used, even when the corresponding Schrödinger problem has no solution: in that case, the algorithm leads to two limit points which are reasonable relaxed solutions, relevant for many applications. More precisely:

- We show that under the only light assumption that they are well defined, the two sequences $(P^n)_{n \in \mathbb{N}^*}$ and $(Q^n)_{n \in \mathbb{N}^*}$ defined in (1.2) converge towards two possibly different matrices P^* and Q^*, each of them being the solution of a Schrödinger problem with modified marginals. Specifically, the matrix P^* is the solution of the problem $\mathrm{Sch}(R; \mu, \nu^*)$, where ν^* minimizes the relative entropy w.r.t. ν within the set of marginals $\overline{\nu}$ for which the Schrödinger problem $\mathrm{Sch}(R; \mu, \overline{\nu})$ admits a solution, and a similar statement holds for Q^*. This result, stated at Theorem 3.2, is the main result of Section 3. Appendix A presents an explicit example where the Schrödinger problem does not admit a solution, but this theorem still applies.
- Then, we show in Section 4 that P^* and Q^* enable to compute the solution of a modified Schrödinger problem where the marginal

constraints are replaced by marginal *penalizations:* as shown at Theorem 4.3, the limit of the solution of the problem

$$\min\left\{H(\overline{R}\,|\,R)+\lambda\left(H\big(\mu^{\overline{R}}\,\big|\,\mu\big)+H\big(\nu^{\overline{R}}\,\big|\,\nu\big)\right)\,\middle|\,\overline{R}\in\mathcal{M}_+(\mathcal{D}\times\mathcal{F})\right\},$$

(where once again, $\mu^{\overline{R}}$ and $\nu^{\overline{R}}$ are the first and second marginal of $\overline{R}$, see Remark 1.1) converges towards the componentwise geometric mean of the two limits P^* and Q^* of the Sinkhorn algorithm as $\lambda\to+\infty$. This goes through an analysis of the dual of the limiting problem, see Proposition 4.5 and the proof of Theorem 4.3.

Unfortunately, similarly to what happens in the approximately scalable case, when the Schrödinger problem has no solution, then typically the Sinkhorn algorithm does not converge at a linear rate, so that one of the reason of its popularity is lost. Therefore, the second part of this work is dedicated to suggesting a way to improve its speed of convergence. Our reasoning goes as follows:

- In Section 5, we explain that calling $\mathcal{S}$ the (common) support of P^* and Q^*, then the Sinkhorn algorithm applied to $\mathrm{Sch}(\mathbb{1}_{\mathcal{S}}R;\mu,\nu)$ instead of $\mathrm{Sch}(R;\mu,\nu)$ leads to the same limit points but does converge at a linear rate (see Proposition 5.3). Therefore, an idea to improve the speed of convergence is to first compute $\mathcal{S}$ and then apply the Sinkhorn algorithm to the modified problem. The main result of Section 5, namely, Proposition 5.6, is a theoretical procedure designed to compute $\mathcal{S}$ without knowing P^* and Q^*. It relies on a well known necessary and sufficient condition on R, μ and ν for $\mathrm{Sch}(R;\mu,\nu)$ to admit a solution that we recall at Theorem 5.1 and for which we provide a new proof in Appendix C.
- Section 6 is an application of the developments made at Section 5. We implement an approximate but fast algorithm, usable in practice, allowing to recover an estimate of the support $\mathcal{S}$. We then compare the Sinkhorn algorithm and the technique coming from [6] with our method consisting in first computing $\mathcal{S}$ with our approximate algorithm and then applying the Sinkhorn algorithm to $\mathrm{Sch}(\mathbb{1}_{\mathcal{S}}R;\mu,\nu)$. We also detail the regimes in which our method is a significant improvement of the other techniques.

Prior works and difficulties

One of the most famous results about the Schrödinger problem states that in the so-called scalable case, even in the continuous setting, its solutions are (f,g)-transforms of the reference measure R (see for instance [21,

Definition 3.2]). More precisely, being an (f, g)-transform is always sufficient to be a solution of a Schrödinger problem, and it is often necessary. For instance, the following theorem is an easy version, in the finite case, of [24, Theorem 2.1]:

THEOREM 1.3. — *Let $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, $a : \mathcal{D} \to \mathbb{R}_+$ and $b : \mathcal{F} \to \mathbb{R}_+$ (seen respectively as a matrix and two vectors). Then the measure R^* defined for all i, j by*

$$R^*_{ij} = a_i b_j R_{ij} \tag{1.4}$$

(or, for short, $R^* = a \otimes bR$) is the solution of $\mathrm{Sch}(R; \mu^{R^*}, \nu^{R^*})$.

In addition, for given $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$, if there exists $\overline{R} \in \Pi(\mu, \nu)$ with $\overline{R} \sim R$ in the sense of measures (which is the case in particular when $R \sim \mu \otimes \nu$), then there exist $a : \mathcal{D} \to \mathbb{R}_+$ and $b : \mathcal{F} \to \mathbb{R}_+$ such that the solution R^* of $\mathrm{Sch}(R; \mu, \nu)$ *satisfies* (1.4).

A lot of works are dedicated to finding conditions on R, μ, ν for this decomposition to hold in various topological settings, and we refer to [21, 24] for reviews about this question. At least formally, these a and b appear as the exponential of the Lagrange multipliers associated with the marginal constraints, called *Schrödinger potentials*. The existence of such multipliers and of the decomposition (1.4) is important for several reasons:

- These a and b, are often limits of the sequences (a^n) and (b^n) from (1.3), and studying the convergence of these sequences can even be a way to prove the convergence of the Sinkhorn algorithm [9, 27].
- Theorem 1.3 or its continuous versions are a keystone in order to find a necessary and sufficient condition (but only working in the scalable case!) for a law R^* to be the solution $\mathrm{Sch}(R; \mu^{R^*}, \nu^{R^*})$ which is stable under weak limits, and which generalizes the cyclical monotonicity in the theory of Optimal Transport (see for instance [30, Chapter 1]). This notion called cyclical invariance can be used in several contexts, for instance to prove stability of the Schrödinger problem in the continuous case (implying convergence of the Sinkhorn algorithm) or convergence towards Optimal Transport [25].
- Exploiting decomposition (1.4) allows to analyze further Schrödinger problems associated with specific dynamical models [21].

Unfortunately, we will almost never be able to use Theorem 1.3 as such because in most cases, we will have to deal with solutions of approximately scalable problems, those for which the solutions are not of form (1.4). Note that indeed, we are precisely interested in cases when $R \sim \mu \otimes \nu$ does not hold, and when the sequences (a^n) and (b^n) from (1.3) do not converge, as illustrated in Appendix A.

In that case, only a few result are known. Among them, [7] plays an important role since it presents a proof of convergence of the Sinkhorn algorithm in the finite case as soon as the Schrödinger problem admits a solution, and this without referring to Schrödinger potentials or (f, g)-transforms. In our proof of convergence, we will mainly use ideas from that work. However, the proof by Csiszár has an interpretation resulting from an analogy with Euclidean geometry. Our proof still has such an interpretation, but which is much heavier. This is why we provide in Section 3 a proof relying on direct computations, and give a geometric interpretation of this proof in Appendix B. We will be more precise about these questions in Remark 3.3.

In the approximately scalable case, there is still something remaining from Theorem 1.3 in the form of a necessary optimality condition. Namely, for all $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$, if R^* is the solution of $\mathrm{Sch}(R; \mu^{R^*}, \nu^{R^*})$, then there exists $a \in \mathscr{F}(\mathcal{D}; \mathbb{R}_+)$ and $b \in \mathscr{F}(\mathcal{F}; \mathbb{R}_+)$ such that

$$R^* = \mathbb{1}_{\mathcal{S}} a \otimes bR, \tag{1.5}$$

where $\mathcal{S}$ is the support of R^* (see for instance [28, Theorem 3] for this result in a rather general continuous context).

Still, there is a gap between the sufficient optimality condition (1.4) and the necessary optimality condition (1.5), and as far as we know, there is no satisfactory necessary and sufficient condition, stable under weak limits, for a law R^* to be the solution of $\mathrm{Sch}(R; \mu^{R^*}, \nu^{R^*})$ unless further assumptions such as $R \sim \mu^{R^*} \otimes \nu^{R^*}$ are made, ensuring the problem to be scalable. For this difficult topic, we refer to the discussion provided in [21, Section 2]. This is a major obstacle for treating the questions of stability or asymptotics of the Schrödinger problem in the general continuous case, and this is the main reason why we could not generalize our proof of convergence beyond the finite case. We give more details on this topic in Remark 3.5.

We believe that this gap might be filled by finding necessary and sufficient conditions on the support $\mathcal{S}$ to be sure that a R^* of the form (1.5) is a solution of the corresponding Schrödinger problem. Therefore, we think that studying the support of P^* and Q^* in the finite case as it is done in Section 5 could be a step towards a better understanding of the Schrödinger problem in more general settings.

Along this text, we keep in mind that we would like to replace our finite spaces $\mathcal{D}$ and $\mathcal{F}$ by general Polish spaces. This is the reason why we will often write $H(P \,|\, R) < +\infty$ instead of $P \ll R$: these are equivalent in the finite case, but not in the continuous one. In the latter case, we would need the stronger entropic assumption.

Before coming up with our contributions, we recall a few facts about the relative entropy functional and about the Schrödinger problem at Section 2.

2. Notations, properties of the entropy and terminology

In this preliminary section, we introduce some notations, provide well known elementary results concerning the entropy, and recall the terminology usually used in the theory of matrix scaling.

2.1. Notations

Let us first give a few notations that will be used systematically in this work. Most of them were already given in the introduction.

- Whenever I is a finite set of labels and $\mathcal{E} = \{u_k,\ k \in I\}$ is a finite set indexed by I, we denote by $\mathcal{M}(\mathcal{E})$ and $\mathcal{M}_+(\mathcal{E})$ the set of measures and nonnegative measures on $\mathcal{E}$ respectively. These set are identified with $\mathbb{R}^I$ and $\mathbb{R}^I_+$ respectively through the correspondence
$$\mathsf{r} \in \mathcal{M}(\mathcal{E}) \leftrightsquigarrow \big(\mathsf{r}_k := \mathsf{r}(\{u_k\})\big)_{k\,\in\, I} \in \mathbb{R}^I.$$
For all $\mathsf{r} \in \mathcal{M}(\mathcal{E})$, we denote by $\mathsf{M}(\mathsf{r}) := \sum_k \mathsf{r}_k$ its total (signed) mass. If $\mathsf{r} \in \mathcal{M}_+(\mathcal{E})$ and $\mathsf{M}(\mathsf{r}) = 1$, we say that r is a probability measure on $\mathcal{E}$, and we write $\mathsf{r} \in \mathcal{P}(\mathcal{E})$. The topology considered on $\mathcal{M}(\mathcal{E})$, $\mathcal{M}_+(\mathcal{E})$ and $\mathcal{P}(\mathcal{E})$ is the one of $\mathbb{R}^I$.
- In the same way, we identify the set $\mathscr{F}(\mathcal{E};\mathbb{R})$ of real functions Z on $\mathcal{E}$ with $\mathbb{R}^I$ through the correspondence
$$Z \in (\mathcal{E};\mathbb{R}) \leftrightsquigarrow \big(Z_k := Z(u_k)\big)_{k\,\in\, I} \in \mathbb{R}^I.$$
Depending on the context, we will either call such functions Z *test functions*, or *random variables*, thinking of $\mathcal{E}$ as a measurable set. The random variables that we will consider will actually often be slightly more general, and be allowed to take the value $-\infty$, in which case we will tell it explicitly.
- Through our identifications, the duality between $\mathcal{M}(\mathcal{E})$ and $\mathscr{F}(\mathcal{E};\mathbb{R})$ is nothing but the usual scalar product on $\mathbb{R}^I$, and denoted for all $Z \in \mathscr{F}(\mathcal{E};\mathbb{R})$ and $\mathsf{r} \in \mathcal{M}(\mathcal{E})$ by
$$\langle Z, \mathsf{r}\rangle := \sum_k Z_k \mathsf{r}_k.$$
When Z possibly takes the value $-\infty$, we always choose by convention $-\infty \times 0 = 0$.

- In the context of the introduction, when $\mathcal{D} = \{x_1, , \dots, x_N\}$ and $\mathcal{F} = \{y_1, \dots, y_M\}$ are two nonempty finite spaces and $\mathcal{E} = \mathcal{D} \times \mathcal{F}$, then the corresponding I is the product space $\{1, \dots, N\} \times \{1, \dots, M\}$, and $\overline{R} \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, is seen as a matrix. We define its marginals $\mu^{\overline{R}} \in \mathcal{M}_+(\mathcal{D})$ and $\nu^{\overline{R}} \in \mathcal{M}_+(\mathcal{F})$ by the formulas
$$\begin{aligned} &\forall\, i = 1, \dots, N, \quad \mu_i^{\overline{R}} := \sum_j \overline{R}_{ij}, \\ \text{and} \quad &\forall\, j = 1, \dots, M, \quad \nu_j^{\overline{R}} := \sum_i \overline{R}_{ij}. \end{aligned} \tag{2.1}$$
Of course, $\mu^{\overline{R}}$ and $\nu^{\overline{R}}$ have the same total mass as $\overline{R}$, that is:
$$\mathsf{M}(\overline{R}) = \mathsf{M}\big(\mu^{\overline{R}}\big) = \mathsf{M}\big(\nu^{\overline{R}}\big). \tag{2.2}$$
In particular, if R is a probability measure, its marginals are probability measures as well.
- As before, if $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$, we call $\Pi(\mu, \nu)$ the set of measures $\overline{R} \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$ such that $\mu^{\overline{R}} = \mu$ and $\nu^{\overline{R}} = \nu$. We also call $\Pi_1(\mu)$ and $\Pi_2(\nu)$ the set of those $\overline{R} \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$ satisfying $\mu^{\overline{R}} = \mu$ for the first one, and $\nu^{\overline{R}} = \nu$ for the second one. In particular, $\Pi(\mu, \nu) = \Pi_1(\mu) \cap \Pi_2(\nu)$.
- For the sake of simplicity, we do not use different notations for the same functions applied in different context. For instance, notations for the total mass M or the relative entropy H (see Definition 2.1 below) might be applied to different sets $\mathcal{E}$ namely $\mathcal{D}$, $\mathcal{F}$ and $\mathcal{D} \times \mathcal{F}$.

2.2. First properties of the relative entropy

This subsection only contains easy and very well known results concerning the relative entropy that will be useful in the sequel. We stick to the finite case as this is the one studied in this paper, and we provide some proofs for the readers who are not acquainted with this notion of entropy. Yet, all the properties given here are known in a much wider context, see for instance [20] for what concerns continuity and duality properties, and [7] for what concerns entropic projections.

Definition and topological properties

As already said in the introduction, the relative entropy is defined as follows.

Definition 2.1. — *Let $\mathcal{E} = \{u_k,\ k \in I\}$ be a finite set and $\mathsf{r} = (\mathsf{r}_k) \in \mathcal{M}_+(\mathcal{E})$. For all $\bar{\mathsf{r}} = (\bar{\mathsf{r}}_k) \in \mathcal{M}_+(\mathcal{E})$, the relative entropy of $\bar{\mathsf{r}}$ w.r.t. r is the value in $[0, +\infty]$ given by*

$$H(\bar{\mathsf{r}} \,|\, \mathsf{r}) := \sum_k \left\{ \bar{\mathsf{r}}_k \log \frac{\bar{\mathsf{r}}_k}{\mathsf{r}_k} + \mathsf{r}_k - \bar{\mathsf{r}}_k \right\} = \sum_k \bar{\mathsf{r}}_k \log \frac{\bar{\mathsf{r}}_k}{\mathsf{r}_k} + \mathsf{M}(\mathsf{r}) - \mathsf{M}(\bar{\mathsf{r}}),$$

with convention $a \log \frac{a}{0} = +\infty$ for all $a > 0$, and $0 \log 0 = 0 \log \frac{0}{0} = 0$.

First, this definition provides a convex function with good continuity properties. We state them in the following proposition, for which we omit the straightforward proof.

Proposition 2.2. — *Let $\mathcal{E}$ be a finite set and $\mathsf{r} \in \mathcal{M}_+(\mathcal{E})$. The functional*

$$\bar{\mathsf{r}} \in \mathcal{M}_+(\mathcal{E}) \mapsto H(\bar{\mathsf{r}} \,|\, \mathsf{r}) \in [0, +\infty]$$

is strictly convex, lower semicontinuous with compact sublevels, and continuous on its domain, which is the closed set $\{\bar{\mathsf{r}} \ll \mathsf{r}\} \subset \mathcal{M}_+(\mathcal{E})$.

For a given $\bar{\mathsf{r}} \in \mathcal{M}_+(\mathcal{E})$, the functional

$$\mathsf{r} \in \mathcal{M}_+(\mathcal{E}) \mapsto H(\bar{\mathsf{r}} \,|\, \mathsf{r}) \in [0, +\infty]$$

is convex and continuous for the canonical topology of $[0, +\infty]$. Its domain is the open set $\{\mathsf{r} \gg \bar{\mathsf{r}}\} \subset \mathcal{M}_+(\mathcal{E})$.

Finally, for all $\mathsf{r}, \bar{\mathsf{r}} \in \mathcal{M}_+(\mathcal{E})$, $H(\bar{\mathsf{r}} \,|\, \mathsf{r})$ cancels if and only if $\bar{\mathsf{r}} = \mathsf{r}$.

For convenience, we give the following notation for the domain of the relative entropy.

Definition 2.3. — *Given $\mathcal{E} := \{u_k, k \in I\}$ be a finite set and $\mathsf{r} \in \mathcal{M}_+(\mathcal{E})$, we denote by $\mathcal{H}_+(\mathsf{r})$ the set of measures $\bar{\mathsf{r}} =\in \mathcal{M}_+(\mathcal{E})$ such that $H(\bar{\mathsf{r}} \,|\, \mathsf{r}) < +\infty$, or equivalently $\bar{\mathsf{r}} \ll \mathsf{r}$. We also define $\mathcal{H}(\mathsf{r})$ as the subset of signed measures $\bar{\mathsf{r}} \in \mathcal{M}(\mathcal{E})$ satisfying $\bar{\mathsf{r}} \ll \mathsf{r}$. In terms of vectors of $\mathbb{R}^I$, it simply means that the zero coordinates of r are also zero coordinates of $\bar{\mathsf{r}}$.*

A duality formula

One of the most useful property of the relative entropy is the computation of its Legendre transform. This property can be stated as follows.

Theorem 2.4. — *Let $\mathcal{E} = \{u_k,\ k \in I\}$ be a finite set, and $\mathsf{r} \in \mathcal{M}_+(\mathcal{E})$. For all test function Z possibly taking the value $-\infty$ on $\mathcal{E}$ and all nonnegative measure $\bar{\mathsf{r}}$ on $\mathcal{E}$, we have*

$$\langle Z, \bar{\mathsf{r}} \rangle \leqslant H(\bar{\mathsf{r}} \,|\, \mathsf{r}) + \left\langle e^Z - 1, \mathsf{r} \right\rangle, \tag{2.3}$$

with conventions $e^{-\infty} = 0$, $-\infty \times 0 = 0$ *and* $+\infty - \infty = +\infty$.

Moreover, equality in $\mathbb{R}$ *holds if and only if* $\bar{\mathsf{r}} \in \mathcal{H}_+(\mathsf{r})$ *and for all* $k \in I$,

$$Z_k = \log \frac{\bar{\mathsf{r}}_k}{\mathsf{r}_k} \in [-\infty, +\infty) \tag{2.4}$$

with convention $\log \frac{0}{a} = -\infty$ *for all* $a \geqslant 0$.

Proof. — Let $\mathsf{r}, \bar{\mathsf{r}}$ and Z be as in the statement of the theorem. If $H(\bar{\mathsf{r}} \,|\, \mathsf{r}) = +\infty$, there is nothing to prove, and we assume $\bar{\mathsf{r}} \in \mathcal{H}_+(\mathsf{r})$.

By direct real computations, with the same conventions as in the statement of the theorem, we find that for all $k \in I$:

$$Z_k \bar{\mathsf{r}}_k \leqslant \bar{\mathsf{r}}_k \log \frac{\bar{\mathsf{r}}_k}{\mathsf{r}_k} + \mathsf{r}_k - \bar{\mathsf{r}}_k + \left(e^{Z_k} - 1\right) \mathsf{r}_k,$$

with equality if and only if $\mathsf{r}_k = \bar{\mathsf{r}}_k = 0$ or $\mathsf{r}_k > 0$ and

$$Z_k = \log \frac{\bar{\mathsf{r}}_k}{\mathsf{r}_k} \in [-\infty, +\infty).$$

We find (2.3) and (2.4) by summing this inequality over k. □

Entropic projections

The seminal paper [7] has developed a very powerful approach for studying entropy minimization problems: the relative entropy shares some properties with squared Euclidean distances. Specifically, the notion of entropic projections is particularly relevant since they lead to an analogue of the Pythagorean law.

Actually, these tools will only be used in Lemma 4.6 and Appendix B, but we still decided to present them to fix some ideas about what we do.

Let $\mathcal{E}$ be a finite set, and $\mathsf{r} \in \mathcal{M}_+(\mathcal{E})$. Because of the properties of the relative entropy given at Proposition 2.2, the following proposition is straightforward

Proposition 2.5. — *For all closed convex set* $\mathcal{C} \subset \mathcal{M}_+(\mathcal{E})$*, the optimization problem*

$$\min\left\{ H(\bar{\mathsf{r}} \,|\, \mathsf{r}) \;\middle|\; \bar{\mathsf{r}} \in \mathcal{C} \right\} \tag{2.5}$$

admits a minimizer if and only if $\mathcal{C} \cap \mathcal{H}_+(\mathsf{r}) \neq \emptyset$*. In that case, this minimizer is unique.*

The solution of this minimization problem is called the entropic projection of r on $\mathcal{C}$.

DEFINITION 2.6. — *Let $\mathcal{C}$ be a closed convex subset of $\mathcal{M}_+(\mathcal{E})$. When it exists, the unique minimizer r^* of (2.5) is called the entropic projection of r on $\mathcal{C}$.*

With this definition and in the context of the introduction, the solution of the Schrödinger problem (1.1) is the entropic projection of R on $\Pi(\mu,\nu)$, and the Sinkhorn algorithm (1.2) consists in alternately projecting R on $\Pi_1(\mu)$ and $\Pi_2(\nu)$ as defined in Section 2.1.

In our setting, the main result of [7] can be stated as follows. Note that identity (2.7) is reminiscent of the Euclidean Pythagorean law.

THEOREM 2.7 (Adapted from [7, Lemma 2.1 and Theorem 2.2]). — *Let $\mathcal{C}$ be a closed convex subset of $\mathcal{M}_+(\mathcal{E})$ and $\mathsf{r}^* \in \mathcal{C} \cap \mathcal{H}_+(\mathsf{r})$. The three following assertions are equivalent:*

(1) *r^* is the entropic projection of r on $\mathcal{C}$.*
(2) *For all $\mathsf{v} \in \mathcal{M}(\mathcal{E})$ for which there exists $\varepsilon > 0$ such that $\mathsf{r}^* + \varepsilon\mathsf{v} \in \mathcal{C} \cap \mathcal{H}_+(\mathsf{r})$, then $\mathsf{v} \in \mathcal{H}(\mathsf{r}^*)$ and*

$$\sum_{k \in \mathcal{E}} \mathsf{v}_k \log \frac{\mathsf{r}^*_k}{\mathsf{r}_k} \geqslant 0. \tag{2.6}$$

(3) *For all $\bar{\mathsf{r}} \in \mathcal{C} \cap \mathcal{H}_+(\mathsf{r})$, then $\bar{\mathsf{r}} \in \mathcal{H}_+(\mathsf{r}^*)$ and*

$$H(\bar{\mathsf{r}} \,|\, \mathsf{r}) \geqslant H(\bar{\mathsf{r}} \,|\, \mathsf{r}^*) + H(\mathsf{r}^* \,|\, \mathsf{r}). \tag{2.7}$$

If in addition, $\mathcal{E} = \mathcal{D} \times \mathcal{F}$ as in the introduction, and $\mathcal{C}$ is $\Pi_1(\mu)$, $\Pi_2(\nu)$ or $\Pi(\mu,\nu)$ for some $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$, then the conclusion holds with an equality sign in (2.6) and (2.7).

Remark 2.8. — The main difference in (2.7) with the Euclidean case is that the relative entropy is not symmetric: the order in which the different measures appear is crucial!

Proof. — (3) $\Rightarrow$ (1) just follows from the positivity of the relative entropy.

Let us now prove (2) $\Rightarrow$ (3). Let us consider $\bar{\mathsf{r}} \in \mathcal{C} \cap \mathcal{H}_+(\mathsf{r})$. As $\mathsf{v} := \bar{\mathsf{r}} - \mathsf{r}^*$ satisfies the conditions of point (2) with $\varepsilon = 1$, we also have $\bar{\mathsf{r}} - \mathsf{r}^* \in \mathcal{H}(\mathsf{r}^*)$, which easily implies $\bar{\mathsf{r}} \in \mathcal{H}(\mathsf{r}^*)$. Then, the crucial observation is that

$$\begin{aligned} H(\bar{\mathsf{r}} \,|\, \mathsf{r}) - H(\bar{\mathsf{r}} \,|\, \mathsf{r}^*) &= \sum_k \bar{\mathsf{r}}_k \left(\log \frac{\bar{\mathsf{r}}_k}{\mathsf{r}_k} - \log \frac{\bar{\mathsf{r}}_k}{\mathsf{r}^*_k} \right) + \mathsf{M}(\mathsf{r}) - \mathsf{M}(\mathsf{r}^*) \\ &= \sum_k \bar{\mathsf{r}}_k \log \frac{\mathsf{r}^*_k}{\mathsf{r}_k} + \mathsf{M}(\mathsf{r}) - \mathsf{M}(\mathsf{r}^*), \end{aligned}$$

which is affine in $\bar{\mathsf{r}}$! Now, applying (2.6) with $\mathsf{v} := \bar{\mathsf{r}} - \mathsf{r}^*$, we find that

$$\sum_k \bar{\mathsf{r}}_k \log \frac{\mathsf{r}^*_k}{\mathsf{r}_k} \geqslant \sum_k \mathsf{r}^*_k \log \frac{\mathsf{r}^*_k}{\mathsf{r}_k}.$$

Plugging this identity in the previous one, we get (2.7).

Finally, let us prove (1) $\Rightarrow$ (2), which is the principal result of Theorem 2.7. Let $\mathsf{v} \in \mathcal{M}(\mathcal{E})$ and $\varepsilon > 0$ satisfy the condition of point (2). By convexity of $\mathcal{C} \cap \mathcal{H}_+(\mathcal{E})$, for all $t \in [0, \varepsilon]$, $\mathsf{r}_t := \mathsf{r} + tv \in \mathcal{C} \cap \mathcal{H}_+(\mathcal{E})$ and by optimality of r^*,

$$H(\mathsf{r}^* \,|\, \mathsf{r}) \leqslant H(\mathsf{r}_t \,|\, \mathsf{r}) = \sum_k (\mathsf{r}^*_k + t\mathsf{v}_k) \log \frac{\mathsf{r}^*_k + t\mathsf{v}_k}{\mathsf{r}_k} + \mathsf{M}(\mathsf{r}) - \mathsf{M}(\mathsf{r}^*) - t\mathsf{M}(\mathsf{v}),$$

with equality when $t = 0$. As $\tau \mapsto \tau \log \tau$ has infinite negative slope at $\tau = 0$, this inequality ensures to have $\mathsf{v} \in \mathcal{H}(\mathsf{r}^*)$, and deriving the inequality at $t = 0$ leads easily to (2.6).

For the last point of the statement, the only thing to prove is that when $\mathcal{C}$ is $\Pi_1(\mu)$, $\Pi_2(\nu)$ or $\Pi(\mu, \nu)$, if R^* is the entropic projection of some $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$ on $\mathcal{C}$, and if $V \in \mathcal{M}(\mathcal{D} \times \mathcal{F})$ and $\varepsilon > 0$ are such that $R^* + \varepsilon V \in \mathcal{C} \cap \mathcal{H}_+(R)$, then there exists $\varepsilon' > 0$ such that $R^* - \varepsilon' V \in \mathcal{C} \cap \mathcal{H}_+(R)$. Indeed, we would conclude that (2.6) holds with both inequality signs, so it would hold with an equality sign, and the rest would follow easily.

The conditions defining $\mathcal{C}$ are affine, so under the condition that $R^* + \varepsilon V \in \mathcal{C} \cap \mathcal{H}(R)$, then $R^* + tV$ satisfy these conditions for all $t \in \mathbb{R}$, and we only need to check that for $\varepsilon' > 0$ sufficiently small, $R^* - \varepsilon' V \in \mathcal{H}_+(R)$. But by point 2., we actually have $V \in \mathcal{H}(R^*)$, so for $\varepsilon' > 0$ sufficiently small, $R^* - \varepsilon' V$ has nonnegative entries, and hence belongs to $\mathcal{H}_+(R^*) \subset \mathcal{H}_+(R)$. The conclusion follows. □

One step of the Sinkhorn algorithm

We have now all the concepts necessary to get a full understanding of one step in the Sinkhorn algorithm (1.2).

Corollary 2.9. — *Let $\mathcal{D}$ and $\mathcal{F}$ be two finite sets, and $\overline{R}, R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$. With the notations of* (2.1)*, we have*

$$H\big(\mu^{\overline{R}} \,\big|\, \mu^R\big) \leqslant H(\overline{R} \,|\, R) \qquad \text{and} \qquad H\big(\nu^{\overline{R}} \,\big|\, \nu^R\big) \leqslant H(\overline{R} \,|\, R). \tag{2.8}$$

In the case where $H(\overline{R} \,|\, R)$ is finite, equality holds if and only if for all i, j, respectively:

$$\overline{R}_{ij} = \frac{\mu_i^{\overline{R}}}{\mu_i^R} R_{ij} \qquad \text{and} \qquad \overline{R}_{ij} = \frac{\nu_j^{\overline{R}}}{\nu_j^R} R_{ij},$$

with convention $\frac{0}{0} = 0$.

In particular, given $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$ *and* $\mu \in \mathcal{M}_+(\mathcal{D})$, R *admits an entropic projection* P *on* $\Pi_1(R)$ *if and only if* $H(\mu \,|\, \mu^R) < +\infty$, *and in this case, we have for all* i, j

$$P_{ij} = \frac{\mu_i}{\mu_i^R} R_{ij} \tag{2.9}$$

with convention $\frac{0}{0} = 0$. *Moreover,* $H(P \,|\, R) = H(\mu \,|\, \mu^R)$.

Similarly, given $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$ *and* $\nu \in \mathcal{M}_+(\mathcal{F})$, R *admits an entropic projection* Q *on* $\Pi_2(\nu)$ *if and only if* $H(\nu \,|\, \nu^R) < +\infty$, *and in this case,we have for all* i, j

$$Q_{ij} = \frac{\nu_j}{\nu_j^R} R_{ij}$$

with convention $\frac{0}{0} = 0$. *Moreover,* $H(Q \,|\, R) = H(\nu \,|\, \nu^R)$.

Proof. — The first inequality in (2.8) is a direct application of (2.3) with $\mathsf{r} = R$, $\bar{\mathsf{r}} = \overline{R}$ and for all i, j,

$$Z_{ij} = \log \frac{\mu_i^{\overline{R}}}{\mu_i^R}.$$

The second inequality is proved in the same way, and the equality case is a consequence of (2.4).

For the second part of the statement, let us observe that for all $P \in \Pi_1(\mu)$, because of (2.8), $H(P \,|\, R) \geqslant H(\mu \,|\, \mu^R)$, which – by the equality case – is attained if and only if (2.9) holds. The problem involving the second marginal is treated in the same way. □

2.3. The Schrödinger problem, assumptions and terminology

Definition of the Schrödinger problem.

Let $\mathcal{D} = \{x_1, \ldots, x_N\}$ and $\mathcal{F} = \{y_1, \ldots, y_M\}$ be two nonempty finite sets, and let us choose a reference measure $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$. Given $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$, the Schrödinger problem, already defined in the introduction aims at finding the entropic projection of R on $\Pi(\mu, \nu)$, which rewrites with the notations of Section 2.1:

$$\mathrm{Sch}(R; \mu, \nu) := \min \left\{ H(\overline{R} \,|\, R) \;\middle|\; \begin{array}{l} \overline{R} \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F}) \\ \text{such that } \mu^{\overline{R}} = \mu \text{ and } \nu^{\overline{R}} = \nu \end{array} \right\}. \tag{2.10}$$

Once again, existence and uniqueness hold as soon as $\Pi(\mu, \nu) \cap \mathcal{H}_+(R)$ is nonempty.

Remark 2.10. — Here, we define $\mathrm{Sch}(R;\mu,\nu)$ as the optimal value of our problem. However, with an abusive terminology, we will refer to the minimizer of the r.h.s. of (2.10) as "the solution of $\mathrm{Sch}(R;\mu,\nu)$". More generally, we will call "the problem $\mathrm{Sch}(R;\mu,\nu)$" the optimization problem consisting in computing the value $\mathrm{Sch}(R;\mu,\nu)$.

Assumptions for this work

As we will see in Theorem 3.2, the Sinkhorn algorithm (1.2) associated with the Schrödinger problem $\mathrm{Sch}(R;\mu,\nu)$ is well defined if and only if the following assumption holds.

Assumption 2.11. — *Let $R \in \mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$, and let us call*

$$\mathcal{E} := \Big\{(x_i,y_j)\in\mathcal{D}\times\mathcal{F} \text{ such that } R_{ij}>0, \mu_i>0 \text{ and } \nu_j>0\Big\}. \tag{2.11}$$

We say that the triple $(R;\mu,\nu)$ satisfies Assumption 2.11 provided $R^0 := \mathbb{1}_{\mathcal{E}}\cdot R$ is such that:

$$H\big(\mu \,\big|\, \mu^{R^0}\big) < +\infty \qquad \text{and} \qquad H\big(\nu \,\big|\, \nu^{R^0}\big) < +\infty. \tag{2.12}$$

This assumption is easily seen to be necessary for $\mathrm{Sch}(R;\mu,\nu)$ to admit a solution. Under Assumption 2.11 either $\mathsf{M}(\mu)=\mathsf{M}(\nu)=0$, or none of them is 0. In the second case, up to replacing $\mathcal{D}$ by $\mathcal{D}'$, the support of μ, $\mathcal{F}$ by $\mathcal{F}'$, the support of ν, and R by its restriction (or equivalently of the one of R^0) on $\mathcal{D}'\times\mathcal{F}'$, we end up with the following assumption, that will often be used in this paper.

Assumption 2.12. — *Let $R \in \mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$. We say that the triple $(R;\mu,\nu)$ satisfies Assumption 2.12 provided the support of μ and μ^R is $\mathcal{D}$ and the support of ν and ν^R is $\mathcal{F}$.*

Terminology

The Schrödinger problem (2.10) consists in minimizing a convex function under linear constraints. Therefore, the functional $(\mu,\nu)\in\mathcal{M}_+(\mathcal{D})\times\mathcal{M}_+(\mathcal{F})\mapsto \mathrm{Sch}(R;\mu,\nu)\in[0,+\infty]$ is convex.

In the case where Assumption 2.12 holds, following the usual terminology of the matrix scaling theory (except for the last item which is more exotic), see [16], we say that:

- The problem is *scalable* if (μ, ν) is in the relative interior of the domain of $\mathrm{Sch}(R; \cdot)$. In this case, $\mathsf{M}(\mu) = \mathsf{M}(\nu)$, the Schrödinger problem admits a unique solution R^*, $R^* \sim R$ in the sense of measures, and the Sinkhorn algorithm converges towards R^*, at a linear rate. In Lemma 5.2, we recall an explicit necessary and sufficient condition on R, μ, ν for $\mathrm{Sch}(R; \mu, \nu)$ to be scalable.
- The problem is *approximately scalable* if (μ, ν) is at the relative boundary of the domain of $\mathrm{Sch}(R; \cdot)$. In this case, $\mathsf{M}(\mu) = \mathsf{M}(\nu)$, the Schrödinger problem admits a unique solution R^*, and the Sinkhorn algorithm converges towards R^*. However, in this case, the support of R^* is strictly included in the support of R (else, we easily see that we are in the scalable case), and the rate cannot be linear anymore: as proved in [1], a linear rate of convergence for the Sinkhorn algorithm is not compatible with the appearance of new zero entries at the limit. We recall at Theorem 5.1 a necessary and sufficient condition on R, μ and ν for $\mathrm{Sch}(R; \mu, \nu)$ to be at least approximately scalable, that is, either approximately scalable or scalable.
- The problem is *non-scalable* if $\mathsf{M}(\mu) = \mathsf{M}(\nu)$, but the Schrödinger problem $\mathrm{Sch}(R; \mu, \nu)$ does not admit a solution. This is the case when the condition of Theorem 5.1 does not hold. This case is the main case of interest in this work.
- The problem is *unbalanced* if $\mathsf{M}(\mu) \neq \mathsf{M}(\nu)$. Calling $\mu' := \mu/\mu(\mathcal{D})$ and $\nu' := \nu/\nu(\mathcal{F})$ their normalized versions, we will say that $\mathrm{Sch}(R; \mu, \nu)$ is respectively unbalanced scalable, unbalanced approximately scalable and unbalanced non-scalable whenever $\mathrm{Sch}(R; \mu', \nu')$ is scalable, approximately scalable or non-scalable.

Yet, with an abuse of terminology, we will often refer to the non-scalable case for results that are true in *any* situation, including the balanced and unbalanced non-scalable ones, which are often the most difficult.

3. The Sinkhorn algorithm in the non-scalable case

In this section, we consider $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ that we identify respectively with a matrix and two vectors, as before.

The goal of this section is to show that under obvious necessary assumptions, then the algorithm given in (1.2) is well defined, and that the sequences $(P^n)_{n \in \mathbb{N}^*}$ and $(Q^n)_{n \in \mathbb{N}^*}$ that it provides converge separately towards matrices P^* and Q^* that we define now. It will be obvious from their definition that these matrices coincide if and only if the problem $\mathrm{Sch}(R; \mu, \nu)$ defined

in (2.10) admits a solution, that is, if it is at least approximately scalable. Hence our proof recovers the classical fact that the Sinkhorn algorithm converges towards the solution of the Schrödinger problem as soon as the latter exists. Actually, in that case, it follows the lines of the one given in [7, Theorem 3.2].

The first step to define P^* and Q^* is to define a pair of new marginals $\mu^* \in \mathcal{M}_+(\mathcal{D})$ and $\nu^* \in \mathcal{M}_+(\mathcal{F})$ as solutions of the following optimization problem:

$$\begin{aligned} \mu^* &:= \arg\min\Big\{ H(\bar{\mu} \,|\, \mu) \,\Big|\, \bar{\mu} = \mu^Q \text{ for some } Q \in \Pi_2(\nu) \cap \mathcal{H}_+(R) \Big\}, \\ \nu^* &:= \arg\min\Big\{ H(\bar{\nu} \,|\, \nu) \,\Big|\, \bar{\nu} = \nu^P \text{ for some } P \in \Pi_1(\mu) \cap \mathcal{H}_+(R) \Big\}. \end{aligned} \tag{3.1}$$

These are the entropic projections of μ and ν on the sets $\{\mu^Q \,|\, Q \in \Pi_2(\nu) \cap \mathcal{H}_+(R)\}$ and $\{\nu^P \,|\, P \in \Pi_1(\mu) \cap \mathcal{H}_+(R)\}$ respectively. The question of existence of μ^* and ν^* is treated in Theorem 3.2 below. Of course, if the problem $\mathrm{Sch}(R; \mu, \nu)$ admits a competitor, then $\mu^* = \mu$ and $\nu^* = \nu$.

Remark 3.1. — In the unbalanced case, notice that the total mass of ν^* is the one of μ, and the total mass of μ^* is the one of ν, that is, $\mathsf{M}(\nu^*) = \mathsf{M}(\mu)$ and $\mathsf{M}(\mu^*) = \mathsf{M}(\nu)$.

Then P^* and Q^* are simply defined as the solutions of the Schrödinger problems $\mathrm{Sch}(R; \mu, \nu^*)$ and $\mathrm{Sch}(R; \mu^*, \nu)$ respectively, that is:

$$\begin{aligned} P^* &:= \arg\min\Big\{ H(P \,|\, R) \,\Big|\, P \in \Pi(\mu, \nu^*) \Big\} \\ \text{and} \quad Q^* &:= \arg\min\Big\{ H(Q \,|\, R) \,\Big|\, Q \in \Pi(\mu^*, \nu) \Big\}. \end{aligned} \tag{3.2}$$

Of course, if the problem $\mathrm{Sch}(R; \mu, \nu)$ admits a competitor, and hence a solution, then both P^* and Q^* coincide with this solution.

Our convergence theorem can be stated as follows.

Theorem 3.2. — *Let $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ satisfy Assumption 2.11. The sequences $(P^n)_{n \in \mathbb{N}^*}$ and $(Q^n)_{n \in \mathbb{N}^*}$ from* (1.2)*, the marginals μ^* and ν^* from* (3.1) *and the matrices P^* and Q^* from* (3.2) *are well defined, and*

$$P^n \underset{n \to +\infty}{\longrightarrow} P^* \qquad \text{and} \qquad Q^n \underset{n \to +\infty}{\longrightarrow} Q^*.$$

Remark 3.3.

- Assumption 2.11 is necessary: it is straightforward to check that if Q^1 from (1.2) is well defined, then $H(Q^1 \,|\, R^0) < +\infty$. In particular, projecting on the second marginal, we conclude that $H(\nu \,|\, \nu^{R^0}) < +\infty$. Arguing in the same way with P^2 in place of Q^1 and the second

marginal in place of the first one, we see that if P^2 is well defined, then $H(\mu \,|\, \mu^{R^0}) < +\infty$. In particular, there is nothing to check before starting the algorithm: if the algorithm is able to compute P^2, then it means that our assumption is satisfied and that the convergence holds.

- Note that the topology for the convergence stated in the theorem does not matter since we are working in finite dimensional spaces. However, we believe that as in the scalable case, the result is still true replacing $\mathcal{D}$ and $\mathcal{F}$ by general Polish spaces (see [24] for a review of this question in the scalable case). In this case, the convergence needs to be understood in the sense of the narrow topology, a topology for which the sequences (P^n) and (Q^n) can be proved to be compact due to the properties of their marginals. We detail in Remark 3.5 below the reasons why the known proofs of convergence in the continuous case do not apply directly in the non-scalable case.
- Surprisingly, we will be able to prove this theorem without deriving the optimality conditions for μ^* and ν^*. However, these optimality conditions will be needed in the next section, and hence written at Proposition 4.5.
- The Sinkhorn algorithm consists in alternately entropically projecting R on $\Pi_1(\mu)$ and $\Pi_2(\nu)$. With this picture in mind, we can give in Figure 3.1 a visual representation of the scalable and non-scalable case. In the scalable case, the two convex sets intersect, and the sequences $(P^n)_{n\in\mathbb{N}^*}$ and $(Q^n)_{n\in\mathbb{N}^*}$ converge towards the point of the intersection that is the closest to R. In the non-scalable case, the two convex sets do not intersect. However, the sequences $(P^n)_{n\in\mathbb{N}^*}$ and $(Q^n)_{n\in\mathbb{N}^*}$ still converge respectively to P^* and Q^*, the two extreme points of the shortest line segment connecting both sets (or more precisely, of the shortest line segment that is the closest to R, as shortest line segments are not unique in general). Theorem 3.2 indeed justifies this type of behaviour for the Sinkhorn algorithm. One should still keep in mind that this analogy and our drawings are only sketchy. In reality, the projections are not orthogonal, and the convex sets are affine.
- Our proof is very similar to the one given in [7, Theorem 3.2] in the scalable or approximately scalable case. There, the proof relies on a fundamental formula (see [7, formula (3.14)]) which describes how the iterates of the Sinkhorn algorithm gets closer and closer to $\Pi(\mu,\nu)$ in an entropic sense. More precisely, for all $\overline{R} \in \Pi(\mu,\nu)$ having finite entropy w.r.t. R and $n \geqslant 0$, this formula rewrites with our notations

$$H(\overline{R}\,|\,R) = H(\overline{R}\,|\,Q^n) + \sum_{k=1}^{n} H(Q^k\,|\,P^k) + H(P^k\,|\,Q^{k-1}), \tag{3.3}$$

so that $(H(\overline{R}\,|\,Q^n))$ is nonincreasing. The main novelty of our proof is our definitions of μ^* and ν^*, and the fact that our crucial formula (3.11) is slightly different. Our new ideas are therefore concentrated in Step 2 below whose goal is to derive (3.11). From this step, it is also obvious why μ^* and ν^* need to be defined as in (3.1).

Remarkably, as explained in [7], formula (3.3) has a geometric interpretation: it can be deduced directly from the Pyhthagorean law stated at Theorem 2.7. This is still true for our formula (3.11), but much more complex. We decided to give here a hand by hand proof not referring explicitly to this geometric interpretation. In Appendix B, we show how (2.7) implies (3.3) in the case studied by Csiszár, and explain precisely how it needs to be changed in the non-scalable case. We also provide there some more realistic even though less visual drawings, see Figure B.1.

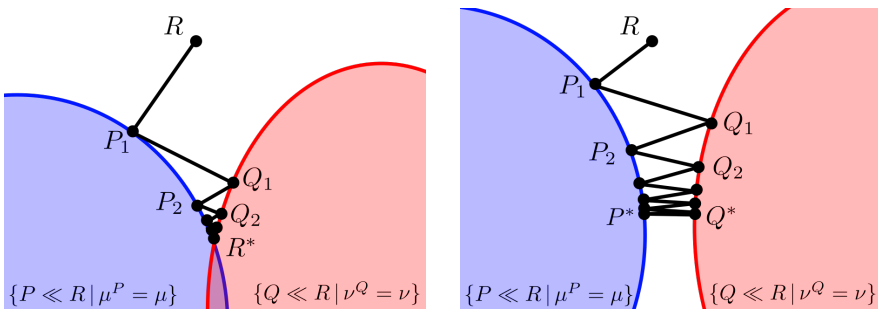

Figure 3.1. Sketchy[(1)] representation of the Sinkhorn algorithm in the scalable case (to the left) and nonscalable case (to the right).

[(1)] We would like to warn the reader that our drawing is not only sketchy, it can also be misleading: it is not always true that alternate orthogonal projections on two intersecting convex sets converge towards the intersection point that is the closest to the original point! However, it is true when the convex sets are affine.

Proof of Theorem 3.2.

Step 1: All the objects are well defined. — Let us first show that under the assumption of the theorem, the sequences $(P^n)_{n\in\mathbb{N}^*}$ and $(Q^n)_{n\in\mathbb{N}^*}$ are well defined. We start with P^1. As by assumption $\mu \ll \mu^{R^0} \ll \mu^R$, Corollary 2.9 shows that P^1 is well defined, and that for all i, j,

$$P^1_{ij} = \frac{\mu_i}{\mu^R_i} R_{ij},$$

with convention $\frac{0}{0} = 0$. Clearly, $R^0 \ll P^1$, as the support of the latter is

$$\Big\{(x_i, y_j) \in \mathcal{D} \times \mathcal{F} \text{ s.t. } R_{ij} > 0 \text{ and } \mu_i > 0\Big\} \supset \mathcal{E}.$$

Therefore, $\nu \ll \nu^{R^0} \ll \nu^{P^1}$. So once again, Corollary 2.9 shows that Q^1 is well defined, and that for all i, j,

$$Q^1_{ij} = \frac{\nu_j}{\nu^{P^1}_j} P^1_{ij},$$

with convention $\frac{0}{0} = 0$. The support of Q^1 is

$$\begin{aligned}\Big\{(x_i, y_j) \in \mathcal{D} \times \mathcal{F} \text{ s.t. } P^1_{ij} > 0 \text{ and } \nu_j > 0\Big\}\\ = \Big\{(x_i, y_j) \in \mathcal{D} \times \mathcal{F} \text{ s.t. } R_{ij} > 0 \text{ and } \mu_i > 0 \text{ and } \nu_j > 0\Big\} = \mathcal{E}.\end{aligned}$$

Then, a direct induction argument relying on the following formulas holding for all $n \in \mathbb{N}$ and all i, j:

$$P^{n+1}_{ij} = \frac{\mu_i}{\mu^{Q^n}_i} Q^n_{ij} \qquad \text{and} \qquad Q^{n+1}_{ij} = \frac{\nu_i}{\nu^{P^{n+1}}_j} P^{n+1}_{ij}, \tag{3.4}$$

with convention $\frac{0}{0} = 0$ show that P^n and Q^n are well defined, and that they admit $\mathcal{E}$ as their common support as soon as $n \geqslant 2$ and $n \geqslant 1$ respectively.

Let us now show that μ^* and ν^* are well defined. Their role are symmetric, so we just need to show that μ^* is well defined. First $Q^1 \in \Pi_2(\nu) \cap \mathcal{H}_+(R)$ and its support $\mathcal{E}$. We easily deduce that $\mu^{Q^1} \in \mathcal{H}_+(\mu)$, so that the closed convex set $\{\bar{\mu} = \mu^Q \text{ for some } Q \in \Pi_2(\nu) \cap \mathcal{H}_+(R)\}$ intersects $\mathcal{H}_+(\mu)$. By Proposition 2.5, the entropic projection μ^* of μ on this set is well defined.

Finally, let us show that P^* and Q^* are well defined. Once again, their role are symmetric, so we only show the existence of Q^*. By definition of μ^*, $\Pi(\mu^*, \nu) \cap \mathcal{H}_+(R) \neq \emptyset$. So by Proposition 2.5, the entropic projection Q^* of R on the closed convex set $\Pi(\mu^*, \nu)$ is well defined.

Step 2: A formula for $H(Q \,|\, R)$, for all $Q \in \Pi(\mu^*, \nu)$ with $H(Q \,|\, R) < +\infty$. — This contains the main new idea of this proof w.r.t. to prior works. The goal is to derive (3.11) below, a generalization of formula (3.3) which is

the main ingredient of the proof in [7]. We recall that a geometric interpretation of this formula is given in Appendix B.

In the classical proof for the scalable or approximately scalable case given in [7], at a given rank $n \geqslant 0$, formula (3.3) is an exact formula expressing as a sum of $2n$ nonnegative terms the difference between $H(\overline{R} \,|\, R)$ and $H(\overline{R} \,|\, Q^n)$, where $\overline{R}$ is any measure in $\Pi(\mu, \nu) \cap \mathcal{H}_+(R)$ (think of it as a potential limit for Q^n). This formula can be obtained thanks to the specific form of the density of Q^n w.r.t. R. The positivity of the terms in the sum ensures Q^n to get closer and closer to $\overline{R}$ in the entropic sense, and forces these terms to go to zero as $n \to +\infty$.

Here the idea is exactly the same, except that there is no $\overline{R} \in \Pi(\mu, \nu) \cap \mathcal{H}_+(R)$. So let us consider any $Q \in \Pi_2(\nu) \cap \mathcal{H}_+(R)$ (once again, we think of Q as a potential limit for Q^n so this condition is natural), compute $H(Q \,|\, R) - H(Q \,|\, Q^n)$, and find conditions on Q so that this difference can be written as a sum of nonnegative terms. This is where μ^* will naturally appear.

Let us fix $n \in \mathbb{N}^*$. Recalling $Q^0 = R$, we infer from (3.4) that for all $(x_i, y_j) \in \mathcal{E}$,

$$Q^n_{ij} = \frac{\nu_j}{\nu_j^{P^n}} \times \frac{\mu_i}{\mu_i^{Q^{n-1}}} \times \cdots \times \frac{\nu_j}{\nu^{P_j^1}} \times \frac{\mu_i}{\mu_i^{Q^0}} \times R_{ij}. \tag{3.5}$$

Observe that in the product in the r.h.s., because we assumed that $(x_i, y_j) \in \mathcal{E}$, the common support of all the iterates of the Sinkhorn algorithm, all the factors are positive.

We want to compute $H(Q \,|\, R) - H(Q \,|\, Q^n)$. For this quantity to be well defined, the support of Q needs to be included in $\mathcal{E}$. Let us assume this is true and pursue the argument. In that case, we have

$$\begin{aligned} &H(Q \,|\, R) - H(Q \,|\, Q^n) \\ &\quad = \sum_{i,j} \left\{ Q_{ij} \log \frac{Q_{ij}}{R_{ij}} + R_{ij} - Q_{ij} - Q_{ij} \log \frac{Q_{ij}}{Q^n_{ij}} - Q^n_{ij} + Q_{ij} \right\} \\ &\quad = \sum_{ij} Q_{ij} \log \frac{Q^n_{ij}}{R_{ij}} + \mathsf{M}(R) - \mathsf{M}(Q^n) \\ &\quad = \sum_{ij} Q_{ij} \log \frac{Q^n_{ij}}{R_{ij}} + \mathsf{M}(R) - \mathsf{M}(\nu) \end{aligned} \tag{3.6}$$

Now we can plug (3.5), use the additivity of the log and the property $\nu^Q = \nu$ to get

$$\begin{aligned}\sum_{i,j} Q_{ij} \log \frac{Q_{ij}}{R_{ij}} &= \sum_{k=1}^{n} \left\{ \sum_{i,j} Q_{ij} \log \frac{\nu_j}{\nu_j^{P^k}} + \sum_{i,j} Q_{ij} \log \frac{\mu_i}{\mu_i^{Q^{k-1}}} \right\} \\ &= \sum_{k=1}^{n} \left\{ \sum_{j} \nu_j \log \frac{\nu_j}{\nu_j^{P^k}} + \sum_{i} \mu_i^Q \log \frac{\mu_i}{\mu_i^{Q^{k-1}}} \right\}.\end{aligned} \tag{3.7}$$

For a given $k \geqslant 1$, being cautious with the different total masses, we have the two identities:

$$\begin{aligned}\sum_{j} \nu_j \log \frac{\nu_j}{\nu_j^{P^k}} &= H\big(\nu \,\big|\, \nu^{P^k}\big) + \mathsf{M}(\nu) - \mathsf{M}\big(\nu^{P^k}\big) \\ &= H\big(\nu \,\big|\, \nu^{P^k}\big) + \mathsf{M}(\nu) - \mathsf{M}(\mu),\end{aligned} \tag{3.8}$$

$$\begin{aligned}&\sum_{i} \mu_i^Q \log \frac{\mu_i}{\mu_i^{Q^{k-1}}} \\ &\quad= \sum_{i} \mu_i^Q \log \frac{\mu_i^Q}{\mu_i^{Q^{k-1}}} - \sum_{i} \mu_i^Q \log \frac{\mu_i^Q}{\mu_i} \\ &\quad= \begin{cases} H\big(\mu^Q \,\big|\, \mu^{Q^{k-1}}\big) - H\left(\mu^Q \,\middle|\, \mu\right) + \mathsf{M}(\nu) - \mathsf{M}(\mu), & \text{if } k \geqslant 2, \\ H\big(\mu^Q \,\big|\, R\big) - H\big(\mu^Q \,\big|\, \mu\big) + \mathsf{M}(\nu) - \mathsf{M}(R), & \text{if } k = 1. \end{cases}\end{aligned} \tag{3.9}$$

So gathering (3.6), (3.7), (3.8) and (3.9), we find

$$\begin{aligned}&H(Q \,|\, R) \\ &\quad= H(Q \,|\, Q^n) + \sum_{k=1}^{n} \Big\{ H\big(\nu \,\big|\, \nu^{P^k}\big) + H\big(\mu^Q \,\big|\, \mu^{Q^{k-1}}\big) - H\left(\mu^Q \,\middle|\, \mu\right) \Big\}.\end{aligned} \tag{3.10}$$

Consequently, we would have reached our goal provided we can show that for all $k \geqslant 1$,

$$H\big(\nu \,\big|\, \nu^{P^k}\big) + H\big(\mu^Q \,\big|\, \mu^{Q^{k-1}}\big) \geqslant H\big(\mu^Q \,\big|\, \mu\big).$$

If we see Q as a potential limit for Q^n, it is natural to expect $H(\mu^Q \,|\, \mu^{Q^{k-1}})$ to converge to 0 as $k \to +\infty$, so we do not expect this term to help, at least when k is large. Hence, we would need to find a condition on Q so that for all $k \geqslant 1$

$$H\big(\nu \,\big|\, \nu^{P^k}\big) \geqslant H\left(\mu^Q \,\middle|\, \mu\right).$$

To guarantee this inequality, it is therefore natural to ask $H(\mu^Q \,|\, \mu)$ to be as small as possible under the condition that $\nu^Q = \nu$, that is $\mu^Q = \mu^*$ as defined in (3.1).

So let $Q \in \Pi(\mu^*, \nu) \cap \mathcal{H}_+(R)$. On the one hand, as $\mu^* \in \mathcal{H}_+(\mu)$, the support of Q is included in $\mathcal{E}$. As a consequence, identity (3.10) applies, that is

$$H(Q \,|\, R) = H(Q \,|\, Q^n) + \sum_{k=1}^{n} \Big\{ H\big(\nu \,\big|\, \nu^{P^k}\big) - H(\mu^* \,|\, \mu) \Big\} + \sum_{k=1}^{n} H\big(\mu^* \,\big|\, \mu^{Q^{k-1}}\big). \quad (3.11)$$

On the other hand, every term in the first sum in the r.h.s. is nonnegative. Indeed, by Corollary 2.9, for all $k \geqslant 1$, we have $H(\nu \,|\, \nu^{P^k}) = H(Q^k \,|\, P^k)$. Moreover, still by Corollary 2.9,

$$H\left(Q^k \,\middle|\, P^k\right) \geqslant H\big(\mu^{Q^k} \,\big|\, \mu^{P^k}\big) = H\big(\mu^{Q^k} \,\big|\, \mu\big).$$

Finally, $Q^k \in \Pi_2(\nu) \cap \mathcal{H}_+(R)$ (use for instance (3.5) with $n = k$). So by optimality of μ^*, $H(\mu^{Q^k} \,|\, \mu) \geqslant H(\mu^* \,|\, \mu)$. Our claim follows, which concludes this step.

Step 3: Consequence of (3.11)*, convergence of the marginals.* — As a consequence of Step 2, both sums in the r.h.s. of (3.11) are bounded sums of nonnegative terms. Therefore, they converge as $n \to +\infty$, and their terms tend to 0 as $k \to +\infty$. We deduce in particular that

$$H\big(\mu^* \,\big|\, \mu^{Q^n}\big) \underset{n \to +\infty}{\longrightarrow} 0.$$

In particular, by continuity of H w.r.t. its second variable as stated in Proposition 2.2, and by compactness of $\{\bar{\mu} \in \mathcal{M}_+(\mathcal{D}) \text{ s.t. } \mathsf{M}(\bar{\mu}) = \mathsf{M}(\nu)\}$, $\mu^{Q^n} \to \mu^*$. So now let us pick $\overline{Q}$ any limit point of (Q^n). Such a limit point exist by compactness of $\{Q \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F}) \text{ s.t. } \mathsf{M}(Q) = \mathsf{M}(\nu)\}$. It follows from $\mu^{Q^n} \to \mu^*$ that $\mu^{\overline{Q}} = \mu^*$.

Step 4: $\overline{Q} = Q^*$. — Let us show that $\overline{Q} = Q^*$, so that actually the whole sequence (Q^n) converges towards Q^*. On the one hand, passing to the limit $n \to +\infty$ along the subsequences generating $\overline{Q}$ in (3.11) and using the continuity of H w.r.t. the second variable as stated in Proposition 2.2, we find

$$H(Q \,|\, R) = H(Q \,|\, \overline{Q}) + \sum_{k=1}^{+\infty} \Big\{ H\big(\nu \,\big|\, \nu^{P^k}\big) - H(\mu^* \,|\, \mu) \Big\} + \sum_{k=1}^{+\infty} H\big(\mu^* \,\big|\, \mu^{Q^{k-1}}\big). \quad (3.12)$$

On the other hand, as for all $n \in \mathbb{N}^*$, $Q^n \in \mathcal{H}_+(R)$, this is also true for $\overline{Q}$. Therefore, $\overline{Q} \in \Pi(\mu^*, \nu) \cap \mathcal{H}_+(R)$, we can apply (3.12) with $\overline{Q}$ in place of Q,

and find

$$H(\overline{Q}\,|\,R) = \sum_{k=1}^{+\infty}\Big\{H\big(\nu\,\big|\,\nu^{P^k}\big) - H(\mu^*\,|\,\mu)\Big\} + \sum_{k=1}^{+\infty} H\big(\mu^*\,\big|\,\mu^{Q^{k-1}}\big). \qquad (3.13)$$

Now it remains to apply (3.12) with $Q = Q^*$ and to plug the previous equality to find

$$H(Q^*\,|\,R) = H(Q^*\,|\,\overline{Q}) + H(\overline{Q}\,|\,R).$$

As by optimality of Q^*, $H(\overline{Q}\,|\,R) \geqslant H(Q^*\,|\,R)$, we can conclude that $H(Q^*\,|\,\overline{Q}) = 0$. Therefore, $\overline{Q} = Q^*$, as announced.

The proof of $P^n \to P^*$ follows along the same lines. □

As a free output of the proof of Theorem 1.3, we can show that we could have swapped μ and $\overline{\mu}$, and ν and $\overline{\nu}$ in the definitions (3.1) of μ^* and ν^* respectively. This is justified in the following remark.

Remark 3.4. — Observe the following optimization problem, where R, μ and ν are given, and where the competitor is $\overline{\nu}$:

$$\min\Big\{H(\nu\,|\,\overline{\nu})\;\Big|\;\overline{\nu} = \nu^P,\ \text{for some } P \in \Pi_1(\mu)\cap\mathcal{H}_+(R)\Big\}. \qquad (3.14)$$

This problem is almost the same as the one defining ν^* in (3.1), except from the fact that ν and $\overline{\nu}$ are swapped in the relative entropy. In this remark, we justify that the solution of this problem is ν^* as well, and that the corresponding optimal value is $H(\mu^*\,|\,\mu)$.

Provided there exists a competitor $\overline{\nu}$ for this problem with $H(\nu\,|\,\overline{\nu}) < +\infty$, we can find $P \in \Pi(\mu,\overline{\nu})\cap\mathcal{H}_+(R)$, and $Q \in \mathcal{M}_+(\mathcal{D}\times\mathcal{F})$ its entropic projection on $\Pi_2(\nu)$ as in Corollary 2.9. We have then $Q \in \Pi_2(\nu)\cap\mathcal{H}_+(R)$, and using Corollary 2.9 and the definition (3.1) of μ^*, we have

$$H(\nu\,|\,\overline{\nu}) = H(Q\,|\,P) \geqslant H\big(\mu^Q\,\big|\,\mu\big) \geqslant H(\mu^*\,|\,\mu).$$

On the other hand, as soon as the assumption of Theorem 3.2 holds, ν^* is a competitor for the problem in (3.14), and so in particular $H(\nu\,|\,\nu^*) \geqslant H(\mu^*\,|\,\mu)$. But because the terms of the first series in (3.12) tend to 0 and $\nu^{P^k} \to \nu^*$, we conclude that actually, $H(\nu\,|\,\nu^*) = H(\mu^*\,|\,\mu)$ and ν^* is a solution of (3.14). Finally, it is easy to see that a solution $\overline{\nu}$ of (3.14) must satisfy $\overline{\nu} \in \mathcal{H}_+(\nu)$ (because conditioning on the support of ν reduces the entropy), and by strict convexity of $\overline{\nu} \mapsto H(\nu\,|\,\overline{\nu})$ on the set $\mathcal{H}_+(\nu)$, under the assumption of Theorem 3.2, the problem (3.14) admits ν^* as its unique solution, so that (3.14) can be used as an alternative definition of ν^*.

Of course, we could argue in the same way to provide an alternative definition of μ^*, and we have the following equalities:

$$H(\nu\,|\,\nu^*) = H(\mu^*\,|\,\mu) \qquad \text{and} \qquad H(\mu\,|\,\mu^*) = H(\nu^*\,|\,\nu).$$

In particular, $\mu^* \sim \mu$ and $\nu^* \sim \nu$ in the sense of measures.

We also give another remark concerning the generalization of Theorem 3.2 to Polish spaces.

Remark 3.5. — We crucially use the fact that $\mathcal{D}$ and $\mathcal{F}$ are finite in this proof. Yet, the definitions of the Sinkhorn iterates, and of the limiting objects P^* and Q^* are perfectly valid, and we believe that the result is true when $\mathcal{D}$ and $\mathcal{F}$ are Polish spaces.

In that case, the beginning of the proof and the crucial formula (3.11) still hold, and consequently, the marginal convergence stated at Step 3 is still true. So the problem is to show the convergence of (P^n) and (Q^n) once it is known that their marginals converge.

The problem to adapt our proof is that (3.12) and (3.13) do not hold in general in that context. Indeed, as H is not more than lower semicontinuous w.r.t. the second variable, without new ideas, we can only prove

$$H(Q \,|\, R) \geqslant H(Q \,|\, \overline{Q}) + \sum_{k=1}^{+\infty} \Big\{ H\big(\nu \,\big|\, \nu^{P^k}\big) - H(\mu^* \,|\, \mu) \Big\} + \sum_{k=1}^{+\infty} H\big(\mu^* \,\big|\, \mu^{Q^{k-1}}\big)$$

in place of identity (3.12), and

$$H(\overline{Q} \,|\, R) \geqslant \sum_{k=1}^{+\infty} \Big\{ H\big(\nu \,\big|\, \nu^{P^k}\big) - H(\mu^* \,|\, \mu) \Big\} + \sum_{k=1}^{+\infty} H\big(\mu^* \,\big|\, \mu^{Q^{k-1}}\big) \qquad (3.15)$$

in place of (3.13). So we cannot conclude that

$$H(Q \,|\, R) \geqslant H(Q \,|\, \overline{Q}) + H(\overline{Q} \,|\, R),$$

and hence that $\overline{Q}$ is optimal. To conclude, we would need an equality sign in (3.15).

This difficulty already exists in the scalable case. As far as we know, it can be overcome following two approaches, which both fail in the approximately scalable or non-scalable case. The first one consists in proving the convergence of the potentials (a^n) and (b^n) from (1.3). This is for instance the approach developed in [27] or more recently in [9]. This approach has no chance to work as such, since these potentials *do not converge in general* in our setting, as explained for our toy example in Appendix A.

The second approach, developed for instance in [13], relies on the notion of cyclical invariance (see [13, Definition 1.1]). This property is a sufficient condition for a plan $\overline{R}$ to be the solution of $\mathrm{Sch}(R; \mu^{\overline{R}}, \nu^{\overline{R}})$, is true for all the iterates of the Sinkhorn algorithm, and passes to the limit $n \to \infty$ under reasonable assumptions including $R \sim \mu \otimes \nu$, the product measure of μ and ν. In our case, the limits P^* and Q^* *are not cyclically invariant in general* since

they are not equivalent to R. The reason is that we are precisely interested in the cases when $R \sim \mu \otimes \nu$ does not hold (otherwise, the problem would be scalable).

4. Γ-convergence in the marginal penalization problem

In this section, we want to show that when R, μ and ν are such that the Schrödinger problem $\mathrm{Sch}(R; \mu, \nu)$ has no solution, then the limit points P^* and Q^* given by Theorem 3.2 are relevant in view of the possible applications of the Sinkhorn algorithm.

To do so, let us think of R as an imperfect theoretical model describing the coupling between the initial and final positions of the particles of a large system. Also, let us imagine that μ and ν are data obtained by measuring the positions of the particles of the actual system that R is supposed to describe, at the initial and final time. In this situation, if $\mathrm{Sch}(R; \mu, \nu)$ has a solution R^*, this solution is interpreted as the model that is the closest to R that can explain the data.

However, even when R is a rather good model, and when μ and ν are rather precise measurements, it is possible that $\mathrm{Sch}(R; \mu, \nu)$ has no solution for several reasons:

- The first reason could be that our modeling does not take into account some physical phenomena. For instance, in Section 4.1, we will consider the case where the true system allows creation or annihilation of mass with very small probability, whereas the modeling does not.
- Another reason could be that μ and ν are only approximations of the real marginals. This can result from imprecise or biased measurements, or from a restricted amount of collected data. This will be considered in Section 4.2.

In both cases, it is very natural to relax the marginal constraints in (2.10) by introducing a fitting term in the value functional, that cancels when the constraints are satisfied, but which remains finite otherwise.

The main result of this section asserts that in these two situations, that are actually very close, the limit points P^* and Q^* of the Sinkhorn algorithm allow to compute the solution of the relaxed problem when the new fitting term takes the form of an entropy, in the limit where the level of marginal penalization tends to $+\infty$. The second case is a direct consequence of the first one, but that we wanted to keep separated because it does not have the same interpretation.

4.1. Unbalanced problems

In this subsection, we give ourselves $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ as before, and we study the following optimization problem, which is a reasonable modification of $\mathrm{Sch}(R; \mu, \nu)$ where the marginal constraints are replaced with marginal penalizations:

$$\min\Big\{ H(\overline{R} \,|\, R) + \lambda \left(H\big(\mu^{\overline{R}} \,\big|\, \mu\big) + H\big(\nu^{\overline{R}} \,\big|\, \nu\big) \right) \;\Big|\; \overline{R} \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F}) \Big\}, \quad (4.1)$$

where $\lambda > 0$ parametrizes the level of penalization.

This approach is extremely reminiscent of the idea introduced by Liero, Mielke and Savaré in [22] to deal with unbalanced data, that is, when $\mathsf{M}(\mu) \neq \mathsf{M}(\nu)$, in Optimal Transport problems. This was the starting point of the theory of *unbalanced Optimal Transport*, also discovered independently by other teams [6, 17].

More precisely, we will study the limit of the problem in (4.1) as $\lambda \to +\infty$. In this limit, it is actually more convenient to call $\varepsilon = 1/\lambda$ and to multiply the value functional by ε, to find the problem that we call $\mathrm{Sch}^\varepsilon(R; \mu, \nu)$:

$$\mathrm{Sch}^\varepsilon(R; \mu, \nu) := \min\Big\{ \varepsilon H(\overline{R} \,|\, R) + H\big(\mu^{\overline{R}} \,\big|\, \mu\big) + H\big(\nu^{\overline{R}} \,\big|\, \nu\big) \;\Big|\; \overline{R} \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F}) \Big\}.$$

As we want to study the behavior of this problem in the limit $\varepsilon \to 0$, we define the following functionals:

$$\Lambda^\varepsilon : \overline{R} \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F}) \mapsto \varepsilon H(\overline{R} \,|\, R) + H\big(\mu^{\overline{R}} \,\big|\, \mu\big) + H\big(\nu^{\overline{R}} \,\big|\, \nu\big),$$

$$\Lambda : \overline{R} \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F}) \mapsto \chi_{\mathcal{H}_+(R)}(\overline{R}) + H\big(\mu^{\overline{R}} \,\big|\, \mu\big) + H\big(\nu^{\overline{R}} \,\big|\, \nu\big),$$

where $\chi_{\mathcal{H}_+(R)}$ is the convex indicatrix taking value 0 on the set $\mathcal{H}_+(R)$ and $+\infty$ elsewhere.

The following proposition follows from standard arguments in the theory of Γ-convergence, see for instance [3, Theorem 1.47], and from the strict convexity of the relative entropy w.r.t. its first variable. We omit the proof.

Proposition 4.1. — *We have:*

$$\Gamma - \lim_{\varepsilon \to 0} \Lambda^\varepsilon = \Lambda.$$

Moreover, assuming that Λ is not uniformly infinite, let us call R_{opt} one of its minimizers, $\mu^g := \mu^{R_{\mathrm{opt}}}$ and $\nu^g := \nu^{R_{\mathrm{opt}}}$. The marginals μ^g and ν^g do not depend on the choice of R_{opt}, and as $\varepsilon \to 0$, the unique solution R^ε of $\mathrm{Sch}^\varepsilon(R; \mu, \nu)$ exists and converges towards the solution of $\mathrm{Sch}(R; \mu^g, \nu^g)$.

Remark 4.2. — In the notations μ^g and ν^g, the g stands for *geometric*. This is because as shown in Theorem 4.3, μ^g and ν^g are respectively the componentwise geometric means of μ and μ^*, and of ν and ν^*.

Therefore, studying the behavior of $\mathrm{Sch}^{\varepsilon}(R;\mu,\nu)$ in the limit $\varepsilon \to 0$ reduces to the study of the Schrödinger problem with modified marginals μ^g and ν^g. The following theorem shows the link between R^* – the solution of $\mathrm{Sch}(R;\mu^g,\nu^g)$ – on the one hand, and P^* and Q^* from Theorem 3.2 on the other hand.

THEOREM 4.3. — *Let $R \in \mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ satisfy Assumption 2.11. Then the functional Λ is not uniformly infinite. Moreover, considering P^* and Q^* as given by Theorem 3.2, and μ^g and ν^g as given by Proposition 4.1, the solution of* $\mathrm{Sch}(R;\mu^g,\nu^g)$ *is the componentwise geometric mean of P^* and Q^*, that is, the matrix R^* defined for all i,j by*

$$R^*_{ij} := \sqrt{P^*_{ij}Q^*_{ij}}. \tag{4.2}$$

Also, if μ^* and ν^* are defined by (3.1)*, μ^g and ν^g are the componentwise geometric means of μ^* and μ for the first one, and of ν^* and ν for the second one. In other terms, we have for all i,j,*

$$\mu^g_i = \sqrt{\mu^*_i\mu_i} \qquad \text{and} \qquad \nu^g_j = \sqrt{\nu^*_j\nu_j}. \tag{4.3}$$

Remark 4.4.

- Having in mind the approach of [22], we can give the following interpretation of the matrix R^*. In the degenerate case where the Schrödinger problem has no solution, it is necessary to allow creation and annihilation of mass to find solutions. Following [22], we can do this by replacing the balanced problem $\mathrm{Sch}(R;\mu,\nu)$ by the unbalanced problem $\mathrm{Sch}^{\varepsilon}(R;\mu,\nu)$. Following this analogy, $\lambda = \frac{1}{\varepsilon}$ parametrizes the cost of creating particles. The matrix R^* from Theorem 4.3 is therefore the limit of these solutions when the cost of creating or destroying matter tends to $+\infty$.
- A small adaptation of the proof shows that given $\alpha \in [0,1]$, if we replace the problem in (4.1) by

$$\min\Big\{H(\overline{R}\,|\,R) + \lambda\Big((1-\alpha)H\big(\mu^{\overline{R}}\,\big|\,\mu\big) + \alpha H\big(\nu^{\overline{R}}\,\big|\,\nu\big)\Big)\;\Big|\;\overline{R}\in\mathcal{M}_+(\mathcal{D}\times\mathcal{F})\Big\},$$

  and if we call $R^{\alpha,\lambda}$ its solution, then as $\lambda\to+\infty$, we have for all i,j:

$$R^{\alpha,\lambda}_{ij} \underset{\lambda\to+\infty}{\longrightarrow} \big(P^*_{ij}\big)^{1-\alpha}\big(Q^*_{ij}\big)^{\alpha}.$$

To prove this theorem, we will need to study carefully the optimality conditions for μ^* and ν^*. This could be done writing the Karush–Kuhn–Tucker conditions for the corresponding optimalization problems. We will rather adopt a more hand by hand approach, that is more likely to be generalizable in the continuous case. This is done in the following proposition.

Proposition 4.5. — *Assume that the conditions of Theorem 3.2 are fulfilled. Then P^* is the entropic projection of Q^* on $\Pi_1(\mu)$ and Q^* is the entropic projection of P^* on $\Pi_2(\nu)$, that is, for all i, j, we have*

$$P^*_{ij} = \frac{\mu_i}{\mu^*_i} Q^*_{ij} \qquad \text{and} \qquad Q^*_{ij} = \frac{\nu_j}{\nu^*_j} P^*_{ij}, \tag{4.4}$$

with convention $\frac{0}{0} = 0$. In particular, P^* and Q^* are equivalent, and we call $\mathcal{S}$ their common support, which is a subset of $\mathcal{E}$ defined in (2.11)*. Finally, we call for all i, j*

$$\varphi_i := \log \frac{\mu^*_i}{\mu_i} \qquad \text{and} \qquad \psi_j := \log \frac{\nu^*_j}{\nu_j}. \tag{4.5}$$

For all $(x_i, y_j) \in \mathcal{E}$, φ_i and ψ_j are well defined in $\mathbb{R}$, and:

$$\begin{cases} \varphi_i + \psi_j = 0, & \text{if } (x_i, y_j) \in \mathcal{S}, \\ \varphi_i + \psi_j \geqslant 0, & \text{if } (x_i, y_j) \in \mathcal{E}. \end{cases} \tag{4.6}$$

Proof of Proposition 4.5. — To get (4.4), it suffices to let n tend to $+\infty$ in (3.4). The fact that $\mathcal{S} \subset \mathcal{E}$ relies on the closed property of P^n and Q^n defined in (1.2) to have its support included in $\mathcal{E}$ for $n \geqslant 2$. If $(x_i, y_j) \in \mathcal{E}$, let us check that φ_i and ψ_j are well defined. On the one hand, by definition of $\mathcal{E}$, i is in the support of μ and j is in the support of ν. On the other hand, as observed in Remark 3.4, $\mu^* \sim \mu$ and $\nu^* \sim \nu$. Our claim follows.

Now, let $(x_i, y_j) \in \mathcal{S}$. A consequence of (4.4) is

$$P^*_{ij} = \frac{\mu^*_i}{\mu_i} \frac{\nu^*_j}{\nu_j} P^*_{ij} = \exp(\varphi_i + \psi_j) P^*_{ij}.$$

As (x_i, y_j) is in the support of P^* by definition of $\mathcal{S}$, we conclude that $\varphi_i + \psi_j = 0$.

Finally, it remains to prove that for all $(x_i, y_j) \in \mathcal{E}$, $\varphi_i + \psi_j \geqslant 0$. For this we use the optimality of $H(\mu^* \,|\, \mu) = H(\mu^{Q^*} \,|\, \mu)$ over all $Q \in \Pi_2(\nu)$. So let us take $(x_i, y_j) \in \mathcal{E}$. As $\nu_j > 0$, there exists i' such that $(i', j) \in \mathcal{S}$, that is, such that $Q^*_{i'j} > 0$. Let us define for $\varepsilon > 0$

$$Q^\varepsilon = Q^* + \varepsilon \delta_{ij} - \varepsilon \delta_{i'j},$$

where δ_{ij} is the matrix whose only nonzero coefficient is a one at position (i, j), and similarly for $\delta_{i'j}$. If ε is sufficiently small, Q^ε has nonnegative entries, and hence belongs to $\Pi_2(\nu)$, and with obvious notations, $\mu^{Q^\varepsilon} = \mu^* + \varepsilon \delta_i - \varepsilon \delta_{i'}$. Therefore, for such ε,

$$H\big(\mu^{Q^\varepsilon} \,\big|\, \mu\big) \geqslant H(\mu^* \,|\, \mu).$$

derivating to the right this inequality at $\varepsilon = 0$, we find

$$\log \frac{\mu_i^*}{\mu_i} - \log \frac{\mu_{i'}^*}{\mu_{i'}} \geqslant 0,$$

which rewrites $\varphi_i - \varphi_{i'} \geqslant 0$. But $(i', j) \in \mathcal{S}$ so $\varphi_{i'} = -\psi_j$, and so $\varphi_i + \psi_j \geqslant 0$. □

With this proposition at hand, we can prove Theorem 4.3.

Proof of Theorem 4.3. — The fact that under Assumption 2.11, Λ is not uniformly infinite follows from observing that $\Lambda(R^0) < +\infty$, where R_0 was defined Assumption 2.11. Now we reason in two steps. First we will prove using Proposition 4.5 that R^* defined by (4.2) is an optimizer of Λ, and then that it is the solution of the Schrödinger problem between its marginals.

Step 1: R^* is an optimizer of Λ. — To see that R^* is an optimizer of Λ, we first give a formula relating the vectors φ and ψ as defined by formula (4.5) and the marginals μ^{R^*} and ν^{R^*} of R^*. Using (4.4) and the definition (4.2) of R^*, we see that for all i, j,

$$R^*_{ij} = \sqrt{\frac{\nu_j}{\nu_j^*}} P^*_{ij} = \sqrt{\frac{\mu_i}{\mu_i^*}} Q^*_{ij}. \tag{4.7}$$

Summing respectively these identities w.r.t. i and j, we deduce that for all i, j, with convention $\frac{0}{0} = 0$,

$$\mu_i^{R^*} = \sqrt{\mu_i^* \mu_i} = \sqrt{\frac{\mu_i^*}{\mu_i}} \mu_i \qquad \text{and} \qquad \nu_j^{R^*} = \sqrt{\nu_j^* \nu_j} = \sqrt{\frac{\nu_j^*}{\nu_j}} \nu_j.$$

Let us define for all i, j:

$$Z_i^\mu := \log \frac{\mu_i^{R^*}}{\mu_i} = \frac{1}{2}\varphi_i \qquad \text{and} \qquad Z_j^\nu := \log \frac{\nu_j^{R^*}}{\nu_j} = \frac{1}{2}\psi_j.$$

Note that for all $(x_i, y_j) \in \mathcal{E}$, Z_i^μ and Z_j^ν are well defined in $\mathbb{R}$.

Now let $\overline{R}$ be such that $\Lambda(\overline{R}) < +\infty$. Using inequality (2.3) to bound from below each relative entropy, we have

$$\begin{aligned}
\Lambda(\overline{R}) &= H\big(\mu^{\overline{R}} \,\big|\, \mu\big) + H\big(\nu^{\overline{R}} \,\big|\, \nu\big) \\
&\geqslant \left\langle Z^\mu, \mu^{\overline{R}} \right\rangle - \left\langle e^{Z^\mu} - 1, \mu \right\rangle + \left\langle Z^\nu, \nu^{\overline{R}} \right\rangle - \left\langle e^{Z^\nu} - 1, \nu \right\rangle \\
&= \frac{1}{2} \left\langle \varphi, \mu^{\overline{R}} \right\rangle + \frac{1}{2} \left\langle \psi, \nu^{\overline{R}} \right\rangle - \sum_i \left\{ \mu_i^{R^*} - \mu_i \right\} - \sum_j \left\{ \nu_j^{R^*} - \nu_j \right\} \\
&= \frac{1}{2} \left\langle \varphi \oplus \psi, \overline{R} \right\rangle + \mathsf{M}(\mu) + \mathsf{M}(\nu) - 2\mathsf{M}(R^*),
\end{aligned}$$

where $\varphi \oplus \psi$ is the matrix defined for all i, j by $(\varphi \oplus \psi)_{ij} := \varphi_i + \psi_j$. Now, because of the second line of (4.6), as the support of $\overline{R}$ is easily seen to be a subset of $\mathcal{E}$, we get

$$\Lambda(\overline{R}) \geqslant \mathsf{M}(\mu) + \mathsf{M}(\nu) - 2\mathsf{M}(R^*).$$

On the other hand, by definition of Z^μ and Z^ν,

$$\begin{aligned}\Lambda(R^*) &= H\big(\mu^{R^*} \,\big|\, \mu\big) + H\big(\nu^{R^*} \,\big|\, \nu\big) \\ &= \left\langle Z^\mu, \mu^{R^*} \right\rangle + \mathsf{M}(\mu) - \mathsf{M}(R^*) + \left\langle Z^\nu, \nu^{R^*} \right\rangle + \mathsf{M}(\nu) - \mathsf{M}(R^*) \\ &= \frac{1}{2}\langle \varphi \oplus \psi, R^* \rangle + \mathsf{M}(\mu) + \mathsf{M}(\nu) - 2\mathsf{M}(R^*).\end{aligned}$$

But now, as the support of R^* is precisely $\mathcal{S}$, by the first line of (4.6), we get

$$\Lambda(R^*) = \mathsf{M}(\mu) + \mathsf{M}(\nu) - 2\mathsf{M}(R^*).$$

We deduce that $\Lambda(\overline{R}) \geqslant \Lambda(R^*)$ and R^* is indeed an optimizer of Λ. In particular, $\mu^g = \mu^{R^*}$ and $\nu^g = \nu^{R^*}$, which proves (4.3).

Step 2: R^* is the solution of $\mathrm{Sch}(R; \mu^g, \nu^g)$. — This part of the proof relies on very classical arguments. On the one hand, because of (4.7) and Theorem 1.3, R^* is the solution of $\mathrm{Sch}(P^*; \mu^g, \nu^g)$. On the other hand, as P^* is the solution of $\mathrm{Sch}(R; \mu, \nu^*)$, Lemma 4.6 below applies and R^* is the solution of $\mathrm{Sch}(R; \mu^g, \nu^g)$. □

Lemma 4.6. — *Let $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, $\mu, \mu' \in \mathcal{M}_+(\mathcal{D})$ and $\nu, \nu' \in \mathcal{M}_+(\mathcal{F})$. Assume that* $\mathrm{Sch}(R; \mu, \nu)$ *admits a solution P and that* $\mathrm{Sch}(P; \mu', \nu')$ *admits a solution Q. Then the unique solution of* $\mathrm{Sch}(R; \mu', \nu')$ *exists: it is Q.*

Remark 4.7. — Note that with the notations of the statement, if there exists $a, a' \in \mathscr{F}(\mathcal{D}, \mathbb{R}_+)$ and $b, b' \in \mathscr{F}(\mathcal{F}, \mathbb{R}_+)$ such that for all i, j, $P_{ij} = a_j b_j R_{ij}$ and $Q_{ij} = a'_i b'_j P_{ij}$, then we have $Q_{ij} = (a_i a'_i)(b_j b'_j) R_{ij}$ and the result follows from Theorem 1.3. This lemma is therefore an obvious result in the scalable case that still holds in the approximately scalable one. In the approximately scalable case, its proof relies on the Pythagorean law of Theorem 2.7.

Proof. — First, $\mathrm{Sch}(R; \mu', \nu')$ admits a solution. Indeed, $Q \in \mathcal{H}_+(P) \subset \mathcal{H}_+(R)$, so it is a competitor for $\mathrm{Sch}(R; \mu', \nu')$. Let us call Q' this solution, and $V := Q' - Q \in \mathcal{M}(\mathcal{D} \times \mathcal{F})$. We easily check that $V \in \mathcal{H}(R)$. As $Q \in \mathcal{H}_+(P)$, for $\varepsilon > 0$ small enough, $P + \varepsilon V$ has nonnegative entry, and therefore belongs to $\mathcal{H}_+(R)$ and $\Pi(\mu, \nu)$. Therefore, as P is the solution of $\mathrm{Sch}(R; \mu, \nu)$, we have by point 2. of Theorem 2.7 that $V \in \mathcal{H}(P)$ (whence $Q' \in \mathcal{H}_+(P)$), and

$$\sum_{i,j} Q'_{ij} \log \frac{P_{ij}}{R_{ij}} \geqslant \sum_{i,j} Q_{ij} \log \frac{P_{ij}}{R_{ij}}.$$

It is easy to deduce from this identity that

$$H(Q' \,|\, R) - H(Q' \,|\, P) \geqslant H(Q \,|\, R) - H(Q \,|\, P).$$

But Q is optimal for $\mathrm{Sch}(P; \mu', \nu')$, so $H(Q \,|\, P) \leqslant H(Q' \,|\, P)$, and Q' is optimal for $\mathrm{Sch}(R; \mu', \nu')$ so $H(Q' \,|\, R) \leqslant H(Q \,|\, R)$. So actually, these inequalities are equalities and $Q = Q'$. □

4.2. Balanced version

In the last subsection, we interpreted the fact that $\mathrm{Sch}(R; \mu, \nu)$ has no solution by the fact that our model does not incorporate the ability of the real system to create or destroy mass. In that case, the total mass of R^* is not the same as the one of μ and ν in general, even when the latter two coincide. Therefore, R^* cannot be interpreted directly as a joint law for the initial and final positions of the particles. Following the lines of [22], we see that its interpretation is actually rather complicated.

In this subsection, we want to consider the case where the real system under study is truly balanced, that is, no creation of annihilation of mass is possible at all. In this situation, whatever the way we are obtaining the data, μ and ν must have the same mass, and up to renormalizing, we can assume that they are probability measures. We want to interpret the fact that $\mathrm{Sch}(R; \mu, \nu)$ has no solution by the fact that μ and ν are imperfect measurements of the true marginals, and we want to find a *probability* measure $\overline{R}^*$ that is entropically close to R while having its marginals entropically close to μ and ν, that can be interpreted as a joint law.

Therefore, we introduce the following problem that is a slight modification of Sch^ε where the competitor $\overline{R}$ needs to be a probability measure: for all $R \in \mathcal{P}(\mathcal{D} \times \mathcal{F})$, $\mu \in \mathcal{P}(\mathcal{D})$ and $\nu \in \mathcal{P}(\mathcal{F})$,

$$\overline{\mathrm{Sch}}^\varepsilon(R; \mu, \nu) := \min\Big\{ \varepsilon H(\overline{R} \,|\, R) + H\big(\mu^{\overline{R}} \,\big|\, \mu\big) + H\big(\nu^{\overline{R}} \,\big|\, \nu\big) \;\Big|\; \overline{R} \in \mathcal{P}(\mathcal{D} \times \mathcal{F}) \Big\}.$$

The following theorem states the behaviour of this optimization problem as $\varepsilon \to 0$, and is a direct adaptation of Theorem 4.3 to the balanced case.

Theorem 4.8. — *Let $R \in \mathcal{P}(\mathcal{D} \times \mathcal{F})$, $\mu \in \mathcal{P}(\mathcal{D})$ and $\nu \in \mathcal{P}(\mathcal{F})$ satisfy the conditions of Assumption 2.11, and call*

$$\mathcal{Z} := \sum_{ij} \sqrt{P^*_{ij} Q^*_{ij}},$$

where P^* and Q^* are given by Theorem 3.2. Then for all $\varepsilon > 0$, the solution $\overline{R}^\varepsilon$ of $\overline{\mathrm{Sch}}^\varepsilon(R;\mu,\nu)$ exists, is unique, and satisfies for all i,j:

$$\overline{R}^\varepsilon_{ij} \underset{\varepsilon\to 0}{\longrightarrow} \overline{R}^*_{ij} := \frac{\sqrt{P^*_{ij}Q^*_{ij}}}{\mathcal{Z}}.$$

Its marginals are given for all i,j by

$$\mu_i^{\overline{R}^*} = \frac{\sqrt{\mu_i^*\mu_i}}{\mathcal{Z}} \qquad \text{and} \qquad \nu_j^{\overline{R}^*} = \frac{\sqrt{\nu_j^*\nu_j}}{\mathcal{Z}}.$$

Proof. — Theorem 4.8 is a direct consequence of Theorem 4.3 once noticed the following fact. If R,μ,ν are as in the statement of the theorem, if $\varepsilon > 0$ and if R^ε is the solution of $\mathrm{Sch}^\varepsilon(R;\mu,\nu)$, then $R^\varepsilon/\mathsf{M}(R^\varepsilon)$ is the solution of $\overline{\mathrm{Sch}}^\varepsilon(R;\mu,\nu)$. To see this, consider $R' \in \mathcal{P}(\mathcal{D}\times\mathcal{F})$. Direct computations imply

$$H(R' \,|\, R) = \frac{H\left(\mathsf{M}(R^\varepsilon)R' \,|\, R\right)}{\mathsf{M}(R^\varepsilon)} + \log\frac{1}{\mathsf{M}(R^\varepsilon)} + 1 - \frac{1}{\mathsf{M}(R^\varepsilon)},$$
$$H\left(\frac{R^\varepsilon}{\mathsf{M}(R^\varepsilon)}\,\middle|\, R\right) = \frac{H(R^\varepsilon \,|\, R)}{\mathsf{M}(R^\varepsilon)} + \log\frac{1}{\mathsf{M}(R^\varepsilon)} + 1 - \frac{1}{\mathsf{M}(R^\varepsilon)}.$$

By optimality of R^ε, $H(\mathsf{M}(R^\varepsilon)R' \,|\, R) \geqslant H(R^\varepsilon \,|\, R)$, and therefore $H(R' \,|\, R) \geqslant H(R^\varepsilon/\mathsf{M}(R^\varepsilon) \,|\, R)$. Our claims follows, and hence the theorem as M is a continuous functional and $\mathcal{Z} = \mathsf{M}(R^*)$, where R^* is given by Theorem 4.3. □

5. Existence and support of the solutions of Schrödinger problems

In this section, our goal is to give a detailed study of the support of the solution of $\mathrm{Sch}(R;\mu,\nu)$ when the latter exists, or of the common one of P^*, Q^* and R^* from Theorems 3.2 and 4.3 in the non-scalable case. Our motivation for doing so has been explained in the "Contribution" and "Prior works and difficulties" parts of the Introduction. This study will rely on a new interpretation of the well known existence conditions for the Schrödinger problem in finite spaces, for which we refer to [4, 16].

We start with our new formulation of these conditions of existence, which is very close to the ones introduced by Brualdi [4], but has the advantage of helping understanding the shape of the support of the optimizers seen as a bipartite graph.

In the second part of the section, we provide a theoretical procedure allowing to get the support of the optimizers, both in the approximately scalable and non-scalable cases, without using the Sinkhorn algorithm. This

procedure will be used in the next section as a preliminary step, before launching the Sinkhorn algorithm, in order to recover a linear rate for the latter.

5.1. A necessary and sufficient condition of existence for the Schrödinger problem in finite spaces

Let us state a necessary and sufficient condition on R, μ and ν for the existence of a solution R^* of $\mathrm{Sch}(R;\mu,\nu)$, that is, for $\mathrm{Sch}(R;\mu,\nu)$ to be scalable or approximately scalable. In order to do so, we need to give a few definitions. First, we endow the set $\mathcal{D}\cup\mathcal{F}$ with a bipartite graph structure related to R: we set

$$\forall\, i=1,\dots,N \text{ and } j=1,\dots,M, \qquad x_i \triangle y_j \iff R_{ij}>0.$$

We have $x_i \triangle y_j$ whenever it is possible to travel from x_i to y_j under R. We write indifferently $x_i \triangle y_j$ or $y_j \triangle x_i$.

With this structure in hand, we are able to push forward or pull backward subsets of $\mathcal{D}$ and $\mathcal{F}$, that is, we define:

$$\begin{aligned} &\forall\, A\subset\mathcal{D}, \quad F_R(A) := \big\{y\in\mathcal{F}\;\big|\;\exists\, x\in A \text{ s.t. } x\triangle y\big\},\\ &\forall\, B\subset\mathcal{F}, \quad D_R(B) := \big\{x\in\mathcal{D}\;\big|\;\exists\, y\in B \text{ s.t. } x\triangle y\big\}. \end{aligned} \tag{5.1}$$

Heuristically, for all $A\subset\mathcal{D}$, $F_R(A)$ is the set of all possible final positions of particles starting from A, under R. Correspondingly, for all $B\subset\mathcal{F}$, $D_R(B)$ is the set of all possible initial positions of particles arriving in B under R. Notice the explicit mention of R in the notations: in the following, we will allow ourselves to replace R by any other measure $\overline{R}\in\mathcal{M}_+(\mathcal{D}\times\mathcal{F})$.

The main result of this section is the following.

Theorem 5.1. — *Let $R\in\mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, $\mu\in\mathcal{M}_+(\mathcal{D})$ and $\nu\in\mathcal{M}_+(\mathcal{F})$. The three following assertions are equivalent:*

(a) $\mathsf{M}(\mu)=\mathsf{M}(\nu)$ *and for all* $A\subset\mathcal{D}$, $\mu(A)\leqslant\nu(F_R(A))$.
(b) $\mathsf{M}(\mu)=\mathsf{M}(\nu)$ *and for all* $B\subset\mathcal{F}$, $\nu(B)\leqslant\mu(D_R(B))$.
(c) $\mathrm{Sch}(R;\mu,\nu)$ *is scalable or approximately scalable.*

Note that the implications (c) $\Rightarrow$ (a) and (c) $\Rightarrow$ (b) are straightforward, and that only the reverse implications are challenging. Also, we already noticed in Section 2.3 that (c) implies Assumption 2.11. Hence, it is also the case for (a) and (b).

The proof relies on the following Lemma 5.2, which gives a necessary and sufficient condition on R, μ and ν ensuring R^* to have the same support as

R, that is, to be in the scalable case. In this statement, we use the notations μ^R and ν^R as defined in (2.1), and we work under Assumption 2.12, which is always possible under Assumption 2.11 up to considering subspaces of $\mathcal{D}$ and $\mathcal{F}$, see Section 2.3.

LEMMA 5.2. — *Let $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$, satisfying Assumption 2.12. The three following assertions are equivalent:*

- (a′) $\mathsf{M}(\mu) = \mathsf{M}(\nu)$ *and for all $A \subset \mathcal{D}$, $\mu(A) \leqslant \nu(F_R(A))$, with a strict inequality whenever $\mu^R(A) < \nu^R(F_R(A))$.*
- (b′) $\mathsf{M}(\mu) = \mathsf{M}(\nu)$ *and for all $B \subset \mathcal{F}$, $\nu(B) \leqslant \mu(D_R(B))$, with a strict inequality whenever $\nu^R(B) < \mu^R(D_R(B))$.*
- (c′) $\mathrm{Sch}(R; \mu, \nu)$ *is scalable.*

In plain words, it highlights the difference between the approximately scalable and scalable cases, by showing that the scalable case consists in assuming as much strict inequalities in (a) or in (b) as possible. Although both Theorem 5.1 and Lemma 5.2 can be directly deduced from the work of Brualdi [4], we provide in Appendix C a short and independent proof based on topological arguments.

5.2. Theoretical construction of the support

In the scalable case, the Sinkhorn algorithm is known to have a linear rate of convergence. On the other hand, in the approximately scalable case, the algorithm still converges, but the (unknown) convergence rate cannot be linear [1].

In this subsection, we study the support of the solution of the Schrödinger problem in the approximately scalable and non-scalable cases for the following reason. Take R, μ and ν such that $\mathrm{Sch}(R; \mu, \nu)$ is approximately scalable, R^* the solution of this problem, and $\mathcal{S}$ the support of R^*. Then $\mathrm{Sch}(\mathbb{1}_{\mathcal{S}} R; \mu, \nu)$ is scalable and its solution is R^*. In particular, the Sinkhorn algorithm applied to this problem has a linear rate of convergence. Interestingly, a similar reasoning is valid in the non-scalable case, as we show in Proposition 5.3 below.

Without loss of generality, and for the sake of simplicity, in the whole subsection, we work under Assumption 2.12. By Remark 3.4, if μ^* and ν^* are defined by (3.1), we have $\nu \sim \nu^*$ and $\mu \sim \mu^*$. So under Assumption 2.12, they have a full support as well.

PROPOSITION 5.3. — *Let $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ satisfying Assumption 2.12. Let us call $\mathcal{S}$ the common support of*

P^* and Q^* from Theorem 3.2, and R^* from Theorem 4.3. Let $(P^n)_{n\in\mathbb{N}^*}$ and $(Q^n)_{n\in\mathbb{N}^*}$ be given by (1.3) *applied to* $\mathrm{Sch}(\mathbb{1}_{\mathcal{S}}R;\mu,\nu)$. *They converge respectively towards P^* and Q^*, both of them at a linear rate.*

Proof. — Let $(P^n)_{n\in\mathbb{N}^*}$ and $(Q^n)_{n\in\mathbb{N}^*}$ be given by the equivalent formulations (1.2) and (1.3) applied to $\mathrm{Sch}(\mathbb{1}_{\mathcal{S}}R;\mu,\nu)$. Let us show that (P^n) converges towards P^* at a linear rate. The case of (Q^n) follows the same arguments. The idea is that if $(\widetilde{P}^n)_{n\in\mathbb{N}^*}$ and $(\widetilde{Q}^n)_{n\in\mathbb{N}^*}$ are given by (1.2) and (1.3) applied to $\mathrm{Sch}(\mathbb{1}_{\mathcal{S}}R;\mu,\nu^*)$, then for all $n\in\mathbb{N}^*$, $P^n=\widetilde{P}^n$. As the problem $\mathrm{Sch}(\mathbb{1}_{\mathcal{S}}R;\mu,\nu^*)$ is scalable (its solution, P^*, has the same support as $\mathbb{1}_{\mathcal{S}}R$), by [36], the rate of convergence of $(\widetilde{P}^n)$ towards P^* is linear, and the result follows.

So let us prove by induction that for all $n\in\mathbb{N}^*$, $P_n=\widetilde{P}_n$. According to (1.2), P^1 and $\widetilde{P}^1$ are solutions to the same problem, and therefore coincide. Let us now consider $n\in\mathbb{N}^*$ such that $P^n=\widetilde{P}^n$ and show that $P^{n+1}=\widetilde{P}^{n+1}$. By construction, the support of P^{n+1} and $\widetilde{P}^{n+1}$ is $\mathcal{S}$, so we just need to check that for all $(x_j,y_j)\in\mathcal{S}$, $P^{n+1}_{ij}=\widetilde{P}^{n+1}_{ij}$. By the first line of (4.6), for all $(x_i,y_j)\in\mathcal{S}$, we have:

$$\nu_j=\frac{\mu_i^*\nu_j^*}{\mu_i}.$$

Hence, for all $(x_i,y_j)\in\mathcal{S}$:

$$\begin{aligned}
P^{n+1}_{ij}=\frac{\mu_i}{\mu_i^{Q^n}}Q^n_{ij}&=\frac{\mu_i}{\sum_{j'}\frac{\nu_{j'}}{\nu_{j'}^{P^n}}P^n_{ij'}}\times\frac{\nu_j}{\nu_j^{P^n}}P^n_{ij}\\
&=\frac{\mu_i}{\frac{\mu_i^*}{\mu_i}\sum_{j'}\frac{\nu^*_{j'}}{\nu_{j'}^{\widetilde{P}^n}}\widetilde{P}^n_{ij'}}\times\frac{\mu_i^*}{\mu_i}\frac{\nu_j^*}{\nu_j^{\widetilde{P}^n}}\widetilde{P}^n_{ij}\\
&=\frac{\mu_i}{\mu_i^{\widetilde{Q}^n}}\widetilde{Q}^n_{ij}=\widetilde{P}^{n+1}_{ij},
\end{aligned}$$

where the change from P^n to $\widetilde{P}^n$ in the middle coming from the induction assumption $P^n=\widetilde{P}^n$. The result follows. □

Therefore, even in the non-scalable case, a way to improve the Sinkhorn algorithm consists in first finding $\mathcal{S}$, and then computing the solution of a scalable problem. We propose in this subsection a theoretical procedure allowing to get this support without using the Sinkhorn algorithm in both the approximately and non-scalable cases, and we will propose an approximate method for achieving this task numerically at Section 6.

To detail our procedure, we introduce a class of subsets of $\mathcal{D}$ associated with a triple $(R;\mu,\nu)$.

Definition 5.4. — *Let $R \in \mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ satisfying Assumption 2.12. Let us consider R^* from Theorem 4.3. We say that a subset $A \subset \mathcal{D}$ is the* source of an isolated scalable problem *(or for short that A is a SISP set) for $(R;\mu,\nu)$ if $A \neq \emptyset$ and:*

- *The set $(\mathcal{D}\backslash A) \times F_R(A)$ is R^*-negligible, i.e.*
$$R^*\Big((\mathcal{D}\backslash A) \times F_R(A)\Big) = 0. \tag{5.2}$$
- *For all $x_i \in A$ and $y_j \in \mathcal{F}$,*
$$R^*_{ij} > 0 \quad \Longleftrightarrow \quad R_{ij} > 0. \tag{5.3}$$

We show at Figure 5.1 an illustration of what a SISP is.

Of course, as P^* and Q^* from Theorem 3.2 are equivalent to R^* in the sense of measures, we could have replaced R^* in the previous definition by one of them.

$$R = \left[\begin{array}{c|c} \mathfrak{A} & \mathfrak{B} \\ \hline 0 & \mathfrak{C} \end{array}\right] \begin{matrix} \} A \\ \\ \end{matrix}$$
$$\underbrace{\qquad\qquad}_{F_R(A)}$$

Figure 5.1. If $A \subset \mathcal{D}$, up to reordering the lines, we can assume that it corresponds to the last lines. Up to reordering the columns, we can assume that $F_R(A)$ corresponds to the last columns. Then, R has the form given in the picture. In this situation, A is a SISP set for $(R;\mu,\nu)$ if R^* cancels on the block $\mathfrak{B}$ and if the supports of R and R^* coincide on the block $\mathfrak{C}$.

On the one hand, SISP sets always exist, at least under Assumption 2.12, as announced in the following lemma. Its proof is our main task in this part of our work, and is given at the end of the subsection.

Lemma 5.5. — *Let $R \in \mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ satisfying Assumption 2.12. Then there exists a SISP set for $(R;\mu,\nu)$.*

On the other hand, once we know how to find SISP sets, an iterative procedure consisting in finding SISP sets for a sequence of more and more restricted problems makes is possible to reconstruct the whole subset $\mathcal{S}$.

Proposition 5.6. — *Let $R \in \mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ satisfy Assumption 2.12. Let us call $\mathcal{S}$ the common support of P^* and Q^* from Theorem 3.2, and R^* from Theorem 4.3.*

We define by inference $(R^n)_{n\in\mathbb{N}}$ *a sequence in* $\mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, $(\mathcal{D}^n)_{n\in\mathbb{N}}$ *a nonincreasing sequence of subsets of* $\mathcal{D}$ *and* $(\mathcal{F}^n)_{n\in\mathbb{N}}$ *a nonincreasing sequence of subsets of* $\mathcal{F}$ *in the following way:*

- *For* $n=0$, *we set* $R^0 := R$, $\mathcal{D}^0 := \mathcal{D}$ *and* $\mathcal{F}^0 := \mathcal{F}$;
- *For all* $n\in\mathbb{N}$, *if* $\mathcal{D}^n$ *and* $\mathcal{F}^n$ *are nonempty and*

$$(R^n{}_{\llcorner \mathcal{D}^n\times\mathcal{F}^n}; \mu_{\llcorner\mathcal{D}^n}, \nu_{\llcorner\mathcal{F}^n})$$

satisfies Assumption 2.12, we pick M_n *a SISP set as given by Lemma 5.5, and we set:*

$$\mathcal{D}^{n+1} := \mathcal{D}^n\backslash M_n, \qquad \mathcal{F}^{n+1} := \mathcal{F}^n\backslash F_{R^n{}_{\llcorner\mathcal{D}^n\times\mathcal{F}^n}}(M_n),$$

$$\forall\, i,j, \quad R_{ij}^{n+1} := \begin{cases} 0, & \text{if } y_j \in F_{R^n{}_{\llcorner\mathcal{D}^n\times\mathcal{F}^n}}(M_n) \text{ and } x_i\in\mathcal{D}^{n+1},\\ R_{ij}^n, & \text{otherwise.}\end{cases}$$

Otherwise, we set $R^{n+1} := R^n$, $\mathcal{D}^{n+1} := \mathcal{D}^n$ *and* $\mathcal{F}^{n+1} := \mathcal{F}^n$.

With this construction, the sequence $(R^n, \mathcal{D}^n, \mathcal{F}^n)_{n\in\mathbb{N}}$ *is stationary. More precisely, there exists* $N\in\mathbb{N}^*$ *such that for all* $n\geqslant N$,

$$\mathcal{D}^n = \emptyset, \qquad \mathcal{F}^n = \emptyset, \qquad R^n = \mathbb{1}_{\mathcal{S}} R.$$

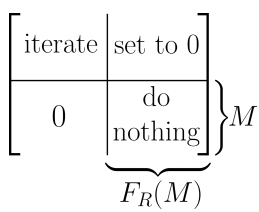

Figure 5.2. In the situation of Figure 5.1 where we have reordered the lines and columns, our procedure consists in recursively add zeros to R at positions where we know thanks to (5.2) that R^* admits a zero. We know that we did not forget any zero in $M\times F_R(M)$ thanks to (5.8).

An illustration of the procedure at each iteration, is provided in Figure 5.2. An illustration of the full procedure in a specific non-scalable case is provided in Figure 5.3.

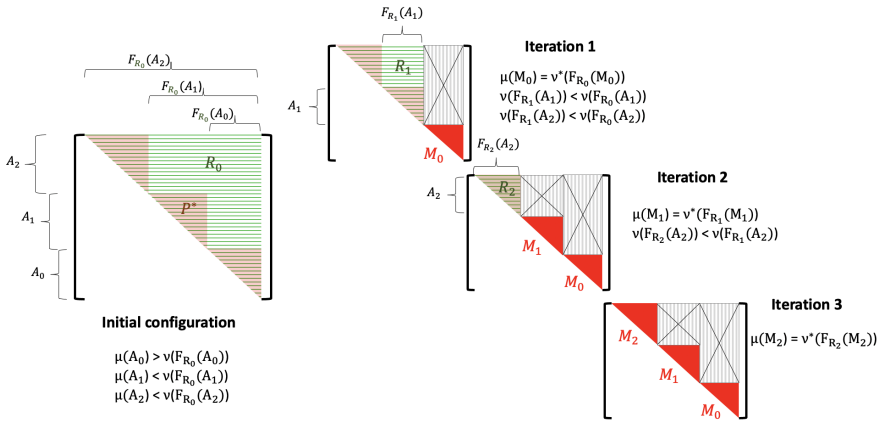

Figure 5.3. Illustration of the procedure of Proposition 5.6 when the matrix R is upper diagonal and R^* a staircase matrix (see Appendix A for more details). Only A_0 is a SISP set for $(R;\mu,\nu)$, but A_1 and A_2 are SISP sets for the restricted problems at the iterations 2 and 3. In this example, the procedure is stationary after 3 steps, and the SISP set at each iteration is the unique maximal θ-set (as defined in Definition 5.8) for the reduced problem. We remark that we can also build ν^*, the second marginal of P^* defined in Theorem 3.2, on the successive SISP sets obtained along the procedure, thanks to the second step of the proof of Lemma 5.5 which ensures that the ratio $\frac{\nu}{\nu^*}$ is constant inside the maximal θ-sets.

Proof. — In this proof, in order to lighten the notations, we call $R_r^n := R^n{}_{\llcorner \mathcal{D}^n \times \mathcal{F}^n}$. We will prove by inference the following facts. For all $n \in \mathbb{N}$:

(1) Calling $\mathcal{S}^n$ the support of R^n, and therefore $\mathcal{S}^0$ the support of R, we have $\mathcal{S} \subset \mathcal{S}^n$, $\mathcal{S}^n \cap (\mathcal{D}^n \times \mathcal{F}^n) = \mathcal{S}^0 \cap (\mathcal{D}^n \times \mathcal{F}^n)$, and

$$\begin{aligned} \mathcal{S}^n \cap \Big((\mathcal{D} \times \mathcal{F}) \backslash (\mathcal{D}^n \times \mathcal{F}^n)\Big) &= \mathcal{S} \cap \Big((\mathcal{D} \times \mathcal{F}) \backslash (\mathcal{D}^n \times \mathcal{F}^n)\Big) \\ &= \mathcal{S}^0 \cap \left(\bigcup_{k=0}^{n-1} M_k \times F_{R_r^k}(M_k) \right). \end{aligned} \tag{5.4}$$

(2) $\mathcal{D}^n$ is empty if and only if $\mathcal{F}^n$ is empty.

(3) If $\mathcal{D}^n$ and $\mathcal{F}^n$ are not empty, $(R_r^n; \mu_{\llcorner \mathcal{D}^n}, \nu_{\llcorner \mathcal{F}^n})$ satisfies Assumption 2.12 and the matrices $P^{n,*}$, $Q^{n,*}$ and $R^{n,*}$ associated with $(R_r^n; \mu_{\llcorner \mathcal{D}^n}, \nu_{\llcorner \mathcal{F}^n})$ through Theorems 3.2 and 4.3 are the restrictions of P^*, Q^* and R^* to $\mathcal{D}^n \times \mathcal{F}^n$.

This is enough to prove the proposition: if the conclusion of the inference is true, then by the third point and Lemma 5.5, as long as $\mathcal{D}^n$ and $\mathcal{F}^n$ are nonempty, $(R_r^n; \mu_{\llcorner \mathcal{D}^n}, \nu_{\llcorner \mathcal{F}^n})$ admits a SISP set M_n, which is not empty by definition. Therefore, $(\mathcal{D}^n)$ is strictly decreasing in the sense of inclusion as long as it is not empty, so it has to reach $\emptyset$ at a certain rank N. At this rank, because of the first point, we also have $\mathcal{F}^N = \emptyset$, and because of (5.4), $\mathcal{S}^N = \mathcal{S}$, so that the conclusion follows. So let us prove the inference.

At rank 0, everything is clear, so let us assume that the conclusions of points one, two and three hold at rank n, and prove them at rank $n+1$. First, if $\mathcal{D}^n$ is empty, by assumption $\mathcal{F}^n$ is empty as well, so we have reached a stationary point, and everything is still true at rank $n+1$. So we can assume without loss of generality that $\mathcal{D}^n$, and hence $\mathcal{F}^n$, are nonempty. By assumption, $(R_r^n; \mu_{\llcorner \mathcal{D}^n}, \nu_{\llcorner \mathcal{F}^n})$ satisfies Assumption 2.12, and by Lemma 5.5, we can find a SISP set M_n. In this context, let us check the points one by one at rank $n+1$.

First point. — Observing that $\mathcal{D}^n$ is the disjoint union of M_n and $\mathcal{D}^{n+1}$, and that $\mathcal{F}^n$ is the disjoint union of $F_{R_r^n}(M_n)$ and $\mathcal{F}^{n+1}$, we have

$$\begin{aligned}(\mathcal{D}\times\mathcal{F})\backslash\big(\mathcal{D}^{n+1}\times\mathcal{F}^{n+1}\big)&\\ =\Big((\mathcal{D}\times\mathcal{F})\backslash(\mathcal{D}^n\times\mathcal{F}^n)\Big)&\cup\big(\mathcal{D}^{n+1}\times F_{R_r^n}(M_n)\big)\\ &\cup\big(M_n\times\mathcal{F}^{n+1}\big)\cup\big(M_n\times F_{R_r^n}(M_n)\big).\end{aligned}$$

So in order to prove (5.4) at rank $n+1$, we need to show that

$$\mathcal{S}^{n+1}\cap\Big((\mathcal{D}\times\mathcal{F})\backslash(\mathcal{D}^n\times\mathcal{F}^n)\Big)=\mathcal{S}^n\cap\Big((\mathcal{D}\times\mathcal{F})\backslash(\mathcal{D}^n\times\mathcal{F}^n)\Big), \tag{5.5}$$

$$\mathcal{S}^{n+1}\cap\big(M_n\times\mathcal{F}^{n+1}\big)=\mathcal{S}\cap\big(M_n\times\mathcal{F}^{n+1}\big)=\emptyset, \tag{5.6}$$

$$\mathcal{S}^{n+1}\cap\big(\mathcal{D}^{n+1}\times F_{R_r^n}(M_n)\big)=\mathcal{S}\cap\big(\mathcal{D}^{n+1}\times F_{R_r^n}(M_n)\big)=\emptyset, \tag{5.7}$$

$$\begin{aligned}\mathcal{S}^{n+1}\cap\big(M_n\times F_{R_r^n}(M_n)\big)&=\mathcal{S}\cap\big(M_n\times F_{R_r^n}(M_n)\big)\\ &=\mathcal{S}^0\cap\big(M_n\times F_{R_r^n}(M_n)\big).\end{aligned} \tag{5.8}$$

To prove these equalities, the main tool is the following formula which is a direct consequence of the construction:

$$\mathcal{S}^{n+1}=\mathcal{S}^n\backslash\big(\mathcal{D}^{n+1}\times F_{R_r^n}(M_n)\big). \tag{5.9}$$

With this formula at hand, we see that (5.5) follows from $\mathcal{D}^{n+1}\times F_{R_r^n}(M_n)\subset\mathcal{D}^n\times\mathcal{F}^n$. We also deduce very easily that $\mathcal{S}^{n+1}\cap(\mathcal{D}^{n+1}\times\mathcal{F}^{n+1})=\mathcal{S}^n\cap(\mathcal{D}^{n+1}\times\mathcal{F}^{n+1})=\mathcal{S}^0\cap(\mathcal{D}^{n+1}\times\mathcal{F}^{n+1})$, where the last equality follows from the first point at rank n.

Then, to prove (5.6), as both $\mathcal{S}$ (by assumption) and $\mathcal{S}^{n+1}$ (by (5.9)) are included in $\mathcal{S}^n$, it suffices to show that $\mathcal{S}^n\cap(M_n\times\mathcal{F}^{n+1})=\emptyset$. But that

last assertion follows from the definition of $\mathcal{F}^{n+1} = \mathcal{F}^n \backslash F_{R_r^n}(M_n)$: these are precisely the columns where R_r^n has only zero entries on the intersection with the lines M_n.

To prove (5.7), let us observe that the equality $\mathcal{S}^{n+1} \cap (\mathcal{D}^{n+1} \times F_{R_r^n}(M_n)) = \emptyset$ is a direct consequence of (5.9). The other equality, namely, $\mathcal{S} \cap (\mathcal{D}^{n+1} \times F_{R_r^n}(M_n)) = \emptyset$ follows from the fact that M_n is a SISP set for $(R_r^n; \mu_{\llcorner \mathcal{D}^n}, \nu_{\llcorner \mathcal{F}^n})$, so that (5.2) applies with M_n instead of A, R_r^n instead of R, $\mathcal{D}^n$ instead of $\mathcal{D}$ and $R^{n,*} = R^{*}{}_{\llcorner \mathcal{D}^n \times \mathcal{F}^n}$ instead of R^* (here, we use the point three at rank n). Notice that as $\mathcal{S} \cap (\mathcal{D}^{n+1} \times F_{R_r^n}(M_n)) = \emptyset$ and $\mathcal{S} \subset \mathcal{S}^n$, by (5.9), we have also proved that $\mathcal{S} \subset \mathcal{S}^{n+1}$.

Finally, to prove (5.8), as $\mathcal{S} \subset \mathcal{S}^{n+1} \subset \mathcal{S}^0$, we just need to prove that $\mathcal{S} \cap (M_n \times F_{R_r^n}(M_n)) = \mathcal{S}^0 \cap (M_n \times F_{R_r^n}(M_n))$. But this is a direct consequence of the fact that M_n is a SISP set for $(R_r^n; \mu_{\llcorner \mathcal{D}^n}, \nu_{\llcorner \mathcal{F}^n})$, so that (5.3) applies with R_r^n instead of R, $R^{n,*}$ instead of R^*, M_n instead of A and $\mathcal{F}^n$ instead of $\mathcal{F}$.

Second point. — By definition, $\mathcal{D}^{n+1}$ is empty if and only if $M_n = \mathcal{D}^n$. So if $\mathcal{D}^{n+1} = \emptyset$, then by Assumption 2.12 applied to $(R_r^n; \mu_{\llcorner \mathcal{D}^n}, \nu_{\llcorner \mathcal{F}^n})$, we clearly have $F^{R_r^n}(M_n) = \mathcal{F}^n$ and so $\mathcal{F}^{n+1} = \emptyset$. On the other hand, if $\mathcal{D}^{n+1} \neq \emptyset$ we have

$$\begin{aligned} 0 < \mu\big(\mathcal{D}^{n+1}\big) &= P^*\big(\mathcal{D}^{n+1} \times \mathcal{F}\big) = P^*\big(\mathcal{D}^{n+1} \times \mathcal{F}^n\big) \\ &= P^*\big(\mathcal{D}^{n+1} \times F_{R_r^n}(M_n)\big) + P^*\big(\mathcal{D}^{n+1} \times \mathcal{F}^{n+1}\big) \\ &= P^*\big(\mathcal{D}^{n+1} \times \mathcal{F}^{n+1}\big), \end{aligned}$$

where the inequality comes from Assumption 2.12, the second equality is an easy consequence of (5.4) at rank n, and the last one comes from the definition of SISP sets (see (5.2)) and from the fact that P^* and R^* has the same support. So $\mathcal{F}^{n+1}$ cannot be empty, as announced.

Third point. — To check that

$$\big(R_r^{n+1}; \mu_{\llcorner \mathcal{D}^{n+1}}, \nu_{\llcorner \mathcal{F}^{n+1}}\big)$$

satisfies Assumption 2.12, we need only need to show that the support of ${\mu^{R^{n+1}}}_{\llcorner \mathcal{D}^{n+1}}$ is $\mathcal{D}^{n+1}$ and the support of ${\nu^{R^{n+1}}}_{\llcorner \mathcal{F}^{n+1}}$ is $\mathcal{F}^{n+1}$. But as we already proved that $\mathcal{S} \subset \mathcal{S}^{n+1}$, we know that $P^* \ll R^{n+1}$ and $Q^* \ll R^{n+1}$, so the conclusion follows from the fact that μ and ν have full support by Assumption 2.12 applied to $(R; \mu, \nu)$. The last thing to check, namely that the matrices $P^{n+1,*}$, $Q^{n+1,*}$ and $R^{n+1,*}$ associated with $(R_r^{n+1}; \mu_{\llcorner \mathcal{D}^{n+1}}, \nu_{\llcorner \mathcal{F}^{n+1}})$ through Theorems 3.2 and 4.3 are the restrictions of P^*, Q^* and R^* to $\mathcal{D}^{n+1} \times \mathcal{F}^{n+1}$ is a direct consequence of Proposition 5.7 below (that we wanted to separate to the rest of the proof because we will use it again later), and of (5.7). □

PROPOSITION 5.7. — *Let $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ satisfy Assumption 2.11. Let P^*, Q^* and R^* be the matrices associated with the problem $\mathrm{Sch}(R;\mu,\nu)$ by Theorems 3.2 and 4.3. Finally, let $A \subset \mathcal{D}$ be such that*

$$R^*\big((\mathcal{D}\backslash A) \times F_R(A)\big) = 0. \tag{5.10}$$

Then (with slightly sloppy notations),

$$P^*{}_{\llcorner A\times F_R(A)}, Q^*{}_{\llcorner A\times F_R(A)} \text{ and } R^*{}_{\llcorner A\times F_R(A)}$$

are the matrices associated with the restricted problem

$$\mathrm{Sch}\big(R_{\llcorner A\times F_R(A)}, \mu_{\llcorner A}, \nu_{\llcorner F_R(A)}\big)$$

by Theorems 3.2 and 4.3.

Similarly, calling $A' := \mathcal{D}\backslash A$ and $F' := \mathcal{F}\backslash F_R(A)$, $P^*{}_{\llcorner A'\times F'}$, $Q^*{}_{\llcorner A'\times F'}$ and $R^*{}_{\llcorner A'\times F'}$ are the matrices associated with the restricted problem

$$\mathrm{Sch}\big(R_{\llcorner A'\times F'}, \mu_{\llcorner A'}, \nu_{\llcorner F'}\big)$$

by Theorems 3.2 and 4.3.

Proof. — We show the result in the case of P^*, related to the "block" $A \times F_R(A)$. The case of Q^* is similar, the case of R^* easily follows from the two previous ones, and the similar results on $A' \times F'$ follow the same lines. Let ν^* be defined by (3.1). The first thing to prove is

$$\nu^*{}_{\llcorner F_R(A)} = \arg\min \left\{ H\left(\overline{\nu} \,\middle|\, \nu_{\llcorner F_R(A)}\right) \,\middle|\, \begin{array}{l} \overline{\nu} = \nu^P \text{ for some} \\ P \in \Pi_1(\mu_{\llcorner A}) \cap \mathcal{H}_+\left(R_{\llcorner A\times F_R(A)}\right) \end{array} \right\}.$$

The measure $\nu^*{}_{\llcorner F_R(A)}$ is a competitor for the problem in the r.h.s. because it corresponds to $P := P^*{}_{\llcorner A\times F_R(A)}$. Let us show that it is the optimizer. To do this, we call $\overline{\nu}^*$ the optimizer, and we show that $\nu^*{}_{\llcorner F_R(A)} = \overline{\nu}^*$. Let us consider $\overline{P}$ a P corresponding to $\overline{\nu}^*$ in the problem above, $\overline{P}^*$ the matrix obtained by replacing the entries of P^* on $A \times F_R(A)$ by the entries of $\overline{P}$, and $\overline{\nu}^* := \nu^{\overline{P}^*}$. We have

$$\begin{aligned} H(\overline{\nu}^* \,|\, \nu) &= H\big(\overline{\nu} \,\big|\, \nu_{\llcorner F_R(A)}\big) + H\big(\nu^*{}_{\llcorner F'} \,\big|\, \nu_{\llcorner F'}\big) \\ &\leqslant H\big(\nu^*{}_{\llcorner F_R(A)} \,\big|\, \nu_{\llcorner F_R(A)}\big) + H\big(\nu^*{}_{\llcorner F'} \,\big|\, \nu_{\llcorner F'}\big) \\ &= H(\nu^* \,|\, \nu), \end{aligned}$$

where the inequality, being a consequence of the optimality of $\overline{\nu}^*$, is an equality if and only if $\nu^*_{F_R(A)} = \overline{\nu}^*$. But by optimality of ν^* in (3.1) this inequality is indeed an equality, and therefore $\nu^*_{F_R(A)} = \overline{\nu}^*$.

It remains to show that $P^*{}_{\llcorner A\times F_R(A)}$ is the solution of

$$\mathrm{Sch}\big(R_{\llcorner A\times F_R(A)}, \mu_{\llcorner A}, \nu^*{}_{\llcorner F_R(A)}\big).$$

For this, let us consider $\overline{P}$ the solution of $\mathrm{Sch}(R_{\llcorner A\times F_R(A)}, \mu_{\llcorner A}, {\nu^*}_{\llcorner F_R(A)})$, and $\overline{P}^*$ the matrix obtained by replacing the entries of P^* on $A \times F_R(A)$ by the entries of $\overline{P}$. Because of (5.10), we have

$$\begin{aligned} H(\overline{P}^* \,|\, R) &= H\big(\overline{P} \,\big|\, R_{\llcorner A\times F_R(A)}\big) + H\big({P^*}_{\llcorner A'\times F'} \,\big|\, R_{\llcorner A'\times F'}\big) \\ &\leqslant H\big({P^*}_{\llcorner A\times F_R(A)} \,\big|\, R_{\llcorner A\times F_R(A)}\big) + H\big({P^*}_{\llcorner A'\times F'} \,\big|\, R_{\llcorner A'\times F'}\big) \\ &= H(P^* \,|\, R^*), \end{aligned}$$

where the inequality is a consequence of the optimality of $\overline{P}$, and is an equality if and only if $\overline{P} = {P^*}_{\llcorner A\times F_R(A)}$. But by optimality of P^*, this inequality is indeed an equality, and we conclude that $\overline{P} = {P^*}_{\llcorner A\times F_R(A)}$. The Proposition 5.7 is proved. □

Now, we want to prove Lemma 5.5. To do this, we introduce a new class of subsets of $\mathcal{D}$, associated with a triple $(R;\mu,\nu)$.

Definition 5.8. — *Let $R \in \mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ satisfying Assumption 2.12. The maximal θ associated to $(R;\mu,\nu)$ is defined by:*

$$\theta_m := \max_{\substack{A\subset\mathcal{D}\\ A\neq\emptyset}} \frac{\mu(A)}{\nu(F_R(A))}. \tag{5.11}$$

We say that $A \subset \mathcal{D}$ is a maximal θ-set for $(R;\mu,\nu)$ if A is a maximizer of (5.11)*. We say that it is a* smallest *maximal θ-set if in addition, it is a minimal element in the sense of inclusion among all maximal θ-sets associated with $(R;\mu,\nu)$.*

As maximal θ-sets are optimizers of a finite function (thanks to Assumption 2.12) on a finite set (the set of all nonempty subsets of $\mathcal{D}$), any triple $(R;\mu,\nu)$ satisfying Assumption 2.12 admits at least one maximal θ-set. The set of all maximal θ-sets being itself finite, we know that there exists at least one minimal element in this set, so that smallest maximal θ-sets always exist under Assumption 2.12. Hence, Lemma 5.5 is an obvious consequence of the following proposition, whose proof heavily relies on the optimality conditions stated in Proposition 4.5.

Proposition 5.9. — *Let $R \in \mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ satisfying Assumption 2.12. A smallest maximal θ-set for $(R;\mu,\nu)$ is a SISP set for $(R;\mu,\nu)$.*

Proof. — Let μ^* and ν^* be defined by (3.1). We first define the two following quantities

$$\theta^{\mathcal{D}} := \max_{i\in\mathcal{D}} \frac{\mu_i}{\mu_i^*}, \qquad \theta^{\mathcal{F}} := \max_{j\in\mathcal{F}} \frac{\nu_j^*}{\nu_j}.$$

(Recall that under Assumption 2.12, by Remark 3.4, μ^* and ν^* defined by (3.1) have full support.) Then, we define the two following sets, that are nonempty subsets of $\mathcal{D}$ and $\mathcal{F}$ respectively:

$$\overline{M} := \left\{ x_i \in \mathcal{D} \text{ s.t. } \frac{\mu_i}{\mu_i^*} = \theta^{\mathcal{D}} \right\}, \qquad \overline{F} := \left\{ y_j \in \mathcal{F} \text{ s.t. } \frac{\nu_j^*}{\nu_j} = \theta^{\mathcal{F}} \right\}.$$

The main argument of the proof consists in showing that $\theta^{\mathcal{D}}$ and $\theta^{\mathcal{F}}$ coincide with θ_m, the maximal θ for $(R;\mu,\nu)$. Even if $\overline{M}$ is not a smallest maximal-θ set in general (more precisely, it is a maximal θ-set that is not minimal in general), we show at Step 3 below how this information allows to conclude. As before, P^* is the matrix defined by (3.2).

Step 1: $\theta^{\mathcal{D}} = \theta^{\mathcal{F}}$. — Let $x_i \in \overline{M}$ and $y_j \in \mathcal{F}$ be such that $P^*_{ij} > 0$ (such a j exists thanks to Assumption 2.12). By the first line of (4.6), we have

$$\frac{\mu_i^*}{\mu_i}\frac{\nu_j^*}{\nu_j} = 1, \tag{5.12}$$

which implies that $\nu_j^*/\nu_j = \theta^{\mathcal{D}}$, and hence that $\theta^{\mathcal{F}} \geqslant \theta^{\mathcal{D}}$. The other inequality is proved in the same way, and the result follows. From now on, we call

$$\bar{\theta} := \theta^{\mathcal{D}} = \theta^{\mathcal{F}}.$$

Step 2: $\bar{\theta} = \theta_m$. — First, $\bar{\theta} \geqslant \theta_m$. Indeed, for any $A \subset \mathcal{D}$, by Theorem 5.1, as $\mathrm{Sch}(R;\mu,\nu^*)$ is at least approximately scalable, we have $\mu(A) \leqslant \nu^*(F_R(A))$. But on the other hand, by definition of $\bar{\theta}$, we have $\nu^* \leqslant \bar{\theta}\nu$, so that actually, $\mu(A) \leqslant \bar{\theta}\nu(F_R(A))$, and hence $\bar{\theta} \geqslant \theta_m$.

Also, $\bar{\theta} \leqslant \theta_m$. To see this, let us first observe that $P^*((\mathcal{D}\backslash\overline{M}) \times \overline{F}) = P^*(\overline{M} \times (\mathcal{F}\backslash\overline{F})) = 0$. This is because if x_i, y_j are such that $P^*_{ij} > 0$, still by (5.12), $\mu_i/\mu_i^* = \bar{\theta}$ if and only if $\nu_j^*/\nu_j = \bar{\theta}$, so that $x_i \in \overline{M}$ if and only if $y_j \in \overline{F}$. Therefore, on the one hand, $F_R(\overline{M}) \subset \overline{F}$, and on the other hand, projecting on both marginals:

$$\begin{aligned}\mu(\overline{M}) = P^*(\overline{M} \times \mathcal{F}) &= P^*(\overline{M} \times \overline{F}) + \underbrace{P^*(\overline{M} \times (\mathcal{F}\backslash\overline{F}))}_{=0} \\ &= P^*(\overline{M} \times \overline{F}) + \underbrace{P^*((\mathcal{D}\backslash\overline{M}) \times \overline{F})}_{=0} = P^*(\mathcal{D} \times \overline{F}) = \nu^*(\overline{F}).\end{aligned}$$

As by definition of $\overline{F}$ and $\bar{\theta}$, $\nu^*(\overline{F}) = \bar{\theta}\nu(\overline{F})$, we conclude that

$$\bar{\theta}\nu(F_R(\overline{M})) \leqslant \bar{\theta}\nu(\overline{F}) = \mu(\overline{M}),$$

so that $\bar{\theta} \leqslant \theta_m$, as announced.

Step 3: Conclusion. — We are now in position to conclude. Let A be a smallest maximal θ-set. As $\mathrm{Sch}(R;\mu,\nu^*)$ is at least approximately scalable,

we know that $\mu(A) \geqslant \nu^*(F_R(A))$. On the other hand, as A is a maximal θ-set, we know that $\mu(A) = \theta_m \nu(F_R(A)) = \bar{\theta}\nu(F_R(A))$. But by definition of $\bar{\theta}$, we know that $\bar{\theta}\nu \geqslant \nu^*$, so that $\nu^*(F_R(A)) \leqslant \mu(A)$. We conclude that $\nu^*(F_R(A)) = \mu(A)$, that ${\nu^*}_{\llcorner A} = \bar{\theta}\nu_{\llcorner A} = \theta_m \nu_{\llcorner A}$, and hence that

$$\begin{aligned} P^*\big((\mathcal{D}\backslash A) \times F_R(A)\big) &= P^*(\mathcal{D} \times F_R(A)) - P^*(A \times F_R(A)) \\ &= \nu^*(F_R(A)) - \mu(A) = 0, \end{aligned}$$

so that (5.2) holds.

In addition, by Proposition 5.7 the measure ${P^*}_{\llcorner A \times F_R(A)}$ is the solution of the problem $\mathrm{Sch}(R_{\llcorner A \times F_R(A)}; \mu_{\llcorner A}, {\nu^*}_{\llcorner F_R(A)})$. So in order to prove (5.3), it suffices to prove that this problem is scalable. For this purpose, we will use Lemma 5.2. We call $R_r := R_{\llcorner A \times F_R(A)}$. Let B be a nonempty strict subset of A. As A is a minimal element in the set of maximal θ-sets for $(R; \mu, \nu)$, we know that $\mu(B) < \theta_m \nu(F_R(B))$. Then, as $F_R(B) \subset F_R(A)$, we have $F_R(B) = F_{R_r}(B)$, so $\mu(B) < \theta_m \nu(F_{R_r}(B))$. Finally, as ${\nu^*}_{\llcorner A} = \theta_m \nu_{\llcorner A}$, we have $\mu(B) < \nu^*(F_{R_r}(B))$. So Lemma 5.2 applies, and $\mathrm{Sch}(R_{\llcorner A \times F_R(A)}; \mu_{\llcorner A}, {\nu^*}_{\llcorner F_R(A)})$ is scalable, which concludes the proof of Proposition 5.9. □

We close this section with a remark concerning the stability with respect to union of SISP sets.

Remark 5.10. — It is easy to check that SISP sets associated with a triple $(R; \mu, \nu)$ are stable by union. Therefore, there exists an upper bound in the set of all SISP sets for $(R; \mu, \nu)$, that we call the *largest* SISP set. If we want the procedure described in Proposition 5.6 to be as fast as possible, it is logical to look for SISP sets that are as large as possible, in order to minimize the rank N at which the procedure reaches its stationary point. This is what we are going to do in the next section.

6. Numerical applications

A simple consequence of the theoretical procedure described in the previous section is that in a lot of cases, if the problem $\mathrm{Sch}(R; \mu, \nu)$ is non-scalable, the matrices P^* and Q^* from Theorem 3.2, and R^* from Theorem 4.3 have more zero entries than R. For instance, if the problem is balanced (i.e. $\mathsf{M}(\mu) = \mathsf{M}(\nu)$), and if the bipartite graph of R is connected (that is, $\mu(A) = \nu(F_R(A))$ only holds for $A = \emptyset$ or $A = \mathcal{D}$, which is a reasonable assumption in a lot of contexts), we can check that the matrix R^1 from Proposition 5.6 cannot coincide with R.

Therefore, typically, the Sinkhorn algorithm in the non-scalable case does not converge linearly. Also, when running the algorithm on a computer, some

numerical errors may happen before a chosen stopping criterion is verified. This would be due to the explosion of the terms in the sequences of vectors $(a^n)_{n\in\mathbb{N}}, (b^n)_{n\in\mathbb{N}}$ defined by (1.3) when n is large: indeed, for the matrices P^* and Q^* to have more zeros than R, some coordinates of the terms of these sequences necessarily go to $+\infty$. In order to improve the numerical stability of the algorithm with respect to the choice of the stopping criterion while making it faster, we are going to detail an approximate algorithm allowing to find the common support of P^*, Q^* and R^*, and therefore to recover a linear rate of convergence for the Sinkhorn algorithm by Proposition 5.3.

6.1. Stopping criterion

Before any numerical application, we need to define a stopping criterion for the Sinkhorn algorithm when the Schrödinger problem is non-scalable. When the problem is scalable, the classical criterion that is used is the duality gap estimated at each step $n \in \mathbb{N}^*$ of the Sinkhorn algorithm:

$$SC^n = H(P^n \,|\, R) - \langle \log(a^n), \mu\rangle - \left\langle \log\left(b^{n-1}\right), \nu\right\rangle, \tag{6.1}$$

where P^n, a^n and b^{n-1} are defined by (1.3). Indeed, it is known that this quantity is always positive when the relation $P^n_{ij} = a^n_i b^{n-1}_j R_{ij}$ holds for all i, j, and the classical results ensure that $SC^n \to 0$ as $n \to \infty$ when the problem is scalable, i.e. when (a^n) and (b^n) converge.

In the approximately scalable case, numerical instabilities may appear when $n \to \infty$ because (a^n) and (b^n) do not converge, but this criterion may remain useful if the error that is tolerated is not too small. However, in the non-scalable case, this criterion does not hold as the problem $\mathrm{Sch}(R; \mu, \nu)$ has no solution. The results presented in the previous sections allow nevertheless to define an approximate criterion. Indeed, it has been shown in [6] that for a given $\lambda > 0$, the problem defined with the notations of Section 2.1 by

$$\mathrm{Schu}_\lambda(R; \mu, \nu) := \min\Big\{ H(P \,|\, R) + \lambda H\big(\nu^P \,\big|\, \nu\big) \,\Big|\, P \in \Pi_1(\mu)\Big\}, \tag{6.2}$$

can be solved numerically with a generalization of the Sinkhorn algorithm. More precisely, the duality gap defined for all $n \in \mathbb{N}^*$ by

$$\begin{aligned} &SCu^n_\lambda \\ &= H(P^n \,|\, R) + \lambda\left(H\big(\nu^{P^n} \,\big|\, \nu\big) - \left\langle 1 - \left(1/b^{n-1}\right)^{\frac{1}{\lambda}}, \nu\right\rangle\right) - \langle \log(a^n), \mu\rangle, \end{aligned} \tag{6.3}$$

converges to 0 whenever $P^n_{ij} := a^n_i b^{n-1}_j R_{ij}$ for all i, j, and with a^n, b^n defined for all $n \in \mathbb{N}^*$ by the relations:

$$\begin{cases} \forall\, j, \qquad b^0_j := 1, \\ \forall\, n \geqslant 0, \quad \forall\, i, \quad a^{n+1}_i := \dfrac{\mu_i}{\sum_j b^n_j R_{ij}}, \\ \forall\, n \geqslant 0, \quad \forall\, j, \quad b^{n+1}_j := \left(\dfrac{\nu_j}{\sum_i a^{n+1}_i R_{ij}}\right)^{\frac{\lambda}{1+\lambda}}. \end{cases} \tag{6.4}$$

On the other hand, a slight modification of our Γ-convergence result of Proposition 4.1 asserts that P^*, from Theorem 3.2, is the limit of the solution of (6.2) as $\lambda \to +\infty$. So if we now define SCu^n_λ by the formula (6.3) where P^n, a^n and b^{n-1} are computed with the standard Sinkhorn algorithm (1.3), instead of the modified one (6.4), we conclude that for all $\varepsilon > 0$, there exists a threshold λ^ε such that for all $\lambda \geqslant \lambda^\varepsilon$,

$$\limsup_{n \to +\infty} SCu^n_\lambda \leqslant \varepsilon.$$

Therefore, the stopping criterion (6.3) can still be used for the sequence $(P^n)_{n \in \mathbb{N}}$ generated by the classical Sinkhorn algorithm (1.3), as long as λ^ε is chosen to be sufficiently large w.r.t. ε. In practice, we observe that taking $\lambda^\varepsilon = \frac{1}{\varepsilon}$ works well. This is what we are going to do in the following section, considering a level of error $\varepsilon := 10^{-3}$.

6.2. An approximate numerical method for constructing the support of R^*

An interesting application of the theoretical procedures described in Section 5.2 is the construction of an approximate algorithm allowing the identification of the support $\mathcal{S}$ of R^* w.r.t. R, when the problem $\mathrm{Sch}(R; \mu, \nu)$ is approximately scalable or non-scalable. As mentioned at the beginning of this section, our motivation is twofold.

- First, as knowing the support allows to recover a linear rate for the Sinkhorn algorithm thanks to Proposition 5.3, we can hope that if our way to find the support is fast enough, the full procedure to obtain R^* (preprocessing to know the support $\mathcal{S}$ and then the Sinkhorn algorithm applied to $\mathrm{Sch}(\mathbb{1}_{\mathcal{S}} R; \mu, \nu)$) is faster than the Sinkhorn algorithm applied directly to $\mathrm{Sch}(R; \mu, \nu)$.

- Second, even if we have shown that the Sinkhorn algorithm converges in the non-scalable case, the Schrödinger potentials appearing in the Sinkhorn procedure are likely to be too high to be computed numerically before the stopping criterion is satisfied: thus, even without talking of speed, finding the support before running the Sinkhorn algorithm is useful.

Let us now explain our method to find the support of R^*. The idea is of course to apply Proposition 5.6. Therefore, what we need to do is to recursively identify SISP sets, and preferably, the largest possible ones. Actually, as suggested by Proposition 5.9 and Remark 5.10, our approximate algorithm will aim at identifying at each iteration the union of all smallest maximal θ-sets of the corresponding restricted problem. Let us present the first iteration of our algorithm, which will then be used recursively following the procedure described at Proposition 5.6. In other words, let us explain how we identify a SISP set for $(R; \mu, \nu)$. This is done in two steps.

(1) The first step looks for the largest set M' (in the sense of inclusion) such that the supports of R and R^* coincide in $M' \times \mathcal{F}$. Having in mind the notations of Figure 5.1, M' is the largest set A such that the support of R and R^* coincide on $\mathfrak{C} := A \times F_R(A)$. As we will see, this step is the one that is not exact. Notice that this set M' must contain every SISP sets. It is therefore nonempty under Assumption 2.12 in virtue of Lemma 5.5.

(2) However, this M' is not necessarily a SISP set itself: still with the notations of Figure 5.1, if $A = M'$, R^* can perfectly be nonzero on the block $\mathfrak{B}^{(2)}$. Therefore, the second step consists in finding $M \subset M'$, the union of all smallest maximal θ-set for the restricted problem $(R_{\llcorner M' \times F_R(M')}; \mu_{\llcorner M'}, \nu_{\llcorner F_R(M')})$, which as we will see will be much easier once we have restricted the problem to $M' \times F_R(M')$. It will turn out that this M will also be a SISP set for the full problem $(R; \mu, \nu)$.

Finding M'

Let us develop further the first step. In order to find M', we actually use the Sinkhorn algorithm. We initialize at $U := \mathcal{D}$ and $V := \mathcal{F}$, and we launch Sinkhorn. At each iteration, for all $x_i \in U$, if there exists y_j such

[(2)] As a counterexample, take $R := \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$, $\mu = (2, 1, 1)$ and $\nu := (1, 2, 1)$. We find $R^* = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$ and $M' = \{2, 3\}$, which is not a SISP set.

that $R_{ij} > 0$ and the coupling obtained at this step is smaller than a given threshold m_i, we set $U = U\backslash\{x_i\}$. Then, for all $y_j \in V$, if $y_j \notin F_R(U)$, we set $V = V\backslash\{y_j\}$.

By construction, the limiting sets U and V are clearly M' and $F_R(M')$, unless the threshold are not well chosen, as explained below. Since the Sinkhorn algorithm restricted to $U \times V$ converges linearly (thanks to Proposition 5.3), we then rapidly observe the convergence of the procedure.

Finding M

For detailing how to obtain M at the second step, we need to introduce the notion of connected components:

Definition 6.1. — *Let $A \subset U \subset \mathcal{D}$ and $B \subset V \subset \mathcal{F}$. We say that $A \times B$ is a connected component of the graph $(U \cup V, \triangle)$ (using the notation of (5.1)) if:*

- $R(A \times (V\backslash B)) = R((U\backslash A) \times B) = 0$;
- $(A \cup B, \triangle)$ *is connected.*

Finding connected components for such undirected graph is a classical task in Graph theory, for which there exists ready-to-use algorithms [14]. Proposition 6.3 below justifies that we can define at the second step of the procedure

$$M := \bigcup \left\{ U_i \;\middle|\; \frac{\mu(U_i)}{\nu(V_i)} = \max_{j=1,\dots,C} \frac{\mu(U_j)}{\nu(V_j)} \right\}, \tag{6.5}$$

where $U_1 \times V_1, \dots, U_C \times V_C$ are the connected components of the graph $(M' \cup F_R(M'), \triangle)$.

Remark 6.2. — In view of how we identify M', our method can be seen as an improvement over the naive approximate method which consists, at each iteration of the Sinkhorn algorithm applied to $\mathrm{Sch}(R; \mu, \nu)$, to set to zero all the entries of the obtained matrix that are smaller than a certain threshold. With our method, we do not have to identify all theses entries one by one but line by line, which can avoid numerical errors due to entries converging slowly to zero. However, this is done at the cost of identifying the connected components of some subgraphs of $(\mathcal{D} \cup \mathcal{F}, \triangle)$ at every iteration of the procedure, which slows down the algorithm in cases where these subgraphs are large. Note that in typical cases where the matrix R is structured, we do not expect to find more than one connected component at each iteration, as in Figure 5.3.

Proposition 6.3. — *Let $R \in \mathcal{M}_+(\mathcal{D} \times \mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ satisfying Assumption 2.12, and let us denote M' the largest subset of $\mathcal{D}$ (in the sense of inclusion) such that the supports of R and R^* coincide in $M' \times \mathcal{F}$. Let us denote $(U_1 \times V_1, \dots, U_C \times V_C)$ the connected components of the graph $(M' \cup F_R(M'), \triangle)$. Finally, consider $U \subset \mathcal{D}$. Then, the two following assertions are equivalent.*

(1) *U is a smallest maximal θ-set.*
(2) *There is $i \in \{1, \dots, C\}$ such that $U = U_i$ and*

$$\frac{\mu(U_i)}{\nu(V_i)} = \max_{j=1,\dots,C} \frac{\mu(U_j)}{\nu(V_j)}. \tag{6.6}$$

Proof. — Let Θ the map that associates to any nonempty $U \subset \mathcal{D}$ the quantity

$$\frac{\mu(U)}{\nu(F_R(U))}.$$

It is well defined under Assumption 2.12. We only need to prove that the maximizers of Θ that are minimal in the sense of inclusion (i.e., the smallest maximal θ-sets) are of the form U_i, $i = 1, \dots, C$ as defined in the statement. Indeed, these smallest maximal θ-sets will then be the maximizers of Θ within this class, that is, the maximizers of the r.h.s. of (6.6) (as it is clear from the definition of connected components that for all i, $V_i = F_R(U_i)$).

So let us consider a smallest maximal θ-set U. We have to show that $U \times F_R(U)$ is a connected component of $(M' \cup F_R(M'), \triangle)$. By Proposition 5.9, U is a SISP set and so it is in M' and we have necessarily $R^*(U^c \times F_R(U)) = R^*(U \times F_R(U)^c) = 0$. The supports of R and R^* being the same in $M' \times F_R(M')$ by definition of M', we have therefore $R((M' \backslash U) \times F_R(U)) = R(U \times (F_R(M') \backslash F_R(U))) = 0$. We conclude that U is a set of the form

$$\bigcup_{i \in \mathcal{J} \subset \{1,\dots,C\}} U_j.$$

But now, because of the following inequality, valid for all $a, b \in \mathbb{R}_+$ and $c, d \in \mathbb{R}_+^*$:

$$\frac{a+b}{c+d} \leqslant \max\left(\frac{a}{c}, \frac{b}{d}\right),$$

it is clear that $\Theta(U) \leqslant \max_{i \in \mathcal{J}} \Theta(U_i)$. As U is a maximizer of Θ, this is an equality. As it is minimal in the sense of inclusion, the cardinality of $\mathcal{J}$ is 1 and the result is proved. □

Pseudo-code and comments. Now that we have all the ingredients for our procedure, we provide in Algorithm 1 the pseudo-code of this iterative method. Let us make a few comments about this algorithm.

- Easy arguments show that the support of R^* only depends on the support of R and not on its values. For this reason, when identifying $\mathcal{S}$, we decided to consider the problem $\mathrm{Sch}(R'; \mu, \nu)$, where $R' = \mathbb{1}_{R \neq 0}$ instead of $\mathrm{Sch}(R; \mu, \nu)$. This explains why we replace R by $\mathbb{1}_{\mathrm{Supp}}$ at the beginning of Algorithm 1. Our reason for doing so is that in that case, it is fully relevant in our method to find M' to have only one threshold by line, and not by entry: the threshold will not be overpassed just because of small values of R.
- The stopping criterion corresponds to the criterion (6.3) detailed in the previous section, which has to be smaller than a certain threshold ε to be satisfied.
- Choosing in an appropriate way the set of minimal factors $\{m_i, i = 1, \dots, N\}$, is crucial: it determines the level of approximation that is considered as acceptable, i.e. the minimal value at which we can consider that the algorithm should create a new zero entries. In practice, we observe that

$$m_i := \frac{1}{n} \frac{\mu_i}{\mu_i^{R'}}, \tag{6.7}$$

where $R' = \mathbb{1}_{R \neq 0}$ seems to be a good tradeoff between efficiency and security in most of the cases that we explored.

Our algorithm is only approximate

We emphasize the fact that this algorithm is only approximate, and this for two reasons. The first one occurs when the set of thresholds $\{m_i, i = 1, \dots, N\}$ are too large. Then, we can set to zero lines which should not be, just because some of the entries of R^* should be small on this line. This must be avoided as then the algorithm cannot converge towards R^*. For example, in the extreme case where for all $i, m_i \geqslant \max_j a_i b_j$ in Algorithm 1, the first step gives $U = \mathcal{A}$ and then $\mathrm{Supp} = \emptyset$. Note however that by construction of this algorithm, the new support, even if too sparse, gives rise to a new reference matrix $\mathbb{1}_{\mathcal{S}} R$ which is diagonal by block (up to proper permutations of lines and columns), and such that every block $A \times F_{\mathbb{1}_{\mathcal{S}} R}(A)$ is associated to a scalable problem $\mathrm{Sch}(R_{\llcorner A \times F_{\mathbb{1}_{\mathcal{S}} R}(A)}, \mu_{\llcorner A}, \nu_{\llcorner F_{R \mathbb{1}_{\mathcal{S}} R}})$ (and in line with Assumption 2.11, we restrict the problem to the blocks which do not contain only zero entries). Thus, the Sinkhorn procedure applied to the problem $\mathrm{Sch}(\mathbb{1}_{\mathcal{S}} R; \mu, \nu)$ is always going to converge without numerical errors for the same stopping criterion as the one used for the preprocessing, but providing the wrong optimum.

Algorithm 1 Find the support Supp of R^*: return Supp

Require: • A set of minimal factors: $\{m_i, i = 1, \dots, N\}$,
• A stopping criterion: $stop(a, b, R, \mu, \nu)$.
We set $A = \mathcal{D}$, $B = \mathcal{F}$, Supp = Support(R), $R = \mathbb{1}_{\text{Supp}}$.
while $A \neq \emptyset$ **do**
$\overline{R} = R_{\llcorner A \times B}$, $\overline{\mu} = \mu_{\llcorner A}$, $\overline{\nu} = \nu_{\llcorner B}$
$b = \mathbb{1}_B$, $a = \mathbb{1}_A$
$U = A$, $V = B$
while $stop(a, b, \overline{R}, \overline{\mu}, \overline{\nu}) \neq 1$ **do**
for $x_i \in U$ **do**
if $\sum_{y_j \in V} \overline{R}_{ij} b_j = 0$ **then**
$U = U \backslash \{x_i\}$
else
$a_i = \dfrac{\overline{\mu}_i}{\sum_{y_j \in V} \overline{R}_{ij} b_j}$
if $\min_{y_j \in V \text{ s.t. } R_{ij} > 0} a_i \times b_j < m_i$ **then**
$U = U \backslash \{x_i\}$
end if
end if
end for
$\overline{R} = \overline{R}_{\llcorner U \times V}$, $\overline{\mu} = \overline{\mu}_{\llcorner U}$, $a = a_{\llcorner U}$
for $y_j \in V$ **do**
if $\sum_{x_i \in A} \overline{R}_{ij} a_i = 0$ **then**
$V = V \backslash \{y_j\}$
else
$b_j = \dfrac{\overline{\nu}_j}{\sum_{x_i \in A} \overline{R}_{ij} a_i}$
end if
end for
$\overline{R} = \overline{R}_{\llcorner U \times V}$, $\overline{\nu} = \overline{\nu}_{\llcorner V}$, $b = b_{\llcorner V}$
end while
$(U_1 \times V_1, \dots, U_C \times V_C) =$ connected components of the graph $(U \cup V, \triangle)$
$U \times V = \bigcup \left\{ U_i \times V_i \;\middle|\; \frac{\mu(U_i)}{\nu(V_i)} = \max_{j=1,\dots,C} \frac{\mu(U_j)}{\nu(V_j)} \right\}$
$A = A \backslash U$, $B = B \backslash V$
Supp $=$ Supp $\backslash \big((A \backslash U) \times V \big)$
end while

The other case where our algorithm does not identify $\mathcal{S}$ exactly is either when the threshold ε of the stopping criterion is large, or when the thresholds $\{m_i, i = 1, \dots, N\}$ are small. Then, two cases may happen: (i) the Sinkhorn algorithm can satisfy the stopping criterion before all the zeros have been

identified or (ii) some numerical errors may appear at certain steps before the stopping criterion is satisfied (and then before to identify the SISP sets guaranteed by Lemma 5.5). Case (i) is not a big problem, since it means that the algorithm converges well without having to identify the additional zeros of R^*, and the Sinkhorn algorithm with the new support will converge without numerical errors for the same stopping criterion. However, in case (ii) the Sinkhorn algorithm will not converge numerically either, due to similar numerical errors.

With these observations, we conclude that the thresholds $\{m_i, i = 1, \dots, N\}$ need to be taken rather small w.r.t. the level of error ε of the stopping criterion, even though if they are taken too small, efficiency is lost since then the algorithm just behaves as Sinkhorn without any improvement. The formula (6.7) is in general well suited. If numerical errors appear during the preprocessing procedure, we recommend to increase the level of error, and restart the procedure.

Efficiency of the procedure

We illustrate in Figure 6.1 the efficiency of this procedure by representing the number of iterations (represented in the upper figure) and the computation time (in the below figure) for the Sinkhorn algorithm to converge as a function of the number of additional zero entries in R^* w.r.t. to R, and compare when we apply or not the preprocessing described in Algorithm 1.

For this comparison, we considered matrices of size 100×100. For varying the number of zero entries, we take R upper-diagonal, build μ and ν similarly to what we described in Figure 5.3, and then vary the number of blocks (the number N from Proposition 5.6) from 1 (corresponding to the scalable case) to 10. For the case with preprocessing, we consider the sum of the iterations needed for the Sinkhorn-like method described in Algorithm 1 to find the support $\mathcal{S}$, and of the ones needed for the Sinkhorn algorithm then applied to the problem $\mathrm{Sch}(\mathbb{1}_{\mathcal{S}} R; \mu, \nu)$. We observe that the preprocessing makes the number of iterations needed for the convergence to be significantly smaller than for the case without preprocessing when the number of additional zero entries (or equivalently, the number N of blocks) is high. It is also smaller than for the naive approximate method described in Remark 6.2, illustrating the benefit of our approach.

In terms of computation time, our method is slightly slower than the classical Sinkhorn algorithm when the number of additional zeros is small. This is the price to pay to avoid numerical instabilities in the approximately scalable or non-scalable case.

Remark 6.4. — The careful reader has probably noticed that we run Sinkhorn in order to find the support of he limiting matrices, and once it is done, apply Sinkhorn again. So we could wonder if it is possible to both find the support and compute the solution at the same time. Actually, we could, and the only reason why we do not is that for finding the support, as already explained, we preferred to work with the matrix $\mathbb{1}_{R\neq 0}$ instead of R for stability reasons. However, the Sinkhorn algorithm is so fast once the support is found that we do not reduce a lot the number of iterations needed when computing everything at once. To give some order of magnitude, in our numerical experiments, we never had more than a few hundreds of iterations for the whole procedure (as seen in Figure 6.2), among which only a few tens were dedicated to computing the solution. In other terms, in the most favorable case, finding the support reduced the number of iterations of the Sinkhorn algorithm from more than 3000 to a few tens for the chosen stopping criterion. Therefore, in order to still improve the computation time, one should mainly focus in computing the limiting support more quickly.

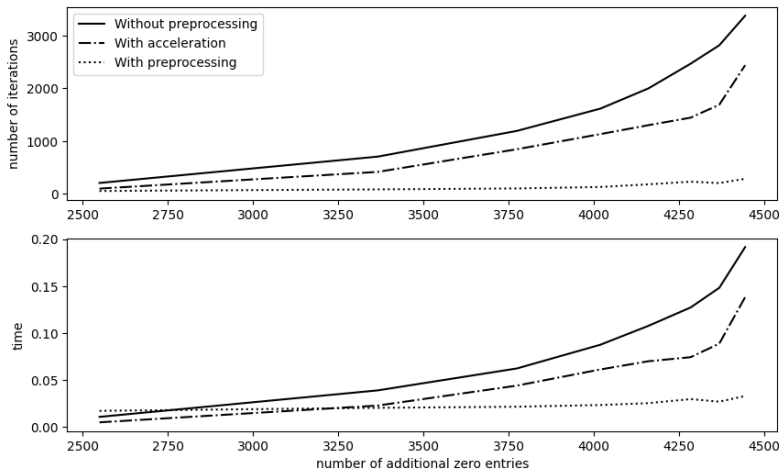

Figure 6.1. Number of iterations (up) and computational time (down) needed for convergence *vs.* number of additional zero entries of the limits P^* and Q^* from Theorem 3.2 w.r.t. R, using: the Sinkhorn algorithm (1.3) (solid line); the Sinkhorn algorithm where at each step, the entries below the threshold (6.7) are set to zero (dashed line). Algorithm 1 to compute the support $\mathcal{S}$ of the limits, and then the Sinkhorn algorithm replacing R by $\mathbb{1}_{\mathcal{S}}R$ (dotted line). When the threshold ε is small enough, as it is the case here, these three methods provide the same limits.

6.3. Comparison of the method with the balanced and unbalanced Sinkhorn algorithms

We compared in Figure 6.2 the outputs of the Sinkhorn algorithm when the problem $\mathrm{Sch}(R;\mu,\nu)$ is non-scalable, given by the geometric mean described in Theorem 4.3, and two alternatives:

- When the reference coupling R is modified such that it has only positive entries. For that, we built a new coupling R_ε by adding on every zero entry of R a small quantity ε. We then found the optimizers R^*_ε of the Schrödinger problems $\mathrm{Sch}(R_\varepsilon;\mu,\nu)$ and compared its distance in total variation of its solution to R^*, for different values of ε.
- When we the marginal constraints are replaced by marginal penalizations, leading to an unbalanced problem of the form (4.1), using the scaling algorithm described in [6]. We then compare the distance in total variation of its solution to R^* for different values of λ.

For the comparison realized here, we took a coupling R of size 100×100 and R,μ,ν as in Figure 5.3 in such a way that R^* has only two blocks $A_1\times B_1$ and $A_2\times B_2$ for which the factor λ appearing in the procedure of Proposition 5.6 is greater than one for the first component, and smaller for the second one. The problem is thus non-scalable. As expected, we observe in Figure 6.2(a) that in the first case it is impossible for the solution R^*_ε of $\mathrm{Sch}(R_\varepsilon;\mu,\nu)$ to be close to R^* (and then to recover the right minimum entropy), and that the faster is the convergence, the further R^*_ε is from R^*. In the second case, we observe in Figure 6.2(b) that the solution R^*_λ of the unbalanced problem with penalization λ converges to R^* when $\lambda\to\infty$. However, the convergence goes faster only for values of λ smaller than 150, for which we still observe a significant difference between R^* and R^*_λ.

Thus, adding a small ε on the zero entries of R and solving the resulting scalable Schrödinger problem provides a solution which is different from the original one even when $\varepsilon\to 0$. Considering an unbalanced problem with parameter λ allows to recover the right optimum when λ is high enough, but the number of iterations required for the unbalanced algorithm to converge, for such a λ, is higher than the number of iterations required for our method.

Note that these results are not limited to Schrödinger problems similar to the one described in Figure 5.3, and that we observed the same type of results for randomly-generated R, μ and ν in non-scalable cases.

Overall, this suggests that when the problem is non-scalable, the Sinkhorn algorithm overcomes these two methods for minimizing the limit of (4.1) when $\lambda\to\infty$.

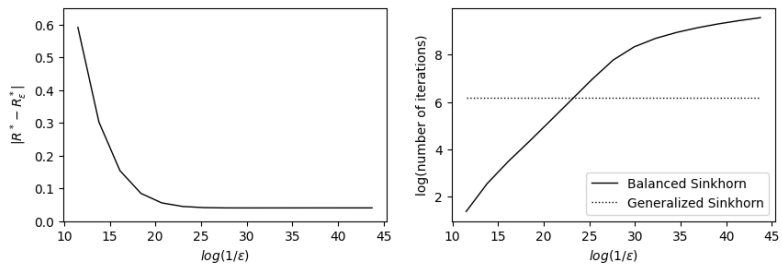

(a)

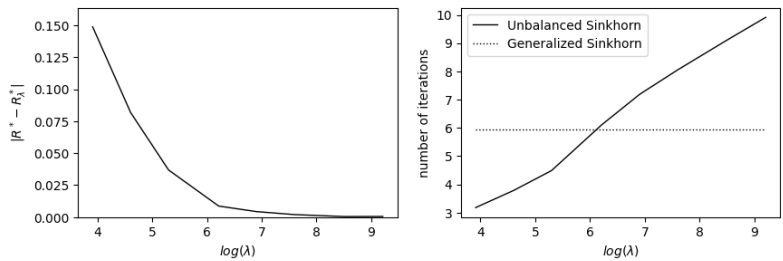

(b)

Figure 6.2. Comparison of the outputs of the Sinkhorn algorithm in a non-scalable case, where R^* is given by Theorem 4.3, and: (a) the Sinkhorn algorithm (1.3) where the zero entries of R are replaced by a small value ε; (b) the unbalanced Sinkhorn algorithm from [6] applied to solve $\mathrm{Schu}_\lambda(R; \mu, \nu)$ for large values of λ.

Acknowledgments

We would like to thank Olivier Gandrillon, Thibault Espinasse, Thomas Lepoutre and more generally the INRIA team Dracula for enlightening discussions. We also thank Thomas Lepoutre and Gauthier Clerc for critical reading of the manuscript. We would like to thank as well the BioSyL Federation, the LabEx Ecofect (ANR-11-LABX-0048) and the LabEx Milyon of the University of Lyon for inspiring scientific events. Finally, we would like to express our gratitude to the anonymous referee. Their careful reading and relevant remarks definitely led to significant improvements of our work.

Appendix A. Example of Schrödinger problems without solutions

There exists a lot of degenerate cases where the problem $\mathrm{Sch}(R;\mu,\nu)$ has no solution. Indeed, in the extreme situation where most of the entries of R cancel, two randomly chosen vectors μ and ν have more chance to be non-scalable than to satisfy the conditions of Theorem 5.1. For example, in the typical example of a squared diagonal reference coupling R, we must necessarily have $\mu = \nu$ for these conditions to be satisfied.

When illustrating our results at Sections 5 and 6, we chose R to be a squared upper-diagonal matrix (see Figure 5.3). This is of particular interest, as it corresponds to a case that typically arises when considering entropy minimization problems in cell biology. Indeed, the dynamics of mRNA levels within a cell, which drives cellular differentiation processes, is often modeled by a piecewise deterministic Markov process, where stochastic bursts of mRNAs compensate their deterministic degradation [37, 39]. Considering the simplest cartoonish but enlightening situation where there is no degradation, a constant number of cells, and where we measure the activity of only one gene, the quantity of mRNAs in the cells corresponding to this gene can only increase with time. Therefore, if R is the matrix whose entry R_{ij} gives the number of cells having i molecules of mRNA at a first timepoint and j molecules at a later timepoint, R must be upper-diagonal.

To give an insight of the behaviour of the Sinkhorn algorithm in the non-scalable case with an upper-diagonal reference matrix, let us treat explicitly a simple example. We consider:

$$R = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{pmatrix}, \quad \mu = (2,2,2), \quad \nu = (2,3,1).$$

In this example, $\mu_3 > \nu_3$ while the image of x_3 by the graph associated to R is reduced to y_3, that is, with the formalism of Section 5, $F_R(\{x_3\}) = \{y_3\}$ and hence $\nu(F_R(\{x_3\}) < \mu(\{x_3\})$. In view of Theorem 5.1 (which is very easy to check in our simple situation), the problem is therefore indeed non-scalable: no matrix can satisfy the marginal constraints and be absolutely continuous w.r.t. R at the same time.

With the notations of (1.3), let us reproduce below the output of the Sinkhorn algorithm at some of the first iterations. Starting at Iteration 5, we only give approximate numerical values.

Iteration 1. —

$$a^1 = (2/3,1,2), \quad b^0 = (1,1,1), \quad P^1 = \begin{pmatrix} 2/3 & 2/3 & 2/3 \\ 0 & 1 & 1 \\ 0 & 0 & 2 \end{pmatrix}.$$

Iteration 2. —

$$a^1 = (2/3, 1, 2), \quad b^1 = (3, 9/5, 3/11), \quad Q^1 = \begin{pmatrix} 2 & 6/5 & 2/11 \\ 0 & 9/5 & 3/11 \\ 0 & 0 & 6/11 \end{pmatrix}.$$

Iteration 5. —

$$a^3 = \left(2.7e^{-1}, 8.6e^{-1}, 1.7e^{1}\right), \quad b^2 = \left(5.0, 2.2, 1.2e^{-1}\right),$$
$$P^3 = \begin{pmatrix} 1.4 & 0.59 & 3.2e^{-2} \\ 0 & 1.9 & 1.0e^{-1} \\ 0 & 0 & 2.0 \end{pmatrix}.$$

Iteration 11. —

$$a^6 = \left(1.2e^{-1}, 5.1e^{-1}, 1.5e^{2}\right), \quad b^5 = \left(1.3e^{1}, 3.9, 1.3e^{-2}\right),$$
$$P^6 = \begin{pmatrix} 1.55 & 0.45 & 1.5e^{-3} \\ 0 & 2.0 & 6.6e^{-3} \\ 0 & 0 & 2.0 \end{pmatrix}.$$

Iteration 80. —

$$a^{40} = \left(5.5e^{-5}, 2.8e^{-4}, 2.7e^{12}\right), \quad b^{40} = \left(3.6e^{4}, 9.1e^{3}, 3.8e^{-13}\right),$$
$$Q^{40} = \begin{pmatrix} 2.0 & 0.50 & 1.7e^{-17} \\ 0 & 2.5 & 8.4e^{-16} \\ 0 & 0 & 1.0 \end{pmatrix}.$$

Iteration 81. —

$$a^{41} = \left(4.4e^{-5}, 2.2e^{-4}, 5.3e^{12}\right), \quad b^{40} = \left(3.6e^{4}, 9.1e^{3}, 3.8e^{-13}\right),$$
$$P^{40} = \begin{pmatrix} 1.6 & 0.40 & 4.2e^{-17} \\ 0 & 2.0 & 2.1e^{-16} \\ 0 & 0 & 2.0 \end{pmatrix}.$$

Of course, in this case, the matrices P^*, Q^*, R^* from Theorems 3.2 and 4.3 are given by

$$P^* = \begin{pmatrix} 8/5 & 2/5 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 2 \end{pmatrix}, \quad Q^* = \begin{pmatrix} 2 & 1/2 & 0 \\ 0 & 5/2 & 0 \\ 0 & 0 & 1 \end{pmatrix},$$
$$R^* = \begin{pmatrix} 4/\sqrt{5} & 1/\sqrt{5} & 0 \\ 0 & \sqrt{5} & 0 \\ 0 & 0 & \sqrt{2} \end{pmatrix}.$$

Finally, to get $\overline{R}^*$ from Theorem 4.8, it suffices to normalize R^*.

This very simple example illustrates the different points developed in this article:

- When R does not have only positive entries, the limits of the sequences (P^n) and (Q^n) given by the Sinkhorn algorithm may be different and have more zero entries than R.
- Because new zero entries appear, the potentials (a^n) and (b^n) that are updated at each iteration of the Sinkhorn algorithm cannot converge: some of their coordinates have to tend to 0 and then some other ones need to diverge to $+\infty$ as the number of iterations increases.
- More precisely, for (i, j) on the common support of P^* and Q^* from Theorem 3.2, the infinitely small and high values of the two potentials are compensated. For (i, j) outside of this common support, but still in the one of R, the multiplication of the two potentials generate infinitely small values. Outside of the support of R, the multiplication of the potentials can diverge. Also, the zero entries of R prevent the sums involved in the computations of (a^n) and (b^n) to diverge: the large values of the potentials are sent to zero in the multiplication with R.
- When the problem is non-scalable, the algorithm still converges to two limits and the algorithm alternates between them. These two limits correspond to solutions of the Schrödinger problem with modified marginals, that is with modified μ or modified ν alternatively (see the iterations 80 and 81).

Going back to the context of the beginning of the section, where the upper-diagonal R models the evolution of the quantity of mRNAs corresponding to one gene in a population of cells between two timepoints, we see that the non-scalable case appears when there exists a threshold such that more cells with less mRNAs than the threshold are measured at the second timepoint than at the first one, which is incompatible with the model where the quantity of mRNAs can only increase. If we believe enough in our model, it is natural to look for a solution with modified marginals – like for instance the law $\overline{R}^*$ described in Theorem 4.8 – and to advocate for a bad sampling or imprecise measurements when collecting data.

If we consider that such incompatibilities between the theoretical model R and the observations μ and ν should be rare, this law $\overline{R}^*$ is a natural choice since among all the solutions of a Schrödinger problem w.r.t. R, it is the one whose marginals are the closest (in a specific entropic sense) to the experimental ones.

Appendix B. The analogy between entropic and Euclidean projections

In this appendix, the goal is to show that similarly to what happens in the scalable or approximately scalable case studied by Csiszár [7], identity (3.11) has a geometric interpretation relying on entropic projections and the Pythagorean law of Theorem 2.7. To achieve that goal we first recall how this analogy leads to the result in [7], and then explain how the argument needs to be changed in our non-scalable case. In each situation, we present an analogous Euclidean situation of which our results are entropic generalizations.

Let us consider as before $\mathcal{D}$ and $\mathcal{F}$ two finite spaces, and $R \in \mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, $\mu \in \mathcal{M}_+(\mathcal{D})$ and $\nu \in \mathcal{M}_+(\mathcal{F})$ satisfying Assumption 2.11.

Alternate orthogonal projections on two intersecting lines

In our entropic context, $\Pi_1(\mu)$ and $\Pi_2(\nu)$ are the intersections of affine subsets of the set of matrices with the cone $\mathcal{M}_+(\mathcal{D}\times\mathcal{F})$, and in the scalable or approximately scalable case, these two sets intersect in $\mathcal{M}_+(\mathcal{D}\times\mathcal{F})$. Therefore, a good analogy in the Euclidean world is to study the sequence built by alternately projecting orthogonally on two affine subspaces that intersect, and to show that this sequence converges towards the intersection point that is the closest to the initial point. The easiest of such problems is the case of two lines intersecting in $\mathbb{R}^2$. We drew this case in the l.h.s. of Figure B.1.

With the notations of the figure and setting $Q^0 := R$, the main idea to show that (Q^n) converges towards the intersection point R^* is to find an inference formula for the sequence $(\|R^* - Q_n\|^2)$. Using twice the Pythagorean theorem we see that for all $n \geqslant 0$,

$$\|R^* - Q^n\|^2 = \left\|R^* - Q^{n+1}\right\|^2 + \left\|Q^{n+1} - P^{n+1}\right\|^2 + \left\|P^{n+1} - Q^n\right\|^2. \quad \text{(B.1)}$$

Therefore, the sequence $(\|Q^n - R^*\|^2)$ is nonincreasing, and the main argument of the proof consists in using iteratively this formula in order to find for all $n \geqslant 0$

$$\|R^* - R\|^2 = \|R^* - Q^n\|^2 + \sum_{k=1}^{n} \left\|Q^k - P^k\right\|^2 + \left\|P^k - Q^{k-1}\right\|^2. \quad \text{(B.2)}$$

This is the Euclidean version of [7, formula (3.14)] (that we recalled at formula (3.3) above).

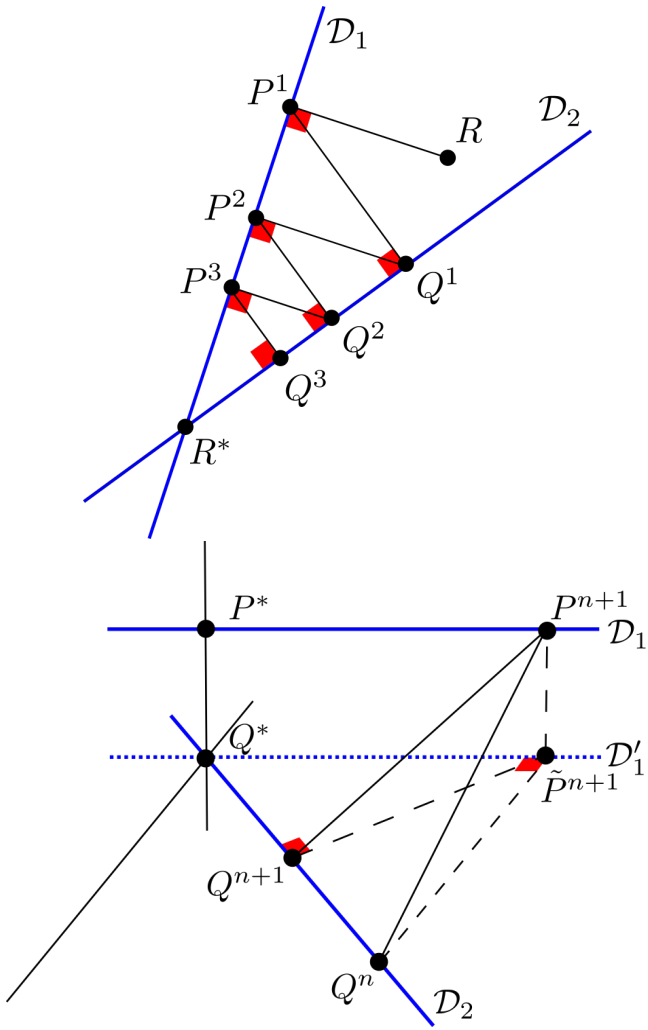

Figure B.1. On the left: a sketchy Euclidean version of the Sinkhorn algorithm in the scalable case. On the right: a sketchy Euclidean version of the Sinkhorn algorithm in the non-scalable case.

Adaptation to the entropic scalable or approximately scalable case

Remarkably, in the approximately scalable or scalable case, this proof works in the entropic context replacing the Pythagorean theorem by Theorem 2.7. Indeed, calling $\overline{R}$ any competitor of $\mathrm{Sch}(R;\mu,\nu)$ (this time, the

intersection point is not unique, but it does not change much the idea), we know that at a given rank $n \geqslant 0$:

- P^{n+1} is the projection of Q^n on $\Pi_1(\mu)$, so that
$$H(\overline{R} \,|\, Q^n) = H\big(\overline{R} \,\big|\, P^{n+1}\big) + H\big(P^{n+1} \,\big|\, Q^n\big);$$
- Q^{n+1} is the projection of P^{n+1} on $\Pi_2(\nu)$, so that
$$H\big(\overline{R} \,\big|\, P^{n+1}\big) = H\big(\overline{R} \,\big|\, Q^{n+1}\big) + H\big(Q^{n+1} \,\big|\, P^{n+1}\big).$$

We deduce the following transparent analogue of formula (B.1):

$$H(\overline{R} \,|\, Q^n) = H\big(\overline{R} \,\big|\, Q^{n+1}\big) + H\big(Q^{n+1} \,\big|\, P^{n+1}\big) + H\big(P^{n+1} \,\big|\, Q^n\big).$$

Thus, the sequence $(H(R^* \,|\, Q^n))$ is nonincreasing, and using iteratively this formula leads to (3.3), the main tool of the proof of convergence of the Sinkhorn algorithm given in [7].

Alternate orthogonal projections on two non-intersecting lines

Now, we would expect the non-scalable case to be similar to the Euclidean case, but when we alternately project orthogonally on two affine subspaces that do not intersect, which is possible for instance for two lines $\mathcal{D}_1$ and $\mathcal{D}_2$ in $\mathbb{R}^3$, as drawn in the r.h.s. of Figure B.1 (note that this figure is completely general up to choosing an appropriate coordinate system of $\mathbb{R}^3$).

In that case, the first thing that we need in order to study the limits of (P^n) and (Q^n) is to define $(P^*, Q^*) \in \mathcal{D}_1 \times \mathcal{D}_2$ as the unique pair such that $\|P^* - Q^*\|$ is the distance between the two lines. The main outcome of this definition is that P^* is the orthogonal projection of Q^* on $\mathcal{D}_1$ and Q^* is the orthogonal projection of P^* on $\mathcal{D}_2$.

Let us give the main step in proving the convergence of (Q^n) to Q^*. This time, the two applications of the Pythagorean theorem relying on the facts that for a given rank $n \geqslant 0$, P^{n+1} is the orthogonal projection of Q^n on $\mathcal{D}_1$ and that Q^{n+1} is the orthogonal projection of P^{n+1} on $\mathcal{D}_2$ are

$$\|P^* - Q^n\|^2 = \big\|P^* - P^{n+1}\big\|^2 + \big\|P^{n+1} - Q^n\big\|^2, \tag{B.3}$$

$$\big\|Q^* - P^{n+1}\big\|^2 = \big\|Q^* - Q^{n+1}\big\|^2 + \big\|Q^{n+1} - P^{n+1}\big\|^2. \tag{B.4}$$

These identities can only be combined thanks to the following further applications of the Pythagorean theorem relying on the fact that P^* is the

orthogonal projection of Q^* on $\mathcal{D}_1$, and that Q^* is the orthogonal projection of P^* on $\mathcal{D}_2$:

$$\|P^* - Q^n\|^2 = \|P^* - Q^*\|^2 + \|Q^* - Q^n\|^2, \tag{B.5}$$

$$\left\|Q^* - P^{n+1}\right\|^2 = \|Q^* - P^*\|^2 + \left\|P^* - P^{n+1}\right\|^2. \tag{B.6}$$

Therefore, we have:

$$\begin{aligned}\|Q^* - Q^n\|^2 = \left\|Q^* - Q^{n+1}\right\|^2 + \left\|Q^{n+1} - P^{n+1}\right\|^2 \\ - \|Q^* - P^*\|^2 + \left\|P^{n+1} - Q^n\right\|^2 - \|P^* - Q^*\|.\end{aligned} \tag{B.7}$$

This is enough to conclude that $(\|Q_n - Q^*\|^2)$ is nonincreasing because by definition of P^* and Q^*,

$$\left\|Q^{n+1} - P^{n+1}\right\|^2 \geqslant \|Q^* - P^*\|^2 \quad \text{and} \quad \left\|P^{n+1} - Q^n\right\|^2 \geqslant \|P^* - Q^*\|.$$

Once again, this formula can be used iteratively to give a formula of type (B.2) and this is the main idea of the proof.

To go further in the interpretation, it is useful to introduce the point $\widetilde{P}^n = P^n + Q^* - P^*$, the orthogonal projection of P^n on the line parallel to $\mathcal{D}_1$ passing through Q^*, that we call $\mathcal{D}_1'$. Using the fact that the line passing through P^{n+1} and $\widetilde{P}^{n+1}$ is orthogonal to the plan containing $\mathcal{D}_2$ and $\mathcal{D}_1'$, we end up with the following more concise formula giving the same information as (B.7)

$$\|Q^* - Q^n\|^2 = \left\|Q^* - Q^{n+1}\right\|^2 + \left\|Q^{n+1} - \widetilde{P}^{n+1}\right\|^2 + \left\|\widetilde{P}^{n+1} - Q^n\right\|^2. \tag{B.8}$$

As we will see, in the entropic case, there is no analogue to (B.7), and there is no analogue to (B.8), but we will find a formula that is between the two.

Actually, what we realize with this proof is that as shown on the picture, the family $(\widetilde{P}^n, Q^n)$ is the sequence of points obtained by alternately projecting orthogonally on $\mathcal{D}_1'$ and $\mathcal{D}_2$, starting from the orthogonal projection of R on the plan containing $\mathcal{D}_1'$ and $\mathcal{D}_2$: *when the lines do not cross, up to projecting orthogonally everything on the plan just described, we are back to the case when they do cross.* Unfortunately, this property is not true in the entropic context, which makes everything more delicate.

Adaptation to the entropic non-scalable case

Let us check what remains of the previous paragraph in the entropic non-scalable context, in spite of the non-symmetry of the relative entropy.

The first difficulty is to find admissible limiting pairs $(\overline{P}, \overline{Q})$, in the sense that $\overline{P}$ is the entropic projection of $\overline{Q}$ on $\Pi_1(\mu)$ and $\overline{Q}$ is the entropic projection of $\overline{P}$ on $\Pi_2(\nu)$. As in the scalable case, these limiting objects are not unique, but we do not enter into the details here. The pair (P^*, Q^*) defined in (3.2) is an instance of such a pair thanks to Proposition 4.5, and we could see that a $\overline{Q} \in \Pi_2(\nu)$ admits a $\overline{P} \in \Pi_1(\mu)$ such that $(\overline{P}, \overline{Q})$ is admissible if and only if $\overline{Q} \in \Pi(\mu^*, \nu)$. So let us give ourselves and admissible limiting pair $(\overline{P}, \overline{Q})$. In particular, let us keep in mind that $\overline{Q} \in \Pi(\mu^*, \nu)$.

Then, the analogues of (B.3), (B.4), (B.5) and (B.6) are at a given rank $n \geqslant 0$:

- as P^{n+1} is the entropic projection of Q^n on $\Pi_1(\mu)$,
$$H(\overline{P} \,|\, Q^n) = H\big(\overline{P} \,\big|\, P^{n+1}\big) + H\big(P^{n+1} \,\big|\, Q^n\big),$$
- as Q^{n+1} is the entropic projection of P^{n+1} on $\Pi_2(\nu)$,
$$H\big(\overline{Q} \,\big|\, P^{n+1}\big) = H\big(\overline{Q} \,\big|\, Q^{n+1}\big) + H\big(Q^{n+1} \,\big|\, P^{n+1}\big), \tag{B.9}$$
- as $\overline{Q}$ is the entropic projection of $\overline{P}$ on $\Pi_2(\nu)$,
$$H\big(Q^{n+1} \,\big|\, \overline{P}\big) = H\big(Q^{n+1} \,\big|\, \overline{Q}\big) + H\big(\overline{Q} \,\big|\, \overline{P}\big),$$
- as $\overline{P}$ is the entropic projection of $\overline{Q}$ on $\Pi_1(\mu)$,
$$H\big(P^{n+1} \,\big|\, \overline{Q}\big) = H\big(P^{n+1} \,\big|\, \overline{P}\big) + H\big(\overline{P} \,\big|\, \overline{Q}\big).$$

Unfortunately, these identities cannot be combined because of the non-symmetry of the relative entropy. Therefore, the Euclidean proof cannot be used as such, and there is no direct analogue of (B.7).

Nevertheless, we can follow the idea of the previous paragraph and define the sequence $(\widetilde{P}^n)$. What plays the role of $\mathcal{D}_1'$ here is $\Pi_1(\mu^*)$: this is the space passing through $\overline{Q}$, and parallel to $\Pi_1(\mu)$ in the sense that for all $P \in \Pi_1(\mu)$, calling $\widetilde{P}$ its entropic projection on $\Pi_1(\mu^*)$, we have $H(\widetilde{P} \,|\, P) = H(\overline{Q} \,|\, \overline{P}) = H(\mu^* \,|\, \mu)$, which does not depend on P.

Following what we did in the previous paragraph, let us call $\widetilde{P}^{n+1}$ the entropic projection of P^{n+1} on $\Pi_1(\mu^*)$. This projection leads to the identities, valid for all $n \geqslant 0$:

$$H\big(\overline{Q} \,\big|\, P^{n+1}\big) = H\big(\overline{Q} \,\big|\, \widetilde{P}^{n+1}\big) + H\big(\widetilde{P}^{n+1} \,\big|\, P^{n+1}\big), \tag{B.10}$$
$$H\big(\widetilde{P}^{n+1} \,\big|\, P^n\big) = H(\overline{Q} \,|\, \overline{P}). \tag{B.11}$$

Similarly to what happens in the Euclidean case, it is easy to check that $\widetilde{P}^{n+1}$ is the entropic projection of Q^n on $\Pi_1(\mu^*)$, so that

$$H(\overline{Q} \,|\, Q^n) = H\big(\overline{Q} \,\big|\, \widetilde{P}^{n+1}\big) + H\big(\widetilde{P}^{n+1} \,\big|\, Q^n\big). \tag{B.12}$$

But, contrarily to the Euclidean case, Q^{n+1} is not the entropic projection of $\widetilde{P}^{n+1}$ on $\Pi_2(\nu)$: this time, $(\widetilde{P}^n, Q^n)$ *is not the family generated by the Sinkhorn algorithm associated with the marginals* μ^* *and* ν.

Actually, this is not a problem and we have still enough identities to be able to conclude. Indeed, gathering (B.9), (B.10), (B.11) and (B.12), we get

$$\begin{aligned} &H(\overline{Q} \,|\, Q^n) \\ &= H\big(\overline{Q} \,\big|\, Q^{n+1}\big) + H\big(Q^{n+1} \,\big|\, P^{n+1}\big) - H(\overline{Q} \,|\, \overline{P}) + H\big(\widetilde{P}^{n+1} \,\big|\, Q^n\big). \end{aligned} \tag{B.13}$$

Since $H(Q^{n+1} \,|\, P^{n+1}) \geqslant H(\overline{Q} \,|\, \overline{P})$ by the properties of $\overline{P}$ and $\overline{Q}$, we get that $(H(\overline{Q} \,|\, Q^n))$ is nonincreasing, which once again is the main argument of the proof. Using iteratively this identity and replacing some of the entropies by their values in terms of the marginals of the measures in play leads to (3.11) and hence to the result.

As a final remark, let us observe that in (B.13), contrarily to the Euclidean case, the quantity $H(Q^{n+1} \,|\, P^{n+1}) - H(\overline{Q} \,|\, \overline{P})$ cannot be simplified as $H(Q^{n+1} \,|\, \widetilde{P}^{n+1})$, and $H(\widetilde{P}^{n+1} \,|\, Q^n)$ cannot be expanded as $H(P^{n+1} \,|\, Q^n) - H(\overline{P} \,|\, \overline{Q})$.

Appendix C. Proof of Theorem 5.1

In this section, we prove Lemma 5.2 and then Theorem 5.1.

Proof of Lemma 5.2. — In Lemma 5.2, we are looking for necessary and sufficient conditions for $\mathrm{Sch}(R; \mu, \nu)$ to be scalable. As $\mathsf{M}(\mu) = \mathsf{M}(\nu)$ is clearly a necessary condition because of Remark 1.2, we assume once for all that it is true. Up to normalizing, we assume that $\mu \in \mathcal{P}(\mathcal{D})$ and $\nu \in \mathcal{P}(\mathcal{F})$.

In order to clarify what (a′) and (b′) mean, we start by considering the case where the reference matrix R is such that the graph $(\mathcal{D} \cup \mathcal{F}, \triangle)$ as defined in Section 5 is connected. In that case, recalling that μ and ν are assumed to be probability measures, the conditions (a′) and (b′) are equivalent to:

- (a″) The measures μ and ν have full support, and for all $\emptyset \subsetneq A \subsetneq \mathcal{D}$, $\mu(A) < \nu(F_R(A))$,
- (b″) The measures μ and ν have full support, and for all $\emptyset \subsetneq B \subsetneq \mathcal{F}$, $\nu(B) < \mu(D_R(B))$.

Indeed, it is easy to see that in the balanced and connected case, the only subsets A of $\mathcal{D}$ for which $\mu^R(A) = \nu^R(F_R(A))$ are $A = \emptyset$ and $A = \mathcal{D}$. Similarly, the only subsets B of $\mathcal{F}$ for which $\nu^R(B) = \mu^R(D_R(B))$ are $B = \emptyset$ and $B = \mathcal{F}$.

We are going to prove Lemma 5.2 under this connectivity assumption. Passing to the general case is direct, up to restricting ourselves to connected components of $(\mathcal{D} \cup \mathcal{F}, \triangle)$.

We only prove $(\mathrm{b}'') \Leftrightarrow (\mathrm{c}')$, as $(\mathrm{a}'') \Leftrightarrow (\mathrm{c}')$ is proved in the same way. The idea of the proof is to fix $\mu \in \mathcal{P}(\mathcal{D})$ of full support, that is, such that for all $i = 1, \dots, N$, $\mu_i > 0$, and to introduce the two following subsets of $\mathcal{P}(\mathcal{F})$:

$$\mathfrak{A} := \left\{ \nu \in \mathcal{P}(\mathcal{F}) \;\middle|\; \begin{array}{c} \forall\, \emptyset \subsetneq B \subsetneq \mathcal{D}, \nu(B) < \mu(D_R(B)) \\ \text{and } \forall\, j = 1, \dots, M, \nu_j > 0 \end{array} \right\},$$
$$\mathfrak{B} := \left\{ \nu := \nu^{\overline{R}} \;\middle|\; \overline{R} \sim R \text{ and } \mu^{\overline{R}} = \mu \right\}.$$

With these definitions, proving $(\mathrm{b}'') \Leftrightarrow (\mathrm{c}')$ exactly means proving

$$\mathfrak{A} = \mathfrak{B},$$

and this is what we will prove now. To do so, we will first show that $\mathfrak{B} \subset \mathfrak{A}$, and then that $\mathfrak{B}$ is open and closed in $\mathfrak{A}$. As $\mathfrak{A}$ is convex, and hence connected, the result will follow.

Step 1: $\mathfrak{B} \subset \mathfrak{A}$. — Let us show that $\mathfrak{B} \subset \mathfrak{A}$. To this end, let us consider $\nu \in \mathfrak{B}$, and $\overline{R}$ such that $R \in \Pi(\mu, \nu)$ and $\overline{R} \sim R$. For all $\emptyset \subsetneq B \subsetneq \mathcal{D}$ we have:

$$\begin{aligned} \nu(B) &= \sum_{x_i \in \mathcal{D}} \sum_{y_j \in B} \overline{R}_{ij} = \sum_{x_i \in D_R(B)} \sum_{y_j \in B} \overline{R}_{ij} \leqslant \sum_{x_i \in D_R(B)} \sum_{y_j \in \mathcal{F}} \overline{R}_{ij} \\ &= \mu(D_R(B)), \end{aligned}$$

and the equality holds only if for all $(x_i, y_j) \in D_R(B) \times B^c$, $\overline{R}_{ij} = 0$ and hence $R_{ij} = 0$. As by definition of $D_R(B)$, for all $(x_i, y_j) \in D_R(B)^c \times B$, we have $R_{ij} = 0$, an equality in the formula above would imply that $D_R(B) \times B$ is not connected to $D_R(B)^c \times B^c$ in $(\mathcal{D} \times \mathcal{F}, \triangle)$, which contradicts our connectivity assumption. We have thus a strict inequality. Finally, the full support of ν is also an consequence of the connectivity of $(\mathcal{D} \times \mathcal{F}, \triangle)$. For all $y_j \in \mathcal{F}$, let $x_i \in \mathcal{D}$ such that $R_{ij} > 0$, and hence $\overline{R}_{ij} > 0$. We have $\nu_j \geqslant \overline{R}_{ij} > 0$. Thus, $\mathfrak{B} \subset \mathfrak{A}$.

Step 2: $\mathfrak{B}$ *is open in* $\mathfrak{A}$. — It is clear by its definition that $\mathfrak{A}$ is open in $\mathcal{P}(\mathcal{F})$. Therefore, to prove that $\mathfrak{B}$ is open in $\mathfrak{A}$, it suffices to prove that $\mathfrak{B}$ is open in $\mathcal{P}(\mathcal{F})$. Let $\nu \in \mathfrak{B}$, and $\overline{R}$ be such that $\overline{R} \sim R$ and $\nu^{\overline{R}} = \nu$. We choose:

$$0 < \varepsilon < \min \left\{ \overline{R}_{ij} \;\middle|\; i, j \text{ s.t. } \overline{R}_{ij} > 0 \right\}, \tag{C.1}$$

which is possible because $\mathcal{D}$ and $\mathcal{F}$ are finite. By convexity, it is enough to prove that for all $j \neq j'$ in $\{1, , \dots, M\}$, $\nu + \varepsilon(\delta_j - \delta_{j'}) \in \mathfrak{B}$. As $(\mathcal{D} \cup \mathcal{F}, \triangle)$

is connected, we can find $j = j_0, i_1, j_1, , \dots, i_p, j_p = j'$ such that

$$y_j = y_{j_0} \triangle x_{i_1} \triangle y_{j_1} \triangle \cdots \triangle x_{i_p} \triangle y_{j_p} = y_{j'}.$$

Then we set

$$P := \overline{R} + \varepsilon \sum_{n=1}^{p} \left(\delta_{i_n j_{n-1}} - \delta_{i_n j_n} \right).$$

It is easy to check that $P \sim R$ (by the definition (C.1) of ε), that $\mu^P = \mu$, and that $\nu^P = \nu + \varepsilon(\delta_j - \delta_{j'})$, which therefore belongs to $\mathfrak{B}$.

Step 3: $\mathfrak{B}$ is closed in $\mathfrak{A}$, strategy of the proof. — Let us introduce the following subset of $\mathcal{P}(\mathcal{F})$:

$$\mathfrak{C} := \left\{ \nu := \nu^{\overline{R}} \,\middle|\, \overline{R} \in \Pi_1(\mu) \cap \mathcal{H}_+(R) \right\}.$$

This set is clearly closed in $\mathcal{P}(\mathcal{F})$, so to prove that $\mathfrak{B}$ is closed in $\mathfrak{A}$, it suffices to prove that $\mathfrak{B} = \mathfrak{A} \cap \mathfrak{C}$. As we have already seen that $\mathfrak{B} \subset \mathfrak{A}$, and as clearly $\mathfrak{B} \subset \mathfrak{C}$, the only inclusion that needs to be justified is $\mathfrak{A} \cap \mathfrak{C} \subset \mathfrak{B}$.

Let us choose $\nu \in \mathfrak{A} \cap \mathfrak{C}$, and let us consider R^* the solution of $\mathrm{Sch}(R; \mu, \nu)$. We will prove by contradiction that $R^* \sim R$, and hence that $\nu \in \mathfrak{B}$.

So let us assume that $R^* \nsim R$. Once again, we choose $0 < \varepsilon < \min\{R^*_{ij} \mid i, j \text{ s.t. } R^*_{ij} > 0\}$. For all i, j, we write

$$x_i \blacktriangle y_j$$

whenever $R^*_{ij} > 0$. We will first prove that $(\mathcal{D} \cup \mathcal{F}, \blacktriangle)$ is connected (this is the hardest part of the proof), and then that it coincides with $(\mathcal{D} \cup \mathcal{F}, \triangle)$, which exactly means that $R^* \sim R$.

Step 4: $(\mathcal{D} \cup \mathcal{F}, \blacktriangle)$ is connected. — We call $\mathcal{D}_1 \cup \mathcal{F}_1, \dots, \mathcal{D}_p \cup \mathcal{F}_p$ the connected components of $(\mathcal{D} \cup \mathcal{F}, \blacktriangle)$. Let us assume that $p > 1$, and show that it leads to a contradiction.

First, we claim that if $p > 1$, there exist $k_1, , \dots, k_l \in \{1, \dots, p\}$ a family of two by two distinct indices, $x_{i^{k_1}}, , \dots, x_{i^{k_l}} \in \mathcal{D}$ and $y_{j^{k_1}}, , \dots, y_{j^{k_l}} \in \mathcal{F}$ such that for all $q = 1, , \dots, l$, $x_{i^{k_q}} \in \mathcal{D}_{k_q}$ and $y_{j^{k_q}} \in \mathcal{F}_{k_q}$, and with the convention $l + 1 = 1$, $y_{j^{k_q}} \triangle x_{i^{k_{q+1}}}$.

For proving this claim, we start by building a directed graph structure on $\{1, , \dots, p\}$, the set of indices of the connected components of $(\mathcal{D} \cup \mathcal{F}, \blacktriangle)$. For all $k, k' \in \{1, , \dots, p\}$, we write $k \rightsquigarrow k'$ whenever $k \neq k'$ and there exists $y_j \in \mathcal{F}_k$ and $x_i \in \mathcal{D}_{k'}$ such that $y_j \triangle x_i$. Of course, our claim precisely means that the directed graph $(\{1, \dots, p\}, \rightsquigarrow)$ admits a cycle. Let us prove that for all $k = 1, \dots, p$, there exists $k' \in \{1, \dots, p\}$ such that $k \rightsquigarrow k'$, which is clearly enough to conclude.

Let us consider $k \in \{1, \dots, p\}$. Because $\nu \in \mathfrak{A}$, we have $\mu(D_R(\mathcal{F}_k)) > \nu(\mathcal{F}_k)$. On the other hand, $\nu(\mathcal{F}_k) = \mu(\mathcal{D}_k)$ (as under R^*, all the mass on $\mathcal{D}_k$ is sent to $\mathcal{F}_k$ and vice versa), and $\mathcal{D}_k \subset D_R(\mathcal{F}_k)$ (as $R^* \ll R$). Therefore, $\mu(D_R(\mathcal{F}_k)\backslash\mathcal{D}_k) = \mu(D_R(\mathcal{F}_k)) - \mu(\mathcal{D}_k) > 0$, and we conclude that $D_R(\mathcal{F}_k)\backslash\mathcal{D}_k \neq \emptyset$. Let $x_i \in D_R(\mathcal{F}_k)\backslash\mathcal{D}_k$, and k' such that $x_i \in \mathcal{D}_{k'}$. As $x_i \in D_R(\mathcal{F}_k)$, there is $y_j \in \mathcal{F}_k$ such that $y_j \triangle x_i$. It follows that $k \rightsquigarrow k'$, and the claim is proved, and we can consider $k_1, \dots, k_l$ satisfying the properties above.

We now show that we reach a contradiction, which allows to conclude that actually, p must be equal to 1 and hence that $(\mathcal{D} \cup \mathcal{F}, \blacktriangle)$ needs to be connected.

For all $q = 1, \dots, l$, as $(\mathcal{D}_{k_q} \cup \mathcal{F}_{k_q}, \blacktriangle)$ is connected, we can find a family of indices $i^{k_q} = i_1^{k_q}, j_1^{k_q}, \dots, i_{n_q}^{k_q}, j_{n_q}^{k_q} = j^{k_q}$ such that

$$x_{i^{k_q}} = x_{i_1^{k_q}} \blacktriangle y_{j_1^{k_q}} \blacktriangle \cdots \blacktriangle x_{i_{n_q}^{k_q}} \blacktriangle y_{j_{n_q}^{k_q}} = y_{j^{k_q}}$$

Now, still keeping the convention $l + 1 = 1$, we set

$$P := R^* + \varepsilon \sum_{q=1}^{l} \delta_{i^{k_{q+1}} j^{k_q}} - \varepsilon \sum_{q=1}^{l} \left(\sum_{n=1}^{n_q - 1} \left(\delta_{i_n^{k_q} j_n^{k_q}} - \delta_{i_{n+1}^{k_q} j_n^{k_q}} \right) + \delta_{i_{n_q}^{k_q} j_{n_q}^{k_q}} \right).$$

The matrix P has less zeros than R^*: by the definition (C.1) of ε, it has no additional zero and we have for instance $P_{i^{k_2} j^{k_1}} > 0$ and $R^*_{i^{k_2} j^{k_1}} = 0$. Moreover, $P \in \Pi(\mu, \nu)$, and the construction of the indices ensures that P has new non-zero entries w.r.t. R^* only on (x_i, y_j) such that $R_{ij} > 0$, which ensures that $P \in \mathcal{H}_+(R)$. In virtue of point (3) of Theorem 2.7, this contradicts the fact that R^* is the solution of $\mathrm{Sch}(R; \mu, \nu)$, and we conclude that $p = 1$.

Step 5: $(\mathcal{D} \cup \mathcal{F}, \blacktriangle) = (\mathcal{D} \cup \mathcal{F}, \triangle)$. — Our last task is to prove that whenever $x_i \triangle y_j$, for some $x_i \in \mathcal{D}$ and $y_j \in \mathcal{F}$, then we also have $x_i \blacktriangle y_j$ (the reciprocal statement follows from $R^* \ll R$). So let us consider $x_i \in \mathcal{D}$ and $y_j \in \mathcal{F}$ with $x_i \triangle y_j$, assume that we do not have $x_i \blacktriangle y_j$, and show that we reach a contradiction. As we know that $(\mathcal{D} \cup \mathcal{F}, \blacktriangle)$ is connected., we can find $i = i_1, j_1, i_2, , \dots, i_p, j_p = j'$ such that

$$x_i = x_{i_1} \blacktriangle y_{j_1} \blacktriangle x_{i_2} \blacktriangle \cdots \blacktriangle x_{i_p} \blacktriangle y_{j_p} = y_j.$$

We set

$$P := R^* + \varepsilon \delta_{ij} - \varepsilon \left(\sum_{n=1}^{p-1} \left(\delta_{i_n j_n} - \delta_{i_{n+1} j_n} \right) + \delta_{i_p j_p} \right).$$

Once again, P has less zeros than R^*, which is a contradiction, and the result follows. □

We are now ready to conclude the proof of Theorem 5.1.

Proof of Theorem 5.1. — It remains to show that if μ and ν are such that (a) is verified (and not (a′)), then the problem is approximately scalable (once again, (b) $\Rightarrow$ (c) is proved in the same way, and (c) $\Rightarrow$ (a) and (c) $\Rightarrow$ (b) are easy).

Let us first remark that the problem $\mathrm{Sch}(R; \mu^R, \nu^R)$ is obviously scalable. Let us consider $\varepsilon \in (0,1)$. We define:

$$\mu^\varepsilon := (1-\varepsilon)\mu + \varepsilon\mu^R \quad \text{and} \quad \nu^\varepsilon = (1-\varepsilon)\nu + \varepsilon\nu^R.$$

The condition (a) implies that for all $A \subset \mathcal{D}$:

$$\mu^\varepsilon(A) \leqslant (1-\varepsilon)\nu(F_R(A)) + \varepsilon\nu^R(F_R(A)) = \nu^\varepsilon(A).$$

with a strict inequality whenever $\mu^R(A) < \nu^R(F_R(A))$.

Let us now assume that Assumption 2.12 holds. In this case, in virtue of Lemma 5.2, the problem $\mathrm{Sch}(R; \mu^\varepsilon, \nu^\varepsilon)$ is scalable. In particular, there exists $R^\varepsilon \in \Pi(\mu^\varepsilon, \nu^\varepsilon)$ such that R^ε and R have the same support. As the family (R^ε) has value in the compact set

$$\left\{R' \in \mathcal{M}_+(\mathcal{D}\times\mathcal{F}) \,\middle|\, \mathsf{M}(R') \leqslant \max(\mathsf{M}(\mu), \mathsf{M}(R))\right\},$$

we can chose one of its limit points $\overline{R}$ as $\varepsilon \to 0$. Obviously, $\overline{R} \ll R$ and $\overline{R} \in \Pi(\mu,\nu)$ so that $\mathrm{Sch}(R;\mu,\nu)$ is approximately scalable.

It remains to prove that (a) $\Rightarrow$ (c) even when Assumption 2.12 does not hold. To do so, we claim that under (a), assuming Assumption 2.12 is not restrictive. The reason is that (a) implies Assumption 2.11, and hence Assumption 2.12 up to restricting the problem to the supports of μ and ν, as explained in Section 2.3.

So let us prove that (a) implies Assumption 2.11. We suppose that (a) holds, and we consider $\mathcal{E}$ and R^0 as defined in Assumption 2.11.

Let us show that $\mu \ll \mu^{R^0}$. Let $x_i \in \mathcal{D}$ be such that $\mu_i > 0$, and let us show that $\mu_i^{R^0} > 0$. By (a), $\nu(F_R(\{x_i\})) \geqslant \mu_i > 0$. Therefore, $F_R(\{x_i\})$ is nonempty, and there exists $y_j \in F_R(\{x_i\})$ such that $\nu_j > 0$. This pair (x_i, y_j) belongs to $\mathcal{E}$, so $R^0_{ij} > 0$, and then $\mu_i^{R^0} > 0$.

Let us show that $\nu \ll \nu^{R^0}$. Let $y_j \in \mathcal{F}$ be such that $\nu_j > 0$, and let us show that $\nu_j^{R^0} > 0$. Let us call $\mathcal{D}'$ the support of μ. By (a), $\mathsf{M}(\nu) \geqslant \nu(F(\mathcal{D}')) \geqslant \mu(\mathcal{D}') = \mathsf{M}(\mu)$. But as $\mathsf{M}(\nu) = \mathsf{M}(\mu)$, we conclude that $\nu(F(\mathcal{D}')) = \mathsf{M}(\nu)$, and so in particular that $y_j \in F(\mathcal{D}')$. So there exists $x_i \in \mathcal{D}'$ such that $R_{ij} > 0$. This pair (x_i, y_j) belongs to $\mathcal{E}$, so $R^0_{ij} > 0$, and then $\nu_j^{R^0} > 0$. □

Bibliography

[1] E. Achilles, "Implications of convergence rates in Sinkhorn balancing", *Linear Algebra Appl.* **187** (1993), p. 109-112.

[2] J.-D. Benamou, G. Carlier, M. Cuturi, L. Nenna & G. Peyré, "Iterative Bergman projections for regularized transportation problems", *SIAM J. Sci. Comput.* **37** (2015), no. 2, p. A1111-A1138.

[3] A. Braides, *Gamma-convergence for Beginners*, Oxford Lecture Series in Mathematics and its Applications, vol. 22, Oxford University Press, 2002.

[4] R. A. Brualdi, "Convex sets of non-negative matrices", *Can. J. Math.* **20** (1968), p. 144-157.

[5] G. Carlier, V. Duval, G. Peyré & B. Schmitzer, "Convergence of entropic schemes for optimal transport and gradient flows", *SIAM J. Math. Anal.* **49** (2017), no. 2, p. 1385-1418.

[6] L. Chizat, G. Peyré, B. Schmitzer & F.-X. Vialard, "Scaling algorithms for unbalanced optimal transport problems", *Math. Comput.* **87** (2018), no. 314, p. 2563-2609.

[7] I. Csiszár, "I-divergence geometry of probability distributions and minimization problems", *Ann. Probab.* **3** (1975), p. 146-158.

[8] M. Cuturi, "Sinkhorn distances: Lightspeed computation of optimal transport", in *NIPS'13: Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, Advances in Neural Information Processing Systems, vol. 26, Curran Associates, Inc., 2013, p. 2292-2300.

[9] S. Di Marino & A. Gerolin, "An optimal transport approach for the Schrödinger bridge problem and convergence of Sinkhorn algorithm", *J. Sci. Comput.* **85** (2020), no. 2, article no. 27.

[10] H. Föllmer, "Random fields and diffusion processes", in *École d'Été de Probabilités de Saint-Flour XV-XVII, 1985-87*, Lecture Notes in Mathematics, vol. 1362, Springer, 1988, p. 101-203.

[11] R. Fortet, "Résolution d'un système d'équations de M. Schrödinger", *J. Math. Pures Appl.* **19** (1940), p. 83-105.

[12] I. Gentil, C. Léonard & L. Ripani, "About the analogy between optimal transport and minimal entropy", *Ann. Fac. Sci. Toulouse, Math. (6)* **26** (2017), no. 3, p. 569-600.

[13] P. Ghosal, M. Nutz & E. Bernton, "Stability of entropic optimal transport and Schrödinger bridges", *J. Funct. Anal.* **283** (2022), no. 9, article no. 109622.

[14] A. Hagberg, P. J. Swart & D. A. Schult, "Exploring network structure, dynamics, and function using NetworkX", in *Proceedings of the 7th Python in Science Conference (SciPy 2008)* (G. Varoquaux, J. Millman & T. Vaught, eds.), Los Alamos National Laboratory (LANL), Los Alamos, NM (United States), 2008, p. 11-16.

[15] U. Herbach, A. Bonnaffoux, T. Espinasse & O. Gandrillon, "Inferring gene regulatory networks from single-cell data: a mechanistic approach", *BMC Syst. Biol.* **11** (2017), article no. 105.

[16] M. Idel, "A review of matrix scaling and Sinkhorn's normal form for matrices and positive maps", 2016, `https://arxiv.org/abs/1609.06349`.

[17] S. Kondratyev, L. Monsaingeon & D. Vorotnikov, "A new optimal transport distance on the space of finite Radon measures", *Adv. Differ. Equ.* **21** (2016), no. 11-12, p. 1117-1164.

[18] H. Lavenant, S. Zhang, Y.-H. Kim & G. Schiebinger, "Towards a mathematical theory of trajectory inference", *Ann. Appl. Probab.* **34** (2024), no. 1A, p. 428-500.

[19] C. Léonard, "From the Schrödinger problem to the Monge–Kantorovich problem", *J. Funct. Anal.* **262** (2012), no. 4, p. 1879-1920.

[20] ———, "Some properties of path measures", in *Séminaire de Probabilités XLVI*, Lecture Notes in Mathematics, vol. 2123, Springer, 2014, p. 207-230.

[21] ———, "A survey of the Schrödinger problem and some of its connections with optimal transport", *Discrete Contin. Dyn. Syst., Ser. A* **34** (2014), no. 4, p. 1533-1574.

[22] M. Liero, A. Mielke & G. Savaré, "Optimal entropy-transport problems and a new Hellinger–Kantorovich distance between positive measures", *Invent. Math.* **211** (2018), no. 3, p. 969-1117.

[23] T. Mikami, "Monge's problem with a quadratic cost by the zero-noise limit of h-path processes", *Probab. Theory Relat. Fields* **129** (2004), no. 2, p. 245-260.

[24] M. Nutz, "Introduction to Entropic Optimal Transport", 2022, `https://www.math.columbia.edu/~mnutz/docs/EOT_lecture_notes.pdf`.

[25] M. Nutz & J. Wiesel, "Entropic optimal transport: convergence of potentials", *Probab. Theory Relat. Fields* **184** (2022), no. 1-2, p. 401-424.

[26] G. Peyré & M. Cuturi, "Computational Optimal Transport With Applications to Data Science", *Found. Trends Mach. Learn.* **11** (2019), no. 5-6, p. 355-607.

[27] L. Rüschendorf, "Convergence of the iterative proportional fitting procedure", *Ann. Stat.* **23** (1995), no. 4, p. 1160-1174.

[28] L. Rüschendorf & W. Thomsen, "Note on the Schrödinger equation and I-projections", *Stat. Probab. Lett.* **17** (1993), no. 5, p. 369-375.

[29] I. N. Sanov, "On the probability of large deviations of random variables", Tech. report, North Carolina State University. Dept. of Statistics, 1958.

[30] F. Santambrogio, *Optimal transport for applied mathematicians*, Progress in Nonlinear Differential Equations and their Applications, vol. 87, Birkäuser/Springer, 2015.

[31] G. Schiebinger, J. Shu, M. Tabaka, B. Cleary, V. Subramanian, A. Solomon, J. Gould, S. Liu, S. Lin, P. Berube, L. Lee, J. Chen, J. Brumbaugh, P. Rigollet, K. Hochedlinger, R. Jaenisch, A. Regev & E. S. Lander, "Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming", *Cell* **176** (2019), p. 928-943.

[32] E. Schrödinger, "Über die Umkehrung der Naturgesetze", *Sitzungsber. Preuß. Akad. Wiss., Phys.-Math. Kl.* **1931** (1931), no. 8-9, p. 144-153.

[33] ———, "Sur la théorie relativiste de l'électron et l'interprétation de la mécanique quantique", *Ann. Inst. Henri Poincaré* **2** (1932), no. 4, p. 269-310.

[34] R. Sinkhorn, "A relationship between arbitrary positive matrices and doubly stochastic matrices", *Ann. Math. Stat.* **35** (1964), no. 2, p. 876-879.

[35] ———, "Diagonal equivalence to matrices with prescribed row and column sums", *Am. Math. Mon.* **74** (1967), no. 4, p. 402-405.

[36] G. W. Soules, "The rate of convergence of Sinkhorn balancing", *Linear Algebra Appl.* **150** (1991), p. 3-40.

[37] E. Ventre, "Reverse engineering of a mechanistic model of gene expression using metastability and temporal dynamics", *In Silico Biol.* **14** (2021), no. 3–4, p. 89-113.

[38] ———, "Analyse, calibration et évaluation de modèles stochastiques d'expression des gènes", PhD Thesis, École Normale Supérieure de Lyon, France, 2022, `https://inserm.hal.science/INSA-LYON-THESES/tel-03848137v1`.

[39] E. Ventre, U. Herbach, T. Espinasse, G. Benoit & O. Gandrillon, "One model fits all: combining inference and simulation of gene regulatory networks", *PLoS Comput. Biol.* **19** (2023), no. 3, article no. e1010962.

[40] J.-C. Zambrini, "Variational processes and stochastic versions of mechanics", *J. Math. Phys.* **27** (1986), no. 9, p. 2307-2330.