From Knothe's transport to Brenier's map and a continuation method for optimal transport

G. Carlier[*], A. Galichon[†], F. Santambrogio[‡]

October 22, 2008

Abstract

A simple procedure to map two probability measures in $\mathbb{R}^d$ is the so-called *Knothe-Rosenblatt rearrangement*, which consists in rearranging monotonically the marginal distributions of the last coordinate, and then the conditional distributions, iteratively. We show that this mapping is the limit of solutions to a class of Monge-Kantorovich mass transportation problems with quadratic costs, with the weights of the coordinates asymptotically dominating one another. This enables us to design a continuation method for numerically solving the optimal transport problem.

Keywords: optimal transport, rearrangement of vector-valued maps, Knothe-Rosenblatt transport, continuation methods.

1 Introduction

Given two Borel probability measures μ and ν on $\mathbb{R}^d$, a Borel map $S : \mathbb{R}^d \to \mathbb{R}^d$ is said to be a transport map between μ and ν if $S\sharp\mu = \nu$ where $S\sharp\mu$ denotes the push-forward (or image measure) of μ through ν (i.e. $S\sharp\mu(B) = \mu(S^{-1}(B))$ for every Borel B). In the present article, we will focus on two particular transport maps: the Knothe-Rosenblatt transport and the Brenier's map.

[*]CEREMADE, Université Paris IX Dauphine carlier@ceremade.dauphine.fr

[†]Département d'Economie, Ecole Polytechnique, alfred.galichon@polytechnique.edu

[‡]CEREMADE, Université Paris IX Dauphine, filippo@ceremade.dauphine.fr

The Knothe-Rosenblatt transport. The Knothe-Rosenblatt rearrangement was independently proposed by Rosenblatt [6] for statistical purposes and by Knothe [4] in order to extend the Brunn-Minkowski inequalities. The principle is the following, as explained in Villani [8]. Let μ and ν be two Borel probability measures on $\mathbb{R}^d$ and assume for simplicity for the moment that μ is absolutely continuous with respect to the Lebesgue measure. Let us denote by μ^d (respectively ν^d) the d-th marginal of μ (respectively ν) and $\mu^{d-1}_{x_d}$, $\mu^{d-2}_{(x_d,x_{d-1})},\ldots,$ $\mu^1_{(x_d,\ldots,x_2)}$ (respectively $\nu^{d-1}_{y_d}$, $\nu^{d-2}_{(y_d,y_{d-1})},\ldots,$ $\nu^1_{(y_d,\ldots,y_2)}$) the successive disintegrations (or conditional measures) of μ (respectively ν) given x_d, (x_d,x_{d-1}),..., $(x_d,...,x_2)$ (respectively given y_d, (y_d,y_{d-1}), ..., $(y_d,...,y_2)$). Now let $T_d = T_d(x_d)$ be the monotone nondecreasing map transporting μ^d to ν^d, such a map is well-defined and unique as soon as μ_d has no atoms and in this case, it is explicitly given by $T_d = G_d^{-1} \circ F_d$ (with $F_d(\alpha) := \mu((-\infty,\alpha])$ and $G_d(\alpha) := \nu((-\infty,\alpha])$). Then let $T_{d-1} = T_{d-1}(x_{d-1},x_d)$ be such that $T_d(.,x_d)$ is monotone and maps $\mu^{d-1}_{x_d}$ to $\nu^{d-1}_{T_d(x_d)}$. One repeats the construction (well-known by statisticians under the name of conditional quantile transforms) iteratively until we define $T_1(x_1,x_2,...,x_d)$, which is monotone in x_1 and transports $\mu^1_{(x_d,\ldots,x_2)}$ onto $\nu^1_{T_2(x_2,\ldots,x_d)}$. Finally, the *Knothe-Rosenblatt rearrangement* T is defined by $T(x) = (T_1(x_1,x_2,...,x_d),...,T_{d-1}(x_{d-1},x_d),T_d(x_d))$. Obviously, T is a transport map from μ to ν, i.e. $T\sharp\mu = \nu$. By construction, the Knothe transport T has a triangular Jacobian matrix with nonnegative entries on its diagonal. Note also that the computation of the Knothe transport only involves one-dimensional monotone rearrangements and that it is well defined as soon the measures one transports have no atoms. The precise assumption is the following.

Assumption *(H-source)***:** the measure μ^d, as well as μ^d-almost all the measures $\mu^{d-1}_{x_d}$, and the measures $\mu^{d-2}_{x_d,x_{d-1}}$ for μ^d-a.e. x_d and $\mu^{d-1}_{x_d}-$a.e. $x_{d-1}\ldots$up to almost all the measures $\mu^1_{x_d,\ldots,x_2}$, which are all measures on the real line, must have no atoms.

Notice that *(H-source)* is satisfied as soon as μ is absolutely continuous with respect to the Lebesgue measure.

The Monge-Kantorovich problem and the Brenier's map. Optimal transportation theory provides an alternative way to transport μ to ν. We recall that in the case of the quadratic cost, the Monge-Kantorovich problem reads as

$$\inf_{\pi\in\Gamma(\mu,\nu)} \int_{\mathbb{R}^d\times\mathbb{R}^d} |x-y|^2 d\pi(x,y) \tag{1.1}$$

where $\Gamma(\mu,\nu)$ denotes the set of transport plans between μ and ν i.e. the set of probability measure on $\mathbb{R}^d\times\mathbb{R}^d$ with marginals μ and ν. We refer to

the books of Villani [7], [8] for a modern account of optimal transportation theory. The linear problem (1.1) is a relaxation of the Monge problem

$$\inf_{S\,:\,S\sharp\mu=\nu}\int_{\mathbb{R}^d}|x-S(x)|^2d\mu(x) \tag{1.2}$$

When μ is absolutely continuous with respect to the Lebesgue measure, Brenier [2] proved that (1.2) has a unique solution which is characterized by the fact that it is the gradient of some convex function. More precisely, there exists a unique (up to constants and μ-a.e. equivalence) convex function $V:\ \mathbb{R}^d\to\mathbb{R}$ such that $\nabla V\sharp\mu=\nu$. Also $(id\times\nabla V)\sharp\mu$ is the only solution of (1.1) and ∇V is called the Brenier's map between μ and ν.

The Knothe-Rosenblatt as a limit of optimal transportation plans. Let us slightly modify the quadratic cost in (1.1) and replace it with the weighted quadratic cost

$$c_\varepsilon(x,y):=\sum_{i=1}^d\lambda_i(\varepsilon)(x_i-y_i)^2$$

where the $\lambda_i(\varepsilon)$'s are positive scalars depending on a parameter $\varepsilon>0$. If μ is absolutely continuous with respect to the Lebesgue measure, the corresponding optimal transportation problem admits a unique solution T_ε. When in addition, for all $k\in\{1,...,d-1\}$, $\lambda_k(\varepsilon)/\lambda_{k+1}(\varepsilon)\to 0$ as $\varepsilon\to 0$, it is natural to expect the convergence of T_ε to the Knothe transport T. We will show that this convergence holds provided ν satisfies some additional condition, and namely

Assumption *(H-target)***:** the measure ν^d, as well as ν^d-almost all the measures $\nu^{d-1}_{y_d}$, and the measures $\nu^{d-2}_{y_d,y_{d-1}}$ for ν^d-a.e. y_d and $\nu^{d-1}_{y_d}-$a.e. $y_{d-1}\dots$ up to almost all the measures $\nu^2_{y_d,\dots,y_3}$, which are all measures on the real line, must have no atoms.

Notice that *(H-target)* is not natural as *(H-source)* is. Yet, we will show a counter-example to the convergence result when it is not satisfied. *(H-target)* as well is satisfied should ν be absolutely continuous (actually, this assumption is slightly weaker then *(H-source)*, since the last disintegration measures are not concerned).

This convergence result was conjectured by Y. Brenier as a very natural one, and actually its proof is not hard. Yet, it was not known before that extra assumptions on ν were needed. This makes one of the point of interest of this paper.

The other point is what we investigate later in the paper, i.e. the other direction: from Knothe to Brenier. We will study the dependence $\varepsilon\mapsto T_\varepsilon$ by

means of the evolution with respect to ε of the dual variables. This will enable us, to design a numericaly strategy to approximate all the optimal transports T_ε taking as initial condition the (cheap to compute) Knothe transport T.

An example. To illustrate the problem in a particular case where explicit solutions are available, take $d = 2$, and μ and ν two Gaussian measures where $\mu = N(0, I_2)$ and $\nu = N\left(0, \begin{pmatrix} a & b \\ b & c \end{pmatrix}\right)$. Take $\lambda_1(\varepsilon) = \varepsilon$ and $\lambda_2(\varepsilon) = 1$. Then it can be verified that T_ε is linear, and that its matrix in the canonical basis of $\mathbb{R}^2$ is

$$T_\varepsilon = \frac{1}{\sqrt{a\varepsilon^2 + c + 2\varepsilon\sqrt{ac - b^2}}} \begin{pmatrix} a\varepsilon + \sqrt{ac - b^2} & b \\ b\varepsilon & c + \varepsilon\sqrt{ac - b^2} \end{pmatrix}$$

which converges as $\varepsilon \to 0$ to $T = \begin{pmatrix} \sqrt{a - b^2/c} & b/\sqrt{c} \\ 0 & \sqrt{c} \end{pmatrix}$, which is precisely the matrix of the Knothe transport from μ to ν.

Organization of the paper. In section 2, we show, under suitable assumptions, that the optimal transportation maps for the cost c_ε converge to Knothe's transport map as the parameter ε goes to 0, we will also emphasize that some conditions are to be imposed on ν for the convergence to hold. In section 3, we show that the evolution of the dual variables in the optimal transportation problem for cost the c_ε is given by a well-posed ordinary differential equation. Finally in section 4, we discretize this equation and give several numerical results.

2 Knothe transport as a limit of quadratic optimal transports

We directly state our first result, whose proof, in the spirit of $\Gamma-$convergence developments (see [1]), will require several steps.

Theorem 2.1. *Let μ and ν be two probability measures on $\mathbb{R}^d$ satisfying (H-source) and (H-target), respectively, with finite second moments, and γ_ε be an optimal transport plan for the costs $c_\varepsilon(x, y) = \sum_{i=1}^d \lambda_i(\varepsilon)(x_i - y_i)^2$, for some weights $\lambda_k(\varepsilon) > 0$. Suppose that for all $k \in \{1, ..., d-1\}$, $\lambda_k(\varepsilon)/\lambda_{k+1}(\varepsilon) \to 0$ as $\varepsilon \to 0$. Let T be the Knothe-Rosenblatt map between μ and ν and $\gamma_K \in \mathcal{P}(\mathbb{R}^d \times \mathbb{R}^d)$ the associated transport plan (i.e. $\gamma_K := (id \times T)\sharp\mu$). Then $\gamma_\varepsilon \rightharpoonup \gamma_K$ as $\varepsilon \to 0$.*

Moreover, should the plans γ_ε be induced by transport maps T_ε, then these maps would converge to T in $L^2(\mu)$ as $\varepsilon \to 0$.

Proof. Take the plans γ_ε that are optimal for the Brenier-like cost c_ε given by

$$c_\varepsilon(x,y) = \sum_{i=1}^{d} \lambda_i(\varepsilon)(x_i - y_i)^2$$

(we suppose for simplicity $\lambda_d(\varepsilon) = 1$ and $\lambda_i(\varepsilon)/\lambda_{i+1}(\varepsilon) \to 0$). Suppose (which is possible, up to subsequences) $\gamma_\varepsilon \rightharpoonup \gamma$. We want to prove $\gamma = \gamma_K$.

By comparing γ_ε to γ_K and using optimality we first get

$$\int c_\varepsilon \, d\gamma_\varepsilon \leq \int c_\varepsilon \, d\gamma_K \tag{2.1}$$

and, passing to the limit as $\varepsilon \to 0$, since c_ε converges locally uniformly to $c^{(d)}(x,y) = (x_d - y_d)^2$, we get

$$\int c^{(d)} d\gamma \leq \int c^{(d)} d\gamma_K.$$

Yet, the function $c^{(d)}$ only depends on the variables x_d and y_d and this shows that the measure $(\pi_d)_\sharp \gamma$ gets a better result than $(\pi_d)_\sharp \gamma_K$ with respect to the quadratic cost (π_d being the projection onto the last coordinates, i.e. $\pi_d(x,y) = (x_d, y_d)$). Yet, the measure γ_K has been chosen on purpose to get optimality from μ_d to ν_d with respect to this cost, and the two measures share the same marginals. Moreover, thanks to the assumptions on μ_d, this optimal transport plan (which is actually induced by a transport map) is unique. This implies $(\pi_d)_\sharp \gamma = (\pi_d)_\sharp \gamma_K$. Let us call γ^d this common measure.

We go back to (2.1) and go on by noticing that all the measures γ_ε have the same marginals as γ_K and hence their (separate) projection onto x_d and y_d are μ_d and ν_d, respectively. This implies that $(\pi_d)_\sharp \gamma_\varepsilon$ must realize a result which is worse than $(\pi_d)_\sharp \gamma_K$ as far as the quadratic cost is concerned and consequently we have

$$\begin{aligned}
&\int |x_d - y_d|^2 d(\pi_d)_\sharp \gamma_K(x_d, y_d) + \sum_{i=1}^{d-1} \lambda_i(\varepsilon) \int (x_i - y_i)^2 d\gamma_\varepsilon \\
&\leq \int c_\varepsilon \, d\gamma_\varepsilon \leq \int c_\varepsilon \, d\gamma_K \\
&= \int |x_d - y_d|^2 d(\pi_d)_\sharp \gamma_K(x_d, y_d) + \sum_{i=1}^{d-1} \lambda_i(\varepsilon) \int (x_i - y_i)^2 d\gamma_K,
\end{aligned}$$

which implies, by simplifying the common term in $d(\pi_d)_\sharp \gamma_K$, dividing by $\lambda_{d-1}(\varepsilon)$ and passing to the limit,

$$\int c^{(d-1)} d\gamma \leq \int c^{(d-1)} d\gamma_K$$

(we use the general notation $c^k(x,y) = |x_k - y_k|^2$). We can notice that both integrals depend on the variables x_{d-1} and y_{d-1} only. Anyway, we can project onto the variables (x_{d-1}, x_d) and (y_{d-1}, y_d) (obtaining measures $(\pi_{d-1})_\sharp\gamma$ and $(\pi_{d-1})_\sharp\gamma_K$) so that we disintegrate with respect to the measure γ^d. We have

$$\begin{aligned}&\int d\gamma^d(x_d,y_d)\int |x_{d-1}-y_{d-1}|^2 d\gamma^{d-1}_{(x_d,y_d)}(x_{d-1},y_{d-1})\\ \leq &\int d\gamma^d(x_d,y_d)\int |x_{d-1}-y_{d-1}|^2 d\gamma^{d-1}_{(x_d,y_d),K}(x_{d-1},y_{d-1}).\end{aligned} \tag{2.2}$$

It is is sufficient to prove that the measures $\gamma^{d-1}_{(x_d,y_d)}$ share the same marginals on x_{d-1} and y_{d-1} as the corresponding $\gamma^{d-1}_{(x_d,y_d),K}$ to get that their quadratic performance should be worse than the corresponding performance of $\gamma^{d-1}_{(x_d,y_d),K}$ (this is because the Knothe measure has been chosen exactly with the intention of being quadratically optimal on (x_{d-1}, y_{d-1}) once x_d and y_d are fixed). Yet, (2.2) shows that, on average, the result given by the those measures is not worse than the results of the optimal ones. Thus, the two results coincide for almost any pair (x_d, y_d) and, by uniqueness of the optimal transports (this relies on the assumptions on the measures $\mu^{d-1}_{x_d}$), we get $\gamma^{d-1}_{(x_d,y_d)} = \gamma^{d-1}_{(x_d,y_d),K}$. To let this proof work it is sufficient to prove that the projections of the two measures coincide for γ^d-a.e. pair (x_d, y_d). For fixed (x_d, y_d) we would like to prove, for any ϕ

$$\int \phi(x_{d-1})d\gamma^{d-1}_{(x_d,y_d)} = \int \phi(x_{d-1})d\gamma^{d-1}_{(x_d,y_d),K}$$

(and to prove an analogous equality for functions of y_{d-1}). Since we accept to prove it for a.e. pair (x_d, y_d), it is sufficient to prove this equality:

$$\begin{aligned}&\int d\gamma^d(x_d,y_d)\psi(x_d,y_d)\int \phi(x_{d-1})d\gamma^{d-1}_{(x_d,y_d)}\\ =&\int d\gamma^d(x_d,y_d)\psi(x_d,y_d)\int \phi(x_{d-1})d\gamma^{d-1}_{(x_d,y_d),K}\end{aligned}$$

for any ϕ and any ψ. This means proving

$$\int \psi(x_d,y_d)\phi(x_{d-1})d\gamma^{d-1} = \int \psi(x_d,y_d)\phi(x_{d-1})d\gamma^{d-1}_K,$$

which is not trivial since we only know that the two measures γ^{d-1} and γ^{d-1}_K have the same marginals with respect to the pairs (x_{d-1}, x_d), (y_{d-1}, y_d) (since they have the same projections onto x and onto y) and (x_d, y_d) (since we just proved it). But here there is a function of the three variables (x_{d-1}, x_d, y_d).

Yet, we know that the measure γ^d is concentrated on the set $y_d = T_d(x_d)$ for a certain map T_d, and this allows to replace the expression of y_d, thus getting rid of one variable. This proves that the function $\psi(x_d, y_d)\phi(x_{d-1})$ is actually a function of (x_{d-1}, x_d) only, and that equality holds when passing from γ to γ_K.The same can be performed on functions $\psi(x_d, y_d)\phi(y_{d-1})$ but we have in this case to ensure that we can replace x_d with a function of y_d, i.e. that we can invert T_d. This is possible thanks to the assumption on ν_d, since T_d is the optimal transport from μ_d to ν_d, but an optimal transport exists in the other direction as well and it gives the same optimal plan (thanks to uniqueness). These facts prove that the measures $\gamma^{d-1}_{(x_d,y_d)}$ and $\gamma^{d-1}_{(x_d,y_d),K}$ have the same marginals and hence, since they are both optimal, they coincide for a.e. pair (x_d, y_d). This implies $\gamma^{d-1} = \gamma^{d-1}_K$.

Now, it is possible to go on by steps: once we have proven that $\gamma^h = \gamma^h_K$, let us take (2.1) and estimate all the terms with $|x_i - y_i|^2$ and $i \geq h$ thanks to the optimality of γ_K, thus getting

$$\begin{aligned}
&\sum_{i\geq h} \lambda_i(\varepsilon) \int |x_i - y_i|^2 d\gamma_K + \sum_{i=1}^{h-1} \lambda_i(\varepsilon) \int (x_i - y_i)^2 d\gamma_\varepsilon \\
&\leq \int c_\varepsilon \, d\gamma_\varepsilon \leq \int c_\varepsilon \, d\gamma_K \\
&= \sum_{i\geq h} \lambda_i(\varepsilon) \int |x_i - y_i|^2 d\gamma_K + \sum_{i=1}^{h-1} \lambda_i(\varepsilon) \int (x_i - y_i)^2 d\gamma_K,
\end{aligned}$$

and consequently, by dividing by $\lambda_{h-1}(\varepsilon)$ and passing to the limit,

$$\int c^{(h-1)} d\gamma \leq \int c^{(h-1)} d\gamma_K.$$

We disintegrate with respect to γ^h and we act exacly as before: proving that the marginals of the disintegrations coincide is sufficient to prove equality of the measures. Here we will use test-functions fo the form

$$\psi(x_h, x_{h+1}, \dots, x_d, y_h, y_{h+1}, \dots, y_d)\phi(x_{h-1})$$

and

$$\psi(x_h, x_{h+1}, \dots, x_d, y_h, y_{h+1}, \dots, y_d)\phi(y_{h-1}).$$

The same trick as before, i.e. replacing the variables y with functions of the variables x is again possible. To invert the trick and replace x with y one needs to invert part of Knothe's transport. This is possible since our assumptions imply that all the monotone transports we get are invertible.

In the end we get, as before, $\gamma^{h-1} = \gamma_K^{h-1}$. This procedure may go on up to $h = 2$, thus arriving at $\gamma = \gamma_K$.

We have now proven $\gamma_\varepsilon \rightharpoonup \gamma_K$. Yet, if all these transport plans come from transport maps, it is well known that $(T_\varepsilon \times id)_\sharp \mu \rightharpoonup (T \times id)_\sharp \mu$ implies $T_\varepsilon \to T$ in $L^p(\mu)$, for any $p > 1$, as far as T_ε is bounded in $L^p(\mu)$. Actually, weak convergence is a simple consequence of boundedness: to go on, we can look at Young's measures. The assumption (the limit is a transport map as well) exactly means that all the Young measures are dirac masses, which implies strong convergence. In particular we get $L^2(\mu)$ convergence and μ-a.e. convergence on a subsequence. □

Let us remark here that if instead of considering the quadratic cost c_ε, one considers the more general separable cost

$$c_\varepsilon(x, y) := \sum_{i=1}^{d} \lambda_i(\varepsilon) c_i(x_i - y_i)$$

where each c_i is a smooth strictly convex function (with suitable growth), then the previous convergence proof carries over.

A counterexample when the measures have atoms We now show that interestingly, and perhaps counterintuitively, the hypothesis of absence of atoms in theorem 2.1 is necessary not only for μ, but also for ν. We propose a very simple example in $\mathbb{R}^2$ where μ is absolutely continuous with respect to the Lebesgue measure but ν does not satisfy *(H-target)*, and we show that the conclusion of theorem 2.1 fails to hold. On the square $\Omega := [-1, 1] \times [-1, 1]$, define μ such that $\mu(dx) = 1_{\{x_1 x_2 < 0\}} dx/2$ so that the measure μ is uniformly spread on the upper left and the lower right quadrants, and $\nu = \mathcal{H}^1_{|S}/2$, being S the segment $[-1, 1] \times \{0\}$.

The Knothe-Rosenblatt map is easily computed as $(y_1, y_2) = T(x) := (2(x_1 + sgn(x_2)), 0)$. The solution of any symmetric transportation problem with $\lambda^\varepsilon = (\varepsilon, 1)$ is $(y_1, y_2) = T^0(x) := (x_1, 0)$ (no transport may do better than this one, which projects on the support of ν). Therefore, in this example the optimal transportation maps fail to tend to the Knothe-Rosenblatt map. The reason is the atom in the measure $\nu^2 = \delta_0$.

Convergence even with atoms The convergence result of theorem 2.1 requires the absence of atoms in the projections of ν. This is obviously not the case if ν itself is purely atomic! Yet, this will precisely be the case we will consider in the algorithm we propose in the sequel. The same proof may be extended to this case under the following assumption. Keep the same assumptions on μ but suppose that ν is concentrated on a set S with the property

$$y, z \in S, \quad y \neq z \Rightarrow y_d \neq z_d.$$

This means that, if we restrict ourselves to S, then all the variables y_i for $i < d$ are actually a function of the last variable y_d. This is particularly useful when ν is purely atomic, concentrated on a finite (or countable) set of points with different y_d components.

Just come back to the proof. The equality

$$\begin{aligned}&\int d\gamma^d(x_d, y_d)\psi(x_d, y_d) \int \phi(x_{d-1}) d\gamma^{d-1}_{(x_d,y_d)}\\ &= \int d\gamma^d(x_d, y_d)\psi(x_d, y_d) \int \phi(x_{d-1}) d\gamma^{d-1}_{(x_d,y_d),K}\end{aligned}$$

only relied on y_d being a function of x_d, which is still true. The other equality, namely

$$\begin{aligned}&\int d\gamma^d(x_d, y_d)\psi(x_d, y_d) \int \phi(y_{d-1}) d\gamma^{d-1}_{(x_d,y_d)}\\ &= \int d\gamma^d(x_d, y_d)\psi(x_d, y_d) \int \phi(y_{d-1}) d\gamma^{d-1}_{(x_d,y_d),K}\end{aligned}$$

gives some extra troubles. It is not any more true that x_d is a function of y_d. Yet, it is true that y_{d-1} is a function of y_d and this allows us to reduce the expression to functions of (x_d, y_d) only, which is sufficient to get equality. The same procedure may be performed at subsequent steps as well.

3 An ODE for the dual variables

In this section, we consider for simplicity the case $d = 2$ (although our method extends to higher dimensions), μ uniform on some convex polyhedron Ω (for the sake of simplicity we will assume $|\Omega| = 1$) and $\nu = \frac{1}{N}\sum_{i=1}^N \delta_{y_i}$ where all the points $y_i \in \Omega$ have a different second coordinate $y_i^{(2)}$. For every $\varepsilon \geq 0$, let A_ε be the diagonal 2×2 matrix with diagonal entries $(\varepsilon, 1)$ and let c_ε be the quadratic cost defined by $c_\varepsilon(x, y) = A_\varepsilon(x - y)(x - y)$. We are interested in solving the family of optimal transportation problems

$$\inf_{\pi\in\Gamma(\mu,\nu)} \int_{\mathbb{R}^d\times\mathbb{R}^d} c_\varepsilon(x, y) d\pi(x, y) \tag{3.1}$$

for all values of the parameter $\varepsilon \in [0, 1]$. It is well-known, that (3.1) can be conveniently solved by the dual problem formulated in terms of prices:

$$\sup_p \Phi(p, \varepsilon) := \frac{1}{N}\sum_{i=1}^N p_i + \int_\Omega p^*_\varepsilon(x) dx, \tag{3.2}$$

where $p^*_\varepsilon(x) = \min_i\{c_\varepsilon(x, y_i) - p_i\}$ and we impose as a normalization $p_1 = 0$. For each ε, there is a unique maximizer $p(\varepsilon)$. For each (p, ε) we define $C(p, \varepsilon)_i = \{x \in \Omega \,:\, \inf_j c_\varepsilon(x, y_j) - p_j = c_\varepsilon(x, y_i) - p_i\}$. It is easy to check that $\Phi_\varepsilon := \Phi(., \varepsilon)$ is concave differentiable and that the gradient of Φ_ε is given by

$$\frac{\partial \Phi_\varepsilon}{\partial p_i}(p) = \frac{1}{N} - |C(p, \varepsilon)_i|.$$

By concavity of Φ_ε, the solution $p(\varepsilon)$ of (3.2) is characterized by the equation $\nabla\Phi_\varepsilon(p(\varepsilon)) = 0$. The optimal transportation between μ and ν for the cost c_ε is then the piecewise map taking value y_i in the cell $C(p(\varepsilon), \varepsilon))_i$. Our aim is to charcacterize the evolution of $p(\varepsilon)$ as ε varies. Formally, differentiating the optimality condition $\nabla\Phi(p(\varepsilon), \varepsilon) = 0$, we obtain a differential equation for the evolution of $p(\varepsilon)$:

$$\frac{\partial}{\partial \varepsilon}\nabla_p\Phi(p(\varepsilon), \varepsilon) + D^2_{p,p}\Phi(p(\varepsilon), \varepsilon) \cdot \frac{dp}{d\varepsilon}(\varepsilon) = 0. \tag{3.3}$$

Our aim now is to show that the equation (3.3) is well-posed starting with the initial condition $p(0)$ (corresponding to horizontal cells of area $1/N$); this will involve computing the second derivatives of Φ in (3.3), proving their Lipschitz behavior as well as obtaining a negative bound on the larger eigenvalue of the negative semidefinite matrix $D^2_{p,p}\Phi$.

The price vector $p(\varepsilon)$, along the evolution, will always be such that all the areas $|C(p, \varepsilon)_i|$ are equal (and are equal to $1/N$). Yet, we need to prove that the differential equation is well posed and we will set it in an open set,

$$\mathcal{O} = \left\{(p, \varepsilon) \,:\, \frac{1}{2N} < |C(p, \varepsilon)_i| < \frac{2}{N} \text{ for every } i\right\}. \tag{3.4}$$

The initial datum of the equation will be such that $|C(p(0), 0)_i| = 1/N$ and we will look at the solution only up to the first moment where it exits $\mathcal{O}$. Yet inside the set it will be well-posed and it will imply conservation of the areas. Hence we will never exit $\mathcal{O}$.

All the quantities we are interested in depend on the position of the vertices of the cells $C(p, \varepsilon)_i$, which are all polygons. Let us call $x(p, \varepsilon)^\pm_{i,j}$ the two extremal points of the common boundary between $C(p, \varepsilon)_i$ and $C(p, \varepsilon)_j$, that we call $D(\varepsilon, p)_{i,j}$ (if such a common boundary exists; should it be a single point we consider the two points as coinciding). Each one of this points is obtained as the intersection of at least this common boundary with another one, or with the boundary of Ω (which is supposed to be a polygon as well, so that in the neighbourhood of almost any point locally the boundary is a line). We want to investigate the dependence of these points with respect

to (p,ε) and prove that this dependence is Lipschitz. Notice that each point $x(p,\varepsilon)^{\pm}_{i,j}$ is not defined for any value of (p,ε) but only on a certain (closed) subset of the space $\mathbb{R}^N \times (0,1)$.

Lemma 3.1. *The positions of the vertices $x(p,\varepsilon)^{\pm}_{i,j}$ depend in a Lipschitz way on p and ε.*

Proof. Locally it is true that the same point $x(p,\varepsilon)^{\pm}_{i,j}$ is defined either by a system of equations

$$\begin{cases} A_\varepsilon(x-y_i)(x-y_i)-p_i = A_\varepsilon(x-y_j)(x-y_j)-p_j, \\ A_\varepsilon(x-y_i)(x-y_i)-p_i = A_\varepsilon(x-y_h)(x-y_h)-p_h \end{cases} \tag{3.5}$$

in the case of a point on the intersection of two common boundaries, or by a system

$$\begin{cases} A_\varepsilon(x-y_i)(x-y_i)-p_i = A_\varepsilon(x-y_j)(x-y_j)-p_j, \\ Lx = l_0 \end{cases} \tag{3.6}$$

in the case of intersection with the boundary $\partial\Omega$ (locally being given by the equation $Lx = l_0$) . The first system, after simplifying, reads as

$$\begin{cases} 2A_\varepsilon(y_j-y_i)(x) = A_\varepsilon(y_j)(y_j) - A_\varepsilon(y_i)(y_i) - p_j + p_i, \\ 2A_\varepsilon(y_h-y_i)(x) = A_\varepsilon(y_h)(y_h) - A_\varepsilon(y_i)(y_i) - p_h + p_i, \end{cases} \tag{3.7}$$

and the second as well may be simplified the same way. This means that they are of the form

$$M(\varepsilon)x = V(\varepsilon,p)$$

for a matrix $M(\varepsilon)$ which reads, in usual coordinates,

$$M(\varepsilon) = \begin{pmatrix} \varepsilon(y_j^{(1)}-y_i^{(1)}) & y_j^{(2)}-y_i^{(2)} \\ \varepsilon(y_h^{(1)}-y_i^{(1)}) & y_h^{(2)}-y_i^{(2)} \end{pmatrix} \text{ in the first case and}$$

$$M(\varepsilon) = \begin{pmatrix} \varepsilon(y_j^{(1)}-y_i^{(1)}) & y_j^{(2)}-y_i^{(2)} \\ l_1 & l_2 \end{pmatrix} \text{ in the second.}$$

The vector $V(\varepsilon,p)$ is obtained regarding the right hand sides of the system, as in (3.7). Both M and V depend Lipschitzly on (ε,p), with uniformly bounded Lipschitz constants. Hence, to check that the dependence of x on (ε,p) is Lipschitz, we only need to bound (away from zero) $\det M(\varepsilon)$. The determinant of a 2×2 matrix is given by the product of the modulus of the two vector composing its lines, times the sinus of the angle between them.

In the first case the vectors are $A_\varepsilon(y_j - y_i)$ and $A_\varepsilon(y_h - y_i)$, while in the second they are $A_\varepsilon(y_j - y_i)$ and $l = (l_1, l_2)$. In both cases they are the normal vectors to the sides of the cell we are considering. This implies, thanks to the lower bound on the areas of the cells, that the angles between these vectors may not be too small. Actually, since Ω is bounded and the cells are convex, the area of each cell is smaller than $(\mathrm{diam}\Omega)^2 \sin\alpha/2$, α being any angle between two neighbour sides. This implies a lower bound on $\sin\alpha$. The lower bound on the moduli comes from the fact that the vectors l are chosen with modulus one and the modulus of $A_\varepsilon(y_j - y_i)$ is always greater than its vertical component, which is $y_j^{(2)} - y_i^{(2)}$, which is supposed different from zero for any pair (i, j). Notice that in the first case the matrix has ε in the determinant, even if we proved a lower bound on such a determinant, independent on ε: this agrees with the fact that actually, this kind of crossing (between two common boundaries of two cells) will only happen for $\varepsilon \geq \varepsilon_0$ (for $\varepsilon < \varepsilon_0$ we only have almost horizontal strips crossing non-horizontal sides of Ω). □

Lemma 3.2. *The function Φ admits pure second derivatives with respect to p and mixed second derivatives with respect to p and ε, and these second derivatives are Lipschitz continuous.*

Proof. We have proven that the positions of the points $x(p,\varepsilon)^{\pm}_{i,j}$ depend Lipschitzly, with uniformly bounded Lipschitz constants, on p and ε. Since the volumes of the cells $C(p,\varepsilon)_i$ are Lipschitz functions of these points, this implies that $\nabla_p\Phi$ is $C^{0,1}$. Hence it admits derivatives almost everywhere and we can compute them in the following way.

The derivative of a volume of a polygonal cell is given, side by side, by the length of the side times the average of the components which are normal to such a side of the two derivatives $\dot{x}(p,\varepsilon)^+_{i,j}$ and $\dot{x}(p,\varepsilon)^-_{i,j}$ (the other terms - which are mainly the terms at the corners - are of higher order). The equation of the side $D(\varepsilon,p)_{i,j}$ is, as we know,

$$2A_\varepsilon(y_j - y_i)(x) = A_\varepsilon(y_j)(y_j) - A_\varepsilon(y_i)(y_i) - p_j + p_i \tag{3.8}$$

and the normal unit vector to the side is

$$n = \frac{A_\varepsilon(y_j - y_i)}{|A_\varepsilon(y_j - y_i)|}.$$

Let us start from the derivatives of the cell $C(p,\varepsilon)_i$ with respect to a variable p_j with $j \neq i$. We differentiate (3.8) with respect to p_j and we get

$$2A_\varepsilon(y_j - y_i)(\dot{x}) = -1.$$

This formula only works for $x = x(p,\varepsilon)^+_{i,j}$ and $x = x(p,\varepsilon)^-_{i,j}$. Obviously it ony works where they are differentiable, i.e. almost everywhere. Hence the derivative, by summing up and rescaling the normal vector, is given by

$$\frac{\partial |C(p,\varepsilon)_i|}{\partial p_j} = -\frac{l_{i,j}}{2|A_\varepsilon(y_j - y_i)|}, \tag{3.9}$$

where $l_{i,j}$ is the length of $D(\varepsilon,p)_{i,j}$.

As far as the derivative with respect to p_i is concerned, it is not difficult to check that we have (by summing up the results on every side)

$$\frac{\partial |C(p,\varepsilon)_i|}{\partial p_i} = \sum_j \frac{l_{i,j}}{2|A_\varepsilon(y_j - y_i)|}, \tag{3.10}$$

where the sum is performed on all the indices j such that the cell $C(p,\varepsilon)_j$ is in contact with the cell $C(p,\varepsilon)_i$.

Automatically, since these derivatives only depend on the values of $l_{i,j}$, which depend in a Lipschitz manner on the positions of $x = x(p,\varepsilon)^\pm_{i,j}$, they are Lipschitz as well. This proves that Φ_ε is actually $C^{2,1}$ and that these derivatives are well defined and admit the previous expressions (3.9)-(3.10) everywhere.

The computation of the derivatives with respect to ε is a bit trickier. We derive again (3.8), but with respect to ε. Since $dA_\varepsilon/d\varepsilon = B$, we get

$$\begin{aligned} 2A_\varepsilon(y_j - y_i)(\dot{x}) &= -2B(y_j - y_i)(x) + B(y_j)(y_j) - B(y_i)(y_i) \\ &= 2B(y_j - y_i)\left(\frac{y_j + y_i}{2} - x\right). \end{aligned}$$

Then we renormalize the normal vector, sum up the results for $x = x(p,\varepsilon)^+_{i,j}$ and $x = x(p,\varepsilon)^-_{i,j}$, multiply by the lengths and sum up the results for all the sides, and get

$$\frac{\partial |C(p,\varepsilon)_i|}{\partial \varepsilon} = \sum_j l_{i,j} \frac{B(y_j - y_i)(y_j + y_i - x(p,\varepsilon)^+_{i,j} - x(p,\varepsilon)^-_{i,j})}{2|A_\varepsilon(y_j - y_i)|}. \tag{3.11}$$

In this case as well the result is Lipschitz in (p,ε) and hence $\nabla_p \Phi$ is differentiable everywhere with respect to ε, with Lipschitz derivative.

□

We can come now back to the evolution of $p = p(\varepsilon)$ and consider again the differential equation (3.3). To solve this equation we need to prove that the matrix $D^2_{p,p}\Phi$ is actually invertible (for numerical purpose, we will also

need to bound its eigenvalues away from zero). It is important to recall that we look at the evolution of the vector $p = (p_2, \dots, p_N)$, since we may assume $p_1(\varepsilon) = 0$ for all ε. Hence, we will not look at the entries 1 in the vectors or the matrices. The matrix we consider is $M := -(D^2_{p,p}\Phi)_{i,j=2,\dots,N}$ has the following properties:

- on each line, outside the diagonal we have negative terms $M_{i,j} = -\frac{l_{i,j}}{2|A_\varepsilon(y_j - y_i)|}$;
- each element on the diagonal is the sum of minus all the others on the same line (hence it is positive), and possibly of the term which should be in the same line at the first column;
- an entry (i, j) of the matrix is non-zero if and only if the cells $C(p, \varepsilon)_i$ and $C(p, \varepsilon)_j$ share a common boundary with positive length;
- in particular, for any pair (i, j), even if the entry at place (i, j) is zero, it is possible to find a path $i = i_0, i_1, i_2, \dots, i_k = j$ so that the matrix has non-zero values at all the positions (i_h, i_{h+1});
- the total of the entries of the first column (the one which is not present in the matrix) is strictly positive.

The invertibility of M is ensured by the following:

Lemma 3.3. *Let the matrix M satisfy the following properties*

$(H1)$ *for all* i, $\quad M_{i,i} \geq \sum_{j \neq i} |M_{i,j}|$,

$(H2)$ *there exists* i *such that* $M_{i,i} > \sum_{j \neq i} |M_{i,j}|$,

$(H3)$ *for any pair* (i, j) *there is a sequence* $i_0, i_1, i_2, \dots, i_k$ *with* $i_1 = i$, $i_k = j$, *and* $M_{i_h, i_{h+1}} \neq 0$.

then M is invertible.

Proof. Let $x \in \mathrm{Ker}(M)$ and let $\bar{i}$ be an index such that $|x_i|$ is maximal. We may suppose for simplicity that $x_{\bar{i}}$ is positive. Then we have

$$0 = M_{\bar{i},\bar{i}} x_{\bar{i}} - \sum_j M_{\bar{i},j} x_j \geq M_{\bar{i},\bar{i}} x_{\bar{i}} - \sum_j M_{\bar{i},j} x_{\bar{i}} = x_{\bar{i}} \left(M_{\bar{i},\bar{i}} - \sum_j M_{\bar{i},j} \right) \geq 0.$$

This implies that all inequalities are equalities and in particular $x_j = x_{\overline{i}}$ whenever $M_{\overline{i},j} \neq 0$. Hence, the entries of x on all the indices which are "neighbours" of $\overline{i}$ equal $x_{\overline{i}}$ (and they are maximal as well). This allows to repeat the argument replacing $\overline{i}$ with another maximizing index j and so on... since any index is connected by a chain of neighbours to $\overline{i}$, we get that all the entries are equal. But this implies that the vector in the kernel we selected must be a multiple of the vector $(1, 1, \ldots, 1)$. Yet, this vector is not in the kernel since the sum of the elements on each line is not zero for all lines, by assumption $(H2)$. This proves that M is invertible. □

Finding a lower bound for the modulus of the eigenvalues of M, i.e. quantifying its inversibility is not straightforward. Indeed, the properties $(H1)$, $(H2)$, $(H3)$ are not sufficient to get this bound, even if we fix the norm of the remainding column as the following counter-example shows. The determinant of the matrices

$$M_\varepsilon = \begin{pmatrix} 1 & -\varepsilon & 0 \\ -\varepsilon & 1 & -(1-\varepsilon) \\ 0 & -(1-\varepsilon) & 1 \end{pmatrix}$$

is $2\varepsilon(1-\varepsilon) \to 0$, which implies that some eigenvalue as well goes to 0.

We will obtain a positive lower bound by a compactness argument, but we will use something stronger than simply assumptions $(H1)$, $(H2)$, $(H3)$. The idea is that assumption $(H3)$ is not closed, but it stays closed when we replace it with the stronger condition of the matrix being associated to a partition (as is the case for $M = -D^2_{p,p}\Phi$). In this case if one connection degenerates (i.e. a common boundary reduces to a point), some other connections will play the role.

Lemma 3.4. *There is a positive uniform lower bound on the least eigenvalue of any matrix M associated to the cell partition corresponding to a pair $(p, \varepsilon) \in \mathcal{O}$.*

Proof. The proof will be obtained by contradiction. To this aim, take a sequence of partitions of Ω into sets $(\Omega_i^n)_{i=1,\ldots,N}$. We assume these sets to be convex polygons with a bounded number of sides. This is the case for the cells associated to pairs (p, ε). We also know that their areas are always bounded between $1/2N$ and $2/N$). These partitions give rise to a certain topology of connections between the cells. Up to subsequences, we may suppose that this topology is always the same on all the partitions of the sequence (since the number of possible topologies is finite). Up to subsequences, we also have convergence in the Hausdorff distance. This means that for any i we

have $\Omega_i^n \to \Omega_i$, and this convergence, which is the same as the convergence of all the vertices, preserves the areas, the convexity, the upper bound on the number of sides, the fact of being a partition... The matrices associated to these partitions depend continuously on these sets (with respect to this convergence, since they actually depend on the positions of the vertices). Notice that it is possible that a side reduces its length along the sequence up to becoming a single point in the limit. Yet, two cells which share a boundary along the sequence will do it along the whole sequence (thanks to our choice of not changing the topology) and at the limit, either they share a side as well, or they share a point only, but in this case the terms $l_{i,j}$ converged to zero. Hence we can associate to all the partitions their matrices and this correspondence is continuous. We have a sequence of matrices $M_n \to M$ and let us suppose that some eigenvalue $\lambda_1^{(n)}$ goes to zero. This would imply that the matrix M is associated to a partition but has a zero eigenvalue. This is not possible, thanks to the proof of lemma 3.3 above. M is associated to a partition and hence satisfies assumptions $(H1)$ and $(H3)$. To check $(H2)$ we observe that the column that we remove, the first one, is associated to the first cell and, up to some rescaling but bounded factor $|A_\varepsilon(y_1 - y_j)|$, its entries are the lengths of its sides. Yet, this cell conserves the area bounds we had on the sequence and its entry cannot be all zero. □

From the previous results on the form and the regularity of the derivatives of $\nabla_p \Phi$, we deduce from the Cauchy-Lipschitz Theorem that the ODE (3.3) governing the evolution of the dual variables is well posed and actually characterizes the optimal prices:

Theorem 3.5. *Let $p(\varepsilon)$ be the solution of the dual problem (3.2) (recall the normalization $p_1(\varepsilon) = 0$), then it is the only solution of the ODE:*

$$\frac{dp}{d\varepsilon}(\varepsilon) = -D^2_{p,p}\Phi(p(\varepsilon), \varepsilon)^{-1} \left(\frac{\partial}{\partial \varepsilon} \nabla_p \Phi(p(\varepsilon), \varepsilon) \right) \tag{3.12}$$

with initial condition $p(0)$ such that all the horizontal strips $C(p(0), 0)_i$ have area $1/N$.

4 Numerical results

4.1 Algorithm

The algorithm we propose consists simply in discretizing (3.12) (together with the initial condition $p(0)$ determined as in Theorem 3.5) by an explicit Euler scheme. Let n be some positive integer and $h := n^{-1}$ be the step size. Let us set $p_0 = p(0)$ and define prices inductively as follows.

- While (p_k, kh) belongs to the open set $\mathcal{O}$ defined by (3.4), compute:

$$A_k := -D^2_{p,p}\Phi(p_k; kh),\ \delta_k := \frac{\partial}{\partial\varepsilon}\nabla_p\Phi(p_k, kh).$$

  Note that computing A_k and δ_k by formulas (3.9)-(3.10) and (3.11), requires to construct the cells $C(kh, p_k)_i$.

- Solve the linear system $A_k z = \delta_k$; by taking advantage of A_k being positive definite, we use the conjugate gradient algorithm for minimizing $J_k(z) = A_k(z)(z) - 2\delta_k \cdot z$ to solve this system exactly in $N-1$ steps. We denote by z_k the solution.

- Update the prices by setting

$$p_{k+1} = p_k + hz_k.$$

Thanks to the Lipschitz properties established in section 3, it is easy to check that for h small enough, (p_k, kh) always remain in $\mathcal{O}$ and then p_k is well-defined for every k up to n. For such an h and since (3.12) is generated by a Lipschitz function on $\mathcal{O}$, it is well-known that the convergence of the Euler scheme is linear (see for instance [3]). Denoting by p^h the piecewise constant function having values p_k on intervals $[kh, (k+1)h)$, we thus get the following convergence:

Theorem 4.1. *For h small enough, the algorithm above is well-defined and the uniform error between p^h and the optimal price p is $O(h)$.*

4.2 Numerical experiments

The construction of the cells at each step is achieved efficiently by an implementation in Matlab. In the setting described above where $d = 2$, $\Omega = [0,1]^2$, μ is the uniform distribution on Ω, $\nu = \frac{1}{N}\sum_{k=1}^N \delta_{y_k}$, our algorithm computes the cells $\Omega_k^\varepsilon = \{x \in \Omega : T_\varepsilon(x) = y_k\}$ as well as the prices p_k^ε of the cell y_k for $\varepsilon = 0$ to $\varepsilon = 1$. For $\varepsilon = 0$ which is the case where the transportation plan is the Knothe one, the computational task amounts to sorting the second component of the y_k's. Appropriate discretization steps are then chosen for the transition $\varepsilon = 0$ to $\varepsilon = 1$. At each step, the tesselation of Ω into the polyhedral cells Ω_k^ε is computed based on the prices p_k^ε. Adjacency information on these cells is computed, as well as the length of the facet between two cells and the coordinates of its extreme points. This information allows one to formulate a discretized version of ODE (3.12) using an Euler discretization scheme.

For geometric computations we use the Multi-Parametric Toolbox library, available online at http://control.ee.ethz.ch/ mpt. In particular, `polytope` computes the tesselation of Ω into the polyhedral cells Ω_k^ε. (A slighlty modified version of) `mpt_buildAdjacency` extracts adjacency information on this tesselation, from which the vertices and the lengths of the sides between two cells can be deduced. The library also incorporates convenient graphical routines.

Error analysis Some numerical examples are presented below, for which relative errors in cell areas (i.e. deviation from the optimality conditions) are given as well as a comparison between the tesselations obtained with our method and the true solution. Another way to test our method is as follows. Let us consider the set

$$C := \left\{ z(\gamma) := \left(\int_{\mathbb{R}^2\times\mathbb{R}^2} x_1 y_1 d\gamma, \int_{\mathbb{R}^2\times\mathbb{R}^2} x_2 y_2 d\gamma \right),\ \gamma \in \Gamma(\mu,\nu) \right\}.$$

C is a closed convex subset of $\mathbb{R}^2$ and it is strictly convex in the sense that its boundary contains no line segment. Denoting by γ_ε the solution of (3.1), it is easy to check that $z(\varepsilon) := z(\gamma_\varepsilon)$ is an extreme point of C and that $\varepsilon z_1 + z_2 = \varepsilon z_1(\varepsilon) + z_2(\varepsilon)$ is the equation of a supporting line of C at $z(\varepsilon)$ and this supporting line intersects C only at $z(\varepsilon)$. If we consider the *correlation* curve $\varepsilon \in (0,1) \mapsto z(\varepsilon)$ it can be represented as the graph of a concave decreasing function whose slope (when it exists) at point $z(\varepsilon)$ is $-\varepsilon$. In our numerical test, we will also present graphs comparing the true concave correlation curve to the one computed by our method.

We give three instances of executions of our algorithm, with samples of respectively 5, 10, and 15 points. Taking weights $(\varepsilon, \varepsilon^{-1})$ with $\varepsilon \in (0, +\infty)$ (rather than $(\varepsilon, 1)$ with $\varepsilon \in (0,1)$) we get the full evolution of the optimal transports from one Knothe's transport (horizontal strips) to the other (vertical strips). Further examples as well as videos can be found at `http://alfred.galichon.googlepages.com/anisotropic`. It should also be pointed that our method does not approximate the solution of a single optimal transportation problem but a whole family of such problems (which actually explains relatively high running times).

Five sample points. We take as our sample set a sample of five points. We get the following errors:

# steps	**Relative errors in cell areas**					**Time**
100	-4.41%	2.66%	3.41%	-2.46%	0.80%	66 s
500	-0.88%	0.54%	0.68%	-0.49%	0.16%	349 s

for which we draw in Figure 1 the partition obtained for $\varepsilon = 1$ using an exact method, as well as the true evolutions of the componentwise correlations of the x's and the y's from $\varepsilon = 0$ until $\varepsilon = 1$.

Running the continuation algorithm with 100 and 500 discretization steps we obtained partition sketched below, where the relative error on the cells area is inferior to 5%. The evolution of the componentwise correlations from $\varepsilon = 0$ until $\varepsilon = 1$ are also sketched and compared with their exact conterparts (the above curve). Running the algorithm with 500 discretization steps on a standard laptop took 349 seconds, and yields to the results below, where the maximal relative error on the cells area is 1%. The evolution of the componentwise correlations from $\varepsilon = 0$ until $\varepsilon = 1$ are also sketched and compared with their exact counterparts.

Ten sample points. Taking a sample of ten points, we obtain the following errors:

# steps	**Relative errors in cell areas**					**Time**
2,000	-1.37%	-12%	-2.44%	-0.09%	2.27%	$\leq$ 1 h
3,000	-0.92%	-8.23%	1.63%	0.06%	1.52%	$\leq$ 1,5 h
	(continued)					
2,000	0.27%	5.71%	2.12%	7.36%	-1.58%	$\leq$ 1 h
3,000	0.18%	3.84%	1.42%	4.94%	-1.05%	$\leq$ 1,5 h

We draw in Figure 3 the partition obtained for $\varepsilon = 1$ using an exact method, as well as the true evolutions of the componentwise correlations of the x's and the y's from $\varepsilon = 0$ until $\varepsilon = 1$.

Fifteen sample points. With a sample of 15 points, we get the following errors:

# steps	**Relative errors in cell areas**					**Time**
3,000	-2.24%	-0.12%	-0.58%	0.20%	-3.20%	$\simeq$ 2 h
10,000	-0.67%	-0.04%	-0.17%	0.06%	-0.97%	$\leq$ 4 h
	(continued)					
3,000	0.17%	-2.66%	4.25%	-2.43%	28.85%	$\simeq$ 2 h
10,000	0.05%	-0.82%	1.28%	-0.74%	9.42%	$\leq$ 4 h
	(continued)					
3,000	-30.11%	2.51%	1.82%	2.92%	0.63%	$\simeq$ 2 h
10,000	-9.79%	0.76%	0.55%	0.89%	0.19%	$\leq$ 4 h

Acknowledgements G.C. and F.S. gratefully acknowledge the support of the Agence Nationale de la Recherche via the research project OTARIE. A.G. gratefully acknowledge the support of Chaire EDF-Calyon "Finance et

Développement Durable," of Chaire AXA "Assurance et risques majeurs," and of Chaire Société Générale "Risques Financiers". The authors wish to warmly thank Yann Brenier and Alessio Figalli for stimulating discussions.

References

[1] A. Braides, Γ-*convergence for beginners*, Oxford University Press, Oxford, 2002.

[2] Brenier, Y., Polar factorization and monotone rearrangement of vector-valued functions, *Communications on Pure and Applied Mathematics* 44, 375-417, 1991.

[3] Crouzeix M., Mignot A.L, *Analyse numérique des équations différentielles*, Masson, Paris, 1984.

[4] Knothe, H., Contributions to the theory of convex bodies, *Michigan Mathematical Journal* 4, pp. 39–52, 1957.

[5] Rachev S.T., Rüschendorf L., *Mass Transportation Problems. Vol. I: Theory; Vol. II : Applications*, Springer-Verlag 1998.

[6] Rosenblatt, M., Remarks on a multivariate transformation, *Annals of Mathematical Statistics* 23, pp. 470–472, 1952.

[7] Villani, C., *Topics in Optimal Transportation*, Providence: American Mathematical Society, 2003.

[8] Villani, C., *Optimal transport: Old and New*, lecture notes, Ecole d'été de probabilités de Saint-Flour, to appear, New York: Springer, 2008.

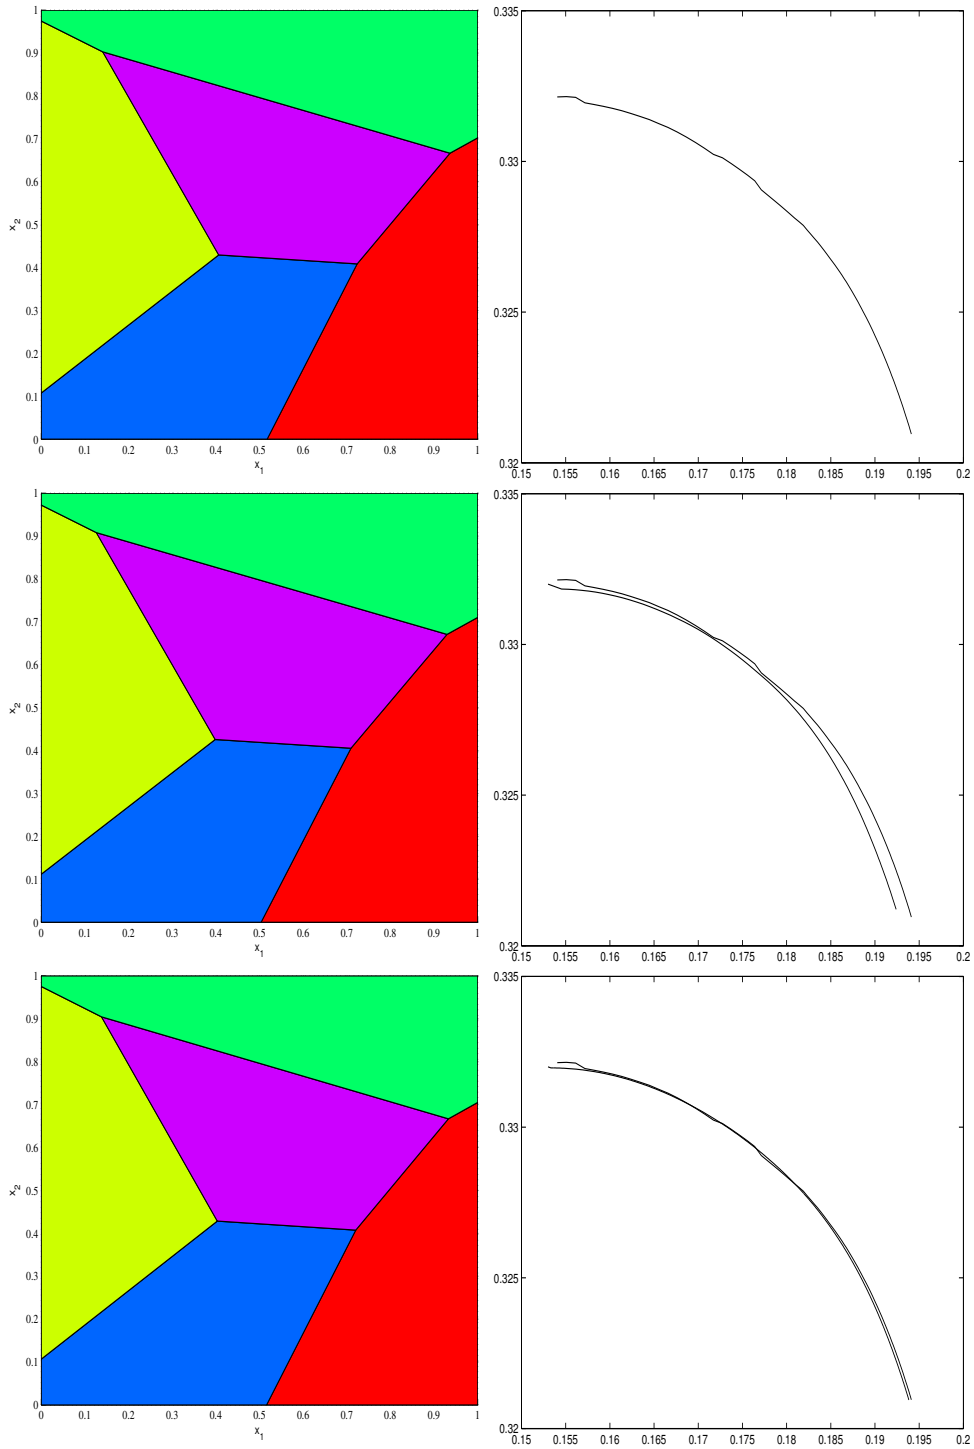


Figure 1: Five sample points. Top row: exact algorithm (gradient method). Middle row: continuation algorithm, 100 steps. Bottom row: continuation algorithm, 500 steps.

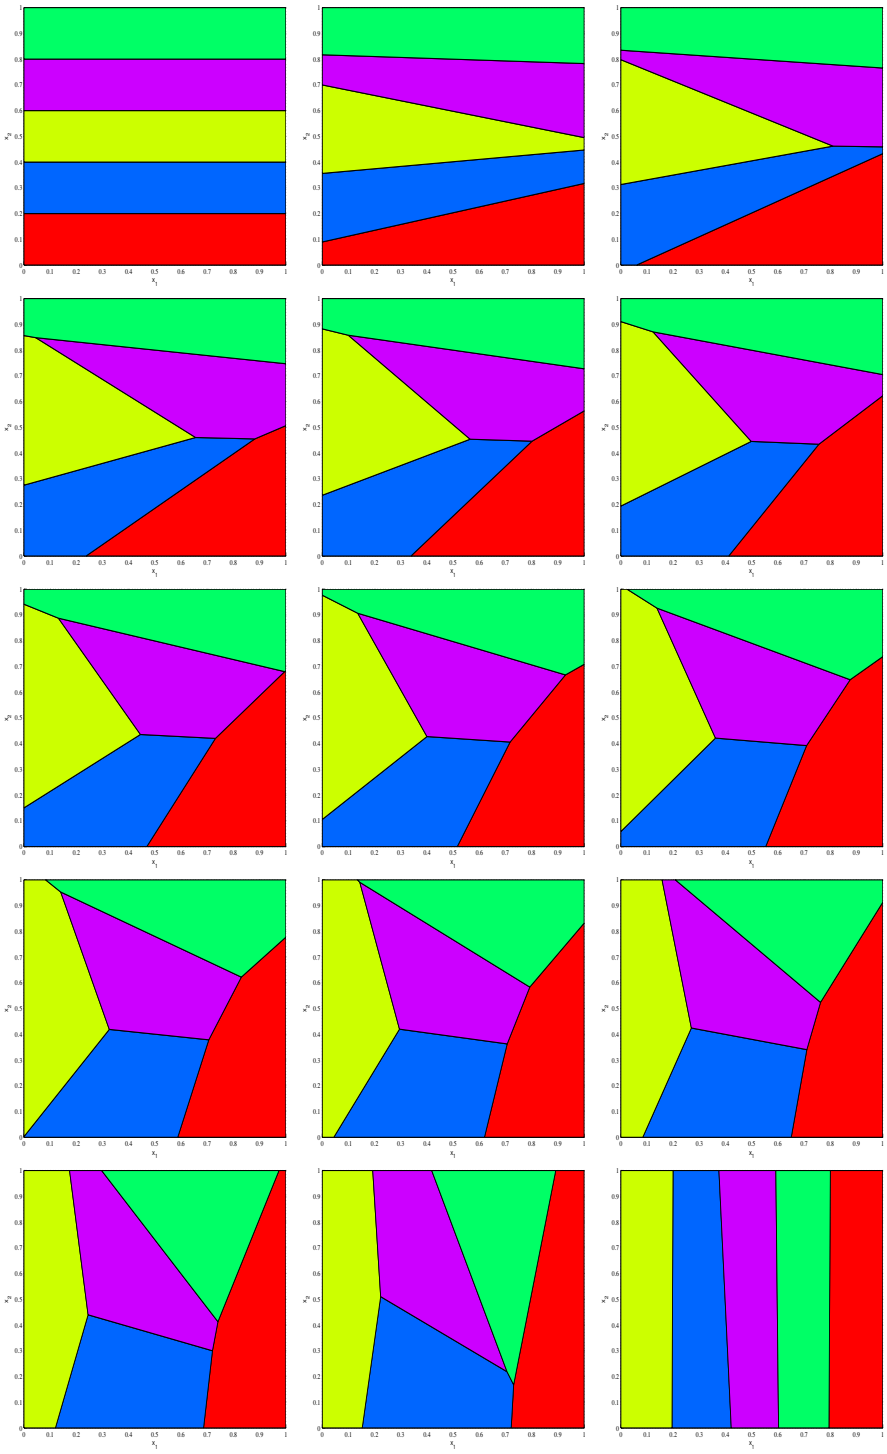


Figure 2: Five sample points: evolution of the tesselation for $\varepsilon = 0$ to $\varepsilon = +\infty$ (from top left to bottom right).

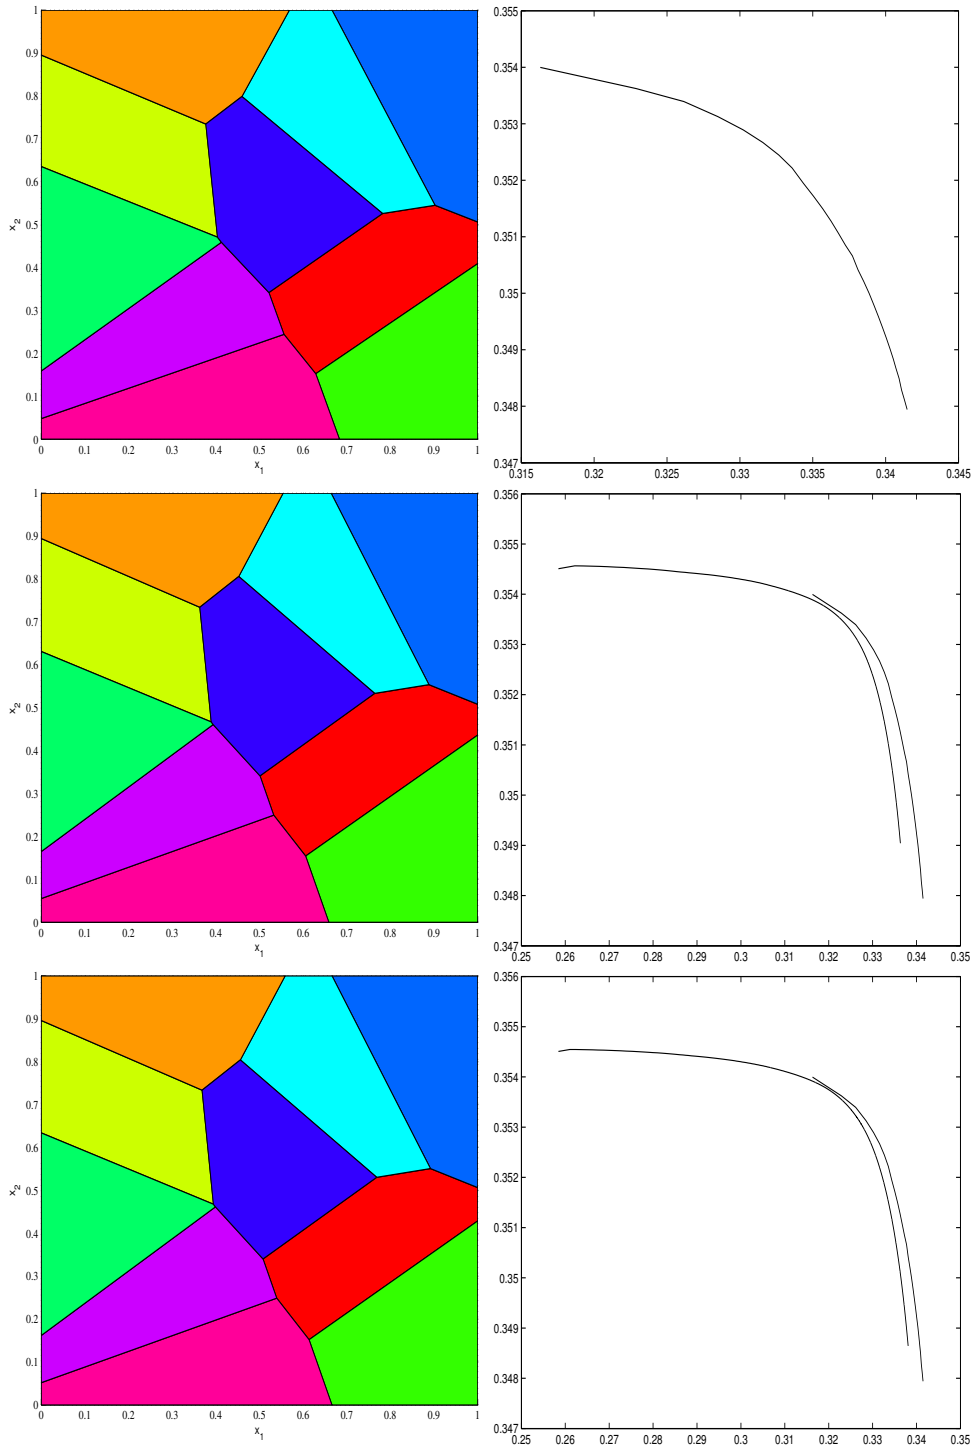


Figure 3: Ten sample points. Top row: exact algorithm (gradient method). Middle row: continuation algorithm, 2000 steps. Bottom row: continuation algorithm, 3000 steps.

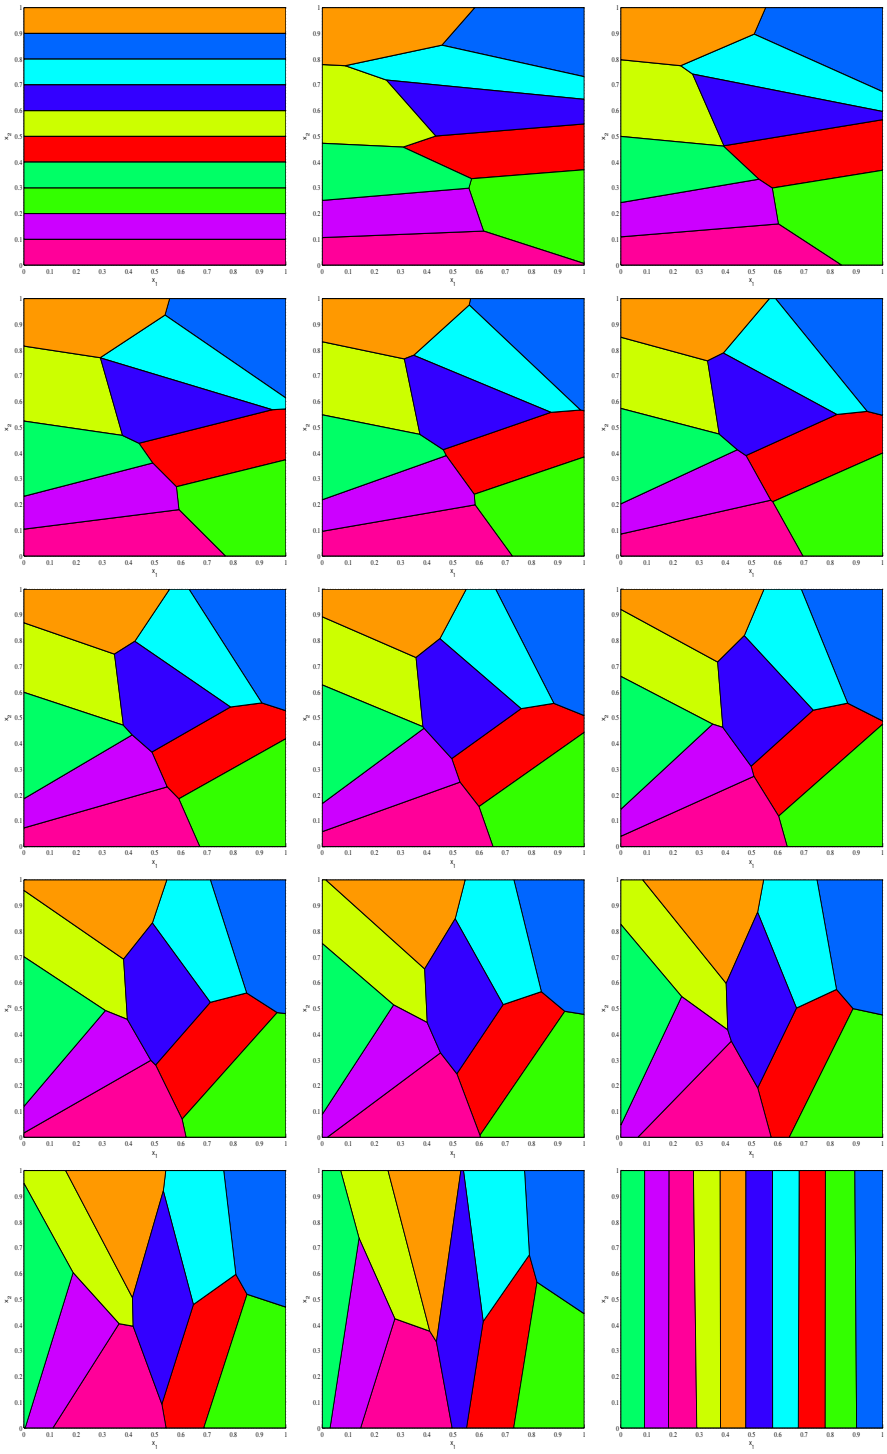


Figure 4: Ten sample points: evolution of the tesselation for $\varepsilon = 0$ to $\varepsilon = +\infty$ (from top left to bottom right).

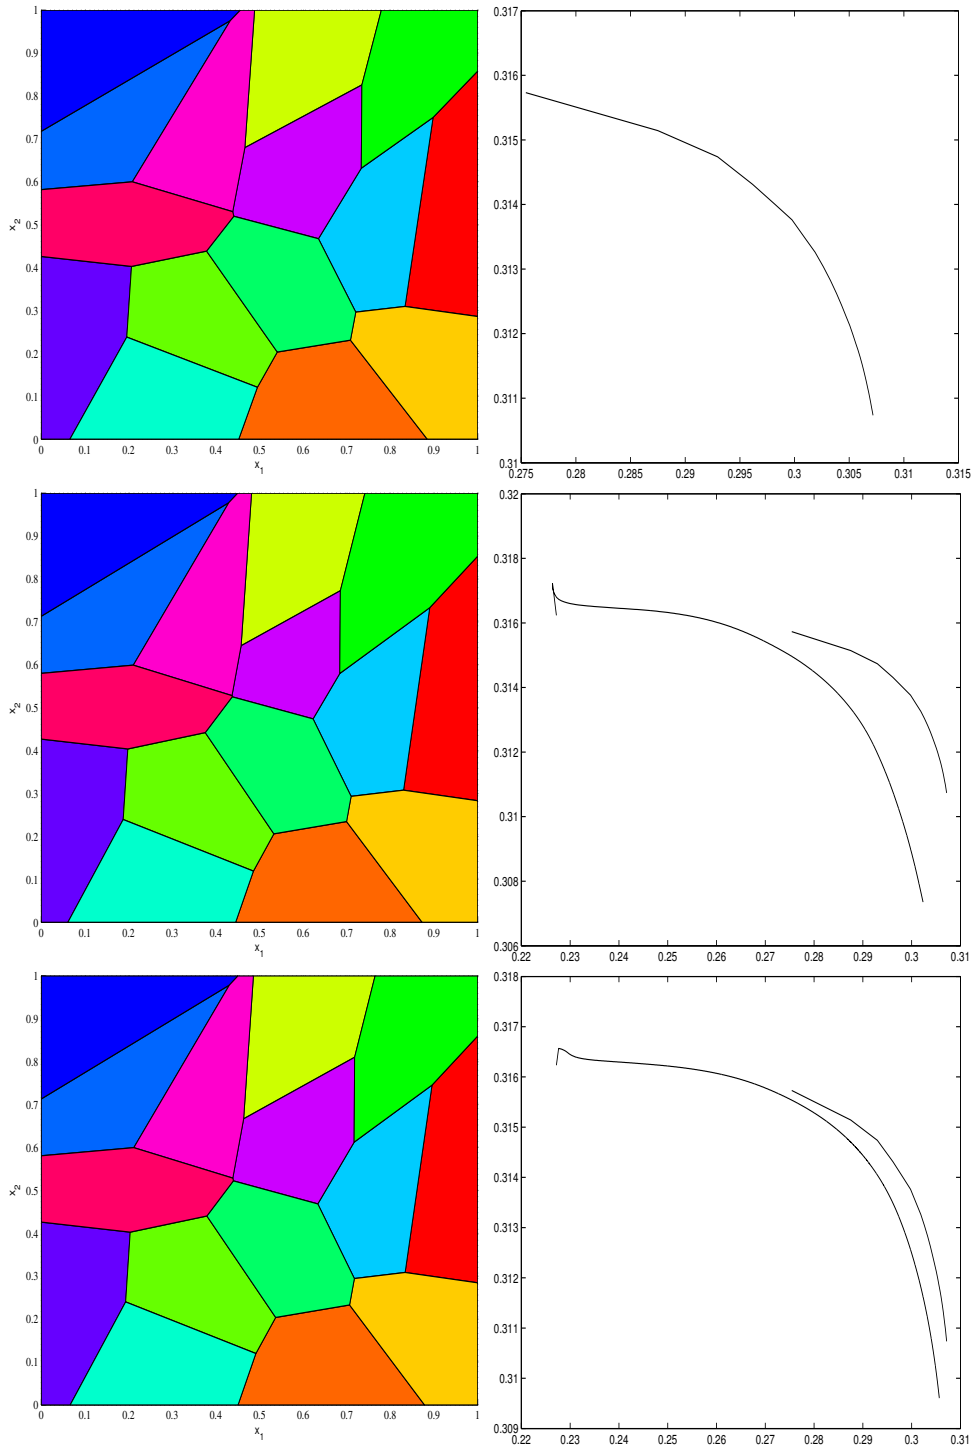


Figure 5: Fifteen sample points. Top row: exact algorithm (gradient method). Middle row: continuation algorithm, 3000 steps. Bottom row: continuation algorithm, 10,000 steps.

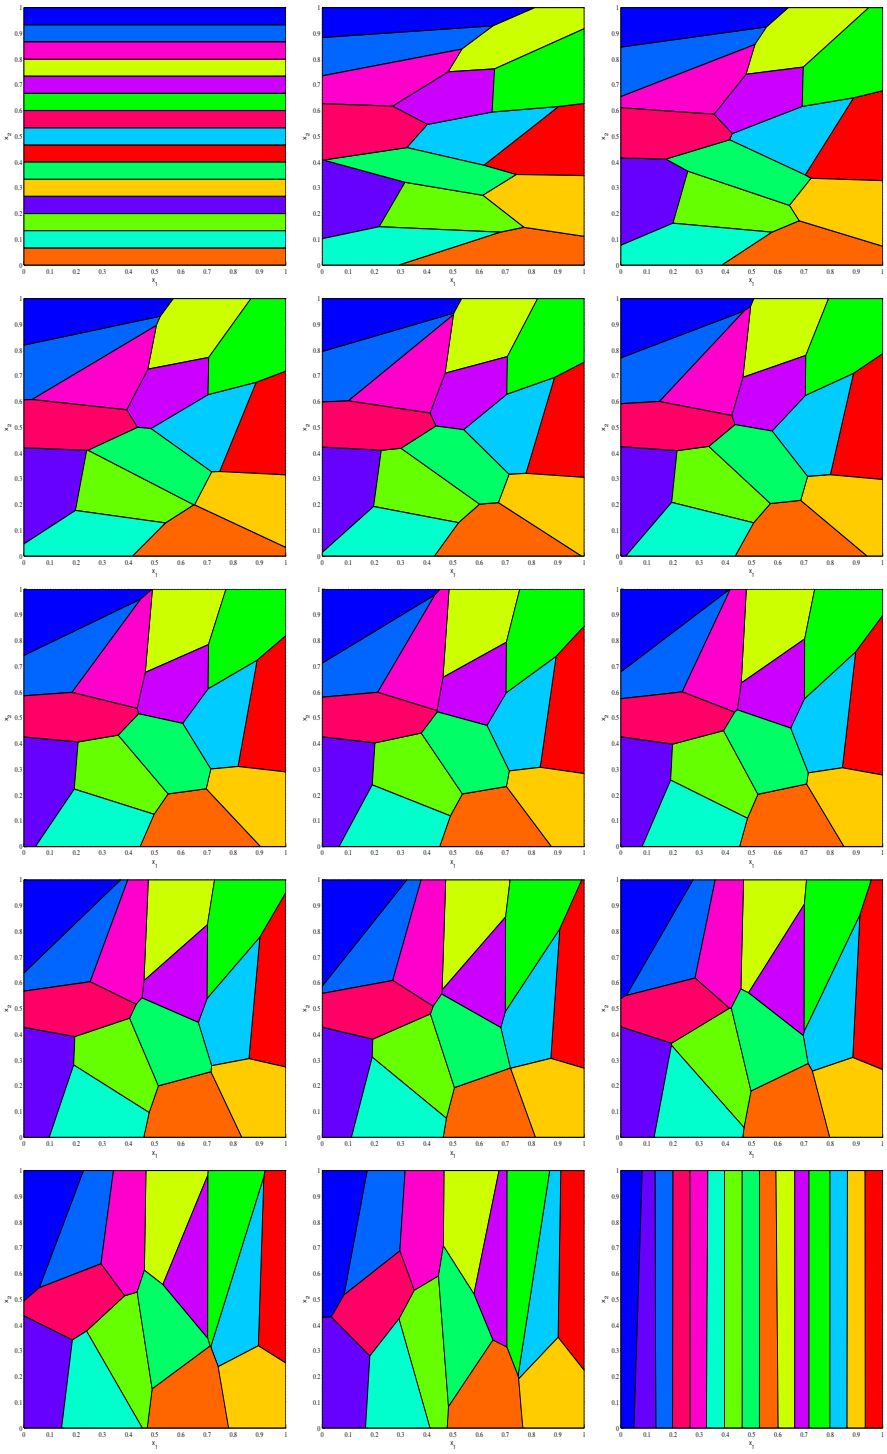


Figure 6: Fifteen sample points: evolution of the tesselation for $\varepsilon = 0$ to $\varepsilon = +\infty$ (from top left to bottom right).