DOUGLAS–RACHFORD SPLITTING FOR HIGH-ACCURACY UNREGULARISED MULTI-MARGINAL OPTIMAL TRANSPORT

Anonymous authors
Paper under double-blind review

ABSTRACT

Multi-Marginal Optimal Transport (MMOT) is commonly solved with entropic regularization because iterative scaling is fast and stable at moderate accuracy. The same regularization, however, introduces a non-vanishing bias with respect to the original linear program. We study **DRMOT**, a Douglas–Rachford splitting method for unregularised MMOT. Its key technical ingredient is a closed-form orthogonal projection onto the affine set of prescribed tensor marginals, which extends the known matrix projection to arbitrary-order tensors. We further introduce a scale-aware default step size and a cost-aware feasible recovery tailored to the sparse iterates produced by DR; on identical DRMOT iterates, this recovery reduces the final reported error by roughly 5–$10\times$ relative to the standard rank-one correction. We re-evaluate DRMOT under equal wall-clock budgets as a complete finite-time solver: the entropic baselines retain their published recovery rules, while DRMOT uses the proposed cost-aware recovery. On Wasserstein-barycenter grids, the best DRMOT variant is $2.1\times$–$15.7\times$ more accurate than the best tuned entropic baseline within the same 20-second budget. In within-family cost diagnostics over the two evaluated instances per family, the best observed DRMOT-s runs are the only tested ones to reach an absolute objective gap of 10^{-4} within 20 seconds on both Coulomb and Brenier–Euler costs, while the dimension-only DRMOT variant gives the best observed 10^{-3}-level time on risk pooling. On model-independent option pricing, DRMOT agrees with a high-accuracy HiGHS reference down to float64 numerical precision whereas the entropic methods plateau near 2×10^{-4}. These results identify the regime in which solving MMOT without entropic bias is practically beneficial.

1 INTRODUCTION

Multi-Marginal Optimal Transport (MMOT) generalises classical optimal transport from a coupling of two distributions to a joint coupling of $m \geq 2$ marginals. This higher-order view appears in Wasserstein barycenters (Benamou et al., 2015), generative modelling (Albergo et al., 2024), robust risk management (Ennaji et al., 2024), matching for teams, and strictly correlated-electron models in density functional theory (Buttazzo et al., 2012). Its flexibility comes at a cost: a dense transport plan contains $\prod_k n_k$ entries, and even enforcing all marginal constraints becomes expensive as either support size or the number of marginals grows. Figure 1 makes this transition explicit: the transport plan changes from a matrix to a higher-order tensor, while every new mode introduces another family of marginal constraints that must be satisfied simultaneously.

Regularised versus unregularised MMOT. Most practical MMOT solvers add entropy and exploit iterative scaling. Multi-Sinkhorn (MSK) and its accelerated variant (Lin et al., 2022), as well as PD-AAM (Tupitsa et al., 2020), are effective when moderate accuracy is sufficient. The regularization parameter, however, couples two roles: it controls both numerical behaviour and the bias relative to the original unregularised linear program. For any fixed positive regularization, the solver targets a regularised problem whose optimum generally differs from that of the original MMOT linear program. This is benign in many applications, but it becomes visible when the desired error is well below 10^{-3}–10^{-4} or when the transport plan itself must satisfy marginal constraints to high precision.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

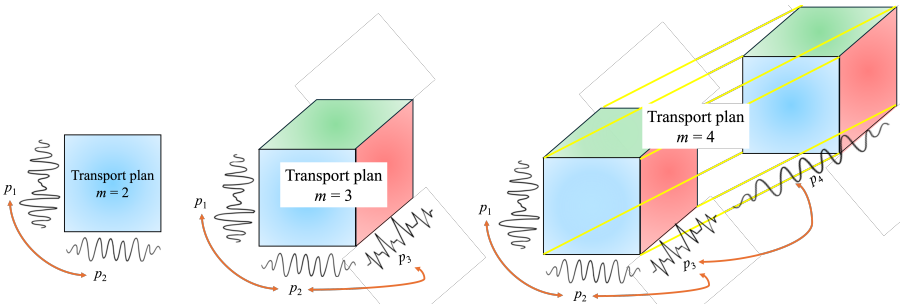


Figure 1: Progressive illustration of MMOT for $m = 2, 3, 4$. The transport plan becomes a higher-order tensor while retaining prescribed marginals along every mode. DRMOT is enabled by projecting onto these simultaneous affine marginal constraints in closed form.

Douglas–Rachford (DR) splitting provides a complementary route. In the bi-marginal setting, Mai et al. (2022) showed that DR can solve the unregularised problem without an entropic approximation. Extending the same idea to MMOT requires repeated projection onto the affine set defined by all marginal equalities. For matrices, a closed form is known (Bauschke et al., 2021); for higher-order tensors, the corresponding formula is the key missing ingredient addressed here. We target the general dense-tensor setting rather than a problem-specific structured MMOT solver; accordingly, DRMOT retains the intrinsic n^m dependence, but its projection and update can be evaluated with simple tensor operations once the closed form is available.

From an optimization iterate to a feasible reported solution. A second issue is empirical rather than purely algorithmic. A finite DR shadow iterate need not yet satisfy every marginal exactly. Reporting its raw objective can therefore be misleading: an infeasible point can have an objective below the true feasible optimum. We consequently treat feasible recovery as part of the finite-time solver. The standard rank-one correction used in the MMOT literature is cost agnostic and works naturally with dense entropic iterates, but it can distort the sparse iterates produced by DRMOT. We therefore introduce a cost-aware correction as a component of DRMOT. In the empirical comparison, each method is evaluated as a complete finite-time pipeline under the same wall-clock budget: competing solvers retain the rounding/recovery procedures specified in their original algorithms, while DRMOT uses the proposed recovery.

Contributions.

1. We derive a **closed-form orthogonal projection onto the affine marginal-consistency set** for tensors of arbitrary order, generalising the bi-marginal matrix result of Bauschke et al. (2021).
2. Substituting this projection into DR yields **DRMOT**, whose dense per-iteration arithmetic is $\mathcal{O}(mn^m)$ for equal support size n. Together with the standard ergodic DR objective bound, this gives the $\mathcal{O}(mn^m/\varepsilon)$ fixed-instance arithmetic scaling analysed in Section 4 for an ergodic shadow iterate (with the instance-dependent DR constant suppressed), while the polyhedral structure yields a problem-dependent linear convergence rate in distance to the solution set.
3. We propose two practical components required for high-accuracy use: a **scale-aware default step size** that removes per-instance step-size tuning, and a **cost-aware feasible recovery** that reduces the error of reported DRMOT plans by roughly 5–10$\times$ relative to the standard rank-one correction on the same iterates.
4. We **evaluate complete finite-time solver pipelines under equal wall-clock budgets**. DRMOT gives the lowest error on all barycenter grid sizes. In the within-family diagnostics over the two evaluated instances per cost family, the best observed DRMOT-s runs reach an absolute objective gap of 10^{-4} within 20 seconds on Coulomb and Brenier–Euler costs where none of the tested entropic configurations does; DRMOT also agrees with a high-accuracy LP reference down to float64 numerical precision on a model-independent option-pricing problem.

Concurrent theoretical work. Augustino et al. (2026) independently analyse the complexity of MMOT. Their classical algorithm reduces (1) to a positive packing LP and completes it with a rank-one correction, reaching an ε-additive coupling in $\mathcal{O}(m^2 n^m \|C\|_\infty/\varepsilon)$ arithmetic operations (Table 1). Its explicit dependence on tensor size and ε can be compared with our fixed-instance

Table 1: Representative arithmetic complexity bounds for dense MMOT methods. The formulations and accuracy notions follow the corresponding analyses in the cited works and are not directly interchangeable. The DRMOT row suppresses an instance-dependent constant from the ergodic DR bound. We use the table as asymptotic context rather than as a substitute for the equal-time empirical comparison in Section 5.

Algorithm	Complexity
Multi-Sinkhorn (Lin et al., 2022)	$\mathcal{O}(m^3 n^m/\varepsilon^2)$
Accelerated multi-Sinkhorn (Lin et al., 2022)	$\mathcal{O}(m^3 n^{m+1/3}/\varepsilon^{4/3})$
PD-AAM (Tupitsa et al., 2020)	$\mathcal{O}(m^3 n^{m+1/2}/\varepsilon)$
Packing-LP MOT (Augustino et al., 2026)	$\mathcal{O}(m^2 n^m \|C\|_\infty/\varepsilon)$
DRMOT (ergodic shadow; fixed instance)	$\mathcal{O}(mn^m/\varepsilon)$

ergodic scaling, but the accuracy notions and problem-dependent constants in the two guarantees are not directly comparable. They further develop quantum algorithms and establish corresponding classical and quantum cost-query lower bounds. Their contribution is primarily asymptotic and does not provide the equal-wall-clock numerical evaluation considered here. DRMOT is complementary: it provides a concrete GPU-friendly solver with a scale-aware default step size and feasible finite-time outputs, together with an empirical evaluation of the high-accuracy unregularised regime.

2 Preliminaries

Notation. $[N] = \{1, \dots, N\}$, $\mathbf{1}_N$ is the all-ones vector, and $\Delta^N = \{u \in \mathbb{R}^N_+ : \mathbf{1}_N^\top u = 1\}$. For $X \in \mathbb{R}^{n_1 \times \cdots \times n_m}$, the k-th marginal is $[r_k(X)]_j = \sum_{i_l \in [n_l],\, l \neq k} X(i_1, \dots, i_{k-1}, j, i_{k+1}, \dots, i_m)$.

Given marginals $\{p_k \in \Delta^{n_k}\}_{k \in [m]}$, define the affine marginal-consistency set $\mathcal{A} = \{X : r_k(X) = p_k, \forall k\}$ and the feasible transportation polytope $\Pi = \mathcal{A} \cap \mathbb{R}^{n_1 \times \cdots \times n_m}_+$. This distinction is useful below: the DR splitting projects onto $\mathcal{A}$, while other proximal operators handle the non-negativity.

Definition 2.1 (MMOT)**.** For $C \in \mathbb{R}^{n_1 \times \cdots \times n_m}$, the discrete MMOT problem is

$$\text{minimize}_{X \in \mathbb{R}^{n_1 \times \cdots \times n_m}_+} \ \langle C, X \rangle \quad \text{subject to} \quad r_k(X) = p_k, \ \forall k \in [m]. \tag{1}$$

Douglas–Rachford splitting. For $\min_x f(x) + g(x)$ with proper closed convex f, g, DR alternates $x^{(t+1)} = \text{prox}_{\lambda f}(y^{(t)})$, $z^{(t+1)} = \text{prox}_{\lambda g}(2x^{(t+1)} - y^{(t)})$, and $y^{(t+1)} = y^{(t)} + z^{(t+1)} - x^{(t+1)}$ (Douglas and Rachford, 1956; Bauschke and Combettes, 2017).

3 The DRMOT Algorithm

3.1 Reformulation and framework

Let $\mathcal{A} = \{X : r_k(X) = p_k, \ \forall k\}$ denote the affine set satisfying the marginal equalities and let $\Pi = \mathcal{A} \cap \mathbb{R}^{n_1 \times \cdots \times n_m}_+$ be the MMOT transportation polytope. We split the objective as

$$f(X) = \langle C, X \rangle + \mathbb{I}_{\mathbb{R}_+}(X), \qquad g(X) = \mathbb{I}_{\mathcal{A}}(X). \tag{2}$$

The proximal maps are $\text{prox}_{\lambda f}(Y) = [Y - \lambda C]_+$ and $\text{prox}_{\lambda g}(W) = \mathsf{Proj}_{\mathcal{A}}(W)$, so a DR step is The initialization makes the first thresholded iterate equal to the product coupling $\bigotimes_k p_k$. The only non-elementwise proximal operation is therefore the orthogonal projection onto $\mathcal{A}$. The next section derives it in closed form.

3.2 Closed-form projection onto the affine marginal set

The only non-elementwise proximal operation in the DR update is the simultaneous projection onto all marginal equalities. This projection is therefore the key structural obstacle in extending the bi-marginal DR formulation to general MMOT. We now show that it admits a closed form for tensors of arbitrary order.

Algorithm 1 Douglas–Rachford splitting for MMOT

Require: Cost tensor C, marginals $\{p_k\}_{k\in[m]}$, step λ, iterations T
1: $Y^{(0)} \leftarrow \bigotimes_{k=1}^{m} p_k + \lambda C$
2: **for** $t = 1, \dots, T$ **do**
3: $\quad X^{(t)} \leftarrow [Y^{(t-1)} - \lambda C]_+$
4: $\quad Y^{(t)} \leftarrow \mathsf{Proj}_{\mathcal{A}}(2X^{(t)} - Y^{(t-1)}) + Y^{(t-1)} - X^{(t)}$
5: **end for**
6: **return** $X^{(T)}$

Let $\mathcal{M}(W) = \big(r_1(W), \dots, r_m(W)\big), \qquad p = (p_1, \dots, p_m)$, so that the affine marginal-consistency set is $\mathcal{A} = \{W : \mathcal{M}(W) = p\}$. We obtain the following result.

Theorem 3.1 (Closed-form affine marginal projection)**.** *For any $W \in \mathbb{R}^{n_1 \times \cdots \times n_m}$, the orthogonal projection onto $\mathcal{A}$ is*

$$\mathsf{Proj}_{\mathcal{A}}(W) = W - \mathcal{M}^{\dagger}\big(\mathcal{M}(W) - p\big), \tag{3}$$

where $\mathcal{M}^{\dagger}$ is the Moore–Penrose pseudoinverse of the marginal operator. Equivalently, it admits the explicit tensor formula

$$\mathsf{Proj}_{\mathcal{A}}(W) = W - \sum_{k=1}^{m} \frac{1}{\prod_{i\neq k} n_i} \left(\bigotimes_{i<k} \mathbf{1}_{n_i} \right) \otimes \left(r_k(W) - p_k - \frac{\left(\sum_{i\neq k} \prod_{j\neq i,k} n_j\right) \langle \mathbf{1}_{n_k}, r_k(W) - p_k \rangle}{\sum_i \prod_{j\neq i} n_j} \mathbf{1}_{n_k} \right) \otimes \left(\bigotimes_{i>k} \mathbf{1}_{n_i} \right). \tag{4}$$

The proof follows by expressing the marginal constraints as an affine linear system and computing the minimum-norm correction induced by its Moore–Penrose pseudoinverse; the complete derivation is given in Appendix 7. For $m = 2$, Theorem 3.1 reduces to the matrix projection derived by Bauschke et al. (2021), whereas (4) applies directly to arbitrary tensor order. Importantly, (3) is a characterization rather than a prescription to form a dense pseudoinverse numerically. The explicit expression in (4) requires only marginal reductions, vector operations, and broadcast tensor corrections. Hence the projection can be evaluated directly in $\mathcal{O}(mn^m)$ arithmetic for equal support size n, without solving an auxiliary projection problem at every DR iteration. Substituting (4) into the DR recurrence yields the specialized DRMOT update used in our implementation. The marginal residuals $\mu_k^{(t)} = r_k(Y^{(t)}) - p_k$ and $\nu_k^{(t)} = r_k(X^{(t)}) - p_k$ obey a direct recursion and can be cached for arbitrary support sizes. Thus the implementation never forms a pseudoinverse or solves an auxiliary constrained projection, and each iteration needs only one full-tensor marginalization. The full closed-form recurrence and residual recursion are given in Appendix 8.

3.3 Scale-aware default step size

The DR thresholding step depends on the product λC. The original DRMOT rule uses only the marginal dimensions. With $\beta = \sum_{k=1}^{m} \prod_{i\neq k} n_i$, we denote this dimension-only default by $\lambda_{\mathrm{dim}} = \frac{1}{\beta}$. This rule is simple, but it is not equivariant to a positive change of cost units: for $s > 0$, if $C \mapsto sC$, preserving the same DR threshold requires $\lambda \mapsto \lambda/s$.

Motivated by this observation, for nonzero cost tensors ($\|C\|_\infty > 0$) we use the scale-aware rule

$$\lambda_{\mathrm{scaled}} = \frac{c}{\|C\|_\infty \sqrt{\beta}}, \qquad c = 0.7, \tag{5}$$

which has the required transformation by construction. The normalized quantity $\lambda^\star \sqrt{\beta} \|C\|_\infty$ is also substantially more stable than $\lambda^\star \beta$ in the representative step-size sweeps reported in Appendix 13; these sweeps provide supporting evidence for the default rather than a scale-wide optimality claim. We use $c = 0.7$ as one global default throughout the reported experiments. For clarity, throughout the remainder of the paper, **DRMOT** denotes the solver with the dimension-only step λ_{dim}, whereas **DRMOT-s** denotes the scaled-step variant using $\lambda_{\mathrm{scaled}}$ in (5). Both variants use the cost-aware feasible recovery of Section 3.4 unless stated otherwise. The scale-aware rule removes per-instance step-size tuning; it does not imply that DRMOT has no algorithmic hyperparameters.

3.4 COST-AWARE RECOVERY OF A FEASIBLE PLAN

DR shadow iterates are only asymptotically feasible, so objective values should not be reported directly from $X^{(t)}$. A finite-time solver must instead map the current iterate to the MMOT feasible polytope Π. The standard recovery of Lin et al. (2022) first multiplicatively rescales the tensor so that each marginal is dominated by its target and then adds a rank-one correction $\bigotimes_k \mathrm{err}_k / \|\mathrm{err}_1\|^{m-1}$. This construction is cost agnostic.

For DRMOT, the distinction matters because the positive-part proximal step $[Y - \lambda C]_+$ often produces sparse iterates. A dense rank-one correction can therefore place a noticeable amount of residual mass on high-cost cells. We retain the multiplicative down-scaling step and define the cost-aware residual transport problem

$$K^\eta \in \arg\min_{K \geq 0} \langle C, K \rangle - \eta \|C\|_\infty H(K), \qquad r_k(K) = \mathrm{err}_k \ \forall k, \tag{6}$$

where $H(K) = -\sum_{\mathbf{i}} K_{\mathbf{i}}(\log K_{\mathbf{i}} - 1)$. Scaling the entropy weight by $\|C\|_\infty$ makes this auxiliary recovery invariant to positive changes of cost units. In the finite-time implementation we approximate K^η from the initialization $\exp(-C/(\eta\|C\|_\infty))$ using 40 multi-marginal Sinkhorn scaling sweeps, followed by a final Lin-style polish that enforces the marginals to numerical precision. We use the same $\eta = 5 \times 10^{-3}$ in all experiments. Importantly, entropy is used only to distribute the finite residual mass during feasible recovery; it does not regularise the MMOT objective optimized by DRMOT. As the DR iterate approaches feasibility, the residual mass tends to zero and this auxiliary correction vanishes.

This recovery is part of DRMOT rather than a proposed modification of the entropic baselines. Consequently, all baseline results in Section 5 use their published recovery rules. To isolate the effect of the new component without changing any competing method, Table 2 compares the two recovery operators on *the same DRMOT iterates*. The negative raw errors are particularly informative: they do not indicate a "better-than-optimal" solution, but expose that the shadow iterate is not yet feasible. Recovery is therefore part of the finite-time solver output rather than a cosmetic post-processing step.

Table 2: Same-iterate ablation of feasible recovery for DRMOT. "raw" is the objective before recovery and can lie below the feasible optimum because the iterate is not yet feasible. Cost-aware recovery reduces the final reported error by roughly 5–10× on these representative cases.

Benchmark	Solver	raw	Lin ROUND	Cost-aware	improvement
ellipse 20×20	DRMOT-s	-0.049%	$+0.69\%$	$+0.13\%$	$5.2\times$
	DRMOT	$+0.042\%$	$+0.36\%$	$+0.06\%$	$6.4\times$
Random m=3, n=15	DRMOT-s	-0.033%	$+0.57\%$	$+0.09\%$	$6.4\times$
SCE (m=3)	DRMOT	-0.068%	$+2.59\%$	$+0.27\%$	$9.5\times$

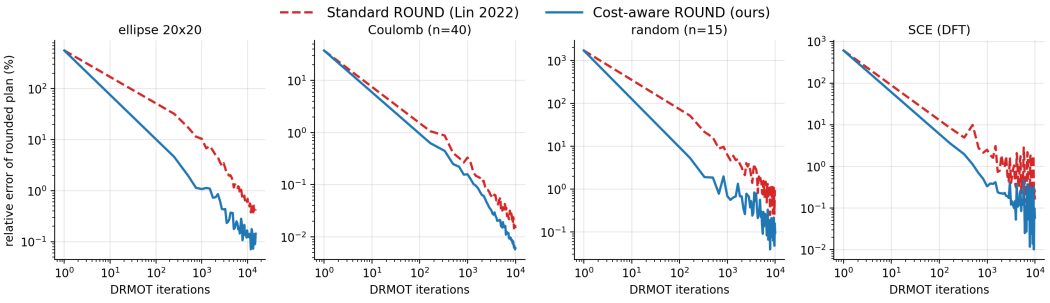


Figure 2: Relative error of the recovered DRMOT plan versus DRMOT iteration on four benchmarks. Cost-aware recovery preserves more of the progress made by the sparse DR iterate and reduces the reported error by up to roughly one order of magnitude across the shown horizons.

Figure 2 shows that Table 2 is not an artifact of one stopping time: the separation between the two recovery rules persists over the optimization trajectory and across several cost structures.

4 CONVERGENCE RATES

Theorem 4.1 (Ergodic objective complexity). *Let $\bar{X}^{(T)} := T^{-1}\sum_{t=1}^{T} X^{(t)}$ be the ergodic average of the DR shadow iterates generated by Algorithm 1. For a fixed MMOT instance, the standard ergodic Douglas–Rachford bound (Mai et al., 2022) implies that there is an instance-dependent constant $\Gamma_{\mathrm{DR}} > 0$ such that*

$$|\langle C, \bar{X}^{(T)} - X^{\star}\rangle| \le \frac{\Gamma_{\mathrm{DR}}}{T}. \tag{7}$$

Consequently, the arithmetic required to obtain an ergodic shadow iterate with objective error below ε is $\mathcal{O}(m\prod_k n_k\,\Gamma_{\mathrm{DR}}/\varepsilon)$. Suppressing this fixed-instance constant gives the familiar $\mathcal{O}(m\prod_k n_k/\varepsilon)$ scaling, or $\mathcal{O}(mn^m/\varepsilon)$ for equal support size n.

Proof. Each iteration applies pointwise operations to the full tensor and computes all m marginals, for a total arithmetic cost $\mathcal{O}(m\prod_k n_k)$. Combining this cost with (7) gives the result. The constant Γ_{DR} absorbs fixed-instance quantities from the DR bound, including the initialization and step size. Appendix 9 gives the arithmetic count and states the distinction between this ergodic shadow guarantee and the feasible finite-time plans used in the experiments. □

Theorem 4.1 is an ergodic optimization-iterate bound rather than a complexity guarantee for the recovered finite-time plan. Because finite DR shadow iterates need not satisfy all marginal constraints exactly, Section 3.4 separately specifies the feasible plan reported in experiments; its measured wall-clock cost is included in every empirical comparison.

Theorem 4.2 (Linear convergence). *Write Algorithm 1 as $Y^{(t+1)} = F(Y^{(t)})$ with $F(Y) = \mathrm{prox}_{\lambda g}(2\mathrm{prox}_{\lambda f}(Y) - Y) + Y - \mathrm{prox}_{\lambda f}(Y)$. Let $\mathcal{Y}^*$ be the fixed-point set and $\mathcal{X}^*$ the MMOT optimal set. For any fixed problem instance and initialization, there exists a problem-dependent $r \in (0,1)$ such that $\mathrm{dist}(Y^{(t)}, \mathcal{Y}^*) \le r^t \mathrm{dist}(Y^{(0)}, \mathcal{Y}^*)$ and $\mathrm{dist}(X^{(t)}, \mathcal{X}^*) \le r^{t-1}\mathrm{dist}(Y^{(0)}, \mathcal{Y}^*)$.*

Appendix 10 provides the detailed proof of Theorem 4.2. The contraction factor is problem dependent and can be close to one; therefore, the finite-budget curves may exhibit a long pre-asymptotic regime even though the polyhedral DR iteration admits a geometric distance bound.

5 EXPERIMENTAL EVALUATIONS

5.1 PROTOCOL

Every method is evaluated through a numerically feasible reported transport plan, with marginal residuals reduced to float64 numerical precision. We compare complete finite-time solver pipelines rather than isolated update rules: DRMOT uses the cost-aware recovery of Section 3.4, while the entropic baselines retain the recovery prescribed in their published algorithms. Thus the main comparison measures end-to-end solver output rather than attempting to isolate the optimization trajectory from solver-specific feasible recovery. The incremental effect of the proposed recovery is instead isolated on identical DRMOT iterates in Table 2. Each (solver, ε) configuration receives the same end-to-end wall-clock budget. The timed pipeline includes the feasible recovery used to produce the reported plan; monitoring and logging are excluded. Unless stated otherwise, relative error is measured against a high-accuracy LP reference from HiGHS, $|\langle C, \hat{X}\rangle - \langle C, X^{\star}\rangle|/|\langle C, X^{\star}\rangle|$. Section 5.4 retains the absolute objective-gap diagnostic used by the original cost-family protocol and labels it explicitly as such. All solvers are implemented in PyTorch float64 on an NVIDIA RTX 5070 Ti. We also verify the baseline update rules against direct transcriptions of their published pseudocode and use wall-clock rather than iteration count because the work per update differs substantially across methods.

5.2 WASSERSTEIN BARYCENTER BENCHMARK UNDER EQUAL TIME

We first revisit the grid-size Wasserstein-barycenter benchmark used in the original DRMOT study. The task uses three nested elliptical marginals ($m = 3$) on image grids from 20×20 to 50×50. The grid width is not the support size of each marginal: each marginal is restricted to the non-zero pixels of its image. Across these instances, n_k ranges from 44 to 186, giving dense transport tensors from

0.17M to 5.1M entries (at most about 41 MB in float64 for the transport tensor itself). Each (solver, ε) configuration receives the same 20-second end-to-end timed budget, including the feasible recovery used to produce the reported plan, and the entropic methods are evaluated over $\varepsilon \in \{0.01, 0.1, 1, 10\}$. DRMOT and DRMOT-s use the cost-aware recovery from Section 3.4; MSK, AccMSK, and PD-AAM use the recovery specified in their published implementations. The important pattern is not uniform domination at every tolerance. At 20×20, MSK is still competitive (0.04% versus 0.019%), whereas at 30×30 the best entropic result is 0.47% and DRMOT-s reaches 0.030%. Thus the gap widens precisely as the requested solution enters the high-accuracy regime. The best entropic ε also changes with grid size, illustrating the finite-budget trade-off between regularization bias and conditioning.

Table 3: Grid-size benchmark under an equal 20-second wall-clock budget per (solver, ε). Every entry is the relative objective error of a numerically feasible reported plan. Best results are shown in bold and second-best results are underlined.

Algorithm	ε, λ	20×20	30×30	40×40	50×50
DRMOT	λ_{dim}	**0.019%**	<u>0.031%</u>	**0.041%**	<u>0.042%</u>
DRMOT-s	λ_{scaled}	<u>0.034%</u>	**0.030%**	<u>0.043%</u>	**0.033%**
MSK	0.01	14.62%	44.76%	73.96%	120.87%
	0.1	0.04%	0.59%	4.33%	16.97%
	1	1.51%	0.47%	0.22%	0.16%
	10	36.14%	12.88%	6.79%	4.03%
AccMSK	0.01	13.25%	79.19%	149.26%	174.80%
	0.1	0.70%	3.48%	14.32%	41.30%
	1	1.51%	0.48%	0.30%	1.17%
	10	36.14%	12.88%	6.79%	4.03%
PD-AAM	0.01	2.13%	5.06%	12.91%	58.33%
	0.1	0.20%	0.47%	1.01%	4.22%
	1	1.54%	0.52%	0.32%	0.53%
	10	36.22%	12.94%	6.83%	4.07%

5.3 High-accuracy model-independent option pricing

Model-independent option pricing (Ennaji et al., 2024) formulates no-arbitrage bounds on the price of a basket option as $\sup_{X \in \Pi} \mathbb{E}_X[\text{payoff}] = -\min_{X \in \Pi} \langle -\text{payoff}, X \rangle$. We use three real daily series from FRED (S&P 500, Dow Jones, and VIX), discretise their 5-day log-returns into equiprobable quantiles ($n = 30$), and set $C = -\text{payoff}$ for a basket call. A relative error of 10^{-4} corresponds to one basis point on the reported option-price bound. Figure 3 shows a regime in which the distinction between regularised and unregularised MMOT becomes visible. DRMOT agrees with the high-accuracy HiGHS reference down to float64 numerical precision in about 0.1 seconds, whereas all tested entropic configurations plateau near 2×10^{-4}. We interpret values below the numerical accuracy of the LP reference only as agreement between two float64 computations, not as a certificate of the exact mathematical optimum. The experiment nevertheless shows a clear practical separation: reducing the entropic parameter improves bias but also changes conditioning, while DRMOT continues to optimize the unregularised objective directly. The point of this example is not that every pricing task requires machine-level agreement, but that here the regularization floor has a directly interpretable scale: 10^{-4} corresponds to one basis point on the reported bound.

5.4 Three cost families beyond the barycenter

We next test whether the high-accuracy behaviour is specific to barycenters. We consider three structurally different MMOT costs with $m = 3$ and $n = 40$: Coulomb cost from strictly correlated-electron models, Brenier–Euler flow (Brenier, 1989), and risk pooling with exponential utility (Carlier and Ekeland, 2010). Each (solver, ε) configuration receives the same 20-second wall-clock budget, with $\varepsilon \in \{1, 0.5, 0.3, 0.2, 0.1, 0.05, 0.02\}$ for the entropic methods. Figure 4 plots, for each method, the best observed absolute objective gap $|\langle C, \hat{X} \rangle - \langle C, X^\star \rangle|$ over its hyperparameter grid and the two evaluated instances. Across the two evaluated instances per family, the best-observed envelopes show

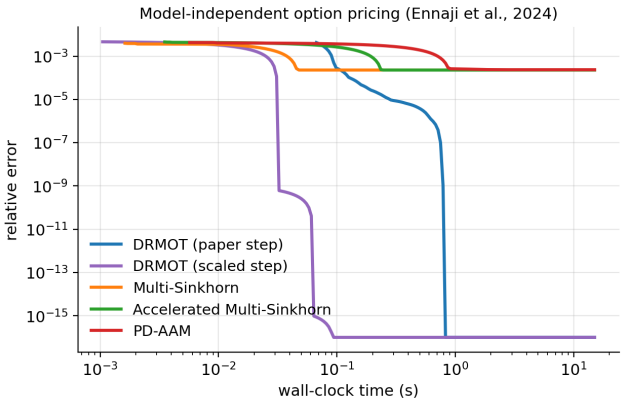


Figure 3: Model-independent option pricing (Ennaji et al., 2024) with real FRED risk factors (S&P 500 / Dow Jones / VIX), $m = 3$, $n = 30$, and 15 seconds per (solver, ε).

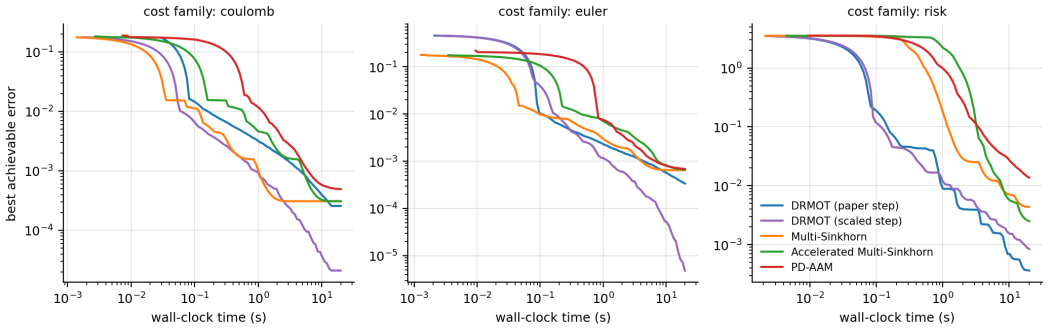


Figure 4: Best observed absolute objective gap versus wall-clock time on three cost families. Each curve is the lower envelope over the corresponding method's hyperparameter grid and the two evaluated instances.

a separation in the high-accuracy regime. At moderate accuracy, entropic scaling remains competitive and can match or outperform DRMOT early in the run. As the objective gap decreases, however, DRMOT continues to improve while the tested entropic configurations exhibit parameter-dependent plateaus. In particular, the best observed DRMOT-s runs reach an absolute objective gap of 10^{-4} on both Coulomb and Brenier–Euler within the 20-second budget, whereas none of the tested entropic configurations does. On risk pooling, no method reaches 10^{-4} within the budget, although DRMOT achieves the smallest observed gap among the tested methods at the 10^{-3} level. These diagnostics are consistent with the regime-dependent picture in the barycenter experiment without being intended as a statistical characterization of each cost family.

5.5 Scaling with support size and number of marginals

The preceding accuracy experiments use $m = 3$, so we separately examine the per-iteration computational scaling for $m \in \{3, 4, 5\}$. Figure 5 reports runtime as the support size grows, with no instrumentation inside the iteration loop. DRMOT tracks Multi-Sinkhorn closely, whereas AccMSK and PD-AAM incur noticeably larger per-iteration overhead. To isolate the effect of the number of marginals, Figure 6 (a) fixes $n = 40$ for all $m \in \{3, 4, 5\}$ and reports the median over five measured repeats after GPU warm-up, with min–max error bars. At $m = 3, 4$, fixed CUDA kernel-launch overhead masks much of the asymptotic tensor growth; once the dense tensor reaches $m = 5, n = 40$, memory traffic dominates and the rapid increase in per-iteration cost becomes visible. DRMOT does not remove the intrinsic n^m dependence of dense MMOT, but its closed-form projection keeps the per-iteration cost close to MSK and well below the accelerated variants. The $m = 5$ regime is also memory intensive: on a single 16GB GPU, PD-AAM's workspace is already about 8.6GB at $n = 40$

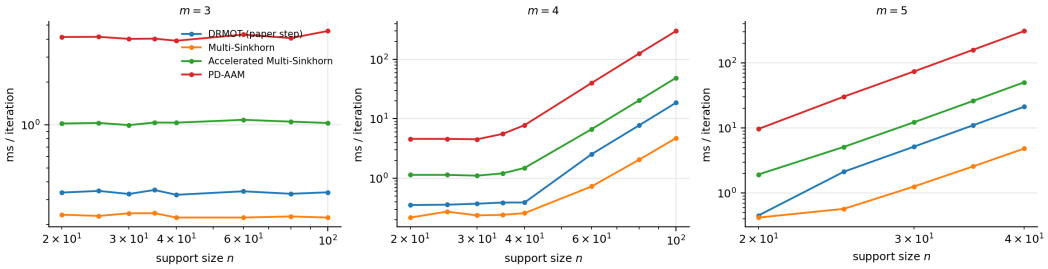
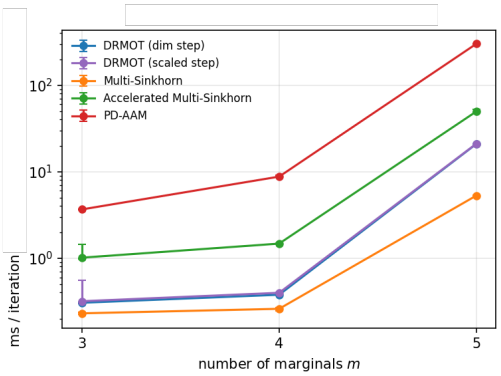
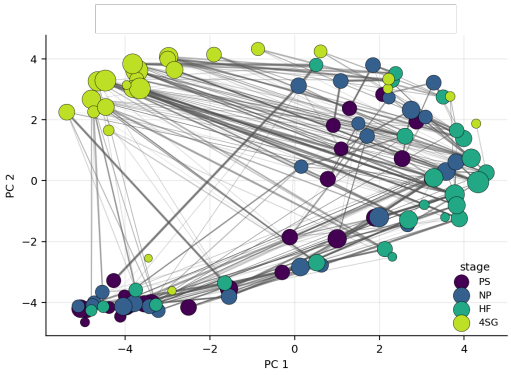


Figure 5: Per-iteration computational scaling with support size n for $m = 3, 4, 5$. DRMOT remains close to the inexpensive Multi-Sinkhorn baseline, while AccMSK and PD-AAM incur substantially larger update overhead.



(a) Per-iteration time versus the number of marginals m at fixed $n = 40$.



(b) Joint 4-marginal DRMOT coupling on the single-cell blood-development time course.

Figure 6: Higher-order behaviour of DRMOT.

and about 15GB at $n = 45$. Most quantitative accuracy experiments above use $m = 3$; Figure 6(b) therefore complements the scaling result with a qualitative four-marginal coupling on the single-cell blood-development time course of Moignard et al. (2015), showing the higher-order object returned by DRMOT without making a biological benchmarking claim.

6 Conclusion

We revisited Douglas–Rachford splitting for unregularised MMOT with a focus on high-accuracy *feasible* solutions under fixed computational budgets. DRMOT is enabled by a closed-form projection onto the affine marginal set and is complemented by a scale-aware step size and a cost-aware recovery tailored to sparse DR iterates. While the method retains the intrinsic $\mathcal{O}(n^m)$ scaling of general dense MMOT, it avoids the accuracy floor caused by fixed entropic regularization. Hence, DRMOT is most relevant in the high-accuracy regime; for moderate accuracy or strongly structured large-scale problems, entropic or specialized methods may remain preferable. Under equal wall-clock budgets, DRMOT achieves the lowest error on the Wasserstein-barycenter grids considered here. On the additional cost-family diagnostics, the best observed DRMOT variants reach 10^{-4} on Coulomb and Brenier–Euler and the best 10^{-3}-level time on risk pooling. On option pricing, DRMOT agrees with the high-accuracy LP reference down to float64 precision, while the entropic methods plateau near 2×10^{-4}. Overall, DRMOT provides a practical high-accuracy alternative when entropic bias becomes the limiting factor, while still inheriting the memory scaling of dense MMOT. A systematic runtime crossover study against general-purpose LP solvers is left for future work.

Reproducibility statement

The implementation is available at the anonymous repository `https://anonymous.4open.science/r/DRMOT-0036`. Implementation details, experimental settings, and evaluation protocols are described in the main text and appendices.

AI use statement

Generative AI tools were used in preparing this work, and we disclose that use here. The authors defined the research questions, problem formulation, mathematical results, and proofs; no dataset in this study was generated by AI. AI tools assisted with experiment implementation and debugging, evaluation and figure-generation scripts, manuscript drafting and editing, literature search, and critical review of intermediate drafts. All citations were checked against their original sources, and the authors decided independently which AI-generated suggestions to adopt. The authors remain fully responsible for the experimental methodology, interpretation of results, final claims, text, and released artifacts. Reported numbers were checked against saved experimental outputs and the corresponding verification scripts.

References

Michael S. Albergo, Nicholas M. Boffi, Michael Lindsey, and Eric Vanden-Eijnden. Multimarginal generative modeling with stochastic interpolants. *International Conference on Learning Representations (ICLR)*, 2024.

Brandon Augustino, Yue Sun, Atithi Acharya, Shouvanik Chakrabarti, Junhyung Lyle Kim, Shree Hari Sureshbabu, and Charlie Che. Faster algorithms for multimarginal optimal transport. *arXiv preprint arXiv:2608.09513*, 2026.

Heinz H. Bauschke and Patrick L. Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2017.

Heinz H. Bauschke, Shambhavi Singh, and Xianfu Wang. Projecting onto rectangular matrices with prescribed row and column sums. *Fixed Point Theory and Algorithms for Sciences and Engineering*, pages 1–20, 2021.

Jean-David Benamou, Guillaume Carlier, Marco Cuturi, Luca Nenna, and Gabriel Peyré. Iterative bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.

Yann Brenier. The least action principle and the related concept of generalized flows for incompressible perfect fluids. *Journal of the American Mathematical Society*, 2(2):225–255, 1989.

Giuseppe Buttazzo, Luigi De Pasquale, and Paola Gori-Giorgi. Optimal-transport formulation of electronic density-functional theory. *Physical Review A*, 85(6):062502, 2012.

Guillaume Carlier and Ivar Ekeland. Matching for teams. *Economic Theory*, 42(2):397–418, 2010.

Jim Douglas and H. H. Rachford. On the numerical solution of heat conduction problems in two and three space variables. *Transactions of the American Mathematical Society*, 82(2):421–439, 1956.

H. Ennaji, Q. Mérigot, L. Nenna, and B. Pass. Robust risk management via multi-marginal optimal transport. *Journal of Optimization Theory and Applications*, 202(2):554–581, 2024.

Wu Li. The sharp Lipschitz constants for feasible and optimal solutions of a perturbed linear program. *Linear Algebra and its Applications*, 187:15–40, 1993. ISSN 0024-3795. doi: https://doi.org/10.1016/0024-3795(93)90125-8. URL `https://www.sciencedirect.com/science/article/pii/0024379593901258`.

Tianyi Lin, Nhat Ho, Marco Cuturi, and Michael I. Jordan. On the complexity of approximating multimarginal optimal transport. *Journal of Machine Learning Research*, 23(1):2835–2877, 2022.

Pierre-Louis Lions and Bertrand Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.

Vien V. Mai, Jacob Lindbäck, and Mikael Johansson. A fast and accurate splitting method for optimal transport: analysis and implementation. *International Conference on Learning Representations (ICLR)*, 2022.

Victoria Moignard, Steven Woodhouse, Laleh Haghverdi, Andrew J. Lilly, Yosuke Tanaka, Adam C. Wilkinson, Florian Buettner, Iain C. Macaulay, Wajid Jawaid, Evangelia Diamanti, et al. Decoding the regulatory network of early blood development from single-cell gene expression measurements. *Nature Biotechnology*, 33(3):269–276, 2015.

Nazarii Tupitsa, Pavel Dvurechensky, Alexander Gasnikov, and Sergey Guminov. Multimarginal optimal transport by accelerated alternating minimization. In *IEEE Conference on Decision and Control (CDC)*, pages 6132–6137, 2020.

7 Proof of Theorem 3.1: closed-form affine projection

This appendix derives the closed-form projection used by DRMOT. The proof is written directly in terms of marginal and broadcast operators; this avoids the auxiliary contraction notation used in the earlier version of the manuscript. Throughout, the tensor space is equipped with the Frobenius inner product and the product space of marginals with the corresponding block Euclidean inner product.

7.1 Marginal operator and its adjoint

Define the linear marginal operator

$$\mathcal{M} : \mathbb{R}^{n_1 \times \cdots \times n_m} \longrightarrow \mathbb{R}^{n_1} \times \cdots \times \mathbb{R}^{n_m}, \qquad \mathcal{M}(W) = \big(r_1(W), \ldots, r_m(W)\big). \tag{8}$$

For $v_k \in \mathbb{R}^{n_k}$, let

$$\mathcal{B}_k(v_k) := \left(\bigotimes_{i<k} \mathbf{1}_{n_i} \right) \otimes v_k \otimes \left(\bigotimes_{i>k} \mathbf{1}_{n_i} \right) \in \mathbb{R}^{n_1 \times \cdots \times n_m}. \tag{9}$$

The operator $\mathcal{B}_k$ broadcasts a vector along every tensor mode except mode k. For $v = (v_1, \ldots, v_m)$,

$$\langle \mathcal{M}(W), v \rangle = \sum_{k=1}^{m} \langle r_k(W), v_k \rangle = \left\langle W, \sum_{k=1}^{m} \mathcal{B}_k(v_k) \right\rangle. \tag{10}$$

Hence

$$\mathcal{M}^*(v_1, \ldots, v_m) = \sum_{k=1}^{m} \mathcal{B}_k(v_k). \tag{11}$$

For compactness, define

$$d_k := \prod_{i \neq k} n_i, \qquad d_{k\ell} := \prod_{i \neq k,\ell} n_i \quad (k \neq \ell), \qquad D := \sum_{i=1}^{m} d_i, \qquad s_k := \sum_{i \neq k} d_{ik}. \tag{12}$$

A direct marginalization of a broadcast tensor gives

$$r_\ell\big(\mathcal{B}_k(v_k)\big) = \begin{cases} d_k v_k, & k = \ell, \\ d_{k\ell}\, \langle \mathbf{1}_{n_k}, v_k \rangle \mathbf{1}_{n_\ell}, & k \neq \ell. \end{cases} \tag{13}$$

7.2 Moore–Penrose inverse of the marginal operator

Lemma 7.1 (Pseudoinverse of the marginal operator)**.** *For $v = (v_1, \ldots, v_m)$,*

$$\mathcal{M}^\dagger v = \sum_{k=1}^{m} \frac{1}{d_k} \mathcal{B}_k\left(v_k - \frac{s_k}{D} \langle \mathbf{1}_{n_k}, v_k \rangle \mathbf{1}_{n_k} \right). \tag{14}$$

Proof. Set

$$u_k := \frac{1}{d_k}\left(v_k - \frac{s_k}{D}\langle \mathbf{1}_{n_k}, v_k\rangle \mathbf{1}_{n_k}\right), \qquad U := \mathcal{M}^*(u_1, \ldots, u_m). \tag{15}$$

First observe that

$$\langle \mathbf{1}_{n_k}, u_k\rangle = \frac{\langle \mathbf{1}_{n_k}, v_k\rangle}{D}. \tag{16}$$

Indeed, $n_k s_k = D - d_k$, which follows immediately from the definitions in (12). Using (13), the ℓ-th block of $\mathcal{M}\mathcal{M}^* u$ is

$$[\mathcal{M}\mathcal{M}^* u]_\ell = d_\ell u_\ell + \sum_{k\neq \ell} d_{k\ell}\langle \mathbf{1}_{n_k}, u_k\rangle \mathbf{1}_{n_\ell}. \tag{17}$$

Substituting (15)–(16) into (17) gives

$$[\mathcal{M}U - v]_\ell = c_\ell \mathbf{1}_{n_\ell}, \qquad c_\ell = \frac{1}{D}\left(\sum_{k\neq\ell} d_{k\ell}\langle \mathbf{1}_{n_k}, v_k\rangle - s_\ell \langle \mathbf{1}_{n_\ell}, v_\ell\rangle\right). \tag{18}$$

The constant shifts cancel because $\sum_{\ell=1}^m c_\ell = 0$: after exchanging the two sums, the first term is $D^{-1}\sum_k s_k\langle \mathbf{1}_{n_k}, v_k\rangle$, exactly the second term. Therefore $\mathcal{M}^*(\mathcal{M}U - v) = 0$, and hence

$$\mathcal{M}^*\mathcal{M}U = \mathcal{M}^* v. \tag{19}$$

Consequently, $v - \mathcal{M}U \in \ker(\mathcal{M}^*) = \mathrm{ran}(\mathcal{M})^\perp$, so $\mathcal{M}U$ is the orthogonal projection of v onto $\mathrm{ran}(\mathcal{M})$. Moreover, $U \in \mathrm{ran}(\mathcal{M}^*) = \ker(\mathcal{M})^\perp$ by construction. Thus U is the unique minimum-norm preimage of $\mathsf{Proj}_{\mathrm{ran}(\mathcal{M})} v$, which is precisely $\mathcal{M}^\dagger v$. Substituting (15) and (11) gives (14). □

7.3 Projection onto the affine marginal-consistency set

The affine set used by DRMOT is

$$\mathcal{A} = \{W : \mathcal{M}(W) = p\}, \qquad p = (p_1, \ldots, p_m). \tag{20}$$

Because every p_k is a probability vector, the blocks of p have the same total mass and therefore $p \in \mathrm{ran}(\mathcal{M})$. The standard projection formula for an affine inverse image gives

$$\mathsf{Proj}_{\mathcal{A}}(W) = W - \mathcal{M}^\dagger\big(\mathcal{M}(W) - p\big). \tag{21}$$

Applying Lemma 7.1 to $v_k = r_k(W) - p_k$ yields exactly the explicit tensor expression in (4). This proves Theorem 3.1.

For $m = 2$, the same expression reduces to the prescribed-row/column-sum matrix projection of Bauschke et al. (2021). Importantly, the pseudoinverse in (21) is used only to characterise the projection: the implementation evaluates the explicit marginal-and-broadcast formula and never forms a dense pseudoinverse matrix.

8 Derivation and implementation of the DRMOT recurrence

Let

$$Z^{(t)} := 2X^{(t)} - Y^{(t-1)}, \qquad \mu_k^{(t)} := r_k(Y^{(t)}) - p_k, \qquad \nu_k^{(t)} := r_k(X^{(t)}) - p_k. \tag{22}$$

Define the scalar map

$$\tau_k(v) := \frac{s_k}{D}\langle \mathbf{1}_{n_k}, v\rangle, \qquad \phi_k^{(t)} := \tau_k(\mu_k^{(t)}), \qquad \psi_k^{(t)} := \tau_k(\nu_k^{(t)}), \tag{23}$$

where s_k and D are defined in (12). By linearity of each marginal,

$$r_k(Z^{(t)}) - p_k = 2\nu_k^{(t)} - \mu_k^{(t-1)}, \tag{24}$$

and, because τ_k is linear,

$$\tau_k\Big(r_k(Z^{(t)}) - p_k\Big) = 2\psi_k^{(t)} - \phi_k^{(t-1)}. \tag{25}$$

Substituting Theorem 3.1 into the DR update therefore gives

$$Y^{(t)} = X^{(t)} - \sum_{k=1}^m \frac{1}{d_k}\mathcal{B}_k\Big(2\nu_k^{(t)} - \mu_k^{(t-1)} - \big(2\psi_k^{(t)} - \phi_k^{(t-1)}\big)\mathbf{1}_{n_k}\Big). \tag{26}$$

Equation (26) is the closed-form recurrence used by DRMOT. It requires marginal reductions and broadcast corrections only; no constrained projection problem is solved inside the loop.

Residual recursion for arbitrary support sizes. Let $P^{(t)} := \mathsf{Proj}_{\mathcal{A}}(2X^{(t)} - Y^{(t-1)})$. By definition of the projection, $r_k(P^{(t)}) = p_k$ for every k, irrespective of the support sizes. Applying r_k directly to the DR update $Y^{(t)} = P^{(t)} + Y^{(t-1)} - X^{(t)}$ therefore gives

$$\mu_k^{(t)} = \mu_k^{(t-1)} - \nu_k^{(t)}, \qquad \phi_k^{(t)} = \phi_k^{(t-1)} - \psi_k^{(t)}. \tag{27}$$

Thus the residuals of $Y^{(t)}$ can be updated from already available vectors for arbitrary $(n_1, \dots, n_m)$, rather than recomputed by an additional full tensor reduction. In particular, one full-tensor marginalization of $X^{(t)}$ per iteration suffices.

9 Proof of Theorem 4.1: objective complexity

Let $N := \prod_{k=1}^m n_k$ denote the number of entries in the dense transport tensor. The non-negativity proximal step $X^{(t)} = [Y^{(t-1)} - \lambda C]_+$ is elementwise and costs $\mathcal{O}(N)$. Computing one marginal requires a reduction over the full tensor, hence $\mathcal{O}(N)$ arithmetic; computing all m marginals is $\mathcal{O}(mN)$. The vector corrections in (26) cost $\mathcal{O}(\sum_k n_k)$ and broadcasting the m corrected marginal vectors back to the dense tensor costs at most $\mathcal{O}(mN)$. Therefore one DRMOT iteration requires

$$\mathcal{O}(mN) = \mathcal{O}\left(m \prod_{k=1}^m n_k\right) \tag{28}$$

elementary operations. For equal support size n, this is $\mathcal{O}(mn^m)$.

For the ergodic shadow average $\bar{X}^{(T)} = T^{-1} \sum_{t=1}^T X^{(t)}$, the standard Douglas–Rachford ergodic objective bound (Mai et al., 2022) gives, for a fixed problem instance,

$$|\langle C, \bar{X}^{(T)} - X^\star \rangle| \le \frac{\Gamma_{\mathrm{DR}}}{T} \tag{29}$$

for an instance-dependent constant $\Gamma_{\mathrm{DR}} > 0$. Hence $T = \mathcal{O}(\Gamma_{\mathrm{DR}}/\varepsilon)$ iterations suffice for an ergodic shadow objective error below ε. Multiplying by the per-iteration arithmetic gives

$$\mathcal{O}\left(\frac{m \prod_k n_k \, \Gamma_{\mathrm{DR}}}{\varepsilon}\right). \tag{30}$$

Suppressing the fixed-instance constant recovers $\mathcal{O}(m \prod_k n_k/\varepsilon)$, or $\mathcal{O}(mn^m/\varepsilon)$ in the equal-support case. The hidden constant absorbs fixed-instance quantities from the DR bound, including the initialization and step size. This is an ergodic optimization-iterate bound; it is not a complexity guarantee for the recovered finite-time plan. Feasible recovery is evaluated separately, and its measured runtime is included in every wall-clock experiment.

10 Proof of Theorem 4.2: polyhedral error bound

We give the fixed-point/error-bound argument underlying the linear convergence statement. Let F denote the DR fixed-point operator from Theorem 4.2, let $\mathcal{Y}^* := \mathrm{Fix}(F)$, and define the residual $G := I - F$. The DR operator F is firmly non-expansive for proper closed convex f and g; consequently $G = I - F$ is also firmly non-expansive (Lions and Mercier, 1979; Bauschke and Combettes, 2017).

Lemma 10.1 (Fejér decrease)**.** *For the DRMOT fixed-point sequence $Y^{(t+1)} = F(Y^{(t)})$,*

$$\mathrm{dist}(Y^{(t+1)}, \mathcal{Y}^*)^2 \le \mathrm{dist}(Y^{(t)}, \mathcal{Y}^*)^2 - \|Y^{(t+1)} - Y^{(t)}\|^2. \tag{31}$$

In particular, $\{Y^{(t)}\}$ is bounded.

Proof. For any $Y^* \in \mathcal{Y}^*$, $G(Y^*) = 0$ and $Y^{(t+1)} - Y^{(t)} = -G(Y^{(t)})$. Firm non-expansiveness of G gives

$$\|Y^{(t+1)} - Y^*\|^2 = \|Y^{(t)} - Y^* - G(Y^{(t)})\|^2 \tag{32}$$

$$\le \|Y^{(t)} - Y^*\|^2 - \|G(Y^{(t)})\|^2. \tag{33}$$

Choosing a nearest fixed point proves (31). The same inequality shows that the distance to any fixed point is non-increasing, and hence the sequence is bounded. □

We next connect the fixed-point residual to the polyhedral primal–dual optimal set. The MMOT primal and dual linear programs are

$$\min_{X \geq 0} \langle C, X \rangle \qquad \text{s.t. } \mathcal{M}(X) = p, \tag{34}$$

$$\max_{\chi=(\chi_1,\dots,\chi_m)} \sum_{k=1}^{m} \langle \chi_k, p_k \rangle \qquad \text{s.t. } \mathcal{M}^*\chi \leq C. \tag{35}$$

Introduce a copy Z for the affine projection output and let $\mathcal{S}^*$ be the set of triples (X, Z, χ) satisfying

$$\begin{gathered} \mathcal{M}(Z) = p, \qquad Z = X, \qquad X \geq 0, \qquad \mathcal{M}^*\chi \leq C, \\ \langle C, X \rangle - \sum_{k=1}^{m} \langle \chi_k, p_k \rangle = 0. \end{gathered} \tag{36}$$

This is a nonempty polyhedron. Hence a Hoffman error bound (Li, 1993) gives a constant $\theta > 0$ for the distance to $\mathcal{S}^*$ in terms of the equality residuals and positive parts of the inequality residuals.

For the DRMOT iteration, define

$$Z^{(t)} := \mathsf{Proj}_{\mathcal{A}}\big(2X^{(t)} - Y^{(t-1)}\big), \qquad \Delta^{(t)} := Y^{(t)} - Y^{(t-1)} = Z^{(t)} - X^{(t)}. \tag{37}$$

Because the marginal constraints are redundant, projection multipliers need not be unique. We choose $q^{(t)}$ to be the minimum-norm projection multiplier, so that $Z^{(t)} = 2X^{(t)} - Y^{(t-1)} - \mathcal{M}^*q^{(t)}$, and set

$$\chi^{(t)} := -q^{(t)}/\lambda. \tag{38}$$

Using the DR identity in (37), this gives

$$Y^{(t)} = X^{(t)} + \lambda \mathcal{M}^*\chi^{(t)}. \tag{39}$$

The iterate already satisfies $X^{(t)} \geq 0$ and $\mathcal{M}(Z^{(t)}) = p$. The remaining KKT residuals can all be controlled by $\Delta^{(t)}$.

Lemma 10.2 (KKT residuals are controlled by the DR residual)**.** *There is a constant $c_0 > 0$, independent of t, such that*

$$\operatorname{dist}\big((X^{(t)}, Z^{(t)}, \chi^{(t)}), \mathcal{S}^*\big) \leq c_0 \|\Delta^{(t)}\|. \tag{40}$$

Proof. The equality $Z^{(t)} - X^{(t)} = \Delta^{(t)}$ gives the first residual. From $X^{(t)} = [Y^{(t-1)} - \lambda C]_+$, the positive-orthant projection condition implies

$$N^{(t)} := Y^{(t-1)} - \lambda C - X^{(t)} \in N_{\mathbb{R}_+}(X^{(t)}), \qquad \langle N^{(t)}, X^{(t)} \rangle = 0, \tag{41}$$

where $N_{\mathbb{R}_+}$ denotes the normal cone. Combining (39) with $Y^{(t)} = Y^{(t-1)} + \Delta^{(t)}$ yields

$$\mathcal{M}^*\chi^{(t)} - C = \lambda^{-1}\big(N^{(t)} + \Delta^{(t)}\big). \tag{42}$$

Since $N^{(t)}$ is componentwise non-positive on the only coordinates where a dual inequality can be violated,

$$\big\|[\mathcal{M}^*\chi^{(t)} - C]_+\big\| \leq \lambda^{-1}\|\Delta^{(t)}\|. \tag{43}$$

Moreover, because $\mathcal{M}(Z^{(t)}) = p$,

$$\sum_k \langle \chi_k^{(t)}, p_k \rangle = \langle \mathcal{M}^*\chi^{(t)}, Z^{(t)} \rangle. \tag{44}$$

Using (41), $Z^{(t)} = X^{(t)} + \Delta^{(t)}$, and (39), a direct cancellation gives

$$\langle C, X^{(t)} \rangle - \sum_k \langle \chi_k^{(t)}, p_k \rangle = -\lambda^{-1} \langle Y^{(t)}, \Delta^{(t)} \rangle. \tag{45}$$

Lemma 10.1 bounds $\|Y^{(t)}\|$ uniformly. Therefore the absolute primal–dual-gap residual in (45), the dual violation in (43), and the copy residual $Z^{(t)} - X^{(t)}$ are all bounded by a fixed multiple of $\|\Delta^{(t)}\|$. Applying the Hoffman bound to (36) proves (40). □

Finally, an optimal triple $(X^*, X^*, \chi^*) \in \mathcal{S}^*$ induces a DR fixed point

$$Y^* := X^* + \lambda \mathcal{M}^* \chi^* \in \mathcal{Y}^*. \tag{46}$$

This follows from complementary slackness for (34)–(35) and the normal-cone characterizations of the two proximal maps. Furthermore,

$$\|\mathcal{M}^*(\chi - \chi^*)\| \le \sqrt{\sum_{k=1}^{m} d_k} \left(\sum_{k=1}^{m} \|\chi_k - \chi_k^*\|^2 \right)^{1/2}, \tag{47}$$

where $d_k = \prod_{i \neq k} n_i$. Combining (40)–(47) gives a constant $\kappa > 0$ such that

$$\operatorname{dist}(Y^{(t)}, \mathcal{Y}^*) \le \kappa \|Y^{(t)} - Y^{(t-1)}\|. \tag{48}$$

Applying (48) to the Fejér inequality (31) yields

$$\operatorname{dist}(Y^{(t)}, \mathcal{Y}^*)^2 \le \frac{\kappa^2}{\kappa^2 + 1} \operatorname{dist}(Y^{(t-1)}, \mathcal{Y}^*)^2. \tag{49}$$

Thus, with

$$r := \frac{\kappa}{\sqrt{\kappa^2 + 1}} \in (0, 1), \tag{50}$$

we obtain

$$\operatorname{dist}(Y^{(t)}, \mathcal{Y}^*) \le r^t \operatorname{dist}(Y^{(0)}, \mathcal{Y}^*). \tag{51}$$

Because $X^{(t)} = \operatorname{prox}_{\lambda f}(Y^{(t-1)})$ and proximal maps are non-expansive,

$$\operatorname{dist}(X^{(t)}, \mathcal{X}^*) \le r^{t-1} \operatorname{dist}(Y^{(0)}, \mathcal{Y}^*), \tag{52}$$

which proves Theorem 4.2.

11 Cost-aware feasible recovery: implementation details

The finite DR shadow iterate is non-negative but need not yet satisfy every marginal equality. The recovery used in the main experiments begins with the multiplicative down-scaling stage of Lin et al. (2022). Starting from $X \ge 0$, each mode is rescaled so that its current marginal is componentwise no larger than the target marginal. Let $\bar{X}$ denote the resulting tensor and

$$e_k := p_k - r_k(\bar{X}) \ge 0. \tag{53}$$

Because every target marginal has total mass one and every marginal of $\bar{X}$ has total mass $\|\bar{X}\|_1$, all residuals have the same mass

$$\delta := \|e_k\|_1 = 1 - \|\bar{X}\|_1, \qquad \forall k. \tag{54}$$

The rank-one recovery of Lin et al. (2022) adds $\bigotimes_k e_k / \delta^{m-1}$. DRMOT instead defines the low-cost residual transport problem

$$\min_{K \ge 0} \langle C, K \rangle - \eta \|C\|_\infty H(K) \quad \text{s.t.} \quad r_k(K) = e_k, \ \forall k, \tag{55}$$

where $H(K) = -\sum_{\mathbf{i}} K_{\mathbf{i}} (\log K_{\mathbf{i}} - 1)$ and $\eta = 5 \times 10^{-3}$ throughout the experiments. The factor $\|C\|_\infty$ makes η dimensionless and preserves the recovery under positive changes of cost units. The finite-time implementation approximates this subproblem from $\exp[-C/(\eta\|C\|_\infty)]$ using 40 multi-marginal Sinkhorn scaling sweeps; a final Lin-style polish removes the remaining floating-point marginal residual. The returned plan is

$$\widehat{X} = \bar{X} + K, \tag{56}$$

which is non-negative and satisfies $r_k(\widehat{X}) = p_k$ to numerical precision.

The same recovery parameters are used across all DRMOT experiments. Competing entropic methods are not modified: they use the recovery prescribed by their published algorithms, as described in Section 5.1.

Algorithm 2 Cost-aware feasible recovery for a DRMOT iterate

Require: Iterate $X \geq 0$, target marginals $\{p_k\}$, cost C, temperature η, tolerance τ_{fp}

1: Apply the multiplicative marginal down-scaling of Lin et al. (2022) to obtain $\bar{X}$ with $r_k(\bar{X}) \leq p_k$ for every k
2: $e_k \leftarrow p_k - r_k(\bar{X})$ for all k
3: **if** $\|e_1\|_1 \leq \tau_{\text{fp}}$ **then**
4: **return** $\bar{X}$
5: **end if**
6: Initialise $K \propto \exp[-C/(\eta\|C\|_\infty)]$
7: Scale K along its m modes toward the residual marginals $\{e_k\}$ for 40 Sinkhorn sweeps
8: Apply a final Lin-style rank-one polish to the remaining numerical residual
9: **return** $\widehat{X} = \bar{X} + K$

12 Additional experimental and reproducibility details

Baselines. We compare against Multi-Sinkhorn (MSK), accelerated Multi-Sinkhorn (AccMSK) (Lin et al., 2022), and PD-AAM (Tupitsa et al., 2020). The entropic regularization grids are reported in the corresponding main-text experiments. Every baseline retains the update rule and feasible recovery of its source method; the proposed cost-aware recovery is used only inside DRMOT.

Timing and precision. All optimization methods are implemented in PyTorch float64 and run on an NVIDIA RTX 5070 Ti. Each main comparison allocates the same end-to-end wall-clock budget to each $(\text{solver}, \varepsilon)$ configuration. The timed pipeline includes the optimization update and the feasible recovery used to produce a reported solution; monitoring and logging are excluded. Unless stated otherwise, relative objective errors are measured against a high-accuracy HiGHS LP reference. The three cost-family time-to-threshold diagnostic in Section 5.4 instead uses the absolute objective gap and labels it explicitly. As stressed in Section 5.3, agreement below the numerical accuracy of the float64 reference is treated only as agreement between numerical solvers, not as an exact certificate of the mathematical optimum.

Benchmarks. The main experiments include: (i) nested-ellipse Wasserstein barycenters on grids from 20×20 to 50×50; (ii) model-independent basket-option pricing using S&P 500, Dow Jones, and VIX risk-factor marginals; (iii) Coulomb, Brenier–Euler, and risk-pooling cost tensors; (iv) dense-tensor scaling tests for $m \in \{3, 4, 5\}$; and (v) a qualitative four-marginal coupling on the Moignard et al. single-cell blood-development time course of Moignard et al. (2015). The exact seeds, preprocessing, and figure-generation scripts are included in the released reproduction pipeline.

13 Scale-aware step size

The proposed rule follows from the simple equivariance requirement that the DR trajectory should not depend on the arbitrary units in which the cost is expressed. For a positive rescaling $C \mapsto sC$ with $s > 0$, the unregularised MMOT minimizer is unchanged and the DR threshold depends on the product λC. Thus an equivariant rule must transform as $\lambda \mapsto \lambda/s$. For nonzero cost tensors, the dimension-only rule $1/\beta$ cannot satisfy this condition, whereas

$$\lambda_{\text{scaled}} = \frac{c}{\|C\|_\infty \sqrt{\beta}}, \qquad \beta = \sum_{k=1}^{m} \prod_{i \neq k} n_i, \tag{57}$$

has the required transformation by construction.

Across the representative random and Gaussian-barycenter sweeps in Table 4, the normalized quantity $\lambda^\star \sqrt{\beta} \|C\|_\infty$ is substantially more stable than $\lambda^\star \beta$. This empirical observation complements the exact cost-rescaling equivariance of (5).

Table 4: Empirical step-size fit on representative random and Gaussian-barycenter instances ($n = 6$). The scale-aware quantity $\lambda^\star\sqrt{\beta}\|C\|_\infty$ varies much less than the dimension-only quantity $\lambda^\star\beta$.

Family	m	$\lambda^\star$	$\lambda^\star\beta$	$\lambda^\star\sqrt{\beta}\|C\|_\infty$
random	3	2.78×10^{-2}	3.0	0.288
random	4	3.47×10^{-2}	30.0	1.020
random	5	4.63×10^{-3}	30.0	0.373
gaussian	3	9.26×10^{-3}	1.0	0.497
gaussian	4	3.47×10^{-3}	3.0	0.489
gaussian	5	4.63×10^{-3}	30.0	1.813

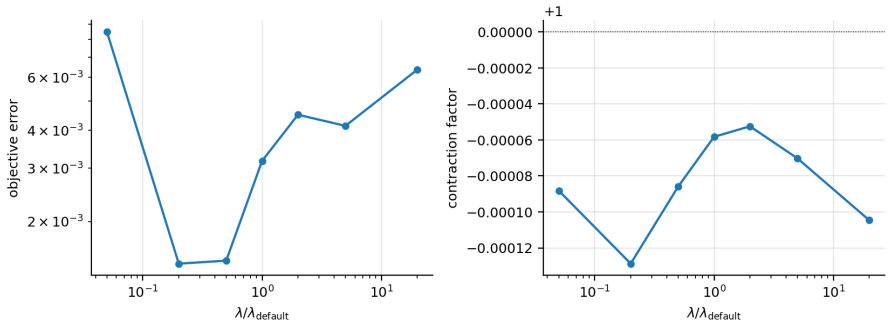


Figure 7: Step-size sensitivity on the ellipse benchmark, $m = 3$. *Left:* objective error after a fixed budget as λ sweeps three decades. *Right:* empirical contraction factor. A broad near-optimal region on this representative benchmark supports the use of a single scale-aware default without claiming instance-wise optimality.

14 Recovery-temperature sensitivity

To check that the fixed recovery temperature is not narrowly tuned, we apply the cost-aware recovery with several values of η to the same saved DRMOT-s iterate on three representative problems. No optimization trajectory is rerun.

Table 5: Relative objective error after cost-aware recovery from the same saved DRMOT-s iterate. The default $\eta = 5 \times 10^{-3}$ remains close to the best tested value across the three cost structures; recovery takes only 23–44 ms in these cases.

Problem	$\eta = 10^{-3}$	$\eta = 5 \times 10^{-3}$	$\eta = 10^{-2}$	$\eta = 5 \times 10^{-2}$
Ellipse 20×20	0.2300%	0.1331%	0.1390%	0.2851%
Coulomb ($n = 40$)	0.0053%	0.0061%	0.0069%	0.0095%
Random ($m = 3, n = 15$)	0.1789%	0.0994%	0.0978%	0.1631%

The dependence on η is problem dependent rather than monotone, but the default 5×10^{-3} is never the poor end of the tested range and is within about 15% of the best tested recovered error on all three problems. This supports using one fixed recovery temperature across the reported experiments without claiming that it is instance-wise optimal.

15 Fixed-iteration comparison

For completeness, Table 6 reports the fixed-iteration protocol used in the earlier version of DRMOT. It is not the primary protocol of this paper because iteration costs differ substantially across solvers. At 50×50, one PD-AAM iteration is approximately $14\times$ more expensive than one DRMOT iteration in our implementation, so equal iteration counts do not represent equal computational budgets.

Table 6: Grid-size benchmark with 2000 iterations for every method. DRMOT uses cost-aware recovery. PD-AAM $\varepsilon = 1$ wins three of four grids under equal iteration counts, whereas the equal-wall-clock comparison in Table 3 favours DRMOT.

Algorithm	ε, λ	20×20	30×30	40×40	50×50
DRMOT	λ_{dim}	0.37%	**0.44%**	0.70%	1.06%
DRMOT-s	λ_{scaled}	0.43%	0.51%	0.66%	1.05%
MSK $\varepsilon = 1$		1.33%	1.68%	5.30%	10.81%
AccMSK $\varepsilon = 1$		1.33%	1.54%	4.74%	10.18%
PD-AAM	0.1	0.27%	0.71%	1.38%	2.31%
	1	**0.27%**	0.48%	**0.32%**	**0.32%**

16 Long-run entropic accuracy floor

Figure 8 extends the Coulomb experiment to 60 seconds of pure solver time. MSK with $\varepsilon = 0.1$ reaches relative error 0.179% by second 8 and remains at essentially the same level for the remaining 52 seconds. Different entropic parameters flatten at different levels, making the bias–conditioning trade-off visible over a longer horizon. DRMOT, which optimizes the unregularised objective directly, continues improving over the same interval.

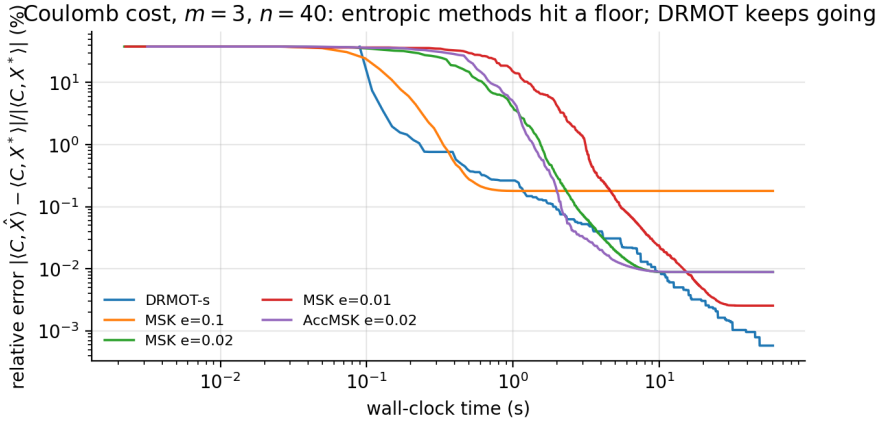


Figure 8: Long-run Coulomb experiment ($m = 3$, $n = 40$). Each tested entropic configuration reaches a parameter-dependent accuracy floor, while DRMOT keeps improving over the 60s horizon.

17 Additional qualitative barycenter

The four-marginal single-cell coupling is shown in the main paper. For completeness, Figure 9 gives the corresponding qualitative Wasserstein-barycenter example throughout the synthetic benchmark.

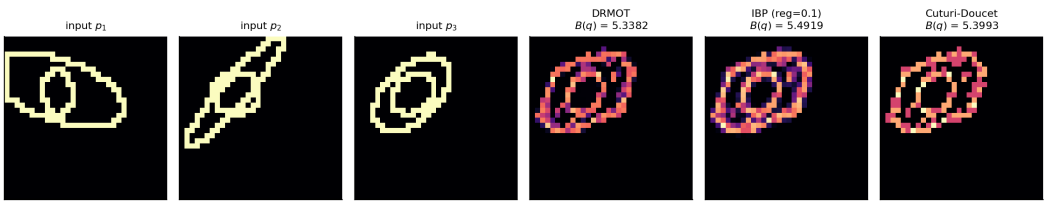


Figure 9: Wasserstein barycenter of three nested ellipses computed by DRMOT (fourth panel). The unregularised solution preserves a sharp geometric structure without the smoothing induced by entropic regularization.