# FORWARD TANGENT-CENTERED SINKHORN DRO: EXACT NOMINAL RECOVERY FROM OPTIMAL TRANSPORT TO KL

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Fixed-reference entropy-regularized transport yields smooth distributionally robust optimization (DRO) objectives, but its value as a function of the adversarial marginal is generally minimized at a kernel-smoothed law rather than at the nominal distribution. Consequently, the smallest raw ambiguity set need not recover nominal risk. Building on prior information-geometric plan-KL orientations, we study the *forward tangent-centered* discrepancy obtained by subtracting the value and first variation at the nominal law. This is the unique affine recentering, up to the simplex gauge; for symmetric costs, its plan-KL representation coincides with the forward candidate noted by Amari et al. (2018). The representation shows that the discrepancy dominates marginal KL, is uniquely zero at the nominal law, and is strongly convex in total variation. The resulting DRO problem retains a scalar log-sum-exp dual and closed-form conditional adversarial weights. On finite supports with symmetric zero-diagonal costs, the discrepancy converges uniformly to the unregularized transport cost at low temperature; for any finite cost, its rescaled value converges uniformly to marginal KL at high temperature. We derive its exact local Hessian, an explicit finite-sample containment radius, and a generalized chi-square calibration for multinomial data. Controlled simulations verify nominal centering, both limiting regimes, and practical plug-in coverage.

## 1 INTRODUCTION

Entropy regularization makes optimal transport (OT) smooth and computationally accessible, and it is now a standard ingredient in transport-based learning and DRO (Cuturi, 2013; Peyré & Cuturi, 2019; Wang et al., 2025). Yet entropy has two distinct effects when one transport marginal is adversarial: it smooths the optimization problem, and it selects a baseline distribution through the reference measure. For fixed-reference Sinkhorn DRO, Wang et al. (2025, Remark 5) identify the zero-budget law as a kernel smoothing of the nominal distribution. Thus a set described as being "centered" at a nominal law $p$ can have its minimum at another law.

The issue is already visible on a finite support. For a cost matrix $C$ and temperature $\gamma > 0$, let

$$T_\gamma(q,p) = \min_{\Pi\mathbf{1}=q,\ \Pi^\mathsf{T}\mathbf{1}=p} \left\{ \langle C, \Pi\rangle + \gamma \sum_{j,i} \Pi_{ji}(\log \Pi_{ji} - 1) \right\}. \tag{1}$$

If $C = 0$, the minimizer over $q$ is uniform, regardless of $p$. Merely shifting the radius to make $p$ feasible does not remove this first-order preference. The natural question is therefore:

> Can one remove the fixed-reference drift while preserving the convexity and log-sum-exp structure that make Sinkhorn DRO useful?

We answer affirmatively by subtracting the tangent plane of $q \mapsto T_\gamma(q,p)$ at $p$, using the forward Bregman orientation dictated by the adversarial marginal. If $\Pi_\gamma^{q,p}$ is the entropic optimal plan and $\Pi_\gamma^{p,p}$ is the nominal self-plan, then

$$D_\gamma^{\mathrm{TC}}(q\|p) = \gamma \mathrm{KL}\big(\Pi_\gamma^{q,p} \big\| \Pi_\gamma^{p,p}\big) \geq \gamma \mathrm{KL}(q\|p). \tag{2}$$



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.

The correction therefore has an information-geometric interpretation at the level of transport plans, controls probability reweighting more sharply than Pinsker's inequality alone, and retains the original cost geometry through the optimal-plan map.

**Contributions.** Our contributions are threefold.

1. For the forward counting-reference orientation, we characterize the raw center, prove that tangent subtraction is the unique affine correction restoring $p$, connect its symmetric-cost plan-KL form to the forward candidate of Amari et al. (2018), and derive its KL chain-rule decomposition.
2. We derive a strong finite-dimensional DRO dual with a single scalar transport multiplier and explicit conditional adversarial weights. For symmetric zero-diagonal discrete costs, we prove a uniform OT limit as $\gamma \downarrow 0$; for any finite cost, we prove a uniform KL limit for $D_\gamma^{\mathrm{TC}}/\gamma$ as $\gamma \to \infty$.
3. We identify the local metric through a conditional-softmax covariance matrix. This yields an explicit nonasymptotic multinomial radius and an asymptotically exact generalized chi-square calibration; controlled simulations evaluate the latter plug-in procedure.

The analysis is intentionally finite-dimensional. It covers categorical laws, scenario reduction, predictive or importance-weighted supports, and other settings in which nominal probabilities are nonuniform. It does *not* provide a continuous-support sampling theorem or account for a learned support or ground cost.

**Closest work.** Amari et al. (2018, Eqs. 49–50) introduce both KL orientations between optimal plans and analyze the self-plan-first orientation. For symmetric costs, their dual candidate with the nominal reference, after transposing plans and matching regularization conventions, coincides with the forward plan-KL quantity in (2). We therefore do not claim that candidate itself as new. Our contribution is to identify it as the unique tangent recentering of fixed/counting-reference transport and develop the consequences needed for DRO. Bregman divergences generated by one-marginal product-reference entropic OT have also appeared as mirror maps (Kerimkulov et al., 2025); the exact relationship is given in (9). General Sinkhorn-DRO duality is due to Wang et al. (2025), and regularized Wasserstein DRO is studied by Azizian et al. (2023b;a).

Symmetric Sinkhorn divergences instead target two-sample comparison and interpolate between OT and a kernel discrepancy (Feydy et al., 2019; Genevay et al., 2019; Séjourné et al., 2019; Aubin-Frankowski et al., 2026). The $n^{-1}$ scale and generalized chi-square mechanism are known for finite-space Sinkhorn losses (Bigot et al., 2019; Goldfeld et al., 2024), and explicit pseudoinverse geometry is available for the symmetric divergence (Lavenant et al., 2026). Our distinctive results are the forward fixed-reference Sinkhorn-DRO reformulation and worst-case law, its assumption-explicit low-temperature OT and general high-temperature KL limits, and the compact conditional-softmax Hessian used for one-sided ambiguity-radius calibration. See the appendix for a precise algebraic comparison. Regularized OT's broader Bregman geometry is surveyed by Dessein et al. (2018).

## 2 Fixed-reference drift and canonical recentering

### 2.1 Setup and the raw center

Let $\Delta_m = \{q \in \mathbb{R}_+^m : \mathbf{1}^\mathsf{T} q = 1\}$ and $\Delta_m^\circ = \{q \in \Delta_m : q_j > 0\}$. Throughout, $p \in \Delta_m^\circ$, $C \in \mathbb{R}_+^{m\times m}$ is finite, and $0 \log 0 = 0$. The coupling polytope is $\Gamma(q,p) = \{\Pi \geq 0 : \Pi\mathbf{1} = q, \Pi^\mathsf{T}\mathbf{1} = p\}$, and $T_\gamma$ is defined by (1). This is entropy relative to counting measure, up to a constant. Multiplying a fixed reference matrix by positive row or column factors only adds affine terms in $q$ and constants in $p$; Theorem 2.4 below shows that tangent centering removes them. A reference that itself depends on the varying $q$, such as $qp^\mathsf{T}$, defines a different geometry.

The raw minimizer is explicit.

**Proposition 2.1** (Raw fixed-reference center)**.** *For $K_{ji} = e^{-C_{ji}/\gamma}$ and $Z_i = \sum_k K_{ki}$,*

$$\Pi^\circ_{ji} = p_i \frac{K_{ji}}{Z_i}, \qquad q^\circ_{\gamma,j}(p) = \sum_i p_i \frac{K_{ji}}{Z_i} \tag{3}$$

*uniquely minimize $T_\gamma(q,p)$ over $\Pi$ and $q$, respectively. Hence $q^\circ_\gamma(p) \to m^{-1}\mathbf{1}$ as $\gamma \to \infty$. If $C_{ii} = 0$ and $\Delta_C = \min_{i\neq j} C_{ji} > 0$, then $\left\|q^\circ_\gamma(p) - p\right\|_1 \le 2(m-1)e^{-\Delta_C/\gamma}$.*

The formula is the discrete kernel smoothing already implicit in the zero-budget representation of Wang et al. (2025). Our goal is not to claim this smoothing as new, but to remove it canonically without sacrificing tractability.

### 2.2 Tangent centering and a plan-level identity

For fixed $p$, write $F_{\gamma,p}(q) = T_\gamma(q,p)$. This function is strictly convex on $\Delta_m$ and smooth on $\Delta_m^\circ$. Its relative gradient is defined up to a multiple of $\mathbf{1}$. Choose any representative $\nu_{\gamma,p} \in \partial F_{\gamma,p}(p)$ and define

$$D^{\mathrm{TC}}_\gamma(q\|p) := F_{\gamma,p}(q) - F_{\gamma,p}(p) - \nu_{\gamma,p}^{\mathsf{T}}(q-p). \tag{4}$$

Thus $D^{\mathrm{TC}}_\gamma$ is an *anchored* Bregman remainder: the convex generator $F_{\gamma,p}$ depends on the anchor $p$.

**Proposition 2.2** (Canonical affine correction)**.** *For $a \in \mathbb{R}^m$, let $D_a(q,p) = F_{\gamma,p}(q) - F_{\gamma,p}(p) - a^{\mathsf{T}}(q-p)$. Then $p$ minimizes $D_a(\cdot,p)$ on $\Delta_m$ if and only if $a - \nu_{\gamma,p} \in \operatorname{span}\{\mathbf{1}\}$. Whenever this holds, strict convexity makes $p$ the unique minimizer. Hence (4) is the unique affine recentering on the simplex, up to gauge.*

The next theorem is the main structural result. For distributions $P, Q$ on a finite set, we use $\mathrm{KL}(P\|Q) = \sum_x P_x \log(P_x/Q_x)$.

**Theorem 2.3** (Exact plan-KL geometry)**.** *Let $\Pi^{q,p}_\gamma$ be the unique optimizer defining $T_\gamma(q,p)$. For every $q \in \Delta_m$,*

$$D^{\mathrm{TC}}_\gamma(q\|p) = \gamma \mathrm{KL}\big(\Pi^{q,p}_\gamma \big\| \Pi^{p,p}_\gamma\big) \tag{5}$$

$$= \gamma \mathrm{KL}(q\|p) + \gamma \sum_{j:q_j>0} q_j \mathrm{KL}\big(\Pi^{q,p}_\gamma(\cdot \mid j) \big\| \Pi^{p,p}_\gamma(\cdot \mid j)\big)\,. \tag{6}$$

*More generally, for $r \in \Delta_m^\circ$, the Bregman remainder of the fixed generator $F_{\gamma,p}$ satisfies*

$$B_{F_{\gamma,p}}(q,r) = \gamma \mathrm{KL}(\Pi^{q,p}_\gamma \| \Pi^{r,p}_\gamma) \ge \gamma \mathrm{KL}(q\|r) \ge \frac{\gamma}{2}\|q-r\|_1^2\,. \tag{7}$$

*Consequently, $D^{\mathrm{TC}}_\gamma(\cdot\|p)$ is convex, uniquely zero at $p$, and*

$$D^{\mathrm{TC}}_\gamma(q\|p) \ge \gamma \mathrm{KL}(q\|p) \ge \frac{\gamma}{2}\|q-p\|_1^2\,. \tag{8}$$

*Thus, for $\rho \ge 0$, $\mathcal{U}^{\mathrm{TC}}_{\rho,\gamma}(p) := \{q : D^{\mathrm{TC}}_\gamma(q\|p) \le \rho\}$ is nonempty, compact, and convex, with $\mathcal{U}^{\mathrm{TC}}_{0,\gamma}(p) = \{p\}$ and $\mathcal{U}^{\mathrm{TC}}_{\rho,\gamma}(p) \subseteq \{q : \mathrm{KL}(q\|p) \le \rho/\gamma\}$.*

Identity (5) says that tangent centering replaces an arbitrary additive convention by a genuine relative entropy between two *optimized* couplings. The residual term in (6) measures how much the nominal scenario conditioned on an adversarial state changes beyond marginal reweighting.

This orientation can also be related exactly to prior product-reference EOT Bregman geometry. Define $h_{\gamma,p}(q) = \min_{\Pi\in\Gamma(q,p)}\{\langle C,\Pi\rangle + \gamma\mathrm{KL}(\Pi\|qp^{\mathsf{T}})\}$ and let $B_{h_{\gamma,p}}(q,p)$ be its forward Bregman remainder. Marginal identities give

$$D^{\mathrm{TC}}_\gamma(q\|p) = B_{h_{\gamma,p}}(q,p) + \gamma\mathrm{KL}(q\|p). \tag{9}$$

Thus the conditional term in (6) is precisely the product-reference remainder, while counting-reference entropy supplies the marginal KL and strong convexity. For symmetric costs, (5) is the dual candidate recorded in Amari et al. (2018, Eq. 50); our results identify its tangent-centering origin and develop its DRO geometry. Equation (9) likewise distinguishes it from the general product-reference mirror map of Kerimkulov et al. (2025).

**Proposition 2.4** (Fixed-reference invariance)**.** *Let $R \in \mathbb{R}^{m\times m}_{++}$ be a fixed entropy reference and let $R' = \mathrm{Diag}(r)R\,\mathrm{Diag}(c)$ for $r,c>0$. The tangent-centered discrepancies generated by $R$ and $R'$ coincide. Equivalently, adding arbitrary row and column potentials to $C$ leaves $D^{\mathrm{TC}}_\gamma$ unchanged.*

Table 1: Three ways to treat fixed-reference entropic bias. "Simple dual" means the raw scalar log-sum-exp form, possibly after a cost shift.

| | Raw $T_\gamma$ | Symmetric Sinkhorn | Tangent-centered |
|---|---|---|---|
| Nominal is zero/minimizer | generally no | yes$^\dagger$ | yes |
| One-sided convexity | yes | conditional$^\dagger$ | yes |
| Simple DRO dual | yes | generally no | yes |
| High-$\gamma$ geometry | reference-driven | kernel/MMD$^\dagger$ | KL after scaling |

$^\dagger$ Under the kernel/cost assumptions of Feydy et al. (2019); Aubin-Frankowski et al. (2026).

## 3 A TRACTABLE DISTRIBUTIONALLY ROBUST OBJECTIVE

Let $L(\theta) = (L_1(\theta), \dots, L_m(\theta))^\mathsf{T}$ and define

$$\Phi^{\mathrm{TC}}_{\rho,\gamma}(\theta) = \sup_{q \in \mathcal{U}^{\mathrm{TC}}_{\rho,\gamma}(p)} q^\mathsf{T} L(\theta). \tag{10}$$

Normalize the gauge of $\nu = \nu_{\gamma,p}$ so that $p^\mathsf{T}\nu = T_\gamma(p,p)$, and set $\bar{C}_{ji} = C_{ji} - \nu_j$. Then the nominal self-plan has corrected budget zero, and (10) is equivalently a maximization over $\Pi^\mathsf{T}\mathbf{1} = p$ subject to $\langle \bar{C}, \Pi \rangle + \gamma \sum_{j,i} \Pi_{ji}(\log \Pi_{ji} - 1) \le \rho$.

**Theorem 3.1** (Scalar strong dual)**.** *Suppose $\rho > 0$ and $L_j(\theta)$ is finite. Define $Z_i(\theta,\lambda) := p_i^{-1} \sum_{j=1}^m \exp\big([L_j(\theta)/\lambda - \bar{C}_{ji}]/\gamma\big)$. Then*

$$\Phi^{\mathrm{TC}}_{\rho,\gamma}(\theta) = \inf_{\lambda > 0} \lambda \left[ \rho + \gamma + \gamma \sum_{i=1}^m p_i \log Z_i(\theta,\lambda) \right]. \tag{11}$$

*If the infimum is attained at $\lambda^* > 0$, the unique worst-case coupling is*

$$\begin{aligned} \Pi^*_{ji} &= p_i \frac{\exp\big((L_j(\theta) - \lambda^* \bar{C}_{ji})/(\lambda^*\gamma)\big)}{\sum_k \exp\big((L_k(\theta) - \lambda^* \bar{C}_{ki})/(\lambda^*\gamma)\big)}, \\ q^*_j &= \sum_i \Pi^*_{ji}. \end{aligned} \tag{12}$$

*If each $L_j$ is convex in $\theta$, the joint objective before minimizing over $\lambda$ is convex in $(\theta,\lambda)$. If the losses are differentiable and $q^*$ is unique, then $\nabla_\theta \Phi^{\mathrm{TC}}_{\rho,\gamma}(\theta) = \sum_j q^*_j \nabla L_j(\theta)$.*

The self-potential is computed once by Sinkhorn scaling. Each subsequent dual evaluation costs $O(m^2)$ for a dense kernel and uses stable columnwise log-sum-exp operations. Formula (12) is conditional on the scalar multiplier; the $\lambda \downarrow 0$ boundary has the usual hard-max interpretation. At $\rho = 0$, Theorem 2.3 gives the exact identity $\Phi^{\mathrm{TC}}_{0,\gamma}(\theta) = p^\mathsf{T} L(\theta)$ for every $\gamma > 0$.

## 4 A CONTINUUM FROM DISCRETE OT TO KL

Let $W_C(q,p) = \min_{\Pi \in \Gamma(q,p)} \langle C, \Pi \rangle$. We call this a discrete OT cost; it is a Wasserstein cost when $C$ is induced by an appropriate ground metric. The two temperature limits are uniform, which is important because the adversary also varies with $\gamma$.

**Theorem 4.1** (Uniform limiting geometry)**.** *Assume first that $C = C^\mathsf{T}$, $C_{ii} = 0$, and $p_i \ge p_{\min} > 0$. Set $\varepsilon_\gamma := 2\gamma[1 + 2\log m + \log(\sqrt{m}/p_{\min})]$. Then, for every $\gamma > 0$,*

$$\sup_{q \in \Delta_m} \big| D^{\mathrm{TC}}_\gamma(q\|p) - W_C(q,p) \big| \le \varepsilon_\gamma. \tag{13}$$

*For any finite $C$ and $p \in \Delta^\circ_m$,*

$$\sup_{q \in \Delta_m} \big| \gamma^{-1} D^{\mathrm{TC}}_\gamma(q\|p) - \mathrm{KL}(q\|p) \big| \longrightarrow 0 \qquad (\gamma \to \infty). \tag{14}$$

The scaling in (14) is essential. At fixed $\rho$, (8) implies that $\mathcal{U}^{\mathrm{TC}}_{\rho,\gamma}(p)$ collapses to $\{p\}$ as $\gamma \to \infty$. A nondegenerate KL limit uses $\rho = \gamma\tau$: for every $\tau > 0$, the sets $\mathcal{U}^{\mathrm{TC}}_{\gamma\tau,\gamma}(p)$ converge in Hausdorff distance to $\{q : \mathrm{KL}(q\|p) \le \tau\}$.

The uniform bounds also transfer directly to robust risks. If $\sup_\theta \|L(\theta)\|_\infty \le B$ and $\varepsilon_\gamma < \rho$, then

$$\sup_\theta \left| \Phi^{\mathrm{TC}}_{\rho,\gamma}(\theta) - \sup_{q: W_C(q,p) \le \rho} q^\mathsf{T} L(\theta) \right| \le \frac{2B\varepsilon_\gamma}{\rho}. \tag{15}$$

Likewise, at radius $\gamma\tau$, bounded-loss robust risks converge uniformly to KL-DRO risks. Thus $\gamma$ controls a genuine modeling continuum, not merely numerical smoothing.

## 5 Local geometry and statistical calibration

The semi-dual log-partition function is

$$A_{\gamma,p}(u) = \gamma \sum_i p_i \log \sum_j \exp((u_j - C_{ji})/\gamma), \quad s_{ji}(u) = \frac{e^{(u_j - C_{ji})/\gamma}}{\sum_k e^{(u_k - C_{ki})/\gamma}}. \tag{16}$$

The self-potential satisfies $\sum_i p_i s_i(\nu_{\gamma,p}) = p$. Define

$$M_{\gamma,p} = \sum_i p_i \left[ \mathrm{Diag}(s_i(\nu_{\gamma,p})) - s_i(\nu_{\gamma,p}) s_i(\nu_{\gamma,p})^\mathsf{T} \right], \qquad \mathcal{T} = \{v : \mathbf{1}^\mathsf{T} v = 0\}. \tag{17}$$

**Theorem 5.1** (Local metric). *$M_{\gamma,p}$ is positive definite on $\mathcal{T}$, and the intrinsic Hessian of $q \mapsto T_\gamma(q,p)$ at $q = p$ is*

$$H_{\gamma,p} = \gamma M^\dagger_{\gamma,p}. \tag{18}$$

*Here $\dagger$ denotes the Moore–Penrose pseudoinverse, equivalently inversion on $\mathcal{T}$. Consequently, $D^{\mathrm{TC}}_\gamma(q\|p) = \frac{1}{2}(q-p)^\mathsf{T} H_{\gamma,p}(q-p) + o(\|q-p\|_2^2)$. Moreover, for every $v \in \mathcal{T}$,*

$$v^\mathsf{T} H_{\gamma,p} v \ge \gamma \sum_j \frac{v_j^2}{p_j}. \tag{19}$$

The lower bound is the local counterpart of $D^{\mathrm{TC}}_\gamma \ge \gamma\mathrm{KL}$ and follows also from the law of total covariance.

Now suppose $p^\star \in \Delta^\circ_m$ generates $N \sim \mathrm{Multinomial}(n, p^\star)$. To keep the nominal law interior, use the add-$a$ estimator $\widetilde{p}_{n,j} = (N_j + a)/(n + ma)$, $a > 0$.

**Theorem 5.2** (Finite-sample and asymptotic radii). *Assume $p^\star_j \ge \underline{p} > 0$ and $0 < \delta < 1$, let $M_C = \max_{j,i} C_{ji}$, and define*

$$r_{n,\delta,a} = \frac{n}{n+ma}\sqrt{\frac{m}{2n}\log\frac{2m}{\delta}} + \frac{a\sqrt{m(m-1)}}{n+ma}. \tag{20}$$

*If $r_{n,\delta,a} \le \underline{p}/2$, then with probability at least $1-\delta$,*

$$D^{\mathrm{TC}}_\gamma(p^\star\|\widetilde{p}_n) \le \rho^{\mathrm{FS}}_{n,\delta,\gamma} := \frac{\gamma e^{2M_C/\gamma}}{\underline{p}} r^2_{n,\delta,a}. \tag{21}$$

*For fixed $(m, a, \gamma)$,*

$$\begin{aligned} 2nD^{\mathrm{TC}}_\gamma(p^\star\|\widetilde{p}_n) \;&\Rightarrow\; Z^\mathsf{T} H_{\gamma,p^\star} Z, \\ Z \sim \mathcal{N}(0, \Sigma^\star), \qquad &\Sigma^\star = \mathrm{Diag}(p^\star) - p^\star (p^\star)^\mathsf{T}. \end{aligned} \tag{22}$$

*If $c_{1-\delta}(\gamma, p^\star)$ is the corresponding quantile, then $c_{1-\delta}/(2n)$ has asymptotic coverage $1-\delta$. Replacing $(p^\star, H, \Sigma^\star)$ by their add-$a$ plug-in estimates is consistent. If the conditional quantile is estimated by Monte Carlo, the number of Gaussian draws must satisfy $B_n \to \infty$.*

The nonzero eigenvalues of $(\Sigma^\star)^{1/2} H_{\gamma,p^\star} (\Sigma^\star)^{1/2}$ give the representation $Z^\mathsf{T} H Z \overset{d}{=} \sum_{r=1}^{m-1} \lambda_r \chi^2_{1,r}$. The second-order delta-method mechanism follows Bigot et al. (2019); the distinctive ingredient here is the explicit forward Hessian (18) and its use in calibrating a one-sided DRO set. As $\gamma \to \infty$, $H_{\gamma,p}/\gamma$ approaches the Fisher metric and the normalized limit becomes $\chi^2_{m-1}$; for $C = 0$, this statement is exact at every $\gamma$.

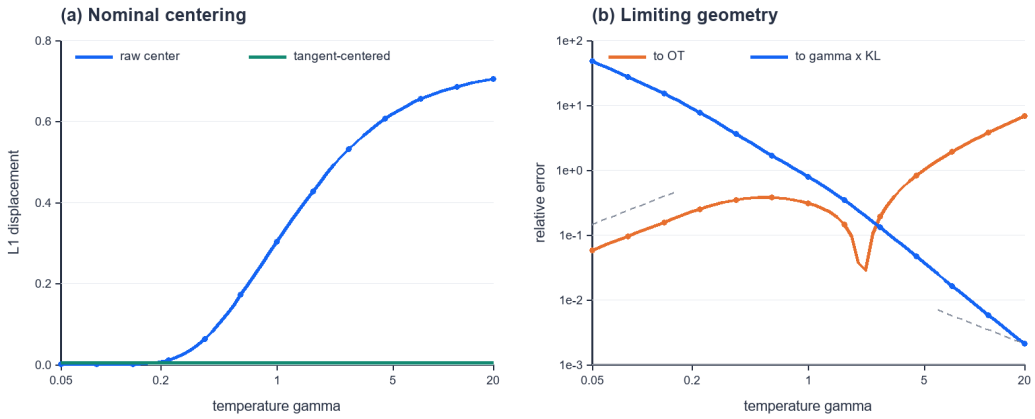



Figure 1: Controlled validation. (a) Entropy moves the raw fixed-reference center toward uniform, whereas tangent centering fixes the unique minimizer at $p$. (b) The orange curve is $\left|D_\gamma^{\mathrm{TC}} - W_C\right| / W_C$ and the blue curve is $\left|D_\gamma^{\mathrm{TC}}/\gamma - \mathrm{KL}\right| / \mathrm{KL}$; dashed lines are slope-$+1$ and slope-$-1$ guides.

Table 2: Plug-in coverage of the nominal 95% radius. Parentheses give Monte Carlo standard errors; the final column is the median data-dependent radius.

| $\gamma$ | $n$ | Coverage (s.e.) | Median radius |
|---|---|---|---|
| 0.2 | 50 | 0.9520 (0.0048) | 0.02499 |
| 0.2 | 100 | 0.9545 (0.0047) | 0.01244 |
| 0.2 | 500 | 0.9470 (0.0050) | 0.00248 |
| 1.0 | 50 | 0.9395 (0.0053) | 0.09606 |
| 1.0 | 100 | 0.9495 (0.0049) | 0.04803 |
| 1.0 | 500 | 0.9565 (0.0046) | 0.00961 |
| 5.0 | 50 | 0.9355 (0.0055) | 0.47440 |
| 5.0 | 100 | 0.9535 (0.0047) | 0.23720 |
| 5.0 | 500 | 0.9475 (0.0050) | 0.04744 |

## 6 Controlled validation

We test centering, the temperature limits, and the asymptotic plug-in calibration. All reported values were generated by `code/reproduce.py` with seed 20260725.

**Centering and limiting geometry.** For $p = (0.70, 0.20, 0.10)$ on locations $(0, 1, 2)$ with squared-distance cost, Figure 1a shows the raw-center displacement from (3). It grows from $2.1 \times 10^{-9}$ at $\gamma = 0.05$ to $0.704$ at $\gamma = 20$, while the tangent-centered minimizer remains exactly $p$. For $q = (0.45, 0.35, 0.20)$, Figure 1b solves the entropic transport plans and evaluates the exact plan-KL identity: the corrected discrepancy approaches $W_C$ at low temperature and $\gamma \mathrm{KL}(q\|p)$ at high temperature, as predicted by Theorem 4.1.

**Practical plug-in coverage.** We use $p^\star = (0.40, 0.25, 0.15, 0.12, 0.08)$, $x_j = (j-1)/4$, $C_{ji} = (x_j - x_i)^2$, and $a = 1/2$. For each sample, both $H$ and $\Sigma$ are evaluated at $\widetilde{p}_n$; the 0.95 generalized chi-square quantile uses a fixed bank of 5,000 Gaussian draws. Table 2 reports 2,000 independent multinomial repetitions. Coverage is within roughly three Monte Carlo standard errors of 0.95 in every configuration, including the more curved $\gamma = 5$ regime.

**DRO primal–dual consistency.** For $p = (0.55, 0.30, 0.15)$, $L = (0.10, 1.20, -0.40)$, $\gamma = 0.7$, and $\rho = 0.12$, scalar minimization of (11) gives $\lambda^* = 1.310730393$. Equation (12) recovers $q^* = (0.395776, 0.541719, 0.062505)$ with $D_\gamma^{\mathrm{TC}}(q^*\|p) = 0.119999996$. The primal and dual values are $0.664638662$ and $0.664638668$, an absolute gap of $5.8 \times 10^{-9}$.

This table evaluates the asymptotic plug-in radius, not the conservative finite-sample radius (21). At $\delta = 0.05$, (20) gives $r_{n,\delta,a} = 0.533, 0.377, 0.166$ for $n = 50, 100, 500$, all larger than the sufficient threshold $\underline{p}/2 = 0.04$. The experiments validate algebra, limits, and practical asymptotic calibration; they do not establish superiority on a downstream benchmark. The method is most relevant when the finite support is decision-relevant and its nominal weights are nonuniform.

## 7 Limitations and conclusion

Tangent centering cleanly separates nominal weights from fixed-reference entropic drift. It is canonical among affine corrections, admits the exact plan-KL representation (5), preserves a scalar Sinkhorn-DRO dual, and connects discrete OT to KL through temperature limits under the stated cost assumptions. Its local metric also supports both conservative and asymptotically exact ambiguity radii.

The guarantees are conditional on a known finite support and fixed cost. Support discretization, learned representations, and cost estimation require additional error terms or sample splitting. Dense evaluation costs $O(m^2)$, while a full Hessian pseudoinverse costs $O(m^3)$; sparse kernels and Hessian-vector methods are needed at large $m$. Finally, tangent centering is deliberately asymmetric. Symmetric Sinkhorn divergences remain preferable for two-sample comparison, whereas DRO has a distinguished nominal anchor.

### Reproducibility Statement

All assumptions are stated in the main text, complete proofs are provided in the appendix, and the supplemental script regenerates every numerical value and figure from a fixed seed using only NumPy and Pillow. The appendix records tolerances, repetition counts, and quantile construction.

## REFERENCES


Shun-ichi Amari, Ryo Karakida, and Masafumi Oizumi. Information Geometry connecting Wasserstein distance and Kullback–Leibler divergence via the entropy-relaxed transportation problem. *Information Geometry*, 1(1):13–37, 2018. doi: 10.1007/s41884-018-0002-8.

Pierre-Cyril Aubin-Frankowski, Virginie Ehrlacher, and Gabriele Todeschi. Debiasing optimal transport: Classical and entropic. *arXiv preprint arXiv:2604.19981*, 2026.

Waïss Azizian, Franck Iutzeler, and Jérôme Malick. Exact generalization guarantees for (regularized) Wasserstein distributionally robust models. In *Advances in Neural Information Processing Systems*, volume 36, 2023a.

Waïss Azizian, Franck Iutzeler, and Jérôme Malick. Regularization for Wasserstein distributionally robust optimization. *ESAIM: Control, Optimisation and Calculus of Variations*, 29:33, 2023b. doi: 10.1051/cocv/2023019.

Jérémie Bigot, Elsa Cazelles, and Nicolas Papadakis. Central limit theorems for entropy-regularized optimal transport on finite spaces and statistical applications. *Electronic Journal of Statistics*, 13 (2):5120–5150, 2019. doi: 10.1214/19-EJS1637.

Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, volume 26, pp. 2292–2300, 2013.

Arnaud Dessein, Nicolas Papadakis, and Jean-Luc Rouas. Regularized optimal transport and the ROT mover's distance. *Journal of Machine Learning Research*, 19(15):1–53, 2018.

Jean Feydy, Thibault Séjourné, François-Xavier Vialard, Shun-ichi Amari, Alain Trouvé, and Gabriel Peyré. Interpolating between optimal transport and MMD using Sinkhorn divergences. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 2681–2690. PMLR, 2019.

Aude Genevay, Lénaïc Chizat, Francis Bach, Marco Cuturi, and Gabriel Peyré. Sample complexity of Sinkhorn divergences. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pp. 1574–1583. PMLR, 2019.

Ziv Goldfeld, Kengo Kato, Gabriel Rioux, and Ritwik Sadhu. Limit theorems for entropic optimal transport maps and Sinkhorn divergence. *Electronic Journal of Statistics*, 18(1):980–1041, 2024. doi: 10.1214/24-EJS2217.

Bekzhan Kerimkulov, David Šiška, Łukasz Szpruch, and Yufei Zhang. Mirror descent for stochastic control problems with measure-valued controls. *Stochastic Processes and their Applications*, 190: 104765, 2025. doi: 10.1016/j.spa.2025.104765.

Hugo Lavenant, Jonas Luckhardt, Gilles Mordant, Bernhard Schmitzer, and Luca Tamanini. The Riemannian geometry of Sinkhorn divergences. *Annales de l'Institut Henri Poincaré C, Analyse non linéaire*, 43(5):1079–1144, 2026. doi: 10.4171/AIHPC/165.

Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5–6):355–607, 2019. doi: 10.1561/2200000073.

Thibault Séjourné, Jean Feydy, François-Xavier Vialard, Alain Trouvé, and Gabriel Peyré. Sinkhorn divergences for unbalanced optimal transport. *arXiv preprint arXiv:1910.12958*, 2019.

Jie Wang, Rui Gao, and Yao Xie. Sinkhorn distributionally robust optimization. *Operations Research*, 74(3):1581–1603, 2025. doi: 10.1287/opre.2023.0294.

# A CONVEX-ANALYTIC PRELIMINARIES

This appendix supplies complete proofs and implementation details. All derivatives on the simplex are relative to its affine hull. Adding a multiple of $\mathbf{1}$ to a potential never changes a pairing with $q - p$.

## A.1 SEMI-DUAL AND SMOOTHNESS

Define $A_{\gamma,p}$ and $s_i(u)$ as in (16).

**Lemma A.1** (Semi-dual representation)**.** *For $q, p \in \Delta_m^\circ$,*

$$T_\gamma(q,p) = \sup_{u \in \mathbb{R}^m} \{u^\mathsf{T} q - A_{\gamma,p}(u)\} + \gamma \sum_i p_i \log p_i - \gamma. \tag{23}$$

*Moreover,*

$$\nabla A_{\gamma,p}(u) = \sum_i p_i s_i(u), \qquad \nabla^2 A_{\gamma,p}(u) = \frac{1}{\gamma} \sum_i p_i [\mathrm{Diag}(s_i(u)) - s_i(u) s_i(u)^\mathsf{T}]. \tag{24}$$

*The maximizing potential is unique modulo* $\mathrm{span}\{\mathbf{1}\}$.

*Proof.* Introduce row and column potentials $u, v$ for the two marginal constraints. The scalar identity

$$\inf_{x \geq 0} \{\gamma x(\log x - 1) + (C_{ji} - u_j - v_i)x\} = -\gamma \exp((u_j + v_i - C_{ji})/\gamma)$$

gives

$$T_\gamma(q,p) = \sup_{u,v} \left\{ u^\mathsf{T} q + v^\mathsf{T} p - \gamma \sum_{j,i} e^{(u_j + v_i - C_{ji})/\gamma} \right\}.$$

For fixed $u$, first-order optimality in $v_i$ gives $v_i = \gamma \log p_i - \gamma \log \sum_j e^{(u_j - C_{ji})/\gamma}$. Substitution proves (23); differentiation gives (24).

Every $s_i(u)$ has full support. Hence for any nonconstant $z$,

$$z^\mathsf{T} \nabla^2 A_{\gamma,p}(u) z = \gamma^{-1} \sum_i p_i \, \mathrm{Var}_{s_i(u)}(z) > 0.$$

The Hessian nullspace is exactly $\mathrm{span}\{\mathbf{1}\}$. After fixing, for example, the gauge $\mathbf{1}^\mathsf{T} u = 0$, strict convexity and the usual finite Sinkhorn scaling theorem yield a unique optimizer. Equivalently, the unique positive entropic plan determines both potentials up to the common additive gauge. □

The strict positivity of the Gibbs kernel and $p \in \Delta_m^\circ$ imply that the self-plan $\Pi_\gamma^{p,p}$ has strictly positive entries. The value $T_\gamma(q,p)$ and its unique optimizer are continuous in $(q,p)$ on the closed product of simplices; this follows directly from compactness of the coupling correspondence, continuity of the objective with $0 \log 0 = 0$, and uniqueness of the optimizer. On compact subsets of the interior, the value is $C^\infty$ by Theorem A.1 and the implicit-function theorem on the tangent space.

# B PROOFS FOR THE RAW CENTER AND TANGENT GEOMETRY

## B.1 PROOF OF THEOREM 2.1

Minimizing over $q$ removes the row-sum constraint, so the problem separates over columns subject only to $\Pi^\mathsf{T} \mathbf{1} = p$. For $p_i > 0$, write $\Pi_{ji} = p_i s_j$ with $s \in \Delta_m$. The column objective is

$$p_i \sum_j s_j C_{ji} + \gamma p_i \left( \log p_i - 1 + \sum_j s_j \log s_j \right).$$

Its unique minimizer is $s_j = K_{ji}/Z_i$. (A column with $p_i = 0$ is identically zero.) Averaging the resulting column probabilities with weights $p_i$ proves (3). The minimum value, useful for checking implementations, is

$$\min_q T_\gamma(q,p) = \gamma \sum_i p_i (\log p_i - 1 - \log Z_i), \tag{25}$$

with the convention $0 \log 0 = 0$. Since $K_{ji} \to 1$ as $\gamma \to \infty$, every normalized column converges to $m^{-1}\mathbf{1}$.

For the low-temperature bound, the mass that the $i$th Gibbs column assigns away from $i$ is at most $\sum_{j \neq i} e^{-C_{ji}/\gamma} \leq (m-1)e^{-\Delta_C/\gamma}$. The $\ell_1$ distance between that column and the $i$th point mass is twice its off-diagonal mass. Averaging over $p_i$ proves the claim.

### B.2 Proof of Theorem 2.2

Because $p$ lies in the relative interior, it minimizes the differentiable convex function $D_a(\cdot, p)$ on $\Delta_m$ if and only if its relative directional derivative vanishes in every $v \in \mathcal{T}$. That condition is $(\nu_{\gamma,p} - a)^{\mathsf{T}} v = 0$ for every $v \in \mathcal{T}$, equivalently $a - \nu_{\gamma,p} \in \mathcal{T}^{\perp} = \operatorname{span}\{\mathbf{1}\}$. Strict convexity of $F_{\gamma,p}$ then makes the minimizer unique. The constant term is already forced by the normalization $D_a(p,p) = 0$.

### B.3 Proof of Theorem 2.3

First let $q \in \Delta_m^{\circ}$, and let $u_q$ and $\nu = \nu_{\gamma,p}$ be semi-dual potentials satisfying $\nabla A_{\gamma,p}(u_q) = q$ and $\nabla A_{\gamma,p}(\nu) = p$. Up to the constant in (23), $F_{\gamma,p} = A^*_{\gamma,p}$. The standard conjugate-Bregman identity gives

$$D_{F_{\gamma,p}}(q,p) = D_{A_{\gamma,p}}(\nu, u_q). \tag{26}$$

For a single column, direct expansion of the log-partition function gives

$$\begin{aligned} &\gamma \log \sum_j e^{(\nu_j - C_{ji})/\gamma} - \gamma \log \sum_j e^{(u_{q,j} - C_{ji})/\gamma} - s_i(u_q)^{\mathsf{T}}(\nu - u_q) \\ &\qquad = \gamma \mathrm{KL}(s_i(u_q) \| s_i(\nu)). \end{aligned}$$

Multiplying by $p_i$ and summing over columns shows

$$D_\gamma^{\mathrm{TC}}(q\|p) = \gamma \sum_i p_i \mathrm{KL}(s_i(u_q) \| s_i(\nu)). \tag{27}$$

The optimal plans satisfy $\Pi^{q,p}_{\gamma,ji} = p_i s_{ji}(u_q)$ and $\Pi^{p,p}_{\gamma,ji} = p_i s_{ji}(\nu)$, so (27) is exactly (5).

The identical calculation with any $r \in \Delta_m^{\circ}$ and its semi-dual potential $u_r$ in place of $(p, \nu)$ gives

$$B_{F_{\gamma,p}}(q,r) = \gamma \mathrm{KL}(\Pi^{q,p}_\gamma \| \Pi^{r,p}_\gamma).$$

Data processing under the row-marginal map gives the first inequality in (7); Pinsker gives the second. Hence $F_{\gamma,p}$, and therefore its affine shift $D_\gamma^{\mathrm{TC}}(\cdot\|p)$, is $\gamma$-strongly convex with respect to $\ell_1$ in the Bregman sense.

For boundary $q$, take $q_t = (1-t)q + tp$ and let $t \downarrow 0$. Continuity of the entropic value and optimizer gives $\Pi^{q_t,p}_\gamma \to \Pi^{q,p}_\gamma$. Since the reference self-plan is strictly positive and $x \mapsto x \log x$ is continuous at zero, both sides of (5) converge, proving the identity on the closed simplex.

Apply the KL chain rule to the row marginal map $(j,i) \mapsto j$:

$$\mathrm{KL}(\Pi^{q,p}_\gamma \| \Pi^{p,p}_\gamma) = \mathrm{KL}(q\|p) + \sum_{j: q_j > 0} q_j \mathrm{KL}(\Pi^{q,p}_\gamma(\cdot \mid j) \| \Pi^{p,p}_\gamma(\cdot \mid j)).$$

This proves (6) and the first inequality in (8); Pinsker's inequality proves the second. Convexity follows directly from the defining Bregman remainder in its first argument. Continuity and compactness of the simplex give compact sublevel sets, and strict positivity of KL away from equality gives exact centering.

### B.4 Proof of Theorem 2.4 and arbitrary fixed references

For a fixed $R \in \mathbb{R}^{m \times m}_{++}$, define

$$T_{\gamma,R}(q,p) = \min_{\Pi \in \Gamma(q,p)} \left\{ \langle C, \Pi \rangle + \gamma \sum_{j,i} \Pi_{ji} \left( \log \frac{\Pi_{ji}}{R_{ji}} - 1 \right) \right\}. \tag{28}$$

If $R'_{ji} = r_j R_{ji} c_i$, then

$$T_{\gamma,R'}(q,p) = T_{\gamma,R}(q,p) - \gamma \sum_j q_j \log r_j - \gamma \sum_i p_i \log c_i. \tag{29}$$

The first added term is affine in the varying marginal and the second is constant. Both disappear from the anchored Bregman remainder. Adding a cost $a_j + b_i$ is equivalent to multiplying $R$ by $e^{-a_j/\gamma} e^{-b_i/\gamma}$ and gives the same conclusion.

The semi-dual for (28) is obtained by replacing each exponential in (16) by $R_{ji}$ times that exponential. The proof of Theorem 2.3 is unchanged, so the exact plan-KL identity holds for every fixed positive $R$. The high-temperature KL limit in Theorem 4.1, however, uses the counting/separable-reference structure and should not be transferred verbatim to an arbitrary nonseparable $R$.

## C Proof of the robust dual

With the normalized gauge $p^\mathsf{T}\nu = T_\gamma(p,p)$ and $\bar{C}_{ji} = C_{ji} - \nu_j$, the constraint $D_\gamma^{\mathrm{TC}}(q\|p) \le \rho$ is equivalent to the existence of $\Pi^\mathsf{T}\mathbf{1} = p$ such that

$$\langle \bar{C}, \Pi \rangle + \gamma \sum_{j,i} \Pi_{ji}(\log \Pi_{ji} - 1) \le \rho. \tag{30}$$

Indeed, minimizing the left side over plans with first marginal $q$ gives $T_\gamma(q,p) - \nu^\mathsf{T} q$, while the gauge makes the right-side offset $T_\gamma(p,p) - \nu^\mathsf{T} p$ equal to zero.

Attach $\lambda \ge 0$ to (30) and $\beta_i$ to each column constraint. The Lagrangian for the maximization problem is

$$\begin{aligned}\mathcal{L}(\Pi, \lambda, \beta) = \lambda\rho + p^\mathsf{T}\beta \\ + \sum_{j,i} \left[\Pi_{ji}(L_j - \beta_i - \lambda \bar{C}_{ji}) - \lambda\gamma \Pi_{ji}(\log \Pi_{ji} - 1)\right].\end{aligned}$$

For $\lambda > 0$,

$$\sup_{x \ge 0} \{ax - \lambda\gamma x(\log x - 1)\} = \lambda\gamma e^{a/(\lambda\gamma)}.$$

Therefore the dual before eliminating $\beta$ is

$$\inf_{\lambda \ge 0,\, \beta \in \mathbb{R}^m} \left\{ \lambda\rho + p^\mathsf{T}\beta + \lambda\gamma \sum_{j,i} e^{(L_j - \beta_i - \lambda \bar{C}_{ji})/(\lambda\gamma)} \right\}, \tag{31}$$

with the lower-semicontinuous boundary interpretation at $\lambda = 0$. The nominal self-plan is strictly positive and has budget zero, so $\rho > 0$ gives Slater slack. Strong duality follows.

For fixed $(\theta, \lambda)$, set $S_i = \sum_j e^{(L_j - \lambda \bar{C}_{ji})/(\lambda\gamma)}$. The terms involving $\beta_i$ are $p_i\beta_i + \lambda\gamma S_i e^{-\beta_i/(\lambda\gamma)}$. Their minimizer is $\beta_i^* = \lambda\gamma \log(S_i/p_i)$ and their minimum is $\lambda\gamma p_i[1 + \log(S_i/p_i)]$. Summing proves (11).

At a positive optimal multiplier, maximizing each scalar Lagrangian term gives $\Pi^*_{ji} = e^{(L_j - \beta_i^* - \lambda^* \bar{C}_{ji})/(\lambda^*\gamma)}$. Substituting $\beta_i^*$ proves (12); strict concavity in $\Pi$ gives uniqueness. The exponential perspective in (31) is jointly convex. Composing its increasing argument with convex $L_j(\theta)$ and partially minimizing over $\beta$ preserves convexity in $(\theta, \lambda)$. The gradient statement follows from Danskin's theorem; in the nonunique case, under convex losses, every maximizing distribution gives a convex subgradient.

For the fixed reference $R$ in (28), the same derivation inserts $R_{ji}$ in every exponential and in the numerator and denominator of (12).

# D PROOFS OF THE TEMPERATURE LIMITS

## D.1 LOW TEMPERATURE

Every coupling has total mass one on $m^2$ entries, so $-2\log m \le \sum_{j,i} \Pi_{ji} \log \Pi_{ji} \le 0$. Hence

$$W_C(q,p) - \gamma(1 + 2\log m) \le T_\gamma(q,p) \le W_C(q,p) - \gamma. \tag{32}$$

**Lemma D.1** (Bounded self-potential)**.** *Under the low-temperature assumptions of Theorem 4.1, the self-plan has a symmetric representation*

$$\Pi^{p,p}_{\gamma,ji} = \exp((u_{\gamma,j} + u_{\gamma,i} - C_{ji})/\gamma)$$

*with*

$$\gamma(\log p_j - \tfrac{1}{2}\log m) \le u_{\gamma,j} \le \tfrac{\gamma}{2}\log p_j, \qquad \|u_\gamma\|_\infty \le \gamma \log(\sqrt{m}/p_{\min}). \tag{33}$$

*Proof.* Symmetry and uniqueness of the self-plan allow equal row and column potentials. Put $a_j = e^{u_{\gamma,j}/\gamma}$. Then $\Pi_{ji} = a_j a_i e^{-C_{ji}/\gamma}$ and $p_j = a_j \sum_i a_i e^{-C_{ji}/\gamma}$. The diagonal term gives $p_j \ge a_j^2$, so $a_j \le \sqrt{p_j}$. Conversely,

$$p_j \le a_j \sum_i a_i \le a_j \sum_i \sqrt{p_i} \le a_j \sqrt{m}.$$

Thus $p_j/\sqrt{m} \le a_j \le \sqrt{p_j}$. Taking logarithms gives the result. □

The row self-potential is a valid representative of $\nu_{\gamma,p}$. Combining (32), $W_C(p,p) = 0$, (33), and $\|q-p\|_1 \le 2$ gives

$$\begin{aligned}\left|D^{\mathrm{TC}}_\gamma(q\|p) - W_C(q,p)\right| &\le |T_\gamma(q,p) - W_C(q,p)| + |T_\gamma(p,p)| + \left|u_\gamma^\mathsf{T}(q-p)\right| \\ &\le 2\gamma[1 + 2\log m + \log(\sqrt{m}/p_{\min})],\end{aligned}$$

which proves (13).

For completeness, let $V_W(r;L) = \sup_{W_C(q,p)\le r} q^\mathsf{T} L$. Uniform closeness of the two discrepancies gives

$$\{W_C \le \rho - \varepsilon_\gamma\} \subseteq \mathcal{U}^{\mathrm{TC}}_{\rho,\gamma}(p) \subseteq \{W_C \le \rho + \varepsilon_\gamma\}.$$

If $\|L\|_\infty \le B$, contract any $q$ feasible at radius $\rho + s$ toward $p$ by the factor $\rho/(\rho+s)$. Convexity of $W_C(\cdot, p)$ makes the contraction feasible at radius $\rho$, while its objective changes by at most $2Bs/\rho$. The analogous contraction from $\rho$ to $\rho - s$ gives $|V_W(\rho \pm s; L) - V_W(\rho; L)| \le 2Bs/\rho$ for $s < \rho$. Taking $s = \varepsilon_\gamma$ proves (15).

## D.2 HIGH TEMPERATURE

Let $h(q) = \sum_j q_j \log q_j$. For any $\Pi \in \Gamma(q,p)$,

$$\sum_{j,i} \Pi_{ji}(\log \Pi_{ji} - 1) = h(q) + h(p) - 1 + \mathrm{KL}(\Pi \| qp^\mathsf{T}). \tag{34}$$

Therefore

$$T_\gamma(q,p) = \gamma[h(q) + h(p) - 1] + R_\gamma(q,p), \tag{35}$$

where

$$R_\gamma(q,p) = \min_{\Pi \in \Gamma(q,p)} \{\langle C, \Pi\rangle + \gamma \mathrm{KL}(\Pi \| qp^\mathsf{T})\}.$$

Writing $M_C = \max_{j,i} C_{ji}$, the independent coupling gives the upper bound $R_\gamma(q,p) \le q^\mathsf{T} C p$. For any feasible $\Pi$, Pinsker gives

$$\langle C, \Pi\rangle + \gamma \mathrm{KL}(\Pi \| qp^\mathsf{T}) \ge q^\mathsf{T} C p - M_C t + \tfrac{\gamma}{2} t^2, \quad t = \left\|\Pi - qp^\mathsf{T}\right\|_1.$$

Minimizing the right side over $t \ge 0$ yields

$$q^\mathsf{T} C p - \frac{M_C^2}{2\gamma} \le R_\gamma(q,p) \le q^\mathsf{T} C p \tag{36}$$

uniformly on the closed simplex.

Set $f_\gamma(q) = \gamma^{-1} T_\gamma(q, p)$. Equations (35)–(36) give

$$f_\gamma(q) = h(q) + h(p) - 1 + \gamma^{-1} q^\mathsf{T} C p + O(\gamma^{-2}) \tag{37}$$

uniformly in $q$. In particular, $f_\gamma \to f_\infty := h(\cdot) + h(p) - 1$ uniformly. We next justify convergence of normalized relative gradients at the fixed interior point $p$. For $v \in \mathcal{T}$, choose $t > 0$ with $p \pm tv \in \Delta_m^\circ$. Convexity gives

$$\frac{f_\gamma(p) - f_\gamma(p - tv)}{t} \le \nabla f_\gamma(p)^\mathsf{T} v \le \frac{f_\gamma(p + tv) - f_\gamma(p)}{t}.$$

First let $\gamma \to \infty$ using uniform convergence, then let $t \downarrow 0$. The two bounds meet at $\nabla h(p)^\mathsf{T} v$. Finite dimensionality implies convergence of the normalized relative gradients on $\mathcal{T}$. Hence

$$\sup_{q \in \Delta_m} \Big| [f_\gamma(q) - f_\gamma(p) - \nabla f_\gamma(p)^\mathsf{T}(q - p)] \\ - [h(q) - h(p) - \nabla h(p)^\mathsf{T}(q - p)] \Big| \longrightarrow 0.$$

The second bracket is $\mathrm{KL}(q\|p)$, including at boundary points by continuity. This proves (14).

Let $g_\gamma = D_\gamma^{\mathrm{TC}}/\gamma$ and $g = \mathrm{KL}(\cdot\|p)$. Uniform convergence on compact $\Delta_m$ implies that any limit point of $\{g_\gamma \le \tau\}$ lies in $\{g \le \tau\}$. Conversely, if $g(q) < \tau$, then $q$ is eventually feasible. If $g(q) = \tau > 0$, use $q_t = (1 - t)q + tp$: strict convexity along this segment gives $g(q_t) < \tau$, and one may choose $t = t_\gamma \downarrow 0$ slowly enough that $q_{t_\gamma}$ is feasible for $g_\gamma$. Compactness turns these two facts into Hausdorff convergence. Uniform convergence of bounded linear support functions follows by the same compactness argument.

# E Proofs for local geometry and calibration

## E.1 Proof of Theorem 5.1

By (24), $\nabla^2 A_{\gamma,p}(\nu) = \gamma^{-1} M_{\gamma,p}$. For nonzero $v \in \mathcal{T}$, every $s_i(\nu)$ has full support and

$$v^\mathsf{T} M_{\gamma,p} v = \sum_i p_i \operatorname{Var}_{s_i(\nu)}(v) > 0.$$

Since $T_\gamma(\cdot, p)$ is $A^*_{\gamma,p}$ up to a constant, their intrinsic Hessians are inverses, proving $H_{\gamma,p} = \gamma M^\dagger_{\gamma,p}$ and the Taylor expansion.

There are two equivalent proofs of (19). Differentiating the global inequality $D_\gamma^{\mathrm{TC}}(q\|p) \ge \gamma \mathrm{KL}(q\|p)$ twice at $q = p$ gives it immediately. Alternatively, under the self-plan let $J$ be the row state and $I$ the column state. Both marginals are $p$, and $M_{\gamma,p} = \mathbb{E}[\mathrm{Cov}(e_J \mid I)]$. The law of total covariance gives

$$M_{\gamma,p} \preceq \mathrm{Cov}(e_J) = \mathrm{Diag}(p) - pp^\mathsf{T}.$$

Inversion on $\mathcal{T}$ reverses the Loewner order, yielding $v^\mathsf{T} H_{\gamma,p} v \ge \gamma \sum_j v_j^2/p_j$.

## E.2 An explicit curvature bound

**Lemma E.1** (Uniform interior curvature). *Let $\Delta_m(\epsilon) = \{q \in \Delta_m : q_j \ge \epsilon\}$ and $M_C = \max C_{ji}$. For $q, p \in \Delta_m(\epsilon)$,*

$$\left\|\nabla^2_{11} T_\gamma(q, p)\right\|_{\mathrm{op},\mathcal{T}} \le \frac{\gamma e^{2M_C/\gamma}}{\epsilon}. \tag{38}$$

*Proof.* Let $K_{ji} = e^{-C_{ji}/\gamma}$ and $\kappa = e^{-M_C/\gamma}$, so $\kappa \le K_{ji} \le 1$. Write the positive optimal plan as $\Pi = \mathrm{Diag}(a) K \mathrm{Diag}(b)$ and set $A = \sum_j a_j$, $B = \sum_i b_i$. Total mass gives $1 = \sum_{j,i} a_j K_{ji} b_i \ge \kappa AB$, hence $AB \le \kappa^{-1}$. The row constraint gives $q_j = a_j \sum_i K_{ji} b_i \le a_j B$, so $a_j \ge \epsilon/B$. Therefore every conditional column probability satisfies

$$s_{ji} = \frac{a_j K_{ji}}{\sum_k a_k K_{ki}} \ge \frac{(\epsilon/B)\kappa}{A} \ge \epsilon \kappa^2.$$

For $v \in \mathcal{T}$,

$$\mathrm{Var}_{s_i}(v) = \min_c \sum_j s_{ji}(v_j - c)^2 \geq \epsilon\kappa^2 \min_c \sum_j (v_j - c)^2 = \epsilon\kappa^2 \|v\|_2^2 .$$

Thus the log-partition covariance is bounded below by $\epsilon\kappa^2 I$ on $\mathcal{T}$. Inverting it in the Hessian formula proves the claim. □

### E.3 Proof of the finite-sample statement in Theorem 5.2

Let $\widehat{p}_n = N/n$. A coordinatewise Hoeffding bound and union bound give, with probability at least $1 - \delta$,

$$\|\widehat{p}_n - p^\star\|_2 \leq \sqrt{\frac{m}{2n} \log \frac{2m}{\delta}}. \tag{39}$$

The smoothed estimator obeys

$$\widetilde{p}_n - p^\star = \frac{n}{n + ma}(\widehat{p}_n - p^\star) + \frac{a}{n + ma}(\mathbf{1} - mp^\star).$$

The maximum of $\|\mathbf{1} - mp\|_2$ over the simplex is attained at a vertex and equals $\sqrt{m(m-1)}$. Thus (39) implies $\|\widetilde{p}_n - p^\star\|_2 \leq r_{n,\delta,a}$. If $r_{n,\delta,a} \leq \underline{p}/2$, the entire segment between $p^\star$ and $\widetilde{p}_n$ belongs to $\Delta_m(\underline{p}/2)$.

For fixed second marginal $\widetilde{p}_n$, the discrepancy is the Taylor remainder of $q \mapsto T_\gamma(q, \widetilde{p}_n)$ about $q = \widetilde{p}_n$. By Theorem E.1,

$$D_\gamma^{\mathrm{TC}}(p^\star \| \widetilde{p}_n) \leq \frac{1}{2} \frac{\gamma e^{2M_C/\gamma}}{\underline{p}/2} \|p^\star - \widetilde{p}_n\|_2^2 \leq \frac{\gamma e^{2M_C/\gamma}}{\underline{p}} r_{n,\delta,a}^2.$$

This proves (21).

### E.4 Proof of the generalized chi-square limit

The add-$a$ bias is $O(n^{-1})$, so the multinomial central limit theorem gives

$$\sqrt{n}(\widetilde{p}_n - p^\star) \Rightarrow Z \sim \mathcal{N}(0, \Sigma^\star).$$

For each realized interior nominal $r$, $D_\gamma^{\mathrm{TC}}(q\|r)$ is the Bregman remainder in $q$ about $r$. Entropic OT is $C^\infty$ on a common interior neighborhood of $p^\star$, so Taylor expansion uniformly for $r \to p^\star$ gives

$$D_\gamma^{\mathrm{TC}}(p^\star \| \widetilde{p}_n) = \frac{1}{2}(p^\star - \widetilde{p}_n)^\mathsf{T} H_{\gamma,\widetilde{p}_n}(p^\star - \widetilde{p}_n) + o_p(n^{-1}). \tag{40}$$

Continuity of $H_{\gamma,p}$ and Slutsky's theorem prove (22).

Let $\lambda_1, \ldots, \lambda_{m-1}$ be the positive eigenvalues of $(\Sigma^\star)^{1/2} H_{\gamma,p^\star} (\Sigma^\star)^{1/2}$. Diagonalizing this matrix gives $Z^\mathsf{T} H Z \overset{d}{=} \sum_{r=1}^{m-1} \lambda_r \chi_{1,r}^2$. The distribution has a continuous positive density away from zero, so its positive quantiles are continuity points. The plug-in matrices converge in probability, their nonzero eigenvalues converge, and so do the associated quantiles. If the quantile is estimated from $B_n$ conditional Gaussian draws, conditional Glivenko–Cantelli (or the Dvoretzky–Kiefer–Wolfowitz inequality) makes the empirical quantile error vanish provided $B_n \to \infty$. Combining this fact with (40) proves plug-in coverage.

Finally, as $\gamma \to \infty$, the self-plan converges to $pp^\mathsf{T}$, so every conditional vector $s_i(\nu_{\gamma,p})$ converges to $p$. Hence $M_{\gamma,p} \to \mathrm{Diag}(p) - pp^\mathsf{T}$ and

$$\gamma^{-1} H_{\gamma,p} \to (\mathrm{Diag}(p) - pp^\mathsf{T})^\dagger \quad \text{on } \mathcal{T}. \tag{41}$$

The resulting normalized quadratic form is $\chi_{m-1}^2$. When $C = 0$, the optimal plan is $qp^\mathsf{T}$, so $D_\gamma^{\mathrm{TC}} = \gamma \mathrm{KL}(q\|p)$ and (41) is exact for every $\gamma$.

# F Relation to earlier EOT Bregman orientations

For completeness, we verify (9) and make the orientation distinction algebraic. Let

$$h_{\gamma,p}(q) = \min_{\Pi \in \Gamma(q,p)} \left\{ \langle C, \Pi \rangle + \gamma \mathrm{KL}(\Pi \| qp^{\mathsf{T}}) \right\}.$$

For any $q, p \in \Delta_m^\circ$ and $\Pi \in \Gamma(q,p)$,

$$\mathrm{KL}(\Pi \| qp^{\mathsf{T}}) = \sum_{j,i} \Pi_{ji} \log \Pi_{ji} - \sum_j q_j \log q_j - \sum_i p_i \log p_i.$$

Consequently,

$$T_\gamma(q,p) = h_{\gamma,p}(q) + \gamma \sum_j q_j \log q_j + \gamma \sum_i p_i \log p_i - \gamma. \tag{42}$$

Taking the forward Bregman remainder of both sides in the first argument, anchored at $p$, proves

$$D_\gamma^{\mathrm{TC}}(q\|p) = B_{h_{\gamma,p}}(q,p) + \gamma \mathrm{KL}(q\|p).$$

Combining this equality with (6) further yields

$$B_{h_{\gamma,p}}(q,p) = \gamma \sum_{j: q_j > 0} q_j \mathrm{KL}\big(\Pi_\gamma^{q,p}(\cdot \mid j) \big\| \Pi_\gamma^{p,p}(\cdot \mid j)\big).$$

The boundary extension follows by the same interior approximation used in the proof of Theorem 2.3.

The product-reference Bregman construction appears in Kerimkulov et al. (2025). Amari et al. (2018, Eqs. 49–50) explicitly introduce both KL orientations between referenced optimal plans, call the second the dual candidate, and then analyze the self-plan-first orientation obtained from their Eq. 49. For symmetric $C$, setting their reference to $p$, transposing the plans to match our row-$q$/column-$p$ convention, and matching regularization scales turns their Eq. 50 exactly into (5). Thus the forward plan-KL candidate is prior art. The results here identify it with the unique affine tangent correction of fixed/counting-reference EOT and derive its conditional/marginal decomposition, DRO reformulation, limits, and calibration. KL asymmetry still matters: $\gamma^{-1} D_\gamma^{\mathrm{TC}}(q\|p) \to \mathrm{KL}(q\|p)$ and, for symmetric zero-diagonal costs, $D_\gamma^{\mathrm{TC}}(q\|p) \to W_C(q,p)$ without rescaling as $\gamma \downarrow 0$.

# G Relation to symmetric Sinkhorn debiasing

When $C$ is symmetric, define

$$S_\gamma(q,p) = T_\gamma(q,p) - \tfrac{1}{2} T_\gamma(q,q) - \tfrac{1}{2} T_\gamma(p,p), \qquad G_\gamma(q) = -\tfrac{1}{2} T_\gamma(q,q).$$

Symmetry gives equal first- and second-marginal derivatives on the diagonal, hence $\nabla G_\gamma(p) = -\nu_{\gamma,p}$ in the relative tangent space. Direct algebra yields

$$D_\gamma^{\mathrm{TC}}(q\|p) = S_\gamma(q,p) - \left[ G_\gamma(q) - G_\gamma(p) - \nabla G_\gamma(p)^{\mathsf{T}}(q-p) \right]. \tag{43}$$

Thus symmetric self-cost subtraction and one-sided tangent subtraction are not interchangeable. The former targets symmetric comparison and an OT-to-kernel limit under additional assumptions; the latter preserves convexity inherited from $T_\gamma(\cdot, p)$ and produces the plan-KL and marginal-KL geometry proved above.

# H Implementation and reproduction details

## H.1 Stable computation

For production use, the self-plan can be computed by alternating scaling

$$a \leftarrow p/(Kb), \qquad b \leftarrow p/(K^{\mathsf{T}} a), \qquad K_{ji} = e^{-C_{ji}/\gamma},$$

with log-domain stabilization when the dynamic range of $C/\gamma$ is large. A row potential is $\nu = \gamma \log a$; it is shifted by a constant until $p^{\mathsf{T}} \nu = T_\gamma(p,p)$. For fixed $(\theta, \lambda)$ in (11), set

$$x_{ji} = \frac{L_j(\theta) - \lambda \bar{C}_{ji}}{\lambda \gamma}, \qquad m_i = \max_j x_{ji}.$$

Then compute $\log \sum_j e^{x_{ji}} = m_i + \log \sum_j e^{x_{ji}-m_i}$ and normalize the same shifted exponentials to obtain (12).

For reference, differentiating the scalar dual objective $\Psi(\theta, \lambda)$ gives

$$\partial_\lambda \Psi = \rho + \gamma + \gamma \sum_i p_i \log \left[ \frac{1}{p_i} \sum_j e^{x_{ji}} \right] - \frac{1}{\lambda} \sum_i p_i \sum_j \omega_{ji} L_j(\theta), \tag{44}$$

where $\omega_{ji} = e^{x_{ji}} / \sum_k e^{x_{ki}}$. The corrected cost cancels from the last expectation after differentiating the scaled exponent. Searching over $\log \lambda$ enforces positivity.

The local covariance can be read directly from the self-plan:

$$s_i(\nu) = \Pi^{p,p}_{\gamma,\cdot i}/p_i, \qquad M_{\gamma,p} = \sum_i p_i [\mathrm{Diag}(s_i) - s_i s_i^{\mathsf{T}}].$$

Thus no differentiation through the Sinkhorn iterations is required.

### H.2 Controlled simulations

The supplied script `code/reproduce.py` uses NumPy for all transport, linear algebra, random generation, and quantiles, and Pillow to draw the figure. It performs the following checks.

1. It evaluates the raw-center formula (3) on a logarithmic temperature grid.
2. It computes $D^{\mathrm{TC}}_\gamma$ from the plan-KL identity and computes the one-dimensional unregularized OT cost by the monotone greedy coupling. This produces Figure 1.
3. It minimizes the scalar objective (11), recovers the conditional worst-case plan from (12), and reports its budget residual and primal–dual objective gap.
4. For coverage, it draws one fixed $5000 \times 5$ bank of squared standard Gaussians. For each of 2,000 fresh multinomial samples per $(n, \gamma)$, it recomputes the add-$1/2$ nominal, self-plan, plug-in Hessian, covariance eigenvalues, conditional $0.95$ quantile, and containment event. Reusing the Gaussian bank is a common-random-number device; the experiment assesses a finite-$B$ approximation, while the theorem requires $B_n \to \infty$ for exact asymptotics.
5. Alternating-scaling marginal residual tolerance is $2 \times 10^{-13}$, checked every 20 updates, with a maximum of 10,000 iterations. For nearly diagonal symmetric self-plans in the low-temperature figure, the script instead solves the symmetric scaling equations by damped Newton iteration at the same tolerance, with a 1,000-iteration cap. The largest exponent range used is $\max C/\gamma = 80$, so direct exponentiation is safe in these experiments; production code should use the log-domain scheme above for more extreme inputs.

## I Statement on the use of language models

An OpenAI language model assisted with language editing, formatting, manuscript organization, mathematical consistency checks, reference verification, and development of the reproducibility code. The authors independently verified all mathematical arguments, citations, and numerical results and take full responsibility for the manuscript.