SEMIPARAMETRIC EFFICIENCY FOR ENTROPIC OPTIMAL TRANSPORT UNDER STRUCTURED DISTRIBUTION SHIFT

Anonymous authors
Paper under double-blind review

ABSTRACT

Entropic optimal transport (EOT) compares distributions and learned representations, but few target observations can make empirical discrepancy estimates noisy. Estimating the two marginals separately can also waste information when a credible relationship links the domains. We study fixed-regularization EOT under exponential density-ratio shifts with a nonparametric baseline. We show that the variance gain equals the weighted squared residual left after projecting the unrestricted transport influence contrast onto the tilt space. An empirical-likelihood estimator, DRM-EOT, attains this bound on compact support, with consistent analytic variance estimation and a stratified refit bootstrap. The same residual governs first-order sensitivity to local misspecification. Under exact semi-synthetic shifts on real OfficeHome representations, fixed-proportion scaling connects empirical variance to the oracle prediction: at $n = 16000$ and ESS ratio 0.5, the variance ratio is 0.950, compared with the oracle value of 0.952. Under natural shifts, a low-dimensional DRM improves target-scarce RMSE mainly by reducing bias rather than variance. Increasing model complexity reveals finite-sample separation and persistent off-model bias. These diagnostics distinguish exact-model efficiency from structural regularization. Our results show when structural information sharing can improve EOT estimation, when it cannot, and how that gain trades off against model misspecification.

1 INTRODUCTION

Few target observations can make source–target discrepancy estimates noisy, even when source data are abundant. This noise complicates comparisons of learned representations across domains and the diagnosis of distribution shift. In our asymptotic theory, target scarcity is represented by an unbalanced proportional design with a small but strictly positive limiting target fraction; regimes in which $n_1/n \to 0$ require a different asymptotic analysis. Entropic optimal transport (EOT) provides a geometry-aware comparison, but standard empirical EOT estimates the two marginals separately. As a result, it treats the domains as unrelated even when a scientifically plausible relationship permits information sharing. How much precision is lost by treating structurally related domains as unrelated?

We study this question under a density ratio model (DRM), $dF_1/dF_0 = \exp\{\theta^\top Q\}$, with a fixed low-dimensional basis Q. This relationship constrains the shift, while the baseline distribution remains nonparametric. We study the statistical efficiency of a fixed EOT functional rather than introduce a new transport discrepancy, representation learner, or domain-adaptation system.

An exact projection identity answers this question. Let ϕ_0, ϕ_1 be centered transport potentials, ρ_g the sampling fractions, $\omega = \rho_0 + \rho_1(dF_1/dF_0)$, and e the residual after projecting $\zeta = \phi_1 - (\rho_1/\rho_0)\phi_0$ onto $\mathrm{span}(Q)$ in the weighted squared norm $E_1(\cdot^2/\omega)$. Then

$$\boxed{V_{\mathrm{NP}} - V_{\mathrm{DRM}} = \frac{\rho_0}{\rho_1} E_1(e^2/\omega).} \tag{1}$$

The valid restriction removes an EOT-relevant direction precisely when $e \neq 0$ in $L^2(F_1)$, and only then does it yield a strict first-order gain. Correct specification alone does not imply improvement.

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.

Moreover, the same residual determines sensitivity to local departures from the DRM: information sharing comes with an explicit efficiency–robustness trade-off.

Our contributions are fourfold. (1) We derive the DRM canonical gradient—the efficient first-order sensitivity to local changes—and an exact efficiency identity that characterizes both strict gain and zero gain. (2) We establish regularity and semiparametric efficiency for an empirical-likelihood plug-in method, *DRM-EOT*, together with consistent analytic variance estimation and a valid stratified refit bootstrap. (3) We identify the residual that controls sensitivity to local misspecification, connecting valid structural restrictions to off-model vulnerability. (4) We provide evidence from frozen OfficeHome neural representations: exact semi-synthetic shifts exhibit oracle variance gains and empirical-to-oracle scaling, while natural shifts show target-scarce bias reduction and complexity-dependent instability. At $n = 16000$, the exact-DRM variance ratio of 0.950 closely matches the oracle value of 0.952; natural shifts must instead be interpreted as potentially misspecified structural regularization.

2 Related Work

Regularized OT and its statistical properties. Entropic regularization enables transport computation by matrix scaling (Cuturi, 2013). Sinkhorn divergences debias regularized transport and connect geometric and kernel-based distribution comparisons (Feydy et al., 2019). Genevay et al. (2019) study their sample complexity; Mena & Niles-Weed (2019) establish statistical bounds and a central limit theorem centered at the expected empirical EOT under sub-Gaussian marginals. Goldfeld et al. (2024) develop differentiability, efficiency, and bootstrap theory, while stability of the transport potentials supports plug-in analysis (Nutz & Wiesel, 2023). Existing statistical analyses characterize regularized OT under unrestricted marginals. We quantify and attain the additional efficiency available when a semiparametric density-ratio restriction links the marginals. We study the product-marginal EOT functional because its first-order pathwise derivative leads directly to the projection geometry here. We do not prove inference for its debiased counterpart. Extending the analysis to the Sinkhorn divergence requires separate treatment of its first-order degeneracy at identical marginals.

Transport and domain shift. Transport has been used to align representations across domains (Courty et al., 2017b) and jointly learn a target predictor and a coupling (Courty et al., 2017a). Such methods use transport as an alignment objective. Here representations are frozen, and we ask how efficiently their fixed EOT discrepancy can be estimated. We use OfficeHome (Venkateswara et al., 2017) to obtain realistic representation geometry, without claiming state-of-the-art domain-adaptation accuracy.

Density ratios and structured information sharing. Direct density-ratio estimation supports importance weighting under covariate shift (Sugiyama et al., 2007). DRM empirical likelihood instead shares information while retaining a nonparametric baseline (Qin, 1998); related functional inference includes quantiles and dominance indices (Chen & Liu, 2013; Zhuang et al., 2019), with bootstrap theory developed by Zhuang et al. (forthcoming). We target neither reweighted classification risk nor density-ratio prediction itself. We derive the attainable semiparametric efficiency bound for a nonlinear, two-marginal transport functional.

Recent work also combines regularized OT with semiparametric inference in causal settings. Agarwal & Luedtke (2026) introduce a Sinkhorn-based discrepancy between counterfactual outcome distributions and develop debiased inference for this causal functional. Burman et al. (2026) use entropic OT and calibration to recover a density ratio for transporting randomized treatment effects, with semiparametric efficiency for the resulting causal estimator. These works embed OT in causal estimands or weighting procedures. We instead study the fixed-regularization EOT functional of two observed marginals linked by a finite-dimensional DRM. We characterize how this restriction changes the EOT efficiency bound, derive the exact restricted-versus-unrestricted projection identity, construct an estimator that attains the bound, and relate the same projection residual to local misspecification sensitivity.

3 STRUCTURED EOT ESTIMATION

3.1 DENSITY-RATIO SHIFT AND TRANSPORT TARGET

Observe independent source samples $X_1, \dots, X_{n_0} \sim F_0$ and target samples $Y_1, \dots, Y_{n_1} \sim F_1$. Group sizes are nonrandom, $n = n_0 + n_1$, and $\rho_{gn} = n_g/n \to \rho_g \in (0,1)$, without a convergence-rate assumption. Write $E_g f = \int f \, dF_g$, and let $q : K \to \mathbb{R}^p$ be fixed, with $Q = (1, q^\top)^\top \in \mathbb{R}^{p+1}$. Then

$$dF_1 = r \, dF_0, \qquad r = \exp\{\theta_0^\top Q\}, \qquad E_0 r = 1. \tag{2}$$

For fixed $\varepsilon > 0$ and $c(x,y) = \|x - y\|^2/2$, define

$$\Psi_\varepsilon(F_0, F_1) = \inf_{\pi \in \Pi(F_0, F_1)} \left\{ \int c \, d\pi + \varepsilon \, \mathrm{KL}(\pi \| F_0 \otimes F_1) \right\}. \tag{3}$$

This product-marginal entropy convention matters: the target need not vanish for identical marginals and is not the Sinkhorn divergence. NP-EOT evaluates (3) at the two raw empirical marginals P_{0n}, P_{1n}.

Assumption 1. *The baseline support $K \subset \mathbb{R}^d$ is compact; q is continuous on K with fixed dimension; θ_0 is interior; and $M = E_1(QQ^\top/\omega)$ is positive definite, where $\omega = \rho_0 + \rho_1 r$.*

The positive tilt is bounded above and away from zero, so both distributions have common compact support. Appendix A records the EOT regularity and empirical-process interface, using Goldfeld et al. (2024, Appendix E, Lemmas E.22–E.24) and their delta-method results. All theory conditions on a fixed representation and basis.

3.2 DRM-EOT ESTIMATOR

Pool observations as $Z_1, \dots, Z_n$, retaining multiplicities. For $\theta \in \mathbb{R}^{p+1}$, set $r_\theta = \exp(\theta^\top Q)$ and $\omega_{\theta,n} = \rho_{0n} + \rho_{1n} r_\theta$. Fit

$$\widehat{\theta} = \arg\max_\theta \left\{ \sum_{j=1}^{n_1} \theta^\top Q(Y_j) - \sum_{i=1}^{n} \log \omega_{\theta,n}(Z_i) \right\}. \tag{4}$$

Given an interior solution, define

$$\widehat{F}_0 = \sum_i \frac{\delta_{Z_i}}{n \omega_{\widehat{\theta},n}(Z_i)}, \qquad \widehat{F}_1 = \sum_i \frac{r_{\widehat{\theta}}(Z_i) \delta_{Z_i}}{n \omega_{\widehat{\theta},n}(Z_i)}, \tag{5}$$

$$\widehat{\Psi}_\varepsilon^{\mathrm{DRM}} = \Psi_\varepsilon(\widehat{F}_0, \widehat{F}_1).$$

The intercept score normalizes both measures. The DRM profile likelihood is closely related to discriminative logistic estimation with a sampling-fraction offset. Its empirical-likelihood formulation also yields normalized pooled fitted marginals for the EOT solver. Pooling the supports allows information to flow across groups through the tilt. Under the assumptions, an interior maximizer exists with probability tending to one. On the exceptional event that it does not exist, define DRM-EOT by NP-EOT without changing the limits (Appendix D). Numerical failure rates are reported separately in the experiments.

Algorithm 1: DRM-EOT
Input: source X, target Y, fixed basis Q, regularization ε.
1. Fit the exponential tilt by empirical likelihood, (4).
2. Construct pooled fitted marginals, (5).
3. Compute entropic OT between the fitted marginals.
4. Optionally estimate analytic standard errors or run a stratified refit bootstrap, refitting the tilt and transport in every replicate. Bootstrap refits use the same NP-EOT fallback convention when the DRM fit fails.

DRM-EOT adds a low-dimensional density-ratio fit and applies the same Sinkhorn algorithm to the pooled fitted marginals. Its transport matrix can be larger: pooling gives up to n support points per

marginal, compared with n_0 and n_1 for NP-EOT. Dense storage and each Sinkhorn iteration therefore cost $O(n^2)$ for DRM-EOT versus $O(n_0 n_1)$ for NP-EOT, before duplicate-support compression. The additional tilt fit is low-dimensional. DRM-EOT targets statistical rather than computational efficiency; representative timings and process-memory measurements for the existing compressed implementation are reported in Appendix H. We retain the exact sample fractions throughout, avoiding deterministic errors that may not be negligible at root-n scale. Numerical objective errors must be $o_p(n^{-1/2})$ for the inferential limits, with analogous conditional accuracy in bootstrap replicates.

4 Efficiency Geometry

4.1 Unrestricted influence function and projection

Let (u, v) be optimal dual potentials and $\phi_0 = u - E_0 u$, $\phi_1 = v - E_1 v$. Their separate centering represents the derivative on zero-mass perturbations:

$$\dot{\Psi}_\varepsilon(s_0, s_1) = E_0(\phi_0 s_0) + E_1(\phi_1 s_1). \tag{6}$$

In $\mathcal{H} = L_0^2(F_0) \oplus L_0^2(F_1)$ with norm squared $\rho_0 E_0 a_0^2 + \rho_1 E_1 a_1^2$, the unrestricted gradient is $(\phi_0/\rho_0, \phi_1/\rho_1)$ and

$$V_{\mathrm{NP}} = E_0\phi_0^2/\rho_0 + E_1\phi_1^2/\rho_1. \tag{7}$$

Project the joint potential contrast, not either potential separately:

$$\zeta = \phi_1 - \frac{\rho_1}{\rho_0}\phi_0, \quad b^\star = M^{-1}E_1(Q\zeta/\omega), \quad e = \zeta - b^{\star\top}Q. \tag{8}$$

Here $b^\star$ minimizes $E_1\{(\zeta - b^\top Q)^2/\omega\}$, so $E_1(Qe/\omega) = 0$.

4.2 Canonical gradient and exact efficiency identity

Theorem 1 (Canonical gradient and variance reduction)**.** *Under Assumption 1, the DRM canonical gradient is*

$$g_0 = \frac{\phi_0}{\rho_0} + \frac{re}{\omega}, \qquad g_1 = \frac{\phi_1}{\rho_1} - \frac{\rho_0}{\rho_1}\frac{e}{\omega}. \tag{9}$$

Its efficiency bound is

$$\boxed{V_{\mathrm{DRM}} := \rho_0 E_0 g_0^2 + \rho_1 E_1 g_1^2 = V_{\mathrm{NP}} - \frac{\rho_0}{\rho_1}E_1(e^2/\omega).} \tag{10}$$

Thus $V_{\mathrm{DRM}} < V_{\mathrm{NP}}$ if and only if $e \neq 0$ in $L^2(F_1)$; equality holds if and only if $\zeta \in \mathrm{span}(Q)$ F_1-almost surely.

Proof sketch. The tangent space comprises the allowed local distribution changes. Normalized DRM paths have scores $(h, h - E_1 h + a^\top(q - E_1 q))$, with $E_0 h = 0$. The tangent-space orthogonal complement consists of $k(w) = (-\rho_1 r w/\rho_0, w)$ with $E_1(Qw) = 0$. The normal equations project the unrestricted gradient onto that complement as $k(\rho_0 e/(\rho_1\omega))$. Subtraction gives (9); Pythagoras gives (10). Appendix B and Appendix C establish the path, closure, and projection details. □

$$\underbrace{\zeta}_{\text{unrestricted contrast}} \longrightarrow \underbrace{\Pi_Q\zeta}_{\text{weighted projection}} \longrightarrow \underbrace{e = \zeta - \Pi_Q\zeta}_{\text{excluded residual}} \longrightarrow \underbrace{(\rho_0/\rho_1)E_1(e^2/\omega)}_{\text{variance gain}}$$

The restriction helps only if it rules out a nonzero EOT-relevant direction. For nested correctly specified tilt spaces $\mathrm{span}(Q_1) \subset \mathrm{span}(Q_2)$, we have $V_{\mathrm{DRM}}(Q_1) \leq V_{\mathrm{DRM}}(Q_2) \leq V_{\mathrm{NP}}$: a larger model permits more local alternatives and weakens the population gain. This does not justify selecting a basis by minimizing estimated variance.

When $F_0 = F_1$ is symmetric about zero on the real line and $q(x) = x$, the centered potential ϕ is even. The gain becomes $(\rho_0 - \rho_1)^2 \mathrm{Var}(\phi)/(\rho_0\rho_1)$. In this example, imbalance can therefore make valid information sharing especially valuable, whereas balanced sampling yields no gain. This is not a universal monotonicity claim. Appendix C gives the full derivation. More generally, when $F_0 = F_1$ and the sampling fractions are balanced under a symmetric cost, the centered potential contrast satisfies $\zeta \equiv 0$, so no first-order efficiency gain is available regardless of the tilt basis.

5 Efficient Estimation and Robustness

5.1 Attainable efficiency and inference

Regularity requires the estimator's centered first-order limit to remain stable under small changes within the model, not just at one fixed distribution.

Theorem 2 (Efficient DRM-EOT estimation)**.** *Under Assumption 1, DRM-EOT is regular, asymptotically linear with the canonical gradient in* (9)*, and semiparametrically efficient:*

$$\sqrt{n}(\widehat{\Psi}_{\varepsilon}^{\mathrm{DRM}} - \Psi_{\varepsilon}) \Rightarrow N(0, V_{\mathrm{DRM}}). \tag{11}$$

Appendix D supplies the exact finite-n score expansion, fixed-group tilt covariance, uniform fitted-measure process, and transport delta method. In particular, no rate of $\rho_{gn} \to \rho_g$ is imposed. Regularity along differentiable DRM submodels, rather than a pointwise variance match alone, establishes efficiency.

Theorem 3 (Feasible inference)**.** *Under the same assumptions, the analytic estimator $\widehat{V}_{\mathrm{DRM}}$ in Appendix E satisfies $\widehat{V}_{\mathrm{DRM}} \to_p V_{\mathrm{DRM}}$. A stratified refit bootstrap that resamples n_g observations independently within each observed group and refits DRM-EOT has conditional limit $N(0, V_{\mathrm{DRM}})$ for $\sqrt{n}(\widehat{\Psi}_{\varepsilon}^{*} - \widehat{\Psi}_{\varepsilon}^{\mathrm{DRM}})$. If $V_{\mathrm{DRM}} > 0$, Wald and centered basic-bootstrap confidence intervals are asymptotically valid.*

The analytic estimator uses fitted transport potentials and projection coefficients, centering each raw group separately. In finite samples, $\widehat{V}_{\mathrm{DRM}}$ is computed as the sampling-fraction-weighted sum of within-group empirical variances of the fitted efficient influence functions. The Wald interval is $\widehat{\Psi}_{\varepsilon}^{\mathrm{DRM}} \pm \Phi^{-1}(1-\alpha/2)\sqrt{\widehat{V}_{\mathrm{DRM}}/n}$. Here Φ is the standard normal distribution function. Bootstrap refitting is essential: tilt estimation contributes at first order. Resampling the fitted pooled distribution would describe a different experiment. Full variance and conditional-process formulas are in Appendix E.

5.2 Efficiency–Robustness Trade-off

For bounded δ with $E_1\delta = 0$, consider a contiguous departure $dF_{1,n} = (1 + \eta\delta/\sqrt{n})dF_1$, leaving F_0 fixed. Centering at the perturbed transport target gives limiting mean shift ηB_δ, where

$$B_\delta = \rho_1 E_1(g_1\delta) - E_1(\phi_1\delta) = -\rho_0 E_1(e\delta/\omega). \tag{12}$$

Consequently,

$$|B_\delta| \le \sqrt{\rho_0\rho_1(V_{\mathrm{NP}} - V_{\mathrm{DRM}})E_1(\delta^2/\omega)}. \tag{13}$$

This is an unconstrained Cauchy–Schwarz upper relaxation and need not be sharp over admissible perturbations satisfying $E_1\delta = 0$. The equality direction for the unconstrained bound need not define a valid probability perturbation, so (13) should not be interpreted as the exact worst-case sensitivity or as an ordering of robustness across models. The component removed to gain efficiency also determines first-order vulnerability to off-model perturbations. At a fixed weighted perturbation norm, a larger efficiency gain enlarges the displayed upper bound on local bias. The actual bias also depends on the direction δ. The result describes local, direction-specific geometry; it neither guarantees robustness nor provides a quantitative theorem for arbitrary natural shifts. Appendix F provides the contiguous-alternative calculation.

6 Experiments

The experiments have three stages: controlled synthetic validation, exact-model efficiency on real OfficeHome representation geometry, and natural OfficeHome shifts, where the DRM is only approximate and potentially misspecified. The first two stages evaluate the implications of the efficiency theorem; the final stage examines behavior beyond the exact model. Primary OfficeHome summaries use paired-success repetitions to diagnose fitted DRM behavior and disclose fit success, while Appendix H additionally reports unconditional Algorithm 1 performance using the prescribed NP-EOT fallback for failed DRM fits; references, aggregation, and numerical protocols are also in Appendix H.

Table 1: Controlled synthetic validation at $n = 400$. NP-EOT denotes EOT computed from the two raw empirical marginals. DRM-EOT uses the density-ratio empirical-likelihood fitted marginals. The two correctly specified DRM bases are $Q_1(x) = (1, x)^\top$ and $Q_2(x) = (1, x, x^2)^\top$. Variances and Wald coverage use 2,000 repetitions; bootstrap coverage uses 1,000 repetitions, each with 499 stratified refit bootstrap replicates.

Method	Theory V	MC nVar	Theory gain	MC gain	Wald cov.	Boot. cov.
NP-EOT	0.06581	0.07030	0	0	–	–
DRM-EOT Q_1	0.04212	0.04444	36.0%	36.8%	96.05%	97.10%
DRM-EOT Q_2	0.06581	0.07028	≈ 0	≈ 0	95.15%	95.50%

6.1 Controlled synthetic validation

We compare three estimators of the same EOT functional: NP-EOT, based on the two raw empirical marginals, and DRM-EOT under two correctly specified tilt bases, $Q_1(x) = (1, x)^\top$ and $Q_2(x) = (1, x, x^2)^\top$. For $F_0 = F_1 = \mathrm{Unif}[-1, 1]$, $\varepsilon = 1$, and $(n_0, n_1) = (80, 320)$, the target is 0.2799288 (800-node quadrature). Table 1 verifies the projection identity: Q_1 yields a theoretical gain of 36.0% and a Monte Carlo gain of 36.8%, whereas Q_2 yields essentially no first-order gain. Wald and refit-bootstrap coverage are near nominal, with mild bootstrap conservatism for Q_1. Sample-size calibration, complete local-departure plots, and coverage details are in Appendix G.

We next test the efficiency theorem on real neural representation geometry under an exact density-ratio model.

6.2 OfficeHome setup

OfficeHome contains 15,588 images, 65 classes, and four domains: Art, Clipart, Product, and Real World (Venkateswara et al., 2017). We freeze ImageNet-pretrained ResNet-50 penultimate features (He et al., 2016), fit PCA using only the non-target domains, and L2-normalize the retained coordinates. Here $n_S = n_0$ and $n_T = n_1$ denote the source and target sample sizes. The main results use $d = 16$, $Q(z) = (1, z^\top)^\top$, $c(z, z') = \|z - z'\|^2/2$, and $\varepsilon = 0.1$. Labels are never used for discrepancy estimation. The observed and semi-synthetic representations lie on the unit sphere and thus form a bounded representation domain consistent with the compact-support setting of the theory. This does not establish compact support for an unnormalized latent population. The theory conditions on the representation map and excludes uncertainty from representation learning or PCA estimation.

6.3 Exact-model validation on real representation geometry

Denote the semi-synthetic finite-support populations by P_0, P_1. For each base-domain support $\{z_i\}_{i=1}^{N_K}$, set $P_0(z_i) = 1/N_K$ and $P_1(z_i) = P_0(z_i) \exp\{\alpha + \beta^\top z_i\}$. We fix the first-PC tilt direction ex ante, without using variance-ratio or RMSE outcomes. We calibrate the tilt along the first retained PCA coordinate to $\mathrm{ESS}(P_1)/N_K \in \{0.8, 0.5\}$, where $\mathrm{ESS}(P_1) = 1/\sum_i P_1(z_i)^2$. Thus the finite-support DRM is exact, and its population EOT target is computed directly, to numerical tolerance, rather than estimated from a reference sample. The efficiency theorem applies directly conditional on this fixed support. For every evaluated geometry/allocation, $V_{\mathrm{DRM}} < V_{\mathrm{NP}}$; at fixed $(\rho_0, \rho_1) = (0.8, 0.2)$, the four-geometry mean oracle ratios are 0.929 for ESS ratio 0.8 and 0.952 for 0.5.

6.4 Fixed-proportion scaling and high-R confirmation

We keep $(\rho_0, \rho_1) = (0.8, 0.2)$ while increasing n through $500, 1000, 2000, 4000, 8000$. Figure 1 compares empirical variance ratios with the corresponding fixed oracle levels. The ratios move from above one toward the oracle levels below one, although Monte Carlo fluctuations make the paths nonmonotone. This comparison matters because a smaller asymptotic bound need not imply a smaller variance at every finite sample size.

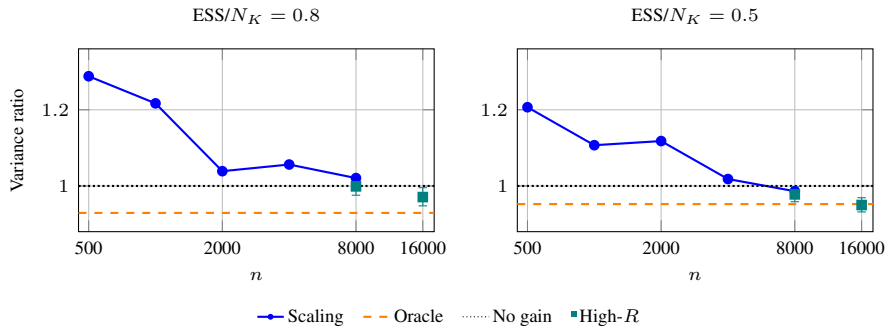


Figure 1: Exact-DRM OfficeHome fixed-proportion scaling on four fixed representation geometries. Each panel compares the mean task-level variance ratio with its oracle mean and $y = 1$. The curve uses the original scaling run; high-R points and 95% Monte Carlo intervals are from an independent confirmation run, summarized in Table 2.

Table 2: High-R exact-DRM confirmation. All ratios are DRM-EOT/NP-EOT, averaged over four geometry-level summaries. $R = 500$ per cell; intervals are 95% Monte Carlo intervals.

ESS/N_K	n	Var. ratio	MC interval	Oracle ratio	RMSE ratio
0.5	8000	0.978	$[0.959, 0.997]$	0.952	0.944
0.5	16000	0.950	$[0.932, 0.969]$	0.952	0.958
0.8	8000	0.999	$[0.976, 1.023]$	0.929	0.867
0.8	16000	0.971	$[0.948, 0.995]$	0.929	0.911

The high-R experiment uses $R = 500$ repetitions per geometry and cell. At ESS ratio 0.5 and $n = 16000$, the empirical ratio of 0.950 has a 95% Monte Carlo interval of $[0.932, 0.969]$ and closely matches the oracle value of 0.952. At ESS ratio 0.8, the ratio of 0.971, with interval $[0.948, 0.995]$, also confirms variance reduction, although it remains above the oracle value of 0.929. The $n = 8000$, ESS-0.8 interval includes one, so we do not claim confirmation for that cell. These intervals quantify simulation uncertainty conditional on the four evaluated geometries, not variation over a population of domains.

6.5 Natural OfficeHome: bias reduction, not variance efficiency

We turn to natural OfficeHome shifts, where the DRM is only approximate and potentially misspecified.

For the 12 directed tasks, we fix $n_S = 1000$ and use $n_T = 50, 100, 200, 400, 800$, with 100 repetitions per task. The target is a full-sample empirical EOT reference, not a known population truth. Ratios of task-averaged RMSE are $0.739, 0.834, 0.918, 1.032, 1.091$ (Figure 2). DRM-EOT improves $12/12$, $12/12$, and $9/12$ tasks at the three smallest target sizes. Yet the corresponding mean task-level variance ratios are $1.123, 1.124, 1.093$, all above one. At $n_T = 50$, mean absolute reference bias falls from 0.0219 to 0.0078. The target-scarce RMSE gain therefore reflects finite-sample bias reduction or structural regularization; it is *not* evidence of variance efficiency or validation of the exact-DRM theorem. Because DRM-EOT replaces the raw target empirical marginal with a pooled fitted marginal, the observed natural-shift bias reduction reflects the net effect of structured information sharing; it cannot be attributed solely to density-ratio parameter estimation. At larger target sizes, the bias advantage weakens and NP-EOT can be preferable.

6.6 Complexity and misspecification

Increasing d from 16 to 32 to 64 reveals two distinct mechanisms. At $n_T = 50$, the fraction of repetitions with $\max_i |\widehat{\theta}^\top Q(Z_i)| > 30$ rises from 0 to 0.098 to 0.600. The mean tilt norm likewise

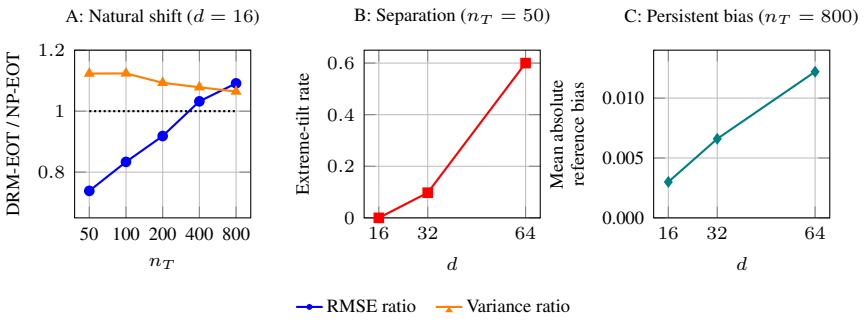


Figure 2: Natural shifts and complexity, with separate diagnostic axes. A: ratio of task-averaged RMSE and mean task-level variance ratio ($d = 16$). B: extreme-tilt frequency as a diagnostic of finite-sample separation, $\widehat{\Pr}\Big(\max_i |\widehat{\theta}^\top Q(Z_i)| > 30\Big)$, at $n_T = 50$. C: mean absolute DRM reference bias at $n_T = 800$, without rescaling. Panel C shows persistent reference bias after extreme-tilt events essentially disappear, consistent with misspecification; natural shifts do not validate the exact-DRM efficiency theorem.

rises from 10.3 to 35.35 to 853.8, while fit success falls from 1.000 to 0.999 to 0.958. At $d = 64$ and $n_T = 50$, including all repetitions through the NP-EOT fallback changes the RMSE ratio from 0.718 to 0.725 and the variance ratio from 1.525 to 1.536, leaving the qualitative conclusion unchanged. These are diagnostics of finite-sample near-separation and overparameterization; the threshold is not an exact test for separation. At $n_T = 800$, extreme-tilt events essentially disappear, but mean absolute DRM reference bias increases from 0.0030 to 0.0066 to 0.0122, indicating persistent reference bias consistent with misspecification over the evaluated range. Higher dimension is not uniformly worse—it can further improve small-target RMSE—but neither additional flexibility nor numerical stability establishes exact specification. The local geometry in (12) explains why structural restrictions can be vulnerable; it does not quantitatively predict these nonlocal natural-shift biases.

As a secondary sensitivity check, we vary ε at $n_T = 50$. The RMSE ratios for $\varepsilon = 0.05, 0.10, 0.20, 0.50$ are $0.566, 0.739, 0.938, 1.001$, respectively. NP-EOT and DRM-EOT become more similar in this observed grid as regularization increases. This is an empirical pattern, not a theorem that smaller ε guarantees a larger gain. Appendix H reports every sensitivity cell, domain-specific interval, oracle bound, and task-level result.

Practical diagnostics. The exact efficiency guarantee requires a credible low-dimensional density-ratio restriction. Large fitted log-ratios, near-separation, low effective sample size, or strong sensitivity to tilt dimension should be treated as specification warnings rather than evidence for a larger efficiency gain. Under natural shifts, DRM-EOT is therefore better interpreted as structural regularization unless exact specification is scientifically defensible. These diagnostics call for scrutiny of the assumed linkage between marginals; they do not constitute a formal model-selection rule or establish a universally optimal tilt dimension.

7 Discussion and Limitations

DRM-EOT is designed to estimate discrepancies between plausibly linked distributions of frozen representations, especially when one group is sample-poor. It is not a new discrepancy, representation learner, or domain-adaptation algorithm. Under exact specification, DRM-EOT attains the efficiency bound. Approximate shifts entail bias–variance and complexity–misspecification tradeoffs.

Our theorems require compact support, fixed $\varepsilon > 0$, a fixed low-dimensional continuous basis, and exact DRM specification. The current theory excludes vanishing target fractions, including fixed target sample sizes and target samples that grow sublinearly relative to the source sample. Natural shifts need not satisfy this model. Representation and PCA uncertainty are excluded; the analysis is conditional on the frozen map. We prove neither shrinking-ε inference nor inference for Sinkhorn

divergence, and we do not justify data-driven basis selection. Studentized intervals require positive first-order variance. Unbounded-support extensions need additional controls (Appendix I). The numerical evidence concerns four evaluated OfficeHome geometries and their directed shifts, not all possible domains. Penalized empirical likelihood or regularized discriminative density-ratio fitting may mitigate finite-sample separation; characterizing penalty sequences that preserve the first-order efficiency result is left for future work.

Reproducibility Statement

Appendices A–F provide the assumptions and proofs. Appendices G–H document the synthetic and OfficeHome protocols, preprocessing, sampling, solver tolerances, aggregation, and Monte Carlo intervals. All theoretical claims condition on the fixed representation map and basis.

AI Use Statement

Generative AI tools were used to assist with language editing, notation and cross-reference checks, literature verification, mathematical proof review, and verification code for this revision. The authors are responsible for the final text, mathematical claims, and reported results.

References

Medha Agarwal and Alex Luedtke. Sinkhorn treatment effects: A causal optimal transport measure. arXiv preprint arXiv:2605.08485, 2026.

Anik Burman, Margaret Gamalo, Promit Ghosal, and Prosenjit Kundu. Transporting randomized trial effects to real-world populations via Riesz-calibrated optimal transport. arXiv preprint arXiv:2608.23453, 2026.

Jiahua Chen and Yukun Liu. Quantile and quantile-function estimations under density ratio model. *The Annals of Statistics*, 41(3):1669–1692, 2013. doi: 10.1214/13-AOS1129.

Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. In *Advances in Neural Information Processing Systems*, volume 30, pp. 3730–3739, 2017a.

Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9):1853–1865, 2017b. doi: 10.1109/TPAMI.2016.2615921.

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

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. Statistical inference with regularized optimal transport. *Information and Inference: A Journal of the IMA*, 13(1):iaad056, 2024. doi: 10.1093/imaiai/iaad056.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016. doi: 10.1109/CVPR.2016.90.

Gonzalo Mena and Jonathan Niles-Weed. Statistical bounds for entropic optimal transport: Sample complexity and the central limit theorem. In *Advances in Neural Information Processing Systems*, volume 32, pp. 4541–4551, 2019.

Marcel Nutz and Johannes Wiesel. Stability of Schrödinger potentials and convergence of Sinkhorn's algorithm. *The Annals of Probability*, 51(2):699–722, 2023. doi: 10.1214/22-AOP1611.

Jing Qin. Inferences for case-control and semiparametric two-sample density ratio models. *Biometrika*, 85(3):619–630, 1998. doi: 10.1093/biomet/85.3.619.

Masashi Sugiyama, Shinichi Nakajima, Hisashi Kashima, Paul von Bünau, and Motoaki Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In *Advances in Neural Information Processing Systems*, volume 20, pp. 1433–1440, 2007.

Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5385–5394, 2017. doi: 10.1109/CVPR.2017.572.

Weiwei Zhuang, B. Y. Hu, and Jiahua Chen. Semiparametric inference for the dominance index under the density ratio model. *Biometrika*, 106(1):229–241, 2019. doi: 10.1093/biomet/asy068.

Weiwei Zhuang, Weiqi Yang, and Jiahua Chen. Bootstrap consistency for empirical likelihood in density ratio models. *Statistica Sinica*, forthcoming. doi: 10.5705/ss.202025.0398. Accepted author version, Preprint No. SS-2025-0398.