ENTROPICALLY REGULARIZED MARTINGALE OPTIMAL TRANSPORT WITH L_1 RELAXATION

Ethan Chen
Georgia Institute of Technology
echen347@gatech.edu

ABSTRACT

Martingale Optimal Transport (MOT) provides a framework for model-free derivative pricing, but the hard martingale constraint becomes statistically unstable under finite-sample estimation. We introduce a regularized formulation that replaces the exact martingale condition with an L_1 penalty on the martingale defect, combined with entropic smoothing of the transport plan. The resulting dual problem admits a natural interpretation in terms of leverage-constrained hedging strategies and is amenable to proximal gradient methods with $O(1/k)$ convergence. Our main theoretical contribution is a concentration analysis establishing that the intrinsic statistical fluctuation in the martingale constraint scales as $O(N^{-1/2})$, yielding a principled lower bound for the relaxation parameter. Experiments on synthetic data confirm this scaling and demonstrate that the relaxed formulation remains tractable as the constraint tightens.

1 INTRODUCTION

Martingale Optimal Transport (MOT) extends classical optimal transport by imposing that the transport plan preserve conditional expectations: $\mathbb{E}[Y|X] = X$. This constraint encodes the absence of arbitrage under model uncertainty (Beiglböck et al., 2013; Galichon et al., 2014). Given marginal distributions of an asset at times t_0 and t_1, MOT yields model-free bounds on path-dependent derivatives.

The strict martingale constraint is problematic for estimation. By Strassen's theorem, a martingale coupling exists only if $\mu \preceq_{cx} \nu$, a condition generically violated when marginals are estimated from finite samples. Even when feasibility holds, the set of exact martingale couplings has measure zero within the space of all couplings, forcing any consistent estimator to interpolate sampling noise.

We develop a relaxed formulation that replaces the hard martingale constraint with a penalty on the martingale defect. The problem remains well-posed when empirical marginals violate convex order. The dual variables correspond to static hedging positions, and the relaxation parameter controls an implicit leverage bound.

Contributions.

1. A primal-dual formulation replacing the martingale constraint with an L_1 penalty, yielding a dual over leverage-bounded hedging strategies (Section 2.2).
2. A proximal gradient algorithm with $O(1/k)$ convergence (Section 3.1).
3. A concentration analysis showing that empirical martingale violations scale as $O(N^{-1/2})$, providing a lower bound for the relaxation parameter (Section 3.2).

Section 4 validates these results experimentally. Section 5 surveys related work.

2 Problem Formulation and Dual Analysis

2.1 Setup and Notation

Let $\mathbf{1}_k$ denote the vector of ones in $\mathbb{R}^k$, and take $\mu \in \mathbb{R}^n_+$ and $\nu \in \mathbb{R}^m_+$ to be discrete probability measures supported on $X = \{x_1, \ldots, x_n\} \subset \mathbb{R}$ and $Y = \{y_1, \ldots, y_m\} \subset \mathbb{R}$, respectively. The set of couplings is

$$\Pi(\mu, \nu) = \{\pi \in \mathbb{R}^{n\times m}_+ \mid \pi \mathbf{1}_m = \mu,\ \pi^\top \mathbf{1}_n = \nu\}.$$

For a coupling π and cost matrix $C \in \mathbb{R}^{n\times m}$, the transport cost is $\langle C, \pi\rangle = \sum_{i,j} C_{ij}\pi_{ij}$.

The *martingale defect* of π at x_i is

$$d_i(\pi) := \mathbb{E}_\pi[Y|X = x_i] - x_i = \frac{1}{\mu_i}\sum_{j=1}^{m} \pi_{ij}(y_j - x_i), \tag{1}$$

with the convention $0/0 = 0$. The martingale constraint $\mathbb{E}_\pi[Y|X] = X$ is equivalent to $d_i(\pi) = 0$ for all i.

We measure aggregate violation in the $L_1(\mu)$ norm:

$$\|d(\pi)\|_{L_1(\mu)} = \sum_{i=1}^{n} \mu_i |d_i(\pi)| = \sum_{i=1}^{n} \left|\sum_{j=1}^{m} \pi_{ij}(y_j - x_i)\right|,$$

where the μ_i weights cancel with the normalization in Equation 1.

The relaxed MOT problem with entropic regularization is:

$$\min_{\pi\in\Pi(\mu,\nu)} \langle C, \pi\rangle + \varepsilon H(\pi) \quad \text{s.t.} \quad \|d(\pi)\|_{L_1(\mu)} \leq \delta, \tag{2}$$

where $H(\pi) = \sum_{i,j} \pi_{ij}(\log \pi_{ij} - 1)$ is the negative entropy, $\varepsilon > 0$ is the regularization strength, and $\delta \geq 0$ is the relaxation parameter. Setting $\delta = 0$ recovers the hard-constrained problem.

Remark 2.1 (Choice of norm)**.** *The $L_1(\mu)$ norm yields a dual with bounded feasible set (Section 2.2). The $L_2(\mu)$ norm yields a smooth constraint but amplifies large pointwise violations. The L_∞ norm controls worst-case error but may be overly restrictive.*

2.2 Dual Analysis

To handle the non-smooth L_1 constraint, we use the variational characterization

$$\|v\|_1 = \sup_{\|h\|_\infty \leq 1} \langle h, v\rangle.$$

Introducing a Lagrange multiplier $\lambda \geq 0$ for the constraint and a dual variable $h \in \mathbb{R}^n$ with $\|h\|_\infty \leq 1$, the Lagrangian is

$$\mathcal{L}(\pi, h, \lambda) = \langle C, \pi\rangle + \varepsilon H(\pi) + \lambda \left(\sum_i h_i \sum_j \pi_{ij}(y_j - x_i) - \delta\|h\|_\infty\right). \tag{3}$$

Define $\tilde{h} = \lambda h$. Since $\lambda \geq 0$ and $\|h\|_\infty \leq 1$, the variable $\tilde{h}$ ranges over all of $\mathbb{R}^n$. Relabeling $\tilde{h} \to h$ and applying strong duality, justified by Slater's condition when $\delta > 0$, we obtain

$$\min_{\pi\in\Pi(\mu,\nu)} \sup_{h\in\mathbb{R}^n} \left\{\langle C, \pi\rangle + \varepsilon H(\pi) + \sum_{i,j} \pi_{ij} h_i (y_j - x_i) - \delta\|h\|_\infty\right\}. \tag{4}$$

The objective is convex in π and concave in h. By Sion's minimax theorem, the inner minimization over π is

$$\min_{\pi\in\Pi(\mu,\nu)} \sum_{i,j} \pi_{ij} \underbrace{(C_{ij} + h_i(y_j - x_i))}_{=:\tilde{C}_{ij}(h)} + \varepsilon H(\pi), \tag{5}$$

which is entropic optimal transport with modified cost $\tilde{C}(h)$. Let $W_\varepsilon(\tilde{C})$ denote its optimal value. The dual problem is

$$\max_{h\in\mathbb{R}^n} J(h) := W_\varepsilon(\tilde{C}(h)) - \delta\|h\|_\infty. \tag{6}$$

Remark 2.2 (Financial interpretation)**.** *The dual variable h corresponds to a static hedge in the underlying asset. In standard MOT, the dual is unconstrained, permitting unbounded positions. Here, the penalty $-\delta\|h\|_\infty$ imposes an implicit leverage bound; as $\delta \to 0$, this bound vanishes and we recover the unconstrained dual. Conversely, $\delta > 0$ allows small martingale violations, corresponding to approximate arbitrage within tolerance δ. This is consistent with pricing under proportional transaction costs (Dolinsky & Soner, 2014).*

3 Algorithm and Statistical Calibration

3.1 Algorithm

The dual objective Equation 6 decomposes as $J(h) = W_\varepsilon(\tilde{C}(h)) - \delta\|h\|_\infty$: a smooth term minus a non-smooth convex penalty. This structure admits proximal gradient ascent.

The function $h \mapsto W_\varepsilon(\tilde{C}(h))$ is differentiable by strict convexity of the entropy. By Danskin's theorem, the gradient is the martingale defect of the optimal plan:

$$\nabla_h W_\varepsilon = d(\pi^\star(h)), \quad \text{i.e.,} \quad [\nabla_h W_\varepsilon]_i = \sum_j \pi^\star_{ij}(h)(y_j - x_i), \tag{7}$$

where $\pi^\star(h) = \arg\min_{\pi \in \Pi(\mu,\nu)} \langle \tilde{C}(h), \pi\rangle + \varepsilon H(\pi)$ is computed via the Sinkhorn algorithm (Cuturi, 2013).

Let $\mathcal{B}_1 = \{v \in \mathbb{R}^n : \|v\|_1 \leq 1\}$ denote the unit L_1 ball. The proximal operator of the L_∞ norm is computed via Moreau decomposition:

$$\operatorname{prox}_{\tau\|\cdot\|_\infty}(v) = v - \tau \cdot \operatorname{proj}_{\mathcal{B}_1}(v/\tau), \tag{8}$$

where $\operatorname{proj}_{\mathcal{B}_1}$ is projection onto $\mathcal{B}_1$, computable in $O(n \log n)$ time (Duchi et al., 2008).

Algorithm 1 Proximal Gradient Ascent for Relaxed MOT

Require: Marginals μ, ν; cost C; parameters $\varepsilon, \delta, \gamma$; max iterations K; tolerance `tol`
1: Initialize $h \leftarrow \mathbf{0}$
2: **for** $k = 1, \ldots, K$ **do**
3: $\quad \pi \leftarrow \textsc{Sinkhorn}(\mu, \nu, C_{ij} + h_i(y_j - x_i), \varepsilon)$
4: $\quad v \leftarrow h + \gamma \cdot d(\pi)$
5: $\quad h_{\text{new}} \leftarrow v - \gamma\delta \cdot \operatorname{proj}_{\mathcal{B}_1}(v/(\gamma\delta))$
6: $\quad$ **if** $\|h_{\text{new}} - h\|_2 <$ `tol` **then break**
7: $\quad$ **end if**
8: $\quad h \leftarrow h_{\text{new}}$
9: **end for**
10: **return** π, h

Proposition 3.1 (Convergence)**.** *Assume $|y_j - x_i| \leq B$ for all i, j. The smooth component $W_\varepsilon(\tilde{C}(h))$ has Lipschitz continuous gradient with constant $L = B^2/\varepsilon$. For step size $\gamma = 1/L$, the sequence generated by Algorithm 1 satisfies*

$$J(h^\star) - J(h^{(K)}) \leq \frac{L\|h^\star\|_2^2}{2K},$$

where $h^\star$ is any dual optimum.

Proof. See Appendix A. □

Remark 3.1. *Accelerated variants (FISTA) achieve $O(1/K^2)$ rates. In practice, backtracking line search can improve convergence, particularly when ε is small.*

3.2 Choosing the Relaxation Parameter

If δ is too small, statistical noise renders the problem infeasible. If δ is too large, the martingale structure is lost. We establish a lower bound based on concentration of the empirical martingale defect.

Consider i.i.d. samples $\{(X_i, Y_i)\}_{i=1}^N$ drawn from a coupling π_0 satisfying $\mathbb{E}[Y|X] = X$, and define the empirical martingale defect

$$D_N := \sum_{i=1}^{N} |Y_i - X_i| \,. \tag{9}$$

When the empirical marginals have full support, the identity coupling $\hat{\pi} = \text{diag}(1/N)$ is a valid transport plan with martingale violation D_N/N.

Assumption 1. *The support satisfies $|Y - X| \leq B$ almost surely.*

Theorem 3.2. *Under Assumption 1, for any $t > 0$:*

$$\mathbb{P}\left(\left|\frac{D_N}{N} - \mathbb{E}|Y - X|\right| \geq t\right) \leq 2\exp\left(-\frac{Nt^2}{2B^2}\right). \tag{10}$$

Proof. Define $f(Z_1, \ldots, Z_N) = \frac{1}{N}\sum_{i=1}^{N} |Y_i - X_i|$ where $Z_i = (X_i, Y_i)$. Changing a single Z_i alters f by at most $\frac{2B}{N}$, so the bounded differences condition holds with $c_i = \frac{2B}{N}$. McDiarmid's inequality (McDiarmid, 1989) gives

$$\mathbb{P}(|f - \mathbb{E}f| \geq t) \leq 2\exp\left(-\frac{2t^2}{\sum_i c_i^2}\right) = 2\exp\left(-\frac{2t^2}{4B^2/N}\right).$$

□

Theorem 3.2 establishes that the empirical martingale defect concentrates at rate $O(N^{-1/2})$. Setting the right-hand side equal to α and solving for t yields the scaling

$$\delta_N = \kappa B \sqrt{\frac{\log(2/\alpha)}{N}} = O(N^{-1/2}), \tag{11}$$

where $\alpha \in (0, 1)$ is the desired confidence level and $\kappa = \sqrt{2}$. This ensures feasibility with probability at least $1 - \alpha$ for finite N, while $\delta_N \to 0$ as $N \to \infty$.

4 Experiments

4.1 Setup

We validate on synthetic 1-D Gaussian data satisfying convex order. The source measure is $\mu = \mathcal{N}(0, 0.5^2)$ and the target is $\nu = \mathcal{N}(0, 1^2)$, discretized on a uniform grid of 50 points per marginal. The cost is quadratic $C_{ij} = (x_i - y_j)^2$. For all experiments, we set $\varepsilon = 0.05$ (approximately 1% of the maximum cost).

Method	**Transport Cost**	**Martingale Violation**
Unconstrained Sinkhorn	0.2637	0.3746
Relaxed MOT ($\delta = 0.1$)	0.5041	0.1099
Relaxed MOT ($\delta = 0.01$)	0.6573	0.0269

Table 1: Unconstrained Sinkhorn vs. Relaxed MOT. Smaller δ reduces martingale violation at the expense of higher transport cost.

4.2 VERIFICATION OF THEORETICAL RATES

We verify the two main quantitative predictions of Sections 3.2 and 3.1. For the concentration rate, we generated data from a martingale coupling $Y = X + \xi$ with $\xi \sim U[-0.5, 0.5]$ and measured the empirical martingale violation across sample sizes $N \in [50, 5000]$, averaged over 50 trials. Figure 1 confirms the $O(N^{-1/2})$ scaling of Theorem 3.2: the log-log slope is approximately -0.53. For the convergence rate, Figure 2 plots the dual optimality gap $J(h^\star) - J(h^{(k)})$ against iteration k; the observed decay matches the $O(1/k)$ rate of Proposition 3.1.

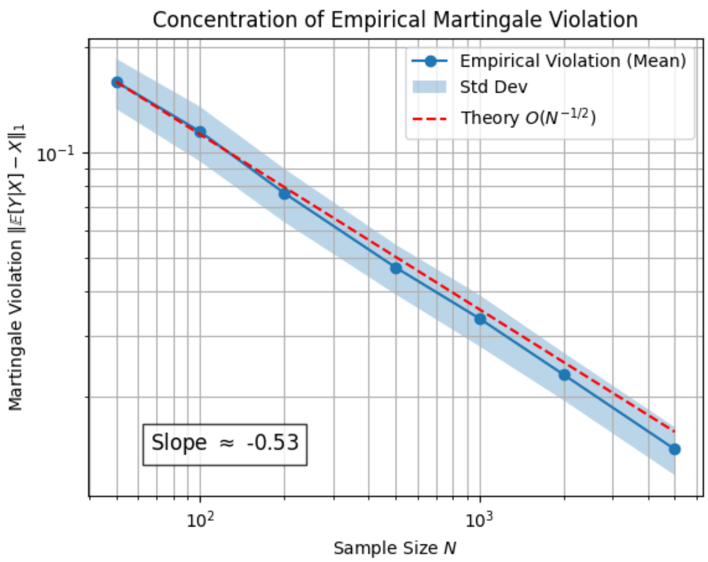

Figure 1: Log-log plot of empirical martingale violation vs. sample size N (averaged over 50 trials). Shaded regions indicate ± 1 standard deviation. The slope ≈ -0.5 confirms the $O(N^{-1/2})$ rate of Theorem 3.2.

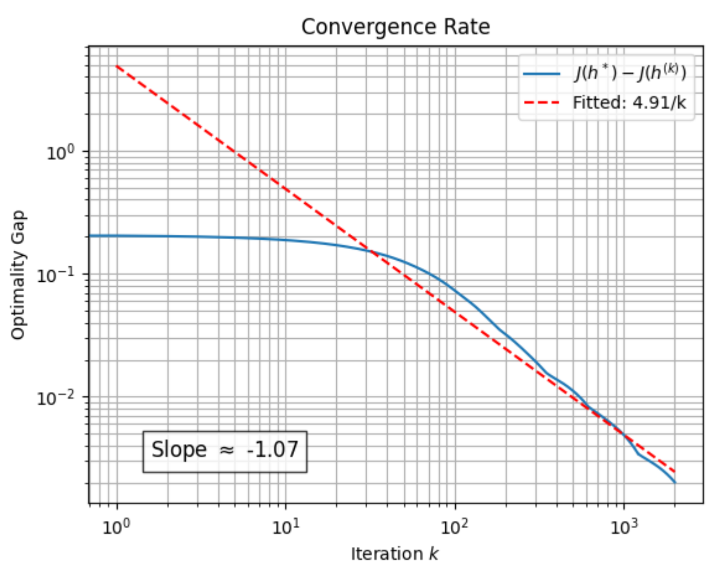

Figure 2: Convergence of the dual objective. The optimality gap decays at rate $O(1/k)$, consistent with Proposition 3.1.

4.3 Relaxation Trade-off

We varied δ from 0 to 0.5. Figure 3 shows the resulting trade-off: as δ decreases, the algorithm enforces the martingale property more tightly, reducing violation at the expense of higher transport cost. The Pareto frontier (middle panel) is smooth, indicating that the relaxation degrades gracefully. The dual variable norm $\|h\|_\infty$ (right panel) grows as $\delta \to 0$, consistent with the leverage-bound interpretation of Remark 2.2.

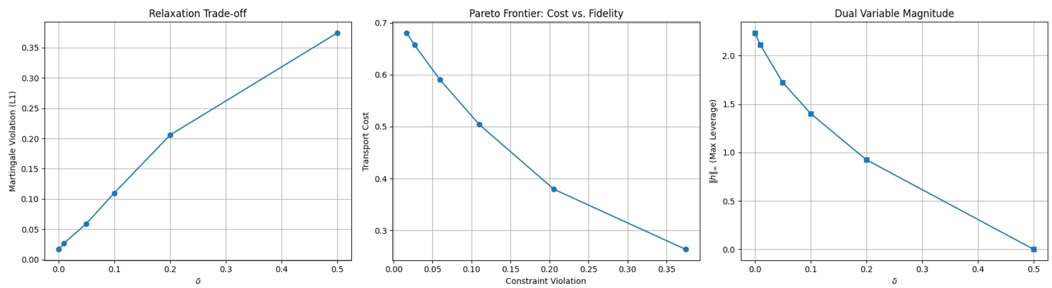

Figure 3: Trade-off curves. (Left) Martingale violation vs. δ. (Middle) Pareto frontier: transport cost vs. martingale violation. (Right) Dual variable norm $\|h\|_\infty$ vs. δ.

5 Related Work

Martingale optimal transport. Beiglböck et al. (2013) and Galichon et al. (2014) established the duality theory connecting MOT to robust finance. Hobson & Neuberger (2012) developed links to Skorokhod embeddings. Dolinsky & Soner (2014) extended the framework to continuous time and characterized the relationship between martingale constraints and transaction costs. Numerical methods include LP relaxations (Guo & Obłój, 2019) and neural network approaches (Eckstein & Kupper, 2021).

Entropic optimal transport. Cuturi (2013) introduced Sinkhorn-based computation for optimal transport. Statistical properties of entropic OT are studied in (Genevay et al., 2019; Mena & Niles-Weed, 2019).

Relaxation and regularization. The L_1 penalty on constraint violations is analogous to penalized regression (Tibshirani, 1996). Related soft-constraint formulations in OT include unbalanced transport (Chizat et al., 2018) and partial transport (Caffarelli & McCann, 2010).

Statistical OT. Sample complexity of OT estimators is studied in (Weed & Bach, 2019; Chewi et al., 2024). The concentration analysis in Section 3.2 uses bounded-difference arguments (McDiarmid, 1989; Wainwright, 2019).

6 Conclusion

We introduced a relaxed formulation of Martingale Optimal Transport that replaces the hard martingale constraint with an L_1 penalty on the conditional expectation defect. The dual problem admits an interpretation in terms of leverage-bounded hedging strategies, and the main theoretical result is a concentration bound establishing that empirical martingale violations scale as $O(N^{-1/2})$, which determines the minimum relaxation required for feasibility under finite-sample estimation.

The experiments confirm the theoretical rates and reveal additional structure. The Pareto frontier between transport cost and martingale fidelity is smooth, suggesting that the relaxation degrades gracefully rather than exhibiting phase transitions. The cost of enforcement is quantitatively mild: tightening δ from 0.1 to 0.01 reduces violation by a factor of four while increasing transport cost by roughly 30% (Table 1). The growth of $\|h\|_\infty$ as $\delta \to 0$ provides empirical confirmation of the leverage-bound interpretation: the dual variable norm diverges as the constraint becomes exact, consistent with the need for unbounded hedging positions in the hard-constrained MOT dual.

Limitations. All experiments use synthetic one-dimensional Gaussian marginals with a fixed entropic regularization $\varepsilon = 0.05$. The behavior under heavy-tailed marginals, higher-dimensional support, or real market option data remains untested. The concentration bound in Theorem 3.2 requires bounded support (Assumption 1), excluding common financial models such as log-normal asset prices; extending to sub-Gaussian or sub-exponential tails would broaden applicability.

Future directions. Extending to multi-marginal MOT introduces additional difficulty: the martingale constraint $\mathbb{E}[X_{t+1}|X_t] = X_t$ for each consecutive pair yields a constraint set that is generally non-convex when more than two marginals are involved. Adaptive selection of δ, for instance via Lepski's method, could yield tighter data-driven calibration than the worst-case bound in Equation 11. Finally, the connection between the relaxed formulation and pricing under transaction costs (Dolinsky & Soner, 2014) suggests links to distributionally robust optimization that merit further investigation.

REFERENCES

Amir Beck and Marc Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM Journal on Imaging Sciences*, 2(1):183–202, 2009.

Mathias Beiglböck, Pierre Henry-Labordère, and Friedrich Penkner. Model-independent bounds for option prices—a mass transport approach. *Finance and Stochastics*, 17(3):477–501, 2013.

Luis Caffarelli and Robert McCann. Free boundaries in optimal transport and Monge-Ampère obstacle problems. *Annals of Mathematics*, 171(2):673–730, 2010.

Sinho Chewi, Aram-Alexandre Pooladian, and Jonathan Niles-Weed. Statistical optimal transport. *arXiv preprint arXiv:2407.18163*, 2024.

Lénaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced optimal transport: dynamic and Kantorovich formulations. *Journal of Functional Analysis*, 274(11): 3090–3123, 2018.

Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, 2013.

Yan Dolinsky and H. Mete Soner. Martingale optimal transport and robust hedging in continuous time. *Probability Theory and Related Fields*, 160(1):391–427, 2014.

John Duchi, Shai Shalev-Shwartz, Yoram Singer, and Tushar Chandra. Efficient projections onto the ℓ_1-ball for learning in high dimensions. In *International Conference on Machine Learning*, 2008.

Stephan Eckstein and Michael Kupper. Computation of optimal transport and related hedging problems via penalization and neural networks. *Applied Mathematics & Optimization*, 83:639–667, 2021.

Alfred Galichon, Pierre Henry-Labordère, and Nizar Touzi. A stochastic control approach to no-arbitrage bounds given marginals. *Annals of Applied Probability*, 24(1):312–336, 2014.

Aude Genevay, Lénaïc Chizat, Francis Bach, Marco Cuturi, and Gabriel Peyré. Sample complexity of Sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, 2019.

Gaoyue Guo and Jan Obłój. Computational methods for martingale optimal transport problems. *Annals of Applied Probability*, 29(6):3311–3347, 2019.

David Hobson and Anthony Neuberger. Robust bounds for forward start options. *Mathematical Finance*, 22(1):31–56, 2012.

Colin McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics, 1989*, volume 141 of *London Math. Soc. Lecture Note Ser.*, pp. 148–188. Cambridge University Press, 1989.

Gonzalo Mena and Jonathan Niles-Weed. Statistical bounds for entropic optimal transport. In *Advances in Neural Information Processing Systems*, 2019.

Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B*, 58(1):267–288, 1996.

Martin J. Wainwright. *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*. Cambridge University Press, 2019.

Jonathan Weed and Francis Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *Bernoulli*, 25(4A):2620–2648, 2019.

A PROOF OF PROPOSITION 3.1

Proof. Let $\phi(h) := W_\varepsilon(\tilde{C}(h))$. The gradient is the martingale defect under the optimal entropic plan:

$$[\nabla\phi(h)]_i = \sum_{j=1}^{m} \pi^\star_{ij}(h)(y_j - x_i).$$

The entropic plan has the form

$$\pi^\star_{ij}(h) \propto \exp(-(C_{ij} + h_i(y_j - x_i))/\varepsilon).$$

Since $\pi^\star_{ij}(h) \propto \exp(-(C_{ij} + h_i(y_j - x_i))/\varepsilon)$, the optimal plan is an exponential family in h with sufficient statistics $(y_j - x_i)$. The Hessian $\nabla^2\phi(h)$ is therefore the covariance matrix of these sufficient statistics under $\pi^\star(h)$, scaled by $1/\varepsilon$. Since $|y_j - x_i| \le B$, Popoviciu's inequality gives $\|\nabla^2\phi(h)\|_2 \le B^2/\varepsilon$. Thus ϕ is L-smooth with $L = B^2/\varepsilon$.

The dual problem is to maximize $J(h) = \phi(h) - \delta\|h\|_\infty$, or equivalently minimize $F(h) = -\phi(h) + \delta\|h\|_\infty$. Algorithm 1 is ISTA applied to F. By Theorem 3.1 of Beck & Teboulle (2009), for step size $\gamma = 1/L$:

$$F(h^{(K)}) - F(h^\star) \le \frac{L\|h^{(0)} - h^\star\|_2^2}{2K}.$$

With initialization $h^{(0)} = \mathbf{0}$ and $L = B^2/\varepsilon$:

$$J(h^\star) - J(h^{(K)}) \le \frac{B^2\|h^\star\|_2^2}{2\varepsilon K}.$$

as desired. □